Assessing Exposures from the *Deepwater Horizon* Oil Spill Response and Clean-up

Patricia Stewart[1,13], Caroline P. Groth[2], Tran B. Huynh[3], Melanie Gorman Ng[4], Gregory C. Pratt[5] , Susan F. Arnold[5], Gurumurthy Ramachandran[6], Sudipto Banerjee[7], John W. Cherrie[8], Kate Christenbury[9], Richard K. Kwok[10,11], Aaron Blair[12], Lawrence S. Engel[10, 13] , Dale P. Sandler[10], Mark R Stenzel[14]

[1]Stewart Exposure Assessments, LLC, 6045 N. 27th. St., Arlington, VA USA 22207

[2]Department of Epidemiology and Biostatistics, School of Public Health, West Virginia University, Morgantown WV 26506 USA

[3]Department of Environmental and Occupational Health, Dornsife School of Public Health, Drexel University, Philadelphia, PA 19104 USA

[4]School of Population and Public Health, Faculty of Medicine, 3rd Floor, 2206 East Mall, VANCOUVER BC V6T 1Z3 Canada

[5]University of Minnesota, School of Public Health, Division of Environmental Health, Minneapolis, MN, USA 55455

[6]Department of Environmental Health and Engineering, Bloomberg School of Public Health, Johns Hopkins University, Baltimore, MD 21205 USA

[7]Department of Biostatistics, Suite: 51-254 CHS. UCLA Fielding School of Public Health, 650 Charles E. Young Drive South, Los Angeles, CA 90095-1772 USA

[8]Insitute of Occupational Medicine, Research Avenue North, Riccarton, Edinburgh, Midlothian EH14 4AP UK

[9]Public Health Sciences Social & Scientific Systems Inc., 4505 Emperor Blvd, Suite 400, Durham, NC, 27703, USA

[10]Epidemiology Branch, National Institute of Environmental Health Sciences, National Institutes of Health, 111 T.W. Alexander Drive – MD A3-05, Research Triangle Park, NC 27709 USA

[11]Office of the Director, National Institute of Environmental Health Sciences, 9000 Rockville Pike, Bethesda, MD 20892, USA

[12]National Cancer Institute, Bethesda, Maryland 20892 USA

[13]Department of Epidemiology, Gillings School of Global Public Health, University of North Carolina, Chapel Hill, NC, 27599 USA

[14]Exposure Assessment Applications, LLC, 6045 N. 27th. St., Arlington, VA, 22207 USA

Table of Contents

Excel File Tables

Table S1 Job-activity descriptions

Table S2 Time Periods

Table S3 THC, BTEX-H Inhalation Estimates by Broad Group

Table S4 $PM_{2.5}$ exposure estimates (µg/m3) May 15-July 15, 2010 from flaring and *in situ* burns

Table S5a. AgDISP generated 1-hour average total aerosol concentrations (µg/m3) at horizontal distances downwind, perpendicular to flight path

Table S5b. AgDISP generated 1-hour average total aerosol concentrations, (µg/m3) at distances downwind of 9500A following surface spraying

Table S6. Oil mist estimates by broad groups (for all time periods)

Table S7a. Estimated GuLF DREAM units (GDUs) from oil for broad exposure groups by chemical and time period

Table S7b. Estimated GuLF DREAM units (GDUs) from tar for broad exposure groups from tar by chemical and tme Period

Appendix A

Determinants of Exposure Used in the Development of the Exposure Groups (EGs)[1]

Job/Activity (N)	Location	Comment
Job title (n= 121)	4 rig vessels	Had collected specific job titles in the questionnaire due to the number of possible study participants (1500) and the identification of job title in the measurement data
Vessel (n=52)	• 14 Vessels piloting remotely operated vehicles (ROV vessels) + "All ROVs" • 3 burner fire control vessels • 33 research vessels + "All RVs"	Substantially smaller pool of workers per vessel (generally 60). All workers on a vessel were in the same EG. Specific vessels were identified from the questionnaire and the measurement data.
Activity-task (n=67)	All water and land areas	Job title often non-specific in measurement data ("responder" or "deck hand"). Used activities identified from the measurement data in the questionnaire. See Supplemental Material (SM), Table S1.
Area on the water (N=7)	• The hot zone (within 1 nautical mile (nmi, 1.8 km) of the wellhead) • source (5 nmi (~9 km) of the wellhead excluding the hot zone) • offshore (the area generally outside the source area to within 3 nmi (~5.5 km) of the shoreline • near shore ($\leq$3 nmi to the shore) • beaches and marshes • ports and docks • "other" land areas	Weathering played a substantial role in the changing composition of the oil. Distance from the source of fresh oil is, in part, reflective of the amount of weathering the oil had undergone.
State (N=4)	• Louisiana (LA), • Mississippi (MS), • Alabama (AL) • Florida (FL)	State indicates distance from the fresh oil source, similar to area on the water and so is reflective, in part, of oil weathering
Time Period (N=7).		Weathering also affects oil composition over time. In addition, different activities occurred over time. See SM, Table S2

