

A BAYESIAN MODEL OF UNDERREPORTING FOR SEXUAL ASSAULT ON COLLEGE CAMPUSES

BY CASEY BRADSHAW^a AND DAVID M. BLEI^b

Department of Statistics, Columbia University, ^acb3431@columbia.edu, ^bdavid.blei@columbia.edu

In an effort to quantify and combat sexual assault, U.S. colleges and universities are required to disclose the number of reported sexual assaults on their campuses each year. However, many instances of sexual assault are never reported to authorities, and consequently, the number of reported assaults does not fully reflect the true total number of assaults that occurred; the reported values could arise from many combinations of reporting rate and true incidence. In this paper we estimate these underlying quantities via a hierarchical Bayesian model of the reported number of assaults. We use informative priors, based on national crime statistics, to act as a tiebreaker to help distinguish between reporting rates and incidence. We outline a Hamiltonian Monte Carlo (HMC) sampling scheme for posterior inference regarding reporting rates and assault incidence at each school and apply this method to campus sexual assault data from 2014–2019. Results suggest an increasing trend in reporting rates for the overall college population during this time. However, the extent of underreporting varies widely across schools. That variation has implications for how individual schools should interpret their reported crime statistics.

1. Introduction. Sexual assault on college campuses is a pressing public health concern in the United States, prompting government action such as the 2014 formation of the White House Task Force to Protect Students from Sexual Assault ([Obama \(2014\)](#)) and the “It’s on us” awareness campaign ([Somanader \(2014\)](#)). Per the Clery Act of 1990, U.S. colleges and universities are required to report annual crime statistics for their campuses and adjacent areas. These crime statistics include the number of reported sexual assaults, but because sexual assault is widely believed to be underreported, these figures do not tell the whole story. The number of reported assaults in the Clery data could arise from many different combinations of reporting rates and true underlying numbers of assaults. Without further assumptions it is equally plausible that a given school has a large number of assaults and a low reporting rate, or few assaults and a high reporting rate.

To estimate the true incidence and reporting rates of sexual assault in the U.S. college population, we construct a hierarchical Bayesian model of the reported data, together with a Markov chain Monte Carlo (MCMC) sampling scheme for posterior inference. More generally, underreported count data arises in many disciplines, with applications including criminology ([Moreno and Girón \(1998\)](#), [Fernández-Fontelo et al. \(2019\)](#)) as well as epidemiology ([Bailey et al. \(2005\)](#), [Bracher and Held \(2021\)](#)), traffic safety ([Kumara and Chin \(2005\)](#), [Ma and Li \(2010\)](#)), and economics ([Winkelmann \(1996\)](#), [Fader and Hardie \(2000\)](#)). The central challenge in such problems is to disentangle the per-school reporting rate from the true latent counts. This issue can be remedied either through the use of validation data not subject to underreporting or by using domain knowledge to assign more informative prior distributions to latent variables. In the case of campus sexual assault, fully observed validation data is unavailable, and so we use national crime statistics as a source of external information. The absence of fully observed validation data also has implications for model assessment.

Without knowledge of the true total number of assaults, we use predictive checks applied to held-out data to demonstrate the fitted model's compatibility with the data.

In general, we find that inference for the latent total number of assaults and reporting rate at a given school offer a clearer picture of the campus environment than the reported figures alone. In particular, these estimates can suggest the relative contributions of reporting rate and incidence to year-over-year changes in the reported number of assaults. This is a relevant line of inquiry, given that *ex ante* an increase in the reported number of assaults could equally be explained by an increase in the reporting rate (a positive outcome, indicating that students are aware of campus support resources) or an increase in the true total number of assaults (a bad outcome).

The remainder of this paper is structured as follows. Section 2 reviews related work, followed by an overview of the dataset in Section 3. The proposed modeling methodology is outlined in Section 4. Section 5 discusses modeling results and insights, and Section 6 offers concluding remarks and directions for future work.

2. Related literature. There are various perspectives on underreported count data. Some studies consider underreporting as a type of censoring. In this formulation each observed value is possibly right-censored and has an accompanying latent censoring indicator. [de Oliveira, Loschi and Assunção \(2017\)](#) take this view and assign a covariate-dependent prior distribution to the latent censoring indicators, including the censoring indicators as a target for posterior inference.

Underreporting also has conceptual overlap with positive-unlabeled (PU) learning. In PU learning, data consist of labeled positive cases and unlabeled cases, which may be positive or negative ([Letouzey, Denis and Gilleron \(2000\)](#)). The recorded number of positive cases is thus a lower bound for the true number of positive cases, which can be viewed as an example of underreporting. The application of PU learning techniques to underreported data is explored in [Shanmugam and Pierson \(2021\)](#) and [Wu et al. \(2023\)](#).

Within the broader literature on underreported count data models, one particular line of work is most relevant to modeling reports of sexual assault. These modeling approaches begin with an assumption that the latent true number of events, z , follows a Poisson distribution with rate parameter λ . It is common to further suppose that each of these z -many events may or may not be reported, independently of all others, with reporting probability p . The number of reported events x , where $x \leq z$, is the data value we observe.

[Moreno and Girón \(1998\)](#) articulate the identifiability issue in this setup: marginally, the observed data x_1, \dots, x_n follow a $\text{Poisson}(\lambda p)$ distribution, so without further constraints or assumptions, one cannot distinguish among the possible pairs (λ, p) with the same product. [Moreno and Girón \(1998\)](#) further demonstrate that placing uninformative prior distributions on λ and p is insufficient to produce useful inference. For instance, the choice of a uniform prior on p and a Jaynes ([Jaynes \(1968\)](#)) or Jeffreys ([Jeffreys \(1946\)](#)) prior on λ results in infinite posterior expectations for λ and z , given the observed data. In some research applications, this identifiability issue is remedied through the use of validation data not subject to underreporting ([Powers, Gerlach and Stamey \(2010\)](#), [Dvorzak and Wagner \(2016\)](#)). An alternative approach is to use external information to assign informative prior distributions, as [Schmertmann and Gonzaga \(2018\)](#) and [Stoner, Economou and Drummond Marques da Silva \(2019\)](#) do for the reporting probability. Because external information is available regarding both the incidence and reporting rate of sexual assault, we take this latter approach, choosing informative priors for both components of the model.

Earlier work on underreporting models explored the somewhat simpler case of i.i.d. observations governed by a single reporting probability and incidence parameter λ ([Moreno and Girón \(1998\)](#), [Fader and Hardie \(2000\)](#)). In this paper we adopt the more recent extension of

works, such as [Dvorzak and Wagner \(2016\)](#) and [Stoner, Economou and Drummond Marques da Silva \(2019\)](#), allowing p and λ to vary across units. When modeling λ_i and p_i , we introduce school features as covariates to account for demographic differences among schools, as detailed in Section 4.

Posterior distributions for latent variables in this type of underreporting model are generally intractable. However, choosing conditionally conjugate priors for λ and p (Gamma and Beta distributions, respectively) leads to analytic expressions for their marginal posterior distributions ([Moreno and Girón \(1998\)](#), [Fader and Hardie \(2000\)](#)). Alternatively, MCMC approaches to posterior inference allow greater flexibility in choice of priors in exchange for a greater computational burden. [Dvorzak and Wagner \(2016\)](#) use a bespoke implementation of a Gibbs-style sampler, while [Stoner, Economou and Drummond Marques da Silva \(2019\)](#) use slice sampling. Our paper extends this line of work by demonstrating the viability of gradient-based samplers such as HMC for posterior inference in this model setting.

With respect to the application context, the most relevant work comes from [Fernández-Fontelo et al. \(2019\)](#), who model underreporting of domestic violence complaints across districts in Spain. Unlike the campus sexual assault data, observations in the dataset compiled by [Fernández-Fontelo et al. \(2019\)](#) are dense in the time domain (quarterly for 10 years), motivating the authors to adopt time series methodology.

3. Dataset. The dataset contains 11,369 records for 1973 U.S. colleges and universities over the 2014–2019 time frame,¹ where a “record” consists of the total number of assaults reported at a particular school in one particular calendar year. Schools located in the 50 U.S. states, the District of Columbia, and Puerto Rico are included in this analysis. Annual campus crime statistics are furnished by the U.S. Department of Education’s Office of Postsecondary Education. All schools that receive federal financial aid funding under Title IV of the Higher Education Act of 1965 are required to submit annual campus crime reports. Most schools have records for all six years in this time frame, though 8% appear in fewer years due to schools opening/closing or beginning/ ceasing to participate in Title IV.

