12-16-2004

Estimating Percentile-Specific Causal Effects: A Case Study of Micronutrient Supplementation, Birth Weight, and Infant Mortality

Francesca Dominici
The Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics, fdominic@jhsph.edu

Scott L. Zeger
The Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics, szeger@jhsph.edu

Giovanni Parmigiani
The Sydney Kimmel Comprehensive Cancer Center, Johns Hopkins University & Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, gp@jimmy.harvard.edu

Joanne Katz
Johns Hopkins Bloomberg School of Public Health, Department of International Health, jkatz@jhsph.edu

Parul Christian
Johns Hopkins Bloomberg School of Public Health, Department of International Health, pchristi@jhsph.edu

ESTIMATING PERCENTILE-SPECIFIC CAUSAL EFFECTS: A CASE STUDY OF MICRONUTRIENT SUPPLEMENTATION, BIRTH WEIGHT, AND INFANT MORTALITY

December 16, 2004

Francesca Dominici, Scott L. Zeger, Giovanni Parmigiani, Joanne Katz, and Parul Christian

Abstract

In developing countries, higher infant mortality is partially caused by poor maternal and fetal nutrition. Clinical trials of micronutrient supplementation are aimed at reducing the risk of infant mortality by increasing birth weight. Because infant mortality is greatest among the low birth weight infants (LBW) ($\leq$ 2500 grams), an effective intervention may need to increase the birth weight among the smallest babies. Although it has been demonstrated that supplementation increases the birth weight in a trial conducted in Nepal, there is inconclusive evidence that the supplementation improves their survival. It has been hypothesized that a potential benefit of the treatment on survival among the LBW is partly compensated by a null or even harmful effects among the largest infants. Thus, two key scientific questions are whether the effect of the treatment on survival differs across the birth weight distribution (e.g. is largest among the LBW), and whether the effect of the treatment on survival is mediated wholly or in part by increases in birth weight.

Motivated by a community trial in Nepal, this paper defines population and causal parameters for estimating the treatment effects on birth weight and on survival as functions of the percentiles of the birth weight distribution. We develop a model with potential outcomes and implement principal stratification for estimating and comparing the causal effects of the treatment on mortality in sub-populations of babies defined by their birth weights. We use a Bayesian approach with data augmentation to approximate the posterior distributions of the parameters taking into account uncertainty associated with the imputation of the counterfactuals. This approach is particularly suitable for exploring the sensitivity of the results to modelling assumptions and other prior beliefs.

Our analysis shows that the average causal effect of the treatment on birth weight is equal to 68 grams (95% posterior regions 25 to 110) and that this causal effect is largest among the LBW. Posterior inferences about average causal effects of the treatment on birth weight are robust to modelling assumptions. However inferences about causal effects for babies at the tails of the birth weight distribution can be highly sensitive to the unverifiable assumption about the correlation between the observed and the counterfactuals birth weights. Among the LBW infants who have a large causal effect of the treatment on birth weight, we found that a baby receiving the treatment has 5% to 7% less chance of death if the same baby had received the control. Among the LBW, we found weak evidence supporting an additional beneficial effect of the treatment on mortality

independent of birth weight.

Key Words: Direct and Mediated Effects, Causal Inference, Post-treatment variables, Percentile-specific effects, Data Augmentation

Affiliations: Francesca Dominici, Associate Professor, Department of Biostatistics; Scott L. Zeger, Professor, Department of Biostatistics; Giovanni Parmigiani, Associate Professor, Department of Oncology Biostatistics; Joanne Katz, Professor, Department of International Health; Parul Christian, Associate Professor, Department of International Health, all at Johns Hopkins University, Baltimore MD 21205. Contact Information: Francesca Dominici, e-mail: fdominic@jhsph.edu, phone: 410- 6145107, fax: 410-9550958.

Acknowledgments: Funding for Francesca Dominici and Scott L Zeger was provided by a grant from the National Institute for Environmental Health Sciences (ES012054-01) and by the by NIEHS Center in Urban Environmental Health (P30 ES 03819). We would like to thank Drs Constantine Frangakis, Thomas Louis for comments.

1 Introduction

The reduction of infant mortality remains a major public health goal (Child Health Reserach Project, 1996), particularly in developing countries where current rates are an order of magnitude higher than in Europe, North America, and Japan. In developing countries, higher infant mortality is partially caused by poor maternal and fetal nutrition as reflected in the distribution of infant birth weights. One intervention trial have attempted to reduce infant mortality by improving maternal micronutrient sufficiency (Christian et al., 2003a). The idea is to improve maternal nutritional status thereby improving fetal growth and reducing the risk of infant mortality. Because infant mortality is greatest among low birth weight (LBW $\leq$ 2500 grams) and very low birth weight (VLBW $\leq$ 1500 grams) infants, it is assumed that an effective intervention must increase birth weight among the smallest babies, that is, in the left tail of the birth weight distribution. That maternal nutritional supplementation increases birth weight has been demonstrated in replicated randomized trials in several countries (Lechtig et al., 1975; Ceesay et al., 1997; Caulfield et al., 1999; Christian et al., 2003a). However, to date, there is limited direct evidence that maternal supplementation causes a reduction in the prevalence of babies born at the smallest weights and that this reduction improves their survival (Garner et al., 1992; McIntire et al., 2001; West et al., 1999; Katz et al., 2000a; Rasmussen, 2001; Christian et al., 2003b).

The methods in this paper are motivated by a double blind randomized community trial in rural Nepal (Christian et al., 2003a). The intervention program provided weekly iron, folic acid and vitamin A while the control was weekly vitamin A alone. The 1051 and 947 pregnant women that were assigned to the control and treatment delivered 866 and 766 live born infants, respectively. Details on the study designs including the rational for the selection and exclusion of the women in the study are detailed in Christian et al. (2003a). The team measured the birth weight within 72 hours of delivery and then followed the infants for one year to determine whether or not they survived. In the motivating study, treatments were randomized to 426 communities rather than to individual women. This can create some correlation among the birth weights and infant deaths within communities. It is a minor extension of the methods discussed in this paper to account for this clustering which turns out to be of negligible magnitude for the infant mortality outcome. To simplify the notation and exposition, we will not address clustering here.

The interesting aspect of this study is that the investigators anticipate that antenatal iron-folic acid supplementation may affect birth weight and ultimately survival differently among the smaller and larger babies. That is, they hypothesize that there could be an interaction between the treatment effect and the birth weight percentiles. Cox (1984) referred to this situation as the most basic form of interaction. Doksum and Sievers (1976) defines a similar form of interaction by allowing the treatment effect to vary as a function of the health response. Koenker and Bassett (1978) introduced quantile regression methods which model the quantile function of an outcome

variable as a function of covariates, and applied this approach to survival times where the regression parameters are allowed to depend on the quantile of interest (Koenker and Geling, 2001). Dominici et al. (2003) recently introduced Smooth Quantile Ratio Estimation (SQUARE), a novel method for estimating the difference in medical expenditures between persons with and without a disease as a function of the medical expenditures percentiles.

The second interesting question from this study is whether the antenatal iron-folic acid supplementation improves survival largely through its positive effect on birth weight. The hypothesis is that supplementation will improve intra-uterine growth, lowering the risk of LBW and thus increasing the chance of survival during the first year after a live birth. Therefore we are interested in investigating the relative importance of different pathways for the antenatal iron-folic acid supplementation on survival. By one pathway, the intervention affects survival only throughout a change in birth weight (the so called "mediated effect"). A second possible pathway is that intervention affects survival over and above its effect on intra-uterine growth, that is through other mechanisms that do not involve birth weight. We refer to this pathway as a "direct effect".

To explore the association between birth weight and mortality, we fit a logistic regression model expressing the log odds of infant death as a separate smooth function of the birth weight for the control and intervention groups. The top panel of Figure 1 shows the smoothed histograms of the birth weights. The bottom panel shows the estimated smooth curves with 95% confidence bands plotted in correspondence to the ranges of the measured birth weights in the two groups. These exploratory plots suggest that: 1) the probability of death decreases as the birth weight increases and tends to rise again for the heaviest babies in the control group; 2) approximately 43% and 34% of the babies in the control and in the intervention groups are LBW, respectively, suggesting that the treatment may reduce the percentage of LBW; and 3) the visual inspection of the two smoothed histograms suggests that the treatment increases birth weight for the smaller babies only, thus indicating that the treatment effect on birth weight might vary with respect to the percentiles of the birth weight distribution.

The statistical literature on surrogate endpoints and causal inference extensively discusses post-treatment variables in clinical trials and observational studies. Prentice (1986) first proposed criteria for a perfect surrogate (e.g. the birth weight), the most important being that the final response is conditionally independent of treatment given the surrogate. When the assumption of conditional independence is violated, related approaches have been proposed that compare results of the regression of the health response on the treatment with and without the adjustment for the intermediate variable (Freedman et al., 1992; Daniels and Hughes, 1997; Buyse and Molenberghs, 1998; Begg and Leung, 2000; Leung, 2001; Molenberghs et al., 2001; Xu and Zeger, 2001; Cowles, 2002). Robins (1989), Robins and Greenland (1992), and Pearl (2000) have developed identifiability results for direct and indirect causal effects under the framework of potential outcomes and they define an "individual direct effect" as the counterfactual effect of a treatment on an outcome

when the intermediate variable is set at the value it would have had if the individual had not been treated (see also Cole and Hernan (2002)). These identifiability results have been recently generalized by van der Laan and Petersen (2004). Frangakis and Rubin (2002) proposed a novel approach for defining causal effects adjusted for post-treatment variables. This approach, known as "principal stratification", is based upon a comparison of treatment effects on the outcome among sub-populations for whom a causal effect of treatment on the post-treatment variable did and did not occur.

