Constructing Inverse Probability Weights for Continuous Exposures

A Comparison of Methods

Ashley I. Naimi,[a] *Erica E.M. Moodie,*[a] *Nathalie Auger,*[b] *and Jay S. Kaufman*[a]

Abstract: Inverse probability–weighted marginal structural models with binary exposures are common in epidemiology. Constructing inverse probability weights for a continuous exposure can be complicated by the presence of outliers, and the need to identify a parametric form for the exposure and account for nonconstant exposure variance. We explored the performance of various methods to construct inverse probability weights for continuous exposures using Monte Carlo simulation. We generated two continuous exposures and binary outcomes using data sampled from a large empirical cohort. The first exposure followed a normal distribution with homoscedastic variance. The second exposure followed a contaminated Poisson distribution, with heteroscedastic variance equal to the conditional mean. We assessed six methods to construct inverse probability weights using: a normal distribution, a normal distribution with heteroscedastic variance, a truncated normal distribution with heteroscedastic variance, a gamma distribution, a *t* distribution (1, 3, and 5 degrees of freedom), and a quantile binning approach (based on 10, 15, and 20 exposure categories). We estimated the marginal odds ratio for a single-unit increase in each simulated exposure in a regression model weighted by the inverse probability weights constructed using each approach, and then computed the bias and mean squared error for each method. For the homoscedastic exposure, the standard normal, gamma, and quantile binning approaches performed best. For the heteroscedastic exposure, the quantile binning, gamma, and heteroscedastic normal approaches performed best. Our results suggest that the quantile binning approach is a simple and versatile way to construct inverse probability weights for continuous exposures.

(*Epidemiology* 2014;25: 292–299)

Inverse probability weighting (IPW) estimators for marginal structural models are common in epidemiology. Such estimators can be used to estimate a standardized measure of effect for time-fixed exposures[1] and account for confounding and selection bias due to measured time-varying confounders affected by prior exposure.[2] The use of IPW has also been extended to address a number of other methodological issues, including time-modified confounding,[3] interaction and effect measure modification,[4,5] spillover effects,[6] and effect decomposition.[7]

Informally, weighted regression models estimate exposure effects in a "pseudo-population" obtained by weighting the observed data by the inverse of the conditional probability of the observed exposure, rendering exposure independent of measured confounders.[8,9] Cole and Hernán[10] have provided guidelines on how to construct inverse probability weights to estimate exposure effects. For a binary exposure, individual weights are constructed by estimating the predicted probability of the observed exposure, often using logistic regression. Weights are constructed by taking the inverse of these predicted probabilities. To obtain asymptotically unbiased exposure effect estimates, correct specification of the model for the weights is required.[8,9,11] To avoid excessively upweighting individuals, weights are usually stabilized by replacing the numerator with the marginal probability of the observed exposure.

Inverse probability weights for a continuous exposure are constructed in a similar fashion.[10–13] However, constructing inverse probability weights for a continuous exposure can be complicated by a number of issues not encountered in the binary exposure setting, including the need to identify a correct distributional form for the exposure, the need to account for nonconstant exposure variance (heteroscedasticity), and the need to deal with outliers that can make highly variable weights more likely.

Submitted 5 April 2013; accepted 6 September 2013.

From the [a]Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, QC, Canada; and [b]Institut national de santé publique du Québec, and Research Centre of the University of Montreal Hospital Centre, Montreal, QC, Canada.

The authors report no conflicts of interest.

Supported by Post-Doctoral Research Award from the Fonds de recherche du Québec-Santé (A.I.N.), Natural Sciences and Engineering Research Council University Faculty Award (E.E.M.M.), Career Award from the Fonds de recherche du Québec-Santé (N.A.), and supported by the Canada Research Chairs Program (J.S.K.).

SDC Supplemental digital content is available through direct URL citations in the HTML and PDF versions of this article (www.epidem.com). This content is not peer-reviewed or copy-edited; it is the sole responsibility of the author.

Correspondence: Ashley I. Naimi, Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Purvis Hall, 1020 Pine Ave West, Montreal, QC, Canada H3A 1A2. E-mail: ashley.naimi@mcgill.ca.

ISSN: 1044-3983/14/2502-0292
DOI: 10.1097/EDE.0000000000000053

We explore the performance of six methods to construct inverse probability weights for continuous exposures: a standard approach (based on a normal distribution), an approach based on the gamma distribution, a heteroscedastic normal approach, a heteroscedastic truncated normal approach, a heavy-tailed approach (based on the *t* distribution), and a quantile binning approach (based on exposure ranks). The gamma and heteroscedastic normal models may be more appropriate for exposure distributions in which the variance changes as a function of measured covariates. The quantile binning and *t* distribution approach may help mitigate issues with highly variable weights because of the way in which they handle outliers in the exposure's distribution. Finally, quantile binning may be more appropriate for markedly nonnormal or heteroscedastic exposure distributions because it does not rely on distributional assumptions. We assess the performance of each method in a simulation study using two exposures generated from data in the Québec birth file.

