Bayesian characterization of uncertainty in species interaction strengths

Considerable effort has been devoted to the empirical estimation of species interaction strengths. This effort has focused primarily on statistical significance testing and on obtaining point estimates of parameters that contribute to interaction strength magnitude, leaving characterizations of estimation uncertainty and distinctions between the deterministic and stochastic contributions to variation largely unconsidered. Here we consider a means of quantifying interaction strength uncertainty by formulating an observational method for estimating per capita attack rates as a Bayesian statistical model. This formulation permits the explicit incorporation of multiple sources of uncertainty. In doing so we highlight the informative nature of several so-called non-informative prior choices in modeling the sparse data typical of predator feeding surveys and provide evidence for the superior performance of a new neutral prior choice. A case study application shows that while Bayesian point estimates may be made to correspond with those obtained by frequentist approaches, estimation uncertainty as described by the 95% intervals is more biologically realistic using the Bayesian method in that the lower bounds of the Bayesian posterior intervals for the attack rates do not include zero when the occurrence of a given predator-prey interaction is in fact observed. This contrasts with bootstrap confidence intervals that often do contain zero in such cases. The Bayesian approach provides a straightforward, probabilistic characterization of interaction strength uncertainty. In doing so it provides a framework for considering both the deterministic and stochastic drivers of species interactions and their impact on food web dynamics.

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

Bayesian characterization of uncertainty in species interaction strengths

(Running head: Bayesian Interaction Strengths)

Christopher Wolf[1*], **Mark Novak**[2†], **and Alix I. Gitelman**[1‡]

[1]Department of Statistics, Oregon State University,
Corvallis, Oregon 97331, United States
[2]Department of Integrative Biology, Oregon State University,
Corvallis, Oregon 97331, United States

September 18, 2015

---

[*] wolfch@science.oregonstate.edu
[†] mark.novak@oregonstate.edu
[‡] gitelman@science.oregonstate.edu

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

Abstract

Considerable effort has been devoted to the empirical estimation of species interaction strengths. This effort has focused primarily on statistical significance testing and on obtaining point estimates of parameters that contribute to interaction strength magnitude, leaving characterizations of estimation uncertainty and distinctions between the deterministic and stochastic contributions to variation largely unconsidered. Here we consider a means of quantifying interaction strength uncertainty by formulating an observational method for estimating per capita attack rates as a Bayesian statistical model. This formulation permits the explicit incorporation of multiple sources of uncertainty. In doing so we highlight the informative nature of several so-called non-informative prior choices in modeling the sparse data typical of predator feeding surveys and provide evidence for the superior performance of a new neutral prior choice. A case study application shows that while Bayesian point estimates may be made to correspond with those obtained by frequentist approaches, estimation uncertainty as described by the 95% intervals is more biologically realistic using the Bayesian method in that the lower bounds of the Bayesian posterior intervals for the attack rates do not include zero when the occurrence of a given predator-prey interaction is in fact observed. This contrasts with bootstrap confidence intervals that often do contain zero in such cases. The Bayesian approach provides a straightforward, probabilistic characterization of interaction strength uncertainty. In doing so it provides a framework for considering both the deterministic and stochastic drivers of species interactions and their impact on food web dynamics.

Keywords: nonlinear interaction strengths, predator-prey, functional response, JAGS, zero-inflated gamma, parameter estimation, non-informative neutral priors.

Introduction

Quantifying the strength of species interactions is an important ecological challenge. Estimates can be used to identify keystone species whose impacts are disproportionate to their abundance (Power et al., 1996), help explain community structure (Wootton, 1994), are key to understanding food web stability (Allesina and Tang, 2012), and underlie forecasts of community dynamics (Yodzis, 1988). Estimates made on a per capita basis are particularly useful in that they underlie all other measures of interaction strength (Laska and Wootton, 1998). In this regard, a fundamental component of predator-prey interaction strengths is the nature of the predator's functional response. For example, with linear Holling type I functional responses and linear density dependence in the prey, per capita interaction strengths correspond to the predators' per capita attack rates. For nonlinear multi-species functional responses, such as the Holling type II functional response which most predators exhibit (Jeschke et al., 2004), the per capita attack rates are parameters reflecting the predators' prey preferences (Chesson, 1983).

Unfortunately, estimating interaction strengths in natural systems is difficult. In most food webs the large number of pairwise interactions—and the large number of weak interactions in particular—makes the use of manipulative experiments logistically prohibitive. Thus, many have resorted to indirect means of estimation, such as using energetic and allometric principles (Wootton and Emmerson, 2005). More effort still has been devoted to estimating interaction strength parameters by characterizing predator functional responses, largely on a pairwise experimental basis or by using observations of predator gut contents and kill rates (Vucetich et al., 2002; Jeschke et al., 2004).

To date, most of the effort spent on measuring interaction strengths has focused on obtaining point estimates of parameters. For example, Paine (1992) used a bootstrapping procedure only to quantify the uncertainty associated with the mean net strength of pairwise species interactions due to variation among experimental replicates. The focus has similarly been on point estimates and their statistical significance in the use of functional response experiments designed to determine the dependence of feeding rates on prey- and/or predator densities. Thus only the "deterministic core" of alternative functional response formulations has been of interest (Vucetich et al., 2002; Jeschke et al., 2004). More specifically, functional responses have often been fit to data using statistical models such as $F = \frac{cN}{1+chN} + \epsilon$ (the Holling type II response) whereby variation in a predator's feeding rate (F) is assumed to be controlled by a deterministic component governed by variation in abundance (N), attack rate (c) and handling time (h), and only a "shell" of stochastic variation (ϵ) is used to describe the variation left unexplained by the deterministic core. This is in contrast to considering the variation that is intrinsic to both the parameters (c and h) and variables (N) by describing each by a distribution that is itself governed by deterministic and stochastic sources of variation.

The distinction between these two approaches to considering variation in interaction strength estimates is important when the uncertainty of estimates itself is of interest. This is particularly true when forecasting the dynamics of species rich communities where indirect effects can rapidly compound even small amounts of uncertainty (Yodzis, 1988; Novak et al., 2011). In such applications, knowledge of the (co-)variation of parameter estimates is essential to assessing the sensitivity of predictions under plausible scenarios of estimation uncertainty. Of course, estimates of uncertainty are also important in comparing the utility and consistency of different interaction strength estimation methods, and for the biological interpretation of the estimates themselves. Estimates derived from allometric relationships, for example, are typically associated with several orders of magnitude of variation (Rall et al., 2012).

Here we extend the observational method for estimating the per capita attack rates of predator-prey interactions presented by Novak and Wootton (2008) to characterize estimation uncertainty. Our interest in such observational methods stems from their ability to more easily accommodate instances of trophic omnivory than experimental and time-series methods, and because they retain the species-specific information lost in allometric and energetic approaches. Furthermore, with the method of Novak and Wootton (2008) attack rates may be estimated for multiple prey simultaneously while accounting for an inherent nonlinearity of predator-prey interactions because a multi-species Holling type II functional response is considered in the method's derivation. Our stochastic formulation connects a deterministic functional response model with each of the sources of empirical data that contribute to the per capita attack rate estimates. The Bayesian framework in which we implement our approach thereby permits us to account for both variation due to sampling effort and the environment, and thus to explicitly incorporate the deterministic and stochastic sources of uncertainty intrinsic to the attack rate estimates.

To assess the Bayesian method's utility we apply it to data collected by Novak (2010), contrasting these estimates with those obtained by non-parametric and parametric bootstrapping procedures. Commensurate with the Bayesian approach we assess the effect of alternative prior choices on posterior point estimates and show that these estimates may be made consistent with those obtained by bootstrapping approaches by the choice of an appropriate non-informative prior. In addition, we show that estimation uncertainty as described by 95% intervals is considerably more constrained and biologically realistic when estimated within the Bayesian framework. Finally, we provide posterior probability distributions on the per capita attack rate estimates that lend themselves to a more useful and descriptive characterization of interaction strength uncertainty.

Methods

Model framework

Novak and Wootton (2008) present a method for estimating a predator's per capita attack rates provided that a multi-species Holling type II functional response,

$$F_i = \frac{c_i N_i}{1 + \sum_{k=1}^{S} c_k h_k N_k}, \qquad \text{eqn 1}$$

describes the per predator feeding rate on the i^{th} prey species ($i = 1, ..., S$). Their estimator for the attack rates may be shown to be equivalent to

$$c_i = \frac{A_i}{A_0} \frac{1}{h_i N_i}, \qquad \text{eqn 2}$$

where h_i and N_i are the handling time and abundance of the i^{th} prey species, A_i is the number of predators feeding on prey i, and A_0 is the number of predators not feeding (see Appendix S1).

Estimates of A_i and A_0 come from one or more predator population "snapshot" feeding surveys in which the number of predator individuals feeding on each prey species is recorded. Other data sources are used to estimate prey abundances and handling times. For example, handling times are more easily measured in laboratory experiments than in the field, while abundance estimates may come from independently performed community surveys. Even if handling time data are based on field observations, they are unlikely to be measured on the predators observed in a feeding survey since the lengths of time those predators have been feeding are unknown due to the snapshot nature of the surveys.

