Timothy Shin Heng Mak*, Nicky Best and Lesley Rushton

Robust Bayesian Sensitivity Analysis for Case–Control Studies with Uncertain Exposure Misclassification Probabilities

Abstract: Exposure misclassification in case–control studies leads to bias in odds ratio estimates. There has been considerable interest recently to account for misclassification in estimation so as to adjust for bias as well as more accurately quantify uncertainty. These methods require users to elicit suitable values or prior distributions for the misclassification probabilities. In the event where exposure misclassification is highly uncertain, these methods are of limited use, because the resulting posterior uncertainty intervals tend to be too wide to be informative. Posterior inference also becomes very dependent on the subjectively elicited prior distribution. In this paper, we propose an alternative “robust Bayesian” approach, where instead of eliciting prior distributions for the misclassification probabilities, a feasible region is given. The extrema of posterior inference within the region are sought using an inequality constrained optimization algorithm. This method enables sensitivity analyses to be conducted in a useful way as we do not need to restrict all of our unknown parameters to fixed values, but can instead consider ranges of values at a time.

Keywords: misclassification, robust Bayes, case–control study, Bayesian

DOI 10.1515/ijb-2013-0044

1 Introduction

Exposure misclassification is a common problem in case–control studies. When exposure misclassification is present, estimates of exposure prevalences and odds ratios are biased. Early references have pointed out that if the extent of misclassification is the same for both cases and controls and a dichotomous exposure is used, then bias is towards the null, and the p -value is not affected [4, 7], such that if a positive/negative relationship is demonstrated between exposure and disease, the evidence for the relationship being positive/negative is not affected by misclassification, under frequentist inference anyway. However, even if this is the case, estimates remain biased, and confidence intervals are inaccurate, and in Bayesian inference, the implication of bias towards the null is also not necessarily true [21]. Moreover, it has been argued that even small departures from the strict non-differential misclassification assumption does not warrant the “bias towards the null” implication [23], and so inference based on the misclassified exposure becomes even less reliable.

Assuming a binary (dichotomous) exposure, then true exposure prevalence (π) and misclassified exposure prevalence (p) are related through a simple formula:

$$p = \pi \text{sens} + (1 - \pi)(1 - \text{spec}) \quad (1)$$

*Corresponding author: Timothy Shin Heng Mak, Centre for Genomic Sciences, University of Hong Kong, 5 Sassoon Road, Pokfulam, Hong Kong, E-mail: tshmak@hku.hk

Nicky Best, Lesley Rushton, Department of Epidemiology and Biostatistics, Imperial College London, London, UK, E-mail: n.best@imperial.ac.uk, l.rushton@imperial.ac.uk

where *sens* is the *sensitivity* and *spec* the *specificity* of the exposure measure and is defined as:

$$\text{sens} = \frac{\left(\begin{array}{l} \text{Number of truly exposed people who} \\ \text{would have been categorized as exposed in the population} \end{array} \right)}{\text{Total number of truly exposed people in the population}}$$

$$\text{spec} = \frac{\left(\begin{array}{l} \text{Number of truly unexposed people who} \\ \text{would have been categorized as unexposed in the population} \end{array} \right)}{\text{Total number of truly unexposed people in the population}}$$

Because of this simple relationship, it has been proposed in the literature [1, 7, 28, 30] that the misclassified exposure prevalence estimate can be corrected in order to estimate the true relative risk/odds ratio, and variance formulae for the adjusted estimate can be obtained using the delta method or maximum likelihood for large samples [14]. In the case of a binary exposure, the adjustment formula is simply the inverse of (1):

$$\hat{\pi} = \frac{\hat{p} - 1 + \text{spec}}{\text{sens} + \text{spec} - 1} \quad (2)$$

where $\hat{p}$ is an estimate of p and $\hat{\pi}$ the *adjusted* estimate of π . Adjustment formulae such as (2), however, assume that the misclassification probabilities *sens* and *spec* are known. This, however, is almost never the case. Even in situations where these probabilities can be estimated from a validation study, estimates are subject to sampling error. In many situations, validation studies are not available, and estimates of *sens* and *spec* will have to depend entirely on intelligent guesses, if the adjustment formulae were to be used at all. Perhaps for this reason, even though exposure misclassification is widespread in epidemiology, adjustment for bias due to misclassification appears to be rarely applied.

Another problem with the adjustment formulae (such as (2)) is that sometimes the adjusted estimate for π (i.e. $\hat{\pi}$) is undefined. For example, if a rare exposure is involved, with prevalence of exposure $p = 0.01$, then so long as *sens* > 0.01, any value of *spec* < 0.99 leads to an undefined estimate for π , since according to (2), $\hat{\pi}$ is either negative or greater than 1. This happens because although *sens* and *spec* are fixed (they are *population* values), $\hat{p}$ is subject to sampling error and may lead to estimates that are incompatible with the population parameters *sens* and *spec*.

