

Statistical Reports

Ecology, 100(7), 2019, e02706
© 2019 by the Ecological Society of America

Statistical modeling of patterns in annual reproductive rates

MOLLIE E. BROOKS,^{1,9} KASPER KRISTENSEN,¹ MARIA ROSA DARRIGO,² PAULO RUBIM,³ MARÍA URIARTE,⁴
EMILIO BRUNA,^{5,6,7} AND BENJAMIN M. BOLKER⁸

¹National Institute of Aquatic Resources, Technical University of Denmark, Lyngby 2800 Denmark

²Greenpeace Brazil, Avenida Joaquim Nabuco, 2367, Centro, Manaus, Amazonas 6902-031 Brazil

³Sítio Cotíngia, Catuçaba, São Luiz do Paraitinga, São Paulo 12140-000 Brazil

⁴Department of Ecology, Evolution & Environmental Biology, Columbia University, New York, New York 10027 USA

⁵Department of Wildlife Ecology and Conservation, University of Florida, Gainesville, Florida 32611-0430 USA

⁶Center for Latin American Studies, University of Florida, Gainesville, Florida 32611-5530 USA

⁷Biological Dynamics of Forest Fragments Project, INPA-PDBFF, CP 478, Manaus, Amazonas 69011-970 Brazil

⁸Department of Mathematics & Statistics and Department of Biology, McMaster University, Hamilton, Ontario L8S 4K1 Canada

Citation: Brooks, M. E., K. Kristensen, M. R. Darrigo, P. Rubim, M. Uriarte, E. Bruna, and B. M. Bolker. 2019. Statistical modeling of patterns in annual reproductive rates. *Ecology* 100(7):e02706. 10.1002/ecy.2706

Abstract. Reproduction by individuals is typically recorded as count data (e.g., number of fledglings from a nest or inflorescences on a plant) and commonly modeled using Poisson or negative binomial distributions, which assume that variance is greater than or equal to the mean. However, distributions of reproductive effort are often underdispersed (i.e., variance < mean). When used in hypothesis tests, models that ignore underdispersion will be overly conservative and may fail to detect significant patterns. Here we show that generalized Poisson (GP) and Conway-Maxwell-Poisson (CMP) distributions are better choices for modeling reproductive effort because they can handle both overdispersion and underdispersion; we provide examples of how ecologists can use GP and CMP distributions in generalized linear models (GLMs) and generalized linear mixed models (GLMMs) to quantify patterns in reproduction. Using a new R package, *glimmTMB*, we construct GLMMs to investigate how rainfall and population density influence the number of fledglings in the warbler *Oreothlypis celata* and how flowering rate of *Heliconia acuminata* differs between fragmented and continuous forest. We also demonstrate how to deal with zero-inflation, which occurs when there are more zeros than expected in the distribution, e.g., due to complete reproductive failure by some individuals.

Key words: Conway-Maxwell-Poisson; fecundity; generalized Poisson; *Heliconia acuminata*; *Oreothlypis celata*; overdispersion; regression; underdispersion; zero-inflation.

INTRODUCTION

Among-individual variation in reproduction, due to both environmental conditions and individual attributes, underpins life-history theory and is central to empirical studies of demography (Lack 1947, van Noordwijk and de Jong 1986). Across taxa, many metrics of individual reproduction used by ecologists (e.g., clutch size, seeds per fruit) are discrete counts with upper limits determined by behavioral or physiological constraints such as

ovule number, resource availability, or parental care strategies. These constraints often lead to underdispersion where the variance is smaller than the mean (Kendall and Wittmann 2010, Lynch et al. 2014). In addition, distributions of reproductive effort are often “zero-inflated”, meaning that there are more zeros observed than would be expected under a given statistical model. These zeros reflect complete reproductive failure – individuals might have insufficient resources to breed, or an entire clutch or brood could be depredated before fledging or germinating. A distribution that is not zero-inflated can still contain many zeroes if it has a low mean (Warton 2005).

Underdispersion and zero-inflation pose major challenges in regression analyses that seek to detect patterns

Manuscript received 4 September 2018; revised 13 February 2019; accepted 21 February 2019. Corresponding Editor: Bruce Kendall.