Both eqn 1 and eqn 2 are implicitly deterministic mathematical models that include no statistical description of their stochastic component. In the next sections of this paper we develop a parameter-based version of eqn 2 that gives a stochastic formulation of the attack rate estimator that can incorporate sampling and environmental variation explicitly.

Bayesian Methods

Frequentist approaches for combining data from multiple sources to estimate functions of parameters generally rely on bootstrap methods or asymptotics like the multivariate delta-method. Both these approaches exhibit poor small-sample performance (Efron and Tibshirani, 1994; Kilian, 1998). This is relevant when dealing with predator feeding surveys as the A_i in eqn 2 are often very small for the rare prey species that typify predator diets (Rossberg et al., 2006). Small values of A_i can be problematic even when the total number of predators surveyed is large (Agresti and Coull, 1998). Ignoring variation in abundance and handling time estimates to focus on the variation within the feeding surveys may avoid this problem, but will lead to underestimation of the uncertainty in the attack rate estimates. The Bayesian framework circumvents this problem.

The Bayesian machinery is built around Bayes theorem:

$$f(\boldsymbol{\theta}|data) \propto f(data|\boldsymbol{\theta}) \cdot f(\boldsymbol{\theta}). \qquad \text{eqn 3}$$

Here, $f(data|\boldsymbol{\theta})$ is the likelihood: a function specifying the likelihood of the observed data in terms of unknown parameters $\boldsymbol{\theta}$. $f(\boldsymbol{\theta})$ is the prior: a probability density function reflecting prior beliefs or uncertainty about the parameters. Together, these inform $f(\boldsymbol{\theta}|data)$: the posterior distribution of the parameters given the data. Here, we consider only objective (also called non-informative) priors, assuming an absence of prior beliefs or information about the parameters in question (Berger, 2006). In other situations, informative priors constructed from previously obtained knowledge or data may be useful. Note that we use f to represent the density function of an arbitrary distribution for convenience only, using the function's argument(s) to indicate the specific distribution being referenced. For example, the density of a random variable X is indicated by $f(x)$, and that of θ by $f(\theta)$, even though these are not necessarily the same density functions. We use bold letters for vectors and bold uppercase letters for matrices of random variables.

A Bayesian Attack Rate Estimator

A parametric formulation of the attack rate estimator (eqn 2) is

$$\xi_i = \frac{\alpha_i}{\alpha_0}\frac{1}{\nu_i \eta_i}. \qquad \text{eqn 4}$$

Here, for the i^{th} prey species, ξ_i is the unknown attack rate, ν_i is the population prey abundance, η_i is the population handling time, α_i is the population proportion of predators feeding, and α_0 is the population proportion of predators that are not feeding on any prey species. In each case, the parameters refer to the broader (statistical) population, rather than sampled data only. By framing the attack rates this way, we are able to estimate them in the context of the broader population about which inference is desired.

If data on prey-specific feeding proportions ($\boldsymbol{F}$), abundances ($\boldsymbol{A}$), and handling times ($\boldsymbol{H}$) are collected independently, the joint likelihood of these distributions may be written as:

$$f(\boldsymbol{\alpha}, \boldsymbol{\nu}, \boldsymbol{\eta} | \boldsymbol{F}, \boldsymbol{A}, \boldsymbol{H}) = f(\boldsymbol{F}|\boldsymbol{\alpha}) f(\boldsymbol{A}|\boldsymbol{\nu}) f(\boldsymbol{H}|\boldsymbol{\eta}). \qquad \text{eqn 5}$$

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

Provided that the corresponding priors are also independent, Bayes theorem implies that

$$f(\boldsymbol{\alpha}|\boldsymbol{F}) \propto f(\boldsymbol{F}|\boldsymbol{\alpha})f(\boldsymbol{\alpha}), \qquad \text{eqn 6a}$$

$$f(\boldsymbol{\nu}|\boldsymbol{A}) \propto f(\boldsymbol{A}|\boldsymbol{\nu})f(\boldsymbol{\nu}), \qquad \text{eqn 6b}$$

$$f(\boldsymbol{\eta}|\boldsymbol{H}) \propto f(\boldsymbol{H}|\boldsymbol{\eta})f(\boldsymbol{\eta}). \qquad \text{eqn 6c}$$

These may therefore be fit with independent models for each component. That is, the posterior distributions of the attack rates in eqn 4 may be estimated using Markov Chain Monte Carlo (MCMC) to obtain samples from each of the three posterior distributions in eqn 6 and combining these using eqn 4. If the three types of data are not gathered independently, then it is necessary to consider likelihood or prior models that account for this dependence (see Appendix S3).

Case study data set

To provide a concrete explanation of the additional details of the Bayesian approach we applied it to a dataset involving the predatory marine intertidal whelk *Haustrum scobina*. *Haustrum* feeds primarily on barnacles and mussels, often by first drilling through the shells of its prey. Handling times, which can be hours to days, are the times needed to drill and ingest a prey individual. The dataset contains information from replicate feeding surveys, quadrat-based abundance surveys, and laboratory-based handling time experiments, which we describe briefly below. Further details may be found in Novak (2010, 2013).

Fifteen feeding surveys were conducted during low tides over two years. In each survey, the number of whelks feeding on each prey species was recorded, as was as the number not feeding (Table 1). The sizes of the predator individuals (both feeding and not

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

feeding) and of the prey being fed upon were also recorded (±1 mm), along with the average temperature of the month in which each survey was conducted. These three covariates contribute to the deterministic variation in per capita attack rate estimates.

Species abundance surveys used 10-15 replicate quadrats randomly distributed along 3 transects, each repeated 3 times over the same time periods as the feeding surveys. As is typical of community abundance surveys, numerous zeros exist in these data as many species did not occur in every quadrat (Table 1). The presence of such zeroes reflects both deterministic variation associated with real variation in species abundances, as well as stochastic variation associated with sampling effort.

Handling times were estimated in laboratory experiments that manipulated predator and prey sizes and temperature. Replicated individuals housed in separate aquaria with different prey species were checked hourly to determine handling time durations. As a result, handling time measurements are interval censored, equally so for prey species with short (hour-long) and long (multi-day) handling times. Such uncertainty, along with variation in the number of replicate experiments that were performed for each of the prey species (Table 1), reflect additional sources of stochastic uncertainty.

Treating the abundances, handling times, and feeding surveys data as independent, we now specify appropriate likelihood and prior models.

Model formulation

Accommodating variation in feeding surveys – Letting P be the total number of predators surveyed, X_i the number observed feeding on prey i, and X_0 the number not feeding, we modeled the combined feeding survey data using a multinomial likelihood with Dirichlet prior:

$$(X_0, X_1, ..., X_S) \sim Mult_P(\alpha_0, \alpha_1, ..., \alpha_S) \qquad \text{eqn 7a}$$

$$(\alpha_0, \alpha_1, ..., \alpha_S) \sim Dirich(c, c, ..., c). \qquad \text{eqn 7b}$$

The resulting posterior distribution is also Dirichlet:

$$(\alpha_0, \alpha_1, ..., \alpha_S)|\boldsymbol{x} \sim Dirich(c + x_0, c + x_1, ..., c + x_S). \qquad \text{eqn 8}$$

We will focus on the posterior medians rather than means as our point estimates of interest since they are the more appropriate measure of a skewed distribution's central tendency (Gelman et al., 2013). The four most commonly used non-informative priors in this setting are Laplace's prior ($c = 1$), Jeffreys' prior ($c = \frac{1}{2}$), Perks' prior ($c = \frac{1}{S+1}$), and Haldane's prior ($c = 0$) (Hutter, 2013). However, these priors all result in posterior medians that may differ substantially from the sample proportions, which are the maximum likelihood estimates (MLEs), especially when any of the X_i are small. This leads to counter-intuitive attack rate point estimates for rarely observed prey species (Fig. 1).

Kerman et al. (2011) showed that when $c = \frac{1}{3}$, the multinomial parameter posterior medians closely match the MLEs, referring to this prior as the non-informative neutral prior. We show that this result applies to the ratios of multinomial parameters as well by letting $\gamma_i = \frac{\alpha_i}{\alpha_0}$ and noting that the posterior distribution of γ_i is the ratio of Dirichlet components, which is the ratio of independent gamma random variables. This may be written as:

$$f\left(\gamma_i | x_i, x_0\right) = \frac{x_0 + c}{x_i + c} \cdot g\left(\frac{x_0 + c}{x_i + c} \cdot \gamma_i; 2(x_i + c), 2(x_0 + c)\right) \qquad \text{eqn 9}$$

where $g(y; d_1, d_2)$ is an F-distribution probability density function with d_1 and d_2 degrees of freedom. Using the approximation for the median of an F-distribution $med(F_n^m) \approx \frac{n}{3n-2}\frac{3m-2}{m}$ (see Appendix S2) and setting it equal to the MLE of $\frac{\alpha_i}{\alpha_0}$, $\frac{x_i}{x_0}$, yields the solution $c = \frac{1}{3}$. Figure 1, which shows the log differences in posterior median relative to the MLE for several values of c, evidences that $c = \frac{1}{3}$ is indeed a reasonable prior to use.