To a certain extent, the above problems can be overcome by adopting a Bayesian approach, as has been demonstrated by Gustafson, Le and Saskin [22], Gustafson [19], Chu, Gustafson and Le [6] and MacLehose and Gustafson [29]. In these authors' approach, prior distributions are given to π_1 , π_0 , *sens*₁, *sens*₀, *spec*₁, *spec*₀, where the subscript 1 refers to the case population and 0 the control population, and the posterior distribution of the odds ratio $= \frac{\pi_1(1-\pi_0)}{\pi_0(1-\pi_1)}$ is sought, through the model

$$Y_i \sim \text{Bin}(N_i, p_i), \quad i = 0, 1 \quad (3)$$

$$p_i = \pi_i \text{sens}_i + (1 - \pi_i)(1 - \text{spec}_i)$$

where Y_i denotes the number of exposed subjects and N_i the total number of cases or controls. Estimation of the posterior distribution can be derived through Markov chain Monte Carlo, which is easily implemented in software such as WinBUGS [27]. This approach does not make use of the estimate $\hat{p}$, and hence is not affected by possible incompatibility between $\hat{p}$ and *sens* and *spec*. There are, however, still a number of potential difficulties, such as:

1. It may not be immediately clear which parameterization we should adopt for the model. In the above-cited studies, prior distributions are given to the parameters π_1 , π_0 , *sens*₁, *sens*₀, *spec*₁, *spec*₀. On the other hand, Chu, Wang, Cole and Greenland [5] assigned prior distributions to π_0 and (logit π_1 - logit π_0) (instead of π_0 and π_1). Still, other vastly different parameterizations could have also been used (e.g. [8, 16]). The choice of parameterization is especially unclear when there are no suitable estimates of *sens*₁, *sens*₀, *spec*₁, *spec*₀ from validation studies.

2. If the extent of misclassification is great, the posterior distribution of the odds ratio becomes especially sensitive to the choice of prior and parameterization (to be demonstrated in Section 2). Moreover, accurate elicitation of prior distribution is made more difficult by the fact that it is often reasonable that prior distributions of the parameters are correlated [5, 6, 9], as elicitation of correlation is difficult [11].
3. Because inference from these analyses can depend so crucially on subjectively elicited probability distributions for the misclassification parameters, it is unsure what conclusions can be drawn by a reader who does not agree with the prior distributions used.

Because of the problems listed above, the goal of this paper is to present an alternative approach to analysing case–control studies subject to exposure misclassification, which is a type of robust Bayes analysis. Robust Bayes inference [2, 3] was originally introduced to examine the robustness of Bayesian estimates to departure from prior distributional assumptions. One type of robust Bayes analysis seeks the maximum and minimum possible inference from a class of prior distributions [12]. This type of analysis overcomes an important limitation in subjective Bayesian analysis, namely that it is rarely possible to specify a *unique* prior distribution for a particular analysis. Although philosophically appealing, this type of analysis has not been widely applied in practice, probably because of computational difficulties. This paper aims to demonstrate its use in solving an epidemiological problem, as well as to show that this type of analysis is not computationally infeasible.

In the rest of this paper, we illustrate some of the deficiencies of the Bayesian approach in Section 2, with reference to the example of a case–control study of childhood leukaemia and electromagnetic fields (EMF) exposure. In Section 3, the approach of this paper is introduced. Section 4 and 5 give further extension of the method. Section 6 discusses the use of the method in sensitivity analyses. Concluding remarks are given in Section 7.

2 Deficiencies of the Bayesian approach in accounting for bias due to exposure misclassification in case–control studies

In this section, we consider how Bayesian inference in the presence of exposure misclassification can be very sensitive to the prior distribution used. First, consider the case–control study by Linet, Hatch, Kleinerman, Robison, Kaune, Friedman, Severson, Haines, Hartsock, Niwa, Wacholder and Tarone [26] on the risk of childhood leukaemia and exposure to high levels of EMF. EMF exposure was assessed by 24 h bedroom measurements at a convenient time (which can be up to a few years) after the diagnosis of leukaemia, where possible, and spot measurements around the residence where this was not possible. Controls were matched to the cases by sex and age and recruited by random-digit dialling. Out of 624 cases, 45 had EMF measurements $> 0.3 \mu\text{T}$ (micro-Tesla). Among the 615 controls, 28 had measurements $> 0.3 \mu\text{T}$. This resulted in an odds ratio estimate of 1.63 comparing the risk of leukaemia in the $> 0.3 \mu\text{T}$ group versus the $< 0.3 \mu\text{T}$ group, with standard large sample 95% confidence interval (1.00, 2.65). Because the exposure measurement was performed in the residences of the children potentially several years after the etiologically relevant period, severe misclassification of exposure is very possible. Bayesian inference for this dataset can be conducted using model (3). In Table 1, we compare the posterior median and 95% credible interval of the (log) odds ratio under three non-differential misclassification scenarios (no misclassification, minor misclassification and severe misclassification), using four different prior distributions. Results are derived using WinBUGS 1.4 [27].

