

Bayesian characterization of uncertainty in species interaction strengths

Christopher Wolf¹  · Mark Novak² · Alix I. Gitelman¹

Received: 28 July 2016 / Accepted: 5 April 2017 / Published online: 19 April 2017
© Springer-Verlag Berlin Heidelberg 2017

Abstract Considerable effort has been devoted to the estimation of species interaction strengths. This effort has focused primarily on statistical significance testing and obtaining point estimates of parameters that contribute to interaction strength magnitudes, leaving the characterization of uncertainty associated with those estimates unconsidered. We consider a means of characterizing the uncertainty of a generalist predator's interaction strengths by formulating an observational method for estimating a predator's prey-specific per capita attack rates as a Bayesian statistical model. This formulation permits the explicit incorporation of multiple sources of uncertainty. A key insight is the informative nature of several so-called non-informative priors that have been used in modeling the sparse data typical of predator feeding surveys. We introduce to ecology a new neutral prior and provide evidence for its superior performance. We use a case study to consider the attack rates in a New Zealand intertidal whelk predator, and we illustrate not only that Bayesian point estimates can be made to correspond with those obtained by frequentist approaches, but also that estimation uncertainty as described by 95% intervals is more useful and biologically realistic using the

Bayesian method. In particular, unlike in bootstrap confidence intervals, the lower bounds of the Bayesian posterior intervals for attack rates do not include zero when a predator–prey interaction is in fact observed. We conclude that the Bayesian framework provides a straightforward, probabilistic characterization of interaction strength uncertainty, enabling future considerations of both the deterministic and stochastic drivers of interaction strength and their impact on food webs.

Keywords Type II functional response · Parameter estimation · Non-informative neutral priors · Stochastic and deterministic variation · Allometric scaling

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 often underlie forecasts of community dynamics (Petchey et al. 2015; Novak et al. 2016). 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 field experiments logistically prohibitive. Thus, many have resorted to indirect means of estimation, such as using energetic principles or allometric relationships from metabolic scaling theory (e.g., Neutel et al. 2002; Rall et al. 2012). More effort still has been devoted to estimating interaction strength parameters by characterizing predator

Communicated by Scott D. Peacor.

Electronic supplementary material The online version of this article (doi:10.1007/s00442-017-3867-7) contains supplementary material, which is available to authorized users.

✉ Christopher Wolf
wolfch@science.oregonstate.edu

¹ Department of Statistics, Oregon State University, Corvallis, OR 97331, USA

² Department of Integrative Biology, Oregon State University, Corvallis, OR 97331, USA

functional responses, largely on a pairwise experimental basis or by tracking predator kills over time (Vucetich et al. 2002; Jeschke et al. 2004).

As a consequence of the difficulty of estimating interaction strengths, most of the effort spent on characterizing species interactions has focused on obtaining point estimates and performing hypothesis tests (see also, Poisot et al. 2016; Wells and O’Hara 2013; Melián et al. 2014). 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 obtaining point estimates and testing whether the associated parameters are different from zero in the use of functional response experiments designed to determine the dependence of feeding rates on prey and/or predator densities (e.g., Jeschke et al. 2004; Ramos-Jiliberto et al. 2016). Thus only the “deterministic core” (i.e. systematic part) of alternative functional response formulations has generally been of interest. More specifically, functional responses have typically been fit to data using statistical models such as $F = \frac{aN}{1+ahN} + \epsilon$ (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 variables such as abundances (N) or parameters such as attack rates (a) and handling times (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 the explicit modeling of the variation intrinsic to both the parameters and variables by describing each by a distribution that is itself governed by deterministic and stochastic sources of variation.

As in many other contexts (e.g., Elderd and Miller 2016), the distinction between these two approaches to considering variation in interaction strengths (i.e., using a single error term, ϵ , or using a more comprehensive consideration of uncertainty) is important when the uncertainty of the 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 (e.g., Wootton 1997) and for the biological interpretation of the estimates themselves. Estimates derived from the allometric relationships underlying metabolic scaling theory, for example, are typically associated with several orders of magnitude in variation left unexplained by species body sizes (Rall et al. 2012; Kalinoski and DeLong 2016).

In this paper 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 observational methods stems from their features of more easily accommodating instances of trophic omnivory than experimental and time-series methods (Novak 2013); retaining the species-specific information lost in allometric and energetic approaches; and, given a sufficient number of feeding observations, estimating the species-specific attack rates for any subset (e.g., size class) of individuals within a focal predator population. Furthermore, with the method of Novak and Wootton (2008), attack rates may be estimated for all of a focal predator’s prey species simultaneously while also accounting for an inherent nonlinearity of predator–prey interactions because a multi-species Holling type II functional response—the most frequently observed functional response form among non-filter feeding consumers (Jeschke et al. 2004)—is assumed in the method’s derivation. The Bayesian formulation we develop here connects the deterministic multispecies type II functional response model with each of the sources of empirical data that contribute to the observational method’s per capita attack rate estimator. We thereby account for variation due to both sampling effort and the environment.

An issue with any Bayesian model is prior selection as posterior results may be sensitive to the choice of prior. This is a particularly sticky issue when data are sparse, wherein the prior weighs more heavily into results. Sparse data occur often in the context of species frequencies, particularly so in regard to species abundances and the diets of predators as these often range across many orders of magnitude (McGill et al. 2007; Novak 2013). Therefore, as a key part of this work, we give careful consideration to prior selection in attack rate estimation, and we assess the effects of alternative prior choices in estimating per capita attack rates. Because of the sensitivity we see to prior selection, we introduce to ecology a new prior that we show has appealing properties.

To demonstrate the Bayesian method’s utility and to explore the sensitivity of our results to prior choice, we apply it to data on the predator–prey interactions of a New Zealand intertidal whelk, contrasting these estimates with those obtained by non-parametric and parametric bootstrapping procedures. We show that posterior results using the sparse data typical of predator diet studies are sensitive to so-called non-informative priors commonly used in the Bayesian literature. We then show that our new ‘neutral’ non-informative prior gives Bayesian posterior point estimates that are intuitively appealing and consistent with estimates obtained by bootstrapping approaches. Finally, we show how estimation uncertainty as described by 95% intervals is considerably more constrained and biologically

realistic within the Bayesian framework; how the species-specific distribution of attack rates in whelks mirrors the skewed distribution of interaction strengths commonly seen at the community scale (Wootton and Emmerson 2005); how a deterministic component of intraspecific predator variation relates to predator body size; and, how our approach provides posterior probability distributions on per capita attack rate estimates that lend themselves to a more useful and descriptive characterization of interaction strengths than do point estimates alone.