Accommodating variation in abundance surveys – We used a zero-inflated gamma (ZIG) model to account for the numerous zeros in the abundances data. Letting $Y_1, ..., Y_n$ denote the abundance measurements, and by conditioning on whether or not a zero occurs, the likelihood density of the ZIG distribution can be written as

$$g(y; \alpha, \beta, \rho) = \rho^{I[y=0]}[(1-\rho)f(y, \alpha, \beta)]^{I[y>0]}, y \geq 0, \qquad \text{eqn 10}$$

where ρ is the probability of a zero, $f(y; \alpha, \beta)$ is the usual gamma density with shape α, rate β, and mean $\frac{\alpha}{\beta}$, and $I[\cdot]$ is the indicator function that equals 1 when its argument is true and 0 otherwise (Ospina and Ferrari, 2012). The ZIG density is separable in ρ and (α, β). It follows that the zero-inflation parameter can be treated separately, provided a separable prior is used. Thus, for each prey species, we modeled the number of observed zeros using a binomial distribution with a uniform prior on ρ and took $log(\alpha) \sim Unif(-100, 100)$ and $log(\beta) \sim Unif(-100, 100)$ to approximate the independent scale-invariant non-informative prior $f(\alpha, \beta) = f(\alpha)f(\beta) \propto \frac{1}{\alpha}\frac{1}{\beta}$ (Syversveen, 1998).

Accommodating variation in handling time experiments – We used regression to model the relationships between handling times and the predator-size, prey-size and temperature covariates of the laboratory experiments. Average handling times for use in the attack rate estimation were obtained by combining these regression coefficients with the same covariate information obtained during feeding surveys.

Specifically, we considered the i^{th} handling time observation for a given prey species to be associated with a covariates vector $\boldsymbol{X_i}$ consisting of 1 followed by temperature, predator size, and prey size (all log transformed). We then modeled the likelihood of the i^{th} handling time with a modified-Normal likelihood written as

$$H_i \sim N_{l_i}(e^{\boldsymbol{X_i^T\beta}}, \sigma^2) \qquad \text{eqn 11}$$

where the subscript l_i refers to the censoring "window" length and indicates that a $Unif(\frac{l_i}{2}, \frac{l_i}{2})$ error was added to the normal distribution corresponding to the interval censoring with which handling times were observed. The exponential link of eqn 11 avoids negative mean handling time estimates.

Treating the field covariates (predator size, prey size, and temperature) as random to account for sampling variability, we modeled the distributions of the (log-transformed) covariate observations $\boldsymbol{X_1}, ..., \boldsymbol{X_N}$, where N is the total number of field observations, as being independent and identically distributed and drawn from a multivariate normal distribution, $N(\boldsymbol{\mu}, \Sigma')$, with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ' . Non-informative multivariate normal and inverse Wishart priors were used for $\boldsymbol{\mu}$ and Σ' respectively (Fink, 1997). Letting $\boldsymbol{X^*}$ follow the posterior predictive distribution (our estimate of the distribution of the covariates), the mean handling time may be written as

$$E(H) = E[E(H|\boldsymbol{X^*})] = E(e^{\boldsymbol{\beta^T X^*}}). \qquad \text{eqn 12}$$

We estimated this expectation by sampling from the regression parameters' posterior distribution, sampling new covariates from their posterior predictive distribution, computing $e^{\boldsymbol{\beta}^T \boldsymbol{X}^*}$ for each sample, and averaging across all samples. The weak law of large numbers ensures convergence to $E(e^{\boldsymbol{\beta}^T \boldsymbol{X}^*})$ as sample size increases (Petrov, 1995).

Model implementation: Putting the pieces together to estimate per capita attack rates

Using the likelihoods and priors of the feeding surveys, abundances and handling times described above, we drew samples from the parameters' posterior distributions using Markov Chain Monte Carlo (MCMC). MCMC sampling was done using JAGS with the R package 'rjags' (Plummer and Stukalov, 2014). Parameter samples were then combined using eqn 4 to produce samples from each prey's attack rate posterior distribution. This use of eqn 4 treated handling times, $\boldsymbol{H}$, as being independent of the predator feeding surveys, $\boldsymbol{F}$, even though covariate observations of predator size, prey size and temperature from the feeding surveys informing $\boldsymbol{F}$ were used to inform $\boldsymbol{H}$ by combining them with the laboratory-based handling time regression coefficients associated with these covariates. We established the validity of this assumption by examining the relationship between feeding proportions and covariate averages between the individual surveys (see Appendix S3).

We verified Markov chain convergence using trace plots, removed pre-convergence samples, and thinned each chain to ensure independence among the remaining samples. Burn-in times and thinning values were selected separately for feeding survey, abundance, and handling time models based on trace plots and autocorrelation function plots. Inferences were based on 1,000 samples after confirming that independent sets of 1,000 samples led to the same conclusions.

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

Results

Comparison of Bayesian and bootstrapping procedures

To assess the utility and performance of our Bayesian approach we contrasted point and 95% interval estimates from (i) the model with Laplace's prior ($c = 1$) on the Dirichlet feeding proportions, (ii) the model with Haldane's prior ($c = 0$), and (iii) the model with the neutral prior ($c = \frac{1}{3}$) to estimates obtained by (iv) non-parametric and (v) parametric bootstrapping procedures. In contrast to the Bayesian 95% credible intervals, which reflect the range of values within which a parameter will occur with 95% probability, the 95% confidence intervals associated with bootstrapping do not have a direct probability-based interpretation. Rather, if we repeatedly constructed 95% confidence intervals on 'new' datasets, about 95% of them would contain the 'true' value of the parameter.

Non-parametric bootstrapping was performed by sampling with replacement from each of the feeding survey, abundance, and handling-time datasets until the same number of samples had been drawn as was present in each dataset (Efron and Tibshirani, 1994). Per capita attack rates were calculated for many sets of such resampled data to estimate the mean and 95% confidence intervals of their bootstrapped distributions.

The parametric bootstrap was implemented using the likelihood functions of the Bayes method. That is, we used the data to estimate the parameters of the three likelihood functions (i.e. eqn 7, eqn 10, eqn 11) by maximum likelihood, used these fit likelihood functions to simulate new datasets, combining samples from the three distributions to calculate per capita attack rates and estimate medians and 95% confidence intervals of the resulting bootstrapped attack rate distributions.

The comparison of the approaches indicates that the model with the neutral prior ($c = \frac{1}{3}$) on the feeding proportions was indeed both sufficient for, and performed best in, describing the variation inherent in the estimated rates with which *Haustrum scobina* attacked its eight prey species (Fig. 2). It exhibited median point estimates most closely matching the point estimates of the two bootstrapping approaches. The two bootstrap distributions however, frequently exhibited lower 95% confidence interval end points of zero; a nonsensical result given that the consumption of these species was in fact observed. As expected (Fig. 1), the models having Laplace's ($c = 1$) or Haldane's ($c = 0$) prior resulted in inflated and depressed attack rate point estimates respectively, particularly on prey species that were observed infrequently in the feeding surveys.

Figure 3 shows the posterior probability distributions of *Haustrum scobina*'s per capita attack rates on each of its prey species as estimated using the neutral prior ($c = \frac{1}{3}$). The distributions are roughly symmetric on the logarithmic scale, indicating right skew and justifying the use of the median as the point estimate of their central tendency.

Discussion

Effort devoted to estimating the strengths of species interactions has centered on obtaining point estimates, leaving the characterization of estimation uncertainty largely unconsidered. This shortcoming reflects not only the logistical difficulty of quantifying interaction strengths in nature's species-rich communities, but is also a consequence of the still nascent integration of the mathematical and statistical methods available to food web ecologists.

The fitting of deterministic mathematical models to data requires that they be formulated as stochastic statistical models whose constants – like the per capita attack rates considered here – be treated as unknown parameters to be estimated. For the observational estimator of Novak and Wootton (2008) the unknown attack rate parameters of interest are functions of other unknown parameters that must themselves be estimated.

Uncertainty in attack rate estimates thus reflects the contributions of both the deterministic and stochastic variation of these component parameters. The propagation of both such forms of variation is inherent to all other experimental and observational approaches as well.

Advantages of the Bayesian approach

Unlike frequentist methods, Bayesian methods offer a relatively straightforward way to estimate parameters that are functions of other parameters using multiple sources of information. This is particularly, though not necessarily, so when the posterior distributions of these parameters are independent. Bayesian methods also permit a more natural interpretation of the uncertainty that accompanies parameter estimates and provide a complete characterization of this uncertainty in the form of posterior probability distributions; frequentist methods provide the moments and intervals of distributions whose interpretation is arguably less intuitive (Clark, 2005).