School characteristics and demographic data are available through the National Center for Education Statistics (NCES), which publishes comprehensive institutional profile information through its IPEDS database.² Preprocessed data is included in the Supplementary Material; raw data is available publicly and from the corresponding author upon reasonable request. Additional dataset details are included in Appendix A in the Supplementary Material ([Bradshaw and Blei \(2024\)](#)).

We exclude institutions with no degree-granting programs and institutions with no programs classified as “academic” in nature. We also exclude institutions with no residential housing facilities. For such schools the school does not represent the students’ community, and we would not reasonably expect students to report sexual assaults to the school.

Institutions in this dataset range in size from to 79,500 in-person students, with a median of 2350. Public institutions comprise 41% of the dataset, religiously affiliated institutions 36%, and private, nonreligiously affiliated institutions 22%. The reported sexual assault data is notably sparse: 42% of records show zero reported assaults, and 22% of institutions show no assaults across all six years. The distribution of reported values, shown in Figure 1, has a long right tail, with a handful of schools reporting more than 100 assaults in a given year. This sparsity pattern is not explained by student population size alone: schools with as many as 36,000 students had zero reported assaults, while assaults *were* reported at schools with as few as several hundred students.

¹Data from 2020 and 2021 are excluded. Because the COVID-19 pandemic significantly disrupted patterns of social interaction, campus sexual assault data from this time are not well described by the same model used for other years.

²<https://nces.ed.gov/ipeds/>.

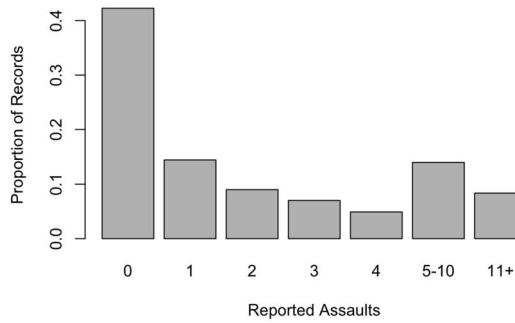


FIG. 1. *School-level reported sexual assault data is sparse, with 42% of records showing zero assaults, and a median of one assault.*

The total number of assaults reported across schools increased steadily from 2014 to 2018, growing by nearly 50% before falling in 2019 (see Figure 2). The total in-person student population was relatively stable over this time period, remaining close to 10.6 million.

Campus crime numbers in this dataset are a direct tally of the individual assaults for which a student voluntarily made a formal report to campus authorities. Note that those reports do not arise from a survey process or broad inquiry about student experiences; on its own this dataset contains no information about underreporting. For complementary information about reporting rates, we consider the National Crime Victimization Survey (NCVS). The NCVS is a large-scale annual survey conducted on behalf of the U.S. Bureau of Justice Statistics to determine incidence rates of personal and property crimes, including sexual assault. This survey dates back 50 years, and one of its primary objectives is to quantify the extent of crime not reported to authorities. By attempting to measure crimes not reported to authorities, the NCVS provides additional context which is absent from the college campus Clery Report data.

4. Methodology. For a particular school i in a particular year j , suppose that the true, unknown number of assaults z_{ij} comes from a Poisson distribution with rate parameter λ_{ij} . Further suppose that each of those z_{ij} assaults is independently reported, or not reported, with probability of reporting equal to p_{ij} . This produces the reported value x_{ij} that we observe, according to a binomial thinning process,

$$(1) \quad x_{ij} | z_{ij}, p_{ij} \sim \text{Binom}(z_{ij}, p_{ij}).$$

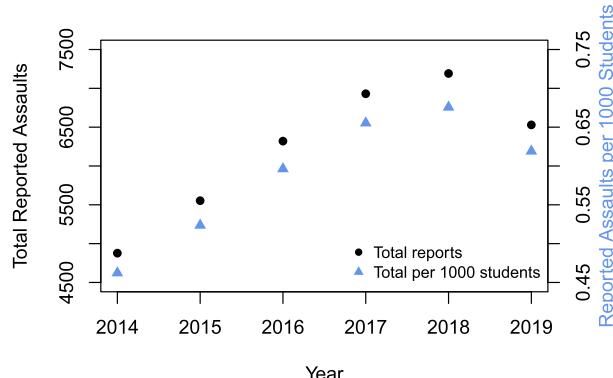


FIG. 2. *The total number of reported assaults grew from 2014 to 2018. Total student population fluctuated relatively little, resulting in per capita reporting trends moving with the overall total number of reports.*

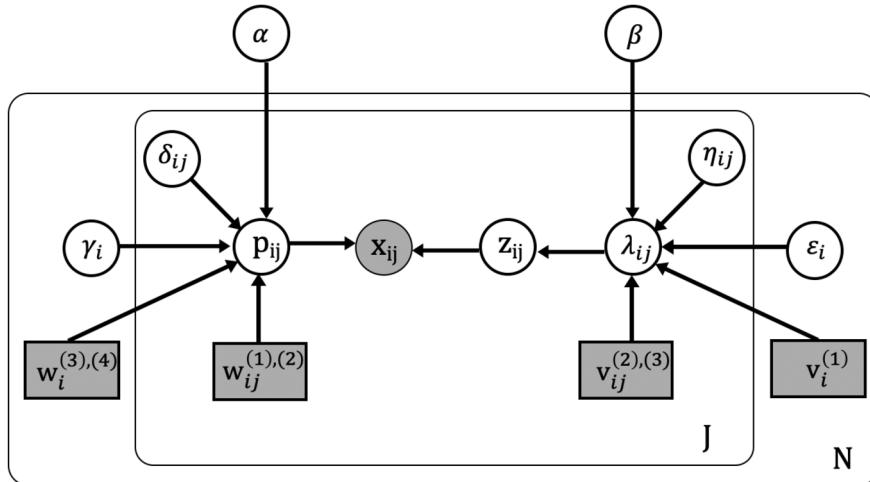


FIG. 3. Graphical model diagram.

With this likelihood model, we construct a hierarchical generative model for the reported campus sexual assault data. The proposed hierarchical model, diagrammed in Figure 3, entails a component describing the reporting probability, and a component describing the Poisson rate governing the true number of assaults. A common approach in the underreporting literature is to model $\log(\lambda_{ij})$ and/or $\log(\frac{p_{ij}}{1-p_{ij}})$ as deterministically equal to a linear function of covariates (Dvorzak and Wagner (2016), Stoner, Economou and Drummond Marques da Silva (2019), de Oliveira et al. (2021)). To provide additional flexibility, we instead model the Poisson rate and reporting probability as random functions of their respective covariates. Full details of each model component are discussed below.

4.1. Modeling z . For each record in the dataset, we suppose the true number of assaults arises from a Poisson distribution parameterized by λ_{ij} , where λ_{ij} depends on observed characteristics \mathbf{v} of school i in year j as follows:

$$(2) \quad z_{ij} | \lambda_{ij} \sim \text{Poisson}(\lambda_{ij}),$$

$$(3) \quad \log(\lambda_{ij}) = \beta_{0,v_i^{(1)}} + \beta_1 v_{ij}^{(2)} + \beta_2 v_{ij}^{(3)} + \varepsilon_i + \eta_{ij},$$

$$\beta_{0,1}, \beta_{0,2}, \beta_{0,3} \sim N(-5.5, 0.5),$$

$$\beta_1 \sim N(1, 0.1),$$

$$\beta_2 \sim N(0, 4),$$

$$\varepsilon_i \sim N(0, 0.75),$$

$$\eta_{ij} \sim N(0, 0.1).$$

The covariates associated with λ are:

$v_i^{(1)}$: degree of urbanization of school i 's campus:

1 = urban;

2 = suburban;

3 = rural.

$v_{ij}^{(2)}$: $\log(\text{number of in-person students})$ for school i in year j .

$v_{ij}^{(3)}$: $(\text{fraction of women in student body at school } i \text{ in year } j - 0.5)^2$.

This generative model considers in-person student population, degree of campus urbanization, and student body gender composition³ as covariates of interest. The functional form of equations (2)–(3) implies that for the expected number of assaults we have

$$(4) \quad \begin{aligned} \mathbb{E}(z_{ij} | \beta, \mathbf{v}) &\propto \exp(\beta_0 v_i^{(1)} + \beta_2 v_i^{(3)}) \exp(\beta_1 \log(\text{student population}_{ij})) \\ &\propto (\text{student population}_{ij})^{\beta_1}. \end{aligned}$$

This modeling choice allows for the possibility that the number of assaults does not scale linearly with the number of in-person students. Because the number of interpersonal interactions does not always scale linearly with community size, this type of “power law” dynamic has been observed for various social phenomena, including crime (Bettencourt et al. (2007), Chang et al. (2018)).