Materials and methods

Model framework

Novak and Wootton (2008) introduced an observational method for obtaining point estimates of a generalist predator's prey-specific per capita attack rates (a_i ; where i indexes prey) using data on prey abundances (N_i), handling times (h_i), and “snapshot” feeding surveys in which the number of predator individuals feeding on each prey species is recorded. Assuming that a multispecies type II functional response

$$F_i = \frac{a_i N_i}{1 + \sum_{k=1}^S a_k h_k N_k}, \quad (1)$$

describes the predator population's per predator feeding rate on the i th prey species, the estimator for the attack rate on the i th prey is equivalent to

$$a_i = \frac{A_i}{A_0} \frac{1}{h_i N_i} \quad (2)$$

(see Appendix S1). Here, A_i is the number of observed predators feeding on prey i , and A_0 is the number of observed predators not feeding, during one or more snapshot surveys of the predator population.

Both Eqs. (1) and (2) are implicitly deterministic mathematical models that include no stochastic component, a deficiency we address below. The observational method capitalizes on the fact that handling times are more easily measured in laboratory experiments than in the field by using a smaller number of longitudinally-followed individuals than may be surveyed during a snapshot survey. Thus, even if handling time data are based on field observations, these will typically not be measured on the individuals observed during the snapshot feeding survey, hence the lengths of time those predators had been feeding are unknown and must be estimated. The feeding counts and species abundances similarly reflect estimates. In acknowledgement of this, we develop a parameter-based version of Eq. (2)—a statistical formulation of the attack rate

estimator that can incorporate sampling and environmental variation explicitly. We describe this next in the context of a specific case study. In Appendix S2, we provide a more general introduction to the principles of the Bayesian framework, as well as more specific details of our implementation than presented in the following main text.

Case study dataset

Our case study dataset pertains to the predatory whelk *Haustrum scobina* of the New Zealand marine intertidal. *Haustrum* feeds primarily on barnacles and mussels but also limpets and snails, 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 we use contains information from replicate feeding surveys and quadrat-based prey species abundance surveys from a single site (Tauranga Head), and laboratory-based handling time experiments. We summarize the relevant attributes of these data below, referring to Novak (2010, 2013) for further details.

Fifteen feeding surveys were conducted during low tides over two years. In each survey, the number of whelks not feeding (x_0) or feeding on each prey species (x_i) was recorded, amounting to a total of 2089 whelks observed. All but two of the eight recorded prey species were observed very rarely (Table 1). The sizes of the predator individuals (both feeding and not 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.

Prey abundance surveys used five replicate quadrats randomly distributed along two transects, each repeated three times over the same time periods in which the feeding surveys were conducted. 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 (i.e., no individuals may be present in the area), as well as stochastic variation associated with sampling effort (i.e., no individuals happened to be found in the quadrat). Aside from these zeroes, the abundance measurements are not integer counts but rather reflect densities derived on the basis of counts (for mobile prey) or percent cover-count relationships (for sessile prey) and the effective area of each quadrat as overlaid on an irregular substrate.

The dataset also contains data on laboratory experiments that Novak (2010) used to build a statistical model for the relationship between handling times and predator size; prey identity and size; and temperature. These experiments housed individual whelks in separate aquaria with

different prey and entailed hourly checks 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. The uncertainty associated with the interval censoring, along with estimation uncertainty associated with the experimental trials that were performed for each prey species (Table 1), reflect sources of variation that are modeled in the error term of the statistical model for handling times.

Bayesian model formulation

Treating the prey abundances, handling times, and feeding surveys data as independent, we now specify likelihood and prior models for each of these components. Following standard statistical practice for notation, we use uppercase letters to denote random variables, lowercase letters for realizations of random variables (the observed data), bold letters for vectors, and bold uppercase letters for matrices of random variables. We use f to represent the density function of an arbitrary distribution, 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 density functions of the same form. A table of notation is given in the Appendix (Table S1).

Modeling the feeding surveys

We model the combined feeding survey data of feeding and not feeding individuals using a multinomial likelihood and a Dirichlet prior having concentration parameters c . The resulting posterior distribution is also Dirichlet (eqn S10), for which we focus on the posterior medians (rather than means) as our point estimates of interest because medians are generally the more appropriate measure of a skewed distribution’s central tendency.

Four concentration parameter values have been used in the past to make Dirichlet priors non-informative: Laplace’s prior ($c = 1$), Jeffreys’ prior ($c = \frac{1}{2}$), Perks’ prior ($c = \frac{1}{S+1}$ where $S + 1$ is the length of the multinomial vector—see Appendix S2)—and Haldane’s prior ($c = 0$) (Hutter 2013). However, our preliminary investigations indicate that these priors result in posterior medians that differ substantially from the sample proportions, x_i/x_0 , particularly for rarely-observed prey. This leads to attack rate point estimates that, for rarely observed prey species, are entirely driven by the choice of the prior (i.e., the priors are not “non-informative” in this case, Fig. 1).

