

A Quantile-Based g-Computation Approach to Addressing the Effects of Exposure Mixtures

Alexander P. Keil,^{1,2} Jessie P. Buckley,^{3,4} Katie M. O'Brien,² Kelly K. Ferguson,² Shanshan Zhao,⁵ and Alexandra J. White²

¹Department of Epidemiology, University of North Carolina, Chapel Hill, North Carolina, USA

²Epidemiology Branch, National Institute of Environmental Health Sciences, National Institutes of Health, Department of Health and Human Services, Research Triangle Park, North Carolina, USA

³Department of Environmental Health and Engineering, Johns Hopkins University, Baltimore, Maryland, USA

⁴Department of Epidemiology, Johns Hopkins University, Baltimore, Maryland, USA

⁵Biostatistics Branch, National Institute of Environmental Health Sciences, National Institutes of Health, Department of Health and Human Services, Research Triangle Park, North Carolina, USA

BACKGROUND: Exposure mixtures frequently occur in data across many domains, particularly in the fields of environmental and nutritional epidemiology. Various strategies have arisen to answer questions about exposure mixtures, including methods such as weighted quantile sum (WQS) regression that estimate a joint effect of the mixture components.

OBJECTIVES: We demonstrate a new approach to estimating the joint effects of a mixture: quantile g-computation. This approach combines the inferential simplicity of WQS regression with the flexibility of g-computation, a method of causal effect estimation. We use simulations to examine whether quantile g-computation and WQS regression can accurately and precisely estimate the effects of mixtures in a variety of common scenarios.

METHODS: We examine the bias, confidence interval (CI) coverage, and bias–variance tradeoff of quantile g-computation and WQS regression and how these quantities are impacted by the presence of noncausal exposures, exposure correlation, unmeasured confounding, and nonlinearity of exposure effects.

RESULTS: Quantile g-computation, unlike WQS regression, allows inference on mixture effects that is unbiased with appropriate CI coverage at sample sizes typically encountered in epidemiologic studies and when the assumptions of WQS regression are not met. Further, WQS regression can magnify bias from unmeasured confounding that might occur if important components of the mixture are omitted from the analysis.

DISCUSSION: Unlike inferential approaches that examine the effects of individual exposures while holding other exposures constant, methods like quantile g-computation that can estimate the effect of a mixture are essential for understanding the effects of potential public health actions that act on exposure sources. Our approach may serve to help bridge gaps between epidemiologic analysis and interventions such as regulations on industrial emissions or mining processes, dietary changes, or consumer behavioral changes that act on multiple exposures simultaneously. <https://doi.org/10.1289/EHP5838>

Introduction

Epidemiologists are increasingly confronted with arrays of unique, yet sometimes highly correlated, exposures of interest that may arise from similar sources and provide unique challenges to inference. The myriad analytic methods for mixtures data have been subject to a number of recent reviews and workshops about mixtures data (Braun et al. 2016; Carlin et al. 2013; Hamra and Buckley 2018). A clear theme has emerged that numerous questions can be answered in mixtures data, and different methods are best suited for different questions. For example, we might be interested in questions about the mixture (clustering of exposures, projecting exposures to a lower dimensional space) or questions about the health effects of individual components of the mixture or joint effects of multiple (perhaps all) components of the mixture. Here, we focus in particular on one of these many questions that can be asked of epidemiologic data on exposure mixtures: How can the mixture as a whole, rather than individual components, influence the health of the populations exposed to

the multitude of components in the mixture? Each of these components may act independently, synergistically, or antagonistically, or they may be inert with respect to the health outcomes of interest. Focusing on the mixture as a whole is advantageous, as it can provide simplicity of inference, integrate over multiple exposures that likely originate from similar sources, and often map directly onto the effects of potential public health interventions (Robins et al. 2004). The mixture effect is a useful metric even in light of policies that target single exposures. Particulate matter, for example, is a target of regulation by the U.S. Environmental Protection Agency, but the actual interventions that have reduced particulate matter occur on the sources of pollution, and such interventions rarely act on only one element of the mixture of air pollutants. Thus, better quantification of the effects of reducing particulate matter would account for joint reductions of other pollutants from particulate matter sources, such as sulfur dioxide and nitrogen oxides (from sources such as coal-fired power plants) or ozone and carbon monoxide (from sources such as automobile emissions).

One of the analytic approaches developed specifically for estimating effects of exposure mixtures, weighted quantile sum (WQS) regression, has become increasingly used as an analytic approach for exposure mixtures in relation to health outcomes (Yorita Christensen et al. 2013; Deyssenroth et al. 2018; Nieves et al. 2016). This approach is based on developing an exposure index that is a weighted average of all exposures of interest after each exposure is transformed into a categorical variable defined by quantiles of the exposures. The index, representing the exposure mixture as a whole, is then used in a generalized linear model to estimate associations with health outcomes. WQS regression has a specific goal of estimating the effect of the mixture as a whole. While several methods from the causal inference literature are capable of estimating the effect of a whole mixture

Address correspondence to A.P. Keil, Department of Epidemiology, CB 7435, University of North Carolina, Chapel Hill, NC 27599-7435. Telephone: 919-966-6652. Fax: 919-966-2089. Email: akeil@unc.edu

Supplemental Material is available online (<https://doi.org/10.1289/EHP5838>).

The authors declare they have no actual or potential competing financial interests.

Received 1 July 2019; Revised 28 February 2020; Accepted 24 March 2020; Published 7 April 2020.

Note to readers with disabilities: *EHP* strives to ensure that all journal content is accessible to all readers. However, some figures and Supplemental Material published in *EHP* articles may not conform to 508 standards due to the complexity of the information being presented. If you need assistance accessing journal content, please contact ehponline@niehs.nih.gov. Our staff will work with you to assess and meet your accessibility needs within 3 working days.

(Howe et al. 2012; Kang and Schafer 2007; Keil et al. 2018a; Robins et al. 2004), WQS regression has the advantage over these methods of a simple implementation. Although such approaches from the causal inference literature (e.g., inverse probability weighting and g-computation) can be implemented with standard regression software (Cole and Hernán 2008; Keil et al. 2014; Snowden et al. 2011), WQS regression is available in a single, self-contained R function that only requires specification of a single model (Czarnota et al. 2015).

Despite the attractive features of this approach, there are two notable reasons to be cautious when applying WQS regression. First, WQS regression requires what we call a directional homogeneity assumption, which assumes that all exposures have coadjusted associations with the outcome that are in the same direction (or can be coded *a priori* to meet this assumption) or are null. Second, WQS regression also assumes the individual exposures have linear and additive effects. Little is known about the demonstrable benefits of these assumptions when they are true and whether there are adverse impacts of these assumptions in realistic epidemiologic data where such assumptions would never be met exactly. Third, little theoretical statistical framework (Carrico et al. 2015) and few simulations (Czarnota et al. 2015) exist that assess the internal validity of effect estimation [e.g., bias, confidence interval (CI) coverage] of WQS. While other statistical approaches have been developed to estimate overall mixture effects that are discussed in previously noted reviews and workshops, we know of no methods other than WQS regression that can provide parsimonious parametric inference for the effect of a mixture of exposures (e.g., a low-dimensional mixture dose–response such as a single coefficient corresponding to the change in the outcome per unit of joint exposure). Thus, there is need for additional methods that give similar inference without making such strong assumptions, as well as an analysis of the conditions under which such assumptions are warranted and possibly beneficial.

In the current manuscript, we demonstrate a new approach to estimating the effects of an exposure mixture, which we call quantile g-computation, that shares the simplicity of interpretation and computational ease of WQS regression while not assuming directional homogeneity. Further, our approach inherits many features of causal inference methods, which allow for nonlinearity and nonadditivity of the effects of individual exposures and the mixture as a whole. We note explicit connections between these two approaches and demonstrate when they give equivalent estimates. We compare, using simulations, the validity of our approach and WQS regression for hypothesis testing, estimating the effects of a mixture, the control of confounding by correlated exposures, the impact of unmeasured confounding, and the impact of nonlinearity and nonadditivity on effect estimates.

Methods

We first describe WQS regression in relation to standard generalized linear models. We then introduce quantile g-computation as a generalization and extension to WQS regression.