Incorporating separate intercepts in (3) allows for different incidence patterns at urban, suburban, and rural schools. This decision is motivated by the national differences in violent crime victimization rates across locales of varying urbanization (Truman and Langton (2014, 2015), Truman and Morgan (2016), Morgan and Kena (2018), Morgan and Truman (2018), Morgan and Thompson (2020)).

Gender composition of the student body is included as a covariate because most perpetrators of sexual assault are men, while most victims are women (Sinozich and Langton (2014)). We center the proportion of women in the student body by subtracting 0.5, and square this value, to allow for the possibility that schools with more extreme gender imbalances may have different rates of assault.

Per-school intercepts ε_i allow for between-school differences in assault incidence patterns which are not otherwise captured by covariates. Per-record noise terms η_{ij} allow for within-school differences across years. Supposing that $\beta_1 = 1$ (linear relationship between student population and number of assaults) and $\beta_2 = 0$ (no relationship between student body gender composition and expected number of assaults), priors on ε , η , and intercepts β_0 are set so that the resulting distribution over assaults per 1000 persons roughly corresponds with external estimates from the National Crime Victimization Survey (NCVS). The resulting distribution is displayed in Figure 4. Note that, conditional on covariates, the Poisson rate λ_{ij} follows a log-normal distribution, leading to a number of assaults z_{ij} whose marginal distribution is overdispersed relative to a Poisson.

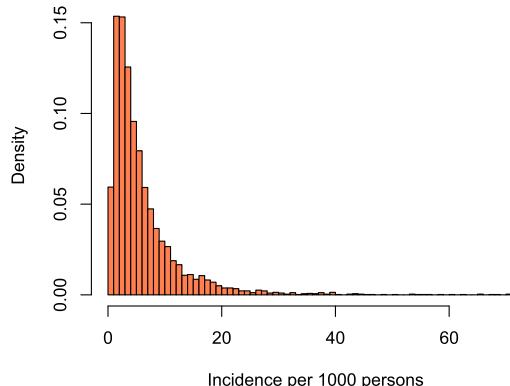


FIG. 4. *Prior distributions on β_0 , ε , and η imply a distribution over incidence centered at 4.1 per 1000 persons, with a central 50% range of 2.3 to 7.4.*

³The gender categories reported by NCES are currently “male” and “female.”.

TABLE 1
Sexual assault incidence per 1000 persons in the U.S., by age group

	18–20	21–24	12+
2014	1.2*	3.8	1.1
2015	3.2	1.6*	1.6
2016	3.8	1.8	1.1
2017	8.6	2.4	1.4
2018	10.1	6.0	2.7
2019	7.1	6.4	1.7

* Per BJS: “[estimate is] based on 10 or fewer sample cases or coefficient of variation is greater than 50%.”

The NCVS is a large sample survey conducted annually on behalf of the U.S. Bureau of Justice Statistics to determine incidence rates of personal and property crimes. This survey includes crimes reported and not reported to the police. Over the 2014–2019 time period, the estimated incidence of sexual assault in the overall population aged 12 and older ranged from 1.1 to 2.7 per 1000 persons. As shown in Table 1, estimated incidence skewed higher in the 18–20 and 21–24 age groups, which are particularly relevant to the college student population.

4.2. Modeling p . As formulated in equation (1), the reported number of assaults arises from the latent true number of assaults z and the reporting probability p . Having discussed a model for z , we now turn to modeling the reporting probability p . Of the total z_{ij} assaults occurring at school i in year j , suppose that each individual assault is reported, independently, with probability p_{ij} , where p_{ij} depends on observed characteristics \mathbf{w} of school i in year j as follows:

$$\log\left(\frac{p_{ij}}{1 - p_{ij}}\right) = \alpha_0 + \alpha_1 w_i^{(1)} + \alpha_2 w_i^{(2)} + \alpha_3 w_{ij}^{(3)} + \alpha_4 w_{ij}^{(4)} + \gamma_i + \delta_{ij},$$

$$\alpha_0 \sim N(-1.25, 0.5),$$

$$(5) \quad \alpha_1, \alpha_2, \sim N(0, 2),$$

$$\alpha_3, \alpha_4 \sim N(0, 4),$$

$$\gamma_i \sim N(0, 1.25),$$

$$\delta_{ij} \sim N(0, 0.5).$$

The covariates associated with p are:

$w_i^{(1)}$: school i issues associate degrees only ($\in \{0, 1\}$).

$w_i^{(2)}$: school i substantially engaged in religious instruction ($\in \{0, 1\}$).

$w_{ij}^{(3)}$: (women as fraction of student body at school i in year $j - 0.5$).

$w_{ij}^{(4)}$: Pell grant recipients as fraction of student body at school i in year j , median centered.

The reporting probability at school i in year j accounts for two school-level covariates that are fixed across years and two that vary. The first fixed school-level covariate is $w_i^{(1)}$, an indicator for whether school i is a junior college, issuing only associate degrees as opposed to bachelor degrees or higher. The reporting probability also incorporates an indicator $w_i^{(2)}$ for whether school i is substantially engaged in religious instruction. A theological seminary, for instance, would satisfy this definition, whereas a religiously affiliated liberal arts college would not. The inclusion of these school-level random effects encourages the probabilities

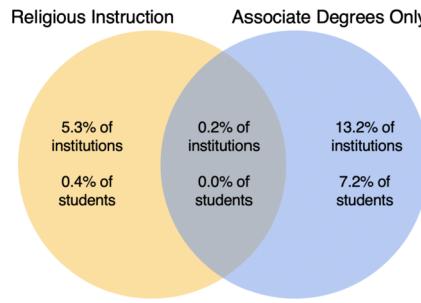


FIG. 5. *Junior colleges and institutions of religious instruction comprise a minority of the dataset.*

for an individual school to exhibit similarity across years. Junior colleges comprise roughly 13% of institutions in the dataset and account for 7% of total in-person students. Institutions of religious instruction make up a relatively smaller portion of the dataset and, generally, have small student populations. These two categories of institution have negligible overlap, as detailed in Figure 5.

Gender composition of a school's student body also contributes to the reporting probability and varies across years. Many institutions in this dataset have approximately gender-balanced student populations. Across records, the proportion of female students is centered around 57%, with smaller peaks at each end of the spectrum corresponding to all-male and all-female schools. The final covariate incorporated in the reporting rate is the fraction of Pell grant recipients in the student body. Pell grants are federal financial aid awarded to students on the basis of exceptional financial need. This variable is included as a proxy for socioeconomic status, which research indicates may be correlated with reporting rates (Fisher et al. (2003), Sabina and Ho (2014)). The median percentage of undergraduates receiving Pell grants is 36%, and values are more widely spread over the unit interval compared to the gender composition covariate.

As in the model of Poisson rate λ_{ij} , per-school intercepts γ_i in equation (5) allow for between-school differences in assault incidence patterns not captured by covariates, while per-record noise terms δ_{ij} allow for within-school differences across years. Supposing that $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$, priors on ε , η , and intercepts β_0 are chosen to be broadly consistent with NCVS estimates of the percentage of sexual assaults reported to police. Over the 2014–2019 time period, the estimated rate of reporting to police ranged from 23% to 40% in the overall 12+ population, while reporting rates for the 18–20 and 21–24 age groups tended to be lower (see Table 2). The induced prior distribution over values of p , when $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$, is depicted in Figure 6 below.

In these models of z and p , global coefficients α and β encourage the sharing of information across schools, while school-level terms γ and ε capture local properties of individual

TABLE 2
Percentage of sexual assaults reported to police, by age group

	18–20	21–24	12+
2014	0.0*	16.48	33.6
2015	10.6*	26.0*	32.5
2016	10.2*	31.0*	23.2
2017	45.3	38.3*	40.4
2018	11.0*	10.2*	24.9
2019	17.5*	24.7*	33.9

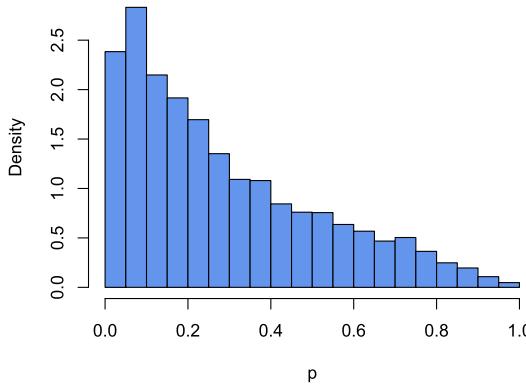


FIG. 6. *Prior distributions on α_0 , γ , and δ imply a distribution over reporting probability centered at 22%, with a central 50% range of 10% to 43%.*

schools. This partial pooling model structure is a compromise between complete pooling, in which no school-level differences are permitted (γ and ε omitted), and no pooling, in which each school i is modeled separately from all others, with its own coefficient vectors α_i and β_i . Compared to the complete-pooling and no-pooling alternatives, the partial pooling model assigns higher likelihood to held-out observational data refer to Appendix C in the Supplementary Material (Bradshaw and Blei (2024)) for details.