Because of the deficiency of the existing “non-informative” priors, we introduce to ecology a neutral prior, $c = \frac{1}{3}$, for modeling the ratios of multinomial parameters (i.e., the ratio of α_i feeding and α_0 not-feeding

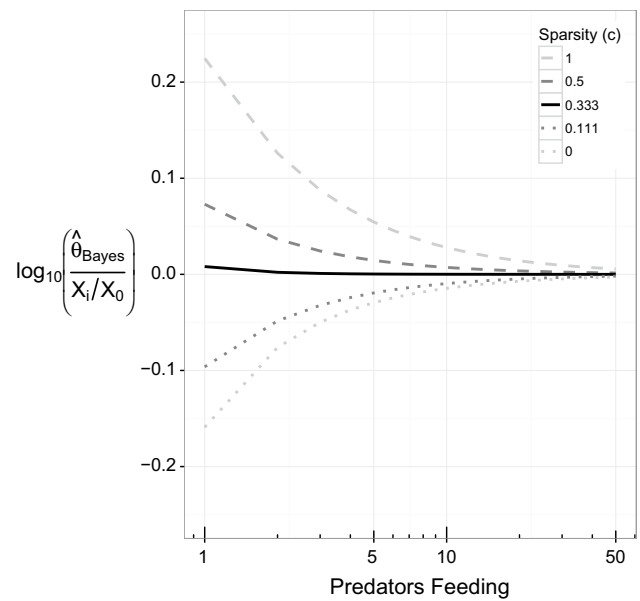


Fig. 1 Comparison of alternative “non-informative” priors in estimating the ratio of the proportions of feeding versus not feeding predator individuals. The x-axis reflects the number of predators observed in the process of feeding on a given prey species, with a total of 1629 individuals assumed to have been not feeding, corresponding to the number not feeding in our case study dataset (Table 1). The y-axis shows the difference in logarithms of the posterior median using a Dirich (c, \dots, c) prior and the maximum likelihood estimate of the ratio (sample proportion). 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). Our neutral prior ($c = \frac{1}{3}$) leads to estimates that most closely match the maximum likelihood estimates

individuals). This prior extends the insight in Kerman et al. (2011) that when $c = \frac{1}{3}$, the multinomial parameter posterior medians closely match the maximum likelihood estimates (MLE), which in this setting correspond to the sample proportions. We derive the prior 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:

$$f(\gamma_i|x_i, x_0) = \frac{x_0 + c}{x_i + c} \cdot g\left(\frac{x_0 + c}{x_i + c} \cdot \gamma_i; 2(x_i + c), 2(x_0 + c)\right), \tag{3}$$

where x_i and x_0 are the observed counts of feeding and non-feeding individuals and $g(y; d_1, d_2)$ 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 $\text{med}(F_n^m) \approx \frac{n}{3n-2} \frac{3m-2}{m}$ and setting it equal to the MLE of $\frac{\alpha_i}{\alpha_0}, \frac{x_i}{x_0}$, yields the solution $c = \frac{1}{3}$ (Appendix S3). Thus, the neutral prior leads to posterior medians that closely match the MLEs for both multinomial parameters, as shown by Kerman et al. (2011), and for the ratios of multinomial parameters.

Modeling the abundance surveys

We use a zero-inflated gamma (ZIG) model to accommodate the numerous zeros in the prey abundance data. By conditioning on whether or not a zero occurs, the likelihood density of the ZIG distribution can be expressed in terms of ρ (the probability of a zero) and $f(y; \alpha, \beta)$ (the usual gamma density with shape α , rate β , and mean $\frac{\alpha}{\beta}$). The ZIG density is separable in ρ and (α, β) , which means that the zero-inflation parameter can be treated separately, provided a separable prior is used. Thus, for each prey species, we model the number of observed zeros using a binomial distribution with a uniform prior on ρ . For the gamma part, we use $\log(\alpha) \sim \text{Unif}(-100, 100)$ and $\log(\beta) \sim \text{Unif}(-100, 100)$ priors 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).

Modeling the handling time experiments

We use regression to model the relationship between handling times and the predator size, prey size and temperature covariates of the laboratory experiments. We obtain average ‘field-estimated’ handling times for use in the attack rate estimation by combining these regression coefficients with the same covariate information obtained during feeding surveys.

Specifically, we consider the i th handling time observation for a given prey species to be associated with a covariate vector, \mathbf{X}_i , consisting of temperature, predator size, and prey size (all log transformed, and a 1 for the intercept term). We then model the likelihood of the i th handling time using a normal distribution (mean $e^{\mathbf{X}_i^T \boldsymbol{\beta}}$, variance σ^2) plus a uniform (minimum $-\frac{l_i}{2}$, maximum $\frac{l_i}{2}$) error corresponding to the interval censoring with which handling times were observed. The exponential link of the normal distribution mean avoids negative mean handling time estimates.

Treating the field covariates (predator size, prey size, and temperature) as random to account for sampling variability, we model the distributions of the (log-transformed) covariate observations $\mathbf{X}_1, \dots, \mathbf{X}_N$, where N is the total number of field observations, as independent, identically distributed, and drawn from a multivariate normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}'$. We use non-informative multivariate normal and inverse Wishart priors for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}'$ respectively (Fink 1997). Letting \mathbf{X}^* follow the posterior predictive distribution (our estimate of the distribution of the covariates), the mean handling time is $E(e^{\boldsymbol{\beta}^T \mathbf{X}^*})$ (Appendix S2). We estimate this expectation by sampling from the posterior distribution of the regression parameters, sampling new covariates from their posterior predictive distribution, computing $e^{\boldsymbol{\beta}^T \mathbf{X}^*}$ for each sample, and averaging across all samples. The weak law of large

numbers ensures convergence to $E(e^{\boldsymbol{\beta}^T \mathbf{X}^*})$ as sample size increases (Petrov 1995).

Accounting for spatio-temporal variation

As the feeding survey, prey abundance, and handling time data all have multiple levels of spatial and temporal structure (e.g. structured variation across replicate quadrats, transects, seasons, and years), we consider several hierarchical models seeking to account for such possible dependencies (Cressie et al. 2009). Unfortunately, insufficient data at the larger scales made fitting such models impractical in the case of the whelk data, with models either failing to converge or resulting in inference similar to that from the non-hierarchical models. We provide an example of how to implement a hierarchical model to incorporate spatial and temporal variation in Appendix S4.

The data of our case study are sufficient to illustrate the utility of the observational approach for assessing the deterministic influence of intraspecific predator variation (in the form of individual body size) on attack rate estimates. To accomplish this we divide all predator observations made during the feeding surveys into eight groups based on predator body size, with approximately equal numbers of predator individuals per group (Table S2). To draw inferences relevant to theory on the allometric scaling between predator size and attack rates (Rall et al. 2012; Kalinoski and DeLong 2016), we convert predator lengths (shell length, L , in mm) to masses (total wet weight, w , in grams) using $W = 1.214 \times 10^{-4} \times L^{3.210}$ (Novak 2013). Our analysis here focuses on the two most common prey species (*C. columna* and *X. pulex*, Table 1) to ensure that sufficient data are available for each predator size group.