In the context of reticulate food webs and predator-prey interactions, the complete probabilistic characterization of uncertainty regarding observational interaction strength estimates opens the door for probabilistic predictions of species effects and population dynamics (Calder et al., 2003; Yeakel et al., 2011). This stands in contrast to the typical use of arbitrarily chosen interaction strength ranges in stochastic simulations and numerical sensitivity analyses (Yodzis, 1988; Novak et al., 2011). An alternative choice to use bootstrapped (frequentist) confidence intervals to inform predictions could lead to additional problems when lower interval bounds extend to zero for prey species that are rarely found in a predator's diet. First, draws of zeros would amount to the outright removal of the predator-prey interaction and could lead to biased predictions through the underestimation of food web complexity. Second, as evidenced by *Haustrum scobina*'s feeding on *Mytilus galloprovincialis* (Fig. 2), prey species whose attack rate confidence intervals extend to zero may in fact experience very high per capita attack rates on average. Treating these interactions as potentially absent would fail to identify strong interactions that are rarely observed only because of strong top-down control of the prey populations' sizes, for example. Such issues do not occur in the Bayesian framework where the Dirichlet prior distribution is conjugate for the multinomial likelihood, thereby producing a Dirichlet posterior from which MCMC samples of zero cannot occur.

Considerations and implications

Bayesian methods offer a powerful tool, but they should not be applied without careful consideration of the prior distribution. The choice of objective ('non-informative') prior is particularly important when data are sparse (Van Dongen, 2006; Boshuizen and Van Baal, 2009). It follows that, for rarely observed prey species, different prior specifications lead to different point estimates of the per capita attack rates (Figs. 1 and 2). That is, while priors with concentration parameters $c > \frac{1}{3}$ (e.g., Laplace's prior) will produce higher attack rate point estimates the less frequently a prey species is observed in the predator's diet, priors with concentration parameters $c < \frac{1}{3}$ (e.g., Haldane's prior) will produce lower attack rate point estimates the less frequently a prey species is observed in the predator's diet (see also Fig. S2). The biological implication of choosing to use one such prior over another is that this choice can alter the relative frequency of weak and strong interactions. Thus, the choice of priors can alter inferences of population dynamics and food web stability when models are parameterized with empirical estimates (Allesina and Tang, 2012). These considerations are avoided only when all prey occur frequently in a predator's diet (see *Mytilus galloprovincialis* and *Chaemaesipho columna* in Fig. 2). In such cases, the large sample sizes mean that the likelihood overwhelms the prior regardless of its information content such that Bayesian and frequentist estimates are similar.

The use of the neutral prior produces posterior distribution median point estimates that are least influenced by the prior and thus most like the point estimates of the frequentist bootstrap methods (Figs. 1, 2, S2). We therefore suggest that this be the preferred objective prior to use. Tuyl et al. (2008) argue against the use of such sparse ($c < 1$) priors for binomial parameters as they put more weight on extreme outcomes. For example, if $Y \sim Bin(n, p)$ and $Y \in \{0, n\}$, the use of sparse priors leads to inappropriately narrow credible intervals. Fortunately, this problem is avoided in our application because all considered prey species (and "not feeding") are observed at least once (i.e. $Y \in \{1, ..., n-1\}$). In hierarchical models, to which our Bayesian framework could be naturally extended (Cressie et al., 2009), $Y \in \{0, n\}$ is more likely for any individual survey, but this is not an issue as inference at the survey level is not desired.

An influence of Bayesian prior choice also occurs in the estimation of prey abundances by means of a zero-inflated gamma likelihood model. Here the assumption that a zero-inflated gamma is descriptive of the abundance structure of all prey species can lead to the inflation of per capita attack rate estimates for species that are ubiquitous. When species occur in all but a few sampled quadrats, relatively little data are available to estimate the probability of obtaining a count of zero. In such situations the influence of even an uninformative uniform prior will be increased, resulting in an inflated estimate of the proportion of zeros and thus a reduced estimate of a species' abundance. Attack rate estimates are thereby inflated because a species' abundance occurs in the denominator of the estimator (eqn 2). For our dataset, where many species were present in all sampled quadrats (Table 1), this inflation effect appears to have been weak as seen by comparing the results of the Bayesian models to the frequentist bootstrapping procedures for which such inflation does not occur (Fig. 2); the probability of obtaining a value of zero during bootstrapping is equal to the sample proportion of zeros in the data, which is zero for species that are always observed. Arguably, however, this inflation effect of the prior that is inherent to the use of the zero-inflated gamma in a Bayesian framework is appropriate because observations of species absences at the spatial scale of quadrats are fundamentally different from observations of species presences when no prior knowledge about the patchiness of species' abundances is available.

Issues of prior choice aside, Bayesian methods offer a more complete characterization of the estimated uncertainty of parameter estimates in the form of posterior probability distributions. Several metrics may be chosen to summarize the shapes of these distributions. For example, means, medians and modes are all commonly used as point estimates to reflect a distribution's typical and most likely value. For strongly skewed distributions – such as those observed here (Fig. 3) – medians are a more representative metric of a distribution's central tendency. Furthermore, a distribution's median, unlike its mean, will always fall within the equal-tailed interval that is typically used to describe the uncertainty or variation surrounding the distribution's estimated central tendency. Of course, point estimates provide little information on a distribution's shape. Confidence or credible intervals provide more such information with which to characterize parameter uncertainty and variation. The typical metrics for these intervals are equal-tailed, but for posterior distributions the highest posterior density (HPD) interval may also be useful (Gelman et al., 2013). While intervals characterized by highest posterior density are more resistant to distribution skewness and will always include the distribution's mode, equal-tailed intervals are invariant under monotone transformations, making them easier to interpret after log-transformation. Log-transformation is frequently necessary in the context of interaction strengths given the wide range of values that the community-wide strengths of species interactions typically exhibit (Wootton and Emmerson, 2005). Ultimately, the entire joint posterior distribution should be presented whenever possible (Chen and Shao, 1999). When this is not practical, the choice of posterior summaries will depend on the goal of the analysis.

Conclusion

While many ecological processes can be described in purely mathematical terms, mathematical models are often most useful when they are linked with real data (Codling and Dumbrell, 2012). Linking models with data is necessary to validate and compare models, and to parameterize them for real-world use in predicting future system dynamics (Bolker, 2008). This has been a challenging task in the study of species rich food webs, not least because of the difficulty of parameter estimation in typical food web models and challenges with integrating data collected across multiple spatial and temporal scales. Statistical models of predator-prey interactions that consider both deterministic and stochastic variation in data are needed to accompany the numerous mathematical models that have been proposed. Our work represents a step in this direction.

Acknowledgments

This work was supported by NSF DEB-1353827 and DEB-0608178. Data and R scripts are available on the Dryad Digital Repository (`doi:10.5061/dryad.6k144`) and on GitHub (`https://github.com/wolfch2/Bayesian-Interaction-Strength`)

References

Agresti, A., and B. A. Coull. 1998. Approximate is better than exact for interval estimation of binomial proportions. The American Statistician **52**:119–126.

Allesina, S., and S. Tang. 2012. Stability criteria for complex ecosystems. Nature **483**:205–208.

Berger, J. 2006. The Case for Objective Bayesian Analysis. Bayesian Analysis **1**:385–402.

Bolker, B. M. 2008. Ecological Models and Data in R. 508 edition edition. Princeton University Press, Princeton, N.J.

Boshuizen, H. C., and P. H. Van Baal. 2009. Probabilistic sensitivity analysis: be a Bayesian. Value in Health **12**:1210–1214.

Calder, C., M. Lavine, P. Muller, and J. S. Clark. 2003. Incorporating Multiple Sources of Stochasticity into Dynamic Population Models. Ecology **84**:1395–1402.

Chen, M.-H., and Q.-M. Shao. 1999. Monte Carlo estimation of Bayesian credible and HPD intervals. Journal of Computational and Graphical Statistics **8**:69–92.

Chesson, J. 1983. The Estimation and Analysis of Preference and Its Relatioship to Foraging Models. Ecology **64**:1297.

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

Codling, E. A., and A. J. Dumbrell. 2012. Mathematical and theoretical ecology: linking models with ecological processes. Interface Focus **2**:144–149.

Cressie, N., C. A. Calder, J. S. Clark, J. M. V. Hoef, and C. K. Wikle. 2009. Accounting for uncertainty in ecological analysis: the strengths and limitations of hierarchical statistical modeling. Ecological Applications **19**:553–570.

Efron, B., and R. J. Tibshirani. 1994. An Introduction to the Bootstrap. CRC Press.

Fink, D. 1997. A compendium of conjugate priors. Technical Report page 46.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. 2013. Bayesian Data Analysis, Third Edition. 3 edition edition. Chapman and Hall/CRC, Boca Raton.

Hutter, M. 2013. Sparse Adaptive Dirichlet-Multinomial-like Processes. arXiv preprint arXiv:1305.3671 .

Jeschke, J. M., M. Kopp, and R. Tollrian. 2004. Consumer-food systems: why type I functional responses are exclusive to filter feeders. Biological Reviews **79**:337–349.

Kerman, J. 2011. A closed-form approximation for the median of the beta distribution. arXiv preprint arXiv:1111.0433 .