4.3. Inference. The model does not admit closed form expressions for posterior quantities of interest. We draw samples from the joint posterior distribution over latent variables via MCMC. The true number of assaults z_{ij} is a discrete latent variable, which presents an obstacle to gradient-based sampling schemes such as Hamiltonian Monte Carlo (HMC). However, this can be circumvented by marginalizing out z_{ij} . Instead of working with z_{ij} , one can take advantage of the fact that, conditional on the event rate λ_{ij} and reporting probability p_{ij} but not the true number of assaults z_{ij} , the reported number of assaults follows a Poisson distribution,

$$x_{ij} | \lambda_{ij}, p_{ij}, \mathbf{w}_{ij}, \mathbf{v}_{ij} \sim \text{Poisson}(\lambda_{ij} p_{ij}).$$

See Appendix B in the Supplementary Material (Bradshaw and Blei (2024)) for derivation. This marginalizes out the discrete latent variable z_{ij} , allowing for gradient-based MCMC sampling from the posterior $p(\lambda, \mathbf{p} | \mathbf{x}, \mathbf{w}, \mathbf{v})$. However, the latent true number of assaults z_{ij} is also a quantity of interest. To address this, first note the conditional distribution of the number of unreported assaults, u_{ij} (where $u_{ij} \stackrel{\text{def}}{=} z_{ij} - x_{ij}$),

$$u_{ij} | x_{ij}, \lambda_{ij}, p_{ij} \sim \text{Poisson}(\lambda_{ij}(1 - p_{ij})).$$

Because u_{ij} is conditionally independent of the covariates \mathbf{v}_{ij} and \mathbf{w}_{ij} , given latent parameters λ_{ij} and p_{ij} , it is equivalently true that

$$u_{ij} | x_{ij}, \lambda_{ij}, p_{ij}, \mathbf{w}_{ij}, \mathbf{v}_{ij} \sim \text{Poisson}(\lambda_{ij}(1 - p_{ij})).$$

Each MCMC sample $(\lambda_{ij}^{(s)}, p_{ij}^{(s)})$ drawn from (approximately) $p(\lambda_{ij}, p_{ij} | x_{ij}, \mathbf{w}_{ij}, \mathbf{v}_{ij})$ can be augmented by sampling a corresponding $u_{ij}^{(s)}$ from $p(u | x_{ij}, \lambda_{ij}, p_{ij})$. The sampled value $u_{ij}^{(s)}$ can then be added to the observed count x_{ij} to produce a sample for the latent count $z_{ij}^{(s)}$. Altogether, the sample $(\lambda^{(s)}, \mathbf{p}^{(s)}, \mathbf{z}^{(s)})$ constitutes an MCMC sample from $p(\lambda, \mathbf{p}, \mathbf{z} | \mathbf{x})$. This sampling procedure is outlined in Algorithm 1 and can be conveniently implemented in a probabilistic programming language. Posterior inference for this work was conducted using HMC in Stan (Stan Development Team (2023)).

Algorithm 1: Posterior Sampling Scheme**Input:**

- Data \mathbf{x}
- Covariates \mathbf{w}, \mathbf{v}
- Index set \mathcal{I} (pairs (i,j) for which the dataset contains a record x_{ij})
- MCMC algorithm SAMPLER (e.g., HMC)
- Markov chain length S

Output: samples from the posterior $p(\mathbf{z}, \boldsymbol{\lambda}, \mathbf{p}|\mathbf{x}, \mathbf{w}, \mathbf{v})$;

for s in $1:S$ **do**

Draw $(\boldsymbol{\lambda}^{(s)}, \mathbf{p}^{(s)})$ from SAMPLER approximating $p(\boldsymbol{\lambda}, \mathbf{p}|\mathbf{x}, \mathbf{w}, \mathbf{v})$ **for** (i, j) in \mathcal{I} **do**

Draw $u_{ij} \sim \text{Poisson}((1 - p_{ij}^{(s)})\lambda_{ij}^{(s)})$;

Set $z_{ij} = u_{ij} + x_{ij}$;

end

end

return $\{(z^{(s)}, \boldsymbol{\lambda}^{(s)}, \mathbf{p}^{(s)})\}_{s=1}^S$

4.4. *Model assessment.* Fully-observed data, including the true total number of assaults, are not available and thus cannot be used for model validation. Instead, the predictive distribution provides information about the quality of the model fit. The posterior distribution of the latent variables implies a predictive distribution over future observations $p(x^{\text{new}}|x)$ and, accordingly, over summary statistics of future observations. This motivates the notion of posterior predictive checks (Guttman (1967), Box (1980), Rubin (1984), Gelman, Meng and Stern (1996)) that a satisfactory model will yield predictive distributions over key summary statistics under which the actual values observed in the dataset are not extreme. As detailed in Moran, Blei and Ranganath (2019) and Li and Huggins (2022), such posterior predictive checks can be overly optimistic about model fit due to double use of the data (both to fit the model and to evaluate it). The aforementioned works propose randomly reserving a portion of the original dataset to use exclusively for model assessment, an approach that demonstrably achieves a more accurate gauge of model performance by avoiding the pitfall of data reuse. In the spirit of this suggestion, roughly 20% of entries in the campus sexual assault dataset are held out, with the remaining 80% used for posterior inference. Held out data points are sampled at the level of (school, year) pairs. The held-out data set includes records from all years in scope (2014–2019) and does not contain data from any additional schools not seen in the main sample. Predictive samples for the held-out data points are generated as follows.

As noted in Section 3, each year many schools had zero reported assaults, and the median number of reports was one. Generating 10,000 resampled datasets according to Algorithm 2 and calculating the proportion of records with zero reported assaults produces the distribution shown below in Figure 7(a). Within the held-out dataset, 42.1% of records had zero reported assaults, while values ranged from roughly 38% to 44% in datasets sampled from the predictive distribution. The distribution of reported assault numbers has been pulled slightly away from the extreme end such that the observed proportion of zeroes falls toward the high end of the predictive distribution. However, the median number of assaults was equal to one in all datasets sampled from the predictive distribution, and the true proportion of held-out records x_{ij} with ≤ 1 assaults sits comfortably within its respective predictive distribution (Figure 7(b)). Overall, the model adequately captures the bottom-heaviness of the distribution of reported assault numbers.

Algorithm 2: Predictive Distribution Sampling**Input:**

- Index set \mathcal{I} (pairs (i, j) for which the held-out dataset contains a record x_{ij}^{new})
- Covariates $\mathbf{w}^{\text{new}}, \mathbf{v}^{\text{new}}$ for held-out dataset
- Posterior distribution $p(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\epsilon} | \mathbf{x}, \mathbf{w}, \mathbf{v})$

Output: Sample dataset from the predictive distribution $p(\mathbf{x}^{\text{new}} | \mathbf{x}, \mathbf{w}, \mathbf{v}, \mathbf{w}^{\text{new}}, \mathbf{v}^{\text{new}})$;

```

Draw  $(\boldsymbol{\alpha}^s, \boldsymbol{\beta}^s, \boldsymbol{\gamma}^s, \boldsymbol{\epsilon}^s)$  ;
for  $(i, j)$  in  $\mathcal{I}$  do
  Draw  $\delta_{ij} \sim N(0, 0.5)$  ;
  Draw  $\eta_{ij} \sim N(0, 0.2)$  ;
  if school  $i \in \text{training data}$  then
     $\gamma_i \leftarrow \gamma_i^s$  ;
     $\varepsilon_i \leftarrow \varepsilon_i^s$  ;
  else
    Draw  $\gamma_i \sim N(0, 1)$  ;
    Draw  $\varepsilon_i \sim N(0, 0.5)$  ;
  end
   $p_{ij} \leftarrow \text{invlogit}(\alpha_0^s + \langle \boldsymbol{\alpha}^s, \mathbf{w}_{ij}^{\text{new}} \rangle + \gamma_i + \delta_{ij})$  ;
   $\lambda_{ij} \leftarrow \exp(\beta_{0, v_i}^s + \langle \boldsymbol{\beta}^s, \mathbf{v}_{ij}^{\text{new}} \rangle + \varepsilon_i + \eta_{ij})$  ;
  Draw  $z_{ij} \sim \text{Poisson}(\lambda_{ij})$  ;
  Draw  $x_{ij}^{\text{new}} \sim \text{Binom}(z_{ij}, p_{ij})$  ;
end
return  $\mathbf{x}^{\text{new}}$ 

```

Other quantities of interest are the total number of reported assaults $\sum_{ij} x_{ij}^{\text{new}}$ and the amount of within-school variability. The predicted number of assaults reported by an individual school should not be too rigid or too flexible over time, relative to the trends seen in the actual data. This summary statistic is quantified as $\sum_i \text{var}(x_i^{\text{new}})$. The total number of reported assaults and the within-school variability observed in the held-out dataset are both plausible under their respective predictive distributions, shown in Figure 7(c)–(d), suggesting that the model adequately captures these characteristics of the data.