Estimating per capita attack rates

Using the likelihoods and priors given above for the feeding surveys, abundances and handling times, we draw samples from the posterior distribution of the parameters using Markov Chain Monte Carlo (MCMC). We use JAGS with the R package ‘rjags’ (Plummer and Stukalov 2016) for the MCMC sampling. We then combine posterior samples using eqn S6 to produce samples from the posterior distribution of attack rates for each prey species. We also plot the individual attack rate components’ posterior distributions separately to show how each part contributes to the attack rates’ posterior distributions. In our approach, we treat handling times, \mathbf{H} , as being independent of the predator feeding surveys, \mathbf{F} , even though we use covariate observations of predator size, prey size and temperature from the feeding surveys informing \mathbf{F} to inform \mathbf{H} by combining them with the laboratory-based handling time regression coefficients

associated with these covariates. We establish the validity of this assumption by examining the relationship between feeding proportions and covariate averages between the individual surveys (Appendix S5).

We verify Markov chain convergence using trace plots and the Gelman and Rubin convergence diagnostic (Gelman and Rubin 1992), remove samples obtained before the chains had converged (i.e., burn-in samples), and thin each chain to ensure independence among the remaining samples. We compute scale reduction factors—a convergence diagnostic that compares ‘within’ versus ‘between’ chain variability—using 250 independent chains with random initial values. With the help of trace plots we determine burn-in lengths separately for feeding survey, prey abundance, and handling time models. We base our final inferences on 1000 samples after confirming that independent sets of 1000 samples led to the same conclusions.

Comparison of Bayesian and bootstrapping procedures

We assess the utility and performance of our Bayesian approach by contrasting point and 95% interval estimates from (1) the model with Laplace’s prior ($c = 1$) on the Dirichlet feeding proportions; (2) the model with Haldane’s prior ($c = 0$); and (3) the model with the neutral prior ($c = \frac{1}{3}$) to estimates obtained using (4) non-parametric and (5) parametric bootstrapping procedures.

For the non-parametric bootstrap, we sample with replacement from each of the feeding survey, prey abundance, and handling time datasets until we draw the same number of samples as was present in each dataset (Efron and Tibshirani 1994). We calculate per capita attack rates for 1000 sets of such resampled data to estimate the mean and 95% confidence intervals of the corresponding bootstrapped distributions.

We implement the parametric bootstrap using the likelihood function that we assume in our Bayesian method. That is, we use the data to estimate the parameters of the three likelihood functions (eqn S9, eqn S11 and eqn S12) by maximum likelihood, use these fit likelihood functions to simulate new datasets, and combine samples from the three distributions to estimate per capita attack rates. We then determine the medians and 95% confidence intervals of the resulting bootstrapped attack rate distributions.

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 simple, empirically useful interpretation. Rather, if 95% confidence intervals were repeatedly constructed using newly collected equivalent datasets, 95% of them would contain the ‘true’ value of the parameter. To highlight this difference, we use our Bayesian results to estimate the probability that

the prey species with the highest posterior median attack rate has an attack rate greater than the species with the next highest posterior median attack rate, and so on. This type of inference is not possible within the frequentist framework (i.e., using the bootstrapped samples).

Results

The comparison of the log differences in the MLE relative to the posterior medians of the Bayesian model for several values of the concentration parameter c of the Dirichlet prior evidenced that our non-informative neutral prior, $c = \frac{1}{3}$, is indeed the most appropriate prior to use on the feeding proportions when the median is the preferred point estimate of the attack rate posterior distribution (Fig. 1). All other “non-informative” priors result in median point estimates that are either considerably greater or less than the MLE at small sample sizes, converging only after significantly more observations are assumed than in fact had been made for most of the prey species in our case study data set. Concentration parameter values greater than $1/3$ result in median point estimates that are inflated relative to the MLE, whereas values less than $1/3$ result in estimates that are less than the MLE.

Similarly, the comparison of the Bayesian and bootstrapping approaches applied to the entire case study data set also indicates that the model with the neutral prior ($c = \frac{1}{3}$) was both sufficient for, and performs best in, describing the variation inherent in the estimated rates with which *Haustrum scobina* attacked its prey species (Fig. 2). With the neutral prior, the model exhibits median point estimates most closely matching the point estimates of the two bootstrapping approaches regardless of the number of feeding observations per prey species. The two bootstrap distributions, in contrast, frequently exhibit lower 95% confidence interval end points of zero; a nonsensical result given that the consumption of these species was in fact observed. Consistent with Fig. 1, the models using Laplace’s ($c = 1$) or Haldane’s ($c = 0$) priors result in inflated and depressed attack rate point estimates respectively, particularly for prey species that were observed infrequently in the feeding surveys.

Figure 3 shows the posterior probability distributions of *Haustrum*’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. Using the posterior distribution of the attack rate, we estimate the probability of the attack rate on *Mytilus galloprovincialis* being greater than on *Xenostrobus pulex* to be 0.68 (relative to the null expectation of 0.5 given no difference in

