Research Article

Statistics in Medicine

Received 13 September 2010, Accepted 14 October 2011 Published online in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.4453

Hierarchical priors for bias parameters in Bayesian sensitivity analysis for unmeasured confounding

Lawrence C. McCandless,[a*†] Paul Gustafson,[b] Adrian R. Levy[c] and Sylvia Richardson[d]

Recent years have witnessed new innovation in Bayesian techniques to adjust for unmeasured confounding. A challenge with existing methods is that the user is often required to elicit prior distributions for high-dimensional parameters that model competing bias scenarios. This can render the methods unwieldy. In this paper, we propose a novel methodology to adjust for unmeasured confounding that derives default priors for bias parameters for observational studies with binary covariates. The confounding effects of measured and unmeasured variables are treated as exchangeable within a Bayesian framework. We model the joint distribution of covariates by using a log-linear model with pairwise interaction terms. Hierarchical priors constrain the magnitude and direction of bias parameters. An appealing property of the method is that the conditional distribution of the unmeasured confounder follows a logistic model, giving a simple equivalence with previously proposed methods. We apply the method in a data example from pharmacoepidemiology and explore the impact of different priors for bias parameters on the analysis results. Copyright © 2011 John Wiley & Sons, Ltd.

Keywords: bias; observational studies; Bayesian statistics; pharmacoepidemiology

1. Introduction

1.1. Unmeasured confounding in pharmacoepidemiology

Bias from unmeasured confounding figures prominently in pharmacoepidemiology, which is concerned with improving our understanding of the effectiveness and safety of medications. A typical pharmacoepidemiology study compares outcome response rates in patients who were prescribed a medication with those who were not. Study findings are often biased without careful adjustment for the factors that influence prescribing. Unfortunately, control of confounding is notoriously difficult because medication prescribing is intimately connected to the disease process that determines the study outcome. The myriad of patient characteristics that influence prescribing can act as powerful confounders and bias effect estimates in a manner that is difficult to predict. Epidemiologists call this *confounding by indication* because the confounders are the clinical indications for treatment [1].

In this paper, we illustrate the problem of unmeasured confounding by using the data example of McCandless *et al.* [2,3]. The authors collected data from a cohort of patients who were discharged from hospital in 1999 and 2000 after treatment for heart failure. The goal of the study was to estimate the association between beta blocker therapy and mortality after 1 year of follow-up. However, the study used healthcare administrative data, which provides only basic information on the many possible confounders. A total of 21 covariates are available in the data, including patient characteristics, disease indicator variables and prescribing of cardiovascular therapies (see Table I for a complete listing).

[a] *Faculty of Health Sciences, Simon Fraser University, Canada*
[b] *Department of Statistics, University of British Columbia, Canada*
[c] *Department of Community Health and Epidemiology, Dalhousie University, Canada*
[d] *Department of Epidemiology and Biostatistics, Imperial College London, UK*
**Correspondence to: Lawrence C. McCandless, Faculty of Health Sciences, Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada.*
[†] *E-mail: mccandless@sfu.ca*

Table I. Adjusted log-odds ratios (95% interval estimates) in outcome models predicting mortality.

		Bayesian sensitivity analysis	
	Naive analysis	$U \perp\!\!\!\perp \boldsymbol{C} \mid X$	$U \not\perp\!\!\!\perp \boldsymbol{C} \mid X$
Beta blocker use	−0.44 (−0.80, −0.07)	−0.41 (−0.87, 0.06)	−0.40 (−0.90, 0.11)
Disease indicator variables			
CVD	0.40 (−0.33, 1.14)	0.22 (−0.45, 0.81)	0.23 (−0.40, 0.79)
COPD	0.01 (−0.44, 0.45)	−0.13 (−0.58, 0.34)	−0.13 (−0.54, 0.20)
HYPNAT	0.05 (−0.39, 0.48)	−0.07 (−0.56, 0.35)	−0.07 (−0.47, 0.25)
MTSTD	1.70 (0.97, 2.44)	1.35 (0.63, 1.97)	1.34 (0.69, 1.82)
MSRD	0.65 (0.24, 1.07)	0.51 (0.05, 0.93)	0.50 (0.12, 0.79)
VENTRAR	0.27 (−0.44, 0.99)	0.12 (−0.49, 0.75)	0.13 (−0.51, 0.64)
MLD	0.39 (−0.23, 1.00)	0.22 (−0.39, 0.74)	0.22 (−0.33, 0.70)
CAN	1.01 (0.52, 1.51)	0.87 (0.37, 1.33)	0.87 (0.41, 1.23)
CARS	−0.09 (−0.73, 0.56)	−0.19 (−0.78, 0.37)	−0.18 (−0.74, 0.32)
Patient characteristics			
Sex: female	−0.35 (−0.60, −0.10)	−0.36 (−0.61, −0.13)	−0.36 (−0.60, −0.12)
Age (in years)	0.03 (0.02, 0.04)	0.03 (0.02, 0.04)	0.03 (0.02, 0.04)
Characteristics of hospitalisation			
Transferred admission	0.75 (0.39, 1.12)	0.77 (0.41, 1.13)	0.76 (0.43, 1.12)
Hospital stay (no. of days)	0.01 (0.00, 0.01)	0.01 (0.00, 0.01)	0.01 (0.00, 0.01)
Heart failure medications			
Digoxin	0.30 (0.00, 0.60)	0.30 (−0.01, 0.59)	0.30 (0.00, 0.58)
Diuretics	−0.04 (−0.32, 0.25)	−0.04 (−0.30, 0.25)	−0.04 (−0.31, 0.24)
CCB	−0.31 (−0.66, 0.04)	−0.32 (−0.67, −0.01)	−0.33 (−0.67, −0.01)
ACE inhibitor	−0.01 (−0.29, 0.27)	−0.02 (−0.28, 0.25)	−0.02 (−0.30, 0.24)
ARB	0.13 (−0.58, 0.84)	0.12 (−0.57, 0.78)	0.11 (−0.59, 0.77)
Statin	−0.62 (−1.10, −0.13)	−0.64 (−1.12, −0.20)	−0.64 (−1.10, −0.19)

Let X and Y denote binary treatment and outcome variables, respectively. We set X equal to one if the patient was dispensed a beta blocker within 30 days of hospital discharge and zero otherwise. Similarly, we let Y denote an indicator variable for death within 1 year of hospital discharge. In pharmacoepidemiological studies of cardiovascular therapies, interest centres on confounding induced by the various patient illnesses. Let $\boldsymbol{C} = (C_1, \ldots, C_p)$ denote the $p = 9$ dimensional vector of disease indicator variables listed in Table I, which include cerebrovascular disease (CVD), chronic obstructive pulmonary disorder (COPD), hyponatremia (HYPNAT), metastatic disorder (MTSTD), renal disease (MSRD), ventricular arrhythmia (VENTRAR), liver disease (MLD), cancer (CAN) and cardiogenic shock (CARS). Additionally, let $\boldsymbol{Z} = (Z_1, \ldots, Z_q)$ denote $q = 21 - 9 = 12$ dimensional vector of the remaining 12 non-disease indicator variables listed in the bottom half of Table I.

In this article, we focus on the data set consisting of $n = 1299$ study participants that have at least one of the co-morbidities in Table I. To estimate the association between beta blocker therapy and mortality while adjusting for confounding, we fit a logistic regression with Y as the dependent variable and $(X, \boldsymbol{C}, \boldsymbol{Z})$ as the independent variables. We present the results in the first column of Table I under the heading 'Naive analysis'. The table displays regression coefficients, which are log-odds ratios, with 95% interval estimates. We estimate the regression coefficient for the treatment effect X as −0.44 with 95% interval estimate (−0.80, −0.07), suggesting that beta blocker therapy reduces mortality. The corresponding odds ratio $\exp(-0.44) = 0.64$ roughly agrees with estimates reported from randomised trials of beta blockers and heart failure. In a scientific review of meta-analyses of randomised trials, Foody *et al.* [4] found that beta blocker use is associated with a 30% reduction in mortality compared with placebo.