⁹E-mail: MolBr@dtu.dk

in individual reproduction. Count data can be modeled using generalized linear models (GLMs) or generalized linear mixed models (GLMMs); GLMMs include random effects to account for correlations due for example to repeated observations of the same parent (Bolker et al. 2009, O'Hara and Kotze 2010, Bolker 2015). When individuals can produce multiple offspring at a time, researchers typically assume either Poisson or negative binomially distributed responses. A Poisson distribution assumes that the variance is equal to the mean, while a negative binomial allows the variance to be greater than the mean (i.e., overdispersed). However, the variance of reproduction is often less than the mean (i.e., underdispersion). Ordinal models relax these constraints on the variation (e.g., Hostetler et al. 2012), but they are only feasible when reproductive counts are small because the number of parameters increases with the number of possible outcomes. Quasi-Poisson models are an alternative, which adjust standard errors, Z , and P values of coefficients from a Poisson regression model based on the observed dispersion; however, it is difficult to combine with zero-inflation. Failure to properly account for underdispersion in models leads to erroneously wide confidence intervals and thus statistical tests are more prone to produce false-negatives, i.e., they have reduced power to detect patterns (Harris et al. 2012). In addition, failing to account for zero-inflation results in biased parameter estimates (Harrison 2014).

Kendall and Wittmann (2010) addressed the problem of underdispersion in distributions of annual reproductive success. They showed that offspring number in vertebrates can usually be described with the generalized Poisson (GP) distribution, which is flexible enough to accommodate both under- and overdispersion (Consul and Famoye 1992, Castellanos et al. 2003, Kendall and Wittmann 2010, Harris et al. 2012). Lynch et al. (2014) subsequently demonstrated that the Conway-Maxwell-Poisson (CMP) distribution, which has the same flexibility to model under- and overdispersion, is also well-suited to modeling reproductive data. Although regression models using the GP distribution commonly fail to converge when residuals are underdispersed (Famoye et al. 2004, Sellers and Shmueli 2010, Lynch et al. 2014), models using the CMP distribution have a far better tendency to converge. Failure to converge can occur if computational difficulties are encountered while finding a valid model fit.

Previous studies elegantly demonstrated that distributions of reproductive rates are frequently underdispersed and that the GP and CMP distributions can be used in this scenario. However, dealing with zero-inflation and random effects remained challenging in this context, in part because the computational tools were unavailable. Here, we describe flexible and easy to use methods for modeling reproductive count data that allow for zero-inflation, underdispersion or overdispersion, and random effects. We first describe the implementation of GP and CMP GLMMs in the recently developed `glmmTMB` package (version 0.2.3; Brooks et al. 2017, R Core Team

2019). We then demonstrate the application of these methods with reproductive data from a tropical perennial herb (*Heliconia acuminata*) and a North American songbird (the Orange-crowned Warbler, *Oreothlypis celata*).

METHODS

glmmTMB implementation

In `glmmTMB`, zero-inflated GLMMs have three components: a model for the conditional mean μ , a model for zero-inflation (i.e., the probability of excess zeros relative to the conditional model), and a dispersion model. The conditional mean and zero-inflation models can contain fixed and random effects to account for covariates and correlation. The conditional mean and dispersion are constrained to be positive using log links. The zero-inflation probability is constrained to be between zero and one using a logit link. `glmmTMB` estimates parameters using maximum likelihood estimation via the TMB package (Kristensen et al. 2016). Random effects are estimated using Laplace approximation (Bolker et al. 2009, Kristensen et al. 2016).

Generalized Poisson distribution

The GP distribution can be specified in a `glmmTMB` model using the argument `family = genpois`. Two parameters control the mean (μ) and variance (σ^2):

$$\mu = \frac{\theta}{1-\lambda}, \sigma^2 = \frac{\theta}{(1-\lambda)^3} = \mu \frac{1}{(1-\lambda)^2}. \quad (1)$$

Thus the index of dispersion, which relates the variance to the mean such that $\sigma^2 = \mu\phi^2$, is $\phi^2 = 1/(1-\lambda)^2$ (Consul and Famoye 1992). With this parameterization, ϕ can take any positive value greater than or equal to $\max(1/2, 1-\mu/4)$, so the distribution can be either overdispersed or underdispersed relative to the Poisson. When the GP distribution is underdispersed (i.e., $0 < \phi < 1$) it has an upper bound at $\mu/(1-\phi)$, which is convenient for characterizing reproductive data that has a finite upper bound (Kendall and Wittmann 2010). The estimate of ϕ^2 can be extracted from a fitted model using the function `sigma()`. It is also possible to model ϕ^2 as a log-linear function of covariates (e.g., x) in a `glmmTMB` model using the argument `dispformula = ~x`.