Kerman, J., et al. 2011. Neutral noninformative and informative conjugate beta and gamma prior distributions. Electronic Journal of Statistics **5**:1450–1470.

Kilian, L. 1998. Small-Sample Confidence Intervals For Impulse Response Functions. The Review of Economics and Statistics **80**:218–230.

Laska, M. S., and J. T. Wootton. 1998. Theoretical concepts and empirical approaches to measuring interaction strength. Ecology **79**:461–476.

Novak, M. 2010. Estimating interaction strengths in nature: experimental support for an observational approach. Ecology **91**:2394–2405.

Novak, M. 2013. Trophic omnivory across a productivity gradient: intraguild predation theory and the structure and strength of species interactions. Proceedings of the Royal Society B: Biological Sciences **280**:20131415.

Novak, M., and J. T. Wootton. 2008. Estimating nonlinear interaction strengths: an observation-based method for species-rich food webs. Ecology **89**:2083–2089.

Novak, M., J. T. Wootton, D. F. Doak, M. Emmerson, J. A. Estes, and M. T. Tinker. 2011. Predicting community responses to perturbations in the face of imperfect knowledge and network complexity. Ecology **92**:836–846.

Ospina, R., and S. L. Ferrari. 2012. A general class of zero-or-one inflated beta regression models. Computational Statistics & Data Analysis **56**:1609–1623.

Paine, R. T. 1992. Food-web analysis through field measurement of per capita interaction strength. Nature **355**:73–75.

Petrov, V. V. 1995. Limit Theorems of Probability Theory: Sequences of Independent Random Variables. Oxford University Press, Oxford : New York.

Plummer, M., and A. Stukalov, 2014. rjags: Bayesian graphical models using MCMC.

Power, M. E., D. Tilman, J. A. Estes, B. A. Menge, W. J. Bond, L. S. Mills, G. Daily, J. C. Castilla, J. Lubchenco, and R. T. Paine. 1996. Challenges in the quest for keystones. BioScience **46**:609–620.

Rall, B. C., U. Brose, M. Hartvig, G. Kalinkat, F. Schwarzmller, O. Vucic-Pestic, and O. L. Petchey. 2012. Universal temperature and body-mass scaling of feeding rates. Philosophical Transactions of the Royal Society B: Biological Sciences **367**:2923–2934.

Rossberg, A. G., K. Yanagi, T. Amemiya, and K. Itoh. 2006. Estimating trophic link density from quantitative but incomplete diet data. Journal of Theoretical Biology **243**:261–272.

Syversveen, A. R. 1998. Noninformative bayesian priors. interpretation and problems with construction and applications .

Tuyl, F., R. Gerlach, and K. Mengersen. 2008. A Comparison of Bayes-Laplace, Jeffreys, and Other Priors: The Case of Zero Events. The American Statistician **62**:40–44.

Van Dongen, S. 2006. Prior specification in Bayesian statistics: Three cautionary tales. Journal of Theoretical Biology **242**:90–100.

Vucetich, J. A., R. O. Peterson, and C. L. Schaefer. 2002. The effect of prey and predator densities on wolf predation. Ecology **83**:3003–3013.

Wootton, J. T. 1994. Predicting Direct and Indirect Effects: An Integrated Approach Using Experiments and Path Analysis. Ecology **75**:151–165.

Wootton, J. T., and M. Emmerson. 2005. Measurement of Interaction Strength in Nature. Annual Review of Ecology, Evolution, and Systematics **36**:419–444.

Yeakel, J. D., M. Novak, J. Guimares, Paulo R., N. J. Dominy, P. L. Koch, E. J. Ward, J. W. Moore, and B. X. Semmens. 2011. Merging Resource Availability with Isotope Mixing Models: The Role of Neutral Interaction Assumptions. PLoS ONE **6**:e22015.

Yodzis, P. 1988. The Indeterminacy of Ecological Interactions as Perceived Through Perturbation Experiments. Ecology **69**:508–515.

527 1. Tables

Prey Species	Feeding Observations	Abundance Zeros	Handling-time Experiments
Chamaesipho columna	265	0	6
Xenostrobus pulex	185	0	52
Austrolittorina antipodum	3	2	0
Austrolittorina cincta	2	0	46
Epopella plicata	2	0	1
Mytilus galloprovincialis	1	24	15
Notoacmea Radialspokes	1	5	66
Risellopsis varia	1	3	68
Not Feeding	1,629		
Total Surveyed	2,089		

Table 1: Summary statistics for the datasets used to estimate the per capita attack rates with which the intertidal predator *Haustrum scobina* feeds on its eight prey species. "Feeding observations" indicates the total frequency with which predator individuals were observed to be feeding on each prey species across all feeding surveys. "Abundance zeros" indicates the number of zeros recorded in the 30 quadrat-based community surveys of species abundances. "Handling-time experiments" indicates the number of laboratory experiments that were used to estimate handling time regression coefficients for each prey species.

2. Figure Legends

Figure 1: Comparison of alternative non-informative priors in estimating the ratio of the proportions of feeding versus not feeding predator individuals. The x-axis reflects the number of predators observed in the process of feeding on a given prey species, with a total of 1,629 individuals assumed to have been not feeding, corresponding to the number not feeding in our dataset (Table 1). The y-axis shows the difference in logarithms of the posterior median using a $Dirich(c, ..., c)$ prior and the maximum likelihood estimate of the ratio. From top to bottom in the graph, the values of c are 1 (Laplace), $\frac{1}{2}$ (Jeffreys'), $\frac{1}{3}$ (neutral), $\frac{1}{S+1} = \frac{1}{9}$ (Perks'), and 0 (Haldane's). The neutral prior ($c = \frac{1}{3}$) leads to estimates that closely match the maximum likelihood estimates.

Figure 2: Comparison of the frequentist and Bayesian approaches to estimating the per capita attack rates with which *Haustrum scobina* consumed its 8 prey species. Variation in attack rate estimates is illustrated for each procedure by the medians and 95% equal-tailed intervals of their distributions. Procedures are organized the same for each prey species as, from top to bottom: (i) non-parametric bootstrap, (ii) parametric bootstrap, (iii-v) Bayesian procedure with sparsity parameters 0 (Haldane's prior), $\frac{1}{3}$ (neutral prior), and 1 (Laplace's prior) respectively. Unlike the 95% confidence intervals for the bootstrap procedures which often span zero ($= 10^{-7}$ for graphical convenience), the 95% posterior posterior intervals of the Bayesian method indicate the regions where attack rates lie with 95% probability.

Figure 3: Posterior distributions for *Haustrum scobina's* per capita attack rates ($prey \cdot predator^{-1} \cdot prey^{-1} \cdot m^{-2} \cdot day^{-1}$) and its components ($\xi_i = \frac{\alpha_i}{\alpha_0} \cdot \frac{1}{\nu_i} \cdot \frac{1}{\eta_i}$) using neutral ($c = \frac{1}{3}$) Dirichlet prior on feeding proportions.