Nonetheless, there are concerns about unmeasured confounding. This analysis uses healthcare administrative data with limited clinical information on the factors that influence prescribing of beta blockers. For example, one unmeasured confounder is severity of heart failure (i.e. the degree of low cardiac output and the resulting impairment in physical activity). Severe heart failure is characterised by systolic dysfunction, which is failure of the pump function of the heart. The severity of heart failure falls on a spectrum from high to low and is measured by the ejection fraction, which is a continuous variable that typically ranges between 50% and 70% but which falls below 50% for severe heart failure. An alternative

measure of the severity of heart failure is the New York Heart Association class of heart failure (see [4] and [5] for a review). This is an ordinal variable with four levels, which measures heart failure symptoms and limitations in physical activity. Severity of heart failure is an important predictor of mortality and treatment. It can either increase or decrease the probability of receiving a beta blocker, depending on the preferences of the prescribing physician [6].

1.2. Bayesian sensitivity analysis for unmeasured confounding and the challenges of prior elicitation for bias parameters

A typical sensitivity analysis for unmeasured confounding posits the existence of an unmeasured binary variable U that confounds the association between X and Y. Paralleling existing modelling frameworks (e.g. [1–3, 7–10]), we model the probability density $P(Y, U|X, \boldsymbol{C}, \boldsymbol{Z}) = P(Y|X, \boldsymbol{C}, \boldsymbol{Z}, U)P(U|X, \boldsymbol{C})$ where

$$\text{logit}[P(Y = 1|X, \boldsymbol{C}, \boldsymbol{Z}, U)] = \beta_0 + \beta_{XY}X + \boldsymbol{\beta_{CY}}^T\boldsymbol{C} + \boldsymbol{\beta_{ZY}}^T\boldsymbol{Z} + \beta_{UY}U \tag{1}$$

$$\text{logit}[P(U = 1|X, \boldsymbol{C})] = \gamma_U + \gamma_{XU}X + \boldsymbol{\gamma_{CU}}^T\boldsymbol{C}. \tag{2}$$

Equation (1) includes U as a missing covariate in the regression model for the outcome. Equation (2) characterises the distribution of the missing confounder within levels of $(X, \boldsymbol{C})$.

The quantity β_{XY} is the parameter of primary interest and is the causal log-odds ratio for the effect of X on Y conditional on $(\boldsymbol{C}, \boldsymbol{Z}, U)$. Provided that all models are correctly specified and that there are no *additional* unmeasured confounders, then the parameter β_{XY} has a causal interpretation. The quantities β_{UY}, γ_U, γ_{XU} and $\boldsymbol{\gamma_{CU}}$ are *bias parameters* because they determine the magnitude of unmeasured confounding. The parameter β_{UY} governs the association between U and Y, conditional on $(X, \boldsymbol{C}, \boldsymbol{Z})$, whereas the parameters $(\gamma_{XU}, \boldsymbol{\gamma_{CU}})$ capture the association between U and $(X, \boldsymbol{C})$. The quantity $\exp(\gamma_U)/(1 + \exp(\gamma_U))$ is the prevalence of $U = 1$ when $X = 0$ and $C_1, \ldots, C_p = 0$ (see Table II for a detailed explanation of the variables and parameters).

In this investigation, the unmeasured confounder is severity of heart failure. We classify patients as having severe ($U = 1$) versus not severe ($U = 0$) heart failure. Thus, the binary variable U is a hypothetical dichotomy to quantify severity, which falls on a spectrum from high to low. Ideally, we would model U as a continuous variable. However, because severity is unmeasured to begin with, there is necessarily a trade-off between realistic modelling versus minimising complexity of the model. Furthermore, we argue that it is reasonable to classify patients into two categories of severity because ejection fraction is dichotomised in clinical practice by using a threshold of 50% (see [5] for a discussion of dichotomisation). In Section 4, we discuss the limitations of assuming a binary unmeasured confounder and extensions to the ordinal case by using a series of indicators.

Because U is unmeasured, the data provide no information about the relationship between U and the measured variables $(Y, X, \boldsymbol{C}, \boldsymbol{Z})$, and the model is non-identifiable. But non-identifiability does not preclude Bayesian model fitting if additional sources of information are incorporated. A Bayesian analysis would start by assigning proper prior distributions to model parameters that translate beliefs about the magnitude and direction of confounding by U. We then study the posterior distribution for the treatment effect β_{XY} integrating over the unmeasured confounder U. Posterior credible intervals for the treatment effect incorporate uncertainty from unmeasured confounding in addition to random error [2, 3, 11–13].

A difficulty with Bayesian analysis is eliciting prior distributions for the bias parameters. In particular, the quantities γ_U, γ_{XU}, $\boldsymbol{\gamma_{CU}}$ consist of $p + 2$ unique parameters that characterise how U is distributed within levels of X and $\boldsymbol{C}$. In many applications, it is burdensome to obtain reasonable prior guesses for $\boldsymbol{\gamma_{CU}}$, which describes the association between $\boldsymbol{C}$ and U given X. For this problem to be mitigated, most sensitivity analysis techniques assume that the unmeasured confounder is independent of measured confounders, conditional on treatment (i.e. that they are not correlated with one another). Mathematically, we write $U \perp\!\!\!\perp \boldsymbol{C}|X$, where '$\perp\!\!\!\perp$' denotes conditional independence ([1–3, 7–9] for examples and [1, 14–16] for a discussion). In Equations (1) and (2), this assumption forces $\boldsymbol{\gamma_{CU}} = \boldsymbol{0}$, where $\boldsymbol{0}$ is a zero vector of length p, and then explores sensitivity for the remaining bias parameters β_{UY}, γ_{XU} and γ_U. Statistical methods that *do not* require $U \perp\!\!\!\perp \boldsymbol{C}|X$ include those of Greenland [10] and Gustafson *et al.* [13], and we review these in Section 1.4.

Hernán and Robins [14] and VanderWeele [15] argue that it is unrealistic to assume that $U \perp\!\!\!\perp \boldsymbol{C}|X$. Furthermore, epidemiologists argue that such assumptions give inferences from sensitivity analysis,

Table II. Description of variables and parameters.

Quantity	Dimension	Description
Data		
Y	1	Binary outcome indicating death within 1 year of hospital discharge
X	1	Binary treatment indicating being dispensed a beta blocker within 30 days of hospital discharge
$\boldsymbol{C}$	p	Vector of p-measured confounding variables
$\boldsymbol{Z}$	q	Vector of q-measured covariates that are not confounding variables
U	1	Binary unmeasured confounding variable
Parameters		
β_0	1	y-intercept in outcome model
β_{XY}	1	Causal log-odds ratio for effect of X on Y conditional on $(\boldsymbol{C}, \boldsymbol{Z}, U)$
β_{UY}	1	Log-odds ratio for association between U and Y given $(X, \boldsymbol{C}, \boldsymbol{Z})$
$\boldsymbol{\beta_{CY}}$	p	Log-odds ratios for association between $\boldsymbol{C}$ and Y given $(X, \boldsymbol{Z}, U)$
$\boldsymbol{\beta_{ZY}}$	q	Log-odds ratios for association between $\boldsymbol{Z}$ and Y given $(X, \boldsymbol{C}, U)$
γ_U	1	Main effect for U in Equation (5). Note that $\frac{\exp(\gamma_U)}{1+\exp(\gamma_U)}$ is the prevalence of $U = 1$ when $X = 0$ and $C_1, \ldots, C_p = 0$
γ_X	1	Main effect for X in Equation (5)
$\boldsymbol{\gamma_C}$	p	Main effects for $\boldsymbol{C}$ in Equation (5)
γ_{XU}	p	Log-odds ratio for association between X and U given $\boldsymbol{C}$
$\boldsymbol{\gamma_{CU}}$	p	Log-odds ratios for association between $\boldsymbol{C}$ and U given X
$\boldsymbol{\gamma_{CX}}$	p	Log-odds ratios for association between $\boldsymbol{C}$ and X given U
$\boldsymbol{\gamma_{C\oplus C}}$	$\binom{p}{2}$	Log-odds ratios for pairwise conditional associations among components of $\boldsymbol{C}$
Hyperparameters		
σ^2_β	1	Prior variance of $\beta_{XY}, \boldsymbol{\beta_{CY}}, \beta_{UY}$
μ_γ	1	Prior mean of $\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}}$
σ^2_γ	1	Prior variance of $\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}}$
μ_0	1	Prior mean of $\gamma_X, \boldsymbol{\gamma_C}, \gamma_U$
σ^2_0	1	Prior variance of $\gamma_X, \boldsymbol{\gamma_C}, \gamma_U$

which are too *pessimistic* [1, 16]. In a simulation study, Fewell *et al.* [16] demonstrate that strong correlations between measured and unmeasured confounders tend to *reduce* bias from unmeasured confounding. Intuitively, the reason is that adjusting for measured variables may control for unmeasured variables because they are correlated with one another. This reasoning suggests that forcing $\boldsymbol{\gamma_{CU}} = \boldsymbol{0}$ for convenience may actually *exaggerate* the sensitivity of the analysis results to unmeasured confounding.