1See Stenzel et al., 2021a for more detail

Appendix B

Statistical Methods

1 Bayesian modeling

To account for the censored values and to provide accurate exposure estimates, we investigated various statistical methods for analyzing censored data and conducted computer simulation studies to identify a method that was most suitable for our needs (Huynh et al. 2014, 2016). A Bayesian method was selected because it was found to provide low relative bias (<15%) and low imprecision (root mean squared error(<65%)) under the conditions of our study (measurements often with small N and/or high censoring and often high geometric standard deviations (GSDs)), while providing statistically credible intervals (the Bayesian analogue of classical confidence intervals) for all parameters of interest: the arithmetic mean (AM), geometric mean (GM), GSD, and the 95th percentile (Huynh et al. 2016).

Total hydrocarbons (THC) was modelled separately from benzene, toluene, ethylbenzene, xylene and n-hexane (BTEX-H). In addition, values for n-hexane were imputed where the chemical was not analyzed on the same samples as THC and the BTEX chemicals. Throughout this document, we assume log-normality of all of these chemicals, which is reasonable based on existing literature in the environmental health sciences and industrial hygiene.

In this document, we provide statistical details of modeling strategies used in Ramachandran et al. (2021) and Huynh et al. (2020 a, b and 2021) to model exposures to THC and BTEX-H. All modeling was done in R using rjags (R Development Core Team 2015; Plummer 2003).

2 THC modeling strategy

For modeling THC, we develop a Bayesian analysis of variance (ANOVA) where we model the mean and variance of each job group (also known as a intercept-only mean model). Bayesian models require the specification of priors; informative uniform priors, i.e. priors with bounds, help stabilize estimates in cases of limited information (i.e. very small sample size and/or high level of censoring). For the purposes of this modeling approach, all job groups were modeled separately. We elucidate the model for one job group.

Let Y_i be the natural log of THC for the it-h observation in a particular job group. Define μ to be the mean of the natural log of THC and σ^2 to be the variance. Let $\text{LOD}_i(Y)$ be the analytic limit of detection (LOD) for observation i of THC (on the natural log scale). Then, let $O_Y = \{i : Y_i > LOD_i(Y)\}$ denote the observed set of measurements of Y and $C_Y = \{i : Y_i \leq LOD_i(Y)\}$ denote the censored set of measurements (measurements below the respective LOD). Then, assuming standard definitions of the normal distribution

($N()$) and the uniform distribution ($Unif(,)$) we can write the following joint distribution for our model:

$$Unif(\mu \mid \ln(0.025), \ln(50)) \times Unif(\sigma \mid \ln(1.01), \ln(12)) \times \prod_{i \in O_Y} N(Y_i \mid \mu, \sigma^2) \times \prod_{i \in C_Y} \Phi\left(\frac{\text{LOD}_i(Y) - \mu}{\sigma}\right) , \quad (1)$$

where $\Phi(Z)$ is the cumulative distribution function (CDF) of Z. The CDF allows us to account for censored measurements. We assign informative priors to μ (corresponding to the natural log of the GM) and σ (corresponding to the natural log of the GSD). We specified the following uniform priors for μ or ln(GM) and σ or ln(GSD) as given below:

$$\mu \sim Unif(ln((0.025), ln(50)) \quad (2)$$

$$\sigma \sim Unif(ln(1.01), ln(12)) \quad (3)$$

The minimum for μ was set at ln(0.025) (a ∼28% reduction of the lowest LOD among the THC measurements (0.035 ppm)). The maximum value of μ was based on the Responsible Party's policy of taking specific actions to reduce workers' exposures when continuous levels of volatile organic compounds (VOCs) measured by direct-reading instruments on the rig vessels exceeded 100 ppm (approximately equivalent to 80 ppm THC) for 15 mins. Therefore, the upper GM prior was conservatively specified as 50 ppm THC as a full-shift exposure. Because a GSD cannot be lower than 1, 1.01 was defined as the minimum GSD. The maximum was set from an early analysis of the rig measurements (Huynh et al., 2016), where we found that approximately 98% of the GSDs were less than 12.