4.5. Capabilities and limitations. Let us pause to examine what this estimation method can and cannot do in terms of resolving the central identifiability issue introduced in Sections 1 and 2. First, consider a toy example in which $z_1, \dots, z_N \stackrel{\text{iid}}{\sim} \text{Poisson}(\lambda_0)$ are latent and $x_i \sim \text{Binom}(z_i, p_0)$, $i = 1, \dots, N$ are observed. The observed dataset identifies the product $\lambda_0 p_0$, while only the prior distributions provide information about λ_0 and p_0 individually. The posterior distribution of the product $\lambda_0 p_0$ concentrates as the number of observations increases. The posterior distributions of λ_0 and p_0 , however, do not become arbitrarily concentrated in the infinite data limit but rather approach $p(\lambda_0 | \lambda_0 p_0)$ and $p(p_0 | \lambda_0 p_0)$, respectively, as demonstrated in Figure 8.

Furthermore, the posterior distribution $p(\lambda_0 | x_1, \dots, x_N)$ will not necessarily be centered at the true value λ_0 (and likewise for p_0); the data provide information about the product $\lambda_0 p_0$, but conditional on $\lambda_0 p_0$, posterior estimates of λ_0 and p_0 are determined by their respective prior distributions.

A similar relationship emerges when the toy model is extended to include covariates for λ and p . Let $\log(\lambda_i) = \beta_0 + \beta v_i$ and $\log(\frac{p_i}{1-p_i}) = \alpha_0 + \alpha w_i$ for $v_i, w_i \in \mathbb{R}$. For $i = 1, \dots, N$,

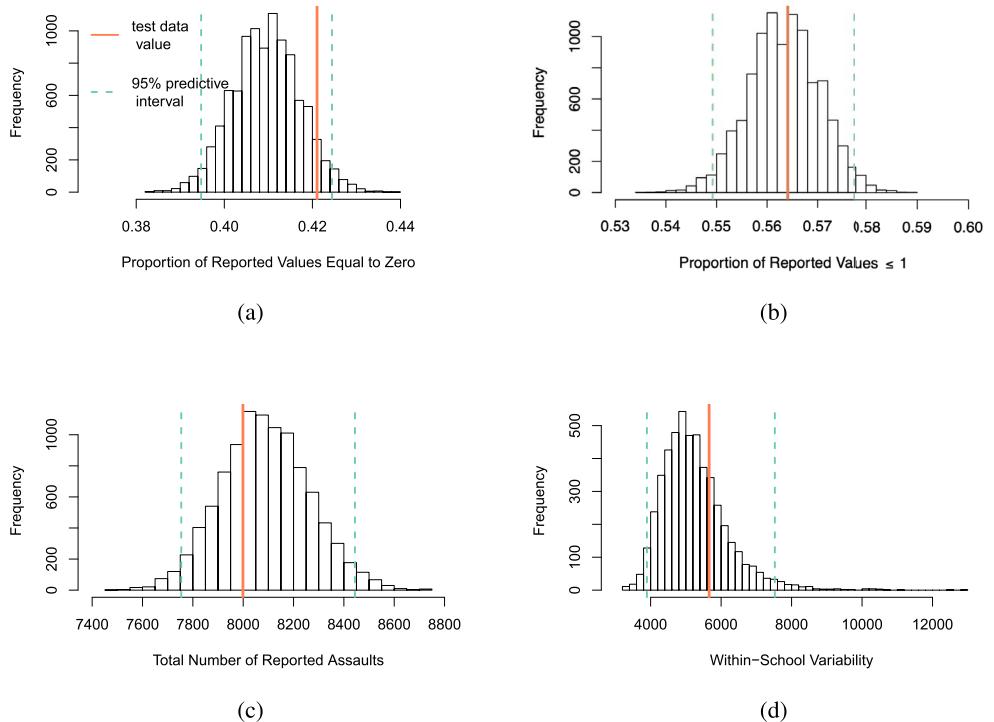


FIG. 7. Summary statistics of the held-out data are plausible under the predictive distribution. Solid lines indicate the value calculated on the held-out data, dashed lines indicate the 2.5th and 97.5th percentiles of the predictive distribution: (a) Proportion of reports indicating zero assaults, (b) Proportion of reports indicating ≤ 1 assaults, (c) Total number of reported assaults across all schools and all years, (d) Within-school variability of reported numbers over time.

latent counts are generated as $z_i \sim \text{Poisson}(\lambda_i)$, and we observe the reported counts $x_i \sim \text{Binom}(z_i, p_i)$. Performing posterior inference on data simulated from this model reveals that increasing the amount of observational data is more effective at reducing posterior uncertainty about the slope parameters β and α than about the intercepts β_0 and α_0 . Figure 9 illustrates the relative difficulty of learning one of the intercept parameters from data.

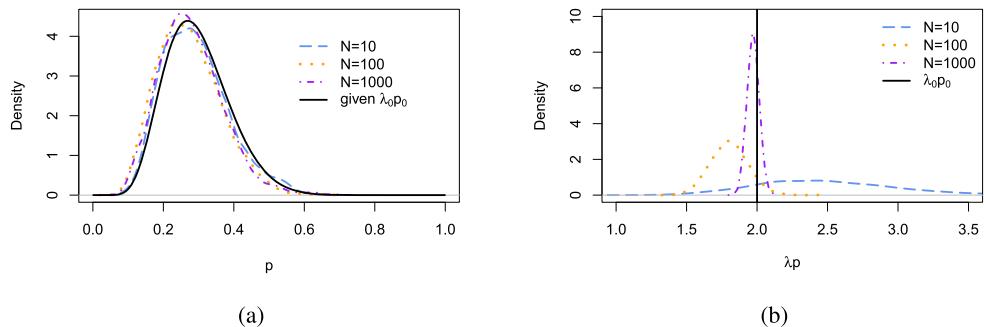


FIG. 8. The posterior distribution of $\lambda_0 p_0$ concentrates as the amount of observational data increases, but the posterior distribution of p_0 has some amount of uncertainty that cannot be eliminated by observational data: (a) Kernel density estimates of posterior samples of p_0 , conditioned on observed datasets of size N , with a solid line representing $p(p_0|\lambda_0 p_0)$, (b) Kernel density estimates of posterior samples of $\lambda_0 p_0$, conditioned on observed datasets of size N .

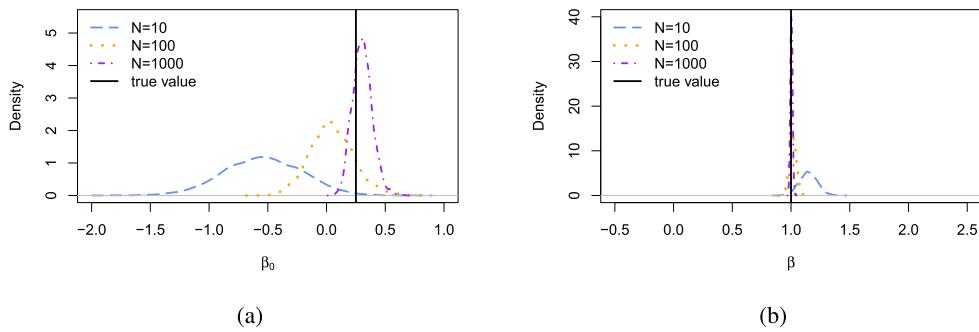


FIG. 9. Additional observational data is more effective at reducing posterior uncertainty about the slope coefficient β than about the intercept: (a) Kernel density estimates of posterior samples of β_0 , conditioned on observed datasets of size N , (b) Kernel density estimates of posterior samples of β_0 , conditioned on observed datasets of size N , displayed on x -axis interval of equal width as left-hand plot.