1.3. Correlations between measured and unmeasured confounders in the beta blocker data

Returning to the beta blocker example, we attempt to elicit judgements about plausible values for the bias parameters β_{UY}, γ_U, γ_{XU} and $\boldsymbol{\gamma_{CU}}$. Table III describes the confounding induced by the $p = 9$ disease indicator variables $\boldsymbol{C} = (C_1, \ldots, C_p)$ listed in Table I. In Section A of Table III, we list the adjusted log-odds ratios for the association between each variable and mortality by copying and pasting from the Naive analysis column of Table I. Section B describes the pairwise conditional associations among the components of $(X, \boldsymbol{C})$. In Section B, we use maximum likelihood to fit the log-linear model

$$P(X, \boldsymbol{C}) = \frac{1}{Q(\gamma_X, \boldsymbol{\gamma_C}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})} \exp\left\{\gamma_X X + \boldsymbol{\gamma_C}^T \boldsymbol{C} + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C} X + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C} \oplus \boldsymbol{C})\right\} \quad (3)$$

for the joint distribution of $(X, \boldsymbol{C})$. The quantities γ_X and $\boldsymbol{\gamma_C}$ are the main effects of $(X, \boldsymbol{C})$ in the log-linear model, whereas the vectors $\boldsymbol{\gamma_{CX}}$ and $\boldsymbol{\gamma_{C\oplus C}}$ are coefficients of the interaction terms, with lengths $p = 9$ and $\binom{9}{2} = 36$, respectively. The quantity $\boldsymbol{C} \oplus \boldsymbol{C}$ denotes the vector of length $\binom{p}{2}$ of pairwise products among the p components of $\boldsymbol{C} = (C_1, \ldots, C_p)$. In other words,

$$\boldsymbol{C} \oplus \boldsymbol{C} = (C_1C_2, C_1C_3, \ldots, C_1C_p, C_2C_3, C_2C_4, \ldots, C_2C_p, \ldots, C_{p-1}C_p).$$

The denominator $Q(\gamma_X, \boldsymbol{\gamma_C}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$ is the constant of normalisation and is the summation of the numerator over the support of $(X, \boldsymbol{C})$, which is a set with 2^{p+1} elements.

Table III. Description of the confounding induced by the $p = 9$ disease indicator variables $\boldsymbol{C} = (C_1, \ldots, C_p)$ in the association between X and Y.

Section A: Adjusted log-odds ratios (standard errors) copied from the Naive analysis column in Table I that describes the association between $(X, \boldsymbol{C})$ and Y.

Beta blocker	CVD	COPD	HYPNAT	MTSTD	MSRD	VENTRAR	MLD	CAN	CARS
−0.44 (0.2)	0.40 (0.4)	0.01 (0.2)	0.05 (0.2)	1.70 (0.4)	0.65 (0.2)	0.27 (0.4)	0.39 (0.3)	1.01 (0.3)	−0.09 (0.3)

Section B: Log-odds ratios (standard errors) for the $\binom{10}{2} = 45$ pairwise conditional associations among individual components of $(X, \boldsymbol{C})$. vector.

	Beta blocker	CVD	COPD	HYPNAT	MTSTD	MSRD	VENTRAR	MLD	CAN	CARS
Beta blocker	—	1.1 (0.4)	−0.1 (0.3)	0.6 (0.3)	0.0 (0.5)	0.9 (0.3)	0.9 (0.4)	−0.4 (0.5)	0.6 (0.3)	0.5 (0.4)
CVD	—	—	2.1 (0.7)	1.6 (0.9)	NA	1.8 (0.7)	NA	NA	2.0 (1.2)	3.3 (0.8)
COPD	—	—	—	1.0 (0.5)	0.5 (0.8)	0.3 (0.6)	1.4 (0.9)	1.3 (0.8)	0.9 (0.7)	1.3 (0.6)
HYPNAT	—	—	—	—	0.9 (0.8)	1.6 (0.6)	1.5 (0.9)	1.1 (0.9)	1.1 (0.7)	2.3 (0.7)
MTSTD	—	—	—	—	—	1.4 (0.7)	NA	NA	4.6 (0.6)	NA
MSRD	—	—	—	—	—	—	1.7 (0.6)	1.5 (0.7)	0.8 (0.6)	1.0 (0.7)
VENTRAR	—	—	—	—	—	—	—	2.2 (1.2)	2.2 (0.8)	2.8 (0.9)
MLD	—	—	—	—	—	—	—	—	2.2 (0.8)	NA
CAN	—	—	—	—	—	—	—	—	—	1.4 (1.2)
CARS	—	—	—	—	—	—	—	—	—	—

Section C: Prevalences of individual components of $(X, \boldsymbol{C})$, given as counts from the 1299 study participants.

211	45	310	224	48	530	45	74	130	67

'NA' denotes interaction terms that were dropped from the log-linear model to obtain a valid maximum likelihood estimate.

Statistics in Medicine

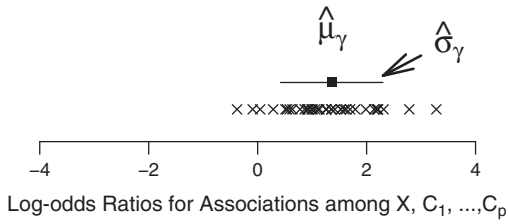

Figure 1. Sample mean ($\hat{\mu}_\gamma = 1.36$) and standard deviation ($\hat{\sigma}_\gamma = 0.94$) of the log-odds ratio estimates listed in Section B of Table III.

Section B contains point estimates and standard errors of the interaction terms ($\boldsymbol{\gamma_{CX}}$, $\boldsymbol{\gamma_{C\oplus C}}$). There is a well-known connection between log-linear models and logistic regression via conditioning. The interaction terms can be interpreted as log-odds ratios [17]. Specifically, the quantities ($\boldsymbol{\gamma_{CX}}$, $\boldsymbol{\gamma_{C\oplus C}}$) are conditional log-odds ratios for pairwise associations among individual components of $(X, \boldsymbol{C})$ (see Section 2.1 for details and an illustration). Consequently, Section B of Table III tells us about how strongly the confounders are correlated with one another. For example, the log-odds ratio for the association between CVD and COPD is 2.1 with standard error 0.7. Elements denoted 'NA' indicate terms that were dropped from the model because of sparsity to obtain a valid maximum likelihood estimator. Section C of Table III describes the prevalences of the disease variables.

Table III suggests that the disease indicator variables are confounders for the effect of X on Y and, furthermore, that they are correlated with one another. Most of the variables show associations with X and Y because the point estimates are positive. Furthermore, previous clinical evidence indicates that they are predictors of mortality in heart failure patients and that they influence prescribing of cardiovascular therapies [6, 18]. Therefore, they induce confounding. But the disease variables are also correlated with one another. In Section B of Table III, most of the log-odds ratios are greater than zero. Figure 1 plots the log-odds ratios. The sample mean is equal to 1.36, which gives an average odds ratio of $\exp(1.36) = 3.90$, and the standard deviation is 0.94. This suggests that patients who have one disease are also likely to have other diseases. In other words, the confounders are correlated with one another.