METHODS

Simulated Data

We used Monte Carlo simulation[14] to assess the performance of each method to construct IPWs for continuous exposures. We used data from the Québec birth file[15] to generate our simulated data. We extracted information on 1.3 million live singleton births that occurred between 1989 and 2008 (inclusive), including gestational age (in completed weeks), maternal and paternal date of birth, and parity (total number of previous deliveries). Gestational age at birth was used to create a preterm birth indicator defined as less than 37 completed weeks. We used the preterm birth indicator to set the marginal probability of simulated outcome.

We generated 3000 Monte Carlo samples of 1500 observations each, using covariate information sampled without replacement from the Québec birth file. Covariates were maternal age, paternal age, and parity, which we collectively denote as C. We simulated two exposures, each with a mean defined as:

$$\mu = 0.025 \times mage + 0.0025 \times page + 0.00125 \times mage \times page - 0.21 \times I(parity = 2) - 0.22 \times I(parity = 3) - 0.45 \times I(parity = 4) - 0.45 \times I(parity \geq 5),$$

where $I(\bullet)$ is an indicator function taking on a value of 1 when the argument $(\bullet)$ is true (0 otherwise), and *mage* and *page* are maternal and paternal age in years. The first exposure (homoscedastic) followed a normal distribution, defined as $X_1 = 15 + \mu + \varepsilon_1$, with $\varepsilon_1 \sim N(0,2)$. The second exposure (heteroscedastic) followed a Poisson distribution with mean = variance = μ. To make this exposure continuous, we added to it a standard normal random variable. We set negative exposure values equal to zero, yielding a zero-inflated distribution.

For each exposure, we generated a binary outcome following a Bernoulli distribution with

$$P(Y_j = 1 \mid X_j, C) = \{1 + \exp(-[\alpha_{0j} + \log(1.25)X_j + \log(1.7)mage^{1/2} + \log(1.5)page^{1/2} + \log(0.75)I(parity = 2) + \log(0.8)I(parity = 3) + \log(0.85)I(parity = 4) + \log(0.9)I(parity = 5)])\}^{-1}$$

where $\alpha_{01} = -11.5$ for exposure X_1 and $\alpha_{02} = -8.05$ for exposure X_2, chosen to render the marginal probability of Y approximately 0.08, which was the proportion of preterm births observed in the birth file. To determine the true marginal value for a single-unit increase in each exposure, we estimated the marginal odds ratio for each scenario using the method of marginal standardization[16] in a simulated data set with 4.5 million observations.

For each exposure scenario in each of the 3000 Monte Carlo samples, we estimated the marginal log-odds ratio for a single-unit increase in the exposure using generalized estimating equations to obtain robust standard error estimates,[17] with IPW to account for confounding.[9]

Statistical Analysis

Stabilized inverse probability weights (*sw*) were defined as:

$$sw = \frac{f_X(X;\mu_1,\sigma_1^2)}{f_{X|C}(X \mid C = c;\mu_2,\sigma_2^2)}, \quad (1)$$

where $f_{\bullet}(\bullet)$ denotes the probability density function with mean μ and variance σ^2, X a continuous exposure, and C the set of confounders. When the exposure is binary, the density functions in the numerator and denominator of the stabilized weights reduce to the probability mass functions for a binary random variable: $P(X = 1 \mid C = c)$ for exposed persons, and $P(X = 0 \mid C = c)$ for unexposed persons.[9] Thus, the stabilized weight equations for a binary exposure are a special case of (1).

The numerator of (1) is used to control the variability of the estimator.[11(p 584)] The marginal exposure density (ie, not conditional on C) is a common choice and yields a marginal estimate of the exposure-outcome effect.[10] When a model for the denominator of the stabilized weights is correctly specified, weighting by *sw* results in an asymptotically unbiased estimator for the true causal effect.[11,12] We constructed weights defined by *sw* using six approaches.

Normal IPWs

First, we assumed that $f_{\bullet}(\bullet)$ was a normal density and used two generalized linear models with an identity link and normal density function to estimate μ and σ^2 for the numerator and denominator of the stabilized weights:

$$\text{Numerator model}: X_i = \hat{\mu}_1 + \hat{\varepsilon}_{i1} \sim N(0, \hat{\sigma}_1^2)$$

$$\text{Denominator model}: X_i = \hat{\mu}_2 + \hat{\varepsilon}_{i2} \sim N(0, \hat{\sigma}_2^2),$$

where $\hat{\mu}_1 = \hat{\mu}_2 = \hat{\beta}_0 + \hat{\beta}C$, and where $\hat{\varepsilon}_{i1}$ and $\hat{\varepsilon}_{i2}$ are the residuals for individual i from the numerator and denominator model, respectively.

Standard IPWs assuming homoscedastic variance were created by using the predicted means $\hat{\mu}$ for each person. We used the deviance statistic divided by its degrees of freedom from a generalized linear model with a normal error distribution and identity link function as an estimate of $\hat{\sigma}$ for all persons. These estimates were used as parameter estimate values for the normal distribution, defined as:

$$\hat{f}_k(x_i) = \frac{1}{\hat{\sigma}_k \sqrt{2\pi}} \exp\left(\hat{\varepsilon}_{ik}^2 \Big/ 2\hat{\sigma}_k^2\right) \quad (2)$$

where (henceforth) $k = 1,2$ denotes the numerator or denominator model.