In light of these dynamics, Appendix E in the Supplementary Material (Bradshaw and Blei (2024)) explores sensitivity to the prior information elicited from NCVS estimates. While some reasons for not reporting a sexual assault are less relevant in a survey setting (such as fear of retaliation or lack of faith that the perpetrator will be held accountable), it is nevertheless plausible that some NCVS respondents choose not to disclose their sexual assault victimization. In line with the discussion in this section, inferences about the role of covariates are relatively more stable if the extent of underreporting is larger than estimated by NCVS, while inferences about the true overall incidence and reporting rate exhibit more sensitivity.

5. Results.

5.1. *Role of covariates.* Sampling from the model in Section 4 yields the posterior means and quartiles given in Table 3 below.

The posterior distribution for β_1 is concentrated around 0.82, which corresponds to sub-linear growth in the expected number of assaults as a function of number of students. The proposed power law relationship between student population and expected number of assaults (4) implies a similar relationship between student population and the expected per-capita rate

TABLE 3
Posterior means and quartiles of latent variables

Variable	Mean	25%	Median	75%	R hat [†]
β_1	0.82	0.81	0.82	0.84	1.00
β_2	-4.05	-4.53	-4.05	-3.56	1.00
$\beta_{0,1}$	-4.54	-4.67	-4.54	-4.42	1.00
$\beta_{0,2}$	-4.40	-4.53	-4.40	-4.28	1.00
$\beta_{0,3}$	-4.32	-4.45	-4.32	-4.21	1.00
α_0	-1.43	-1.50	-1.43	-1.36	1.00
α_1	-1.97	-2.06	-1.97	-1.88	1.00
α_2	-2.48	-2.73	-2.48	-2.23	1.00
α_3	0.49	0.26	0.49	0.73	1.00
α_4	-3.00	-3.18	-3.00	-2.83	1.00

[†]R hat, also known as the potential scale reduction factor, is a convergence diagnostic proposed in Gelman and Rubin (1992)

of assaults, $\frac{z_{ij}}{\text{studentpopulation}_{ij}}$,

$$\mathbb{E}\left(\frac{z_{ij}}{\text{student population}_{ij}} | \boldsymbol{\beta}, \mathbf{v}\right) \propto (\text{student population}_{ij})^{\beta_1 - 1}.$$

Consequently, the posterior distribution suggests an inverse relationship between the size of the student population and the expected per-capita number of assaults in a given year. With all other covariates held equal, we expect a school to have 13% more assaults per capita than an equivalent school twice its size, and 28% more than a school four times its size. Similarly, because the posterior distribution suggests $\beta_1 < 1$, the variance of the per capita number of assaults, $\text{var}(\frac{z_{ij}}{\text{studentpopulation}_{ij}} | \boldsymbol{\beta}, \mathbf{v})$, scales inversely with the number of students. With other covariates held constant, the expected number of assaults per capita is more variable at small schools than at large ones.

Posterior distributions for the intercepts $\beta_{0,1}$, $\beta_{0,2}$, and $\beta_{0,3}$ are similar to each other. Results are mildly suggestive of higher sexual assault incidence in rural areas than in cities but do not reveal a substantial difference in expected number of assaults based on degree of campus urbanization. The negative values for β_2 suggest that a gender imbalance in a school's student body is associated with lower expected incidence of sexual assault, while $\alpha_3 > 0$ suggests that a higher proportion of female students corresponds to a higher reporting probability. Conversely, a higher proportion of students receiving Pell grants (a proxy for lower socioeconomic status) appears associated with a lower probability of reporting. Negative estimates for α_1 and α_2 signal lower expected reporting probabilities for junior colleges and institutions of religious instruction.

5.2. Systemwide results. Drawing posterior samples of the true number of assaults at each school induces a distribution over the total number of assaults across schools each year. The posterior median for total assaults ranges from 2.6 to 2.8 per 1000 persons over the 2014–2019 time period. Although the median estimated incidence is higher in later years, compared to 2014, posterior uncertainty is large enough that the true trend in incidence could conceivably be flat or even increasing, as can be seen in Figure 10 below.

Posterior estimates of the true reporting rate, however, show more movement over time. Figure 11 depicts posterior summaries for the actual number of reported assaults across schools as a fraction of the estimated true number of assaults. The posterior median reporting rate was lowest in 2014, at 17.3% and highest in 2018 at 24.2%, with reporting rates trending upward over those years.

Taken together, these two results suggest that the increase in total sexual assaults reported nationwide over the 2014–2019 period is more likely attributable to an increase in reporting rates than an increase in the true number of assaults that occurred.

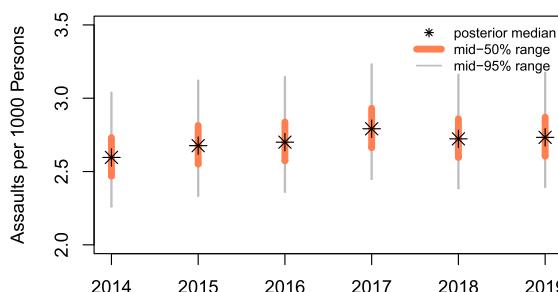


FIG. 10. Posterior distribution of total incidence per 1000 persons in the college population does not exhibit a clear pattern of year-over-year change.

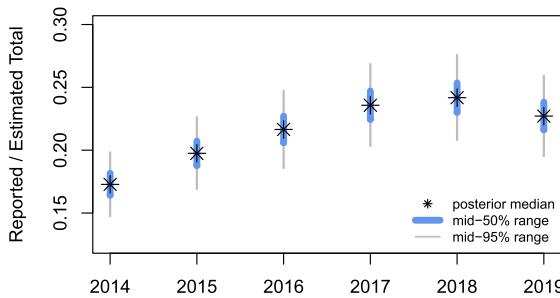


FIG. 11. Posterior estimates of the overall reporting rate in the college population exhibit an increasing trend from 2014 to 2018.

5.3. School-level results. While the systemwide increase in reported assaults may be more easily explained by an increase in reporting rates, aggregate trends are not necessarily representative of the dynamics at an individual school. Examining the posterior medians of the reporting rate p_{ij} and true incidence per 1000 persons across all records in the dataset reveals considerable heterogeneity, as displayed in Figure 12.

While sexual assault is susceptible to underreporting at all schools, the estimated degree of underreporting varies widely. Figure 12(b) illustrates that some schools have extremely low reporting probabilities, while others are considerably higher than the population-level reporting rates published in the NCVS. This variation across schools has implications for how individual schools should interpret their reported crime statistics. With only the prior belief that sexual assault is significantly underreported, one might be equally inclined to attribute an increase in reported assaults at any school to an increase in the reporting probabilities. However, this may lead to flawed conclusions.

Suppose that for a particular school, the true number of assaults in a future year, included in the dataset, remains the same as the true number of assaults in 2019. With an unchanged number of true assaults, how large of a change in the reported number of assaults is plausible due to variation in the reporting probability alone? For school i the predictive distribution for a future year's number of reported assaults can be sampled as follows:

- (1) Draw $(z_{i,2019}^{(s)}, \alpha_0^{(s)}, \boldsymbol{\alpha}^{(s)}, \gamma_i^{(s)})$ from model's posterior distribution.
- (2) Draw $\delta^{\text{new}} \sim N(0, 0.5)$.
- (3) Set $p^{\text{new}} = \text{invlogit}(\alpha_0^{(s)} + \langle \boldsymbol{\alpha}^{(s)}, \mathbf{w}_{i,2019} \rangle + \gamma_i^{(s)} + \delta^{\text{new}})$.
- (4) Draw $x^{\text{new}} \sim \text{Binom}(z_{i,2019}, p^{\text{new}})$.

