

Methodology and Research Practice

Poisson Regressions: A Little Fishy

William H. Ryan¹ ^a, Ellen R. K. Evers¹, Don A. Moore¹

¹ Haas School of Business, University of California, Berkeley, CA, US

Keywords: poisson regression, false positives, type i errors, methodology, fish, statistical methods

<https://doi.org/10.1525/collabra.27242>

Collabra: Psychology

Vol. 7, Issue 1, 2021

When analyzing count data (such as number of questions answered correctly), psychologists often use Poisson regressions. We show through simulations that violating the assumptions of a Poisson distribution even slightly can lead to false positive rates more than doubling, and illustrate this issue with a study that finds a clearly spurious but highly significant connection between seeing the color blue and eating fish candies. In additional simulations we test alternative methods for analyzing count-data and show that these generally do not suffer from the same inflated false positive rate, nor do they result in much higher false negatives in situations where Poisson would be appropriate.

Researchers often analyze count data, such as number of multiple-choice questions answered correctly, the number of different uses of a brick, or the number of pins stuck into a Voodoo doll. Because these data are integers censored at 0, researchers will sometimes use Poisson regressions to examine differences between groups. Poisson regressions test for differences in count data; however, they rely on strong assumptions about the underlying distribution. In particular, Poisson regressions assume that the variance of the distribution is equal to its mean. When the variance is higher than the mean (referred to as data being overdispersed), the risk of false positives increases. We find a) that data in many papers employing Poisson regressions violate these assumptions, and b) that even relatively small violations of Poisson's assumptions can dramatically inflate false positive rates. We demonstrate this issue first in a preregistered study showing that using Poisson regression gives the improbable result that blue shirts prime people to eat more Swedish Fish. We then report a simulation study that shows that under the null, Poisson regressions result in significant p -values more often than the 5% of the time which it should. Additionally, we demonstrate that alternatives to Poisson not only don't lead to significantly more Type I errors when there is no true difference, they also don't lead to significantly more false negatives when the groups actually differ.

While there are papers in specialized journals discussing the risks of using Poisson regressions when its assumptions are not met (e.g., Cox et al., 2009), it appears that many

authors and editors are not aware of these dangers. We find evidence that incorrect use of Poisson is widespread. A review of the *Journal of Personality and Social Psychology* found 18 papers using Poisson regression to analyze count data in the past 10 years; of these 18 papers, 9 appear to have used it incorrectly; using Poisson on data in which the variance is not equal to the mean.¹ Review of two additional top Psychology journals and a top Marketing journal found instances of incorrect use of Poisson in all of them. Based on the simulations presented later, data with overdispersion equivalent to that found in several of these papers can be expected to lead to false positive rates of up to 60%. While the absolute number of papers using Poisson is not large, we nevertheless believe it is important to correct the published literature and do what we can to prevent future publication of false positive results. In addition, based on our findings, we also hope that this paper prevents authors of meta-analyses from unquestioningly including results from Poisson regressions, since, as we will see, these results should be questioned and can dramatically affect meta-analytical estimates. We worry that Poisson analyses may be selected, at least sometimes, precisely because they make the result appear more impressive or statistically significant. In addition, unlike other, more familiar, forms of p-hacking such as conducting multiple tests, pre-registration does not prevent inflated rates of false positives when Poisson regressions are used with overdispersed data.

It is also important to remember that we were only able to review a few of the hundreds of journals in Psychology,

a wryan@berkeley.edu

¹ To carry out this search we searched for all instances of the string "Poisson" in the relevant journal's archives from 2008-2018 (journals were reviewed in November-December 2018). We additionally reviewed the *Journal of Experimental Psychology: General*, the *Journal of Personality and Social Psychology*, and the *Journal of Consumer Research*. These journals were selected because they are among the most read and cited journals in Psychology and Marketing respectively. Articles published here are widely read and cited, making incorrect use of Poisson here particularly impactful. Papers were included if at least one of their analyses used Poisson regression, and another method of analysis was not used as well for the same result. Papers which e.g. reported both Poisson and Negative Binomial regressions on the same data are therefore excluded from our search.

and there's little reason to believe that those journals are exceptional in their use of Poisson. Further, the rate of Poisson use should most appropriately be compared to the number of papers that were analyzing count data and thus could have used Poisson, not to the total number of papers published. While it is difficult to determine the exact number of papers using count data, we conducted a conservative test of the proportion of count data papers which may be using Poisson by looking within a specific literature, papers using the voodoo doll task.² Researchers use the voodoo doll task to assess aggression by measuring the number of pins a participant sticks into a doll representing another person, resulting in count data. The paper that introduced this task explicitly recommended against using Poisson regressions (Dewall et al., 2013). Despite this, of 26 papers featuring the voodoo doll task, one fifth of them used only Poisson regressions. Of those which shared sufficient data to evaluate means and variances, we could find only one in which Poisson may have been appropriate – in the others its use likely inflated the risk of false positives. In short, within the literature employing studies using count data as dependent variable, a substantial proportion of papers appears to rely on Poisson distributions to analyze this data, even in many cases where it may yield increased false positives. Future work should employ alternative methods of analysis.

It may not be surprising that so many are unaware of the weaknesses of Poisson. Many analytic approaches are robust to violations of their assumptions. We are used to safely treating Likert scales as continuous data, not running Levene's test for ANOVAs, or using linear regression on integer data. Poisson regression is a stark exception. Authors, editors, and reviewers either do not realize the limitations of Poisson or they do not realize how dangerous it is to violate its assumptions. This paper aims to raise awareness of this issue by giving a clear and memorable demonstration of Poisson's vulnerabilities and offering useful alternatives. Next, we discuss the distributions involved in more detail. We begin by explaining why linear regression may seem inappropriate for count data, and then explain the assumptions of Poisson regressions, Negative Binomial regressions, and permutation tests. Readers who are not interested in a review of these distributions can proceed directly to the "Experiment" section of the paper.