Inverse probability weights with heteroscedastic variance were implemented by further specifying the variance in (2) as a function of confounders C. Specifically, we defined a variance function[18] as $\hat{\sigma}_{ik} = \hat{\sigma}_k \exp(\hat{\theta}_k C_i)$ for the denominator models, and we generated a stabilized weight for each person using the predicted mean and variance conditional on covariates. We implemented two techniques to obtain heteroscedastic normal IPWs: a two-stage approach as in previous work[13] and a single-stage approach. We also used a heteroscedastic truncated normal model to construct weights by bounding (2) between a_k and b_k, and dividing (2) by $\Phi\{(b_k - \mu_k)/\sigma_k\} - \Phi\{(a_k - \mu_k)/\sigma_k\}$, where $\Phi\{\bullet\}$ is the normal cumulative distribution function. For each exposure, we chose $\min(X_k)-1$ and $\max(X_k)+1$ as the lower (a_k) and upper (b_k) bound, respectively.

Heavy-tailed IPWs

Heavy-tailed IPWs assuming homoscedastic variance were created by standardizing residuals from the normal generalized linear models for the numerator and denominator (defined above) and used to generate a numerator and denominator density value for each person from the t distribution, defined as:

$$\hat{f}_k^{(\eta)}(x_i) = \frac{\Gamma\left\{\frac{\eta+1}{2}\right\}}{\sqrt{\eta\pi}\,\Gamma\left\{\frac{\eta}{2}\right\}} \left(1 + \frac{\hat{e}_{ik}^2}{\eta}\right)^{-\frac{\eta+1}{2}} \quad (3)$$

where $\hat{e}_{ik}$ is the standardized residual, η represents the degrees of freedom for the t distribution, and $\Gamma\{\bullet\} = (\bullet - 1)!$ represents the gamma function for argument $\{\bullet\}$. For our analysis, we used $\eta = 1$, 3, and 5. Relative to the normal distribution, the t distribution's heavy tails should mitigate the large variation in the weights due to outliers in the far tails of the distribution.[19,20(p 465)] For example, as observed in the birth file, a person whose standardized residual value from the numerator and denominator model is 5.4 and 6.8, respectively, will have a weight of approximately 5,115 using the standard normal distribution, but a weight of 1.7 using the t distribution with 1 degree of freedom.

Gamma IPWs

Gamma IPWs were obtained by fitting a generalized linear model with log-link and gamma distribution to estimate a gamma scale parameter and each person's predicted exposure value given they were exposed. The density function for the gamma IPWs was defined as:

$$\hat{f}_k(x_i) = \frac{1}{x_i \Gamma\left(1/\hat{\kappa}_k^2\right)} \left(\frac{x_i}{\hat{\kappa}_k^2 \hat{\mu}_k}\right)^{1/\hat{\kappa}_k^2} \exp\left(-\frac{x_i}{\hat{\kappa}_k^2 \hat{\mu}_k}\right) \text{if } x > 0, \quad (4)$$

where κ is the gamma scale parameter. Importantly, the variance of a random variable following the gamma distribution is defined as a function of the mean ($\sigma^2 = \kappa^2 \mu^2$), and the gamma model may thus better account for heteroscedastic exposure distributions.

Quantile Binning IPWs

Finally, we used a quantile binning approach to approximate $f_{\bullet}(\bullet)$ by ranking X into j quantiles, and fitting a conditional cumulative logistic model to estimate the predicted probability of being in a given quantile.[21] Specifically, we estimated the denominator value using a model defined as

$$\log\left\{\frac{\alpha_{m+1} + \cdots + \alpha_j}{\alpha_1 + \cdots + \alpha_m}\right\} = \gamma_{0m} + \gamma C,$$

where α_m is the probability of being in category m. Predicted probabilities for being in category m were calculated as $\text{expit}\{\gamma_{0m} + \gamma C\} - \text{expit}\{\gamma_{0(m-1)} + \gamma C\}$, where $\text{expit}\{\bullet\} = 1/\{1 + \exp(-\bullet)\}$. In our study, we ranked X using $j = 10$, 15, and 20 quantiles. Therefore, the numerator for the weights (the marginal probability of falling into category j) is calculated as $1/j$.