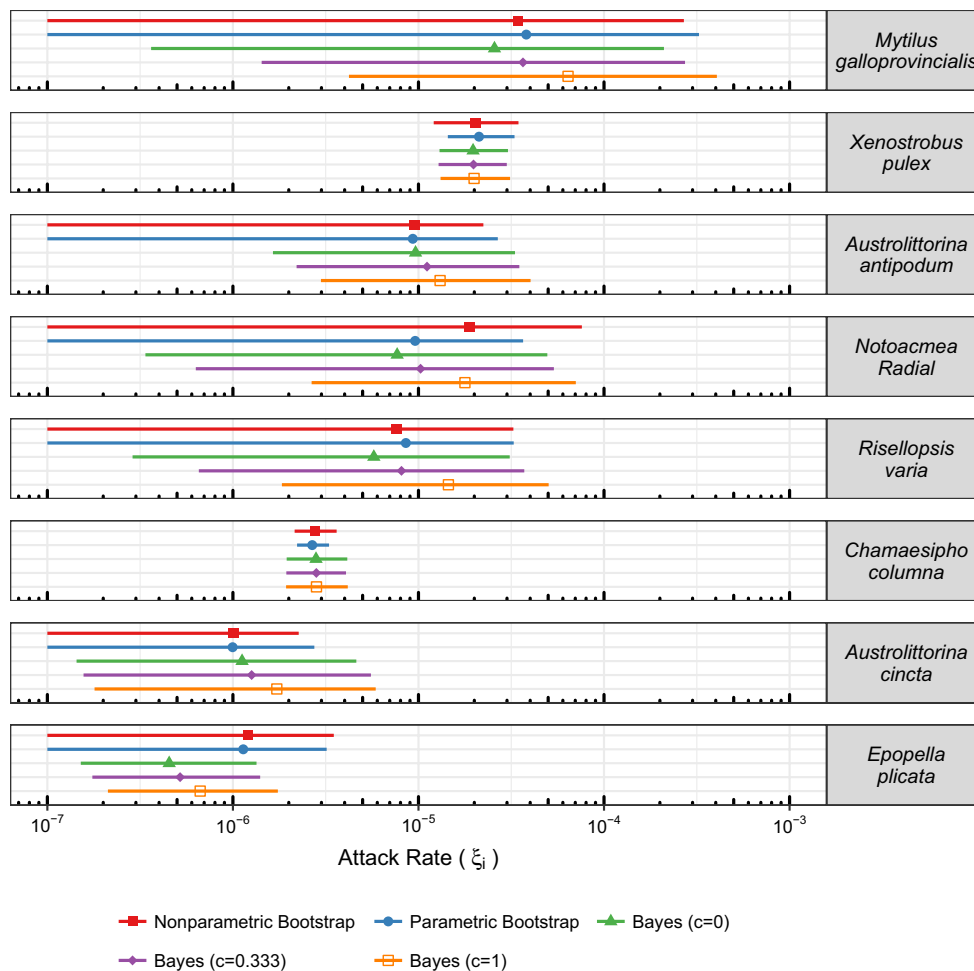


Fig. 2 Comparison of the frequentist and Bayesian approaches to estimating the per capita attack rates with which *Haustrum scobina* consumed its 8 prey species. Variation in attack rate estimates is illustrated for each by the medians and 95% equal-tailed intervals of their distributions. The approaches are organized the same for each prey species as, from top to bottom: (1) non-parametric bootstrap (filled square), (2) parametric bootstrap (circle), (3–5) Bayesian procedure

with sparsity parameters $c = 0$ (Haldane’s prior; triangle), $\frac{1}{3}$ (neutral prior; diamond), and 1 (Laplace’s prior; open square) respectively. Unlike the 95% confidence intervals for the bootstrap procedures which often span zero ($=10^{-7}$ for graphical convenience), the 95% posterior intervals of the Bayesian method indicate the regions where attack rates lie with 95% probability. A color version of this figure is available online

probabilities), even though the 95% posterior interval for *X. pulex* is completely contained within the 95% interval for *Mytilus galloprovincialis* (Fig. 3). In addition, we estimate the probability of the attack rate on *X. pulex* being greater than on *Austrolittorina antipodum* to be 0.80. Probability estimates such as these are unavailable from the bootstrap samples.

Figure 3 also illustrates an additional utility of modeling the three components of the attack rates estimator—the abundances, handling times, and feeding ratios—explicitly, permitting the decomposition of their contributions to the variation of inferred attack rate estimates. For example, while the interspecific variation observed in handling times spans less than two orders of magnitude, the interspecific variation in feeding ratios and prey abundances each span more than four orders of magnitude (Fig. 3). Hence the

latter two components contributed more to the interspecific variation seen in the estimated attack rates than did variation in handling times. That said, nearly all of the interspecific variation in feeding ratios was driven by the two most commonly observed species, *Xenostrobus pulex* and *Chamaesipho columna* (Fig. 3). Thus, for most prey species, interspecific variation in the attack rates may be inferred to have been driven primarily by variation in their abundances.

The partitioning of the predator observations into eight body size groups reveals a generally positive relationships between a whelk’s body size and its attack rate on both of *Haustrum*’s two primary prey species (Fig. 4). The point estimates for the allometric slopes of these intraspecific relationships were 0.068 and 0.256 for *C. columna* and *X. pulex* respectively. These slopes