Analyses for Count Data and Their Assumptions

Many scholars in behavioral science analyze their experimental data using linear regressions.³ Linear regression assumes continuous dependent variables which can hold values below zero. To see how these assumptions might be violated with count data, imagine that we are predicting the number of croissants eaten per day. A linear regression

assumes that the values are drawn from a continuous and symmetrical distribution around a mean, meaning that this model might predict that someone eats 1.87 croissants, or even somehow consumes a negative number of croissants. Further, linear regression assumes homoskedasticity – that the error term is similar across values of the independent variable, an assumption which count data frequently (though not always) violates. In short, count data exhibits characteristics that clearly violate assumptions of linear regressions and based on this it may appear that using a linear regression to analyze count data may be inappropriate. However, even though Poisson regressions do not assume the same characteristics of the underlying data as linear regressions, Poisson regressions bring a new set of assumptions to the table, assumptions that, as we will see, are often violated. Thus, to decide which test to use on count data, we need to explore the assumptions of different tests and investigate the consequences of violating these assumptions.

Poisson Regressions

Poisson regression is a form of the generalized linear model which accommodates non-normal distributions of the dependent variable, and instead assumes that the dependent variable has a Poisson distribution. This distribution expresses the probability that a given number of events will occur in a fixed interval, assuming that these events occur at a known constant rate on average and that each event is independent of the others. For example, Poisson might model the number of fish you catch in the Seine after fishing for a given period of time, provided the rate at which you catch fish is constant over time. Constant rate processes like these naturally have identical means and variances, so the distribution has a single parameter, λ , which is equal to both its mean and variance. This rigidity can cause problems when Poisson distributions are fit to data which didn't actually come from a Poisson process, where mean and variance are often not actually equal. [Figure 1](#) presents examples of Poisson distributions with varying λ fit to data with increasingly high ratios of variance to mean. As the figure shows, Poisson distributions fit to these data correspond poorly to overdispersed data. When fitting a Poisson distribution to data like this, the λ parameter is equal to the mean of the data being fit to, and then both the mean and the variance of the Poisson distribution are set equal to λ . When the variance is actually higher than the mean, fitted Poisson distributions will have lower variance than the actual data. Additionally, as outliers disproportionately affect the mean of the sample, they can strongly affect the distribution being imposed. In short, unlike linear regressions, Poisson regressions allow for dependent variables to be censored at zero, and allow for non-continuous depen-

2 The voodoo doll task papers were reviewed in April–June 2019. The voodoo doll task was selected because it is one of the subfields in which count data is most widely used. We analyzed all papers citing the paper introducing the voodoo doll task which also used the task, and then coded them for if they used Poisson regressions alone to analyze their data or not.

3 Linear, negative binomial, and Poisson regression are all versions of the generalized linear model. Readers interested in more detail on how these regressions relate to one another can refer to Gardner et al. (1995).