2.1 Markov Chain Monte Carlo Sampling

Markov Chain Monte Carlo (MCMC) sampling enables us to quantify uncertainty around the desired summary statistic (AM, GM, GSD, and 95th Percentile). In order to ensure convergence, a *burn-in* period, which is a set number of iterations, is designed to allow the parameters to converge. To develop the THC estimates we used 5000 posterior estimates of the GMs and GSDs. From each set of 5000 GM and GSD posterior samples, we calculated the associated (5000) AMs and 95th percentiles (Finney, 1941). The (posterior) median values and the 2.5 and 97.5 quantiles were selected to represent the statistic and the credible intervals (uncertainty interval in Bayesian statistics).

3 BTEX-H modeling strategy

To estimate BTEX-H, we used a bivariate Bayesian regression framework to jointly model THC and each of the BTEX-H chemicals individually (Groth et al., 2017). This approach allowed us to use THC, which had the lowest level of censoring, to predict exposures for each of the BTEX-H chemicals using a simple linear regression relationship of the natural log-transformed values while accounting for censoring in both

the predictor (THC) and the response (each of the BTEX-H chemicals). Censored BTEX-H measurements were estimated at each iteration using information in THC, the censoring status of other (THC and BTEX-H) measurements, and the linear relationship of the log transformed THC and each log transformed chemical (used as the Bayesian priors) identified by the model (Groth et al. 2017; 2018). (It should be noted that these models also will each produce estimates for THC. That said, we only present (and used) estimates for THC obtained from the aforementioned univariate analysis in Section 2).

3.1 Modeling details

To build this model, we start with a simple linear regression. Let Y_i be the natural log of the chemical of interest and X_i be the natural log of THC (both chemicals in ppb) for the ith measurement in a particular job group. All job groups were modeled separately. Then we can write the following linear regression expression with intercept β_0 and slope β_1 (collectively known as $\boldsymbol{\beta}$) as

$$Y_i = \beta_0 + \beta_1 X_i + \epsilon_i \tag{4}$$

where we assume $\epsilon_i \overset{iid}{\sim} N(0, \sigma^2_{Y|X})$ where $\sigma^2_{Y|X}$ is the variance of error terms ϵ_i.

Since our X is the natural log of THC, it appears reasonable to assume that X is also normally distributed. Based on this assumption, we can write the following joint Bayesian framework for the model assuming all units are measured, i.e., no censoring:

$$p(\mu_X) \times p(\sigma^2_X) \times p(\boldsymbol{\beta}) \times p(\sigma^2_{Y|X}) \times \prod_{i=1}^{N} N(X_i \,|\, \mu, \sigma^2_X) \times \prod_{i=1}^{N} N(Y_i \,|\, \beta_0 + \beta_1 X_i, \sigma^2_{Y|X}) \,, \tag{5}$$

where $p(\theta)$ denotes the prior of θ, μ_X is the mean of X, and σ^2_X is the variance of X.

In this model, we would like to account for censored measurements. Therefore, we extend the framework above to consider censored and observed sets of X and Y. Let $\text{LOD}_i(Y)$ and $\text{LOD}_i(X)$ be the LODs of Y_i and X_i, respectively, for measurement i. Let $O_Y = \{i : Y_i > LOD_i(Y)\}$ and $O_X = \{i : X_i > LOD_i(X)\}$ be the observed sets of observations of Y and X, respectively. Similarly, let $C_Y = \{i : Y_i \leq LOD_i(Y)\}$ and $C_X = \{i : X_i \leq LOD_i(X)\}$ be the censored measurements (measurements below the LOD). Then, we can write the following joint distribution

$$\begin{aligned} &p(\mu_X) \times p(\sigma^2_X) \times p(\boldsymbol{\beta}) \times p(\sigma^2_{Y|X}) \times \prod_{i \in O_X} N(X_i \,|\, \mu, \sigma^2_X) \times \prod_{i \in O_Y} N(Y_i \,|\, \beta_0 + \beta_1 X_i, \sigma^2_{Y|X}) \\ &\times \prod_{i \in C_X} \Phi\left(\frac{\text{LOD}_i(X) - \mu_X}{\sigma_X}\right) \times \prod_{i \in C_Y} \Phi\left(\frac{\text{LOD}_i(Y) - (\beta_0 + \beta_1 X_i)}{\sigma_{Y|X}}\right) \end{aligned} \tag{6}$$

where $\Phi(Z)$ is the cumulative distribution function (CDF) of Z.

For more rigorous details on this modeling strategy, see Groth et al. (2017) and Groth et al. (2018). These manuscripts also provide copies of the `R` code for developing each model using raw `R` code and also with the help of the JAGS language for Bayesian hierarchical modeling.