Conway-Maxwell-Poisson distribution

The CMP distribution can be specified in a `glmmTMB` model using the argument `family = compois`. Unlike the GP, the CMP distribution has no closed-form equation for the mean and variance. Approximations exist to allow for regression models on the approximate mean or the mode, and a method of doing regression on the true mean was recently developed (Brooks et al. 2017, Huang 2017). The

approximation to the mean is problematic because it fails when the data are underdispersed: typically the case in studies of reproduction (Shmueli et al. 2005, Lord et al. 2008, Lynch et al. 2014). The approximation to the mode is problematic because results, such as effect sizes, are hard to compare to studies that use Poisson, negative binomial, or GP distributions (Huang 2017). Comparability is essential for making results useful in meta-analyses. Thus, we use the mean-parameterized CMP distribution, which has a mean $\mu \geq 0$ and a parameter to control variation $\nu \geq 0$ (Huang 2017). When $\nu = 1$, the CMP is equal to the Poisson distribution. It is underdispersed when $\nu < 1$, and overdispersed when $1/\nu > 1$. We report results in terms of $1/\nu$, the approximate index of dispersion, i.e. $(\sigma^2)/\mu$ to facilitate comparison to the GP. As with the GP, it is possible to model the index of dispersion as a function of covariates in `glmmTMB`.

Because there is no closed-form equation for the mean and variance of the CMP distribution, estimating mean-parameterized CMP models requires `glmmTMB` to perform an extra numerical step to solve for a normalizing coefficient; this makes estimation considerably slower than models using other distributions in `glmmTMB`. Due to this computational inefficiency, other distributions may be preferable where appropriate, such as the negative binomial distribution for overdispersed data.

Application to inflorescence production

Heliconia acuminata L. C. (Rich.) is a long-lived perennial plant common to the understory of the Amazonian lowland forests. It is a self-incompatible, simultaneous hermaphrodite with almost no asexual reproduction, making it a simple system in which to study reproductive allocation. Recruitment is exclusively via bird-dispersed seeds (Bruna and Kress 2002). Seed dispersal is spatially restricted (Uriarte et al. 2011), providing an opportunity for local adaptation.

Our data were collected at the Biological Dynamics of Forest Fragments Project, in which experimental fragments of Amazonian forest were isolated in the early 1980s (Gascon and Bierregaard 2001). In 1998, Bruna and Kress initiated a long-term study to examine how fragmentation affected the demography of *H. acuminata* (Bruna and Kress 2002). Plots of size 50×100 m were established in four 1-ha fragments and six locations in continuous forest. All *H. acuminata* in each of these plots were marked with a numbered tag and their size was measured in two ways: the height to the tallest leaf and by counting the number of vegetative shoots. Plots were then censused annually to assess individual survival, growth, and inflorescence production. New recruits were also marked and measured. Average plant densities at the study site in 2009 were 892 ± 653 plants/ha (mean \pm SD) in continuous forest and 344 ± 49 plants/ha in fragments. The data set consists of 37,756 measurements across 12 yr and 5,290 individuals (1,075 in fragments and 4,215 in continuous forest), excluding

observations of possibly dormant individuals that produced zero shoots in a year.

We modeled *Heliconia* inflorescence production with the following distributions: Poisson (P), zero-inflated Poisson (ZIP), GP, zero-inflated GP (ZIGP), CMP, and zero-inflated CMP (ZICMP). We fit full models, with each of these distributions, that allowed the conditional mean to vary with the log of plant size (number of shoots), habitat type (fragment vs. continuous forest), and their interaction. Because the conditional mean models use log links, using log plant size as a covariate is equivalent to fitting a power-law (allometric) relationship between conditional inflorescence number and plant size (Bolker 2008). For ZIP, ZIGP, and ZICMP distributions, the full models allow zero-inflation to vary with log plant size, habitat type, and their interaction. For each distribution, we fit all submodels from the full model to represent null hypotheses. We used small-sample-size corrected Akaike information criterion (AIC_c) via the `AICcTab()` function from the `bbmle` package (Burnham and Anderson 2002, Bolker and R Core Team 2017) to compare all submodels across the six distributions. It is necessary to choose the distribution and covariate model simultaneously because different covariate models will result in different residuals and thus different distributions. To demonstrate how to possibly avoid fitting the slower GP, CMP, ZIGP, and ZICMP models, we tested for underdispersion in the residuals of the top Poisson or ZIP model using the `testDispersion()` function from the `DHARMA` package, which tests the quantiles of scaled simulated residuals against a uniform distribution (Hartig 2018). One can also test whether the expected number of zeros based on a fitted model differs from the number of observed zeros, using the `testZeroInflation()` function from the `DHARMA` package (Hartig 2018). R code is available in Appendix S1. Data on *Heliconia* inflorescence production and size is available at https://github.com/embruna/Brooks_etal_2019_HeliconiaData.