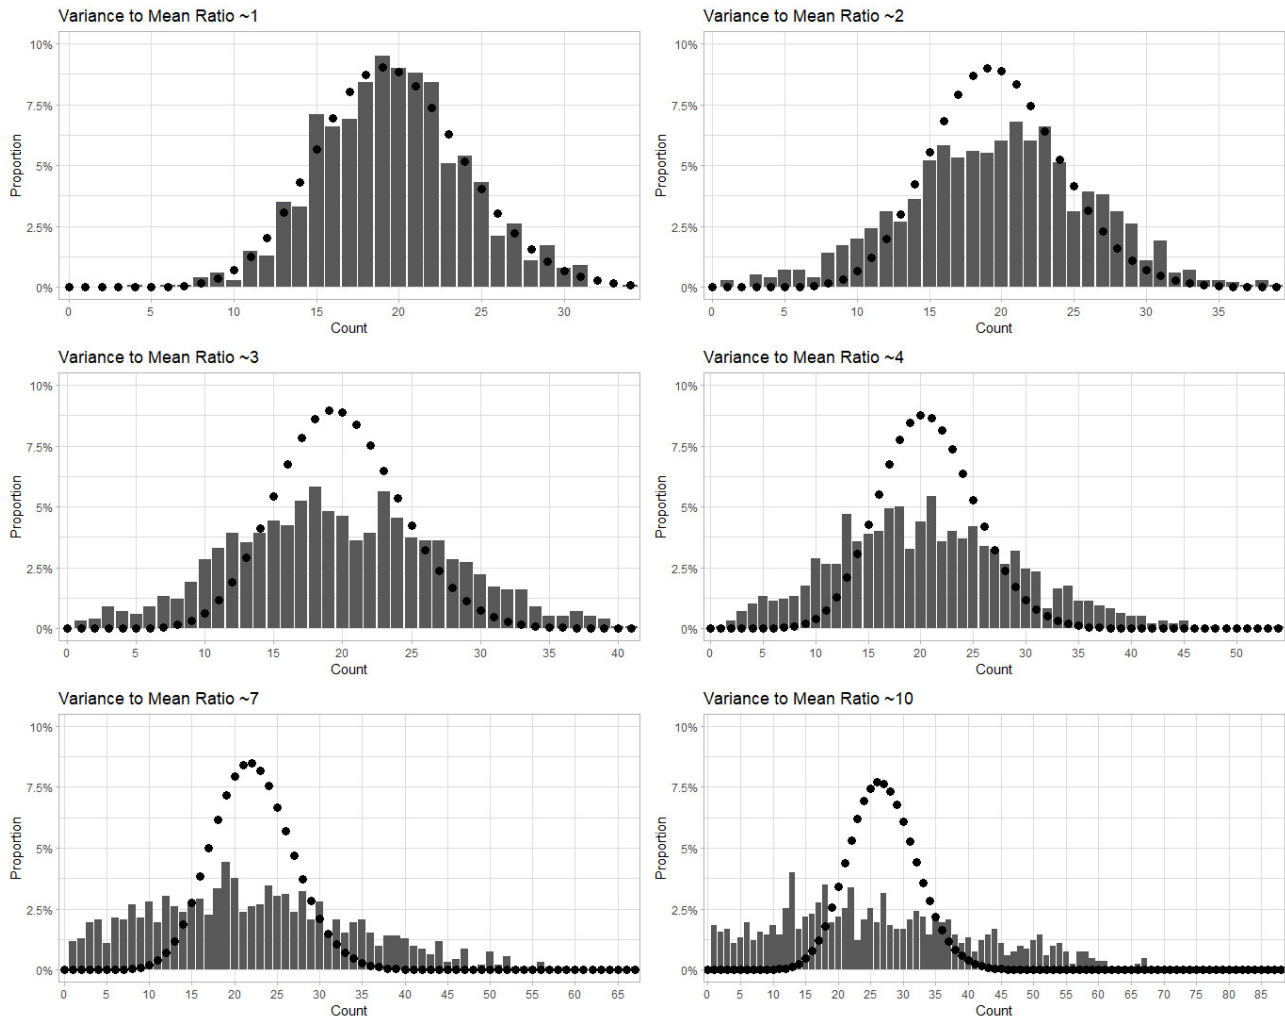


Figure 1. Illustrates actual data with increasing variance to mean ratios (bars), as well as the Poisson distribution fit to that data (black dots).

dent variables, but at the same time, they make the strong assumption that the variance is equal to the mean.

Negative Binomial Regression

Like Poisson regression, Negative Binomial regression is another form of the generalized linear model that accommodates non-normal distributions of the dependent variable, but it assumes that the dependent variable takes a different form (for further discussion, see Cox et al., 2009). The Negative Binomial distribution represents the number of successes in a sequence of identical, independent Bernoulli trials before a given number of failures occur. For example, a Negative Binomial distribution could estimate the number of times you can roll a die before seeing a “3” four times. It has two parameters: r , the number of failures before stopping, and p , the probability of success in each trial. For our purposes, it offers a useful feature: Instead of requiring the mean equal the variance, the variance is a quadratic function of the mean and can differ from it. This allows it to more faithfully model data which is overdispersed – that is, where the variance is greater than the mean. While it is pos-

sible for data to become so overdispersed that they violate the assumptions of the Negative Binomial, there is more flexibility than Poisson accommodates. [Figure 2](#) illustrates this, by fitting Negative Binomial distributions to the same sets of data as in [Figure 1](#). As the figure shows, Negative Binomial regressions can better account for the increased variance up to a much higher variance to mean ratio, although eventually negative binomial distributions also begin to fit the data poorly if the data has such a high variance to mean ratio that it cannot be captured well by the quadratic relationship Negative Binomial assumes.

Permutation Tests

Finally, permutation tests are non-parametric, and do not impose any distributional assumptions on the data. To carry out a simple type of permutation test, we simply repeatedly randomly shuffle the data between conditions and record the percentage of mean differences that are found which are greater than the one actually observed. This percentage is essentially a p -value – it tells you how frequently a value equal or greater to your observed mean difference