WQS Regression

WQS regression developed gradually out of methods designed to estimate the relative contributions of exposures in a mixture to a single health outcome. WQS regression requires inputs similar to a standard regression model: an outcome (Y), a set of exposures of interest (X), and a set of other covariates or confounders of interest (Z). WQS starts by transforming X into a set of categorical variables in which the categories are created using quantiles of X (i.e., quantiles of each variable) as cut points, which we denote as

X^q . The output of WQS regression consists of two parts: a regression model between the outcome of interest and an index exposure (possibly adjusting for covariates), and a set of weights that describe the contribution of each exposure to the single index exposure and overall effect estimate.

For continuous Y measured in individual i , the regression model part of WQS regression can be expressed as

$$Y_i = \beta_0 + \psi S_i + \varepsilon_i,$$

where β_0 is the model intercept, S_i is the exposure index (defined below), ψ is the coefficient representing the incremental change in the expected value of Y per unit increase in the S_i , and ε_i is the error term. The exposure index is defined as in Equation 1 where w_j are the weights for each exposure (here we have d exposures) and X_{ji}^q is the quantized version of the j th exposure for the i th individual.

$$S_i = \sum_{j=1}^d w_j X_{ji}^q. \quad (1)$$

That is, if we use quartiles, then X_{ji}^q will equal 0, 1, 2, or 3 for any participant, corresponding to whether the exposure X_{ji} falls into the 0–25th, 25th–50th, 50th–75th, or 75th–100th percentile of that exposure. The weights are estimated as the mean weight across bootstrap samples of the WQS regression model (often in a distinct training set, which is discussed below) with all weights forced to sum to 1.0 and have the same sign [referred to as a non-negativity/nonpositivity constraint, which is enforced via constrained optimization or post hoc selection of positive coefficients from an unconstrained model (Carrico et al. 2015)]. These constraints imply that all exposures contribute to S_i in the same direction. As a consequence, for ψ to be an unbiased and consistent estimator of an overall effect of X^q on Y , we must assume directional homogeneity: All exposures must have the same direction of effect (inclusive of the null) with the outcome.

For intuition purposes, we can express a WQS regression as a standard linear model for the quantized exposures, given by Equation 2:

$$Y = \beta_0 + \psi \sum_{j=1}^d w_j X_j^q + \varepsilon = \beta_0 + \sum_{j=1}^d \beta_j X_j^q + \varepsilon, \quad (2)$$

where we have suppressed subscript notation for individuals for clarity. In fact, if the directional homogeneity assumption holds, then all β_j are, in fact, positive and need not be constrained (ignoring sampling variability), and the WQS regression approach is equivalent to a generalized linear model in very large samples [as suggested by Carrico et al. (2015)]. This equivalence is useful because generalized linear models do not require directional homogeneity, so this suggests, at the very least, that WQS regression will have similar inference to a generalized linear model, so generalized linear models might form the basis of an alternative to WQS regression.

There is little guidance on exactly how WQS regression results can be interpreted because the index effect is not well-defined in the sense that it does not map onto real-world quantities. However, we can make some progress on what WQS estimates in large samples when the underlying model (Equation 2) is correct and exposure effects are in the same direction. We start by noting that an increase in the exposure index value from S_i to $S_i + 1$ (the effect of which is estimated in a regression of the outcome on the exposure index) can be equated to a generalized linear model under directional homogeneity in the following derivation:

$$\begin{aligned}
\psi &= \psi(1 + S_i) - \psi S_i \\
&= \psi \left(\sum_{j=1}^d w_j + \sum_{j=1}^d w_j X_{ji} \right) - \psi \sum_{j=1}^d w_j X_{ji} \\
&= \left(\sum_{j=1}^d \psi w_j + \sum_{j=1}^d \psi w_j X_{ji} - \sum_{j=1}^d \psi w_j X_{ji} \right) \\
&= \sum_{j=1}^d \beta_j.
\end{aligned} \tag{3}$$

Thus, by Equation 3, the mixture effect ψ is equivalent to the sum of the β coefficients in the underlying linear model, and the extension to alternative link functions is trivial. Such an interpretation suggests that we might be able to estimate ψ using existing tools from causal inference. In small or moderate samples, WQS regression may not act as the large sample results suggest, even under directional homogeneity. This occurs due to sampling variability and because WQS regression utilizes sample splitting, whereby it first estimates the weights in the training set and then estimates the mixture effect in the validation set, given those weights. Thus, small/moderate sample performance of WQS is poorly understood because it has not been explored sufficiently in the literature.

Quantile g-Computation

One side effect of the large sample equivalence between generalized linear models and WQS regression (assuming directional homogeneity) is that, assuming there is no unmeasured confounding, WQS regression can be used to estimate causal effects if the linear model above is correct (which implies that quantization of X is appropriate). This holds because generalized linear models are often used as the basis of causal effect estimation (Cole et al. 2013; Daniel et al. 2011; Edwards et al. 2014; Keil and Richardson 2017; Keil et al. 2018b; Neophytou et al. 2016; Richardson et al. 2018; Robins 1986; Snowden et al. 2011) or the model parameters may be considered causal estimands (Daniel et al. 2013). From an intervention perspective (and assumptions noted below), each β_j is interpretable as the effect of increasing X_j by one unit, so ψ is interpretable as the effect of increasing all X_j by one unit at the same time. In fact, ψ is equivalent to the g-computation estimator (Snowden et al. 2011) of a joint marginal structural model for quantized exposures, which estimates the effect of increasing every exposure simultaneously by one quantile. That is, ψ can estimate a causal dose-response parameter of the entire exposure mixture. While g-computation has been described in detail by multiple authors in both frequentist and Bayesian settings (Keil et al. 2014, 2018a; Snowden et al. 2011; Taubman et al. 2009; Westreich et al. 2012; Young et al. 2011), we note here that it is a generalization of standardization that uses the law of total probability to compute estimates of the (usually population average) expected outcome distribution under specific exposure patterns. It is often combined with causal inference assumptions to select specific variables and models that, in practice, are used to predict, or simulate, outcome distributions that we would expect under different interventions on an exposure that may depend on time-fixed or time-varying factors.

Under standard causal identification assumptions [including causal consistency, no interference, and no unmeasured confounding, outlined in detail in Hernan and Robins (2010)] and correct model specification, g-computation (or the g-formula) can yield the expected outcome, had we been able to intervene on all exposures of interest (Robins 1986; Robins et al. 2004). Under these assumptions, g-computation can be used to estimate causal effects of time-fixed or time-varying exposures. As described later in this

section, g-computation can be used to estimate the parameters of a marginal structural model, which can quantify causal dose-response parameters such as the change in the expected outcome expected as exposures are increased (i.e., by manipulating exposures as in an intervention). The causal assumptions are universal to all inferential approaches (Greenland 2017); any analysis should strive to measure all confounders and specify models as accurately as is reasonably possible. We mention them here to be clear about how ψ can be interpreted. In the special case of time-fixed exposures that enter into the model only with linear terms (additivity and linearity), the g-computation estimator of ψ is given as the sum of all regression coefficients of the exposures of interest. ψ corresponds to the change in Y expected for a one-unit change in all exposures. Variance can be obtained using standard rules for estimating the sum of random variables and the covariance matrix of a linear model, which means that (in this simple setting) this approach requires little more computational time than a standard linear model.

When g-computation is performed using quantized exposures, we refer to the approach as quantile g-computation. Quantile g-computation allows us to estimate both ψ and the weights when the directional homogeneity assumption holds, but we will demonstrate that it allows valid inference regarding the effect of the whole exposure mixture, and individual contributions to that mixture, when directional homogeneity does not hold.