Priors 1 and 2 are both weakly informative. Prior 1, in particular, was used by Gustafson et al. [22]. It has been argued that prior distributions should be given independently to π_0 and θ , rather than π_0 and π_1 [15], where

$$\theta = \text{logit } \pi_1 - \text{logit } \pi_0$$

Table 1: Posterior median and 95% credible intervals for θ under different non-differential exposure misclassification scenarios and different prior distributions.

Log odds ratio (θ)	No misclassification	Minor misclassification	Severe misclassification
Odds ratio (exp θ)	<i>sens</i> = 1, <i>spec</i> = 1	<i>sens</i> = 0.7, <i>spec</i> = 0.98	<i>sens</i> = 0.3, <i>spec</i> = 0.95
Prior 1: $\pi_0 \sim U(0, 1), \pi_1 \sim U(0, 1)$	0.49 (0.004, 0.97) 1.61 (1.00, 2.63)	0.74 (0.01, 1.65) 2.09 (1.01, 5.18)	1.58 (−0.49, 4.81) 4.87 (0.61, 122.9)
Prior 2: $\pi_0 \sim U(0, 1), \theta \sim N(0, 100)$	0.45 (−0.03, 0.94) 1.57 (0.97, 2.56)	0.69 (−0.04, 1.61) 1.99 (0.96, 5.01)	0.45 (−18.16, 4.08) 1.57 (0.00, 58.99)
Prior 3: $\text{logit } \pi_0 \sim N(\text{logit } 0.046, 0.1), \theta \sim N(0, 1)$	0.46 (0.02, 0.89) 1.58 (1.02, 2.42)	0.58 (−0.03, 1.16) 1.78 (0.97, 3.19)	0.42 (−1.30, 1.48) 1.52 (0.27, 4.38)
Prior 4: $\text{logit } \pi_0 \sim N(\text{logit } 0.046, 0.1), \theta \sim N(0, 0.5)$	0.44 (0.01, 0.86) 1.55 (1.01, 2.36)	0.53 (−0.04, 1.11) 1.70 (0.96, 3.02)	0.32 (−0.95, 1.30) 1.38 (0.39, 3.68)

This is the case for priors 2, 3 and 4. Prior 2 has nearly flat prior distribution for both π_0 and θ , whereas in priors 3 and 4, the distribution for π_0 is made to have mean equal to the mean of $p_0 = Y_0/N_0 = 0.046$, the observed prevalence of exposure among the controls, on the log-odds scale. Prior 3 differs from prior 4 in that for prior 3, the variance of θ is 1, while for prior 4 the variance is 0.5. The former corresponds to a 95% credible interval for the odds ratio = (0.14, 7.1). The latter corresponds to a 95% credible interval of (0.25, 4.0).

It can be seen that posterior inference for θ , the log odds ratio, is similar using all four different prior distributions in the “no misclassification” and the “minor misclassification” scenario. However, the posterior inference under the “severe misclassification” scenario is very sensitive to the prior distributions given.

The above results show that in the presence of severe exposure misclassification, the prior distribution of θ has a large impact on the posterior distribution. This is the case even if the misclassification probabilities $\text{sens}_i, \text{spec}_i$ are known. An intuitive explanation of this may be seen by considering the variance of π_i in comparison to the variance of p_i (supposing fixed sens_i and spec_i):

$$\text{Var } \pi_i = \text{Var} \left(\frac{p_i - 1 + \text{spec}_i}{\text{sens}_i + \text{spec}_i - 1} \right) = \frac{\text{Var}(p_i)}{(\text{sens}_i + \text{spec}_i - 1)^2} \quad (\text{cf eq. (2)}) \quad (4)$$

We see that variance of π_i is the variance of p_i divided by $(\text{sens}_i + \text{spec}_i - 1)^2$ and is likely to be much greater than the variance of p_i if $\text{sens}_i + \text{spec}_i$ is close to 1. Because $\theta = \text{logit } \pi_1 - \text{logit } \pi_0$, the variance of θ is also likely to be much larger than the variance of the log-observed odds ratio $\text{logit } p_1 - \text{logit } p_0$, if $\text{sens}_i + \text{spec}_i$ is close to 1. Thus we see that the data are “diluted” and the posterior distribution becomes more influenced by the prior distribution.