The broad objectives of this paper are to develop and apply a statistical model with counterfactual variables for this birth weight-mortality study. The contributions of this paper are to: 1) define and compare population and causal parameters (Holland, 1986) that measure the effects of an intervention on a clinical outcome (infant mortality) that are allowed to vary with the percentiles of the post-treatment variable (birth weight); 2) extend and apply a causal statistical framework to compare the causal "direct" effect of the treatment on mortality, from the causal effect of the treatment on mortality that is "mediated" by post-treatment changes in birth weight; 3) develop a Bayesian approach with data augmentation (Tanner and Wong, 1987; Tanner, 1991; Albert and Chib, 1993; Chib and Greenberg, 1998) for approximating the marginal posterior distributions of all parameters of interest accounting for the uncertainty about the missing counterfactuals; 4) quantify the sensitivity of causal inferences to key assumptions for which there are not direct observations in the data set.

In Section 2, we introduce notation, specify our model, and define the population and causal parameters. In section 3, we define the complete likelihood function for the observed data and the missing counterfactual data. In this section, we also describe our Monte Carlo Markov Chain with data-augmentation algorithm (Tanner and Wong, 1987; Tanner, 1991) for approximating the posterior distributions of all the unknown parameters and the unobservable variables. In Section 4, we summarize the results by comparing causal and population parameter estimates. We explore sensitivity of the causal parameter estimates to the unverifiable assumptions about counterfactuals, to model specification, and to distributional assumptions. In Section 5, we discuss future research opportunities.

2 Definition of Population and Causal Parameters

In this section, we define population and causal parameters of scientific interest in terms of counterfactual variables. To establish notation, let Z_i be the treatment indicator for live birth i that takes values 0 or 1 to indicate the control and the treatment groups, respectively. Let W_i^{obs} be the observed birth weight measurement within the 72 hours of the delivery, and let Y_i^{obs} be the observed mortality indicator within one year. Let $n_0 = 866$ and $n_1 = 766$ be the number of live

births for the control and the treatment groups respectively and let $N = n_0 + n_1 = 1632$ be the total number of live births.

Adopting a causal model with potential outcomes (Rubin, 1978; Holland, 1986), let $\boldsymbol{Z}$ be the N-dimensional vector of treatment assignments with ith element Z_i, and $W_i(\boldsymbol{Z})$ be the birth weight of baby i given the randomly allocated vector $\boldsymbol{Z}$. We define $Y_i(\boldsymbol{Z}, \boldsymbol{W})$ to be the binary random variable for the mortality indicator for baby i corresponding to the vector of birth weights $\boldsymbol{W}$ and the vector of treatment assignments $\boldsymbol{Z}$. We refer to $Y_i(\boldsymbol{W}, \boldsymbol{Z})$ and $W_i(\boldsymbol{Z})$ as potential outcomes. To assure a valid causal interpretation of the causal estimands defined below, we make the following usual assumptions:

1. *Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1978):* the potential outcomes of each baby i are unrelated to the treatment status of other babies. That is, the birth weight and the mortality potential outcomes of each baby are not affected by the treatment assignment of others. Therefore we can write $Y_i(\boldsymbol{Z}, \boldsymbol{W})$ and $W_i(\boldsymbol{Z})$ as $Y_i(Z_i, W_i)$ and $W_i(Z_i)$, respectively;

2. *Ignorable Assignment:* assignment to the supplementation is at random;

3. *Perfect compliance with the treatment:* all mothers take the assigned dose in both treatment groups.

Note that $Y_i(0)$ and $W_i(0)$ are defined for all N babies, but only observed for the n_0 babies in the control group of the study. Similarly, $Y_i(1)$ and $W_i(1)$ are defined for all N babies, but only observed for the n_1 babies in the intervention group. Thus we denote the observed and the missing data as $Y_i^{obs} = \{Y_i(z),\ \text{if } z = Z_i\}$ and $Y_i^{mis} = \{Y_i(z),\ \text{if } z \neq Z_i\}$, respectively. Similar definitions apply for W_i^{obs} and W_i^{mis}.

The parameters of interest are defined in Tables 1 and 2 for birth weight and mortality respectively. The first two rows of Table 1 indicate population parameters measuring difference between the means (Δ^W) and the percentiles (Δ_p^W) of the population of birth weights for the two treatments. Note that the parameter Δ_p^W is defined as $Q_1(p) - Q_0(p)$ where $Q_1(p)$ and $Q_0(p)$ are the quantile functions of the marginal distributions of $W_i(1)$ and $W_{i'}(0)$ respectively.

The last two rows of Table 1 define the causal parameters measuring the effects of the treatment on birth weight, on average (τ^W), and specific to the percentiles of the birth weight distribution (τ_p^W). Note that Δ_p^W is a population parameter, whereas τ_p^W is a causal parameter: in the definition of Δ_p^W, we consider the difference in percentiles of two different distributions of birth weights. In the definition of τ_p^W, we consider the expected difference in birth weights $W_i(1) - W_i(0)$ for the same baby (Holland, 1986) whose control value $W_i(0)$ is at the p-percentile of the control distribution.

Table 2 summarizes the population and causal parameters for the treatment effect on infant mortality. Prior to defining these parameters, we need to specify a functional relationship between death and birth weight. Substantive knowledge and our exploratory analyses indicates that the following logistic regression model is a reasonable approximation to the actual mortality process:

$$\text{logit}Pr\{Y_i(Z_i) = 1 \mid Z_i, W_i(Z_i)\} \quad = \quad \beta_0 + \beta_1 Z_i + s(W_i(Z_i), 3),\ Z_i = 0, 1. \tag{1}$$

where $s()$ denotes a natural cubic splines with 3 knots.

By specifying this parametric model we make two key assumptions.

4. *Conditional independence of survival from the counterfactual birth weight given the treatment assignment and the observed birth weight:* For each baby, we assume that the probability of death under the treatment depends only on the birth weight under that treatment, and it does not depend on what birth weight would have been had the same baby been randomized to the other group. That is we assume, $Pr\{Y_i(Z_i) = 1 \mid Z_i, W_i(Z_i), W_i(1 - Z_i)\} = Pr\{Y_i(Z_i) = 1 \mid Z_i, W_i(Z_i)\}$;

5. *No interaction between the direct treatment effect on survival and the birth weight:* We assume that the direct effect of the treatment on mortality is the same for all babies and does not vary with respect to the birth weight distribution. That is we can write:

   $$\text{logit}Pr\{Y_i(1) = 1 \mid Z_i = 1, W_i(1) = w\} - \text{logit}Pr\{Y_i(0) = 1 \mid Z_i = 0, W_i(0) = w\} = \beta_1.$$

   This assumption can be relaxed by assuming a linear or non-linear interaction between the treatment and the birth weight, for example by replacing $\beta_1 Z_i$ with $\beta_1(Z_i \times W_i(Z_i))$ – or more generally with $Z_i \times s_2(W_i(Z_i), 3)$ – in model (1).

The first two rows of Table 2 indicate population parameters measuring treatment effect on mortality, on average (Δ^Y), and conditional on a specific percentile of the birth weight distribution (Δ_p^Y). Note that Δ_p^Y defines the difference in the probability of death between treated and non treated infants who are at the same percentiles of their respective birth weight distribution. Thus Δ_p^Y is not a causal parameter, because these differences correspond to two different sub-populations of babies.

The last two rows of Table 2 indicate the causal parameters measuring the effects of treatment on infant mortality, on average (τ^Y), and specific to the percentiles of the birth weight distribution (τ_p^Y). Thus, for a specific p, τ_p^Y can be interpreted as a causal effect which compares the probability of death for the same baby i given that the assumption that his/her birth weight under the control ($W_i(0)$) is at the p-th percentile.

In the last row of Table 2, we use the idea of principal stratification by Frangakis and Rubin (2002) for defining causal parameters of the effects of treatment on infant mortality that are

"adjusted" and "mediated" by post-treatment changes in birth weight. More specifically, τ_1^Y and τ_2^Y are the effects of treatment on mortality in the two sub-populations of LBW babies for whom the treatment effect on birth weight was smaller and larger than 50 grams, respectively. Thus a comparison between τ_1^Y and τ_2^Y measures the degree to which a causal effect of treatment on mortality occurs together with a causal effect of treatment on the birth weight among the LBW. The parameters τ_3^Y and τ_4^Y are the analogues of τ_1^Y and τ_2^Y for the not-LBW infants, that is for the infants with birth weight larger than 2500 grams.

All causal parameters (τ) depend upon unverifiable assumptions about the joint distribution of the counterfactual pairs of variables $\{W_i(0)$ and $W_i(1)\}$, and $\{Y_i(0)$ and $Y_i(1)\}$. In order to estimate the average causal effects (τs), we make the following key but unverifiable assumptions about the correlation between the observed outcomes and their counterfactuals:

6. *Correlation between the observed and the counterfactual birth weight:* we assume that the correlation between $W_i(Z_i)$ and $W_i(1-Z_i)$, denoted by ρ, varies between 0.5 and 0.9.

7. *Odds ratio between the observed and counterfactual mortality given birth weight:* Let $\mu_i(11)$ be the joint probability that the same baby i would die in both groups defined as $P(Y_i(Z_i) = 1, Y_i(1-Z_i) = 1 \mid W_i(Z_i), W_i(1-Z_i))$. We assume that the odds ratio $\psi = (\mu(11) \times \mu(00))/(\mu(10) \times \mu(01))$ varies between 3 and 20.

These choices are arbitrary but based on prior knowledge. As a guide for reasonable choices of ρ, we have used data from this randomized trial and from other data sources (Rahmathullah et al., 2003; Katz et al., 2000b, 2001) to estimate the correlations of birth weights for two successive children born to the same mother and birth weights for twins. We found that these correlations range from 0.45 to 0.7. The analogous odds ratios for mortality were estimated to be 1.8 and 52 respectively. As detailed in the next section, we will study the dependence of our causal inferences to the prior choices on the correlation coefficient ρ and on the odds-ratio ψ.