Following standard practice,[10] we explored the impact of truncating the normal, normal heteroscedastic, and gamma weights at three levels: 1st and 99th, 5th and 95th, and 10th and 90th percentiles. In total, 26 weighting schemes were considered, based on six modeling approaches: a normal distribution (with no truncation and truncation at the 1st and 99th, 5th and 95th, and 10th and 90th percentiles), a t distribution (with 1, 3, and 5 degrees of freedom), a gamma model (untruncated and truncated at previously specified percentiles), a normal model with heteroscedastic variance (untruncated and truncated at previously specified percentiles), a truncated normal

model with heteroscedastic variance (untruncated and truncated at previously specified percentiles), and a cumulative logistic model for discretized outcomes (based on 10, 15, and 20 categories). Because the heteroscedastic exposure was characterized by an excess of zero values, we modified the above density functions to include a zero-inflated portion.[22]

We assessed the performance of each method using bias (defined as the average of the difference between the point estimate and the true value across Monte Carlo samples) and mean squared error (defined as squared bias plus the Monte Carlo variance of the estimate). Furthermore, we compare the distributions of the weights across all Monte Carlo samples. SAS version 9.3 (SAS Institute, Cary, NC) was used for all analyses. In the eAppendix (http://links.lww.com/EDE/A756), we provide SAS code illustrating how to construct inverse probability weights for a continuous exposure using all six approaches, as well as the SAS code used to conduct all simulations.

RESULTS

Table 1 shows descriptive statistics for the simulated data, including the four covariates sampled from the Québec birth file. These covariate statistic values can be used to replicate our simulation study.[23] Figure 1 shows histograms of the simulated exposures following a normal (panel A) and contaminated Poisson (panel B) distribution, with normal and nonparametric kernel density curves overlaid. In panel A, the kernel density curve and the normal density curve lay directly on top of one another. In panel B, the zero-inflated density portion is clearly seen at the *x*-axis origin. Figure 2 shows scatterplots of the residuals obtained from two confounder-adjusted ordinary least squares regression models for each simulated exposure. As expected, Figure 2A shows no pattern in the residuals, confirming a homoscedastic exposure distribution. Figure 2B depicts a marked fan shape opening to the right, confirming a heteroscedastic exposure distribution. Moreover, the set of data points below the fan-shaped portion of the plot reveals the zero-inflated nature of the heteroscedastic exposure points.

TABLE 1. Descriptive Statistics for the Simulation Data

Variable (Distribution)	Mean	Variance
Maternal age (normal)	29.84	21.60
Paternal age (normal)	32.52	30.45
Parity (Poisson)		
2	0.24	0.18
3	0.07	0.07
4	0.02	0.02
5+	0.02	0.02
X_1 (normal)	16.98	2.47
X_2 (Poisson)	2.06	2.76
Y_1 (Bernoulli)	0.08	0.07
Y_2 (Bernoulli)	0.09	0.08

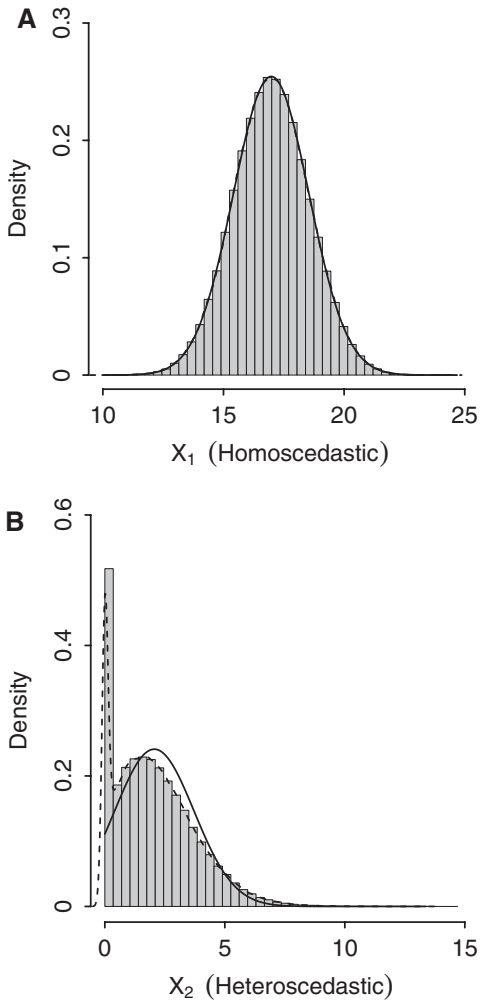

FIGURE 1. Histograms for both simulated exposures. A, normally distributed exposure (homoscedastic); B, contaminated Poisson exposure (heteroscedastic). Solid curve indicates normal density estimate; dashed curve, Kernel density estimate.

Table 2 describes the inverse probability weights for all 3000 Monte Carlo samples in our simulation study from each weighting scheme. With the exception of the weights from the heteroscedastic normal model for the heteroscedastic exposure, the mean of the weights obtained from all methods and both exposures was close to one. The mean of the weights from the heteroscedastic normal model for the heteroscedastic exposure was 1.56 (minimum = 0.00, maximum = 1.2×10^6). These extreme values were due to the influence of fewer than 1% of the data points in each Monte Carlo sample. Truncating the distribution of these heteroscedastic normal weights at the 1st and 99th percentile resulted in weights with a mean of 1.00 (minimum = 0.21, maximum = 4.01). For the heteroscedastic exposure, over all 4.5 million observations, 288 data points had heteroscedastic normal weights greater than the 25% of the Monte Carlo sample size ($1500\times0.25 = 375$). Of all other methods assessed, only the two-stage heteroscedastic normal