The first step of quantile g-computation is to transform the exposures X_j into the quantized versions X_j^q . Second, we fit a linear model (where we have omitted other confounders Z for notational simplicity, but they could also be included):

$$Y_i = \beta_0 + \sum_{j=1}^d \beta_j X_{ji}^q + \varepsilon_i.$$

Third, assuming directional homogeneity (for now), ψ is given as $\sum_{j=1}^d \beta_j$ (where β_j is the effect size for exposure j) and the weights for each exposure (indexed by k) are given as $w_k = \beta_k / \sum_j \beta_j$ (i.e., the weights are defined to sum to 1.0). When directional homogeneity does not hold, quantile g-computation redefines the weights to be negative or positive weights, which are interpreted as the proportion of the negative or positive partial effect due to a specific exposure (and positive and negative weights are defined to both sum to 1.0). When exposures may have nonlinear effects on the outcome, we can extend this approach to include (for example) polynomial terms for X_j^q in a model such as $Y_i = \beta_0 + \sum_{j=1}^d \beta_j X_{ji}^q + \sum_{j=1}^d \beta_{j+d} X_{ji}^q X_{ji}^q + \varepsilon_i$ or a model that uses indicator variables for each level of a quantized exposure variable. Under nonlinearity or when we include product terms between exposures (or between exposures and confounders), the weights themselves are not well-defined (because the proportional contribution of an exposure to the overall effect would then vary according to levels of other variables and no longer function as weights), and ψ is no longer simply a sum of the β coefficients. However, ψ is still easily estimable via standard g-computation algorithms [ψ is a parameter of a marginal structural model estimated via g-computation, which generalizes the approach of Snowden et al. to a multilevel joint exposure (Snowden et al. 2011)], and the variance can be estimated with a nonparametric bootstrap. Briefly [as described in Snowden et al. (2011)], this approach involves the following steps: *a*) fit an underlying model that allows individual effects of exposures on the outcome, including interaction and

nonlinear terms; *b*) make predictions from that model at set levels of the exposures; and *c*) fit a marginal structural model to those predictions. Under linearity and additivity, this algorithm yields equivalent estimates of ψ to the algorithm given above but at a higher computational burden. This further extension allows the effect of the whole mixture to be nonlinear as well (which is generally good practice if individual exposures are expected to have nonlinear/nonadditive effects). Extension to binary outcomes is based on a logistic model for $\Pr(Y=1|X^q)$, where ψ can represent log odds ratios or log risk ratios, and this approach can also be extended to survival outcomes by using the Cox proportional hazards model as the underlying model to estimate hazard ratios that quantify a mixture effect. These methods are implemented in the R package `qgcomp`.

As we show below, quantile g-computation can be used to consistently estimate effects of the exposure mixture in settings in which WQS regression may be biased or inconsistent but also yield equivalent estimates with WQS regression in large samples when its assumptions hold. In addition to avoiding the directional homogeneity assumption, quantile g-computation can be extended to settings in which the effects of exposures may not be additive (e.g., we might wish to include interaction terms among X^q) or nonlinear (e.g., we may wish to include polynomial terms among X^q), and the exposure mixture effect may also be nonlinear. WQS regression, on the other hand, assumes additivity among exposures and allows nonlinear effects that are restricted to polynomial terms for S_i . The index exposure is still derived under a linear model, however, so it is not strictly equal to the quantile g-computation estimate under nonlinearity of the mixture effect, which is based on the fit of a marginal structural model and retains its interpretation as the effect of the mixture.

To demonstrate aspects of quantile g-computation and to compare the performance of quantile g-computation and WQS regression across a range of scenarios, we performed a number of simulation analyses that assess bias and CI coverage for effect estimates, and power and type I error for hypothesis tests.

Simulation Methods

We simulate data on a mixture of d exposures and a single, continuous outcome, where we let d equal 4, 9, or 14 in sample sizes of 100 or 500 (and up to 5,000, when noted). We simulate exposures such that they exactly equal the quantiles. For example, X_1 in each setting is simulated as a multinomial variable that takes on values 0, 1, 2, and 3, each with probability 0.25, such that $X_1 = X_1^q$ in each case ($q=4$ in all scenarios). Note that this simulation scheme does not allow us to quantify potential bias from quantizing exposures that are measured on a continuous scale but enables us to isolate the factors we are interested in and avoid conflating bias from the estimation approach with bias from model misspecification.

Unless otherwise specified, we simulate the outcome according to the following model:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \varepsilon = \beta_0 + \psi w_1 X_1 + \psi w_2 X_2 + \varepsilon.$$

All other exposures $X_3 \dots X_d$ are assumed to be noise exposures (except in scenario 4 below) with no effect on the outcome (equivalently, $\beta_3 = \dots = \beta_d = 0$), and ε is simulated from a standard normal distribution [mean = 0, standard deviation (SD) = 1] for every scenario. The term noise exposures is used because, unless otherwise specified, they are not correlated with other exposures and do not contribute to the outcome. By varying d , we emulate exploratory studies of mixtures when some exposures may not affect the outcome but the number of noncausal exposures may vary by context.

Large sample simulation. We demonstrate, empirically, that WQS regression (under directional homogeneity) is equivalent to a generalized linear model fit with quantized exposures (and thus quantile g-computation) in large samples. To demonstrate the appropriate interpretation of WQS regression output and large sample equivalence between WQS regression and generalized linear when the directional homogeneity assumption, linearity, and additivity all hold, we performed a simulation in a single large data set ($n=100,000$) with four exposures, $\psi=5.0$, and $w_1 = \frac{\beta_1}{\psi} = 0.5$, $\frac{\beta_2}{\psi} = 0.25$, $\frac{\beta_3}{\psi} = 0.15$, $\frac{\beta_4}{\psi} = 0.1$. We analyze these data using WQS regression and quantile g-computation and report estimates of ψ , β , and w for each approach.

Small and moderate sample simulations. We also contrast our new method with WQS regression in simulation settings with small or moderate sample sizes (i.e., typical sizes of observational studies) and when the necessary assumptions of WQS regression may be violated.

To address our study questions, we analyze data simulated under the scenarios given in [Table 1](#) and described here:

1. $\beta_1 = \beta_2 = \psi = 0$, i.e., are null hypothesis tests ($H_0: \psi = 0$) valid in each approach when no exposures have effects on the outcome?
2. $\beta_1 = 0.25$, $\beta_2 = -0.25$, and $\psi = 0$, i.e., are null hypothesis tests ($H_0: \psi = 0$) valid in each approach when exposures have counteracting effects on the outcome (directional homogeneity does not exist)?
3. $\beta_1 = 0.25$, $\beta_2 = 0$, $\psi = 0.25$, $w_1 = 1.0$, and $w_2 = 0.0$, i.e., are estimates and CIs valid in each approach when directional homogeneity does exist and only a single exposure is causal?
4. $\beta_1 = \beta_2 = \dots = \beta_d = \frac{0.25}{d}$, $\psi = 0.25$, and $w_1 = w_2 = \dots = w_p = \frac{1}{p}$, i.e., are estimates and CIs valid in each approach when directional homogeneity does exist and all exposures have causal effects?
5. $\beta_1 = 0.25$, $\beta_2 \in \{-0.2, -0.1, -0.05\}$, $\psi \in \{0.05, 0.15, 0.25\}$, $\rho_{X_1, X_2} \in \{0, 0.4, 0.75\}$, where ρ_{X_1, X_2} is the Pearson correlation coefficient between X_1 and X_2 , i.e., are estimates and CIs valid in each approach when directional homogeneity does not exist due to negative copollutant confounding?
6. $\beta_1 \in \{0, 0.25\}$, $\beta_2 = 0.0$, $\beta_C = 0.5$, $\psi \in \{0, 0.25\}$, and $\rho_{X, C} \in \{0, 0.25, 0.75\}$, where β_C is the effect size of an unmeasured confounder C , and $\rho_{X, C}$ is the Pearson correlation coefficient between X_1 and C (where C is generated similarly to other exposures), i.e., what are the impacts of unmeasured confounding on the validity of each approach?
7. $\beta_1 = 0.25$, $\beta_2 = 0.25$, $\beta_{1,2} = -0.15$, $\psi_1 = 0.50$, and $\psi_2 = -0.15$, where $\beta_{1,2}$ is the coefficient for the product term $X_1 X_2$ in the underlying model, and ψ_1 and ψ_2 are the overall exposure coefficients for a linear and squared term for overall exposure, i.e., are estimates and CIs valid in each approach when exposures interact on the model scale and the overall exposure effect is nonlinear? For this analysis, the underlying model for quantile g-computation included a product term for $X_1 X_2$, and the overall exposure effect was allowed to be nonlinear by including linear and quadratic terms.
8. $\beta_1 = 0.25$, $\beta_2 = 0.25$, $\beta_{1,1} = -0.15$, $\psi_1 = 0.50$, and $\psi_2 = -0.15$, where $\beta_{1,1}$ is the coefficient for the product term $X_1 X_1$ in the underlying model, and ψ_1 and ψ_2 are the overall exposure coefficients for a linear and squared term for overall exposure, i.e., are estimates and CIs valid in each approach when exposures and the overall exposure effect are nonlinear? For this analysis, the underlying model for quantile g-computation included a term for $X_1 X_1$, and the overall exposure effect was allowed to be nonlinear by including linear and quadratic terms.