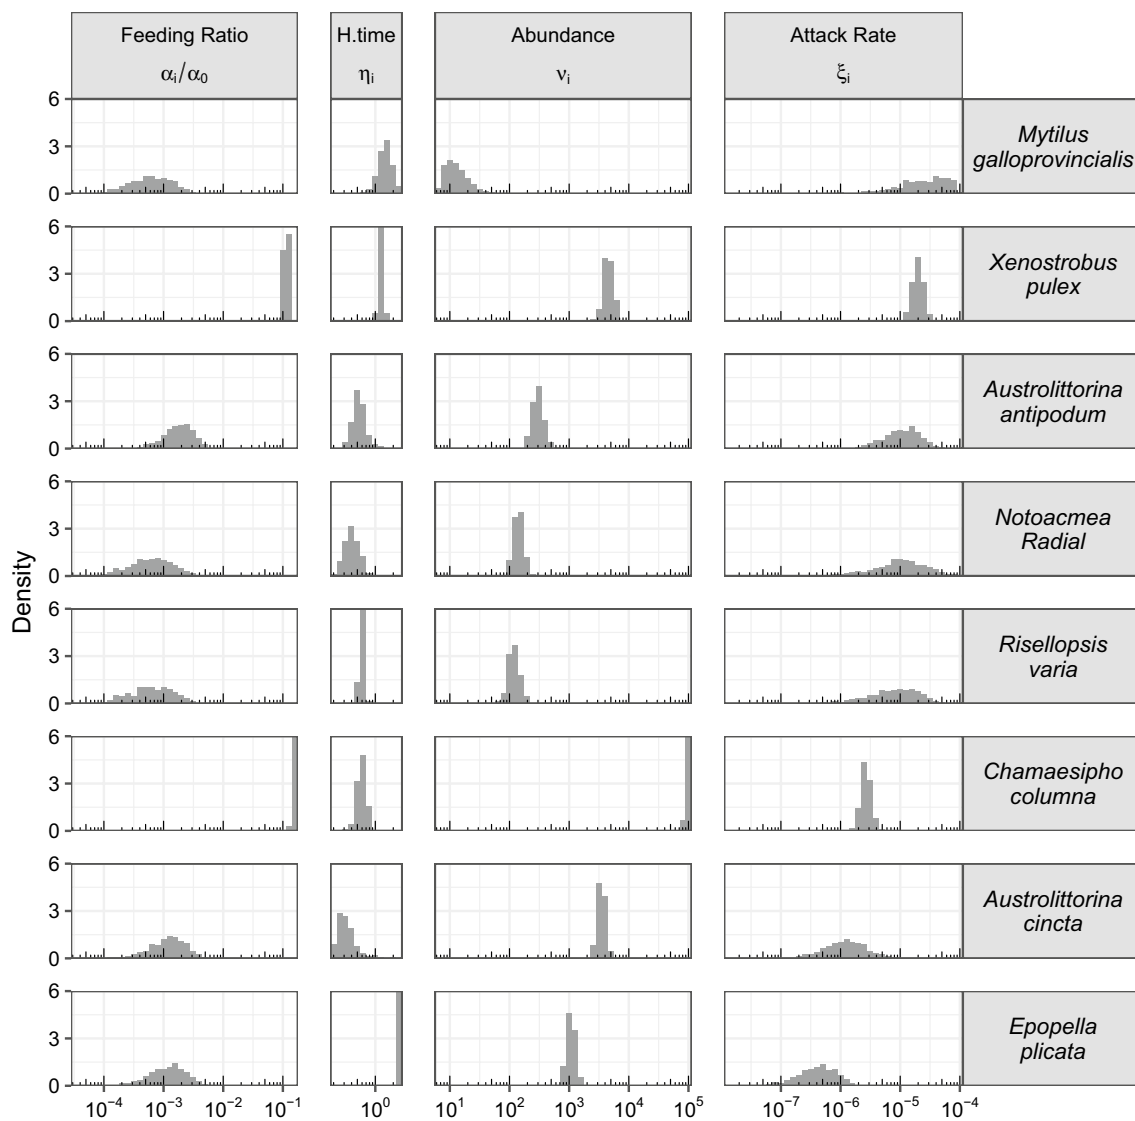


Fig. 3 Posterior distributions for *Haustrum scobina*'s per capita attack rate parameters ($\text{prey} \cdot \text{predator}^{-1} \cdot \text{prey}^{-1} \cdot \text{m}^{-2} \cdot \text{day}^{-1}$) and their components ($\xi_i = \frac{\alpha_i}{\alpha_0} \cdot \frac{1}{v_i} \cdot \frac{1}{\eta_i}$) using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

believe substantial variation surrounding each size group's median attack rates on the two prey species. Thus, while the probability that the attack rate of the largest whelk group (17–28 mm) was greater than the corresponding attack rate of the smallest group (6–10 mm) on *X. pulex* was 0.89, it was only 0.60 for their attack rates on *C. columna* (versus the null expectation of 0.5).

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

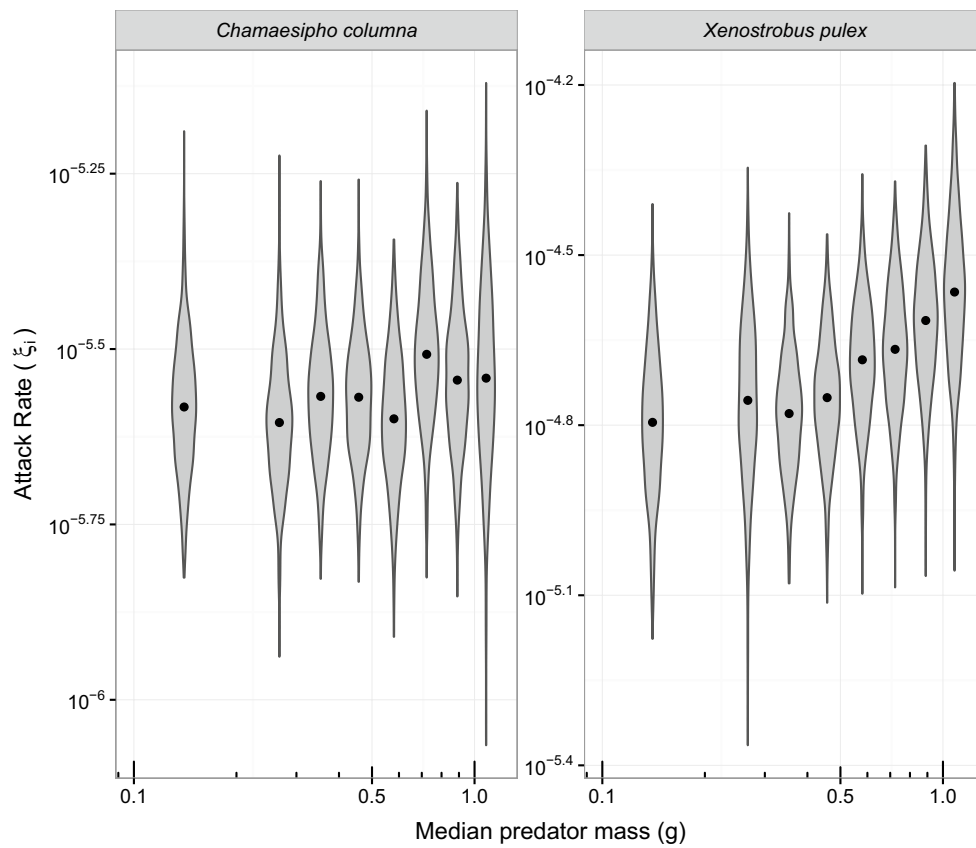


Fig. 4 Deterministic variation in per capita attack rates due to predator body size for the two prey species consumed by *Haustrum scobina* most frequently. Points indicate posterior medians and violin widths reflect posterior probabilities of the attack rate magnitudes. Predator individuals were split into body mass size classes of roughly

equally-numbered counts (Table S2) with class median body weights (W , total wet weight in grams) estimated on the basis of each whelk’s shell length (L , in mm) ($W = 1.214 \times 10^{-4} \times L^{3.210}$, from Novak (2013))