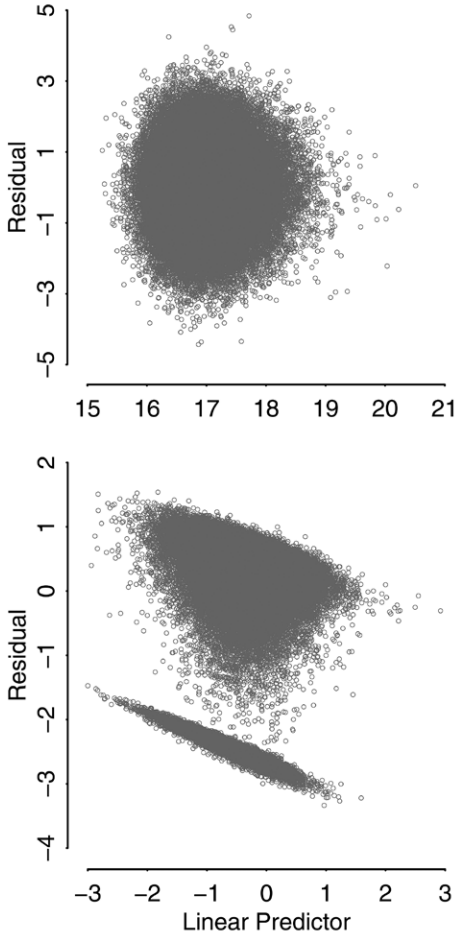

FIGURE 2. Standardized residual plots from ordinary least squares regression models conditional on all confounders for both simulated exposures. A, Homoscedastic exposure; B, heteroscedastic exposure.

approach yielded two observations (of all 4.5 million) with weights greater than 25% of the sample size, with values of 1,306.98 and 1,838.08.

Figure 3 shows the bias of various IPW methods from two weighted regression models for both simulated scenarios. For the heteroscedastic normal weights, we present results from the single-stage approach with weights truncated at the 1st and 99th percentile of the distribution. The eFigure (http://links.lww.com/EDE/A756) reproduces Figure 3 with all assessed methods included. For the normally distributed homoscedastic exposure, the standard normal and the heteroscedastic normal approaches were unbiased. All other methods were slightly biased, with the gamma and quantile binning approaches least so, as indicated by means (denoted by ×) that were closest to zero. For the heteroscedastic exposure, all methods were slightly biased, with the heteroscedastic normal (untruncated) and quantile binning approach (with 10, 15, and 20 categories) least so: the magnitude of bias for these approaches was 0.022, 0.011, 0.013, and 0.012, respectively.

TABLE 2. Distribution of Inverse Probability Weights for Both Simulated Exposures

	Homoscedastic	Heteroscedastic
Method	**Mean (min, max)**	**Mean (min, max)**
Normal	1.00 (0.00, 131.70)	1.01 (0.34, 10.13)
Weight truncation		
(1,99)	0.99 (0.24, 3.41)	1.00 (0.61, 2.22)
(5,95)	0.98 (0.49, 1.77)	1.00 (0.73, 1.50)
(10,90)	0.97 (0.61, 1.40)	1.00 (0.76, 1.30)
Heteroscedastic normal	1.00 (0.00, 64.94)	1.56 (0.00, 1.2×10^6)
t Distribution		
Degrees of freedom		
1	1.03 (0.14, 7.20)	1.01 (0.35, 3.26)
3	1.03 (0.07, 11.04)	1.01 (0.25, 3.82)
5	1.04 (0.04, 16.15)	1.01 (0.20, 4.14)
Gamma distribution	1.00 (0.00, 97.65)	0.99 (0.01, 41.20)
Quantile binning categories		
10	1.00 (0.11, 28.65)	1.00 (0.11, 16.41)
15	1.00 (0.08, 30.29)	0.98 (0.08, 264.83)
20	1.00 (0.06, 28.91)	0.95 (0.06, 74.58)

As expected, truncating the distribution of the weights at higher percentiles resulted in an increase in bias and a decrease in variance.[10] However, a decrease in the overall mean squared error was not observed with increasing truncation. As observed in Table 3, truncating the normal weights beyond the 1st and 99th percentile yielded no further decrease in mean squared error for either the homoscedastic or heteroscedastic exposures. The same was observed for the other heteroscedastic normal weights and gamma distribution weights.

DISCUSSION

We compared six methods for constructing inverse probability weights for a continuous exposure with simulated data generated from an empirical cohort. The first method assumed that the exposure followed a normal distribution with constant variance. This approach has been used in prior research to estimate the effect of epoetin dose on hematocrit levels,[24] and mortality in hemodialysis patients,[25] or the relation between loneliness and symptoms of depression.[26] We found that when the exposure followed a normal distribution with constant variance, this method was unbiased for the true exposure effect. When the exposure was heteroscedastic, using the homoscedastic normal approach resulted in a small bias.