The missing confounder U is a binary indicator of the severity of heart failure. In formulating judgements about U, it is possible that U is correlated with $\boldsymbol{C}$. Vasan *et al.* [5] studied ejection fraction in heart failure patients and showed that patients with low ejection fraction were more likely to have diabetes, hypertension, high blood pressure and other chronic illnesses. This suggests that adjustment for $\boldsymbol{C}$ in the Naive analysis column of Table I may also control for some of the confounding from U because X and $\boldsymbol{C}$ are correlated with one another. Therefore, if we do a sensitivity analysis assuming that $\boldsymbol{\gamma_{CU}} = \boldsymbol{0}$ (i.e. assuming that $U \perp\!\!\!\perp \boldsymbol{C} | X$), then this may exaggerate the bias from U. Thus, we are faced with a conundrum. On the one hand, it seems unrealistic to assume that $\boldsymbol{\gamma_{CU}} = \boldsymbol{0}$ in the sensitivity analysis. On the other hand, it is not clear how to formulate a prior for $\boldsymbol{\gamma_{CU}}$ because it is a p-dimensional vector and because there is only limited information available about U.

1.4. Plan of the paper

One way to elicit priors for the bias parameters is to assume that the confounding effects of measured and unmeasured confounders are exchangeable in a Bayesian analysis, that is, to assume that the confounding induced by U is similar in magnitude to the confounding induced by $\boldsymbol{C}$. The assumption of exchangeability is a strong one; however, it is has been used previously in epidemiology to form qualitative judgements about unmeasured confounding. For example, in a 2000 review paper on confounding by indication, Joffe [19] wrote that '... one can learn about unmeasured confounders and confounding from measured factors. The argument is sometimes advanced that if adjustment for known covariates fails to change the measure of effect, there must be little residual confounding.... When control for measured factors reveals confounding, it is then more likely that there is residual confounding'. This logic rests on the assumption that the measured and unmeasured confounders are similar. If the investigator collects enough covariate information on the study participants, then this can be used to characterise the bias that would be produced from a confounder that is missing. Stampfer and Colditz [20] give a high profile example of this reasoning, concerning the potential confounding from socio-economic status to explain the observation that women taking postmenopausal estrogens are at decreased risk of coronary heart disease.

McCandless *et al.* [3] and Gustafson *et al.* [13] describe Bayesian methods that assume exchangeability in the confounding effects of measured and unmeasured confounders. In the paper of McCandless *et al.* [3], the authors analyse the beta blocker data, but they ignore the bias parameter $\boldsymbol{\gamma_{CU}}$ altogether because of the difficulties of prior specification. Gustafson *et al.* [13] consider the specific case where all of the measured covariates are continuous and are assumed to have a multivariate Gaussian distribution. However, the method of Gustafson *et al.* [13] relies heavily on the Gaussian framework and cannot be extended to binary covariates, as is the case with the beta blocker data. A different approach due to Greenland [10] uses random effects to characterise heterogeneity in the conditional association between U and a categorical covariate, given X. The variance of the random effects is specified by the investigation based on prior considerations. However, this does not permit data-driven learning about the correlations between measured and unmeasured confounders.

In this article, we propose a new formulation that models exchangeability among measured and unmeasured confounders. Our method builds on previous work [2, 3] because it is able to handle the case of binary covariates as in the beta blocker data. We extend the models in [2,3] by using a log-linear model for the joint distribution of $(X, \boldsymbol{C}, U)$ with pairwise interactions. Hierarchical priors borrow information from $\boldsymbol{C}$ to learn about bias from U. The method has the appealing property that conditioning on $(X, \boldsymbol{C})$ yields a logistic model for unmeasured confounding that is identical to that of McCandless *et al.* [2, 3] and Lin *et al.* [9]. Section 2 describes the method including the model, prior distributions and posterior computation by using MCMC. In Section 3, we apply the method to the beta blocker data. A key objective of this article was to investigate the impact of the prevailing approach to sensitivity analysis that assumes zero correlation between measured and unmeasured confounders. We study the results when using degenerate zero mass priors that force $\boldsymbol{\gamma_{CU}} = \mathbf{0}$. Following the logic of Schneeweiss [1] and Fewell *et al.* [16], we illustrate that if U and $\boldsymbol{C}$ are strongly correlated then confounding from U tends to diminish. In the beta blocker data, setting $\boldsymbol{\gamma_{CU}} = \mathbf{0}$ for convenience gives conclusions that are too pessimistic. The exchangeability assumption is a strong one, and we discuss the merits and trade-offs of our approach in Section 4.

2. Bayesian adjustment for unmeasured confounding

2.1. Model

We model the joint probability density $P(Y, X, \boldsymbol{C}, U|\boldsymbol{Z}) = P(Y|X, \boldsymbol{C}, \boldsymbol{Z}, U)P(X, \boldsymbol{C}, U)$ as

$$\text{logit}[Pr(Y = 1|X, \boldsymbol{C}, \boldsymbol{Z}, U)] = \beta_0 + \beta_{XY}X + \boldsymbol{\beta_{CY}}^T\boldsymbol{C} + \boldsymbol{\beta_{ZY}}^T\boldsymbol{Z} + \beta_{UY}U \tag{4}$$

$$\begin{aligned} P(X, \boldsymbol{C}, U) = & \frac{1}{Q(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})} \\ & \times \exp\{\gamma_X X + \boldsymbol{\gamma_C}^T\boldsymbol{C} + \gamma_U U \\ & \quad + \gamma_{XU}XU + \boldsymbol{\gamma_{CU}}^T\boldsymbol{C}U + \boldsymbol{\gamma_{CX}}^T\boldsymbol{C}X + \boldsymbol{\gamma_{C\oplus C}}^T(\boldsymbol{C}\oplus\boldsymbol{C})\}. \end{aligned} \tag{5}$$

Equation (4) is identical to Equation (1) and models the log-odds of the outcome as a function of $(X, \boldsymbol{C}, \boldsymbol{Z}, U)$. Equation (5) is an extension of Equation (3) to incorporate U into the joint distribution of $(X, \boldsymbol{C})$, and they are equivalent if the bias parameters $(\gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}})$ are equal to zero. The denominator $Q(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$ is the constant of normalisation.

As discussed in Section 3.1, there is a well-known connection between logistic and log-linear models through conditioning [17]. The interaction terms γ_{XU} and $\boldsymbol{\gamma_{CU}}$ are conditional log-odds ratios for the association between $(X, \boldsymbol{C})$ and U. If we take $P(X, \boldsymbol{C}, U)$ from Equation (5) and condition on $(X, \boldsymbol{C})$, then $P(U|X, \boldsymbol{C})$ obeys Equation (2). We have

$$\begin{aligned} \text{logit}[P(U = 1|X, \boldsymbol{C})] &= \log\left[\frac{P(U = 1|X, \boldsymbol{C})}{P(U = 0|X, \boldsymbol{C})}\right] \\ &= \log\left[\frac{P(U = 1, X, \boldsymbol{C})}{P(U = 0, X, \boldsymbol{C})}\right] \end{aligned}$$

$$
\begin{aligned}
&= \log\left[\exp\left\{\gamma_X X + \boldsymbol{\gamma_C}^T \boldsymbol{C} + \gamma_U + \right.\right. \\
&\qquad \gamma_{XU} X + \boldsymbol{\gamma_{CU}}^T \boldsymbol{C} + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C} X + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C} \oplus \boldsymbol{C})\big\} / \\
&\qquad \left.\left. \exp\left\{\gamma_X X + \boldsymbol{\gamma_C}^T \boldsymbol{C} + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C} X + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C} \oplus \boldsymbol{C})\right\}\right] . \\
&= \gamma_U + \gamma_{XU} X + \boldsymbol{\gamma_{CU}}^T \boldsymbol{C} .
\end{aligned}
$$

This gives an equivalence between our proposed model and previously proposed models for unmeasured confounding given by McCandless *et al.* [2, 3] and Lin *et al.* [9].