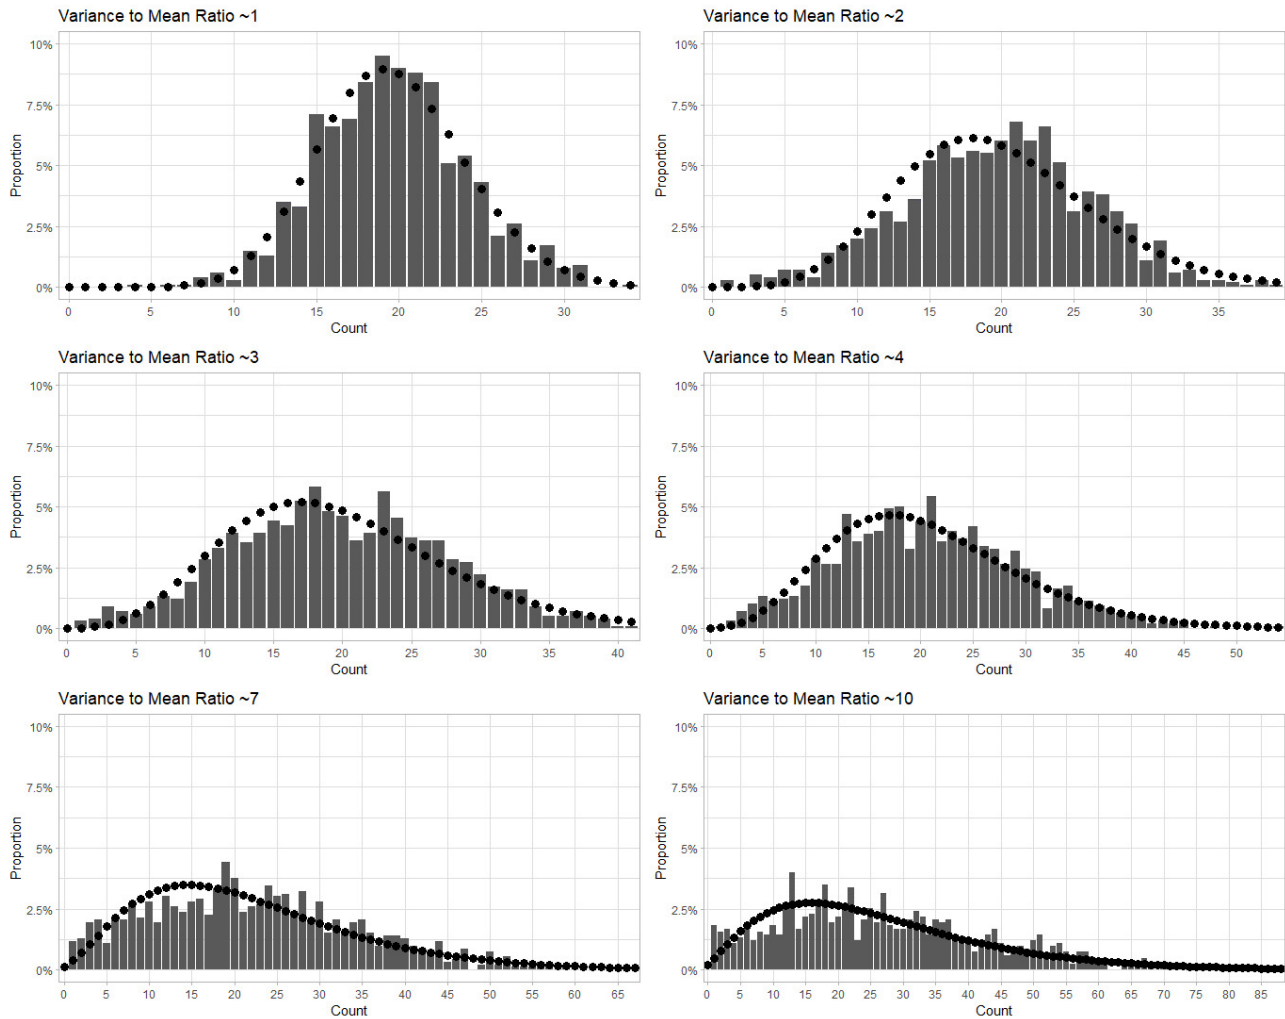


Figure 2. Illustrates actual data with increasing variance to mean ratios (bars), as well as the Negative Binomial distribution fit to that data (overlaid points). The Negative Binomial distribution generally provides a better fit to the data than the Poisson distribution does as variance to mean ratios increase, although with a high enough variance to mean ratio it will begin to fit the data more poorly.

would be found if the null hypothesis was true. Because this makes no distributional assumptions about the data, only about the randomization process, permutation tests are often more robust than regressions when the exact distribution of the data is not known.

In short, when handling count data, different analyses are available, all with specific assumptions about the distribution of the underlying data, that are likely to be violated to some degree. To decide which analysis to use, we need to know the consequences of violating the specific assumptions of these analyses. Frequently Poisson regressions are used, since there are some ways count data corresponds well to its assumptions – data can be censored at 0, and can be non-continuous. However, if the ways that count data do not meet its assumptions – e.g. that data should have a variance equal to the mean – counteracts these benefits, then scholars are erroneously turning to an analysis tool that may make their conclusions less rather than more reliable. This is what we investigate in our experiment and simulations.

Our experimental analysis and simulations focus on the above three regressions: linear, Poisson, and Negative Binomial, as well as permutation tests, but additional options exist. For example, quasi-Poisson regressions adjust statistical results to accommodate overdispersion. A zero-inflated model can accommodate a large number of zero responses. We discuss these alternate models, as well as formal statistical tests for overdispersion which can help you determine which model to use, in the Appendix.

Experiment

To demonstrate how the improper use of Poisson regression can lead to misleading results, we set out to test the highly implausible theory that seeing a blue shirt might prime thoughts of water, thereby affecting their consumption of Swedish Fish gummy candies. We pre-registered opposing hypotheses that the color of the experimenter's shirt would either increase or decrease Swedish fish consumption (the preregistration can be found at <http://aspredicted.org/>

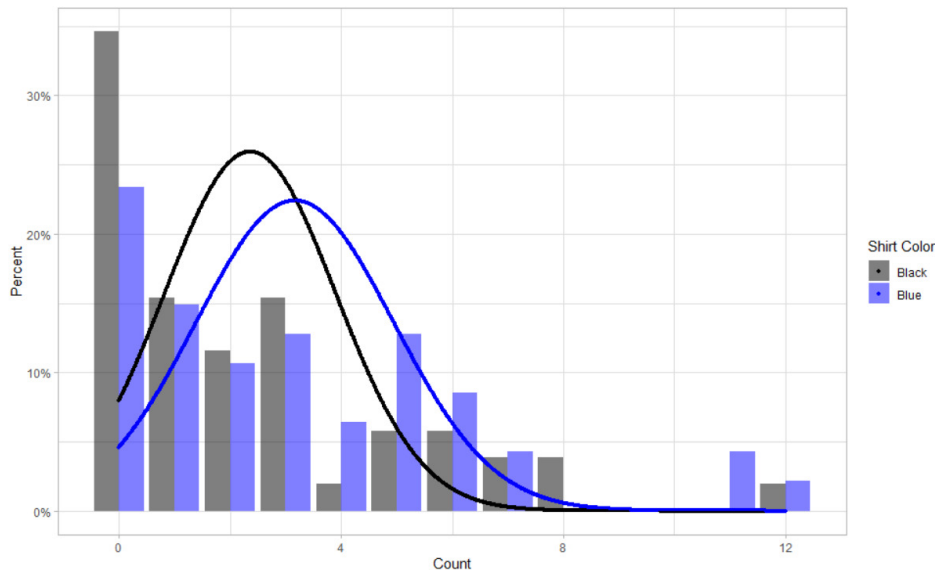


Figure 3. Histogram of the percentage of participants choosing different numbers of fish for both the blue and black shirt conditions. Overlaid on this histogram are two Poisson distributions, which are each fit to the underlying data.

[blind.php?x=di2sg4](https://osf.io/kcgjb/view_only=9282cabbc9044bdbbf101cd87e4a6f6e)). Materials, data, preregistrations, and all code for experimental analysis, simulations, and tutorials can be found on our OSF page: https://osf.io/kcgjb/view_only=9282cabbc9044bdbbf101cd87e4a6f6e.

Method

We approached participants on Sproul Plaza at the University of California, Berkeley ($N = 99$, 51 females, 46 males, $M_{age} = 21.6$)⁴ and asked them to complete a survey for an unrelated study. After completing the survey, participants chose how many Swedish Fish they would like as compensation for their participation, and learned the experimenter would give them whatever number they requested. Our experimental manipulation varied the color of the experimenter's shirt, which was either black or blue. Our randomization procedure had experimenters switch shirts every 15 minutes. The sample size was based on the requirements of the unrelated study, but we believe it reflects common sample sizes in Psychology.

Results

A Poisson regression revealed that participants in the blue shirt conditions ate significantly more Swedish fish ($M = 3.2$, $SD = 3.1$) than participants in the black shirt condition ($M = 2.4$, $SD = 2.7$, $p = .016$). Unlike the Poisson regression, alternative methods of analysis do not reveal a significant difference ($p > .17$ for ANOVA and Negative Binomial regression).