We also used a heteroscedastic normal model to construct weights for both exposures. We assessed the performance of several techniques to implement this approach. The first was a two-stage approach previously implemented to assess the relation between the proportion of residents in a given neighborhood living in poverty and individual alcohol use.[13] This approach can easily be implemented with two regression

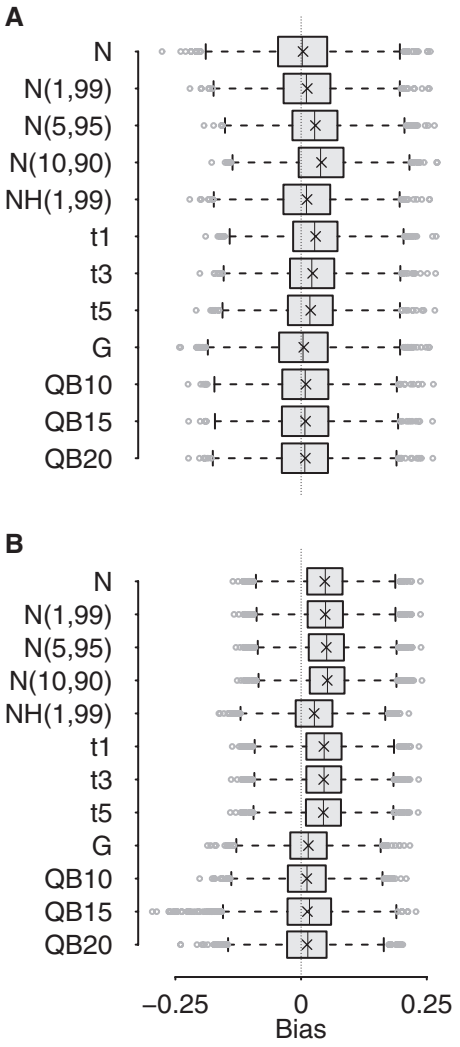

FIGURE 3. Bias of IPW methods from two marginal structural models in two simulated scenarios. A, Homoscedastic exposure; B, heteroscedastic exposure. N, standard normal distribution; N(1,99), truncated at 1st and 99th percentile; N(5,95), truncated at 5th and 95th percentile; N(10,90), truncated at 10th and 90th percentile; NH, normal distribution with heteroscedasticity; t1, *t* distribution with 1df; t3, *t* distribution with 3df; t5, *t* distribution with 5df; G, gamma distribution; QB10, quantile binning 10 categories; QB15, quantile binning 15 categories; QB20, quantile binning 20 categories; ×, mean value.

models: one for the conditional mean of the exposure, and a second for the mean of the squared residuals from the first model. The second was a single-stage approach based on joint maximization of the likelihood for the mean and variance models. This approach can be easily implemented using standard software (eg, SAS procedure) and is equivalent to the two-stage approach when the variance is independent of the mean function.[18] The third technique is based on constructing weights using a truncated normal distribution and joint maximization of the likelihood for the mean and variance models.[27(§12.4)] In our study, each of these methods yielded similar bias, and the major differences in mean squared error were due to very few influential observations with large weights.

TABLE 3. Mean Squared Error for Weighted Regression Model Parameters Obtained Using Different Inverse Probability–weighted Estimators for Two Simulated Exposures

	Homoscedastic	Heteroscedastic
Normal	0.0054	0.0049
Truncated		
(1,99)	0.0048	0.0050
(5,95)	0.0051	0.0052
(10,90)	0.0058	0.0053
Heteroscedastic normal	0.0054	0.0434
t Distribution		
Degrees of freedom		
1	0.0051	0.0047
3	0.0049	0.0047
5	0.0048	0.0047
Gamma distribution	0.0053	0.0033
Quantile binning categories		
10	0.0048	0.0032
15	0.0048	0.0052
20	0.0048	0.0036

Because the variance of a gamma random variable is defined as a function of the mean, the gamma distribution may, in principle, accommodate heteroscedasticity in the exposure. However, while the bias observed using the gamma distribution for the homoscedastic exposure was negligible, bias was more pronounced for the heteroscedastic distribution. One possible explanation lies in the fact that the gamma distribution approach requires a constant coefficient of variation (defined as the standard deviation of the exposure divided by the mean).[28 (p 285)] Indeed, further analyses revealed that the constant coefficient of variation assumption was met for the homoscedastic exposure but not for the heteroscedastic exposure. Thus, while the gamma distribution may not be generally useful to account for heteroscedasticity, it may be useful when the constant coefficient of variation assumption is met.

In principle, constructing inverse probability weights using the *t* distribution should mitigate the impact of outliers on the variation in the weights.[19,20(p 465)] Still, we found little difference between the approach using the *t* distribution and the standard normal approach after truncating the weights. This may be due to the fact that our analyses focused on the performance of methods under homoscedastic and heteroscedastic exposures. Thus, while outliers were present in the heteroscedastic exposure, they may not have been dramatic enough to observe important differences with the *t* distribution. Second, when dealing with outliers, it is possible that the benefit gained from truncating the distribution of the weights obtained from a standard or heteroscedastic normal approach is equal to that gained from using the *t* distribution. Future

research could compare the performance of weight truncation, the *t* distribution, and other robust methods or mixture models that may better account for the presence of outliers.