3 A Bayesian Implementation of Causal Inference

In this section, we define a Bayesian approach for approximating the marginal posterior distributions of the population and the causal parameters defined in Section 2. We start by defining the likelihood function for the complete data as:

$$L(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2) = \prod_{i=1}^{N} Pr(Y_i(1), Y_i(0) \mid W_i(1), W_i(0), \boldsymbol{\eta}_1) \times f(W_i(1), W_i(0) \mid \boldsymbol{\eta}_2). \quad (2)$$

In section 3.1, we specify $f(W_i(1), W_i(0) \mid \boldsymbol{\eta}_2)$ as a mixture of normal distributions. In Section 3.2, we specify an odds-ratio association model for bivariate binary variables $P(Y_i(1), Y_i(0) \mid W_i(1), W_i(0), \boldsymbol{\eta}_1)$ (Liang et al., 1992). This model will be consistent with equation (1). In section 3.3, we then detail the elicitation of the prior distributions and the implementation of the MCMC methods with data augmentation to obtain posterior samples of all the unknown parameters and the missing counterfactuals variables.

3.1 Statistical model for birth weight

We begin our specification of the joint distribution in Equation (2), by assuming that the marginal distributions of the random variables $W_i(z),\ z = 0, 1,\ i = 1, \ldots, N$ are a mixture of $J(= 3)$ normal distributions:

$$f_z(W_i(z) \mid \boldsymbol{\mu}_z, \boldsymbol{\sigma}_z^2, \boldsymbol{\gamma}_z) \quad = \quad \prod_{j=1}^{J} \gamma_{zj} \phi(W_i(z); \mu_{zj}, \sigma_{zj}^2),\ z = 0, 1 \tag{3}$$

where $\phi(x; \mu, \sigma^2)$ is the density of a normal distribution with mean μ and variance σ^2, $\boldsymbol{\mu}_z = (\mu_{1z}, \mu_{2z}, \mu_{3z})$, $\boldsymbol{\sigma}_z = (\sigma_{1z}, \sigma_{2z}, \sigma_{3z})$, and $\boldsymbol{\gamma}_z = (\gamma_{1z}, \gamma_{2z}, \gamma_{3z})$, where γ_{jz} are the mixing probabilities with $\sum_{j=1}^{J} \gamma_{jz} = 1$. To identify the mixture we set the constraint $\mu_{1z} < \mu_{2z} < \mu_{3z}$ (Kadane, 1974). We further assume that $\sigma_{1z}^2 = \sigma_{3z}^2 = 2 \times \sigma_{2z}^2$: assigning a larger variance to the outside components of the mixture is designed to flexibly capture heavy-tailed distributions. For ease of notation, we will set $\sigma_z^2 = \sigma_{2z}^2$.

This distributional assumption allows the parameters Δ_p^W and τ_p^W to vary flexibly as functions of the percentiles (p) of the birth weight distribution. If instead of the mixture model (3), we assumed that $W_i(z) \sim N(\mu_z, \sigma_z)$, then $\Delta_p^W = (\mu_1 - \mu_0) + (\sigma_1 - \sigma_0)\Phi^{-1}(p)$. Therefore, the simpler assumption of normality for the marginal distributions of $W_i(0)$ and $W_i(1)$ imposes a specific parametric form for Δ_p^W which does not depend on p for $\sigma_1 = \sigma_0$. In the results section, we will calculate the posterior probability of $\sigma_0^2 \neq \sigma_1^2$ to provide evidence in favor of the assumption that Δ_p^W depends on p, and we will explore the sensitivity of the posterior distribution of Δ_p^W as a function of p, under the mixture model and under the simpler assumption of normality with $\sigma_0 = \sigma_1$.

To allow for a correlation between $W_i(0)$ and $W_i(1)$, we assume that the standardized variables $\Phi^{-1}[F_z(W_i(z))],\ z = 0, 1$ have a bivariate normal distribution with mean zero, variance 1 and correlation ρ, where Φ is the cdf of a standard normal distribution and F_z is the cdf of $W_i(z)$.

In this formulation for the joint distribution of $(W_i(0), W_i(1))$, letting $\rho = 1$ corresponds to the rank preservation assumption used by Efron and Feldam (1991). Our specification allows for a single interpretable parameter capturing the correlation between $W_i(0)$ and $W_i(1)$, while allowing for a flexible representation of the two marginal distributions. An alternate stochastic generalization of the rank preservation assumption, obtained by specifying a probabilistic distribution on the ranks, has also been developed by Dobbin and Louis (2003).

3.2 Statistical model for infant mortality given birth weight

We specify the causal model for the joint distribution of the two outcome indicators conditional on the birth weights. Following Liang et al. (1992), we parametrize the 2×2 joint distribution $[Y_i(0), Y_i(1) \mid W_i(0), W_i(1)]$ in terms of the margins and the odds ratio. Specifically, we assume that:

$$\begin{aligned} P(Y_i(0) = y_i(0), Y_i(1) = y_i(1) \mid W_i(0), W_i(1), \boldsymbol{\eta}_1) &= \mu_i(0)^{y_i(0)}(1-\mu_i(0))^{1-y_i(0)} \times \\ &\times \ \mu_i(1)^{y_i(1)}(1-\mu_i(1))^{1-y_i(1)} + \\ &\quad (-1)^{y_i(0)-y_i(1)}(\mu_i(11) - \mu_i(0)\mu_i(1)) \end{aligned} \tag{4}$$

where $\mu_i(1) = Pr(Y_i(Z_i) = 1 \mid Z_i, W_i(Z_i))$ is defined in Equation (1). The parameter $\mu_i(11) = Pr(Y_i(0) = Y_i(1) = 1 \mid W_i(0), W_i(1))$ is a known function of the marginal probabilities $\mu_i(1), \mu_i(0)$ and of the odds ratio ψ.

3.3 Prior Distributions and Computation

Distributional assumptions in Sections 3.1 and 3.2 involve the following vectors of unknown parameters: 1) $\boldsymbol{\eta}_1 = (\boldsymbol{\beta}, \psi)$ where $\boldsymbol{\beta}$ includes β_0, β_1 and the spline coefficients defined in the regression model (1); and 2) $\boldsymbol{\eta}_2 = (\boldsymbol{\mu}_0, \boldsymbol{\mu}_1, \sigma_0, \sigma_1, \boldsymbol{\gamma}_0, \boldsymbol{\gamma}_1, \rho)$ denoting all the unknown parameters of the mixture (3). As stated in assumptions 6 and 7, the parameters ρ and ψ measure the association between the observed outcomes and their counterfactuals and they cannot be identified from the observed data. We specify prior distributions on the parameter of the mixture that are proper but vague enough to achieve goodness of fit to the observed birth weights. These choices are summarized in Table 3. In the results section, we explore the sensitivity of our results with respect to different values of ρ and ψ and we evaluate the goodness of fit of the empirical distributions of the observed birth weights.

To investigate the posterior distributions of all parameter of interest we implement Monte Carlo Markov Chain methods with data augmentation for imputing the missing data (Tanner, 1991; Gelman et al., 1995). Bayesian sampling of parameters of normal mixture distributions is typically handled by introducing auxiliary variables representing mixture component indicators, which results in closed form full conditionals (Diebolt and Robert, 1994). In our case, this option was not practical because of the special correlation structure we used, and because the unobserved birth weight variables enter the logistic component of the likelihood as well. We thus implemented a Metropolis-within-Gibbs (Tierney, 1994) approach, in which both the parameters and the counterfactual variables are sampled using a random walk proposal, truncated to the region defined by the constraints wherever applicable.

For each posterior sample of the unknown parameters and counterfactuals, we obtain a posterior sample of the p-specific parameters as follows. To obtain a posterior sample of Δ_p^W, we sort $W_i(0)$

and $W_{i'}(1)$ within the two groups of treated and untreated babies separately and then we take their difference. To calculate a posterior sample of τ_p^W, we sort by $W_i(0)$ and then we take the difference between the sorted $W_i(0)$ and its matched $W_i(1)$ for the same infant i. To calculate a posterior sample of Δ_p^Y, we first sort $Y_i(0)$ with respect to $W_i(0)$ and $Y_{i'}(1)$ with respect to $W_{i'}(1)$ within each of the two groups separately, and then we take the difference. Finally to calculate a posterior sample of τ_p^Y, we sort $Y_i(0)$ with respect to $W_i(0)$, and then we take the difference between the sorted $Y_i(0)$ and its matched $Y_i(1)$ for the same baby i. We smoothed the posterior samples of these percentile-specific parameters to reduce monte carlo variability in the posterior probability bounds.

4 Results

Figure 2 shows the posterior means and 95% posterior regions of the p-specific treatment differences in birth weight (Δ_p^W) comparing the treatment and control populations, plotted with respect to p under the two modelling assumptions for $(W_i(0), W_i(1))$. In Panel (a) $W_i(0), W_i(1)$ have a bivariate normal distribution with equal variances. In Panel (b) $W_i(0), W_i(1)$ have a mixture of normal distributions with correlation ρ as defined in Section 2. The triangles denote the difference between the empirical quantile functions for the observed data. The black dots denote the posterior means of Δ_p^W as a function of p.