3.2 MCMC sampling

This model was implemented in R (R Development Core Team 2015) using `rjags` (JAGS: Just Another Gibbs Sampler; Plummer 2003). This program allows us to implement MCMC methods such as Gibbs sampling, Metropolis Hastings, etc. For further details on different sampling methods, we refer the reader to Gilks et al. (1996), Marin and Robert, (2007), Carlin and Louis, (2008), Gelman et al. (2013), and Brooks et al. (2011).

Specifically, this program uses a Gibbs sampler and full conditional distributions of the parameters to sample from the posterior distribution. For more information on Gibbs sampling, see Gelman et al. (2013). First, the model samples the censored measurements using the parameter estimates (for variances, means, and regression coefficients) found in the previous iteration. This creates a complete data set of measurements (no missing values) at the particular iteration. Next, the model estimates the values of the regression coefficients ($\boldsymbol{\beta}$), variances (σ^2_X and $\sigma^2_{Y|X}$), and means (μ_X) of this now complete data set (using full conditional distributions). We repeat this process thousands of times to obtain posterior samples of the regression coefficients, means, and variances (which are later reported as AMs, GMs, and GSDs). Due to the simultaneous estimation of the censored values and the parameter estimates, Bayesian methods allow us to account for the uncertainty in the censored values in the parameter estimates.

In practice we have seen that this model has immediate convergence under various prior settings. Convergence here was assessed with the help of Gelman Rubin Diagnostics, Monte Carlo Standard Errors, and trace plots.

3.3 Development of priors

Because many exposure groups ($\sim$45%) had a small number of measurements (i.e. 5-20) and relatively high censoring, we developed priors for the mean and variance parameters for each BTEX-H chemical. Because of the presence of our chemicals in the crude oil, we were able to form the priors using the relationship between the THC and the chemical of interest in broad overarching groups. This allowed us to leverage information from all jobs within each job group to obtain more precise parameter estimates of BTEX-H for the job group (reported as AM, GM, GSD, and 95th percentiles). The development of correlations for use as priors is further described in Groth et al. (2021 b).

At first glance, this may raise concerns, as we are technically specifying our prior distributions using data from the very jobs we expect to analyze – i.e. the common "using the data twice" critique of some empirical Bayesian methods. It should be noted, however, that (a) the contribution of any given job toward the specification of these priors was generally small; and (b) because these priors were based on large amounts of data, we used the approach of Quick et al (2017) to downweight the informativeness of our prior distributions to be equivalent to just five samples. For instance, it can be shown that the precision of a regression coefficient's estimate in a classical frequentist analysis increases linearly with the sample size. If we multiply this precision by the ratio of our desired prior sample size, n_0, and the original sample size, N – i.e., $\frac{n_0}{N}$ – we obtain a downweighted precision (i.e., inflated variance) analogous to that produced by a frequentist analysis of only n_0 samples. This downweighted precision can then be used to construct a normal prior distribution

for each regression coefficient with the same mean as the full sample posterior but with the desired informativeness. A similar approach can be used to construct a downweighted inverse gamma prior for all variance parameters. This approach ensured that the data remained the primary driver of the posterior distribution and were not overpowered by the prior. As a result, we believe our approach avoids the common pitfall of inflated precision associated with some empirical Bayesian methods.

Once the prior information had been incorporated, we ran the bivariate linear modeling framework for BTEX-H on each job group separately. We used 20,000 iterations (without burn-in) to obtain estimates of the GMs and GSDs. From these we calculated 20,000 AMs, and 95th percentiles of BTEX-H for each job group (Finney, 1941). The posterior median values and the 2.5, and 97.5th percentiles (quantiles forming a 95% credible interval) were selected to represent each statistic (AM, GM, etc). These estimates are reported in Ramachandran et al. (2021) and Huynh et al. (2020 a, b and 2021).

3.4 Prior Downweighting

To describe the downweighting approach we use (based on the prior sample size approach of Quick et al., 2017), consider the following example in which we develop estimates for xylene (a similar process would occur for all BTEX-H chemicals). In Groth et al. (2021 b) we developed a correlation/linear relationship between xylene and THC (both on the log scale) for an overarching group (A) which was developed based on particular exposure determinants (see Groth et al. (2021 b) for more information on the overarching relationships/correlations).