Table 1 Summary statistics for the datasets we use to estimate the per capita attack rates with which the intertidal predator *Haustrum scobina* fed on its eight prey species

Prey species	Feeding observations	Abundance zeros	Handling-time experiments
<i>Chamaesipho columna</i>	265	0	6
<i>Xenostrobus pulex</i>	185	0	52
<i>Austrolittorina antipodum</i>	3	2	0
<i>Austrolittorina cincta</i>	2	0	46
<i>Epopella plicata</i>	2	0	1
<i>Mytilus galloprovincialis</i>	1	24	15
<i>Notoacmea Radialspokes</i>	1	5	66
<i>Risellopsis varia</i>	1	3	68
Not feeding	1629		
Total surveyed	2089		

“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 prey species’ abundances. “Handling-time experiments” indicates the number of laboratory experiment trials that were used to estimate handling time regression coefficients for each prey species

variation of these component parameters. The propagation of both such forms of variation is inherent to all other experimental and observational approaches as well.

Our case study serves to illustrate how the explicit modeling of the components that go into estimating attack rates using the observational approach (i.e. feeding ratios,

handling times, and prey abundances) provides insight into the drivers of variation in the attack rates among prey species. It is worth highlighting, however, that the attack rate posterior distributions that we estimate here represent population-level uncertainty (i.e. uncertainty in the overall attack rates) rather than intraspecific variation in the diets of individuals. Thus the wide posterior interval for *Mytilus galloprovincialis*, for example, means that there is a high degree of overall uncertainty about the attack rate on this species, rather than being a reflection of the individual specialization that many generalist predators, including whelk species, are known to exhibit (Bolnick et al. 2003).

Nevertheless, just as for interspecific comparisons of variation, insight into the species-specific uncertainty is also possible by the partitioning of its components (which combine additively on the logarithmic scale). Thus for prey species like *Mytilus* that occur infrequently in feeding surveys (all prey species except *Chamaesipho columna* and *Xenostrobus pulex*, Table 1), the largest source of uncertainty comes from the estimation of the feeding ratios (Fig. 3). Additional feeding surveys would therefore provide the most important data for better constraining this stochastic source of uncertainty. In contrast, a simple way to obtain further insight into potential additional deterministic drivers of species-specific variation is to apply the estimation approach to subsets of the data as in Fig. 4. In our case study, the point estimates for the allometric slopes of the intraspecific relationships between body size and attack rate are both far lower than predicted by theory for the dependence of a specialist predator's interspecific attack rates on predator body size (Rall et al. 2012), possibly suggesting size-dependent changes in prey preferences.

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. 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 our case study, we are able to infer the probability of the attack rate on *Mytilus galloprovincialis* being greater than on *Xenostrobus pulex*, and so on.

In the context of food webs and predator–prey interactions, this 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; Koslicki and Novak 2016). 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 (Poisot et al. 2016). 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 'non-informative' (objective) priors is particularly important when data are sparse (Van Dongen 2006). It follows that, for rarely observed prey species, different prior specifications lead to different point estimates of the per capita attack rates (Figs. 1, 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 (Allesina and Tang 2012). These considerations are avoided only when all prey occur frequently in a predator's diet (see *Xenostrobus pulex* and *Chamaesipho 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 $Y \sim \text{Bin}(n, p)$ 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. $Y \in \{1, \dots, n - 1\}$). It is true that in hierarchical models, to which our framework is naturally extended (Appendix S4), $Y \in \{0, n\}$ is more likely for any individual survey, but this is not an issue as inference at the survey level is typically not desired in the absence of additional covariates.

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 (Eq. (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 prey 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 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 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 skew 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.

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 the deterministic and stochastic variation in data are needed to accompany the numerous mathematical models that have been proposed (Poisot et al. 2015). Our work represents a step in this direction, and has the potential to be extended in several directions to address questions of optimal foraging theory, prey profitability, and functional response formulations other than the multispecies type II response considered here (Pyke 1984; Berger and Pericchi 1996; Novak et al. 2017).

Acknowledgements This work was supported by NSF DEB-1353827 and DEB-0608178. We thank the editor and two anonymous reviewers whose comments substantially improved the manuscript.

Author contribution statement MN and AIG originally formulated the idea for the project, CW analyzed the data, and CW, MN, and AIG wrote the manuscript.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Data accessibility 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