Table 1. Summary of simulation scenarios used to explore performance of quantile g-computation and WQS regression for small ($n = 100$)- or moderate ($n = 500$)-sized samples.

Simulation scenario ^a	β_1	β_2	$\beta_{1,1}$	$\beta_{1,2}$	β_C	ψ_1	ψ_2	ρ_{X_1, X_2}	$\rho_{X, C}$
1	0	0	0	0	0	0	0	0	0
2	0.25	-0.25	0	0	0	0	0	0	0
3	0.25	0	0	0	0	0.25	0	0	0
4 ^b	0.25/ d	0.25/ d	0	0	0	0.25	0	0	0
5	0.25	-0.2, -0.1, -0.05	0	0	0	0.05, 0.15, 0.2	0	0.0, 0.4, 0.75	0
6	0.25	0	0	0	0.5	0.25	0	0	0.75
7	0.25	0.25	0	-0.15	0	0.5	-0.15	0	0
8	0.25	0.25	-0.15	0	0	0.5	-0.15	0	0

Note: Table columns are as follows: β_C , true coefficient for unmeasured confounder C ; β_1 , true coefficient for X_1 ; β_2 , true coefficient for X_2 ; $\beta_{1,1}$, true coefficient for interaction term X_1X_2 ; $\rho_{X, C}$, true correlation between X_1 and unmeasured confounder C ; ρ_{X_1, X_2} , true correlation between X_1 and X_2 ; ψ_1 , true mixture effect (main term); ψ_2 true mixture effect (quadratic term).

^aEach scenario was repeated for sample sizes of 100 and 500 and a total number of exposures of 4, 9, and 14. Outcomes are simulated according to the model

$$Y = 0 + \sum_{j=1}^d \beta_j X_j + \beta_{1,1} X_1 X_1 + \beta_{1,2} X_1 X_2 + \beta_C C + \varepsilon; \text{ where } \varepsilon \sim N(0, 1).$$

^b d refers to the total number of exposures.

Simulating correlated, quantized exposures. For scenario 6, we used a novel approach to induce a correlation (ρ) between two exposures X_1 and X_2 of ρ_{X_1, X_2} by drawing values of $X_{2i} = X_{1i}$ with probability $\sqrt{(\rho_{X_1, X_2})}$ and otherwise drawing from other values of X_1 . This sampling scheme ensured that exposures took on values of 0, 1, 2, or 3 with equal probability so that modeling assumptions would be met in all simulations. The simulation scheme was identical when considering unmeasured confounding, with the exception that we did not adjust for the unmeasured confounder in analyses.

For all scenarios, we simulated z data sets and analyzed each of them using WQS regression [using the gWQS (version 2.0; Renzetti et al.) package defaults in R (version 2.13; R Development Core Team), weights assumed positive] and quantile g-computation [using the qqcomp (version 1.3; Keil) package defaults in R]. We report statistics relating to the ψ parameter from each approach: bias (the mean estimate minus the true value of ψ , which is known in advance), the square root of the mean variance estimate from each method across the z data sets [root mean variance (RMVAR)], the SD of the bias across all z data sets [Monte Carlo standard error (MCSE)], the 95% CI coverage (proportion of estimated CIs that included the true value), and the type I error (when the null hypothesis is true) or statistical power (when the null hypothesis is false). Our main results are for sample sizes of 500, but we evaluated alternative sample sizes ($n = 100, 2,000, \text{ and } 5,000$) in sensitivity analyses. A sample R code for simulation analyses is given in the Supplemental Material.

Results

Large Sample Simulation

In large samples ($n = 100,000$) in which the assumptions of WQS hold, WQS regression and quantile g-computation give identical, unbiased estimates of the overall exposure effect and the weights, although WQS regression yields a smaller test statistic (larger

p -value) of the null hypothesis test due to the use of sample splitting (Table 2). Weights are displayed graphically in Figures S1 and S2.

Small and Moderate Sample Simulations

Validity under the null hypothesis $\psi = 0$ when exposures a) have no effect (scenario 1), or b) exposures have counteracting effects (scenario 2). Under the null hypothesis when no exposures affect the outcome, both quantile g-computation and WQS regression yield type 1 error rates close to the alpha level of 0.05 at all values of d we examined (Table 3). When some exposures cause the outcome but there is no overall exposure effect due to counteracting effects of exposures ($\beta_1 = -\beta_2 = 0.25$), quantile g-computation provided a valid test of the null, whereas WQS regression was biased away from the null, and this bias appeared to increase with the number of noise exposures included in the model, which led to type I error rates $>90\%$. Results were similar at $n = 100$ (Table S1). We note that the results when $\psi = 0$ do not depend on whether effects are constrained to be positive or negative in WQS regression due to symmetry of the linear model.

Validity when $\psi \neq 0$ (scenarios 3 and 4). When all exposure effects were either positive or null, quantile g-computation provided unbiased effect estimates for the overall exposure effect, whereas WQS regression was biased away from the null (Table 3). For a single causal exposure (scenario 3), weighed quantile sum regression was more powerful, with power $>90\%$ for up to 14 total exposures, but 95% CIs had poor coverage (57–83%), while quantile g-computation provided valid CIs but at reduced power. When all exposures had positive effects (scenario 4), quantile g-computation was more powerful, and WQS regression was biased toward the null. Results were generally similar at $n = 100$, but RMVAR was lower than the MCSE for WQS regression at $n = 100$, indicating the standard error estimates were biased to be too small (Table S1).

Validity under copollutant and unmeasured confounding (scenarios 5 and 6). Under negative copollutant confounding, quantile g-computation was unbiased for all examined values of the total exposure effect and correlation between the causal exposures (Figure 1). Results were similar for both $n = 100$ and $n = 500$ and across the number of noise exposures included (Figures S3–S7). WQS regression was biased at all studied levels of confounding, and the bias increased with the strength of the negative confounder–outcome association and decreased with the correlation between the two exposures. After noting that the standard error of quantile g-computation estimates of ψ appeared to decrease at higher levels of exposure correlation, we repeated a simpler version of this scenario at higher

Table 2. Single simulation demonstrating equivalence between WQS and quantile g-computation in large samples ($n = 100,000$) when all exposures have effects in the same direction (true $\psi = 5.0$; true weights = 0.5, 0.25, 0.15, and 0.1).

Method	Mixture effect		Estimated weights			
	ψ	t -Statistic	w_1	w_2	w_3	w_4
WQS ^a	5.00	806	0.50	0.25	0.15	0.10
Q-gcomp ^b	5.00	884	0.50	0.25	0.15	0.10

^aWeighted quantile sum regression (R package gWQS version 2.0 defaults).

^bQuantile g-computation (R package qqcomp version 1.3 defaults).

Table 3. Validity of WQS regression and quantile g-computation under the null (no exposures affect the outcome or exposures counteract) and nonnull estimates when directional homogeneity holds; 1,000 simulated samples of $n = 500$. Corresponding estimates for $n = 100$ are provided in Table S1.