2.2. *Prior distributions*

Suppose that $\theta_1, \theta_2, \ldots, \theta_J$ are a collection of J unknown parameters. In Bayesian analysis, we say that $\theta_1, \theta_2, \ldots, \theta_J$ are *exchangeable* in their joint distribution if $P(\theta_1, \theta_2, \ldots, \theta_J)$ is invariant to the permutation of the indices $\{1, \ldots, J\}$ [21]. An exchangeable prior distribution is plausible if, on the basis of available information, we are unable to distinguish one parameter from another. Gelman *et al.* [21] wrote that 'In practice, ignorance implies exchangeability. Generally, the less we know about a problem, the more confidently we can make claims about of exchangeability'.

Now suppose that $\theta_1, \theta_2, \ldots, \theta_J$ are exchangeable. Then, following standard principles of Bayesian analysis, we can apply de Finetti's theorem, which states that in the limit as $J \to \infty$, then under certain regularity conditions, any exchangeable distribution for $\theta_1, \theta_2, \ldots, \theta_J$ can be expressed as an independent and identically distributed mixture of random variables conditional on some latent variable [21]. In other words, $\theta_1, \theta_2, \ldots, \theta_J$ can be modelled as a random sample from a distribution. Technically, the theorem does not apply for finite J (see [22] for further discussion of exchangeability).

Building on the discussion of unmeasured confounding in Section 1, we model the confounding effects of U and $\boldsymbol{C}$ as exchangeable (for further discussion of the merits and trade-offs of assuming exchangeability, see Section 4). For the outcome model, we assign a diffuse normal prior to β_0 and $\boldsymbol{\beta_{ZY}}$ with mean zero and variance 10^3. We assign

$$
\begin{aligned}
\beta_{XY}, \boldsymbol{\beta_{CY}}, \beta_{UY} &\overset{\text{i.i.d.}}{\sim} \mathrm{N}(0, \sigma_\beta^2) \\
\sigma_\beta^2 &\sim \text{Inv-}\chi^2\left(10^{-3}, 10^{-3}\right) .
\end{aligned} \tag{6}
$$

The left-hand side of Equation (6) refers to the individual components of β_{XY}, $\boldsymbol{\beta_{CY}}$ and β_{UY}, and Inv-$\chi^2\{.\}$ is an inverse χ^2 distribution with degrees of freedom 10^{-3} and scale parameter 10^{-3}. This choice of hyperparameters gives priors that are proper but uninformative. Equation (6) models the conditional associations between $(X, \boldsymbol{C}, U)$ and Y as exchangeable. The variance parameter σ_β^2 shares information between $\boldsymbol{C}$ and U. If σ_β^2 is small, then this shrinks the posterior for the bias parameter β_{UY} towards zero to reflect that there is less unmeasured confounding.

Eliciting a prior for γ_{XU} and $\boldsymbol{\gamma_{CU}}$ is more challenging because the parameters describe the manner in which U is distributed within levels of X and $\boldsymbol{C}$. We assign

$$
\begin{aligned}
\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}} &\overset{\text{i.i.d.}}{\sim} \mathrm{N}(\mu_\gamma, \sigma_\gamma^2) \\
\mu_\gamma &\sim N(0, 10^3) \\
\sigma_\gamma^2 &\sim \text{Inv-}\chi^2\left(10^{-3}, 10^{-3}\right) ,
\end{aligned} \tag{7}
$$

where the left-hand side of Equation (7) refers to the individual components of γ_{XU}, $\boldsymbol{\gamma_{CU}}$, $\boldsymbol{\gamma_{CX}}$ and $\boldsymbol{\gamma_{C\oplus C}}$. Equation (7) models the pairwise conditional log-odds ratios among $(X, \boldsymbol{C}, U)$ as exchangeable. The parameter μ_γ is the mean log-odds ratio, and σ_γ^2 is the variance.

Finally, we assign priors to the remaining model parameters γ_X, $\boldsymbol{\gamma_C}$ and γ_U, which are the main effects in the log-linear model of Equation (5). The quantity $\exp(\gamma_U)/(1 + \exp(\gamma_U))$ is the prevalence of $U = 1$ when $X = 0$ and $C_1, \ldots, C_p = 0$. We model the quantities γ_X, $\boldsymbol{\gamma_C}$ and γ_U as exchangeable and assign

$$
\begin{aligned}
\gamma_X, \boldsymbol{\gamma_C}, \gamma_U &\overset{\text{i.i.d.}}{\sim} \mathrm{N}(\mu_0, \sigma_0^2) \\
\mu_0 &\sim N(0, 10^3) \\
\sigma_0^2 &\sim \text{Inv-}\chi^2\left(10^{-3}, 10^{-3}\right) .
\end{aligned} \tag{8}
$$

Note that the Inv-χ^2 priors have the convenience of conjugacy; however, they are known to be inappropriately informative when the true variance is small or the number of observations is small. Alternatively, we could use a uniform prior on the standard deviation. We refer the reader to the study by Gelman *et al.* [21] for a discussion.

2.3. *Model fitting and computation*

Denote data as $data = \{(Y_i, X_i, \boldsymbol{C}_i, \boldsymbol{Z}_i); i = 1, \ldots, n\}$. If $U_1, \ldots, U_n$ were measured, then the likelihood function would be

$$\prod_{i=1}^{n} P(Y_i|X_i, \boldsymbol{C}_i, \boldsymbol{Z}_i, U_i) P(X_i, \boldsymbol{C}_i, U_i)$$
$$= \prod_{i=1}^{n} \left[\frac{\exp\{Y_i(\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i + \beta_{UY} U_i)\}}{1 + \exp\{\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i + \beta_{UY} U_i\}} \right.$$
$$\left. \times \frac{\exp\{\gamma_X X_i + \boldsymbol{\gamma_C}^T \boldsymbol{C}_i + \gamma_U U_i + \gamma_{XU} X_i U_i + \boldsymbol{\gamma_{CU}}^T \boldsymbol{C}_i U_i + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C}_i X_i + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C}_i \oplus \boldsymbol{C}_i)\}}{Q(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})} \right].$$

Because U is unmeasured, we obtain the likelihood for the observed data by integrating over the binary U. We obtain

$$L(\beta_0, \beta_{XY}, \boldsymbol{\beta_{CY}}, \boldsymbol{\beta_{ZY}}, \beta_{UY}, \gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$$
$$= \prod_{i=1}^{n} [P(Y_i|X_i, \boldsymbol{C}_i, \boldsymbol{Z}_i, U = 0) P(X_i, \boldsymbol{C}_i, U = 0) + P(Y_i|X_i, \boldsymbol{C}_i, \boldsymbol{Z}_i, U = 1) P(X_i, \boldsymbol{C}_i, U = 1)]$$
$$= \prod_{i=1}^{n} \left[\frac{\exp\{Y_i(\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i)\}}{1 + \exp\{\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i\}} \right.$$
$$\times \frac{\exp\{\gamma_X X_i + \boldsymbol{\gamma_C}^T \boldsymbol{C}_i + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C}_i X_i + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C}_i \oplus \boldsymbol{C}_i)\}}{Q(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})} \quad (9)$$
$$\times \frac{\exp\{Y_i(\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i) + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i + \beta_{UY})\}}{1 + \exp\{\beta_0 + \beta_{XY} X_i + \boldsymbol{\beta_{CY}}^T \boldsymbol{C}_i + \boldsymbol{\beta_{ZY}}^T \boldsymbol{Z}_i + \beta_{UY}\}}$$
$$\left. \times \frac{\exp\{\gamma_U + (\gamma_X + \gamma_{XU}) X_i + (\boldsymbol{\gamma_C} + \boldsymbol{\gamma_{CU}})^T \boldsymbol{C}_i + \boldsymbol{\gamma_{CX}}^T \boldsymbol{C}_i X_i + \boldsymbol{\gamma_{C\oplus C}}^T (\boldsymbol{C}_i \oplus \boldsymbol{C}_i)\}}{Q(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})} \right].$$