As another example, we examined the sensitivity of posterior inference to specification of the prior distribution for the misclassification probabilities ($\text{sens}_1, \text{sens}_0, \text{spec}_1, \text{spec}_0$). For this example, we fixed the prior distribution for π_0 and θ at $\text{logit } \pi_0 \sim N(\text{logit } 0.046, 0.1), \theta \sim N(0, 1)$ (prior 3). In Table 2, priors A and B assume non-differential misclassification and consist of uniform priors given to sens and spec , though with different but largely overlapping ranges. Priors C and D assign bivariate Normal distributions to $(\text{logit } \text{sens}_1, \text{logit } \text{sens}_0)$ and $(\text{logit } \text{spec}_1, \text{logit } \text{spec}_0)$, with a correlation of 0.8. Means and variance of these bivariate Normal distributions are chosen such that their 2.5% and 97.5%-ile match the limits of priors A and B. As can be seen, although the variation is not as great as those seen in the last column of Table 1, there are still considerable differences between the results. Thus, when presented with a Bayesian analysis of a case-control study with potentially severe misclassification, one set of results is rarely enough. We generally need to examine various situations to see how sensitive the results are to prior distributions used.

In view of these limitations, one may ask whether there is an alternative way to quantify our uncertainties over such Bayesian analysis. In the next section, we present the robust Bayesian method of this paper as one such alternative.

Table 2: Posterior median and 95% credible intervals for θ using different prior distributions for $sens_1$, $sens_0$, $spec_1$, $spec_0$. Prior distribution for π_0 and θ is the same as prior 3 of Table 1.

		Log odds ratio (θ)	Odds ratio (exp θ)
Prior A:	$sens_1 = sens_0 \sim U(0.1, 0.4)$ $spec_1 = spec_0 \sim U(0.95, 0.98)$	0.95 (−0.51, 2.07)	2.58 (0.60, 7.95)
Prior B:	$sens_1 = sens_0 \sim U(0.2, 0.5)$ $spec_1 = spec_0 \sim U(0.96, 0.99)$	0.89 (−0.01, 1.75)	2.43 (0.99, 5.74)
Prior C:	$\text{logit } sens_1, \text{logit } sens_0 \sim \text{MVN} \left(\begin{pmatrix} -1.30 \\ -1.30 \end{pmatrix}, \begin{pmatrix} 0.21 & 0.17 \\ 0.17 & 0.21 \end{pmatrix} \right)$ $\text{logit } spec_1, \text{logit } spec_0 \sim \text{MVN} \left(\begin{pmatrix} 3.42 \\ 3.42 \end{pmatrix}, \begin{pmatrix} 0.06 & 0.05 \\ 0.05 & 0.06 \end{pmatrix} \right)$	1.08 (−0.48, 2.22)	2.93 (0.62, 9.24)
Prior D:	$\text{logit } sens_1, \text{logit } sens_0 \sim \text{MVN} \left(\begin{pmatrix} -0.90 \\ -0.90 \end{pmatrix}, \begin{pmatrix} 0.13 & 0.10 \\ 0.10 & 0.13 \end{pmatrix} \right)$ $\text{logit } spec_1, \text{logit } spec_0 \sim \text{MVN} \left(\begin{pmatrix} 3.89 \\ 3.89 \end{pmatrix}, \begin{pmatrix} 0.13 & 0.10 \\ 0.10 & 0.13 \end{pmatrix} \right)$	0.94 (−0.09, 1.77)	2.57 (0.91, 5.89)

3 A proposed method for robust Bayesian analysis for case–control studies with potentially severe exposure misclassification

Assume that our target parameter of interest is θ , given a suitable range of possible values for η , a set of nuisance parameters, the goal of the method of this paper is to find the feasible range of $\hat{\theta}(\eta)$. In the exposure misclassification example, for example, $\eta = \{sens_1, sens_0, spec_1, spec_0, \pi_0\}$. Here, we have denoted by $\hat{\theta}(\eta)$ an estimate of θ , given η , such that the goal can be written as

$$\text{minimize/maximize } \hat{\theta}(\eta) \text{ over } \eta \text{ subject to } \eta \in \mathcal{E}$$

In this paper, $\hat{\theta}(\eta)$ represent certain percentile values of the posterior distribution of θ given η and the data. Thus, if $\hat{\theta}(\eta)$ is the posterior median, then

$$[\min_{\eta \in \mathcal{E}} \hat{\theta}(\eta), \max_{\eta \in \mathcal{E}} \hat{\theta}(\eta)]$$

defines the feasible range of the posterior medians of θ .