Application to warbler fecundity

We obtained data from the Dryad Digital Repository on the fecundity of a population of Orange-crowned Warblers nesting on Santa Catalina Island, California, USA (Sofaer et al. 2014a,b). During 7 yr of observation ($n = 181$), fledglings per nest ranged from 0 to 6 with a median of 0 and a mean of 1.4. The original analysis used a ZIP model to test for effects of local breeding density and precipitation. We tested the original ZIP model for underdispersion using `DHARMA`'s `testDispersion()` function (Hartig 2018). We fit all models described in the original publication with ZIP, ZIGP, and ZICMP distributions (Table 2). These models included all submodels of the full model, which contained main effects of breeding density and precipitation in both the conditional model and zero-inflation model. As in the original study, all models contained a random effect of year in the zero-inflation model. We

compared the models using AIC_c (Burnham and Anderson 2002). For the most parsimonious model chosen by AIC_c , we examined the residuals using the function `testUniformity()` from the DHARMA package (Hartig 2018). R code is available in Appendix S2.

RESULTS

Inflorescence production increases with size, faster in fragments

The most parsimonious model of *H. acuminata* inflorescence production had a ZICMP distribution and was 33.9 ΔAIC_c units better than the most parsimonious CMP model (Burnham and Anderson 2002). The ZIP, GP, and Poisson distributions followed with 49.6, 60.3, and 60.6 ΔAIC_c units respectively for their top models (Table 1; Appendix S3: Table S1). The top ZICMP and ZIP models had the same formulas and similar coefficient estimates (Appendix S3: Fig. S1). All of the ZIGP, two of the ZICMP, and two of the ZIP models failed to converge. The top ZIP model had underdispersed residuals ($P = 0.046$). The most parsimonious model indicated that the conditional mean and zero-inflation both varied with plant size and habitat type. Plants in fragments had higher size-dependent conditional means and thus higher predicted modes and fitted values at a given size compared to plants in continuous forest (Fig. 1).

Warbler fecundity decreases with density and increases with precipitation

The ZIP model from the original publication of this dataset (Sofaer et al. 2014a,b) had underdispersed residuals ($P = 0.003$). The most parsimonious model of warbler fecundity had a ZICMP distribution, which was 4.9 ΔAIC_c units better than the most parsimonious ZIGP and 34.1 ΔAIC_c units better than the most parsimonious ZIP model (Table 2). For models with the ZIP distribution, AIC_c gave the same ranking and almost identical AIC_c values as in the original publication of this data (Sofaer et al. 2014a) (Appendix S3: Table S2). However, models with the ZICMP and ZIGP distribution had a different ranking and lower AIC_c values than those with the ZIP distribution (Table 2). The top five models involved the same predictor variables whether they had ZICMP or ZIGP distributions (Table 2). The best model showed that average fecundity decreased with population density (slope = -0.14 , SE = 0.05) and increased with precipitation (slope = 0.005, SE = 0.002), while nest failure was more likely in years with high population density (slope = 1.66, SE = 0.42) and in drier years (slope = -0.6 , SE = 0.01). The top ZIGP model had the same conditional and zero-inflation models and nearly identical coefficient estimates as the best ZICMP model; the only near exception was the coefficient for the effect of precipitation in the conditional model (slope = 0.006,

TABLE 1. Inflorescence model selection results.

Conditional model	Zero-inflation model	Distribution	$(\sigma^2)/\mu$	k	ΔAIC_c
~ lsize + habitat	~ lsize + habitat	ZICMP	0.48	10	0
~ lsize + habitat	~ lsize \times habitat	ZICMP	0.48	11	2
~ lsize \times habitat	~ 0	CMP	0.58	8	33.9
~ lsize	~ 0	CMP	0.58	6	45.6
~ lsize + habitat	~ lsize + habitat	ZIP	1	9	49.6
~ lsize \times habitat	~ lsize	ZIP	1	9	50.2
~ lsize \times habitat	~ 0	GP	0.9	8	60.3
~ lsize \times habitat	~ 0	P	1	7	60.6
~ lsize	~ 0	P	1	5	72.8
~ lsize	~ 0	GP	1	6	74.7

Notes: Inflorescence production in *Heliconia acuminata* varies with log of plant size (lsize) and habitat type. Each conditional model contained random effects of individual, plot, and year. For brevity, we only present the top two models from each distribution (see Supplementary Material for extended results). Values of $(\sigma^2)/\mu < 1$ indicate underdispersion. The value of $(\sigma^2)/\mu$ is fixed at 1 for the P and ZIP, but estimated by glmmTMB for other distributions and k is the number of parameters. Distributions are ZICMP, zero-inflated Conway-Maxwell-Poisson; CMP, Conway-Maxwell-Poisson; ZIP, zero-inflated Poisson; P, Poisson; GP, generalized Poisson. ΔAIC_c is the change attributable to the model in the Akaike information criterion corrected for sample size.