The posterior distribution is

$$P(\beta_0, \beta_{XY}, \boldsymbol{\beta_{CY}}, \boldsymbol{\beta_{ZY}}, \beta_{UY}, \gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}}, \sigma^2_\beta, \mu_\gamma, \sigma^2_\gamma, \mu_0, \sigma^2_0 | data) \quad (10)$$
$$\propto L(\beta_0, \beta_{XY}, \boldsymbol{\beta_{CY}}, \boldsymbol{\beta_{ZY}}, \beta_{UY}, \gamma_X, \boldsymbol{\gamma_C}, \gamma_U, \gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$$
$$\times P(\beta_0) P(\boldsymbol{\beta_{ZY}}) P(\beta_{XY}, \boldsymbol{\beta_{CY}}, \beta_{UY} | \sigma^2_\beta) P(\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}} | \mu_\gamma, \sigma^2_\gamma) P(\gamma_X, \boldsymbol{\gamma_C}, \gamma_U | \mu_0, \sigma^2_0) \quad (11)$$
$$\times P(\sigma^2_\beta) P(\mu_\gamma) P(\sigma^2_\gamma) P(\mu_0) P(\sigma^2_0) \quad (12)$$

where lines (11) and (12) refer to the prior distributions for model parameters.

We sample from the posterior distribution by using MCMC and the Metropolis Hastings algorithm. We update sequentially from the conditional densities

$$[\beta_0, \beta_{XY}, \boldsymbol{\beta_{CY}}, \boldsymbol{\beta_{ZY}}|.] \quad [\beta_{UY}|.] \quad [\gamma_x, \boldsymbol{\gamma_C}, \gamma_U|.] \quad [\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}}|.]$$

$$[\sigma^2_\beta|.] \quad [\mu_\gamma, \sigma^2_\gamma|.] \quad [\mu_0, \sigma^2_0|.],$$

where '[|.]' means conditional on the data and remaining model parameters. To update each of the sets of parameters $(\beta_0, \beta_{XY}, \boldsymbol{\beta_{CY}}, \boldsymbol{\beta_{ZY}})$, β_{UY}, $(\gamma_x, \boldsymbol{\gamma_C}, \gamma_U)$ and $(\gamma_{XU}, \boldsymbol{\gamma_{CU}}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$, we use a multivariate random walk with proposal distributions that are multivariate t-distributed with small degrees

of freedom and with scale matrix equal to the identity matrix multiplied by a tuning parameter that is set by trial MCMC runs. To update the hyperparameters, we note that the prior distributions for σ_β^2, $(\mu_\gamma, \sigma_\gamma^2)$ and (μ_0, σ_0^2) are conditionally conjugate, and we can update them by using the computational algorithm described by Gelman *et al.* [21].

However, a computational difficulty is that the model for unmeasured confounding is not identifiable, which means that different points in the parameter space give identical likelihood functions for the data. This can lead to slow MCMC convergence, particularly if the data set is large. In Section 3.1 and in the Supplementary Appendix,[‡] we discuss an alternative method for fast simulation from the posterior distribution. Computer code using the software R [23] is available from the author's website (see [2, 3, 10, 13] for a discussion of non-identifiable models for unmeasured confounding).

3. Analysis results for the beta blocker data

3.1. *Full Bayesian analysis*

We fit the model in Equations (4) and (5) to the beta blocker data and estimate the association between X and Y while adjusting for $(\boldsymbol{C}, \boldsymbol{Z})$ and exploring sensitivity to the unmeasured confounder U.

As discussed in Section 2.3, the model for unmeasured confounding is not identifiable, and this gives poor MCMC mixing. In principle, it is possible to get satisfactory convergence by using long MCMC chains. However, in the Supplementary Appendix, we describe an alternative procedure for sampling approximately from the posterior. Briefly, the idea is to proceed in two stages. First, we estimate the hyperparameters $(\sigma_\beta, \mu_\gamma, \sigma_\gamma, \mu_0, \sigma_0)$ by using maximum likelihood and plug them in place of true parameter values. This reduces the dimension of the parameter space and the overall computational burden. However, it also ignores uncertainty in the hyperparameter estimates (see Section 4 for a discussion). Second, we use the technique of *modularisation* proposed by Liu, Bayarri and Berger [24], which accelerates MCMC computation by sampling from approximations to the full conditional distributions for model parameters. It is equivalent to the computational technique of *cutting feedback*, which has been implemented in the BUGS software [25]. R code for MCMC sampling is available from the author's website.

For the hyperparameters, we obtain estimates $(\hat{\sigma}_\beta, \hat{\mu}_\gamma, \hat{\sigma}_\gamma, \hat{\mu}_0, \hat{\sigma}_0) = (0.62, 1.36, 0.94, -2.23, 1.04)$ (see the Supplementary Appendix for details on how they are calculated and for the interpretation). We run a single MCMC chain of length 100,000 after 10,000 burn-in iterations. We assess sampler convergence by using separate simulation runs with overdispersed starting values and the diagnostics tools included in the R package CODA [23].

We give the results in the second column of Table I. The column has the heading '$U \not\perp\!\!\!\perp \boldsymbol{C} \,|\, X$' to indicate that the components of $\boldsymbol{\gamma_{CU}}$ are modelled as exchangeable with the components of $(\gamma_{XC}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}})$ in Equation (7) and, therefore, that the analysis does not assume that $\boldsymbol{\gamma_{CU}} = \mathbf{0}$. The log-odds ratio for the beta blocker effect parameter β_{XY} is -0.41 with 95% credible interval $(-0.87, 0.06)$. This point estimate is slightly shrunk towards zero compared with the naive analysis because we have assumed an informative prior distribution $N(0, 0.62^2)$ on β_{XY}. The interval estimate is wider because the Bayesian analysis acknowledges uncertainty from unmeasured confounding. The prior distribution in Equation (6) assumes that the bias parameter β_{UY} has a prior mean zero and standard deviation $\hat{\sigma}_\beta = 0.62$. In other words, the sensitivity analysis assumes that U is associated with Y, given $(X, \boldsymbol{C}, \boldsymbol{Z})$, and this association may either increase or decrease the probability of Y. Because the prior is symmetric at zero, this means that the posterior mean for β_{XY} is similar to that of the naive analysis but that the interval estimate is wider. McCandless *et al.* [2, 3] report similar results.

3.2. *Assessing prior sensitivity for* $\boldsymbol{\gamma_{CU}}$

Our modelling framework gives the opportunity to study the role of $\boldsymbol{\gamma_{CU}}$ in sensitivity analysis for unmeasured confounding. One issue is assessing the impact of the usual assumption that $\boldsymbol{\gamma_{CU}} = \mathbf{0}$. Recall from Section 1.2 that most sensitivity analysis techniques assume that measured and unmeasured confounders are uncorrelated (i.e. $U \perp\!\!\!\perp \boldsymbol{C} \,|\, X$) to reduce the burden of prior elicitation.

[‡] *Supplementary appendix may be found in the online version of this article.*

To study the effect of this assumption, we redo the Bayesian analysis in exactly the same way as Section 3.1 but change the prior in Equation (7) to be

$$\begin{aligned} \boldsymbol{\gamma_{CU}} &= \mathbf{0} \\ \gamma_{XU}, \boldsymbol{\gamma_{CX}}, \boldsymbol{\gamma_{C\oplus C}} &\overset{\text{i.i.d.}}{\sim} \text{N}\{\hat{\mu}_\gamma, \hat{\sigma}^2_\gamma\}, \end{aligned} \tag{13}$$

where $\hat{\mu}_\gamma = 1.36$ and $\hat{\sigma}_\gamma = 0.94$. This sets each component $\boldsymbol{\gamma_{CU}}$ equal to zero and guarantees that $U \perp\!\!\!\perp \boldsymbol{C} \,|\, X$. However, it permits the other three bias parameters β_{UY}, γ_{UX} and γ_U to be non-zero. Thus, using the prior in Equation (13) instead of (7) allows U to be an unmeasured confounder for the effect of X and Y, despite the fact that $U \perp\!\!\!\perp \boldsymbol{C} \,|\, X$.