Specifically, for the linear relationship in the overarching group, we use the bivariate left-censored Bayesian model as described above in (6) which accounts for measurements below the LOD in both our response (in this case $Y = ln(xylene)$) and predictor ($X = ln(THC)$). Priors on the mean of X and the regression coefficients $\boldsymbol{\beta}$ were set as wide normal distributions (with mean 0, and variance 100,000), while the variance terms ($\sigma^2_{Y|X}$ and σ^2_X) were given inverse gamma priors (IG(shape=0.01, scale=0.01) using the notation of Gelman et al. 2013). From this model, we obtain posterior distributions of the regression parameters ($\boldsymbol{\beta}$ and $\sigma^2_{Y|X}$) which we use to develop the informative priors. We will call these parameters from this overarching relationship $\boldsymbol{\beta}_A$ and σ^2_A.

Quick et al. (2017) showed that it is sufficient to have summaries of this posterior distribution to construct the priors. Therefore, from each overarching relationship, or in this example the overarching group A, we ultimately use the 2.5, 97.5th, median, and mean estimates of the posterior distribution of each parameter.

In most Bayesian analyses, the priors are based on N observations (number of observations in the dataset). In those cases, we would be concerned that the priors would contain much more information about μ and σ^2 than the data contributes via the likelihood. This is where the motivation for *downweighting* originates — i.e., we *choose* a prior sample size, $n_0 \ll N$, and instead consider a prior that depends on n_0 rather than N.

In order to make the priors for the regression coefficients β_0 and β_1 depend on n_0 rather than N, we will scale the variance of the previous posterior distribution by a particular factor. Specifically, in this example, we obtain posterior estimates of β_0 and β_1 for overarching group A. Thus we can determine the mean ($\mu_{\beta_0}, \mu_{\beta_1}$) and variance ($\sigma^2_{\beta_0}, \sigma^2_{\beta_1}$) of the distributions of β_0 and β_1 that we previously found. Based on the

parameterization used in the bivariate framework, the priors on β_0 and β_1 will also be Normal with the same means identified of the posterior ($\mu_{\beta_0}, \mu_{\beta_1}$), but we will adjust the variance/inflate the variance to allow the distribution to be wider (and depend on n_0 sample size; for exact details of this adjustment see Quick et al. (2017). We perform a similar adjustment on the variance σ_A^2 to obtain a prior on the conditional variance of the natural log of xylene given the information in ln(THC). For the variances, we add some small adjustments are made at the end to restrict the variances in particular ranges (GSDs between 1.01 and 12).

We did this for all BTEX-H chemicals under each overarching relationship. This downweighting process provided the priors for β_0, β_1, and $\sigma^2_{Y|X}$ which were used in the final bivariate left-censored Bayesian model that was used for exposure estimation. We selected the closest overarching group for development of priors for each exposure group (unless sufficient sample sizes were not available).

In the final exposure estimation process, the data provided to the model consisted of only those observations in that exposure group. The model forms a relationship between THC and the BTEX-H chemical (both on the natural log scale) in order to better estimate exposure for the BTEX-H chemical. The priors serve to help inform this relationship, and thus the exposure estimates if there are few observed measurements (either because of low sample size, or because of censoring). However, if sufficient data is present, the data will continue to drive the inference. Outputs of the final exposure model are posterior distributions of the AM, GM, GSD, and 95th percentile of the BTEX-H chemical exposures for that exposure group.

3.5 Considering a multivariate modeling strategy

We considered using a multivariate modeling strategy, where we include THC and all of the BTEX-H chemicals in one model. This question is addressed in Groth et al. (2018), but we briefly address this issue here. For the bivariate modeling strategy to be appropriate we are expecting the crude oil mixture to be the primary source of exposure. When correlations are low between multiple chemicals, it has been suggested that multiple mixtures are the primary sources of exposure and a multivariate modeling strategy would be appropriate (Groth et al., 2018). We have, however, generally observed high correlations (greater than 0.5) between THC and the BTEX-H chemicals and among the BTEX-H chemicals themselves. Thus, if these chemicals were included in the same multivariate model, there would be multicollinearity (collinearity) which could lead to increased variances on the regression estimates. As we are using the regression equation to estimate the mean exposures, increased variances would be problematic because the variance is used to estimate the AM.

In the GuLF STUDY, we have seen that THC by far has the highest consistent correlations observed with the individual BTEX-H chemicals. Similarly, it also has the lowest censoring, meaning that it could provide inference (more information to guide estimation) above another chemical with greater censoring.

We chose not to obtain THC estimates from the bivariate modeling strategy since we would have had to select a chemical to "trust" as a covariate (or chose one of the bivariate models previously run). Since exposures to the BTEX-H chemicals varied greatly, we would have wanted to select different chemicals under different conditions in order to promote accurate estimation. This would lead to a high level of complexity when what we desire are estimates that can be compared with one another (i.e. using same methodology). Therefore, we used the Bayesian ANOVA model to estimate exposure to THC and THC to estimate exposures

for the BTEX-H chemicals.