A similar aim has been described by Vansteelandt, Goetghebeur, Kenward and Molenberghs [31], although these authors did not use a Bayesian estimate or credible interval for $\hat{\theta}(\eta)$. In principle, the methods of this paper can be applied to other non-Bayesian estimators for $\hat{\theta}(\eta)$. However, Bayesian estimators are used here in order that we may take advantage of the use of informative prior distributions, which can greatly aid the extraction of meaningful information from data with severe misclassification, as we have seen in Section 2.

Furthermore, by seeking $[\min_{\eta \in \mathcal{E}} \hat{\theta}(\eta), \max_{\eta \in \mathcal{E}} \hat{\theta}(\eta)]$, we provide bounds to the set of posterior median/credible intervals for the set of priors for (θ, η) where θ and η are independent and that the density of η is 0 outside $\mathcal{E}$, and thus provides a means of carrying robust Bayes analysis (see Appendix A).

In the exposure misclassification problem considered here, $\hat{\theta}(\eta)$ is a posterior percentile of the log odds ratio θ , which is the inverse function of the cumulative posterior distribution function of θ . Thus, denoting the cumulative distribution by $F(\theta|X, \eta)$ and its inverse by $F_{\theta|X, \eta}^{-1}(p)$ for a percentile p , we have:

$$\hat{\theta}(\eta) = F_{\theta|X, \eta}^{-1}(p)$$

$$F(\theta|X, \eta) = \int_{-\infty}^{\theta} p(\theta|X, \eta) d\theta$$

$$p(\theta|X, \eta) = \frac{\text{Lik}(X|\theta, \eta)p(\theta|\eta)}{p(X|\eta)}$$

$$\text{Lik}(X|\theta, \eta) = \binom{N_1}{Y_1} \binom{N_0}{Y_0} p_1^{Y_1} (1 - p_1)^{N_1 - Y_1} p_0^{Y_0} (1 - p_0)^{N_0 - Y_0}$$

$$p_i = \pi_i \text{sens}_i + (1 - \pi_i)(1 - \text{spec}_i), \quad i = 0, 1$$

$$\text{logit } \pi_1 = \text{logit } \pi_0 + \theta$$

where X denote the data: $X = \{Y_1, Y_0, N_1, N_0\}$.

It can be seen that as a function of θ , the likelihood $\text{Lik}(X|\theta, \eta)$ only depends on $p_1^{Y_1}(1 - p_1)^{N_1 - Y_1}$, and $p_1^{Y_1}(1 - p_1)^{N_1 - Y_1}$ only depends on sens_1 , spec_1 and π_0 , and not on sens_0 and spec_0 . Thus, sens_0 and spec_0 cannot affect the posterior percentile function $\hat{\theta}(\eta)$. Indeed, given π_0 , the likelihood does not even depend on Y_0 and N_0 . We are therefore considering a probabilistic modelling of the case data $\{Y_1, N_1\}$ only. This may appear a very unusual way of analysing case-control data, but it has been suggested in Zelen and Parker [32] that in a case-control study, we may have *a priori* information over the prevalence of exposure among the controls, if the controls are representative of the general population, and thus control data are not always necessary. In the presence of severe misclassification, the control data can tell us very little about the true exposure prevalence due to data dilution as discussed above, and *a priori* information, such as the degree to which the prevalence of misclassified exposure differs from the prevalence of true exposure, can be as important or more important than the control data.

As an example, consider the control data of the above example, where we have $Y_0 = 28, N_0 = 615$. Assuming an uninformative prior distribution of $\pi_0 \sim U(0, 1)$ and supposing $\text{sens}_0 = 0.3, \text{spec}_0 = 0.95$ (severe misclassification), the posterior distribution of π_0 has median and 95% credible interval 0.02 (0.0008, 0.08), while those for p_0 is 0.055 (0.050, 0.070). Thus, the control data itself tells us very little about the prevalence π_0 (even though it gives us considerable information for the misclassified prevalence p_0).

As the information concerning π_0 given by the data is limited, it is sensible to seek alternative sources of information. One potential source of information can be gained by considering the measurement error model that relates the true exposure to the observed exposure. For example, assuming a classical measurement error model, the true exposure ought to have lower prevalence than the observed prevalence. This can be seen by writing the observed EMF exposure levels (Z) as their true EMF levels (ζ) plus an error term (ε):

$$Z = \zeta + \varepsilon$$

where ε is typically 0 centred, and ε is independent of ζ . Our observed prevalence estimate ($p = Y/N$) estimates $\Pr(Z > c)$, and our true prevalence is $\Pr(\zeta > c)$, where c denotes a certain threshold, 0.3 μT in this particular case. If ζ and ε were both Normal, and ε has 0 mean, then $\Pr(Z > c)$ will always be greater than $\Pr(\zeta > c)$ if $\Pr(\zeta > c) < 0.5$. Since 0.3 μT is quite a high threshold for exposure, $\Pr(\zeta > c) \ll 0.5$, and hence we would expect the observed prevalence to be greater than the true prevalence. Given that the 95% confidence intervals for p_0 is (0.03, 0.065), let us tentatively assume that a reasonable range for π_0 is

$$0.01 \leq \pi_0 \leq 0.05 \quad (5)$$

3.1 Finding $\left[\min_{\eta \in \mathcal{L}} \hat{\theta}(\eta), \max_{\eta \in \mathcal{L}} \hat{\theta}(\eta) \right]$