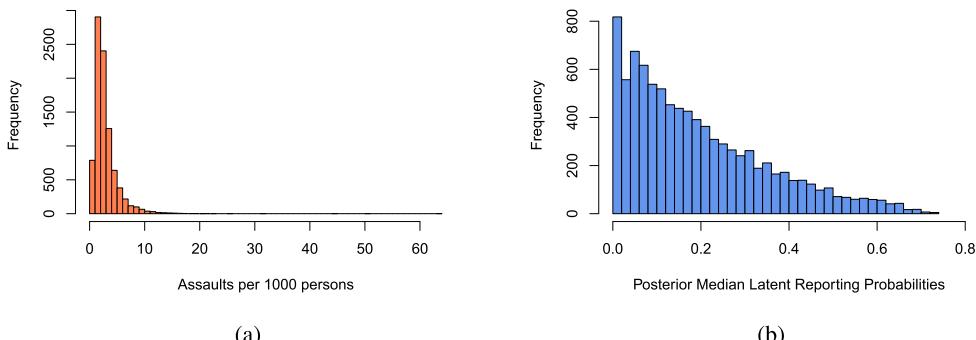


FIG. 12. Per-record posterior medians for incidence and reporting rate exhibit heterogeneity: (a) Posterior medians for incidence are centered slightly below the overall population value at 2.3 per 1000, but 25% of records have posterior median incidence greater than 3.4, and another 25% below 1.6. (b) Posterior medians for reporting probabilities are predominantly below 30%, but for some records are as high as 74%.

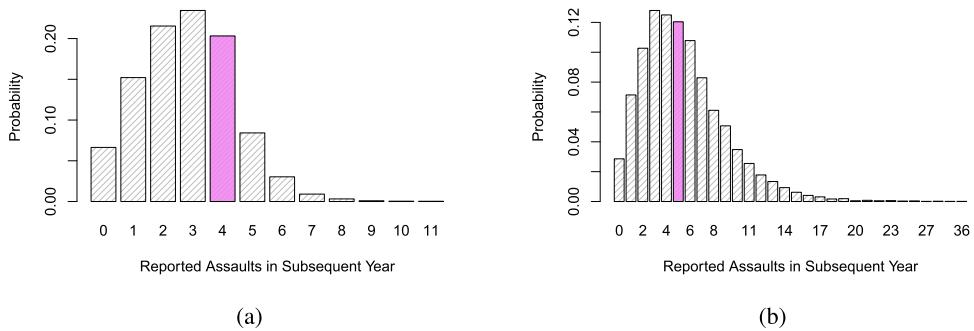


FIG. 13. *Variation in reporting probability alone is less likely to produce an increased number of reported assaults in subsequent year at: (a) Randolph College than at (b) University of South Carolina-Columbia. Solid-colored bar corresponds to number of reported assaults in 2019.*

Randolph College and the University of South Carolina-Columbia (USCC) had a similar number of reported assaults in 2019 (four and five, respectively). However, Randolph College's estimated reporting rate was considerably higher, with a posterior median of 53% compared to only 14% at USCC. Assuming that the true number of assaults is held constant, the predictive distributions for the number of reported assaults are shown in Figure 13.

Under this scenario there is a roughly 13% chance of an increased number of reported assaults at Randolph College in a subsequent year, due to variation in the reporting rate, and less than a 1% chance of the reported number to double to eight or more. At USCC, variation in reporting rate has a comparatively larger chance of producing an increased number of reported assaults (42%) and a 12% chance of doubling the number of reported assaults to 10 or more.

Ursinus College and Temple University provide another illustration of such school-level differences. Of the two, Ursinus had a much higher estimated reporting probability in 2019, and changes in reporting probability are comparatively less likely to drive an increase in reported assaults at Ursinus compared to Temple, as can be seen from their respective predictive distributions in Figure 14.

These disparities underscore the opaque nature of reported campus sexual assault data and the potential pitfalls of a one-size-fits-all approach to interpreting school-level reports over time.

6. Discussion. Colleges and universities prioritize reducing the incidence of sexual assault on their campuses, but underreporting diminishes the interpretability of their published

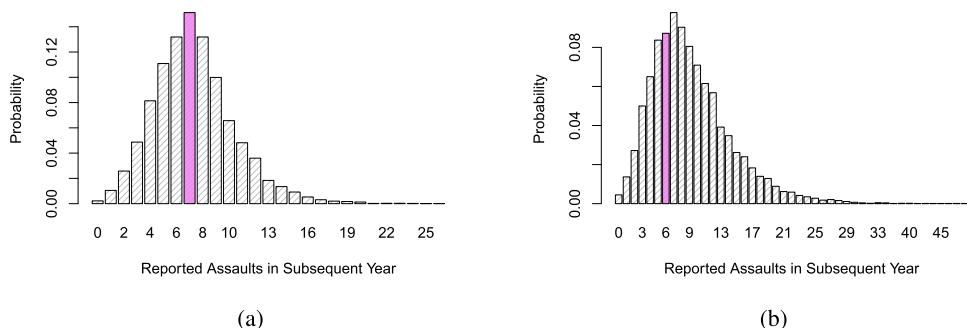


FIG. 14. *Without a change in the true number of assaults, the number of reported assaults has a 44% chance of increasing in a subsequent year at: (a) Ursinus College, and a 67% chance at (b) Temple University.*

sexual assault incidence data. This paper constructs a generative model of underreported campus sexual assault data, which allows for estimation of true incidence and reporting rates. Fitting this model reveals that lower socioeconomic status, as measured by the percentage of a school's undergraduates receiving Pell grants, appears associated with a lower probability of reporting. Status as a junior college or institution of religious instruction appear similarly associated with lower reporting probability, while gender imbalance in a school's student body is associated with lower true incidence of assault.

For the overall college population, estimated incidence of sexual assault remained fairly stable over 2014–2019, while reporting probabilities increased. Results at individual schools, however, varied widely. Estimates of school-level incidence and reporting probabilities may help university officials; for instance, Title IX coordinators assess the effectiveness of initiatives to reduce incidence of sexual assault or to familiarize students with reporting resources and decide where to make changes and improvements moving forward.

One avenue for future work concerns repeat victimizations of the same individual. High-frequency series victimizations, for instance, in the case of patterns of intimate partner violence, can produce extreme values in the reported number of assaults. As an example, in 2017 the University of Nebraska-Lincoln received 104 reports of sexual assault corresponding to the same victim and perpetrator. Such extreme outliers are not easily accommodated and, in particular, may violate the modeling assumption that the reporting decisions for each assault are independent (conditional on reporting probability). Violations of the independence assumption can degrade the quality of posterior inferences; see Appendix D in the Supplementary Material (Bradshaw and Blei (2024)) for further details. More broadly, the Bureau of Justice Statistics has acknowledged the difficulty of accurately incorporating series victimization in its national crime estimates (Lauritsen et al. (2012)).

Acknowledgments. The authors would like to thank the referees, Associate Editor, and the Editor for their constructive comments and suggestions that improved the quality of this paper.

David M. Blei is also affiliated with the Department of Computer Science at Columbia University.

Funding. The second author is supported by NSF IIS-2127869, NSF DMS-2311108, ONR N000142412243, and the Simons Foundation.

SUPPLEMENTARY MATERIAL

Appendices (DOI: [10.1214/24-AOAS1928SUPPA](https://doi.org/10.1214/24-AOAS1928SUPPA); .pdf). This file contains Appendices A–E to the main text.

Code and data (DOI: [10.1214/24-AOAS1928SUPPB](https://doi.org/10.1214/24-AOAS1928SUPPB); .zip). This file contains R and Stan implementations of methods from the main text, along with preprocessed data.

REFERENCES

BAILEY, T. C., CARVALHO, M. S., LAPA, T. M., SOUZA, W. V. and BREWER, M. J. (2005). Modeling of under-detection of cases in disease surveillance. *Ann. Epidemiol.* **15** 335–343. <https://doi.org/10.1016/j.annepidem.2004.09.013>

BETTENCOURT, L. M. A., LOBO, J., HELBING, D., KÜHNERT, C. and WEST, G. B. (2007). Growth, innovation, scaling, and the pace of life in cities. *Proc. Natl. Acad. Sci. USA* **104** 7301–7306.

BOX, G. E. P. (1980). Sampling and Bayes' inference in scientific modelling and robustness. *J. Roy. Statist. Soc. Ser. A* **143** 383–430. With discussion. [MR0603745](https://doi.org/10.2307/2982063) <https://doi.org/10.2307/2982063>

BRACHER, J. and HELD, L. (2021). A marginal moment matching approach for fitting endemic-epidemic models to underreported disease surveillance counts. *Biometrics* **77** 1202–1214. [MR4357831](https://doi.org/10.1111/biom.13371) <https://doi.org/10.1111/biom.13371>

BRADSHAW, C. and BLEI, D. M. (2024). Supplement to “A Bayesian Model of Underreporting for Sexual Assault on College Campuses.” <https://doi.org/10.1214/24-AOAS1928SUPPA>, <https://doi.org/10.1214/24-AOAS1928SUPPB>

CHANG, Y., CHOI, S. B., LEE, J. and JIN, W. C. (2018). Population size vs. Number of crimes: Is the relationship superlinear? *Int. J. Inf. Syst. Soc. Change* **9** 26–39.