Under the two modelling assumptions for the birth weights, the posterior means of Δ_p^W are generally consistent with the observed differences. However these two sets of estimates are very different at the smallest and at the largest percentiles. In fact in Panel (a), Δ_p^W is a constant function of p as is reflected in the flat line relationship in the lower left panel. If we fit a bivariate normal distribution without the constraint of equal variances, the posterior probability that $\log \sigma_1^2 - \log \sigma_0^2$ is less than zero is 97%, thus providing strong evidence that Δ_p^W varies with respect to p. Panel (b) shows that, when a more flexible mixture model is used, the effect of the intervention on the birth weight appears to vary by percentiles of the birth weight distribution. Therefore, estimating the posterior means of Δ_p^W by use of summaries of the posterior samples of Ws without imposing the normality assumption, provides a useful diagnostic tools for the performance of the algorithm and indicates that our mixture model with unequal variance in preferred and it will be used to report the results described below. Under the mixture model, we estimated a difference in birth weights quantiles between groups equal to 100 grams (95% posterior interval: 30 to 190) for the smallest babies ($p \simeq 0.05$) and that the treatment difference was close to zero for the largest babies ($p \simeq .95$). This is an ideal improvement as it has its greatest effect where the need is greatest.

Figure 3, Panel (a), shows the posterior means and 95% posterior regions of the p-specific causal effects of treatment on birth weight (τ_p^W) under the mixture model, $\rho = 0.9$ and $\rho = 0.5$ (darker

line). The vertical line is placed at the 0.42 percentile corresponding to 2500 grams in the control sample. Note that under the hypothesis of rank preservation ($\rho = 1$), then $\tau_p^W = \Delta_p^W$. For ρ different than one, population and causal parameter inferences differ by an amount that increases towards the tails of the birth weight distribution. Among LBW infants, we found that the average causal effects of the intervention on the birth weight are equal to 150 grams (95% posterior regions 100 to 300 grams) and to 410 grams (95% posterior regions 230 to 750 grams) for $\rho = 0.9$ and 0.5, respectively.

Figure 3, Panel (b), shows the sensitivity of the posterior distributions of the causal effect of treatment on birth weight (τ_p^W) separately for three sub-populations of babies ($W_i(0) \leq 1500, 1500 < W_i(0) \leq 2500, W_i(0) < 2500$), and overall for all babies, with respect to (ρ, ψ). The horizontal dotted line is placed at the sample mean difference (Δ^W). Within each sub-population, these causal effects are very sensitive to ρ but not to ψ. However the average causal effect of supplementation on birth weight (τ^W) – estimated to be 68 grams (95% posterior regions 25 to 110) – is robust to modelling assumptions about both ρ and ψ.

Figure 4, Panel (a), shows the posterior means and 95% posterior regions of the p-specific difference in infant mortality rates between the treatment and control populations (Δ_p^Y) plotted with respect to the percentiles of the birth weight distributions. For a specific p, Δ_p^Y is the difference in the probability of death between the babies with birth weights $W_i(1), W_{i'}(0)$, each at the p-percentile of their respective birth weight distributions. The vertical dotted line is placed at the 0.42 percentiles corresponding to 2500 grams in the control sample. There is no convincing evidence of a difference in the probabilities of death across the entire birth weight distribution.

Figure 4, Panel (b), shows the posterior means and 95% posterior regions of the p-specific causal effect of the treatment on infant mortality (τ_p^Y) plotted with respect to the percentiles of $W_i(0)$ for $\rho = 0.9$ and $\rho = 0.5$. For a specific p, τ_p^Y is defined as the difference in the probability of death for the same baby i whose control birth weight $\{W_i(0)\}$ is at the p-th percentile. For $\rho = 0.9$, we found that the causal effect of supplementation on mortality adjusted by birth weight is negative (intervention better) for the smaller babies and that this effect diminished for the larger babies, although with wide posterior regions. Posterior inferences on τ_p^Y are sensitive to ρ at the smallest percentiles. For $\rho = 0.5$ (darker line), there is stronger support for a beneficial "direct" effect of the supplementation on mortality among the very small babies only.

Figure 4, Panel (c), shows the posterior distributions of the causal effects of treatment on mortality for different values of (ρ, ψ) among different sub-population of babies. More specifically the posterior distributions are shown separately for four sub-populations of babies: 1) babies with a birth weight smaller than 2500 grams (LBW infants) for whom there is a causal effect of treatment on birth weight smaller than 50 grams (τ_1^Y); 2) LBW infants for whom there is a causal effect of treatment on birth weight larger than 50 grams (τ_2^Y); 3) babies with a birth weight larger than 2500

grams (not-LBW infants) for whom there is a causal effect of treatment on birth weight smaller than 50 grams (τ_3^Y); and 4) not-LBW infants for whom there is a causal effect of treatment on birth weight larger than 50 grams (τ_4^Y). The four boxplots at the far right show the posterior distributions of the total (direct plus mediated) causal effect of supplementation on mortality on average for all babies (τ^Y).

The four boxplots on the left (posterior distributions of τ_1^Y) indicate that, among the LBW babies with little change in birth weight after the supplementation, there is only weak evidence that antenatal iron-folic acid supplementation affects survival. The second set of four boxplots (posterior distributions of τ_2^Y) suggest that, among the LBW babies with absolute changes in birth weight after the supplementation larger than 50 grams, there is much stronger evidence that the antenatal iron-folic acid supplementation affects survival. The posterior means of these "mediated" causal effects for $\rho = 0.9$ and $\rho = 0.5$ are equal to -0.046 and -0.071 (95% posterior regions -0.11 to 0.02 and -0.13 to -0.02), respectively. These results indicate that a LBW infant receiving the intervention has 5% to 7% smaller chance of death than if the same baby had received the control intervention. This higher chance of death is due to changes in birth weight from the control to the treatment larger than 50 grams. The posterior distributions of the parameters τ_3^Y and τ_4^Y indicate that there is little evidence of a beneficial effect of supplementation on infant mortality for the not-LBW babies. The average causal effect of supplementation on mortality is robust to modelling assumptions and to (ρ, ψ).

Finally, we evaluate the consistency of the model assumptions and prior distributions with the patterns in the observed data. Figure 5 (top) shows 95% posterior regions of $F_z(W_i(z), \boldsymbol{\theta}_z^{(j)}),\ z = 0, 1$ where F_z are the cumulative distribution functions (cdfs) from the mixture model defined in Equation (3) and $\boldsymbol{\theta}_z^{(j)}$ are the j^{th} posterior samples of the parameters of the mixture. The black lines are the corresponding empirical cdfs, estimated directly from the observed birth weights. We see that the assumed model is reasonably consistent with the data.

5 Discussion

A micronutrient supplementation trial is considered effective if the treatment reduces the risk of infant mortality either directly or through increases in birth weight. Because infant mortality is greatest among low birth weight infants (LBW), an effective intervention must increase birth weight mainly among the smallest babies. In addition, it has been hypothesized that the supplementation could be harmful if it increases birth weight among the largest babies. A community-based trial in Nepal has shown that a multiple micronutrient supplementation increases birth weight but the limitation in the study size have to date prevented us from establishing that this translates into a mortality benefit (Christian et al., 2003b).

In this paper we develop a causal model to evaluate the efficacy of micronutrient supplementation trials in developing countries. We focus on whether the supplementation increases birth weight and ultimately survival differently among the smaller and the larger babies, and whether the supplementation improves survival largely through its positive effect on birth weight (mediated effect) or it improves survival even without affecting the birth weight (direct effect). Addressing these scientific questions is challenging because birth weight is a post-treatment variable (i.e. intermediate variable) that is in the causal pathway between nutritional supplementation and infant mortality.

Although average causal effects are robust to unverifiable assumptions about counterfactuals, posterior inferences on causal effects toward the tails of the birth weight distribution (for example among LBW infants) can be highly sensitive to ρ. More specifically we found that: among LBW infants, the effect of micronutrient supplementation on birth weight is greatest and its estimates size is highly sensitive of ρ: lower values of ρ correspond to a larger causal increase in birth weight.

The posterior distributions of the population and causal parameters are evaluated by using Bayesian inferences with data-augmentation methods (Tanner and Wong, 1987; Tanner, 1991; Albert and Chib, 1993; Chib and Greenberg, 1998). A nice feature of this inferential approach is that we can evaluate the posterior distributions of the quantities of interest taking into account uncertainty in the imputation of the the missing counterfactuals. In addition, we can easily explore the sensitivity of the posterior inferences to unverifiable assumptions about the correlation between the observed and the counterfactual variables.

To implement our approach we make several important assumptions. The first two (SUTVA, random assignment), are justified by the randomization of the treatment assignment and the independence of the sampling units. Third we assume perfect compliance. The compliance for this trial was very good and did not depend on the treatment (Christian et al., 2003a). The fourth and the fifth assumptions are in the logistic regression model for the probability of infant mortality as a function of the treatment indicator and the birth weight for the treatment received. Under the fourth assumption, we expect that that the risk of mortality under the treatment would depend only on the actual birth weight and not on the birth weight for the intervention not received. The fifth assumption, that the direct effect of the intervention on mortality is common to babies of all sizes, is consistent with the patterns in Figure 1 but there is little statistical power to show otherwise. Finally the sixth and the seventh assumptions are about the associations between the observed and the missing counterfactuals and these associations cannot be estimated from the data. To deal with this unidentified problem we: a) use data on siblings to estimate lower bounds for ρ and ψ and use those as a guide for our prior choices and sensitivity analyses; b) explore the sensitivity of estimated causal parameters with respect to choices for ρ and ψ; and c) compare inferences on causal parameters versus inferences on population parameters which are not affected by ρ and ψ.