Scenario	Method	d^a	Truth ^b	Bias ^c	MCSE ^d	RMVAR ^e	Coverage ^f	Power/type 1 error ^g
1. Validity under the null, no exposures are causal	WQS ^h	4	0	0.00	0.09	0.09	0.95	0.05
		9	0	-0.01	0.13	0.12	0.94	0.06
		14	0	-0.01	0.15	0.15	0.95	0.05
	Q-gcomp ⁱ	4	0	0.00	0.08	0.08	0.94	0.06
		9	0	0.00	0.12	0.12	0.95	0.05
		14	0	-0.01	0.16	0.15	0.95	0.05
2. Validity under the null, causal exposures counteract	WQS ^h	4	0	0.32	0.08	0.08	0.02	0.98
		9	0	0.41	0.11	0.11	0.04	0.96
		14	0	0.46	0.14	0.14	0.09	0.91
	Q-gcomp ⁱ	4	0	0.00	0.09	0.09	0.95	0.05
		9	0	0.00	0.13	0.13	0.96	0.04
		14	0	-0.01	0.16	0.16	0.96	0.04
3. Validity under single nonnull effect	WQS ^h	4	0.25	0.07	0.07	0.07	0.83	1.00
		9	0.25	0.15	0.10	0.10	0.67	0.98
		14	0.25	0.21	0.14	0.13	0.57	0.94
	Q-gcomp ⁱ	4	0.25	0.00	0.08	0.08	0.94	0.88
		9	0.25	0.00	0.12	0.12	0.95	0.52
		14	0.25	-0.01	0.16	0.15	0.95	0.36
4. Validity under all nonnull effects with directional homogeneity	WQS ^h	4	0.25	-0.06	0.10	0.09	0.87	0.58
		9	0.25	-0.09	0.13	0.13	0.87	0.26
		14	0.25	-0.10	0.17	0.15	0.88	0.19
	Q-gcomp ⁱ	4	0.25	0.00	0.08	0.08	0.95	0.86
		9	0.25	0.00	0.12	0.12	0.95	0.55
		14	0.25	-0.01	0.15	0.15	0.95	0.37

Note: MCSE, Monte Carlo standard error; RMVAR, root mean variance.

^aTotal number of exposures in the model.

^bTrue value of ψ , the net effect of the exposure mixture.

^cEstimate of ψ minus the true value.

^dStandard deviation of the bias across 1,000 iterations.

^eSquare root of the mean of the variance estimates from the 1,000 simulations, which should equal the MCSE if the variance estimator is unbiased.

^fProportion of simulations in which the estimated 95% confidence interval contained the truth.

^gPower when the effect is nonnull (scenarios 3 and 4); otherwise (in scenarios 1 and 2), it is the type 1 error rate (false rejection of null), which should equal alpha (0.05 here) under a valid test.

^hWeighted quantile sum regression (R package gWQS defaults).

ⁱQuantile g-computation (R package qgcomp defaults).

correlations (up to 0.9; $n = 500$, $d = 4$, and $\beta_1 = \psi = 0.25$) for quantile g-computation only and observed that, following intuition, the CI width ($3.92 \times$ standard error) of β_1 increased with exposure correlation, but, counterintuitively, the CI width of the overall effect ψ decreased with exposure correlation (Figure 2).

With increasing numbers of noise exposures, bias due to unmeasured confounding increased with WQS regression, whereas bias was stable with increasing noise exposures across all sample sizes studied for quantile g-computation (Figure 3). After noting that the magnitude of bias for WQS regression seemed to depend on sample size, we increased the sample size in our analysis (up to 5,000). The difference between the two approaches diminished as the sample size increased but was present at all sample sizes.

Validity under nonlinearity and nonadditivity (scenarios 7 and 8). When exposure effects were nonlinear and nonadditive, WQS regression (even allowing for quadratic effects of the exposure index) yielded biased estimates of the quadratic exposure effects (the main effect of all exposures and the effect of all exposures squared), whereas quantile g-computation yielded unbiased estimates of exposure effects and associated variances (Table 4). Quantile g-computation also provided more precise estimates.

Discussion

One of the remarkable and unique aspects of WQS regression in the context of exposure mixtures is that it specifically estimates a

joint effect of the entire mixture: the effect of increasing all exposures by a single quantile. We show this approach to mixtures can leverage the existing correlation among exposures, where it may be difficult, if not impossible, to identify independent effects among correlated exposures (Snowden et al. 2011). Further, for exposures such as hazardous air pollutants (White et al. 2019), phthalates and parabens (Harley et al. 2016; Rudel et al. 2011), and metals (O'Brien et al. 2019), feasible interventions to address any individual exposure would likely affect multiple exposures, leading to natural interest in joint effects. We built on this framework of quantized exposures and a single joint effect by combining these elements with existing causal inference approaches. While this bridge is enabled by the large sample equivalence between WQS regression and quantile g-computation when the necessary assumptions of WQS regression are met, we were primarily motivated by developing an approach for estimating effects of mixtures in realistic sample sizes and when these crucial assumptions may not be met. To these ends, quantile g-computation maintains the simple inferential framework of WQS while providing effect estimates that are robust to routine problems of exposure mixtures.

Under the directional homogeneity, linearity, and additivity assumptions, we demonstrate that WQS regression can be interpreted as an ordinary least-squares linear regression model with a coefficient that corresponds to the expected change in an outcome from a simultaneous increase in all exposures by a single




Figure 1. Scenario 5: Impact of copollutant confounding on the bias of the overall exposure effect estimate ($n = 500$; $d = 9$) for quantile g-computation (q-gcomp) and weighted quantile sum (WQS) regression at varying exposure correlations (ρ_{X_1, X_2} of 0.0, 0.4, and 0.75) and varying total effect sizes ($\psi = \beta_1 + \beta_2 \in 0.2, 0.15, 0.05$). Boxes represent the median (center line) and interquartile range (outer lines of box) and outliers [points outside of the $1.5 \times$ interquartile range (IQR) length whiskers] across 1,000 simulations. Corresponding figures for $d = 4$ and 14 and $n = 100$ ($d = 4, 9, \text{ and } 14$) are provided in Figures S3–S7.

quantile. Further, the approach is quite powerful when a single exposure has an effect in the expected direction. However, this power seems illusory since it results from effect estimates that are biased away from the null and, in small samples, downwardly biased variance estimates. Further, we showed that if the multiple exposures affect the outcome, then WQS regression is biased downward and has reduced power relative to quantile g-computation. For investigators choosing between WQS and quantile g-computation, we note that quantile g-computation was less biased than WQS under every scenario we examined using simulations. Any set of simulations cannot be considered exhaustive, as individual substantive contexts vary widely. For example, in some scenarios, WQS demonstrated lower variance, suggesting that there may be a bias–

variance tradeoff when choosing between the methods under certain contexts (e.g., when no exposures have an effect, WQS may be more precise). We emphasize, however, that the information driving the choice between methods is not often known in the context of exposure mixtures, where individual effects are rarely known for all mixtures components, *a priori*. In such settings, quantile g-computation is less prone to bias and appears robust in the sense that it does not produce or enhance spurious results when the assumptions of WQS are incorrect.

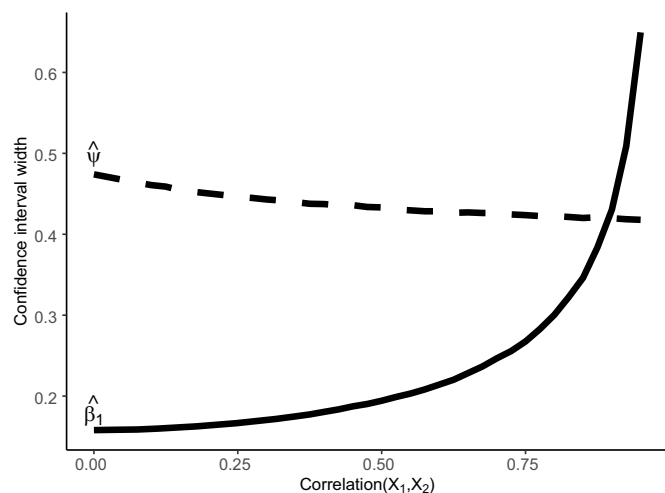


Figure 2. Scenario 5: Impact of copollutant confounding on the confidence interval width of individual exposure estimates (β) and the overall exposure effect estimate (ψ) for quantile g-computation (q-gcomp) under exposure correlations (ρ_{X_1, X_2}) from 0.0 to 0.9 ($n = 500$, $d = 9$, and $\psi = \beta_1 = 0.25$).