Finally, quantile binning is a relatively simple approach that makes no assumptions about the distribution of the exposure. This approach first ranks the exposure into quantiles. For zero-inflated exposures, one can also create a separate category for the zero portion and rank the continuous portion into quantiles (in our study, both approaches yielded similar results). One then fits a multinomial or polytomous logistic regression model to these categories to obtain the predicted probabilities of falling into the observed exposure category. We used a cumulative logistic regression model, which compares the probability of being above a certain threshold to the probability of being below.[29(p 301)] We also used a baseline-category logistic regression model, which compares the probability of being in a certain category to being in a chosen referent category.[29(p 293)] Both methods yielded similar results.

Marginal structural models with a binary exposure are common in epidemiologic research. Westreich et al[30] explored the finite-sample properties of marginal structural Cox proportional hazards regression models with a binary exposure under various confounding scenarios and found their performance followed expectations based on asymptotic theory. For a binary exposure, Lee et al[31,32] compared the performance of machine learning algorithms (including standard, pruned, bagged, and boosted classification and regression trees, and random forests) to standard logistic regression (with and without truncating) to construct inverse probability weights in a simulation study. They found that boosted classification and regression trees were superior to other methods in a majority of the scenarios they explored. While continuous exposures are common in epidemiology, no studies have explored the performance of various methods to construct inverse probability weights for a continuous exposure.

Continuous exposures follow a wide variety of parametric forms, including complex zero-inflated and other mixture distributions. Examples include asbestos exposure in occupational epidemiology,[33] epoetin dose among elderly hemodialysis patients,[25] and occlusion therapy for the treatment of childhood amblyopia.[34] To obtain an asymptotically unbiased and fully efficient inverse probability–weighted estimator with such exposures would require the use of complex mixture models that are not readily fit using most standard software packages. Moreover, the presence of heteroscedasticity would require accounting for the variance of the exposure as a function of relevant covariates. The absence of a relatively straightforward approach to construct inverse probability weights for complex, possibly heteroscedastic, exposure distributions may preclude the use of marginal structural models or lead to biased estimation in such settings. We show that a simple quantile binning approach provides relatively unbiased effect estimates with an acceptable mean squared error under various conditions compared with other methods.

It is important to note that, for the heteroscedastic exposure, all methods were biased for the true causal effect. However, the magnitude of the observed bias is comparable to the known finite-sample bias that characterizes the Cox proportional hazards regression model. Johnson et al[35] found that the Cox partial likelihood estimator was biased up to 6% in finite samples. This is comparable with the maximum bias observe in our study (5.5% for the truncated normal approach with weights truncated at the 10th and 90th percentile, eFigure, http://links.lww.com/EDE/A756). For the quantile binning approach, the largest bias we observed was 1.3% for the heteroscedastic exposure with 15 categories.

Our estimand of interest is the marginal effect of a single-unit increase in the simulated exposures. To estimate this parameter, we stabilized the weights using only the marginal probability density (distribution) function of the exposure. We could have achieved additional stability in the weights by conditioning the numerator on a selection of confounders, and then adjusting for these confounders in our regression model for the outcome. However, this would have changed our estimand to a covariate-conditional effect.[36] For purposes of interpretation, unconfounded marginal effect estimates are arguably of greater public health relevance.[37] Moreover, they may also be used to plot weighted dose-response functions that avoid complications that can arise in more standard practices.

Our simulation analyses were limited, in that we did not explore the performance of each method under a range of causal effects, confounding patterns, and sample sizes. Moreover, we restricted our attention to the time-fixed exposure setting. As with all simulation research, our results apply to the scenarios explored.[38] We did, however, use empirical data sampled from the Québec birth file to generate our simulation data, and we assessed a wide range of approaches to construct inverse probability weights for a continuous exposure. We show how sensitive these estimates can be when the true exposure distribution is heteroscedastic, but when methods that assume homoscedasticity are used. We found that quantile binning is a simple way to obtain inverse probability weights for marginal structural models and performs well whether the exposure distribution is normal with constant variance or markedly nonnormal with nonconstant variance.

ACKNOWLEDGMENTS

We thank Richard MacLehose and two anonymous reviewers for helpful comments on previous versions of the manuscript. We thank Miguel Hernán for providing SAS code to fit the heteroscedastic truncated normal model.

REFERENCES

1. Sato T, Matsuyama Y. Marginal structural models as a tool for standardization. *Epidemiology*. 2003;14:680–686.
2. Robins JM. Structural nested failure time models. In: Andersen P, Keiding N, eds. *The Encyclopedia of Biostatistics*. Chichester: John Wiley and Sons; 1998;4372–4389.
3. Platt RW, Schisterman EF, Cole SR. Time-modified confounding. *Am J Epidemiol*. 2009;170:687–694.