The methodological development of this paper cuts across several contributions in quantile regression and causal inference literature. For example, we could have estimated the p-specific parameter Δ_p^W by use of a quantile regression model of the form $Q(p) = \alpha_p + \Delta_p^W Z_i$ (Koenker and Bassett, 1978) where $Q(p)$ is the quantile function of W_i^{obs}, and Z_i represents the treatment assignment. However in this paper we extend the traditional definition of p-specific regression coefficients in two ways: 1) we introduce p-specific regression coefficients in presence of post-treatment variables where the treatment effect on the dependent variable is allowed to vary with respect to the percentiles of an intermediate variable (Δ_p^Y); 2) we introduce causal analogues of p-specific regression coefficients which vary with respect to the percentile of the counterfactual $W_i(0)$ (τ_p^W, τ_p^Y).

Estimation methods in quantile regression are based upon finding the solution of a quantile regression minimization problem with a pre-specified loss function (Koenker and Bassett, 1978). Bayesian analogues are described by Yu and Moyeed (2001). Our estimation approach for the p-specific parameters is simply based upon transformations of the posterior samples of $(W_i(0), W_i(1))$. By modelling the marginal distributions of $W_i(0)$ and $W_i(1)$ as a mixture of normals instead of a single normal distribution, we allow very flexible shapes for the p-specific treatment effects. This gain in goodness of fit, especially at the tails of the birth weight distribution is clearly shown in Figure 2 and supported by the posterior inferences on the variances components of the mixture model.

In the causal inference literature, Angrist et al. (1996) showed how instrumental variables (IV) can be embedded within the Rubin Causal Model for estimating an average causal effect in the presence of a binary post-treatment variable. These authors introduced five assumptions under which an IV-estimator (Durbin, 1954) can be interpreted as the average causal effect. The first two assumptions are the SUTVA and the random assignment. The third assumption, called exclusion restriction, assumes that any effect of the treatment on the health outcome must be via an effect of the treatment on the post-treatment variable, that is, there is no direct effect. We are not making this assumption: we use principal stratification to compare the different causal pathways on how the supplementation affects survival. In addition Angrist et al. (1996) assume monotonicity in the post-treatment variable, that is that $W_i(1) \geq W_i(0)$, we instead define a joint model for $(W_i(0), W_i(1))$.

By specifying a joint model for for $(W_i(0), W_i(1))$, which allows for the correlation ρ between the normalized percentiles of $W_i(0)$ and $W_i(1)$, we provide a stochastic generalization of the rank preservation assumption (Efron and Feldam, 1991) similar to the one recently proposed by Dobbin and Louis (2003). More specifically, the hypothesis of rank preservation (also called percentiles invariance) implies that, for any group of participants the birth weight percentiles would not be permuted if the group had been assigned to another treatment. In our model specification for the birth weights, the percentile invariance assumption leads to $\rho = 1$ which also implies that all the

population parameters Δs are equal to the to causal parameters τs.

The methodology we described has broad applicability to a variety of situation in which one investigates a continuous post-treatment variable that potentially mediates a binary response of interest. For example, similar issues arise in cancer trials that evaluate both tumor growth and survival. In these studies it is likely that there are both mediated and direct effects of treatments, and that these effects may vary across the distribution of tumor growths. In summary, we have provided an inferential framework for estimating causal effects in a randomized trial with a continuous post-treatment variable. By comparing population with causal parameter estimates, carrying out sensitivity analyses, and implementing principal stratification, we have characterized the amount of evidence supporting the scientific questions of interest and their sources of uncertainty.

The estimation of treatment effects by percentile of the birth weight distribution has public health significance. In the case study presented here, the treatment increased the birth weight of smaller babies and had no apparent effect on larger babies. Had it increased the size of the larger infants, both the infants and their mothers might have been at higher risk of mortality given the absence of obstetrical care in rural communities. In such a situation, it would be necessary to predict those mothers who are likely to have larger infants and to exclude them from intervention programs. However, while maternal pre-pregnancy nutritional status, weight gain during pregnancy and other factors are strong determinants of low birth weight, their ability to predict infants likely to be born with low birth weight is still uncertain.

Currently recommendations exist for supplementing women with iron-folic acid during pregnancy in developing countries. The Nepal study (Christian et al., 2003a) demonstrates that beyond reducing anemia, iron can result in an improvement in birth weight primarily through moving the lower tail of the birth weight distribution to the right. Presumably, this effect is mediated through improving the iron status of those pregnant women who are the most iron deficient. These data from Nepal reveal that when evaluating public health interventions it is important to be, at the very least, cognizant of the differential beneficial effects of an intervention depending on where in the distribution the program participants fall and that an overall effect size may: 1) under-estimate the maximum likely benefit in the most malnourished individuals; and 2) incorrectly assume benefits where none exist and potentially mask harm in the more well-nourished individuals.

References

Albert, J. H. and Chib, S. (1993). "Bayesian Analysis of Binary and Polychotomous Response Data." *Journal of the American Statistical Association*, 88, 669–679.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). "Identification of Causal Effects Using Instrumental Variables (Disc: P456-472)." *Journal of the American Statistical Association*, 91, 444–455.

Begg, C. B. and Leung, D. H. Y. (2000). "On the Use of Surrogate End Points in Randomized Trials (Pkg: P15-28)." *Journal of the Royal Statistical Society, Series A, General*, 163, 1, 15–24.

Buyse, M. and Molenberghs, G. (1998). "Criteria for the Validation of Surrogate Endpoints in Randomized Experiments." *Biometrics*, 54, 1014–1029.

Caulfield, L., Zavaleta, N., Figueroa, A., and Leon, Z. (1999). "Maternal Zinc Supplementation does not affect size at birth and pregnancy duration in Peru." *Journal of Nutrion*, 129, 1563–8.

Ceesay, S., Prenctice, A., Cole, T., Foord, F., Weaver, L., and Poskitt, E. e. a. (1997). "Effects on birth weight and perinatal mortality of maternal dietary supplements in rural Gambia: 5 randomized controlled trials." *British Medical Journal*, 315, 786–790.

Chib, S. and Greenberg, E. (1998). "Analysis of Multivariate Probit Models." *Biometrika*, 85, 347–361.

Child Health Reserach Project (1996). "Special Report, Reducing perinatal and neonatal mortality, Report of a meeting." *Child Health Research Project, Baltimore MD*.

Christian, P., Khatry, S., Katz, J., Pradhan, E., LeClerq, S., Shrestha, S., Adhikari, R., Sommer, A., and West, K. (2003a). "Effects of alternative maternal micronutrient supplements on low birth weight in rural Nepal: double blind randomised community trial." *British Medical Journal*, 326, 1–6.

Christian, P., West, K., Khatry, S., Leclerq, S., Pradhan, E., Katz, J., Shrestha, S., and Sommer, A. (2003b). "Effects of maternal micronutrient supplementation on fetal loss and infant mortality: a cluster-randomized trial in Nepal." *American Journal of Clinical Nutrition*, 78, 1194–1202.

Cole, S. and Hernan, M. (2002). "Fallibility in Estimating Direct Effects." *International Journal of Epidemiology*, 31, 163–165.

Cowles, M. K. (2002). "Bayesian Estimation of the Proportion of Treatment Effect Captured By a Surrogate Marker." *Statistics in Medicine*, 21, 6, 811–834.

Cox, D. (1984). "Interaction." *International Statistical Review*, 52, 1–31.

Daniels, M. J. and Hughes, M. D. (1997). "Meta-analysis for the Evaluation of Potential Surrogate Markers." *Statistics in Medicine*, 16, 1965–1982.

Diebolt, J. and Robert, C. P. (1994). "Estimation of Finite Mixture Distributions Through Bayesian Sampling." *Journal of the Royal Statistical Society, Series B, Methodological*, 56, 363–375.

Dobbin, K. and Louis, T. (2003). "Accomodating stochastic departures from percentiles invariance in causal models." *Journal of the Royal Statistical Association*, 65, 837–850.

Doksum, K. A. and Sievers, G. L. (1976). "Plotting With Confidence: Graphical Comparisons of Two Populations." *Biometrika*, 63, 421–434.

Dominici, F., Cope, L., Naiman, D., and Zeger, S. L. (2003). "Smooth Quantile Ratio Estimation (SQUARE)." Technical report, Johns Hopkins University, Baltimore MD.

Durbin, J. (1954). "Errors in Variables." *Review of the International Statistical Institute*, 22, 23–32.

Efron, B. and Feldam, D. (1991). "Compliance as an Explanatory Variable in Clinical Trials." *Journal of the American Statistical Association*, 86, 9–17.

Frangakis, C. E. and Rubin, D. B. (2002). "Principal Stratification in Causal Inference." *Biometrics*, 58, 1, 21–29.

Freedman, L., Graubard, B., and Schatzkin, A. (1992). "Statistical validation of intermediate endpoints for chronic diseases." *Statistics in Medicine*, 11, 167–179.

Garner, P., Kramer, M., and Chalmers, L. (1992). "Might efforts to increase birth weight in undernourished women do more harm than good?" *Lancet*, 340, 1021–1022.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (1995). *Bayesian Data Analysis*. London: Chapman and Hall.

Holland, P. (1986). "Statistics and Causal Inference." *Journal of American Statistical Association*, 81, 945–960.

Kadane, J. (1974). "The role of identification in Bayesian theory." In *Studies in Bayesian Econometrics and Statistics*, 175–191. Amsterdam: North-Holland.

Katz, J., West, J. J., Khatry, S., Pradhan, E., and LeClerq, S. (2000a). "Low-dose witamin A or beta-carotene supplementation does not reduce early infant mortality: a double masked, randomized controlled community trial in Nepal." *American Journal of Clinical Nutrition*, 71, 1570–1576.

Katz, J., West, K., Khatry, S., LeClerq, S., Christian, P., Pradhan, E., and Shrestha, S. (2001). "Twinning rates and survival of twins in rural Nepal." *International Journal of Epidemiology*, 30, 802–7.