- Allesina S, Tang S (2012) Stability criteria for complex ecosystems. *Nature* 483:205–208
- Berger JO, Pericchi LR (1996) The intrinsic Bayes factor for model selection and prediction. *J Am Stat Assoc* 91:109–122
- Bolker BM (2008) *Ecological models and data in R*, 508th edn. Princeton University Press, Princeton
- Bolnick DI, Svanbck R, Fordyce JA, Yang LH, Davis JM, Hulsey CD, Forister ML (2003) The ecology of individuals: incidence and implications of individual specialization. *Am Nat* 161:1–28
- Calder C, Lavine M, Muller P, Clark JS (2003) Incorporating multiple sources of stochasticity into dynamic population models. *Ecology* 84:1395–1402
- Clark JS (2005) Why environmental scientists are becoming Bayesians. *Ecol Lett* 8:2–14
- Codling EA, Dumbrell AJ (2012) Mathematical and theoretical ecology: linking models with ecological processes. *Interface Focus* 2:144–149
- Cressie N, Calder CA, Clark JS, Hoef JMV, Wikle CK (2009) Accounting for uncertainty in ecological analysis: the strengths and limitations of hierarchical statistical modeling. *Ecol Appl* 19:553–570
- Efron B, Tibshirani RJ (1994) *An introduction to the bootstrap*. CRC Press, Boca Raton
- Elder BD, Miller TEX (2016) Quantifying demographic uncertainty: Bayesian methods for integral projection models. *Ecol Monogr* 86:125–144
- Fink D (1997) A compendium of conjugate priors. Technical Report
- Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. *Stat Sci* 7:457–472
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2013) *Bayesian data analysis*, 3rd edn. Chapman and Hall/CRC, Boca Raton
- Hutter M (2013) Sparse adaptive Dirichlet-multinomial-like processes. arXiv preprint. [arXiv:1305.3671](https://arxiv.org/abs/1305.3671)
- Jeschke JM, Kopp M, Tollrian R (2004) Consumer-food systems: why type I functional responses are exclusive to filter feeders. *Biol Rev* 79:337–349
- Kalinoski RM, DeLong JP (2016) Beyond body mass: how prey traits improve predictions of functional response parameters. *Oecologia* 180:543–550
- Kerman J et al (2011) Neutral noninformative and informative conjugate beta and gamma prior distributions. *Electron J Stat* 5:1450–1470
- Kosllicki D, Novak M (2016) Exact probabilities for the indeterminacy of complex networks as perceived through press perturbations. arXiv preprint. [arXiv:1610.07705](https://arxiv.org/abs/1610.07705)
- McGill BJ, Etienne RS, Gray JS, Alonso D, Anderson MJ, Benecha HK, Dornelas M, Enquist BJ, Green JL, He F et al (2007) Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. *Ecol Lett* 10:995–1015
- Melián CJ, Baldó F, Matthews B, Vilas C, González-Ortegón E, Drake P, Williams RJ (2014) Individual trait variation and diversity in food webs. *Adv Ecol Res* 50:207–241
- Neutel AM, Heesterbeek JAP, Ruiters PCd (2002) Stability in real food webs: weak links in long loops. *Science* 296:1120–1123
- 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. *Proc R Soc B: Biol Sci* 280:20131415
- Novak M, Wootton JT (2008) Estimating nonlinear interaction strengths: an observation-based method for species-rich food webs. *Ecology* 89:2083–2089
- Novak M, Wootton JT, Doak DF, Emmerson M, Estes JA, Tinker MT (2011) Predicting community responses to perturbations in the face of imperfect knowledge and network complexity. *Ecology* 92:836–846
- Novak M, Yeakel J, Noble AE, Doak DF, Emmerson M, Estes JA, Jacob U, Tinker MT, Wootton JT (2016) Characterizing species interactions to understand press perturbations: what is the community matrix? *Annu Rev Ecol Evolut Syst* 47:409–432. doi:10.1146/annurev-ecolsys-032416-010215
- Novak M, Wolf C, Coblenz K, Shepard I (2017) Quantifying predator dependence in the functional response of generalist predators. *Ecol Lett*. doi:10.1101/082115
- Paine RT (1992) Food-web analysis through field measurement of per capita interaction strength. *Nature* 355:73–75
- Petchey OL, Pontarp M, Massie TM, Kfi S, Ozgul A, Weilenmann M, Palamara GM, Altermatt F, Matthews B, Levine JM, Childs DZ, McGill BJ, Schaepman ME, Schmid B, Spaak P, Beckerman AP, Pennekamp F, Pearse IS (2015) The ecological forecast horizon, and examples of its uses and determinants. *Ecol Lett* 18:597–611
- Petrov VV (1995) *Limit theorems of probability theory: sequences of independent random variables*. Oxford University Press, Oxford
- Plummer M, Stukalov A, Denwood M (2016) rjags: Bayesian Graphical Models using MCMC. <https://cran.r-project.org/web/packages/rjags/index.html>. Accessed 23 Jan 2017
- Poisot T, Stouffer DB, Gravel D (2015) Beyond species: why ecological interaction networks vary through space and time. *Oikos* 124:243–251
- Poisot T, Cirtwill AR, Cazelles K, Gravel D, Fortin MJ, Stouffer DB (2016) The structure of probabilistic networks. *Methods Ecol Evolut* 7:303–312
- Power ME, Tilman D, Estes JA, Menge BA, Bond WJ, Mills LS, Daily G, Castilla JC, Lubchenco J, Paine RT (1996) Challenges in the quest for keystones. *BioScience* 46:609–620
- Pyke GH (1984) Optimal foraging theory: a critical review. *Annu Rev Ecol Syst* 15:523–575
- Rall BC, Brose U, Hartvig M, Kalinkat G, Schwarzmüller F, Vucic-Pestic O, Petchey OL (2012) Universal temperature and body-mass scaling of feeding rates. *Philos Trans R Soc B: Biol Sci* 367:2923–2934
- Ramos-Jiliberto R, Heine-Fuster I, Reyes CA, González-Barrientos J (2016) Ontogenetic shift in Daphnia-algae interaction strength altered by stressors: revisiting Jensen's inequality. *Ecol Res* 31:811–820
- Syversveen AR (1998) Noninformative Bayesian priors. Interpretation and problems with construction and applications, Preprint statistics 3
- Tuyl F, Gerlach R, Mengersen K (2008) A comparison of Bayes–Laplace, Jeffreys, and other priors: the case of zero events. *Am Stat* 62:40–44

- Van Dongen S (2006) Prior specification in Bayesian statistics: three cautionary tales. *J Theor Biol* 242:90–100
- Vucetich JA, Peterson RO, Schaefer CL (2002) The effect of prey and predator densities on wolf predation. *Ecology* 83:3003–3013
- Wells K, O’Hara RB (2013) Species interactions: estimating per-individual interaction strength and covariates before simplifying data into per-species ecological networks. *Methods Ecol Evolut* 4:1–8
- Wootton JT (1994) Predicting direct and indirect effects: an integrated approach using experiments and path analysis. *Ecology* 75:151–165
- Wootton JT (1997) Estimates and tests of per capita interaction strength: diet, abundance, and impact of intertidally foraging birds. *Ecol Monogr* 67:45
- Wootton JT, Emmerson M (2005) Measurement of interaction strength in nature. *Annu Rev Ecol Evolut Syst* 36:419–444
- Yeakel JD, Novak M, Guimarães PR Jr, Dominy NJ, Koch PL, Ward EJ, Moore JW, Semmens BX (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