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

551 3. Figures

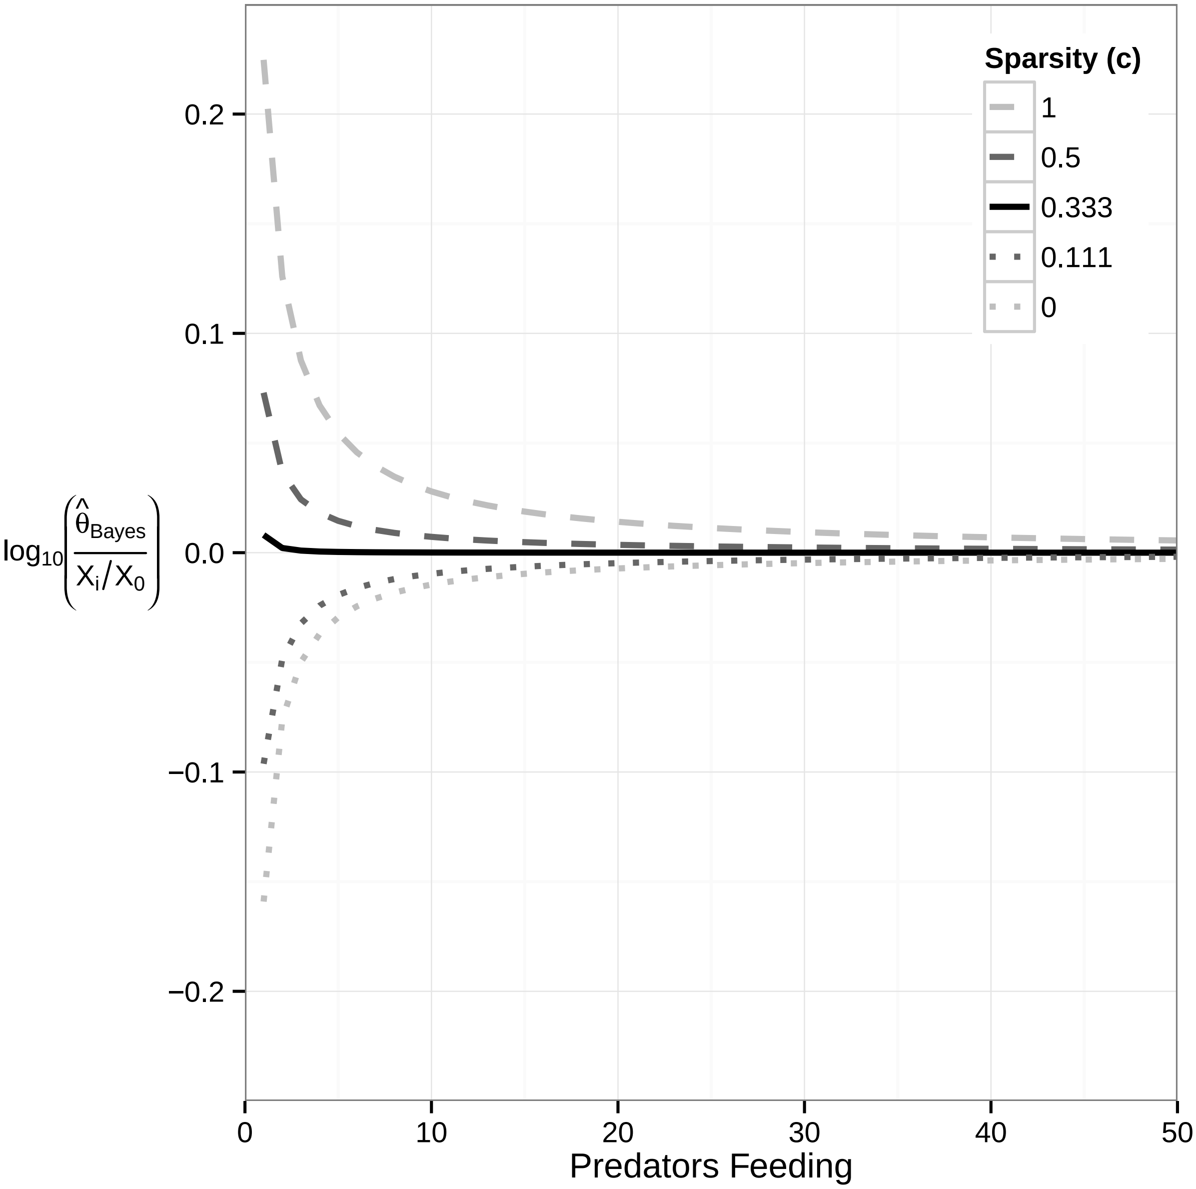

Figure 1

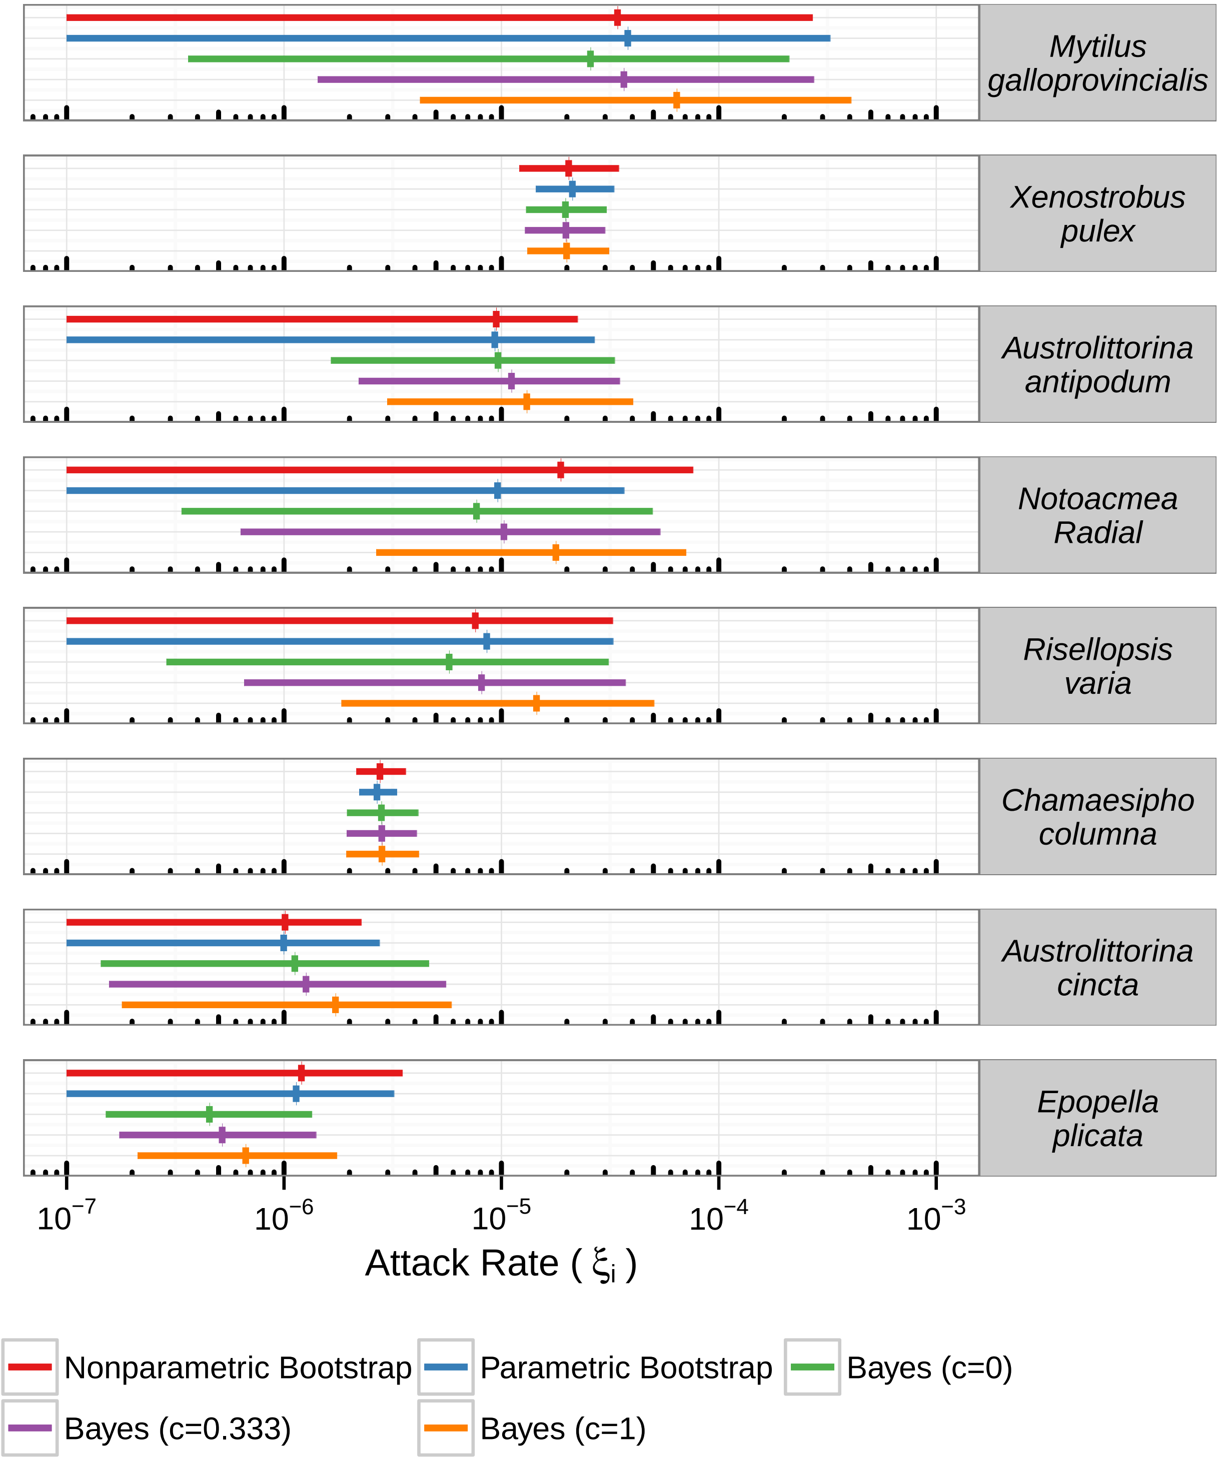
Mytilus galloprovincialis
Xenostrobus pulex
Austrolittorina antipodum
Notoacmea Radial
Risellopsis varia
Chamaesipho columna
Austrolittorina cincta
Epopella plicata
10^{-7}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
Attack Rate (ξ_i)
Nonparametric Bootstrap
Parametric Bootstrap
Bayes (c=0)
Bayes (c=0.333)
Bayes (c=1)