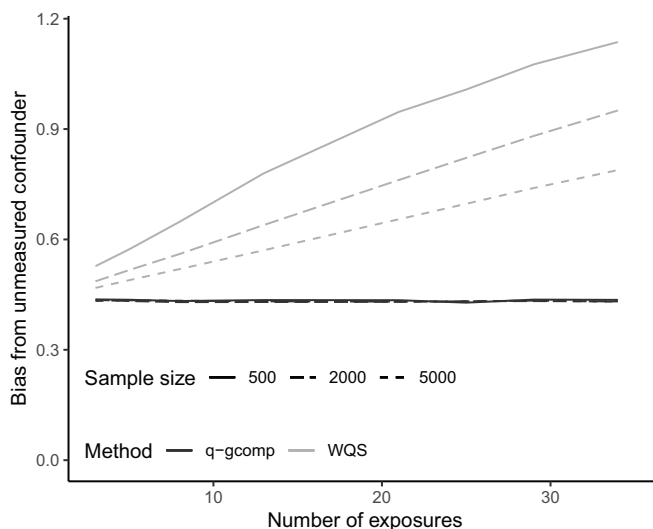


Figure 3. Scenario 6: Impact of unmeasured confounding on the bias of the overall exposure effect estimate (mean across 1,000 simulations; $n = 500$, 2,000, or 5,000) for quantile g-computation (q-gcomp) and weighted quantile sum (WQS) regression with confounder correlation ($\rho_{X, C}$) of 0.75, $\beta_C = 0.5$, and varying the total number of noise exposures ($d \in 4, 7, 9, 14, 22, 29, 35$). Note that all lines for quantile g-computation are overlapping and indicate unmeasured confounding bias is similar across sample sizes and number of exposures.

Table 4. Validity of WQS regression and quantile g-computation under nonnull estimates when directional homogeneity holds, individual exposure effects are nonadditive, and the overall exposure effect includes terms for linear (ψ_1) and squared (ψ_2) exposure (e.g., quadratic polynomial) for 1,000 simulated samples of $n = 500$. Corresponding estimates for $n = 100$ are provided in Table S2.

Scenario	Method	d^e	Bias ^b		MCSE ^c		RMVAR ^d	
			ψ_1	ψ_2	ψ_1	ψ_2	ψ_1	ψ_2
7. Validity when the true exposure effect is nonadditive/nonlinear	WQS ^e	4	0.21	-0.07	0.34	0.11	0.31	0.10
		9	0.21	-0.07	0.73	0.24	0.64	0.21
		14	0.13	-0.04	1.12	0.37	1.02	0.34
	Q-gcomp ^f	4	-0.01	0.00	0.13	0.04	0.13	0.04
		9	0.00	0.00	0.16	0.03	0.16	0.04
		14	0.00	0.00	0.19	0.04	0.18	0.04
8. Validity when the overall exposure effect is nonlinear due to underlying nonlinear effects	WQS ^e	4	-0.20	0.07	0.31	0.12	0.31	0.10
		9	-0.15	0.07	0.61	0.22	0.61	0.20
		14	-0.11	0.05	0.97	0.34	0.98	0.32
	Q-gcomp ^f	4	-0.01	0.00	0.15	0.04	0.15	0.04
		9	0.00	0.00	0.18	0.05	0.18	0.05
		14	0.00	0.00	0.20	0.05	0.20	0.05

Note: MCSE, Monte Carlo standard error; RMVAR, root mean variance.

^aTotal number of exposures in the model.

^bEstimate of ψ_1 or ψ_2 minus the true value.

^cStandard deviation of the bias across 1,000 iterations.

^dSquare root of the mean of the variance estimates from the 1,000 simulations, which should equal MCSE if the variance estimator is unbiased.

^eWeighted quantile sum regression (R package gWQS defaults, allowing for quadratic term for total exposure effect).

^fQuantile g-computation (R package qgcomp defaults, including an interaction term between X_1 and X_2 (scenario 7) or a term for X_1X_1 (scenario 8) as well as quadratic term for total exposure effect).

One aspect we did not address in our simulations was the use of sample splitting in WQS, where the data are divided into training/validation sets and the weights are first estimated in the training set and then applied in the validation set to estimate the overall mixture effect. In small/moderately sized studies, sample splitting has been useful in other domains (such as in neural network-based machine learning) to avoid overfit, and in large samples, sample splitting will not affect inference for consistent estimators of the weights under directional homogeneity (as demonstrated in our large sample simulation). The use of sample splitting allows the constraint on the weights to be on the other side of the null from the overall mixture effect. In the WQS R package (with default 40%/60% training/validation split), one can also estimate the weights and the overall mixture effect in the same data, although such a practice is generally discouraged. We repeated a limited version of our simulation scenario 1 in which we did this (only $n = 500$) and found that estimating the weights in the same data as the exposure effect is estimated in leads to null hypothesis tests for WQS regression (but not quantile g-computation, which does not utilize sample splitting) that are no longer valid with type I error rates >80% under some scenarios and bias that grows with the number of noise exposures included (Table S3).

Contrasted with quantile g-computation, WQS regression (that is, the implementation in the gWQS package) has three features: *a*) a nonnegativity/nonpositivity constraint on weights (necessitating the directional homogeneity assumption), *b*) bootstrap sampling to estimate weights, and *c*) sample splitting, which utilizes a training and validation set for calculation of the weights. As noted by Carrico et al., several of these elements have demonstrated some utility in algorithmic learning, suggesting that their use may improve some aspects of analysis over linear models (Carrico et al. 2015). However, it appears that the nonnegativity constraint (which implies the directional homogeneity assumption) is responsible for the bias away from the null we observed when directional homogeneity was violated or when there was a single causal exposure (scenarios 2, 3, and 5). Furthermore, we speculate that bootstrap sampling is responsible for the bias toward the null when there were multiple causal exposures (scenario 4). This latter result occurs because weights are 1 or 0 for every bootstrap sample

simply due to sampling variability, so we posit there is some bias toward 0.5 for the weights, which would equate to a bias toward the null in this scenario. This scenario suggests that, even if we could *a priori* change the coding of exposures such that all effects are in the same direction, WQS regression would still yield biased results with CIs that are too narrow. We also speculate that sample splitting is responsible for reduced power of WQS regression relative to quantile g-computation in some scenarios.

Exposure correlation in an epidemiologic context can indicate confounding, and we demonstrated that the nonnegativity constraint can magnify confounding bias. This result is intuitive. Using the example of fish consumption and cognitive functioning, it is likely that the consumption of docosahexaenoic acid (DHA) from fish improves cognitive health, but fish are also sources of mercury exposure. If we analyzed these exposures in relation to cognitive functioning together under a nonpositivity constraint, the effect of DHA would likely be forced to be near zero. This would be roughly equivalent to leaving DHA out of the model altogether, which would result in confounding in the negative direction for the effect of mercury. If we reversed the constraints, we would overestimate the benefits of fish consumption because we would induce confounding in the opposite direction. Other examples of such phenomena are abundant in dietary epidemiology, where beneficial and harmful components frequently exist in the same food. Similarly, well water may include beneficial aspects such as trace essential minerals (National Research Council Safe Drinking Water Committee 1980) and harmful constituents such as metals (Sanders et al. 2014). More generally, however, we found little apparent benefit with respect to bias or variance at imposing the directional homogeneity assumption when estimating the effects of mixtures, even when it was true. By avoiding the nonnegativity (or nonpositivity) constraints, quantile g-computation can give a more realistic estimate of the effect of the mixture as a whole. Further, while our analysis focuses on the overall mixture effect, our approach also yields β coefficients for an underlying generalized linear model, which can be interpreted as adjusted, independent effect sizes for quantized exposures (which will suffer variance inflation in the adjusted analysis, like any other method that estimates independent effects). Further, the mixture under consideration

could include only a subset of the measured exposures while controlling for other measured exposures. Such conditional mixture effects may be of interest in some settings where exposure sources can be identified that influence the levels of some, but not all, components of the mixture.

Beyond our simulations in which the mixture was defined to include all simulated exposures, the question “How do we define a mixture?” is not clearly answered. However, our results provide some guidance. If we have correlated exposures that are all causal, then we usually ought to (if possible) include all correlated exposures in the model to avoid copollutant confounding (as in the fish consumption example and our scenario 5 simulation). We should also account for interactions and nonlinearities, if they exist (as in our scenarios 7 and 8 simulations). Outside of copollutant confounding, however, we demonstrated that quantile g-computation is robust to the inclusion of noncausal noise variables, while bias due to unmeasured confounding actually increases with the inclusion of more noise variables for WQS regression. This phenomenon of increasing bias by including more covariates was also observed when no confounding existed and an increasing number of causal exposures were included, although the bias was toward the null in that setting (scenario 4). Thus, the bias magnification observed in scenario 5 is not akin to bias amplification, which is a phenomenon of predictably increased unmeasured confounding bias that occurs when examining independent effects of one exposure and adjusting for a strong correlate of that exposure (e.g., [Weisskopf et al. 2018](#)). Of the two methods, quantile g-computation appears robust to varying definitions of the mixture and so seems appealing in the context of undefined mixtures where we may not have good prior knowledge of the constituent effects.