Why did our Poisson regression reveal significant differ-

ences while the other methods of analysis did not? A crucial assumption of Poisson distributions is that the variance is equal to its mean. Since variance is the standard deviation squared, our data was overdispersed by a factor of about three and thus violates the assumption of Poisson distributions. We show how this can lead to an incorrect fit in [Figure 3](#), which shows histograms of the number of fish taken with the Poisson distributions fit to the data overlaid. This graph makes it evident that the Poisson distributions don't appear to fit the data well – for example, they predict that fewer people will choose to take zero fish, and that more will prefer the mean number of fish, than we observe in the data. The Poisson regression also artificially constrains the variance, making the two distributions appear more different than they actually are.

[Figure 4](#) shows the same histogram, but this time with Negative Binomial distributions fit to the data instead. As the figure makes clear, Negative Binomial regression better captures some aspects of the data, such as the larger percentage of people choosing zero fish, and the increased variance of the data, with some participants choosing large numbers of fish.

We can extend this observation: if overdispersion is what caused t-shirt color to apparently affect Swedish fish consumption, increasing the variance-to-mean ratio of our data even more should lead to even “stronger” effects. In general, multiplying all values by a constant will increase the variance to mean ratio. For example, we might have measured the amounts of rewards taken in grams (5/fish) or in Calories (8/fish) – seemingly inconsequential measure-

⁴ Due to experimenter error, for 4 participants data on condition or fish taken was not recorded, and were necessarily excluded from analysis. For 2 additional participants where condition and fish taken was recorded, age was not, and these participants were not excluded.

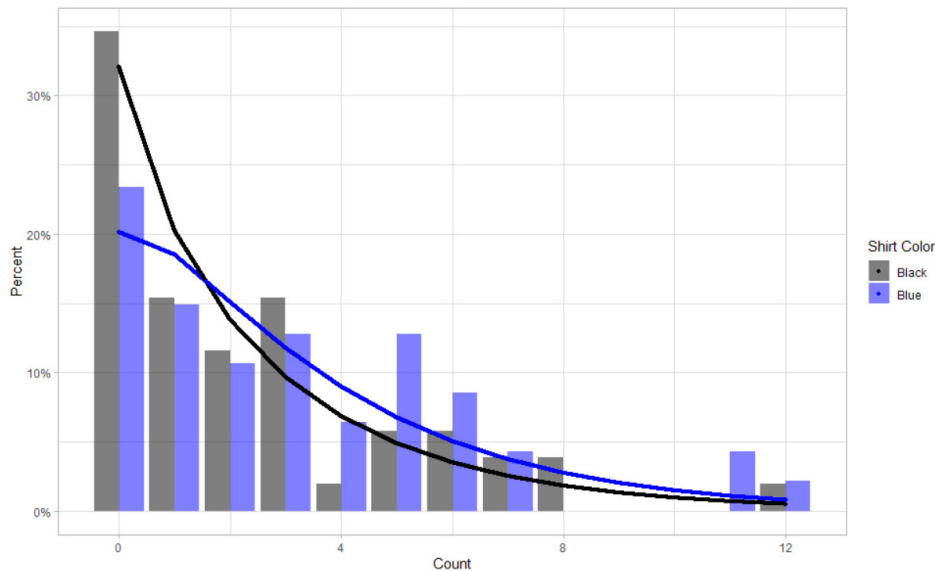


Figure 4. Histogram of participants choosing different numbers of fish for the black and blue shirt conditions. Overlaid on this histogram are two Negative Binomial distributions fit to the underlying data.

ment changes which we will find actually have consequential effects.⁵ The resulting means and variances for each of these counterfactual rewards, as well as the p -values of the resulting Poisson, Negative Binomial, and Linear Regressions, as well as permutation tests, appear in [Table 1](#). As expected, this switch leads to even more significant results when analyzing the data using Poisson regression. Switching from single fish to grams results in $p < 10^{-7}$, switching to Calories in $p < 10^{-11}$. As before, the results were not statistically significant when we repeated this analysis with OLS regression, Negative Binomial regression, or a permutation test⁶ of the difference between conditions.⁷ In short, by arbitrarily changing the expression of the data, without actually changing the actual distribution of the data between conditions, our two conditions appeared to become more significantly different from each other.

Discussion

Our results demonstrate the potential for Poisson regression to exaggerate the statistical significance of small differences in a study testing an implausible hypothesis. In addition, changing the unit of analysis from fish to grams to calories affected the degree to which our statistical analysis claimed our two groups to differ. For this illustration we opted to look at the role of overdispersion by simply multi-

plying our data to investigate the effects of giving out fish in units of five or ten, but in practice researchers face many decisions that may all increase overdispersion such as the number of multiple-choice questions to include in an experiment, or how to score them. We next turn to simulations to assess the degree to which violations of the assumptions of the Poisson distribution increase the risk of false positives, examining the role of overdispersion more comprehensively. We find that even minor violations result in unacceptable increases in false positives. Further, we show that alternatives such as negative binomial or linear regression yield fewer false positives.

Simulated Data

We generated count data randomly drawn from a Poisson distribution. We then modified the simulated data in ways that result in data which violates the assumptions of Poisson (e.g., making each answer count for three points instead of one, or increasing the range of a response scale from seven to 100), and show that it increases the false positive rate. These simulations employ a method known as permutation testing. Essentially, we randomly sample two groups of data points from a Poisson distribution, and then transform the data to give ever-greater violations of Poisson's assumptions. We then use Poisson regressions to test whether

⁵ Although these are no longer technically count data, they still share its critical features: they are integers censored at zero, and thus violate the assumptions of linear regression just as count data does.

⁶ We carried out our permutation test by analyzing participants as though they had been randomly assigned to a group (keeping the overall size of the groups the same) and calculating the difference in means under this random assignment. We repeated this 10,000 times, and then found the percentage of the time that a difference in means greater than that observed in the data was generated by this random shuffling. This percentage was effectively our p -value.