SE = 0.002), which was still within the CI of the ZICMP estimate.

DISCUSSION

Models with a ZICMP distribution gave the best fits for both warbler fecundity and *Heliconia* inflorescence production, resulting in novel insights about the reproductive biology of both species. For instance, the increased power of the underdispersed distributions allowed us to detect a statistically significant effect of warbler population density on conditional fecundity in addition to its effect on nest failure (i.e., zero-inflation). For *Heliconia*, we found that plants of a given size produce more inflorescences in fragments than in continuous forest, an effect that is detectable whether the models account for underdispersion or not. Although this could be the result of changes in above- and below-ground biomass allocation (Bruna and Kress 2002), it could also provide some of the first evidence to date that plants in fragments may “speed up” their life-history in response to higher risk of mortality as previously theorized (Cohen 1971, Kozłowski and Uchmanski 1987, Jacquemyn et al. 2012).

It is difficult to precisely predict the effects on statistical inference of accounting for underdispersion. Predictions of *Heliconia* inflorescence production do not change much based on assumptions about dispersion in

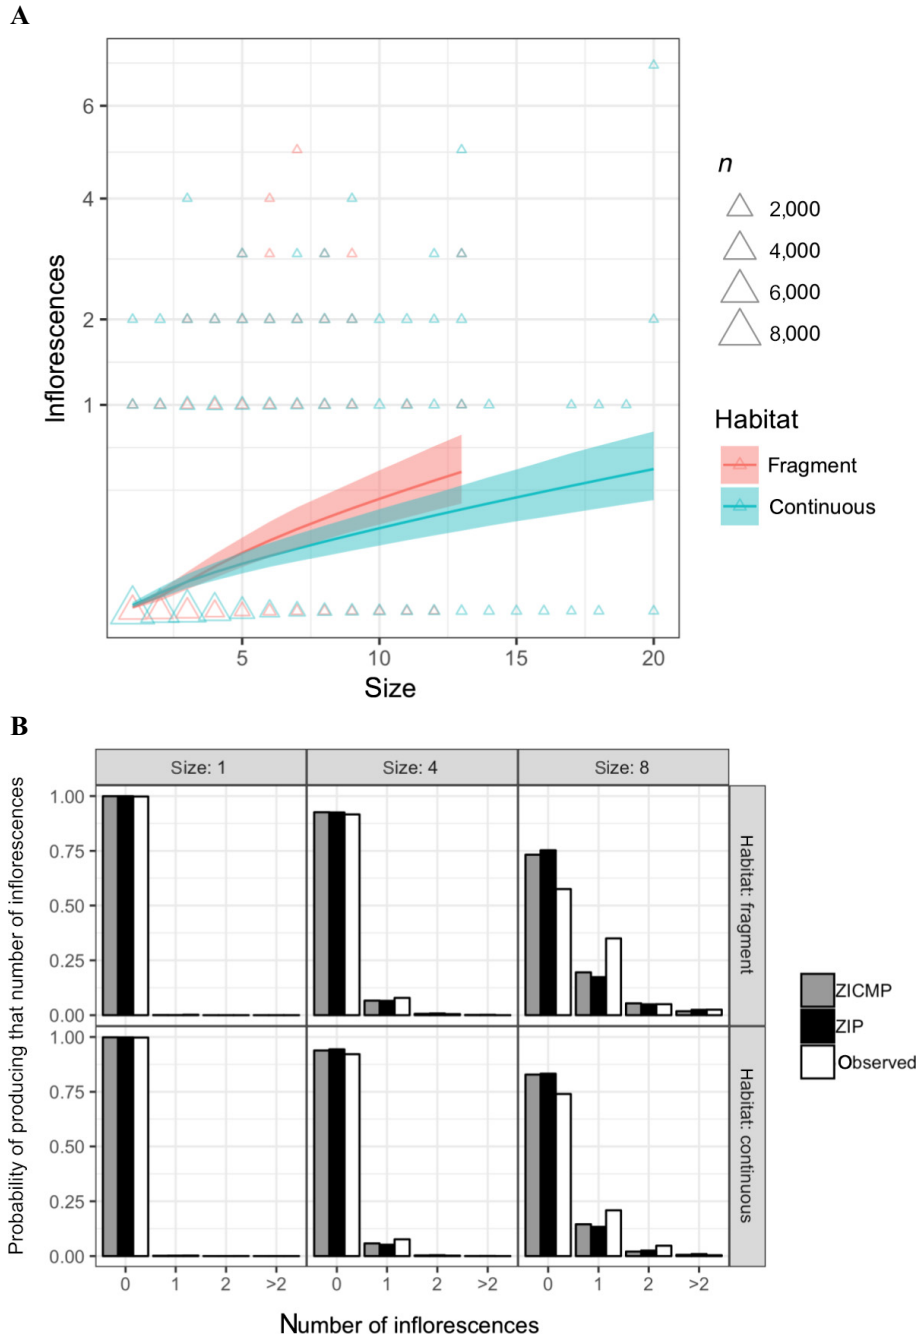


FIG. 1. Inflorescence production in *Heliconia acuminata* depends on plant size (number of shoots) and habitat. (A) Triangles represent observations, with sizes proportional to the number of observations. Lines (and ribbons) represent predicted zero-inflated modes (and 95% CI) from the most parsimonious model. The mode is lower than the mean because random effects are zero at the mode and Jensen’s inequality pulls the mean up; this is true for all GLMMs with log links. (B) Simulations from the most parsimonious model, which had a zero-inflated Conway-Maxwell-Poisson distribution (ZICMP, gray bars), and from the analogous model with a zero-inflated Poisson distribution (ZIP, black bars) are plotted for comparison with the observed data (white bars). Each panel represents a given habitat and size of an individual plant (restricted to a subset of sizes for brevity: 1, 4, and 8). Bar heights represent the probability of producing a certain number of inflorescences, given the size and habitat for that panel. For each model, probabilities were calculated by summarizing 100 simulated data sets of equal size to the original data.