DE OLIVEIRA, G. L., LOSCHI, R. H. and ASSUNÇÃO, R. M. (2017). A random-censoring Poisson model for underreported data. *Stat. Med.* **36** 4873–4892. [MR3734480](https://doi.org/10.1002/sim.7456) <https://doi.org/10.1002/sim.7456>

DE OLIVEIRA, G. L., OLIVEIRA, J. F., PESCARINI, J. M., ANDRADE, R. F. S., NERY, J. S., ICHIHARA, M. Y., SMEETH, L., BRICKLEY, E. B., BARRETO, M. L. et al. (2021). Estimating underreporting of leprosy in Brazil using a Bayesian approach. *PLoS Negl. Trop. Dis.* **15**.

DVORZAK, M. and WAGNER, H. (2016). Sparse Bayesian modelling of underreported count data. *Stat. Model.* **16** 24–46. [MR3457686](https://doi.org/10.1177/1471082X15588398) <https://doi.org/10.1177/1471082X15588398>

FADER, P. S. and HARDIE, B. G. (2000). A note on modelling underreported Poisson counts. *J. Appl. Stat.* **27** 953–964.

FERNÁNDEZ-FONTELO, A., CABANÁ, A., JOE, H., PUIG, P. and MORIÑA, D. (2019). Untangling serially dependent underreported count data for gender-based violence. *Stat. Med.* **38** 4404–4422. [MR4002439](https://doi.org/10.1002/sim.8306) <https://doi.org/10.1002/sim.8306>

FISHER, B. S., DAIGLE, L. E., CULLEN, F. T. and TURNER, M. G. (2003). Reporting sexual victimization to the police and others: Results from a national-level study of college women. *Crim. Justice Behav.* **30** 6–38.

GELMAN, A., MENG, X.-L. and STERN, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statist. Sinica* **6** 733–807. With comments and a rejoinder by the authors. [MR1422404](https://doi.org/10.1162/j.1063-4963.1996.60040001.x)

GELMAN, A. and RUBIN, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statist. Sci.* **457–472**.

GUTTMAN, I. (1967). The use of the concept of a future observation in goodness-of-fit problems. *J. Roy. Statist. Soc. Ser. B* **29** 83–100. [MR0216699](https://doi.org/10.2307/234607002)

JAYNES, E. T. (1968). Prior probabilities. *IEEE Trans. Syst. Sci. Cybern.* **4** 227–241.

JEFFREYS, H. (1946). An invariant form for the prior probability in estimation problems. *Proc. R. Soc. Lond. Ser. A* **186** 453–461. [MR0017504](https://doi.org/10.1098/rspa.1946.0056) <https://doi.org/10.1098/rspa.1946.0056>

KUMARA, S. and CHIN, H. C. (2005). Application of Poisson underreporting model to examine crash frequencies at signalized three-legged intersections. *Transp. Res. Rec.* **1908** 46–50.

LAURITSEN, J. L., GATEWOOD OWENS, J., PLANTY, M., RAND, M. R. and TRUMAN, J. L. (2012). Methods for Counting High-Frequency Repeat Victimization in the National Crime Victimization Survey. Bureau of Justice Statistics NCJ 237308.

LETOUZEY, F., DENIS, F. and GILLERON, R. (2000). Learning from positive and unlabeled examples. In *Algorithmic Learning Theory* (Sydney, 2000). *Lecture Notes in Computer Science* **1968** 71–85. Springer, Berlin. [MR1851968](https://doi.org/10.1007/3-540-40992-0_6) https://doi.org/10.1007/3-540-40992-0_6

LI, J. and HUGGINS, J. H. (2022). Calibrated model criticism using split predictive checks. arXiv preprint. Available at [arXiv:2203.15897](https://arxiv.org/abs/2203.15897).

MA, J. and LI, Z. (2010). Bayesian modeling of frequency-severity indeterminacy with an application to traffic crashes on two-lane highways. In *ICCTP 2010: Integrated Transportation Systems: Green, Intelligent, Reliable* 1022–1033.

MORAN, G. E., BLEI, D. M. and RANGANATH, R. (2019). Population predictive checks. ArXiv preprint. Available at [arXiv:1908.00882](https://arxiv.org/abs/1908.00882).

MORENO, E. and GIRÓN, J. (1998). Estimating with incomplete count data: A Bayesian approach. *J. Statist. Plann. Inference* **66** 147–159. [MR1617002](https://doi.org/10.1016/S0378-3758(97)00073-6) [https://doi.org/10.1016/S0378-3758\(97\)00073-6](https://doi.org/10.1016/S0378-3758(97)00073-6)

MORGAN, R. E. and KENA, G. (2018). Criminal Victimization, 2016: Revised. Bureau of Justice Statistics NCJ 252121.

MORGAN, R. E. and THOMPSON, A. (2020). Criminal Victimization, 2019. Bureau of Justice Statistics NCJ 255113.

MORGAN, R. E. and TRUMAN, J. L. (2018). Criminal Victimization, 2017. Bureau of Justice Statistics NCJ 252472.

OBAMA, B. (2014). Memorandum—Establishing a White House Task Force to Protect Students from Sexual Assault. The White House Office of the Press Secretary. 22 Jan 2014. <https://obamawhitehouse.archives.gov/the-press-office/2014/01/22/memorandum-establishing-white-house-task-force-protect-students-sexual-a> (accessed: 17 Aug 2023).

POWERS, S., GERLACH, R. and STAMEY, J. (2010). Bayesian variable selection for Poisson regression with underreported responses. *Comput. Statist. Data Anal.* **54** 3289–3299. [MR2727752](https://doi.org/10.1016/j.csda.2010.04.003) <https://doi.org/10.1016/j.csda.2010.04.003>

RUBIN, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Ann. Statist.* **12** 1151–1172. [MR0760681](https://doi.org/10.1214/aos/1176346785) <https://doi.org/10.1214/aos/1176346785>

SABINA, C. and HO, L. Y. (2014). Campus and college victim responses to sexual assault and dating violence: Disclosure, service utilization, and service provision. *Trauma Violence Abuse* **15** 201–226. <https://doi.org/10.1177/1524838014521322>

SCHMERTMANN, C. P. and GONZAGA, M. R. (2018). Bayesian estimation of age-specific mortality and life expectancy for small areas with defective vital records. *Demography* **55** 1363–1388. <https://doi.org/10.1007/s13524-018-0695-2>

SHANMUGAM, D. and PIERSON, E. (2021). Quantifying inequality in underreported medical conditions. ArXiv preprint. Available at [arXiv:2110.04133](https://arxiv.org/abs/2110.04133).

SINOZICH, S. and LANGTON, L. (2014). Rape and sexual assault victimization among college-age females, 1995–2013. Bureau of Justice Statistics NCJ 248471.

SOMANADER, T. (2014). President Obama Launches the “It’s On Us” Campaign to End Sexual Assault on Campus. The White House Blog 19 Sept 2014. <https://obamawhitehouse.archives.gov/blog/2014/09/19/president-obama-launches-its-us-campaign-end-sexual-assault-campus> (accessed: 17 Aug 2023).

STAN DEVELOPMENT TEAM (2023). RStan: the R interface to Stan. R package version 2.21.8.

STONER, O., ECONOMOU, T. and DRUMMOND MARQUES DA SILVA, G. (2019). A hierarchical framework for correcting under-reporting in count data. *J. Amer. Statist. Assoc.* **114** 1481–1492. [MR4047275 https://doi.org/10.1080/01621459.2019.1573732](https://doi.org/10.1080/01621459.2019.1573732)

TRUMAN, J. L. and LANGTON, L. (2014). Criminal Victimization, 2013 (Revised). Bureau of Justice Statistics NCJ 247648.

TRUMAN, J. L. and LANGTON, L. (2015). Criminal Victimization, 2014. Bureau of Justice Statistics NCJ 248973.

TRUMAN, J. L. and MORGAN, R. E. (2016). Criminal Victimization, 2015. Bureau of Justice Statistics NCJ 250180.

WINKELMANN, R. (1996). Markov chain Monte Carlo analysis of underreported count data with an application to worker absenteeism. *Empir. Econ.* **21** 575–587.

WU, K., DAHLEM, D., HANE, C., HALPERIN, E. and ZOU, J. (2023). Collecting data when missingness is unknown: A method for improving model performance given under-reporting in patient populations. In *Conference on Health, Inference and Learning* 229–242. PMLR.