We present the results in the third column of Table I under the heading '$U \perp\!\!\!\perp \boldsymbol{C} \,|\, X$'. As intuition suggests, assuming that $\boldsymbol{\gamma_{CU}} = \mathbf{0}$ *increases* the posterior uncertainty about unmeasured confounding in the beta blocker data. The credible interval for the beta blocker effect is nearly 10% wider than in the middle column where $U \not\perp\!\!\!\perp \boldsymbol{C} \,|\, X$. We have -0.40 $(-0.90, 0.11)$ versus -0.41 $(-0.87, 0.06)$. Both analyses acknowledge uncertainty from unmeasured confounding, but only the exchangeable analysis allows the possibility that U and $\boldsymbol{C}$ are correlated. The magnitude of this correlation is driven by the hyperparameter estimates $\hat{\mu}_\gamma = 1.36$ and $\hat{\sigma}_\gamma = 0.94$, which in turn are estimated from the joint distribution of X and $\boldsymbol{C}$.

To illustrate the prior sensitivity more clearly, we repeat the analysis while toying with fixed values for $\boldsymbol{\gamma_{CU}}$. Figure 2 illustrates what happens to the posterior distribution of β_{XY} when we set $\boldsymbol{\gamma_{CU}}$ equal to $\mathbf{0}, \mathbf{1}, \ldots, \mathbf{5}$, which denote vectors of length p. For example, $\mathbf{0} = (0, 0, \ldots, 0)$ and $\mathbf{1} = (1, 1, \ldots, 1)$. When $\boldsymbol{\gamma_{CU}} = \mathbf{5}$, this means that the odds ratios for the conditional association between U and each of component $\boldsymbol{C}$ is equal to $\exp(5) = 148$, which corresponds to roughly perfect correlation between $\boldsymbol{C}$ and U. In Figure 2, we calculate each of the interval estimates for β_{XY} by doing a Bayesian sensitivity analysis with $\boldsymbol{\gamma_{CU}}$ locked at either $\mathbf{0}, \mathbf{1}, \ldots, \mathbf{5}$. The grey-shaded region indicates the width and the positioning of the naive interval estimate for β_{XY}, which is $(-0.80, -0.07)$.

The key observation is that when $\boldsymbol{\gamma_{CU}}$ is large, the interval estimates collapse towards the shaded region, and we obtain inferences that are essentially identical to assuming that there is no unmeasured confounding. Indeed, if U and $\boldsymbol{C}$ are strongly correlated, then regression adjustment for $\boldsymbol{C}$ eliminates confounding from U, despite the fact that U is unmeasured. This occurs even though the prior distributions on the bias parameters β_{UY}, γ_{UX} and γ_U are not degenerate at zero. In other words, U induces essentially no bias upon adjustment for $\boldsymbol{C}$, provided that $\boldsymbol{C}$ is sufficiently correlated with U.

In Figure 2, the interval estimates are shifted slightly towards zero compared with the shaded region. This is because of the informative prior on β_{XY} in Equation (6). The prior has mean zero and variance $\hat{\sigma}_\beta = 0.62$, which tends to shrink point estimates for β_{XY} a little towards the zero. In contrast, we compute the naive interval estimate in Table I by maximum likelihood, which effectively presumes a flat prior on β_{XY}.

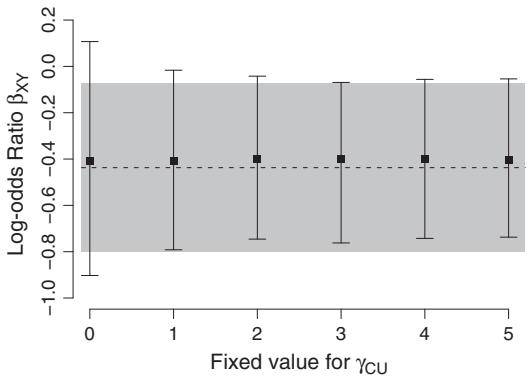

Figure 2. 95% credible intervals for the beta blocker effect β_{XY} when we repeat the Bayesian sensitivity analysis with $\boldsymbol{\gamma_{CU}}$ held fixed at $\mathbf{0}$ versus $\mathbf{1}$ versus $\mathbf{2}, \ldots,$versus $\mathbf{5}$. The grey-shaded region indicates the width and the position of the naive interval estimate for β_{XY}, which is $(-0.80, -0.07)$ and taken from Table I. Notice that the intervals collapse around the grey-shaded region as we move to the right (i.e. as we assume that U is more correlated with $\boldsymbol{C}$).

3.3. Assessing prior sensitivity when incorporating continuous confounders

A limitation of our method is that it assumes that the components of $\boldsymbol{C}$ are dichotomous. In the beta blocker data, the assumption that age and duration of hospital stay are unrelated to severity of heart failure seems unlikely. It would be useful to incorporate them into the analysis so that they inform assumptions about the magnitude of unmeasured confounding. However, we cannot use log-linear models for continuous variables. Log-linear models are only applicable for contingency tables [17].

A pragmatic solution is to dichotomise age and hospital stay at the median when entering them into Equation (5). As an illustration, we fix $\boldsymbol{C}$ to include age and hospital stay *in addition* to the $p = 9$ disease indicator variables listed in Table I. We denote the remaining $21 - 9 - 2 = 10$ covariates by $\boldsymbol{Z}$. We dichotomise age at 77 years and hospital length of stay at 5.0 days. This yields two new indicator variables that take the value 1 if age is greater than 77 years or zero 0 otherwise (versus 1/0 for length of stay that is greater/less than 5.0 days). We include the dichotomised variables in $\boldsymbol{C}$ in Equation (5).

We keep age and hospital length of stay as continuous variables in Equation (4). Following Gelman [21], we rescale both variables to have unit 1 variance prior to analysis. This makes the prior distribution in Equation (6) more realistic. After rescaling, the regression coefficients correspond to log-odds ratios for an increase of roughly one population standard deviation in age or hospital stay. We judge that these coefficients are roughly exchangeable with the coefficients for the dichotomous disease variables.

Dichotomising continuous variables is not recommended because of loss of power and possible residual confounding [26]. However, because we retain continuous version of age and length of stay in the outcome model, this means that residual confounding is less an issue. Nonetheless, Equation (5) will not capture the full richness of dependence in the joint distribution of $(X, \boldsymbol{C})$. This could affect the interpretation and plausibility of the hyperparameter estimates $\hat{\mu}_\gamma, \hat{\sigma}_\gamma, \hat{\mu}_0, \hat{\sigma}_0$. Therefore, we frame this investigation as a sensitivity analysis to study the effect of incorporating continuous variables.

After including age and length of stay, the new hyperparameter estimates are $(\hat{\sigma}_\beta, \hat{\mu}_\gamma, \hat{\sigma}_\gamma, \hat{\mu}_0, \hat{\sigma}_0) = (0.62, 1.06, 1.29, -2.31, 0.95)$ (see the Supplementary Appendix for details on how the estimates are calculated). We repeat the full Bayesian analysis and estimate the exposure effect β_{XY} as -0.40 with 95% credible interval $(-0.86, 0.02)$. The interval estimate is similar, albeit slightly more narrow, than the intervals reported in Table I. For this specific analysis, incorporating continuous confounders does not greatly affect the conclusions. A challenge with interpreting the results (see Section 4 for further discussion) is that we ignore uncertainty in the hyperparameter estimates to speed MCMC computation, and the results may be slightly conservative.

4. Discussion

Recent years have witnessed new innovation in Bayesian techniques to adjust for unmeasured confounding in observational studies. A challenge is that the user is often required to elicit prior distributions for high-dimensional parameters that model competing bias scenarios. This can render the methods unwieldy. In this paper, we propose a novel methodology for settings where the confounding effects of measured and unmeasured variables can be viewed as exchangeable within a Bayesian framework. Exchangeability captures the intuitive idea put forth by Joffe [19] that confounding from measured variables may be informative about unmeasured variables. Our method reduces the burden of prior elicitation in sensitivity analysis because it assigns priors to bias parameters without requiring that the analyst encode assumptions about each parameter individually. It builds on previous work [3, 10, 13] because it is able to accommodate binary covariates in settings where one cannot reasonably assume that $U \perp\!\!\!\perp \boldsymbol{C} \,|\, X$.