the conditional distribution (Fig. 1B); the probabilities of large numbers of offspring are only slightly smaller when allowing for underdispersion. In contrast, models

of warbler fledglings showed stronger underdispersion (see $(\sigma^2)/\mu$ in Table 2 vs. Table 1); this underdispersion is evident in predicted probabilities of large reproductive

TABLE 2. Warbler model selection results.

Conditional model	Zero-inflation model	Distribution	$(\sigma^2)/\mu$	k	ΔAIC_c
~ bd + precip	~ bd + precip	ZICMP	0.26	8	0.0
~ bd	~ bd + precip	ZICMP	0.27	7	2.1
~ 1	~ bd + precip	ZICMP	0.29	6	4.7
~ bd + precip	~ bd + precip	ZIGP	0.36	8	4.9
~ precip	~ bd + precip	ZICMP	0.29	7	5.6
~ bd + precip	~ bd	ZICMP	0.26	7	7.5
~ bd	~ bd + precip	ZIGP	0.41	7	9.2
~ 1	~ bd + precip	ZIGP	0.42	6	10.8
~ precip	~ bd + precip	ZIGP	0.41	7	11.1
~ bd + precip	~ bd	ZIGP	0.36	7	12.4
~ bd	~ bd + precip	ZIP	1.00	6	34.0
~ 1	~ bd + precip	ZIP	1.00	5	34.5
~ bd + precip	~ bd + precip	ZIP	1.00	7	34.6
~ precip	~ bd + precip	ZIP	1.00	6	36.2
~ bd	~ bd	ZIP	1.00	5	41.8

Notes: Fecundity depends on breeding density (bd) and precipitation (precip). Each zero-inflation model contained a random effect of year. For each distribution, we only present the top five models in the main text (see Supplementary Material for extended results). Values of $(\sigma^2)/\mu < 1$ indicate underdispersion. The value of $(\sigma^2)/\mu$ is fixed at 1 for the ZIP, but estimated by glmmTMB for other distributions; k is the number of parameters.

events (Appendix S3: Fig. S2). Having a low maximum reproductive count relative to the mean may make it more important to account for underdispersion as the index of dispersion will be further below one and the distribution will diverge further from the Poisson. In general, smaller data sets (such as the warbler data set) will have lower power and thus the increased power that comes from accounting for underdispersion may be more important in distinguishing clear effects.

Accounting for underdispersion may be important in making population forecasts based on estimated model parameters by reducing predictions of unrealistically large reproductive events. For example, in warblers, simulations from the top model with a ZICMP distribution are 44 times less likely to produce values above the observed range compared to the same model with a ZIP distribution (respectively, 0.0003 vs. 0.0132 overall probability of simulating more than 6 fledglings).