⁷ When individuals 2 standard deviations or more from the mean excluded ($N=4$), no analysis is significant initially, though the Poisson regression still gives the lowest p value ($p = .17$). When grams or Calories are used, the Poisson regression gives a significant result, while no other analysis does except for a Negative Binomial regression using a DV of Calories.

Table 1. Results of Study 1 with Varying Dependent Variable Units of Measurement

Unit	Blue shirt condition		Black shirt condition		p-values			
	M	Var	M	Var	Poisson regression	Negative Binomial regression	Linear regression	Permutation test
Fish	3.2	3.1	2.4	2.7	$p = .02$	$p = .20$	$p = .17$	$p = .17$
Grams (5/fish)	15.9	236.2	11.8	187.3	$p < .001$	$p = .32$	$p = .17$	$p = .17$
Calories (8/fish)	25.4	604.7	18.9	479.4	$p < .001$	$p = .35$	$p = .17$	$p = .17$

the two groups are statistically different from each other. We repeat this process many times, and then record the percentage of the time that the Poisson regression returns a p -value less than .05, indicating statistical significance.

Type I Error Simulations

Our permutation tests examining false positive rates sample both groups of data points from an identical distribution.⁸ Since in this case there is no real difference in the underlying distribution of the two groups, Poisson regressions should result in p -values $< .05$ in only 5% of the simulations (a 5% false positive rate).

Besides a Poisson regression, we also used a Negative Binomial, a non-parametric permutation test, and a linear regression⁹ on the same data so we can compare the false positive rates. We repeated this process 10,000 times for each distribution, to determine how frequently each test returned a p -value below .05. Because the datasets were all drawn from the same underlying distribution, we should expect any properly functioning statistical test to reject the null hypothesis (i.e., find a p -value $< .05$) only 5% of the time. Results of these simulations are plotted in [Figure 5](#). Provided the underlying data are truly Poisson-distributed with a variance to mean ratio of 1, all four tests do equally well with a false-positive rate of 5%. However, when we increased the variance to mean ratio of the underlying data, Poisson distributions resulted in significant p -values (i.e., false positives) more often, over 50% of the time for variance to mean ratios 10 and above. Despite being more robust to slight overdispersion, negative binomial regression also yields increased false positive ratios at high variance to mean ratios.

Count data may be overdispersed naturally. However, there are also many ways that the choices made in research design and analysis affect variance to mean ratios. For example, changing the units data is measured in as in the Swedish Fish study, adding additional easy questions which everyone will get correct, changing the number of options

on a response scale, and more can all change the variance to mean ratio. The wide variety of data and task design choices which can lead to violations of the assumptions of Poisson make it critical to always be aware of the danger of overdispersion.

Type II Error Simulations

We address the risk of false negatives in our next set of simulations. We again ran a permutation simulation, but in this case we simulated data from two Poisson distributions that *are* different (varying the degree of difference plotted on the x-axis in [Figure 6](#)), using 1,000 iterations. When the assumptions of the Poisson are met, Poisson regressions have a rate of Type II errors that is functionally identical to the three alternative methods. We show additional simulations with varying sample sizes in the online experiment code, but results remain similar. In other words, while Poisson regressions (compared to alternative methods) may increase the risk of false positives, not using Poisson regressions does not appear to increase the chance of getting false negatives.

Because the violation of Poisson assumptions increases the likelihood that Poisson analyses will find significant differences, Negative Binomial regressions, permutation tests, and t -tests will lead to higher rates of false negatives when Poisson's assumptions are violated. However, while Poisson regressions do give fewer false negatives when the assumptions are violated, this is because they are generating more false positives in general. For an extreme example of why more false positives means fewer false negatives, consider a hypothetical statistical test which assumed that 100% of relationships were significant. This test would never generate false negatives, but it would be at the cost of a 100% false positive rate when in reality there are no differences. Given the harmful consequences of false positives (Pashler & Harris, 2012), we don't believe this is a valid reason to use Poisson distributions when the data violate the assumptions.

⁸ For all the graphs shown here a Poisson distribution with a $\lambda = 1$ was used, but testing with λ ranging from 1-10 shows no difference in results

⁹ This is equivalent to a t -test in this case




Figure 5. Simulation results plotting the proportion of significant results when in reality there is no difference between the two groups, using a simulated sample size of 100 per cell.




Figure 6. Simulation results showing the proportion of non-significant results when in reality the two groups differ, as a function of the size of the difference (Cohen's d) and conditions of the Poisson regression are met, using a simulated sample size of 100 per cell.

General Discussion

In one experiment and a set of simulations, we find that the use of Poisson regressions can inflate the risks of false positives. Poisson regressions are only appropriate when the data comply with a restrictive set of assumptions. Vi-

olating these assumptions can substantially increase the risk of false-positive statistical results, as we show. These false-positive results are likely to have low p -values, which can bias some types of meta-analyses. For example, p -curve analyses will be thrown off by the addition of low p -values, and may exaggerate evidentiary strength (Simonsohn et al.,

2014). Other meta-analyses may be skewed by inflated effect size estimates (Vosgerau et al., 2019).

Taken as a whole, our results suggest that Poisson results are often fishy. Unless the assumptions of Poisson regression are fully satisfied, it should be avoided. Fortunately, linear regressions and permutation tests offer good alternatives to Poisson regression. Negative Binomial regressions are more broadly useful than Poisson regressions but can also yield false positives when variance to mean ratios become very high. In particular, linear regressions are robust to violations of their assumptions, and permutation tests are robust by virtue of not making assumptions about the distribution of the data to begin with. Appendix I provides additional data on distributions, tests for determining if data are overdispersed, carrying out alternative analyses, and links to further resources for using alternative methods. Tutorial code carrying out these analyses can be found at the paper's OSF page: https://osf.io/kcgjb/?view_only=9282cabb9044bdbbf101cd87e4a6f6e. A simple initial test for overdispersion is to compare the mean of the data with its variance (the standard deviation squared) – if the variance is greater than the mean, overdispersion may render Poisson inappropriate.¹⁰ When interpreting papers presenting Poisson analyses of count data, there is no simple heuristic to adjust their results, and it is likely better to omit papers using Poisson from meta-analyses lest they skew results.