To conduct the robust Bayesian analysis, we need to further specify feasible ranges for the parameters $sens_1$ and $spec_1$. As the misclassification is potentially severe, let us assume:

$$0.1 \leq sens_1 \leq 0.5 \quad (6)$$

$$0.95 \leq spec_1 \leq 1 \quad (7)$$

Equation (6) indicates that sensitivity is unlikely to be high.¹ This is because of the nature of the exposure assessment. For a rare exposure, however, it is often feasible to have a fairly tight bound on the specificity, such as (7). This is because if we believe that classification is better than chance, then $spec_1 > 1 - p_1$.² If misclassification is non-differential, then $spec_1 > 1 - p_0$. Given that the 95% confidence interval for p_0 is (0.03, 0.065), it seems reasonable that $spec_1$ has a lower bound of 0.95.

For the analysis, let us also assume at this stage that the prior distribution of θ is

$$\theta \sim N(0, 0.5)$$

although this can be relaxed later.

Our aim is to find

$$[\min \hat{\theta}(\pi_0, sens_1, spec_1), \max \hat{\theta}(\pi_0, sens_1, spec_1)]$$

subject to constraints (6)–(7) where $\hat{\theta}(\pi_0, sens_1, spec_1)$ denote the posterior median of θ given $(\pi_0, sens_1, spec_1)$. Since

$$\theta = \text{logit } \pi_1 - \text{logit } \pi_0$$

for $\hat{\theta}^+ = \max \hat{\theta}(\pi_0, sens_1, spec_1)$, we want $\text{logit } \pi_1$ to be as positive as possible, while $\text{logit } \pi_0$ should be as negative as possible. We note also that

$$\frac{\partial \pi}{\partial sens} = \frac{1 - spec - p}{(sens + spec - 1)^2} \begin{cases} > 0 & \text{if } spec < 1 - p \\ = 0 & \text{if } spec = 1 - p \\ < 0 & \text{if } spec > 1 - p \end{cases} \quad (8)$$

and

$$\frac{\partial \pi}{\partial spec} = \frac{sens - p}{(sens + spec - 1)^2} \begin{cases} > 0 & \text{if } sens > p \\ = 0 & \text{if } sens = p \\ < 0 & \text{if } sens < p \end{cases} \quad (9)$$

if $sens + spec \neq 1$. Moreover, when $sens + spec > 1$, $sens \geq p$ and $spec \geq 1 - p$. Therefore for $\hat{\theta}^+$, we would expect $sens_1$ to be found at its minimum, $spec_1$ at its maximum and π_0 at its minimum, and the reverse for $\hat{\theta}^- = \min \hat{\theta}(\pi_0, sens_1, spec_1)$. This implies that $[\hat{\theta}^-, \hat{\theta}^+]$ should be

$$[\hat{\theta}(\pi_0 = 0.05, sens_1 = 0.5, spec_1 = 0.95), \hat{\theta}(\pi_0 = 0.01, sens_1 = 0.1, spec_1 = 1)] = [-0.09, 4.47]$$

Of course, the above does not give us a proof that $[\hat{\theta}^-, \hat{\theta}^+] = [-0.09, 4.47]$ since the relationship between $\hat{\theta}(\pi_0, sens_1, spec_1)$ and $(\pi_0, sens_1, spec_1)$ is much more complicated than that between θ and

¹ If we assume the true exposure and the observed exposure levels follow a bivariate Normal distribution with equal mean, with the observed exposure levels having a standard deviation that is 1.2 times that of the true exposure, and correlation in the region of 0.3–0.7, then for an exposure classification defined by the top 5% of underlying exposure, the sensitivity is in the region of 0.1–0.5. The range of correlation (0.3–0.7) comes from a study which estimated r^2 between bedroom measurements and personal monitoring [10].

² This can be seen by rearranging (1) as $spec = 1 - p + \pi(sens + spec - 1)$. Because $0 \leq \pi \leq 1$, when $sens + spec > 1$, $spec \geq 1 - p$.