Models with a ZIGP distribution were estimable for the warbler data set but not for the *Heliconia* data set. It is well known that the GP frequently fails to converge when residuals are underdispersed (Famoye et al. 2004, Sellers et al. 2008, Lynch et al. 2014) as in both of our examples according to tests using the DHARMA package. The higher complexity of the random effects in the *Heliconia* models may have made them harder to estimate than the warbler models. Future studies should address approaches for improved convergence of the GP, but until then the mean-parameterized CMP remains a good alternative.

In the warbler example, ZIGP and ZICMP models gave very similar results; the ranking of the top five models and coefficients from the top model were the same. Note that comparing coefficients across distributions is only possible

with the mean-parameterized CMP (Brooks et al. 2017, Huang 2017). We expect that, when it is possible to fit both distributions to the same data set, results from the GP and CMP will closely match. The main difference between the GP and CMP distributions is that the GP has an upper limit above which the density is exactly zero; this truncation is connected to the convergence problems (Sellers and Shmueli 2010, Huang 2017). It is therefore possible that cases for which truncation matters are exactly the same cases that are numerically problematic, thus preventing a comparison between the GP and CMP. Future studies could test if models consistently agree and if the CMP consistently produces a moderately lower AIC_c , e.g., 4.9 ΔAIC_c units for the warbler data set. If the distributions consistently give the same results, then the relative speed of the GP could outweigh the importance of the small decrease in AIC from the CMP, if the GP model can be fitted.

Beyond reproductive measures, there are many types of count data where the mechanisms that would give rise to Poisson or negative binomial distributions are not particularly compelling (e.g., species abundances, counts of discrete behavioral responses); evaluating underdispersion and considering use of the models discussed here might lead to improvements in data analysis.

CONCLUSION

Distributions of individual reproductive rates are often zero-inflated and underdispersed. When residuals from Poisson models show underdispersion, then either the generalized Poisson or the mean parameterized Conway-Maxwell-Poisson distribution are good choices for use in GLMMs and GLMs. Fitting models with the generalized Poisson will be faster, but with a higher risk of convergence problems.

ACKNOWLEDGMENTS

Many thanks to the biologists who collected data and the developers of the glmmTMB and TMB packages. Thanks to F. Hartig for making the DHARMA package compatible with glmmTMB. Thanks to B. Kendall and two anonymous reviewers for helpful comments on the manuscript. MEB was supported by NSF IGERT#0801544. This is publication number 753 in the BDFFP Technical Series.

LITERATURE CITED

- Bolker, B. 2008. Ecological models and data in R. Princeton University Press, Princeton, New Jersey, USA.
- Bolker, B. M. 2015. Linear and generalized linear mixed models. Pages 309–334 in G. A. Fox, S. Negrete-Yankelevich, and V. J. Sosa, editors. Ecological statistics. Oxford University Press, Oxford, UK.
- Bolker, B. M., and R Development Core Team. 2017. bbmle: tools for general maximum likelihood estimation. R package version 1.0.20. <https://CRAN.R-project.org/package=bbmle>
- Bolker, B. M., M. E. Brooks, C. J. Clark, S. W. Geange, J. R. Poulsen, M. H. H. Stevens, and J.-S. S. White. 2009. Generalized linear mixed models: a practical guide for ecology and evolution. Trends in Ecology and Evolution 24:127–135.