Author Contributions

WH Ryan and ERK Evers designed the experimental study and gathered data. WH Ryan analyzed the experimental data and designed and carried out the simulation study,

with input from ERK Evers and DA Moore. WH Ryan and ERK Evers drafted the manuscript, and DA Moore provided critical revisions. All authors approved final version of the manuscript for submission.

Acknowledgements

Thanks to Stephen Baum and Ekaterina Goncharova for their assistance collecting data for the study, as well as Amelia Dev, Andrew Zheng, Winnie Yan, Mitchell Wong, and Maya Shen for their assistance conducting the literature reviews. Additional thanks to Stephen Antonoplis, Kristin Donnelly, Yoel Inbar, Alex Park, and Sydney Scott (in alphabetical order) for reviewing an early version of this manuscript, and to Kristin Donnelly for suggesting the hypothesis for the illustrative study. Thanks also to anonymous reviewers for their useful comments.

Competing Interests

There are no conflicts of interest to report. Don Moore was an editor at Collabra: Psychology. He was not involved in the review of this manuscript.

Data Accessibility Statement

All relevant materials, participant data, preregistrations, and analysis scripts can be found on this paper's Open Science Foundation project page at https://osf.io/kcgjb/?view_only=9282cabb9044bdbbf101cd87e4a6f6e.

Submitted: July 24, 2020 PDT, Accepted: July 25, 2021 PDT



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-4.0). View this license's legal deed at <http://creativecommons.org/licenses/by/4.0> and legal code at <http://creativecommons.org/licenses/by/4.0/legalcode> for more information.

¹⁰ Even if this ratio is only slightly positive, this should be cause for concern and the use of either an alternative to Poisson or the tests for overdispersion given in the Appendix code

REFERENCES

- Cameron, A. C., & Trivedi, P. K. (1990). Regression-based tests for overdispersion in the Poisson model. *Journal of Econometrics*, *46*(3), 347–364. [https://doi.org/10.1016/0304-4076\(90\)90014-k](https://doi.org/10.1016/0304-4076(90)90014-k)
- Cameron, A. C., & Trivedi, P. K. (2001). Essentials of Count Data Regression. In *A Companion to Theoretical Econometrics* (pp. 331–348). Blackwell Publishing Ltd. <https://doi.org/10.1002/9780470996249.ch16>
- Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics*. Cambridge University Press.
- Cox, S., West, S. G., & Aiken, L. S. (2009). The Analysis of Count Data: A Gentle Introduction to Poisson Regression and Its Alternatives. *Journal of Personality Assessment*, *91*(2), 121–136. <https://doi.org/10.1080/0223890802634175>
- Dewall, C., Finkel, E., Lambert, N., Slotter, E., Bodenhausen, G., Pond, R., Renzetti, C., & Fincham, F. (2013). The voodoo doll task: Introducing and validating a novel method for studying aggressive inclinations. *Aggressive Behavior*, *39*(6), 419–439. <https://doi.org/10.1002/ab.21496>
- Gardner, W., Mulvey, E. P., & Shaw, E. C. (1995). Regression analyses of counts and rates: Poisson, overdispersed Poisson, and negative binomial models. *Psychological Bulletin*, *118*(3), 392–404. <https://doi.org/10.1037/0033-2909.118.3.392>
- Hartig, F. (2016). *DHARMA: Residual diagnostics for hierarchical (multi-level/mixed) regression models* [R Package version 0.4.1.]. <https://CRAN.R-project.org/package=DHARMA>
- Jackman, S. (2020). *pscl: Political Science Computational Laboratory* [R package version 1.5.5]. <https://cran.r-project.org/package=pscl>
- Kleibler, C., & Zeileis, A. (2008). *Applied Econometrics with R*. Springer-Verlag. <https://CRAN.R-project.org/package=AER>
- Pashler, H., & Harris, C. R. (2012). Is the Replicability Crisis Overblown? Three Arguments Examined. *Perspectives on Psychological Science*, *7*(6), 531–536. <https://doi.org/10.1177/1745691612463401>
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-Curve: A Key to the File Drawer. *Journal of Experimental Psychology: General*, *143*(2), 534–547. <https://doi.org/10.1037/a0033242>
- Ver Hoef, J. M., & Boveng, P. L. (2007). Quasi-Poisson vs. Negative Binomial Regression: How Should we Model Overdispersed Count Data? *Ecology*, *88*(11), 2766–2772. <https://doi.org/10.1890/07-0043.1>
- Vosgerau, J., Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2019). 99% impossible: A valid, or falsifiable, internal meta-analysis. *Journal of Experimental Psychology: General*, *148*(9), 1628–1639. <https://doi.org/10.1037/xge0000663>

Online Appendix I: Additional Information on Regressions and Statistical Tests

This appendix provides additional information on alternative regressions for count data, methods for testing for overdispersion, a discussion of how to choose between methods, and references to further resources. Example R code which runs the analyses and tests discussed here can be found in the Code > Appendix Tutorial Code section of this paper's OSF page: https://osf.io/kcgjb/?view_only=9282cabb9044bdbbf101cd87e4a6f6e.

Quasi-Poisson Regression

Unlike Negative Binomial regressions, which use a different statistical distribution which may better fit the data, a quasi-Poisson regression still assumes the Poisson distribution, but adjusts the inferential statistics arising from it to help account for overdispersion. This adjustment adds a scale parameter which allows variance to be a linear function of the mean, meaning that the two do not have to be equal, unlike in the Poisson regression, where the ratio is assumed to be 1 to 1. Standard errors and test statistics are then corrected based on this parameter.