Katz, J., West, K., Khatry, S., Pradhan, E., LeClerq, S., Christian, P., Wu, L., Adhikari, R., Shrestha, S., and Sommer, A. (2000b). "Maternal low-dose vitamin A or beta-carotene supplementation has no effect on fetal loss and early infant mortality: a randomized cluster trial in Nepal." *American Journal of Clinical Nutrition*, 71, 1570–6.

Koenker, R. and Bassett, G. J. (1978). "Regression Quantiles." *Econometrica*, 46, 33–50.

Koenker, R. and Geling, O. (2001). "Reappraising Medfly Longevity: A Quantile Regression Survival Analysis." *Journal of the American Statistical Association*, 96, 454, 458–468.

Lechtig, A., Yarbrough, C., Delgado, H., Habicht, J., Marorelli, R., and Klein, R. (1975). "Influence of maternal nutrition on birth weight." *American Journal of Clinical Nutrition*, 28, 1223–1233.

Leung, D. H.-Y. (2001). "Statistical Methods for Clinical Studies in the Presence of Surrogate End Points." *Journal of the Royal Statistical Society, Series A, General*, 164, 3, 485–503.

Liang, K.-Y., Zeger, S., and Qaqish, B. (1992). "Multivariate regression analyses for categorical data (with discussion)." *Journal of the RoyaL Statistical Society, B*, 54, 3–40.

McIntire, D., Bloom, S., Casey, B., and Leveno, K. (2001). "Birth weight in relation to morbidity and mortality among newborn infants." *New England Journal of Medecine*, 340, 1234–1238.

Molenberghs, G., Geys, H., and Buyse, M. (2001). "Evaluation of Surrogate Endpoints in Randomized Experiments With Mixed Discrete and Continuous Outcomes." *Statistics in Medicine*, 20, 20, 3023–3038.

Pearl, J. (2000). *Casuality: Models, Reasoning, and Inference*. Cambridge University Press.

Prentice, R. L. (1986). "Binary Regression Using An Extended Beta-binomial Distribution, With Discussion of Correlation Induced By Covariate Measurement Errors." *Journal of the American Statistical Association*, 81, 321–327.

Rahmathullah, L., Tielsch, J., Thulasiraj, R., Katz, J., Coles, C., Devi, S., John, R., Sadanand, A., and Edwin, K. (2003). "Impact of Newborn Vitamin A Dosing on Early Infant Mortality: A Community-Based Randomized Trial in South India." *British Medical Journal*, 327, 254–7.

Rasmussen, K. M. (2001). "Is there a causal relationship between iron deficiency anemia and weight at birth, length of gestation and perinatal mortality?" *Journal of Nutrition*, 131, 590–603S.

Robins, J. (1989). "The control for confounding by intermidiate variables." *Statistics in Medicine*, 8, 679–701.

Robins, J. M. and Greenland, S. (1992). "Identifiability and Exchangeability for Direct and Indirect Effects." *Epidemiology*, 3, 143–155.

Rubin, D. B. (1978). "Bayesian Inference for Causal Effects: The Role of Randomization." *The Annals of Statistics*, 6, 34–58.

Tanner, M. A. (1991). *Tools for Statistical Inference – Observed Data and Data Augmentation Methods*, vol. 67 of *Lecture Notes in Statistics*. New York: Springer-Verlag.

Tanner, M. A. and Wong, W. H. (1987). "The calculation of posterior distributions by data augmentation." *Journal of the American Statistical Association*, 82, 398, 528–550.

Tierney, L. (1994). "Markov chains for exploring posterior distributions (with Discussion)." *Annals of Statistics*, 22, 4, 1701–1762.

van der Laan, M. and Petersen, M. (2004). "Estimation of Direct and Indirect Causal Effects in Longitudinal Studies." Tech. rep., No 1155, U.C. Berkeley Division of Biostatistics Working Paper Series from Berkeley Electronic Press.

West, J. J., Katz, J., Khatry, S.K. LeClerq, S., and Pradhan, E. (1999). "Double blind, clustered randomized trial of low dose supplementation with vitamin A or beta-caroten on mortality related to pregnancy in Nepal. The NNIPS-2 study group." *British Medical Journal*, 318, 570–575.

Xu, J. and Zeger, S. L. (2001). "The Evaluation of Multiple Surrogate Endpoints." *Biometrics*, 57, 1, 81–87.

Yu, K. and Moyeed, R. A. (2001). "Bayesian Quantile Regression." *Statistics & Probability Letters*, 54, 4, 437–447.

Table 1: Definition of population and causal parameters for estimating the effects of antenatal iron-folic acid supplementation on birth weight as a function of birth weight percentiles. $Q_1(p)$ and $Q_0(p)$ are the quantile functions of $W_i(1)$ and $W_{i'}(0)$, respectively. The parameters ρ and ψ measure the correlation between $W_i(0)$ and $W_i(1)$ and the odds-ratio between $Y_i(0)$ and $Y_i(1)$. The subscripts i and i' indicate two different infants.

Percentile-specific Effects on Birth Weight	
	Population Parameters
Average	$\Delta^W = E[W_i(1)] - E[W_{i'}(0)] = E[W_i^{obs} \mid Z_i = 1] - E[W_{i'}^{obs} \mid Z_{i'} = 0]$
p-specific	$\Delta_p^W(\rho) = E[W_i(1) \mid F_1(W_i(1)) = p] - E[W_{i'}(0) \mid F_0(W_{i'}(0)) = p] = Q_1(p) - Q_0(p)$
	Causal Parameters
Average	$\tau^W(\rho, \psi) = E[W_i(1) - W_i(0)]$
p-specific	$\tau_p^W(\rho, \psi) = E[W_i(1) - W_i(0) \mid F_0(W_i(0)) = p]$

Table 2: Definition of population and causal parameters for estimating the effects of antenatal iron-folic acid supplementation on infant mortality as a function of the birth weight percentiles. The parameters ρ and ψ measure the correlation between $W_i(0)$ and $W_i(1)$ and the odds-ratio between $Y_i(0)$ and $Y_i(1)$. The subscripts i and i' indicate two different infants.

Percentile-specific Effects on Mortality	
	Population Parameters
Average	$\Delta^Y = E[Y_i(1)] - E[Y_{i'}(0)] = E[Y_i^{obs} \mid Z_i = 1] - E[Y_{i'}^{obs} \mid Z_{i'} = 0]$
p-specific	$\Delta_p^Y = E[Y_i(1) \mid F_1(W_i(1)) = p] - E[Y_{i'}(0) \mid F_0(W_{i'}(0)) = p]$
	Causal Parameters
Average	$\tau^Y(\rho, \psi) = E[Y_i(1) - Y_i(0)]$
p-specific	$\tau_p^Y(\rho, \psi) = E[Y_i(1) - Y_i(0) \mid F_0(W_i(0)) = p]$
P-Stratification	$\begin{cases} \tau_1^Y(\rho, \psi) = E[Y_i(1) - Y_i(0) \text{ given } W_i(0) \leq 2500 \ \& \ \mid W_i(1) - W_i(0) \mid \leq 50] \\ \tau_2^Y(\rho, \psi) = E[Y_i(1) - Y_i(0) \text{ given } W_i(0) \leq 2500 \ \& \ \mid W_i(1) - W_i(0) \mid > 50] \\ \tau_3^Y(\rho, \psi) = E[Y_i(1) - Y_i(0) \text{ given } W_i(0) > 2500 \ \& \ \mid W_i(1) - W_i(0) \mid \leq 50] \\ \tau_4^Y(\rho, \psi) = E[Y_i(1) - Y_i(0) \text{ given } W_i(0) > 2500 \ \& \ \mid W_i(1) - W_i(0) \mid > 50] \end{cases}$

Parameter	Prior distribution
$\boldsymbol{\beta}$	flat
$\boldsymbol{\mu}_0$	$N_3\left[(1500, 2500, 3500), 500^2 I\right]$
$\boldsymbol{\mu}_1$	$N_3\left[(2000, 3000, 3500), 500^2 I\right]$
σ_0^2	$LN(\log(400^2), 0.8)$
σ_1^2	$LN(\log(400^2), 0.8)$
$\boldsymbol{\gamma}_0$	Dirichlet$(10, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
$\boldsymbol{\gamma}_1$	Dirichlet$(10, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

Table 3: Prior distributions on the unknown parameters of the mixture. I denotes a 3×3 identity matrix, 500^2 denotes the prior variance, LN denotes the log-normal distribution with prior mean 400 and prior standard deviation 0.8.