Notably, our results suggest that both quantile g-computation and WQS regression, through their focus on the effect of the mixture as a whole, do not necessarily suffer greatly from exposure correlation. Intuitively, this is true because, as exposures become more correlated, we gain more information on the expected effects of increasing every exposure simultaneously in contrast with focusing on the effects of a single exposure while holding others constant.

Other methods could be used to estimate the effects of the mixture as a whole. For example, if we fit quantile-g computation without using quantized exposures, the approach would yield an effect of increasing every exposure by one unit—this is useful when one unit is meaningful for all exposures. In that case, our approach would simply be g-computation (also known as the g-formula), which has been previously used to estimate the impacts of hypothetical interventions in (among others) environmental and occupational (e.g., [Cole et al. 2013](#); [Edwards et al. 2014](#); [Keil et al. 2018b](#)) settings. More generally, g-computation provides a useful framework to estimate the joint effects of an exposure mixture ([Robins et al. 2004](#)), especially when exposures vary over time ([Keil et al. 2014](#)) or when issues such as exposure measurement error may be important ([Edwards and Keil 2017](#)). Thus, our approach is potentially extensible to such scenarios, and its utility is worthy of future study. There are examples of using such a framework with Bayesian kernel machine regression (BKMR) in a mixtures-specific setting ([Kupsco et al. 2019](#)). Quantile g-computation is a simpler version of these approaches but with the added strengths of being simple to implement and computationally frugal.

We posit that quantile g-computation (and the accompanying R package *qgcomp*) provide a simple framework that allows a flexible approach to the analysis of mixtures data when the overall exposure effect is of interest. Few methods explicitly estimate such effects. One alternative method is BKMR, which uses Gaussian process regression to fit a flexible function of the joint exposure set. In contrast, quantile g-computation forces the user to specify a parametric nonlinear model rather than assuming nonlinearity by default. Whereas quantile g-computation allows estimation of

parsimonious mixture dose–response parameters, BKMR (and many other data-adaptive or machine-learning approaches) does not yield a mixture dose–response parameter that could be compared with quantile g-computation in terms of the bias/variance tradeoffs quantified in our analysis. Instead, BKMR (as implemented in the R package *bkmr*) outputs a flexible prediction of the outcome at quantiles of all exposures, which cannot be expressed via a simple parametric function (which is required to assess the bias of a mixture dose–response parameter). Another way to understand this distinction is that quantile g-computation estimates the parameters of a joint marginal structural model, which quantify the average (or baseline confounder/modifier conditional) effects of modifying all exposures simultaneously. The interpretation of marginal structural models is well described in the epidemiologic literature (e.g., [Robins et al. 2000](#)). A limitation of this approach, as implemented in the *qgcomp* R package, is that if the underlying model is not smooth (e.g., if the dose–response of an individual exposure is a step function as when using indicator functions of individual exposures), the marginal structural model (which, in the current version of the package, can only be a polynomial function of exposures) may not adequately capture the dose–response function. In such settings, standard g-computation may be employed at a much higher computational and programming burden to the analyst. However, quantile g-computation may still provide useful approximations of overall effects of the mixture. Our simulations demonstrate that when the nonlinearities are known, quantile g-computation is unbiased. Whether this holds in realistic settings when nonlinearities are unknown will depend on the subject matter at hand and cannot be effectively explored via simulations without being tied more strongly to specific exposures and outcomes. We note that this is a potential limitation of parametric modeling in general and is not particular to our approach. However, when studying causal effects in general and mixture effects specifically, model specification must be accurate for accurate inference, which underscores the importance of allowing nonlinear and nonadditive effects of individual exposures ([Keil et al. 2014](#)). Understanding nonlinearity in the overall exposure effect is crucial for understanding whether effects are mainly limited to a certain range of joint exposures within the mixture, which informs whether interventions might be most impactful if they focused solely on individuals in a specific exposure range.

The reliance of the *qgcomp* package on existing generalized linear model frameworks allows one to easily include a variety of flexible nonlinear/nonadditive features such as polynomial functions, splines, interaction terms, or indicator functions on the individual quantized exposures (which are each described in the package documentation). While the utility of quantized exposures will depend on the context, the impact of characterizing rich exposure data as ordinal variables that are used as continuous regressors is broadly unknown. Our simulation did not assess the benefits and tradeoffs of using higher numbers of quantiles because the exposures were simulated on the quantized, rather than continuous, basis and little could be learned by increasing the numbers of exposure levels in our simulation. We hypothesize that quantization confers some benefits of similar nonparametric approaches such as rank regression, which are robust to outliers, but may suffer from difficulty in extrapolating results to other populations and reduced power when a linear model fits the data well. This is an area worth more research in the exposure mixtures context, where skewed exposure distributions may result in unduly influential, extreme exposure values, and nonlinear effects may be common.

We propose a method that builds on the desirable, simple output from WQS regression but is appropriate to use when the effects of exposure may be beneficial, harmful, or harmless. In scenarios

where we may not be able to rule out confounding or we may be uncertain about the effect direction of some exposures in the mixture, quantile g-computation is a simple and computationally efficient approach to estimating associations between a mixture of exposures and a health outcome of interest. Thus, our approach may serve as a valuable tool for identifying mixtures with harmful constituents or informing interventions that may prevent or reduce multiple exposures within a mixture.

Acknowledgments

This research was supported in part by the NIH/NIEHS grant numbers R01ES029531, R01ES030078 and the Intramural Research Program of the NIH, NIEHS grant number Z01ES044005.