Zero-Inflated Models

Poisson regressions tend to assume that a very small percentage of all values are zeroes. However, in practice it is often the case that there will be more zero values in the data than you expect. For example, in our experiment looking at the number of Swedish Fish chosen by participants, a significant number of participants were not interested in receiving any Swedish Fish at all. In order to account for this, we model the data generating process in two stages: first, a new first step which determines if a value will be zero or not, and then a step in which we determine what each value is conditional on it not being zero. This correction is often accomplished by first using a logistic regression to predict if values will be zero, and then using a Poisson regression to predict the values of non-zero values. However, this is really a more general procedure – the same logic can apply to other regression models. For example, it would be possible to use some other regression to determine if a value is zero or not, such as probit, or a different regression for determining values conditional on them not being zero, such as a Negative Binomial regression. This may be necessary if zero-inflation isn't the only way that the data violated the assumptions of a Poisson regression – for example, even after accounting for the larger number of zeroes in the data, the data could nevertheless still be overdispersed. One implementation of zero-inflated models can be found in the R package *pscl* (Jackman, 2020).

Tests of Overdispersion

There are a number of ways to test if the assumptions of the Poisson distribution hold. R code on this paper's OSF page demonstrates different tests of over- and under-dispersion. A simple initial heuristic test is to compare the mean and variance of the data to one another and do a simple comparison. If the variance is higher than the mean,

overdispersion is likely, and may be damaging. Another simple method is to fit a Poisson model to the data, and check if residual deviance is greater than residual degrees of freedom. A common rule of thumb is that if the ratio of residual deviance divided by residual degrees of freedom is greater than two, overdispersion may be present. If this is the case, this may indicate overdispersion, although this test can be misleading. However, there are also more formal statistical tests which can be carried out which will provide a more robust test.

One option is using a parametric, regression-based test for overdispersion (Cameron & Trivedi, 1990, see also Cameron & Trivedi, 2001, 2005). What this essentially does is test the hypothesis that the assumption that mean is equal to variance holds in the data against an alternative hypothesis that the variance is actually better described by a function which allows variance to vary from the mean. If the null hypothesis (that mean is equal to variance) is rejected, this signals that there may be under or overdispersion. It can also use an OLS regression to estimate the degree of under- or overdispersion. This is implemented in the *aer* package in R through the function `dispersiontest` (Kleiber & Zeileis, 2008).

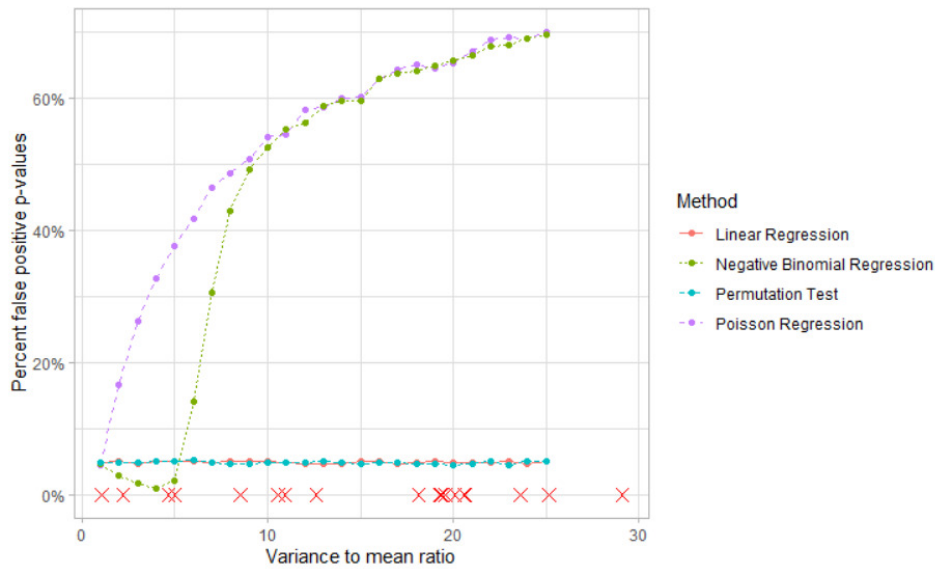
Another option is examining the residuals of a fitted Poisson regression model, and from these determining if over- or under-dispersion is present using a non-parametric test. One implementation comes from the *DHARMA* package in R, using the `simulateResiduals` and `testDispersion` functions as demonstrated in the attached R code (Hartig, 2016).

Deciding which Model to Use

It is not possible to give a complete guide to when to use each type of regression. If you know there is some overdispersion, an OLS regression or non-parametric permutation test may be sufficient, and probably always makes sense as a robustness check. However, if you want to have something which may more closely follow the distribution, then you should probably switch to using quasi-Poisson or Negative Binomial. Which of those two to use is tricky to determine, but there are reasons to choose either. It is not always obvious which of the two methods – Negative Binomial or quasi-Poisson – will provide the best fit to the data. Because quasi-Poisson is technically an adjustment to inferential statistics, not a distribution of its own, it is difficult to compare it with Negative Binomial using likelihood-based measures of model fit. For a useful discussion of how the different weighting parameters in Negative Binomial and Quasi-Poisson models can be more (or less) appropriate for data as well as an example, see Ver Hoef & Boveng (2007).

Online Appendix II: Distributions of Variance to Mean Ratios in Practice

One may ask what the practical significance of the simulations results are – after all, it is possible that high variance to mean ratios are rarely observed in the papers which do use Poisson. To answer this question, we looked at the variance to mean ratios observed in the papers we reviewed for their use of Poisson. [Appendix Figure 1](#) recreates [Figure](#)



Appendix Figure 1. Simulation results plotting the proportion of significant results when in reality there is no difference between the two groups, using a simulated sample size of 100 per cell. Overlaid on the x-axis as red X's which represent the observed variance to mean ratios of studies in papers using Poisson regression to analyze the Voodoo Doll task.

5 from the main paper, but plots along the horizontal axis as red X's the variance to mean ratios above one observed in papers using Poisson to analyze Voodoo Doll task data. The variance to mean ratios span the full length of the axis, providing suggestive evidence that high variance to mean ratios, and the false positive rates which come with them, may be relatively common in practice.

SUPPLEMENTARY MATERIALS

Peer Review History

Download: https://collabra.scholasticahq.com/article/27242-poisson-regressions-a-little-fishy/attachment/67833.docx?auth_token=4WOVLZEHZ852WCgRHNO