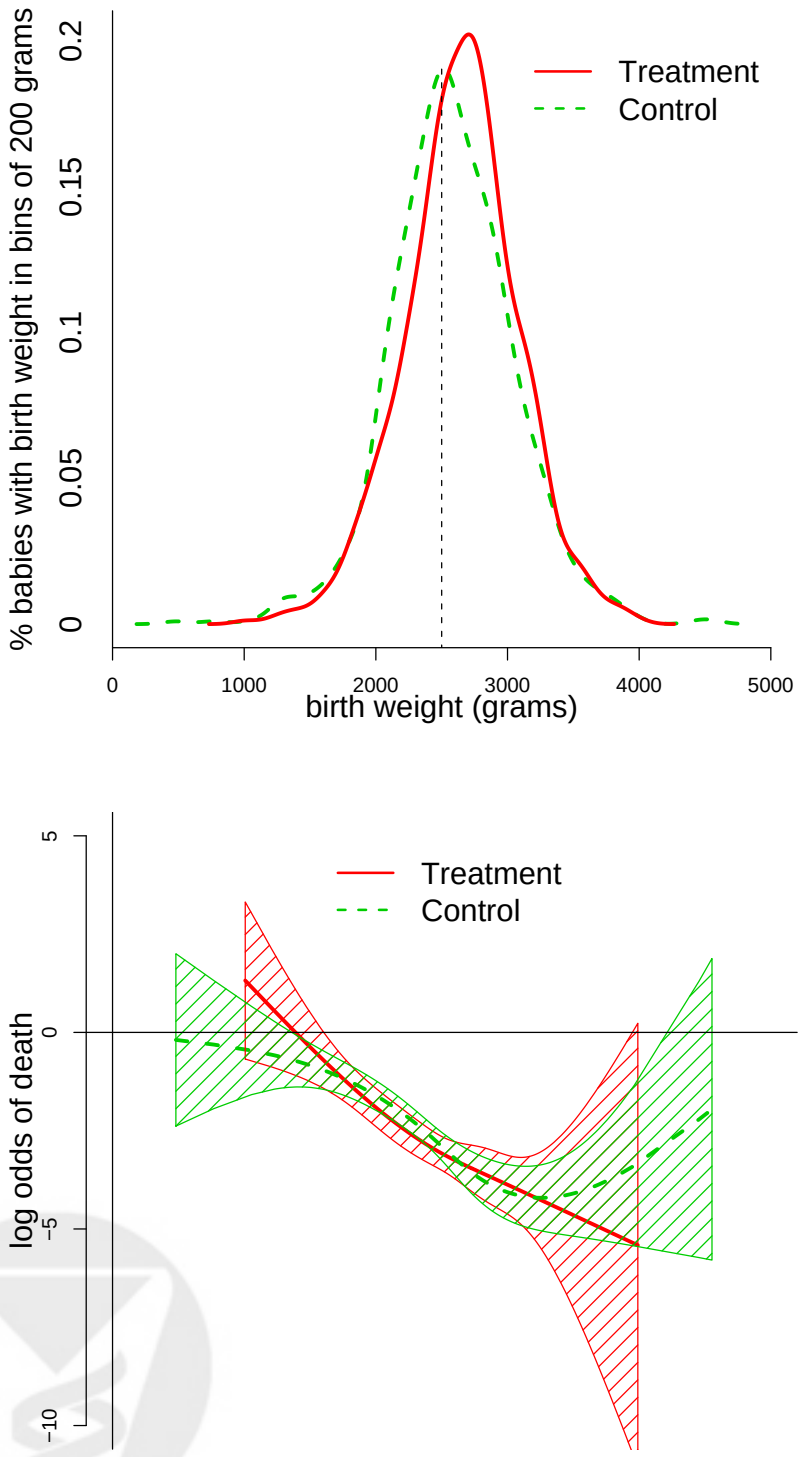

Figure 1: Top: smoothed histograms of the birth weights for the treated and the control groups. Bottom: estimated log-odds of death as smooth function of the birth weight with 95% confidence bands and plotted in correspondence to the observed range of birth weights in the two groups.

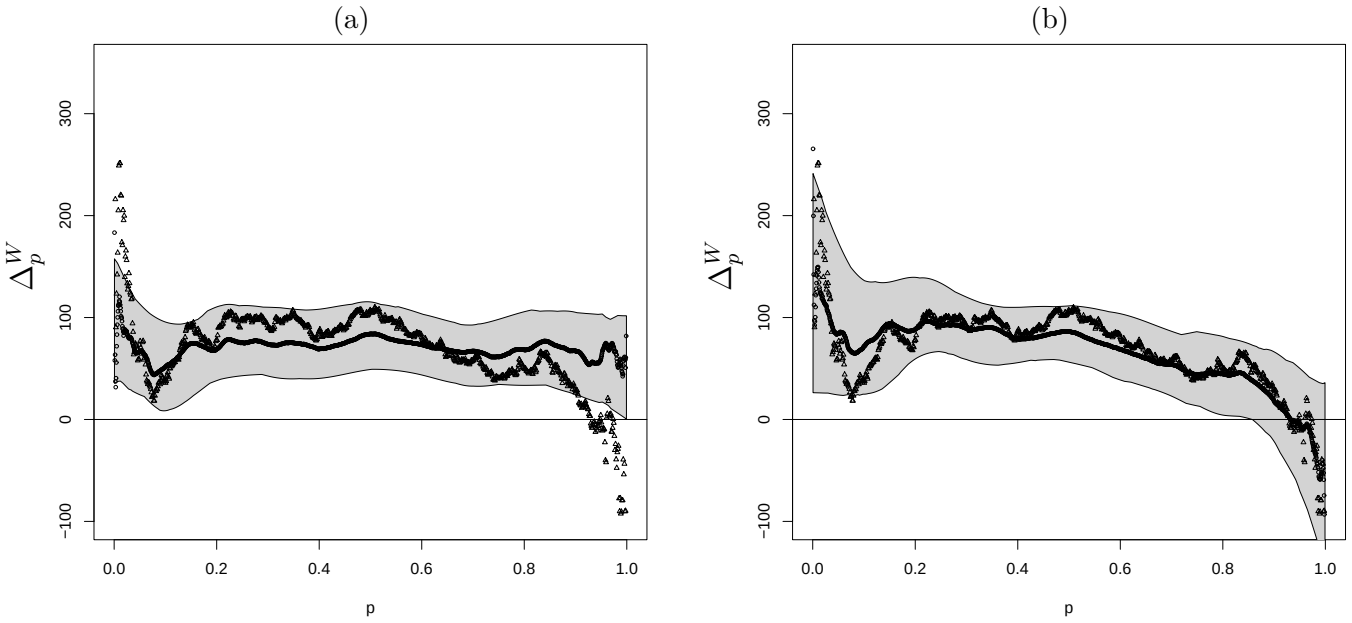

Figure 2: Posterior means and 95% posterior regions of the p-specific effects of treatment on birth weight (Δ_p^W) under the following modelling assumptions for $(W_i(0), W_i(1))$: a) $W_i(0), W_i(1)$ have a bivariate normal distribution with equal variances; b) $W_i(0), W_i(1)$ have a mixture of normal distributions with correlation ρ as defined in Section 2. The triangles denote the differences between the empirical quantile functions for the observed data. The black dots denote the posterior means of Δ_p^W as a function of p.

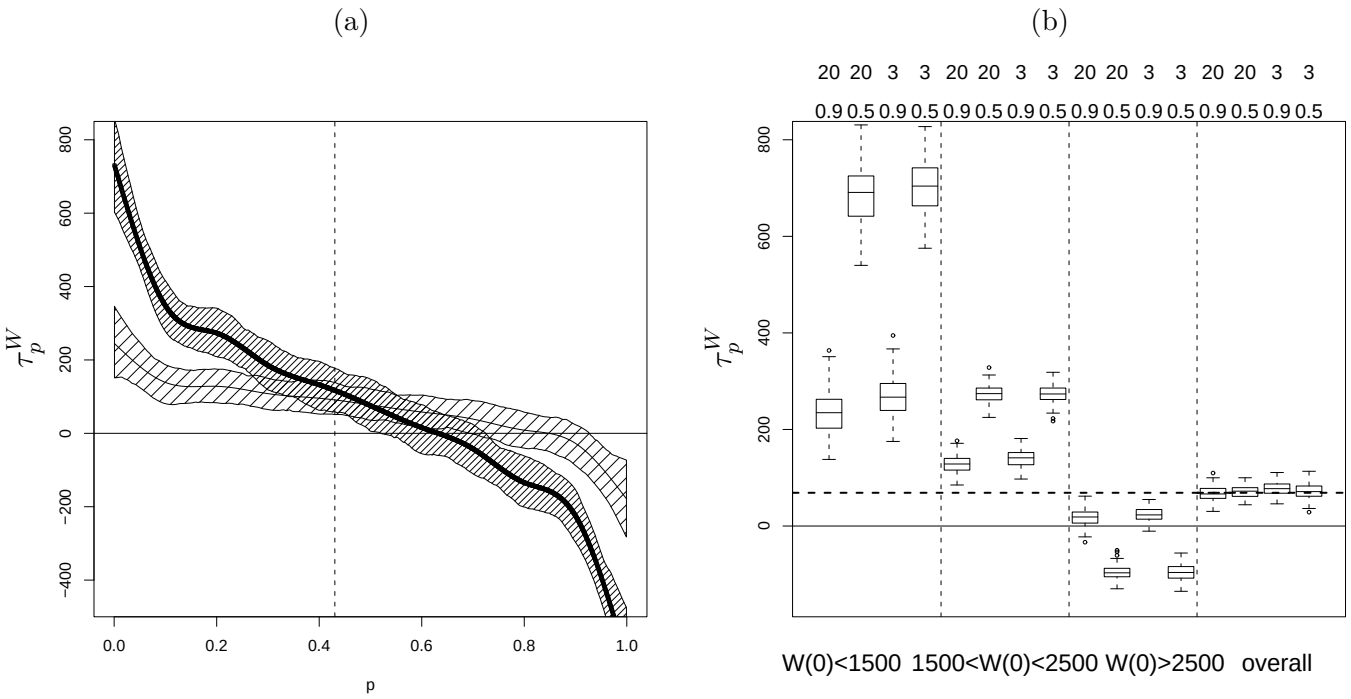

Figure 3: Panel (a): Posterior means and 95% posterior regions of the p-specific causal effects of treatment on birth weight (τ_p^W) for $\rho = 0.9$ and for $\rho = 0.5$ (darker polygon). The vertical dotted line is placed at the 0.42 percentile, corresponding to 2500 grams in the control distribution. Panel (b): sensitivity analysis of the posterior distributions of the causal effect of treatment on birth weight (τ_p^W) separately for three sub-populations of babies $W_i(0) \leq 1500; 1500 < W_i(0) \leq 2500; W_i(0) < 2500$ and overall for all babies with respect to (ρ, ψ). The horizontal dotted line is placed at the sample mean difference (Δ^W).