References

- Braun JM, Gennings C, Hauser R, Webster TF. 2016. What can epidemiological studies tell us about the impact of chemical mixtures on human health? *Environ Health Perspect* 124(1):A6–A9, PMID: 26720830, <https://doi.org/10.1289/ehp.1510569>.
- Carlin DJ, Rider CV, Woychik R, Birnbaum LS. 2013. Unraveling the health effects of environmental mixtures: an NIEHS priority. *Environ Health Perspect* 121(1): A6–A8, PMID: 23409283, <https://doi.org/10.1289/ehp.1206182>.
- Carrico C, Gennings C, Wheeler DC, Factor-Litvak P. 2015. Characterization of weighted quantile sum regression for highly correlated data in a risk analysis setting. *J Agric Biol Environ Stat* 20(1):100–120, PMID: 30505142, <https://doi.org/10.1007/s13253-014-0180-3>.
- Cole SR, Hernán MA. 2008. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol* 168(6):656–664, PMID: 18682488, <https://doi.org/10.1093/aje/kwn164>.
- Cole SR, Richardson DB, Chu H, Naimi AI. 2013. Analysis of occupational asbestos exposure and lung cancer mortality using the g formula. *Am J Epidemiol* 177(9):989–996, PMID: 23558355, <https://doi.org/10.1093/aje/kws343>.
- Czarnota J, Gennings C, Wheeler DC. 2015. Assessment of weighted quantile sum regression for modeling chemical mixtures and cancer risk. *Cancer Inform* 14(suppl 2):159–171, PMID: 26005323, <https://doi.org/10.4137/CIN.S17295>.
- Daniel RM, Cousens SN, De Stavola BL, Kenward MG, Sterne JA. 2013. Methods for dealing with time-dependent confounding. *Statist Med* 32(9):1584–1618, PMID: 23208861, <https://doi.org/10.1002/sim.5686>.
- Daniel RM, De Stavola BL, Cousens SN. 2011. Gformula: estimating causal effects in the presence of time-varying confounding or mediation using the g-computation formula. *Stata J* 11(4):479–517, <https://doi.org/10.1177/1536867X1101100401>.
- Deyssenroth MA, Gennings C, Liu SH, Peng S, Hao K, Lambertini L, et al. 2018. Intrauterine multi-metal exposure is associated with reduced fetal growth through modulation of the placental gene network. *Environ Int* 120:373–381, PMID: 30125854, <https://doi.org/10.1016/j.envint.2018.08.010>.
- Edwards JK, Keil AP. 2017. Measurement error and environmental epidemiology: a policy perspective. *Curr Environ Health Rep* 4(1):79–88, PMID: 28138941, <https://doi.org/10.1007/s40572-017-0125-4>.
- Edwards JK, McGrath LJ, Buckley JP, Schubauer-Berigan MK, Cole SR, Richardson DB. 2014. Occupational radon exposure and lung cancer mortality: estimating intervention effects using the parametric g-formula. *Epidemiology* 25(6):829–834, PMID: 25192403, <https://doi.org/10.1097/EDE.0000000000000164>.
- Greenland S. 2017. For and against methodologies: some perspectives on recent causal and statistical inference debates. *Eur J Epidemiol* 32(1):3–20, PMID: 28220361, <https://doi.org/10.1007/s10654-017-0230-6>.
- Hamra GB, Buckley JP. 2018. Environmental exposure mixtures: questions and methods to address them. *Curr Epidemiol Rep* 5(2):160–165, PMID: 30643709, <https://doi.org/10.1007/s40471-018-0145-0>.
- Harley KG, Kogut K, Madrigal DS, Cardenas M, Vera IA, Meza-Alfaro G, et al. 2016. Reducing phthalate, paraben, and phenol exposure from personal care products in adolescent girls: findings from the HERMOSA Intervention Study. *Environ Health Perspect* 124(10):1600–1607, PMID: 26947464, <https://doi.org/10.1289/ehp.1510514>.
- Hernan MA, Robins JM. 2010. *Causal Inference*. Boca Raton, FL: CRC Press.
- Howe CJ, Cole SR, Mehta SH, Kirk GD. 2012. Estimating the effects of multiple time-varying exposures using joint marginal structural models: alcohol consumption, injection drug use, and HIV acquisition. *Epidemiology* 23(4):574–582, PMID: 22495473, <https://doi.org/10.1097/EDE.0b013e31824d1ccb>.
- Kang JDY, Schafer JL. 2007. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Statist Sci* 22(4):523–539, <https://doi.org/10.1214/07-STS227>.
- Keil AP, Daza EJ, Engel SM, Buckley JP, Edwards JK. 2018a. A Bayesian approach to the g-formula. *Stat Methods Med Res* 27(10):3183–3204, PMID: 29298607, <https://doi.org/10.1177/0962280217694665>.
- Keil AP, Edwards JK, Richardson DB, Naimi AI, Cole SR. 2014. The parametric g-formula for time-to-event data: intuition and a worked example. *Epidemiology* 25(6):889–897, PMID: 25140837, <https://doi.org/10.1097/EDE.0000000000000160>.
- Keil AP, Richardson DB, Westreich D, Steenland K. 2018b. Estimating the impact of changes to occupational standards for silica exposure on lung cancer mortality. *Epidemiology* 29:658–665, PMID: 29870429, <https://doi.org/10.1097/EDE.0000000000000867>.
- Keil AP, Richardson DB. 2017. Reassessing the link between airborne arsenic exposure among anaconda copper smelter workers and multiple causes of death using the parametric g-formula. *Environ Health Perspect* 125(4):608–614, PMID: 27539918, <https://doi.org/10.1289/EHP438>.
- Kupsco A, Kioumourtoglou MA, Just AC, Amarasiriwardena C, Estrada-Gutierrez G, Cantoral A, et al. 2019. Prenatal metal concentrations and childhood cardiometabolic risk using Bayesian kernel machine regression to assess mixture and interaction effects. *Epidemiology* 30(2):263–273, PMID: 30720588, <https://doi.org/10.1097/EDE.0000000000000962>.
- National Research Council Safe Drinking Water Committee. 1980. The contribution of drinking water to mineral nutrition in humans. In: *Drinking Water and Health*, vol. 3. Washington, DC: National Academies Press.
- Neophytou AM, Picciotto S, Costello S, Eisen EA. 2016. Occupational diesel exposure, duration of employment, and lung cancer: an application of the parametric g-formula. *Epidemiology* 27(1):21–28, PMID: 26426944, <https://doi.org/10.1097/EDE.0000000000000389>.
- Nieves JW, Gennings C, Factor-Litvak P, Hupf J, Singleton J, Sharf V, et al. 2016. Association between dietary intake and function in amyotrophic lateral sclerosis. *JAMA Neurol* 73(12):1425–1432, PMID: 27775751, <https://doi.org/10.1001/jamaneurol.2016.3401>.
- O'Brien KM, White AJ, Sandler DP, Jackson BP, Karagas MR, Weinberg CR. 2019. Do post-breast cancer diagnosis toenail trace element concentrations reflect prediagnostic concentrations? *Epidemiology* 30(1):112–119, PMID: 30256233, <https://doi.org/10.1097/EDE.0000000000000927>.
- Richardson DB, Keil AP, Cole SR, Dement J. 2018. Asbestos standards: impact of currently uncounted chrysotile asbestos fibers on lifetime lung cancer risk. *Am J Ind Med* 61(5):383–390, PMID: 29573442, <https://doi.org/10.1002/ajim.22836>.
- Robins J. 1986. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math Model* 7(9–12):1393–1512, [https://doi.org/10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6).
- Robins JM, Hernán MA, Brumback BA. 2000. Marginal structural models and causal inference in epidemiology. *Epidemiology* 11(5):550–560, PMID: 10955408, <https://doi.org/10.1097/00001648-200009000-00011>.
- Robins JM, Hernán MA, Siebert U. 2004. Effects of multiple interventions. In: *Comparative Quantification of Health Risks: The Global and Regional Burden of Disease Attributable to Major Risk Factors*. Ezziati M, Lopez AD, Rodgers A, Murray CJL, eds. Geneva, Switzerland: World Health Organization.
- Rudel RA, Gray JM, Engel CL, Rawsthorne TW, Dodson RE, Ackerman JM, et al. 2011. Food packaging and bisphenol A and bis(2-ethylhexyl) phthalate exposure: findings from a dietary intervention. *Environ Health Perspect* 119(7):914–920, PMID: 21450549, <https://doi.org/10.1289/ehp.1003170>.
- Sanders AP, Desrosiers TA, Warren JL, Herring AH, Enright D, Olshan AF, et al. 2014. Association between arsenic, cadmium, manganese, and lead levels in private wells and birth defects prevalence in North Carolina: a semi-ecologic study. *BMC Public Health* 14(1):955, PMID: 25224535, <https://doi.org/10.1186/1471-2458-14-955>.
- Snowden JM, Rose S, Mortimer KM. 2011. Implementation of G-computation on a simulated data set: demonstration of a causal inference technique. *Am J Epidemiol* 173(7):731–738, PMID: 21415029, <https://doi.org/10.1093/aje/kwq472>.
- Taubman SL, Robins JM, Mittleman MA, Hernán MA. 2009. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *Int J Epidemiol* 38(6):1599–1611, PMID: 19389875, <https://doi.org/10.1093/ije/dyp192>.
- Weisskopf MG, Seals RM, Webster TF. 2018. Bias amplification in epidemiologic analysis of exposure to mixtures. *Environ Health Perspect* 126(4):047003, PMID: 29624292, <https://doi.org/10.1289/EHP2450>.
- Westreich D, Cole SR, Young JG, Palella F, Tien PC, Kingsley L, et al. 2012. The parametric g-formula to estimate the effect of highly active antiretroviral therapy on incident AIDS or death. *Statist Med* 31(18):2000–2009, PMID: 22495733, <https://doi.org/10.1002/sim.5316>.
- White AJ, O'Brien KM, Niehoff NM, Carroll R, Sandler DP. 2019. Metallic air pollutants and breast cancer risk in a nationwide cohort study. *Epidemiology* 30:20–28, PMID: 30198937, <https://doi.org/10.1097/EDE.0000000000000917>.
- Yorita Christensen KL, Carrico CK, Sanyal AJ, Gennings C. 2013. Multiple classes of environmental chemicals are associated with liver disease: NHANES 2003–2004. *Int J Hyg Environ Health* 216(6):703–709, PMID: 23491026, <https://doi.org/10.1016/j.ijheh.2013.01.005>.
- Young JG, Cain LE, Robins JM, O'Reilly EJ, Hernán MA. 2011. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Stat Biosci* 3(1):119–143, PMID: 24039638, <https://doi.org/10.1007/s12561-011-9040-7>.