- Brooks, M. E., K. Kristensen, K. J. van Benthem, A. Magnusson, C. W. Berg, A. Nielsen, H. J. Skaug, M. Mächler, and B. M. Bolker. 2017. glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *R Journal* 9:378–400.
- Bruna, E. M., and W. J. Kress. 2002. Habitat fragmentation and the demographic structure of an Amazonian understory herb (*Heliconia acuminata*). *Conservation Biology* 16:1256–1266.
- Burnham, K., and D. Anderson. 2002. Model selection and multimodel inference: a practical information-theoretic approach. Second edition. Springer, New York, New York, USA.
- Castellanos, M. C., P. Wilson, and J. D. Thomson. 2003. Pollen transfer by hummingbirds and bumblebees, and the divergence of pollination modes in *Penstemon*. *Evolution* 57:2742–2752.
- Cohen, D. 1971. Maximizing final yield when growth is limited by time or by limiting resources. *Journal of Theoretical Biology* 33:299–307.
- Consul, P. C., and F. Famoye. 1992. Generalized Poisson regression model. *Communications in Statistics: Theory and Methods* 21:89–109.
- Famoye, F., J. J. Wulu, and K. P. Singh. 2004. On the generalized Poisson regression model with an application to accident data. *Journal of Data Science* 2:287–295.
- Gascon, C., and R. O. Bierregaard, Jr. 2001. The biological dynamics of forest fragments project. Pages 31–42 in R. O. Bierregaard, C. Gascon, T. E. Lovejoy, and R. Mesquita, editors. *Lessons from Amazonia: the ecology and conservation of a fragmented forest*. Yale University Press, New Haven, Connecticut, USA.
- Harris, T., Z. Yang, and J. W. Hardin. 2012. Modeling underdispersed count data with generalized Poisson regression. *Stata Journal* 12:736–747.
- Harrison, X. A. 2014. Using observation-level random effects to model overdispersion in count data in ecology and evolution. *PeerJ* 2:e616.
- Hartig, F. 2018. DHARMA: residual diagnostics for hierarchical (multi-level/mixed) regression models. R package version 0.2.0. <https://CRAN.R-project.org/package=DHARMA>
- Hostetler, J. A., D. P. Onorato, B. M. Bolker, W. E. Johnson, S. J. O'Brien, D. Jansen, and M. K. Oli. 2012. Does genetic introgression improve female reproductive performance? A test on the endangered Florida panther. *Oecologia* 168:289–300.
- Huang, A. 2017. Mean-parametrized Conway–Maxwell–Poisson regression models for dispersed counts. *Statistical Modelling* 17:1–22.
- Jacquemyn, H., L. De Meester, E. Jongejans, and O. Honnay. 2012. Evolutionary changes in plant reproductive traits following habitat fragmentation and their consequences for population fitness. *Journal of Ecology* 100:76–87.
- Kendall, B. E., and M. E. Wittmann. 2010. A stochastic model for annual reproductive success. *American Naturalist* 175:461–468.
- Kozłowski, J., and J. Uchmanski. 1987. Optimal individual growth and reproduction in perennial species with indeterminate growth. *Evolutionary Ecology* 1:214–230.
- Kristensen, K., A. Nielsen, C. W. Berg, H. J. Skaug, and B. M. Bell. 2016. TMB: automatic differentiation and Laplace approximation. *Journal of Statistical Software* 70:1–21.
- Lack, D. 1947. The significance of clutch size. *Ibis* 89:302–352.
- Lord, D., S. D. Guikema, and S. R. Geedipally. 2008. Application of the Conway–Maxwell–Poisson generalized linear model for analyzing motor vehicle crashes. *Accident Analysis & Prevention* 40:1123–1134.
- Lynch, H. J., J. T. Thorson, and A. O. Shelton. 2014. Dealing with under- and over-dispersed count data in life history, spatial, and community ecology. *Ecology* 95:3173–3180.
- O'Hara, R. B., and D. J. Kotze. 2010. Do not log-transform count data. *Methods in Ecology and Evolution* 1:118–122.
- R Core Team. 2019. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. www.R-project.org
- Sellers, K., and G. Shmueli. 2010. A flexible regression model for count data. *Annals of Applied Statistics* 4:943–961.
- Shmueli, G., T. P. Minka, J. B. Kadane, S. Borle, and P. Boatwright. 2005. A useful distribution for fitting discrete data: revival of the Conway–Maxwell–Poisson distribution. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 54:127–142.
- Sofaer, H. R., T. S. Sillett, K. M. Langin, S. A. Morrison, and C. K. Ghalambor. 2014a. Partitioning the sources of demographic variation reveals density-dependent nest predation in an island bird population. *Ecology and Evolution* 4:2738–2748.
- Sofaer, H. R., T. S. Sillett, K. M. Langin, S. A. Morrison, and C. K. Ghalambor. 2014b. Data from: Partitioning the sources of demographic variation reveals density-dependent nest predation in an island bird population. Dryad Digital Repository. <https://doi.org/10.5061/dryad.m4fs7>
- Uriarte, M., M. Anciaes, M. T. B. da Silva, P. Rubim, E. Johnson, and E. M. Bruna. 2011. Disentangling the drivers of reduced long-distance seed dispersal by birds in an experimentally fragmented landscape. *Ecology* 92:924–937.
- van Noordwijk, A. J., and G. de Jong. 1986. Acquisition and allocation of resources: their influence on variation in life history tactics. *American Naturalist* 128:137–142.
- Warton, D. I. 2005. Many zeros does not mean zero inflation: comparing the goodness-of-fit of parametric models to multivariate abundance data. *Environmetrics* 16:275–289.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at <http://onlinelibrary.wiley.com/doi/10.1002/ecy.2706/supinfo>