3.6 Statistical Assumptions

The use of this model relies on several assumptions. First, we assume each chemical is log-normally distributed such that X and Y are normal after log transformation (for all BTEX-H and THC chemicals). This assumption is valid based on most of the environmental health literature. Second, this model assumes a linear relationship is present between X and Y and that relationship is approximately constant over the time period of interest. This further assumes that this linear relationship continues below the LOD. Third, this model assumes other assumptions of linear regression including independence of observations, normality of error terms, and equal variances. Finally, we modeled each exposure group independently/separately of other exposure groups.

4 VOC Derived THC and BTEX-H Measurements

In Groth et al. (2021 a) and Ramachandran (2021), we describe the process to impute THC measurements using a VOC direct reading area measurement database. These estimates were developed on vessels operating remotely operated vehicles and on marine vessels. While some THC estimates are provided in Ramachandran et al. (2021), these measurements were sometimes incorporated into exposure estimates for large overarching groups. For further information on the BTEX-H imputation from VOC derived THC samples, see Groth et al. (2021 b).

5 N-hexane Missing Data Estimation

Although n-hexane is a component of the crude oil, it was not analyzed in all of the air samples in which THC and BTEX were analyzed. Because the period with the missing hexane measurements (April through May) had many of the highest THC and BTEX exposures, n-hexane levels were imputed for the n-hexane unanalyzed samples in that period. We use the word "imputation/imputed" here to refer to values that were estimated or predicted using a full model based Bayesian inference (not multiple imputation). The imputation of measurements from n-hexane are further discussed in Groth et al. (2021 b).

6 JAGS model code

6.1 THC model

```
model {

for(i in 1:N){
    above.lod[i] ~ dinterval(x[i],llodVec[i])
    x[i]~dnorm(mu[i],tau)
    mu[i] <- alphabeta
 }

alphabeta~dunif(log(25),log(50000))
tau<-1/(sigma*sigma)
sigma~dunif(log(1.01),log(12))
 }
```

6.2 Bivariate left-censored Bayesian model

Note: inverse-gamma priors for the variances are currently given, but this can be adjusted to other priors.

```
model {

for (i in 1:N){
is.notcensoredx[i]~dinterval(X[i],cx[i])
X[i]~dnorm(mux, tausqx)
is.notcensoredy[i]~dinterval(Y[i],cy[i])
Y[i]~dnorm(mu[i],tausqy)
mu[i]<-beta[1]+beta[2]*X[i]   }

mux~ dnorm(0,0.00001)
for (j in 1:2) { beta[j]  ~ dnorm(0,0.00001) }
tausqy~dgamma(0.01,0.01)
tausqx~dgamma(0.01,0.01)
sigmayx<-1/tausqy
sigmax<-1/tausqx
cov<-sigmax*beta[2]
variancey<-sigmayx+cov*tausqx*cov
mu.seg<-beta[1]+beta[2]*mux
corr<-cov/(sqrt(variancey)*sqrt(sigmax))
Rsquared<-corr*corr
```

```
}
```

6.3 Final bivariate model incorporating priors:

This model is built off the model in 5.2 and accounts for downweighting of the priors. This code is available upon request.

Appendix C

Estimation of Exposures to Total Hydrocarbons, Benzene, Ethylbenzene, Toluene, Xylene and n-Hexane for Exposure Groups with Insufficient Measurements and Assignment of Confidence

When job-activity-task/location/time period determinants matched the same job-activity-task/location/time period of the measurement data and there were at least 5 measurements with <80% censoring, Bayesian methods were used to develop the exposure statistics. There remained, however, many other exposure groups (EGs, i.e. job-activity-task, location and time period combinations) that did not meet the measurement criteria. We applied other approaches for estimating exposure levels for EGs with greater censoring (Stenzel et al., 2021b), but always maintained the requirement of at least 5 measurements. Even after applying these approaches, however, there were a number of EGs that had insufficient measurements (n=0-4). We, therefore, developed a set of rules to assign the same type of statistics (arithmetic means, geometric means, geometric standard deviations, 95th percentiles and their 95% credible intervals) to these other EGs as we did for EGs that met the criteria.