Figure 2

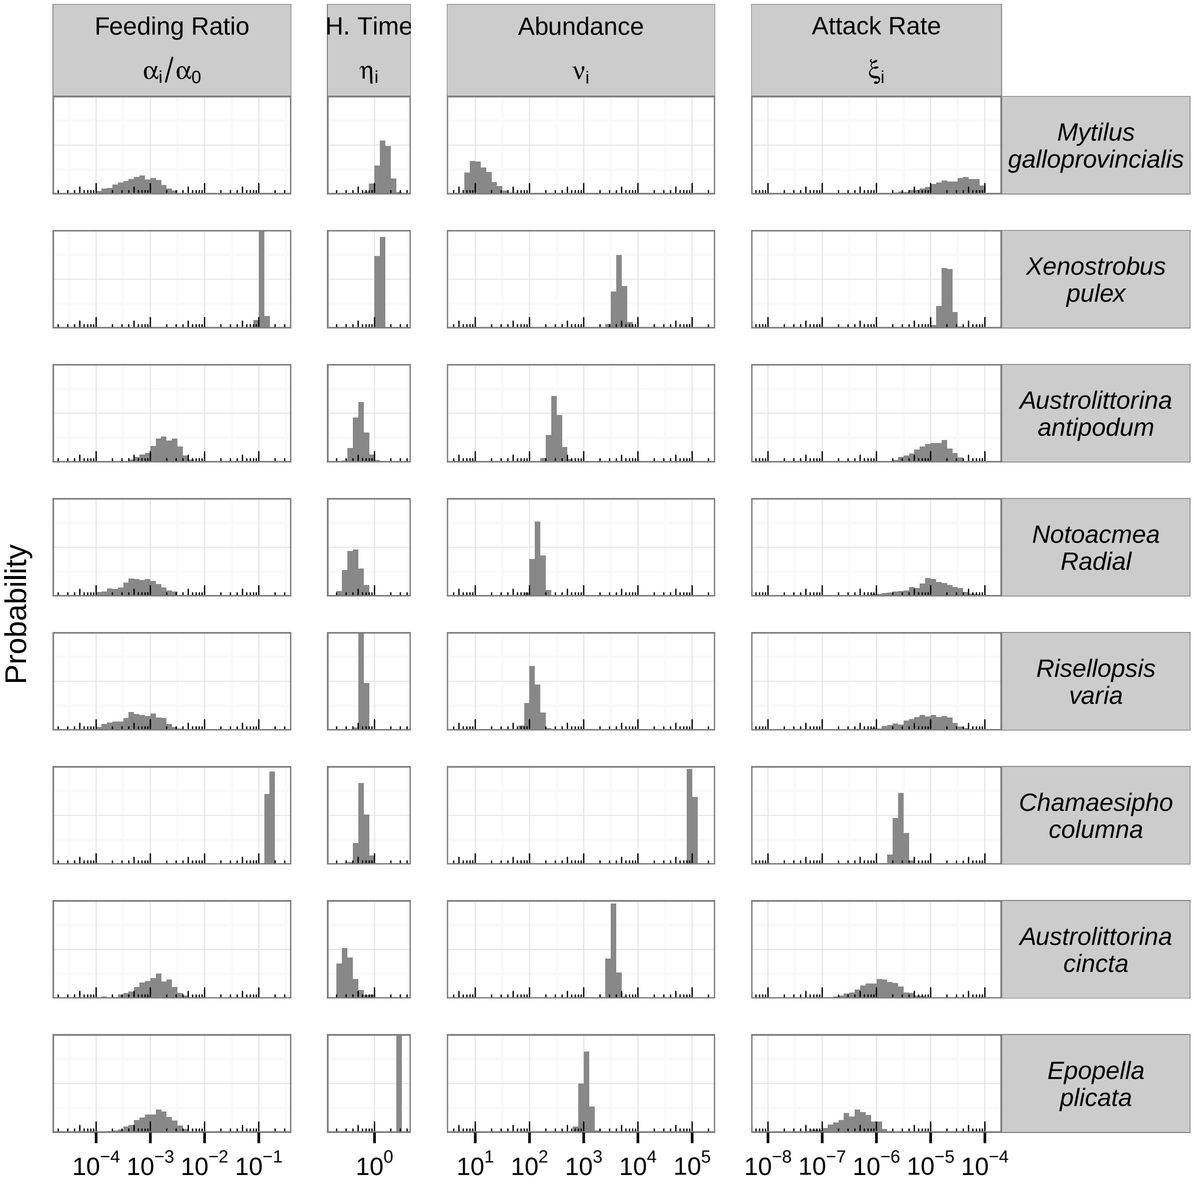
Feeding Ratio
α_i/α_0
H. Time
η_i
Abundance
ν_i
Attack Rate
ξ_i
Probability
Mytilus galloprovincialis
Xenostrobus pulex
Austrolittorina antipodum
Notoacmea Radial
Risellopsis varia
Chamaesipho columna
Austrolittorina cincta
Epopella plicata
10^{-4} 10^{-3} 10^{-2} 10^{-1}
10^{0}
10^{1} 10^{2} 10^{3} 10^{4} 10^{5}
10^{-8} 10^{-7} 10^{-6} 10^{-5} 10^{-4}

Figure 3

Appendix

S1 *On Novak and Wootton's 'Species x'*

This paper expands on the observational method for estimating attack rates presented by Novak and Wootton (2008):

$$c_i = \frac{F_i A_x}{(F_x - A_x) h_i N_i},$$

where c_i is the attack rate, h_i is the handling time, and N_i is the abundance, all for the i^{th} prey species. A_i and F_i are the proportions of all predators and feeding predator respectively feeding on the i^{th} prey species. x refers to an arbitrarily chosen prey species that is the same for all c_i. Here we show that this equation can also be written in a more simplified form, showing that the estimates are not dependent on the choice of species x.

Define A_0 to be the observed proportion of predators that are not feeding, so that $A_0 = 1 - \sum_{i=1}^{S} A_i$. Then, the $F_i's$ can be obtained by normalizing $A_i's$: $F_i = \frac{A_i}{\sum_{j=1}^{S} A_j} = \frac{A_i}{1-A_0}$. Noting that:

$$F_x - A_x = \frac{A_x}{1 - A_0} - A_x = \frac{A_x - A_x(1 - A_0)}{1 - A_0} = \frac{A_x[1 - (1 - A_0)]}{1 - A_0} = \frac{A_x A_0}{1 - A_0}.$$

It follows that

$$\frac{F_i A_x}{F_x - A_x} = \frac{\frac{A_i}{1-A_0} \cdot A_x}{\frac{A_x A_0}{1-A_0}} = \frac{A_i A_x}{A_x A_0} = \frac{A_i}{A_0}.$$

This can be further simplified by noting that the $A_i's$ have a common denominator (total number surveyed). This means that the original attack rate equation can be written as

$$c_i = \frac{A_i}{A_0} \cdot \frac{1}{h_i N_i} = \frac{\#\ feeding\ on\ i}{\#\ not\ feeding} \cdot \frac{1}{h_i N_i}.$$

This shows that the estimate does not involve species x. Moreover, the total number surveyed need not be known to estimate a subset of the attack rates.

S2 *F-distribution median*

In general, the median of the F-distribution does not have a closed form. However, we can derive an approximation by relating the F-distribution to the beta-distribution.

Let $X \sim F_n^m$. We can express X as the ratio of scaled, independent Chi-squared distributions $C_m \sim \chi_m^2$ and $C_n \sim \chi_n^2$:

$$X = \frac{C_m/m}{C_n/n}$$

It follows that we can express X as the ratio of scaled independent gamma distributions $G_m \sim gamma(\frac{m}{2}, 2)$ and $G_n \sim gamma(\frac{n}{2}, 2)$:

$$\begin{aligned} X &= \frac{G_m/m}{G_n/n} \\ &= \frac{n}{m}\frac{G_m}{G_n} \end{aligned}$$

We can then normalize the gamma distributions:

$$X = \frac{n}{m}\frac{\frac{G_m}{G_m+G_n}}{\frac{G_n}{G_m+G_n}}$$

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015

Letting $(D_1, D_2) \sim Dir(\frac{m}{2}, \frac{n}{2})$ and using the relationship between Dirichlet and gamma distributions,

$$X = \frac{n}{m}\frac{D_1}{D_2}$$

Using the marginal distribution for Dirichlet components result and the fact that $D_1 + D_2 = 1$, we have that

$$X = \frac{n}{m}\frac{B}{1-B}$$

where $B \sim Beta(\frac{m}{2}, \frac{n}{2})$. Note that this is a monotone transformation of B, so it preserves the median. When $m > 2$ and $n > 2$, the median of B is approximately $\frac{\frac{m}{2}-\frac{1}{3}}{\frac{m}{2}+\frac{n}{2}-\frac{2}{3}}$ (Kerman, 2011). Substituting this result, we have that

$$\begin{aligned} med(X) &= \frac{n}{m}\frac{med(B)}{1-med(B)} \\ &= \frac{n}{m}\frac{\frac{\frac{m}{2}-\frac{1}{3}}{\frac{m}{2}+\frac{n}{2}-\frac{2}{3}}}{\frac{\frac{n}{2}-\frac{1}{3}}{\frac{m}{2}+\frac{n}{2}-\frac{2}{3}}} \\ &= \frac{n}{m}\frac{\frac{m}{2}-\frac{1}{3}}{\frac{n}{2}-\frac{1}{3}} \\ &= \frac{n}{m}\frac{\frac{m}{2}-\frac{1}{3}}{\frac{n}{2}-\frac{1}{3}} \\ &= \frac{n}{m}\frac{3m-2}{3n-2} \\ &= \frac{n}{2n-2}\frac{3m-2}{m} \end{aligned}$$

S3 *Accounting for dependence among information sources*

In our dataset, predator feeding surveys included covariate information (predator size, prey size, and temperature) that was used to estimate field handling times on the basis of regression models for handling times parameterized using laboratory data. In estimating attack rates we treated the field covariates as part of the handling times data $\boldsymbol{H}$ and assumed they were independent of the feeding proportions data $\boldsymbol{F}$. The validity of this assumption may be assessed by plotting the regression covariates versus the observed feeding proportions, as shown in figure S1. In this figure, every point represents a single feeding survey. The x-axes are the averages of the (log-tranformed) covariate and the y-axes are the proportions of predators feeding. Only two species had sufficient data to be plotted and showed little evidence of a dependence.