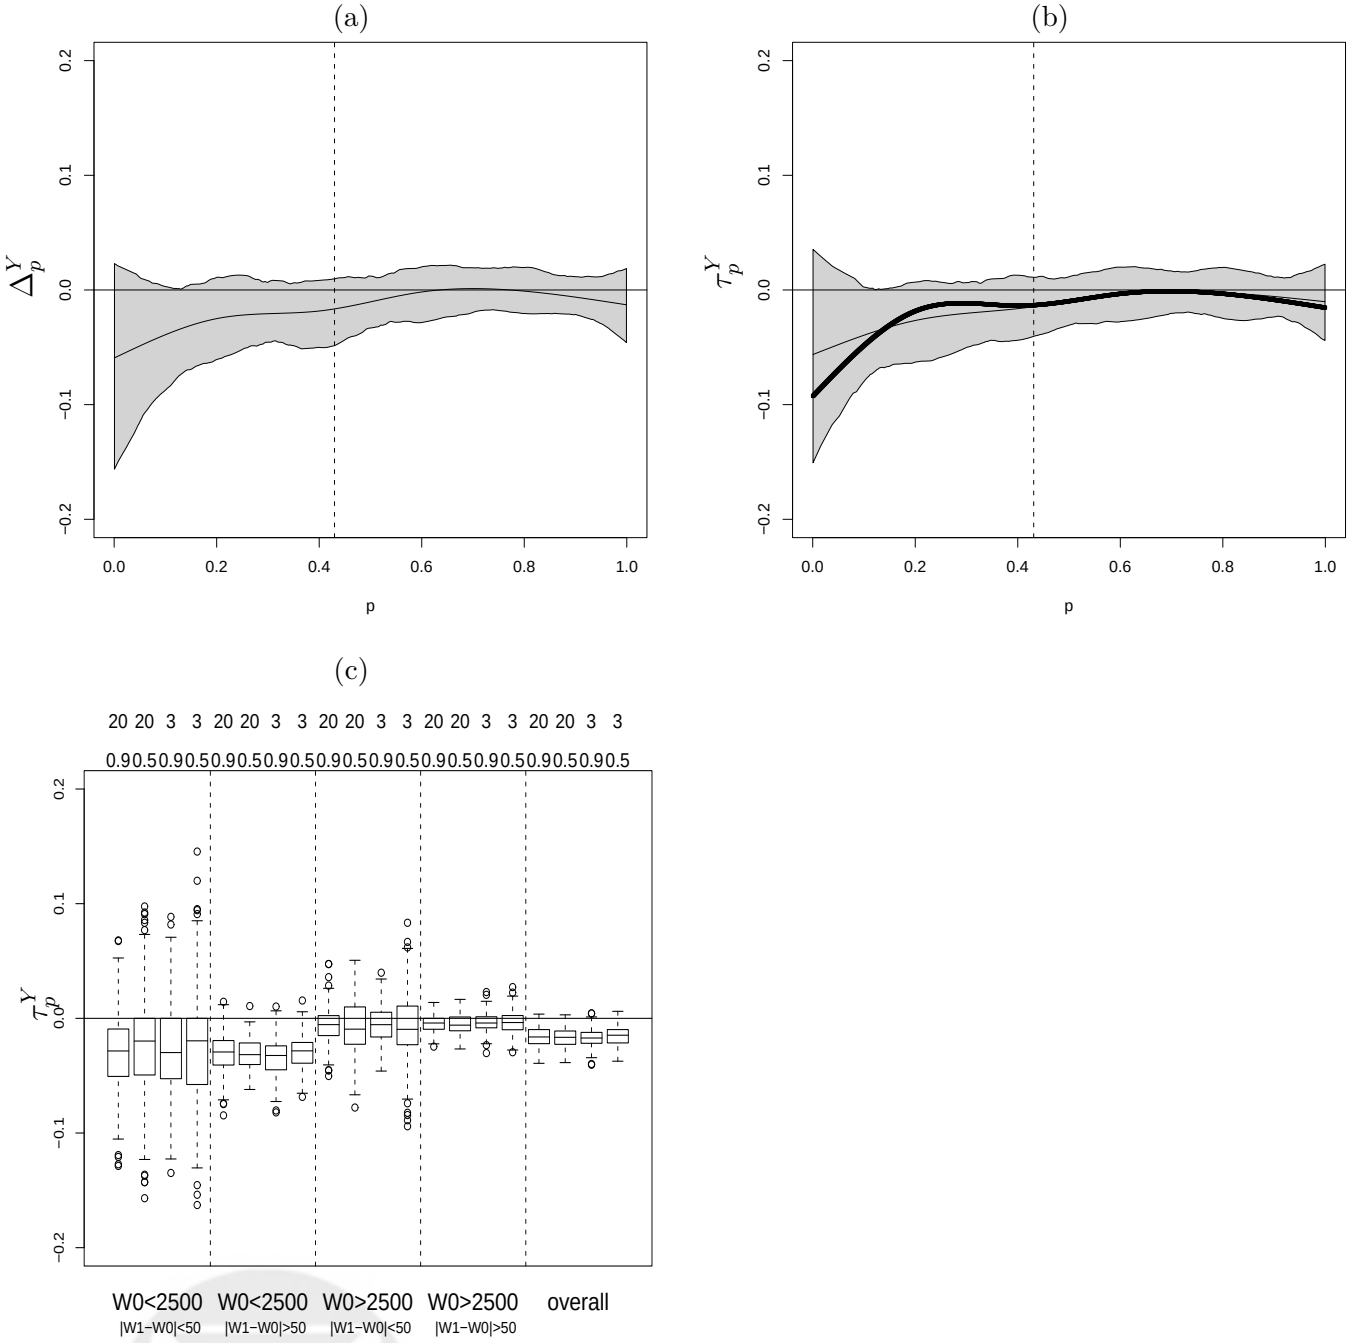

Figure 4: Panel a: posterior means and 95% posterior regions of the p-specific effects of treatment on mortality (Δ_p^Y). Panel b: posterior means and 95% posterior regions of the p-specific causal effects of treatment on mortality (τ_p^Y) for $\rho = 0.9$ and for $\rho = 0.5$ (darker line). Panel c: posterior distributions of the causal effects of treatment on mortality (τ_p^Y) for different values of (ρ, ψ). The posterior distributions are shown separately for five sub-populations of infants: 1) LBW infants for whom there is causal effect of treatment on birth weight smaller than 50 grams; 2) LBW infants for whom there is a causal effect of treatment on birth weight larger than 50 grams; 3) not-LBW for whom there is a causal effect of treatment on birth weight smaller than 50 grams; 4) not-LBW for whom there is a causal effect of treatment on birth weight larger than 50 grams; and 5) all infants.

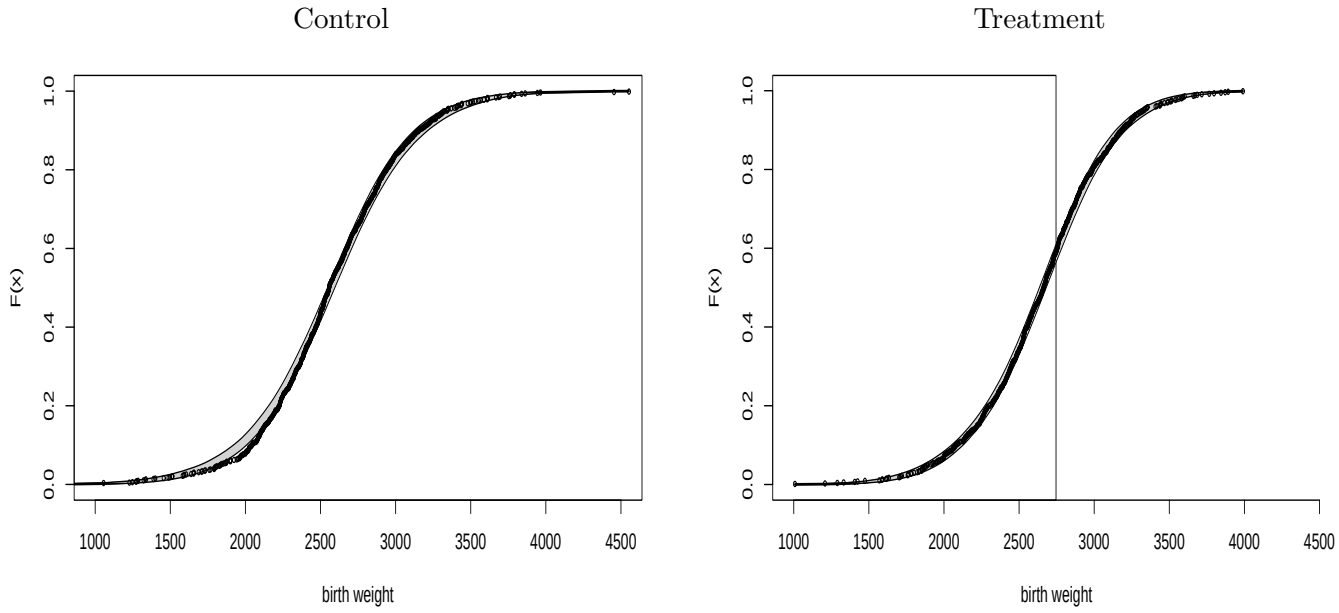

Figure 5: Left and right: 95% posterior regions of $F_0(W_i^{obs}, \boldsymbol{\theta}_0^j)$ and $F_1(W_i^{obs}, \boldsymbol{\theta}_1^j)$ where F_0, F_1 are the cdf of the mixture of three normal distributions, and $\boldsymbol{\theta}_0^{(j)}, \boldsymbol{\theta}_1^{(j)}$ are the the j^{th} posterior sample of the vector of parameters of the mixture. The black lines are the corresponding empirical cdf.