First, we reviewed the jobs-activities-tasks performed and assigned the exposure statistics of a similar job. Generally, on the rigs, we found few jobs-activities-tasks sufficiently comparable to assign exposure statistics to a second job (one example was workers piloting remotely operated vehicles (ROVs) and ROV supervisors). Then, for jobs with insufficient measurements on the 4 rig vessels, we developed 2 broader rig groups. First, the *Discoverer Enterprise* (*Enterpris*e) worked in tandem with the *Helix Q4000 (Q4000)*. They were located within about 500 m from each other at the wellhead. Both were flaring collected oil and gas. The rigs were populated by similar jobs and were present about the same time (although after the oil released had been stopped, the *Q4000* left the area earlier than the *Enterprise*). We therefore considered these two vessels "sister" ships. Similarly, the *Development Driller II* and *Development Driller II* were both drilling a relief well to intersect the existing well casing (~18,000 ft, 5486 m) below the Gulf surface) above the oil and gas formation and plug it with cement and drilling mud to permanently seal and relieve pressure on the damaged well. The two vessels were located within 1 nmi (1852 m) of the wellhead for about the same length of time and had the same jobs performing similar functions. These two vessels were also considered "sister" ships. If a job on a rig with insufficient measurements was the same as a job on its sister ship that had sufficient measurements, the measurements from both jobs were combined into a single set of exposure statistics and assigned to the job with the insufficient measurements. The job with sufficient measurements, however, retained its original estimate. If the job had no measurements and the sister ship job had sufficient measurements, the sister ship's job's set of exposure statistics were assigned to the first job.

For rig jobs that did not meet either of these 2 conditions, we assigned the descriptive statistics of broader groups. Crew members such as roustabouts, roughnecks and tool pushers with insufficient data were assigned the exposure statistics of "Outside crew", comprising all measurements taken on all outside crew jobs. Generally, all crew members on a single vessel were workers of a single employer. If no job title was provided, but a description was provided that was similar to a crew job or if the employer was the crew employer, the job was also assigned "Outside crew". Burner fire control workers and other workers brought in especially for the disaster response were employed by multiple employers. Where measurements were insufficient, these jobs were assigned the exposure statistics of "Outside operations" comprising all measurements taken on all operations jobs. Administrative workers were assigned the

statistics associated with all “Inside/other” measurements. Finally, if the job and the employer were ambiguous, the study participant was assigned to “All groups”, which comprised all measurements on the rig.

Vessel type, e.g., “All ROVs”, was identified for participants who indicated an ROV activity in an open-ended question “Please describe any other tasks you did while on a boat, ship, or barge during the clean-up effort that I did not ask you about,” but did not identify a specific ROV vessel to the question asking for the name of the vessel. Similarly, for participants who indicated taking water or oil samples on an unspecified research vessel (RV) we had a determinant value of “All RVs”.

For workers with insufficient measurements performing activities or tasks on a vessel or on land, we first assigned the exposure statistics of a second activity, such as when we used the descriptive statistics calculated for “General Environment/land” for “Onsite driver”, based on the similarity of the environment. More often, however, we based the assignment on adjacent states or time periods. LA and MS were considered as “sister” states, as were AL and FL. TP1a and TP1b, the 2 periods when the oil was being released, were considered “sister” time periods, as were TP2 and TP3 (work continued on capping the wellhead and drilling the relief well) and TP 4, 5 and 6 (decrease and eventual shut down of most work, except for beach clean-up). Selection of activity, state or time period was based on the likelihood of the similar conditions and the availability of the data. Even fewer measurements resulted in assigning measurements across “All states” or “All time periods”.

Because of the lower confidence in the assignment of some of the exposure statistics, we identified a relative confidence level to each set of statistics. If all 3 determinants of the EG matched those of the measurement data, a confidence of 5 was assigned. If any 2 determinants matched those of the EG, a 4 was assigned. If only 1 determinant matched a determinant of the EG, a 3 was assigned, and if no determinant matched any of those of the EG, a 2 was assigned. The identification of the 3 types of determinants was identified in the JEM (e.g., the confidence for a roustabout on the Enterprise in TP1a might be 4 with the determinant information being “job/rig/TP1b”, meaning the job and rig matched those of the measurement, but the time period did not (i.e. TP1b measurements were used). These confidences were then assigned to the study participants.

References

Brooks S, Gelman A, Jones G, Meng X. (2011) 1st ed. Handbook of Markov Chain Monte Carlo. Boca Raton, FL: Chapman & Hall/CRC. ISBN 1420079417.

Carlin BP, Louis TA. (2008) 3rd ed. Bayesian Methods for Data Analysis. Boca Raton, FL: Chapman & Hall/CRC. ISBN 1584886978.

Finney DJ. (1941) On the distribution of a variate whose logarithm is normally distributed. *J R Stat Soc*; 7 (Suppl.): 155–61.