One can argue that the assumption of exchangeability is not plausible in the beta blocker data. Exchangeability means that there is no reason to believe that the confounding from heart failure severity is any different from the set of assembled disease indicators. However, a recent study [27] estimated 1-year mortality risks in 70-year-old men as 9%, 13%, 21% and 31% for New York Heart Association classes of heart failure 1, 2, 3 and 4, respectively. This suggests that heart failure is a far greater risk factor for mortality than virtually all of the risk factors considered, with the exception of cancer. A further problem is that there are several *additional* confounders that were omitted from our study including smoking, physical activity and blood pressure. If our sensitivity analyses must, for operational reasons, be restricted to a single, dichotomous, unmeasured confounder, then it seems prudent to assume that it has a very strong relationship with the outcome and also that it has a weak relationship with measured confounders.

More generally, exchangeability is only contingent on making judgements about the labelling of the indices of the parameters in the prior distribution (Section 2.2). An exchangeable prior is plausible for a collection of parameters if, on the basis of available information, we are unable to distinguish one parameter for another. We emphasise that care is needed in choosing which variables to include in $\boldsymbol{C}$. If the investigator incorporates weak or marginal confounders, then the hierarchical priors will shrink the hyperparameters towards zero, giving the impression that there is little unmeasured confounding. An alternative extension would be to relax the normal assumption in Equations (6)–(8) to a heavier tailed distribution, such as a t-distribution. This would give less shrinkage and would allow greater variability of the confounder effect.

A limitation of this investigation is the hypothetical dichotomy of U. The unmeasured confounder is severity of heart failure in terms of reduced cardiac output and limitations in physical activities. Severity of heart failure is measured over a range from low to high. Thus, a binary U will not capture the full spectrum of confounding. One possible extension would be to model U as an ordinal variable by using a series of indicators. However, we caution that this adds complexity because it requires that the analyst speculate about how U is distributed over levels of $(X, \boldsymbol{C})$. Additionally, we note that clinical researchers sometimes dichotomise severity by using a threshold of 50% to identify low ejection fraction (see [5] for a discussion of classification of patients as having normal versus low ejection fraction).

An additional limitation of our analysis is that there are other biases at play. The assessment of beta blocker use in the month after hospital discharge can introduce immortal time bias [28]. Ideally, this could be remedied by incorporating time-dependent covariates, although this is beyond the scope of this paper. Additionally, as discussed previously, there may be additional unmeasured confounders apart from heart failure severity. Nonetheless, accounting for several unmeasured confounders is difficult without additional data sources. This is because such a sensitivity analysis requires the analyst to guess the joint confounding effects of several variables acting in unison [1].

A final limitation of our method is that we use MCMC computation based on point estimates for the hyperparameters (see Section 3.1 and the Supplementary Appendix). In our experience, this substantially reduces computational time. However, it also ignores uncertainty in the hyperparameter estimates and may give interval estimates that are too narrow if the number of measured confounders p is small.

References

1. Schneeweiss S. Sensitivity analysis and external adjustment for unmeasured confounders in epidemiologic database studies of therapeutics. *Pharmacoepidemiology and Drug Safety* 2006; **15**:291–303.
2. McCandless LC, Gustafson P, Levy AR. Bayesian sensitivity analysis for unmeasured confounding in observational studies. *Statistics in Medicine* 2007; **26**:2331–2347.
3. McCandless LC, Gustafson P, Levy AR. A sensitivity analysis using information about measured confounders yielded improved assessments of uncertainty from unmeasured confounding. *Journal of Clinical Epidemiology* 2008; **61**:247–255.
4. Foody JAM, Farrell MH, Krumholz HM. β-blocker therapy in heart failure: scientific review. *Journal of the American Medical Association* 2002; **278**:883–889.
5. Vasan RS, Larson MG, Benjamin EJ, Evans JC, Reiss CK, Levy D. Congestive heart failure in subjects with normal versus reduced left ventricular ejection fraction: prevalence and mortality in a population-based cohort. *Journal of the American College of Cardiology* 1999; **33**:1948–1955.
6. Glynn RJ, Knight EL, Levin R, Avorn J. Paradoxical relations of drug treatment with mortality in older persons. *Epidemiology* 2001; **12**:682–689.
7. Rosenbaum PR, Rubin DB. Assessing sensitivity to an unobserved binary covariate in anobservational study with binary outcome. *Journal of the Royal Statistical Society Series B* 1983; **45**:212–218.
8. Yanagawa T. Case–control studies: assessing the effect of a confounding factor. *Biometrika* 1984; **71**:191–194.
9. Lin DY, Psaty BM, Kronmal RA. Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics* 1998; **54**:948–963.
10. Greenland S. The impact of prior distributions for uncontrolled confounding and response bias. *Journal of the American Statistical Association* 2003; **98**:47–54.
11. Jackson C, Best NG, Richardson S. Bayesian graphical models for regression on multiple datasets with different variables. *Biostatistics* 2009; **10**:335–351.
12. Molitor N, Best NB, Jackson C, Richardson S. Using Bayesian graphical models to model biases in observational studies and to combine multiple data sources: application to low birthweight and water disinfection by-products. *Journal of the Royal Statistical Society Series A* 2009; **172**:615–637.
13. Gustafson P, McCandless LC, Levy AR, Richardson SR. Simplified Bayesian sensitivity analysis for mismeasured and unobserved confounders. *Biometrics* 2010; **66**:1129–1137.
14. Hernán MA, Robins JM. Letter to the editor. *Biometrics* 1999; **55**:1316–1317.
15. VanderWeele TJ. Sensitivity analysis: distributional assumptions and confounding assumptions. *Biometrics* 2009; **64**:645–649.

16. Fewell Z, Smith D, Sterne JAC. The impact of residual and unmeasured confounding in epidemiologic studies: a simulation study. *American Journal of Epidemiology* 2007; **166**:646–656.
17. Agresti A. *Categorical Data Analysis*, 2nd Edition. Wiley: New Jersey, 2002.
18. Polanczyk CA, Rohde LE, Philbin EA, Di Salvo TG. A new casemix adjustment index for hospital mortality among patients with congestive heart failure. *Medical Care* 1998; **36**:1489–1499.
19. Joffe MM. Confounding by indication: the case of calcium channel blocker. *Pharmacoepidemiology and Drug Safety* 2000; **9**:37–41.
20. Stampfer MJ, Colditz GA. Estrogen replacement therapy and coronary heart disease: a quantitative assessment of the epidemiological evidence. *International Journal of Epidemiology* 2004; **33**:445–453.
21. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis*, 2nd Edition. Chapman Hall/CRC: New York, 2004.
22. Bernardo JM, Smith AFM. *Bayesian Theory*. Wiley: New York, 1994.
23. R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing: Vienna, 2004. ISBN 3-900051-00-3. URL http://www.R-project.org.
24. Liu F, Bayarri MJ, Berger J. Modularization in Bayesian analysis, with emphasis on analysis of computer models. *Bayesian Analysis* 2009; **4**:119–150.
25. Lunn D, Spiegelhalter D, Thomas A, Best N, Lunn D. The BUGS project: evolution, critique and future directions. *Statistics in Medicine* 2009; **28**:3049–3067.
26. Royston P, Altman DG, Sauerbrei W. Dichotomizing continuous predictors in multiple regression: a bad idea. *Statistics in Medicine* 2006; **25**:127–141.
27. Levy WC, Mozaffarian D, Linker DT, Sutradhar SC, Anker S, Cropp A B, et al. The Seattle Heart Failure Model: prediction of survival in heart failure. *Circulation* 2006; **113**:1424–1433.
28. Suissa S. Immortal time bias in pharmacoepidemiology. *American journal of epidemiology* 2008; **167**:492–499.