4. Chiba Y, Azuma K, Okumura J. Marginal structural models for estimating effect modification. *Ann Epidemiol.* 2009;19:298–303.
5. Vanderweele TJ, Vansteelandt S, Robins JM. Marginal structural models for sufficient cause interactions. *Am J Epidemiol*. 2010;171:506–514.
6. Tchetgen Tchetgen EJ, VanderWeele TJ. On causal inference in the presence of interference. *Stat Methods Med Res*. 2012;21:55–75.
7. VanderWeele TJ. Marginal structural models for direct and indirect effects (erratum in Epidemiology 2009; 20(4):629). *Epidemiol*. 2009;20:18–26.
8. Hernán MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiol*. 2000;11:561–570.
9. Robins JM, Hernán MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11:550–560.
10. Cole SR, Hernán MA. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol*. 2008;168:656–664.
11. Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran M, Berry D, eds. *Statistical Models in Epidemiology: The Environment and Clinical Trials*. Vol. 16. New York, NY: Springer-Verlag; 1999;95–134.
12. Hernán MA, Robins JM. Estimating causal effects from epidemiological data. *J Epidemiol Community Health*. 2006;60:578–586.
13. Cerdá M, Diez-Roux AV, Tchetgen ET, Gordon-Larsen P, Kiefe C. The relationship between neighborhood poverty and alcohol use: estimation by marginal structural models. *Epidemiology*. 2010;21:482–489.
14. Metropolis N, Ulam S. The Monte Carlo method. *J Am Stat Assoc*. 1949;44:335–341.
15. Auger N, Abrahamowicz M, Park AL, Wynant W. Extreme maternal education and preterm birth: time-to-event analysis of age and nativity-dependent risks. *Ann Epidemiol.* 2013;23:1–6.
16. Localio AR, Margolis DJ, Berlin JA. Relative risks and confidence intervals were easily computed indirectly from multivariable logistic regression. *J Clin Epidemiol*. 2007;60:874–882.
17. Liang KYL, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986;73:13–22.
18. Davidian M, Carroll RJ. Variance function estimation. *J Am Stat Assoc*. 1987;82:1079–1091.
19. Lange KL, Little RJA, Taylor JMG. Robust statistical modeling using the *t* distribution. *J Am Stat Assoc*. 1989;84:881–896.
20. Gelman A, Carlin J, Stern H, Rubin D. *Bayesian Data Analysis. Texts in Statistical Science*. Boca Raton, Fla: Chapman & Hall/CRC; 2004.
21. McCullagh P. Regression models for ordinal data. *J R Stat Soc B*. 1980;42:109–142.
22. Lambert D. Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics*. 1992;34:1–14.
23. Peng RD, Dominici F, Zeger SL. Reproducible epidemiologic research. *Am J Epidemiol*. 2006;163:783–789.
24. Cotter D, Zhang Y, Thamer M, Kaufman J, Hernán MA. The effect of epoetin dose on hematocrit. *Kidney Int*. 2008;73:347–353.
25. Zhang Y, Thamer M, Cotter D, Kaufman J, Hernán MA. Estimated effect of epoetin dosage on survival among elderly hemodialysis patients in the United States. *Clin J Am Soc Nephrol*. 2009;4:638–644.
26. VanderWeele TJ, Hawkley LC, Thisted RA, Cacioppo JT. A marginal structural model analysis for loneliness: implications for intervention trials and clinical practice. *J Consult Clin Psychol*. 2011;79:225–235.
27. Hernán MA, Robins J. *Causal Inference. Forthcoming*. Chapman/Hall, Available at: http://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/. Accessed 25 June 2013.
28. McCullagh P, Nelder JA. *Generalized Linear Models*. 2nd ed. London; New York: Chapman and Hall; 2000.
29. Agresti A. *Categorical Data Analysis. Wiley Series in Probability and Statistics*. 3rd ed. Hoboken, NJ: Wiley; 2013.
30. Westreich D, Cole SR, Schisterman EF, Platt RW. A simulation study of finite-sample properties of marginal structural Cox proportional hazards models. *Stat Med*. 2012;31:2098–2109.
31. Lee BK, Lessler J, Stuart EA. Improving propensity score weighting using machine learning. *Stat Med*. 2010;29:337–346.
32. Lee BK, Lessler J, Stuart EA. Weight trimming and propensity score weighting. *PLoS One*. 2011;6:e18174.
33. Naimi AI, Cole SR, Hudgens MG, Brookhart MA, Richardson DB. Assessing the component associations of the healthy worker survivor bias: occupational asbestos exposure and lung cancer mortality. *Ann Epidemiol.* 2013;23:334–341.
34. Moodie EE, Stephens DA. Estimation of dose–response functions for longitudinal data using the generalised propensity score. *Stat Methods in Med Res*. 2012;21:149–166.
35. Johnson ME, Tolley HD, Bryson MC, Goldman AS. Covariate analysis of survival data: a small-sample study of Cox's model. *Biometrics*. 1982;38:685–698.
36. Kaufman JS. Marginalia: comparing adjusted effect measures. *Epidemiol*. 2010;21:490–493.
37. Robins JM, Greenland S, Hu FC. Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome. *J Am Stat Assoc*. 1999;94:687–700.
38. Maldonado G, Greenland S. The importance of critically interpreting simulation studies. *Epidemiol*. 1997;8:453–456.