Ganser GH, Hewitt P. (2010) An accurate substitution method for analyzing censored data. *J of Occup Environ Hyg*; 7(4): 233-44.

Gelfand AE, Ghosh SK. (1998) Model choice: A minimum posterior predictive loss approach. *Biometrika*; 85(1): 1-11.

Gelman A, Carlin B, Stern HS, Dunson DB, Vehtari A, Rubin DB. (2013) 3rd ed. Bayesian Data Analysis. Boca Raton, FL:CRC Press. ISBN 1439840954.

Gilks WR, Richardson S, Spiegelhalter, D. (1996) Markov Chain Monte Carlo in Practice. Boca Raton, FL:Chapman & Hall/CRC. ISBN 0412055511.

Groth C, Banerjee S, Ramachandran G *et al* (2017) Bivariate left-censored Bayesian model for predicting exposure: Preliminary analysis of worker exposure during the Deepwater Horizon oil spill. *Ann Work Expo Health*; 61; 76-86.

Groth CP, Banerjee S, Ramachandran G *et al.* (2018) Multivariate left-censored Bayesian modeling for predicting exposure using multiple chemical predictors. *Environmetrics*; 29: 1-16 doi: 10.1002/env.2505

Groth CP, Banerjee S, Ramachandran G *et al*. (2021 a) Methods for the analysis of 26 million VOC area observations during the *Deepwater Horizon* oil spill response and clean-up. *Ann Work Expo Health*; 1-16. doi: 10.1093/annweh/wxab038

Groth CP, Huynh T, Banerjee S *et al*. (2021 b) Linear relationships between total hydrocarbons and benzene, ethylbenzene, toluene, xylene, and hexane during the *Deepwater Horizon* response and clean-up. *Ann Work Expo Health*; 1-18. doi: 10.1093/annweh/wxab064

Huynh T, Ramachandran G, Banerjee S, Monteiro J, Stenzel G, Sandler DP, et al. (2014) Comparison of methods for analyzing left-censored occupational exposure data. *Ann Occup Hyg*; (18): 1126-42

Huynh T, Quick H, Ramachandran G, Banerjee S, Stenzel G, Sandler DP, *et al* (2016) A comparison of the β substitution and Bayesian approach for handling left-censored data. *Ann Occup Hyg*; (60): 56-73.

Huynh TB, Groth CP, Ramachandran G *et al.* (2020 a). Estimates of occupational inhalation exposures to six oil-related compounds on the four rig vessels responding to the *Deepwater Horizon* oil spill and clean-up. *Ann Work Expo Health*; 1-22. doi: 10.1093/annweh/wxaa072

Huynh TB, Groth CP, Ramachandran G *et al.* (2020 b). Estimates of inhalation exposures to oil-related components on the supporting vessels during the *Deepwater Horizon* oil spill. *Ann Work Expo Health*; 1-13. doi: 10.1093/annweh/wxaa113

Huynh, TB, Groth CP, Ramachandran G *et al.* (2021). Estimates of inhalation exposures among land workers during the *Deepwater Horizon* oil spill clean-up operations. *Annals of Work Exposures and Health*; 1-16. doi: 10.1093/annweh/wxab028

Marin J, Robert C. (2007) 1st ed. Bayesian Core: A Practical Approach to Computational Bayesian Statistics. New York, NY: Springer. ISBN 0387389792.

Plummer M. (2003) JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling, Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003), March 20–22, Vienna, Austria. ISSN 1609-395X.

R Development Core Team. (2017) R: a language and environment for statistical computing. R Foundation for Statistical Computing: Vienna. http://www.R-project.org. Accessed 9 September 2021.

Ramachandran G, Groth CP, Banerjee S *et al.* (2021) Using real-time area VOC measurements to estimate total hydrocarbon exposures to workers involved in the *Deepwater Horizon* oil spill. *Ann Work Expo Health*; 1-16. doi: 10.1093/annweh/wxab066

Stenzel MR, Groth CP, Huynh TB *et al.* (2021 a) Exposure group development in support of the NIEHS GuLF STUDY. *Ann Work Expo Health*; 1-33. doi: 10.1093/annweh/wxab093

Stenzel MR, Groth CP, Banerjee S *et al.* (2021 b) Exposure assessment techniques applied to the highly censored *Deepwater Horizon* Gulf oil spill personal measurements. *Ann Work Expo Health*; 1-15. doi: 10.1093/annweh/wxab060

Quick H, Huynh T, Ramachandran G. (2017) A Method for Constructing Informative Priors for Bayesian Modeling of Occupational Hygiene Data. *Ann of Work Expo Health*; **61**(1):67-75.