If a lack of independence were evident it would need to be accounted for in the covariates distribution model. That is, although our model for the covariates was a multivariate normal, feeding survey level information (specifically proportions of predators feeding on each prey species) could be added to the model to affects its multivariate mean. This way, the mean covariate vector would be a function of the proportion of predators feeding on that prey type. Posterior distribution sampling could then be done by first sampling from the feeding proportions posterior distributions and then using the sampled feeding proportions to obtain samples from the handling times.

601 S4 *Supporting Figures*

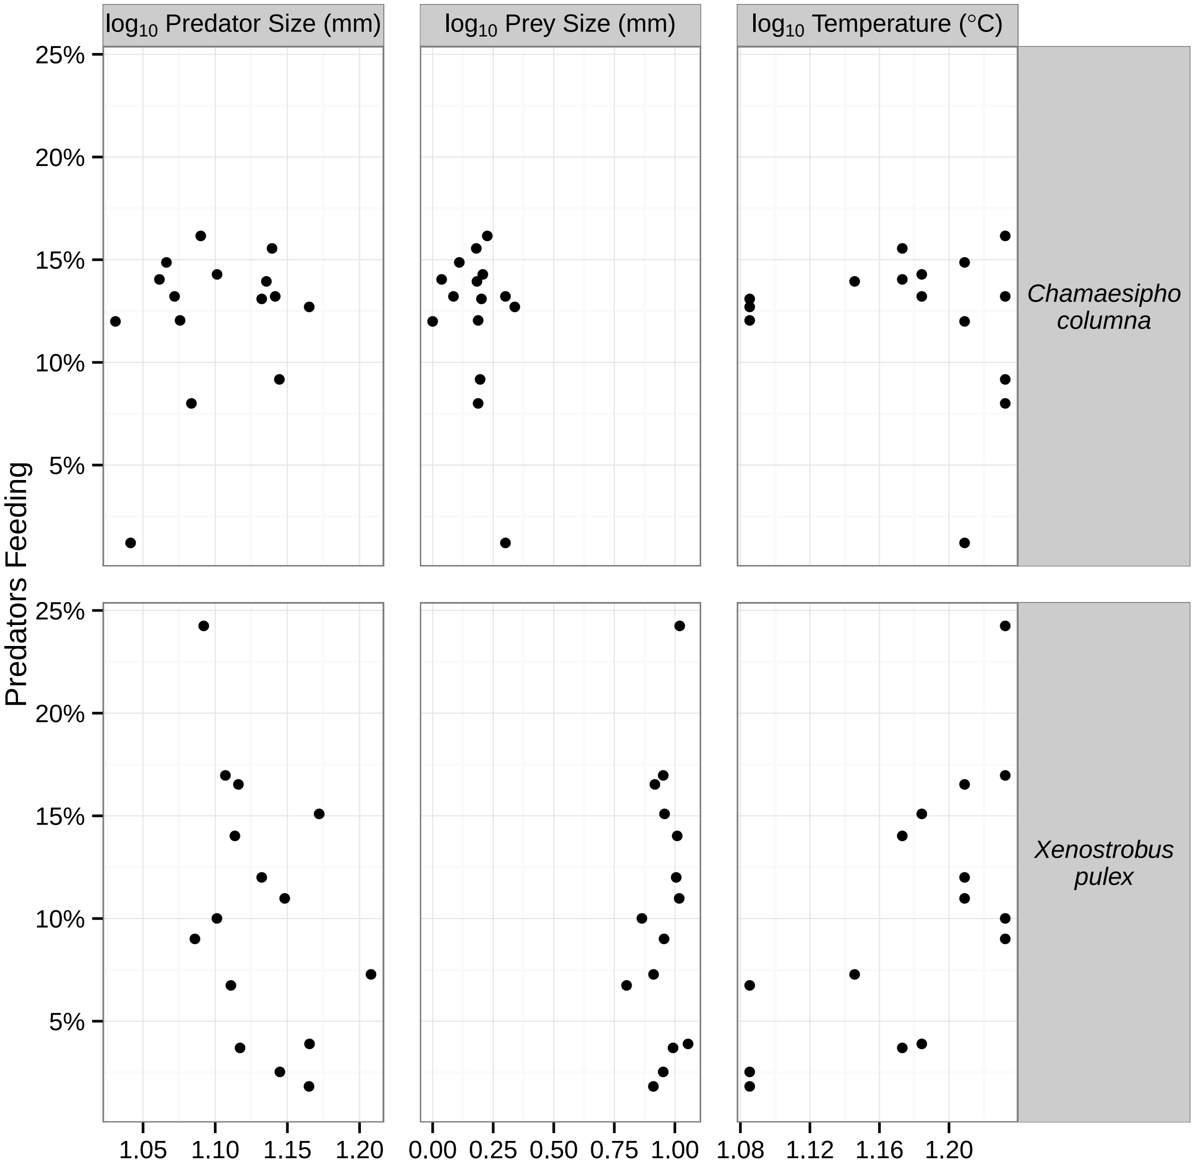

Figure S1: Average field covariates versus feeding proportions. Each point corresponds to a single feeding survey. Only species that appeared in more than three separate feeding surveys are shown. Of the eight species and three covariates, only *Xenostrobus pulex* showed any evidence of a relationship between feeding proportions and feeding covariates (i.e., between $\boldsymbol{F}$ and $\boldsymbol{H}$ in eqn 5).

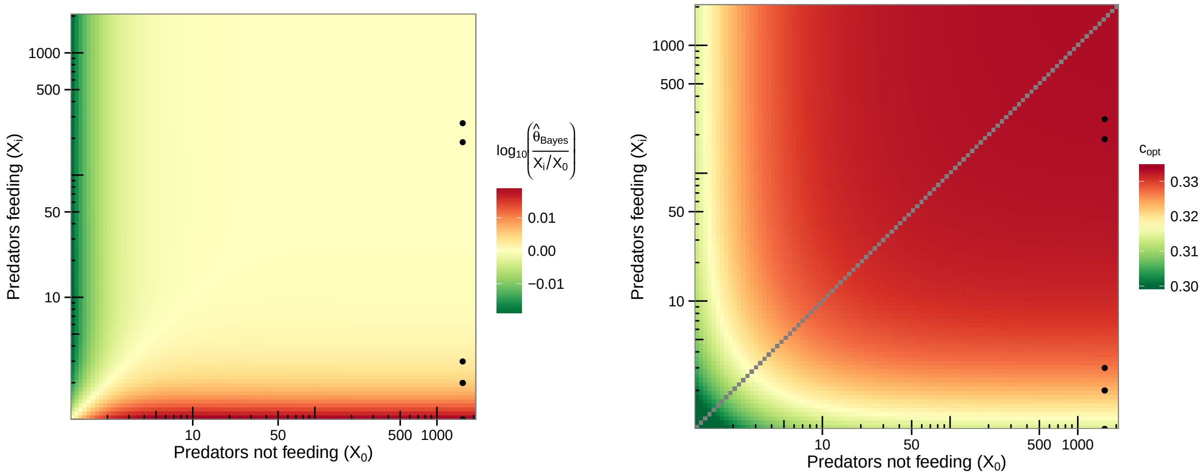

Figure S2: Given the skewed nature of prey-specific per capita attack rate posterior probability distributions, the distribution median serves as a more appropriate point estimate than the mean. Fig. 1 illustrates the difference between the posterior median and maximum likelihood estimate of the ratio of feeding and non-feeding predators as a function of the number of feeding individuals, showing how the neutral ($c = \frac{1}{3}$) prior minimizes this difference. As a generalization of Fig. 1, in the left panel, we illustrate this difference as a function of both the number of predators observed feeding and the number observed not feeding. The right panel shows that the "optimal" value of c that minimizes this difference (a function of both feeding and non-feeding individuals) is typically around $\frac{1}{3}$. In both cases, the survey data from our example are shown as black dots.

PeerJ PrePrints | https://dx.doi.org/10.7287/peerj.preprints.1382v1 | CC-BY 4.0 Open Access | rec: 18 Sep 2015, publ: 18 Sep 2015