$(\pi_0, \text{sens}_1, \text{spec}_1)$. In order to give us more confidence that we have found our extremes, we may use a search algorithm to identify $[\hat{\theta}^-, \hat{\theta}^+]$. For this purpose, we wrote a program in Matlab, which makes use of the **fmincon** algorithm in its Optimization Toolbox. The optimization took less than 1 s and confirmed that $(\pi_0 = 0.05, \text{sens}_1 = 0.5, \text{spec}_1 = 0.95)$ and $(\pi_0 = 0.01, \text{sens}_1 = 0.1, \text{spec}_1 = 0.95)$ are indeed local minimum and maximum.

4 Allowing for uncertainty in the prior distribution of θ

To facilitate exposition, in the previous section, we use the prior distribution of $\theta \sim N(0, 0.5)$ without explanation. This prior distribution places 95% weight in the odds ratio, i.e. $\exp \theta$, being between 0.25 and 4.0. If the relationship between exposure to high levels of EMF and childhood leukaemia is unlikely to be greatly confounded by other factors (c.f.[24], such a prior distribution might be reasonable. It is also the prior distribution used by Greenland and Kheifets [18] in their meta-analysis of childhood leukaemia-EMF studies. However, it is understandable that not all readers may be happy with this prior. Some may find the prior too narrow, and others too wide. One advantage of using the method of this paper is that we can also specify a “range” of prior distribution. For example, it may be supposed that a “reasonable” range of prior variance for the Normal distribution may be from 0.3 to 0.7, such that we have:

$$\theta \sim N(0, \sigma^2)$$

$$0.3 \leq \sigma^2 \leq 0.7$$

Intuitively, the effect of having a prior distribution with a smaller variance is to shrink the estimates further towards zero. Therefore, we would expect $\sigma^2 = 0.7$ for both $\hat{\theta}^+$ and $\hat{\theta}^-$, since previous calculations showed that $\hat{\theta}^+$ is positive and $\hat{\theta}^-$ is negative. Again a search algorithm may be used to confirm this. This is indeed the case for this example, and we have:

$$[\hat{\theta}^-, \hat{\theta}^+] = [-0.11, 4.67] \quad (10)$$

5 Imposing additional constraints in the parameters sens_1 , spec_1 and π_0

So far, we have only given bounds to the parameters p_{i0} , sens_1 , spec_1 and σ^2 . It may be asked whether additional information concerning the sensitivity or specificity in the control population may help us further narrow our feasible region for $\hat{\theta}$. Intuitively, giving bounds to sens_0 and spec_0 may help us narrow the bounds on sens_1 and spec_1 since if misclassification is nearly non-differential, $\text{sens}_1, \text{spec}_1$ should not be too different from $\text{sens}_0, \text{spec}_0$. Furthermore, p_0 is related to sens_0 , spec_0 and π_0 through (1). Therefore, giving bounds to p_0 can also potentially restrict our feasible region of $(\pi_0, \text{sens}_1, \text{spec}_1)$.

In this section, we present a scheme that would enable us to determine whether a particular combination of $\pi_0, \text{sens}_1, \text{spec}_1$ satisfies the constraints:

Given that $\hat{p}_0 = 28/615 = 0.046$ (95% confidence interval = $[0.03, 0.06]$), let us assume

$$[p_0]^- \leq p_0 \leq [p_0]^+ \quad (11)$$

$$^- \leq \text{sens}_1 - \text{sens}_0 \leq [\text{sens}_1 - \text{sens}_0]^+ \quad (12)$$

$$^- \leq \text{spec}_1 - \text{spec}_0 \leq [\text{spec}_1 - \text{spec}_0]^+ \quad (13)$$

for some $\varepsilon_{\text{sens}}$ and $\varepsilon_{\text{spec}}$. In this section, we present a scheme that would enable us to determine whether a particular combination of $\pi_0, \text{sens}_1, \text{spec}_1$ satisfies constraints (11), (12), (13) also.

5.1 Determining if a set of $(\pi_0, \text{sens}_1, \text{spec}_1)$ falls within the feasible region

Denote the lower and upper bounds of a quantity x by $[x]^-$ and $[x]^+$ respectively, e.g. $[\text{sens}_1]^- \leq \text{sens}_1 \leq [\text{sens}_1]^+$, $[\text{spec}_1 - \text{spec}_0]^- \leq \text{spec}_1 - \text{spec}_0 \leq [\text{spec}_1 - \text{spec}_0]^+$. To determine if a set of $(\pi_0, \text{sens}_1, \text{spec}_1)$ is feasible, we can consider the following:

1. Assuming all of the parameters are given feasible ranges, i.e. $[x]^+ \geq [x]^-$ for all x , and assuming also $0 < \pi_0 < 1$, check whether:

$$(a) \quad S(\text{sens}_0)^- = \max([\text{sens}_0]^- , \text{sens}_1 - [\text{sens}_1 - \text{sens}_0]^+) \\ > \min([\text{sens}_0]^+ , \text{sens}_1 - [\text{sens}_1 - \text{sens}_0]^-) = S(\text{sens}_0)^+$$

$$(b) \quad S(\text{spec}_0)^- = \max([\text{spec}_0]^- , \text{spec}_1 - [\text{spec}_1 - \text{spec}_0]^+) \\ > \min([\text{spec}_0]^+ , \text{spec}_1 - [\text{spec}_1 - \text{spec}_0]^-) = S(\text{spec}_0)^+$$

$$(c) \quad S(p_0)^- = [p_0]^- > [p_0]^+ = S(p_0)^+$$

Evidently, if any of (a), (b) or (c) is true, the set of $(\pi_0, \text{sens}_1, \text{spec}_1)$ is outside the feasible region, as it is not possible to find values of $\text{sens}_0, \text{spec}_0, \text{sens}_1 - \text{sens}_0, \text{spec}_1 - \text{spec}_0, p_0, p_0 - \pi_0$ that satisfy the constraints.

2. If all of 1(a), 1(b), 1(c) are untrue, then we further test for feasibility by considering whether:

$$\max \left(\frac{(S(p_0)^- - (1 - \pi_0)(1 - S(\text{spec}_0)^-))/\pi_0,}{S(\text{sens}_0)^-} \right) \\ > \min \left(\frac{(S(p_0)^+ - (1 - \pi_0)(1 - S(\text{spec}_0)^+))/\pi_0,}{S(\text{sens}_0)^+} \right) \quad (14)$$

If this is true then $(\pi_0, \text{sens}_1, \text{spec}_1)$ is not feasible. If it is untrue, then it is feasible. This comes from the fact that $\text{sens}_0, \text{spec}_0, \pi_0$ and p_0 are related by identity (1). Equation (14) is derived by considering the maximum and minimum of sens_0 implied by constraints given to spec_0 and p_0 as well as the location of π_0, sens_1 and spec_1 . The rationale behind the scheme can be seen by considering Figure 1. It can be seen

Figure 1: (a) A situation where the set of $(\pi_0, \text{sens}_1, \text{spec}_1)$ is feasible as it is possible to draw a line through intervals defined by $[1 - S(\text{spec}_0)^+, 1 - S(\text{spec}_0)^-]$, $[S(\pi_0)^-, S(\pi_0)^+]$ and $[S(\text{sens}_1)^-, S(\text{sens}_1)^+]$ is shown. (b) A situation where it is not possible to draw such a line is shown.

that in order that a $(\pi_0, \text{sens}_1, \text{spec}_1)$ be feasible, we must be able to draw a straight line through $[S(p_0)^-, S(p_0)^+]$, $[S(\text{sens}_0)^-, S(\text{sens}_0)^+]$ and $[1 - S(\text{spec}_0)^+, 1 - S(\text{spec}_0)^-]$. (This is because $p_0 = \pi_0 \text{sens}_0 + (1 - \pi_0)(1 - \text{spec}_0)$ represents a linear relationship between p_0 and π_0 .) The range of $[S(p_0)^-, S(p_0)^+]$, $[S(\text{sens}_0)^-, S(\text{sens}_0)^+]$ and $[1 - S(\text{spec}_0)^+, 1 - S(\text{spec}_0)^-]$ is in turn determined by the position of $(\text{sens}_1, \text{spec}_1)$ as well as π_0 .

5.2 Finding $\left[\min_{\eta \in \mathcal{E}} \hat{\theta}(\eta), \max_{\eta \in \mathcal{E}} \hat{\theta}(\eta) \right]$ with additional constraints

Now assume we have the following constraints:

$$0.01 \leq \pi_0 \leq 0.05 \quad (15)$$

$$0.1 \leq \text{sens}_0 \leq 0.5 \quad (16)$$

$$0.25 \leq \text{sens}_1 \leq 0.4 \quad (17)$$

$$0.95 \leq \text{spec}_0 \leq 1 \quad (18)$$

$$0.95 \leq \text{spec}_1 \leq 1 \quad (19)$$

$$0.03 \leq p_0 \leq 0.065 \quad (20)$$

$$-0.05 \leq \text{sens}_1 - \text{sens}_0 \leq 0.05 \quad (21)$$

$$-0.02 \leq \text{spec}_1 - \text{spec}_0 \leq 0.02 \quad (22)$$

$$0.3 \leq \sigma^2 \leq 0.7 \quad (23)$$

The presence of constraints such as (15)–(22) makes it difficult to use simple relationships such as (8) and (9) to locate the extremes of $\hat{\theta}$. For this, an optimization algorithm becomes more useful. Still, it should be borne in mind that a surface such as $\hat{\theta}(\pi_0, \text{sens}_1, \text{spec}_1)$ is not necessarily convex, and multiple local optima may exist. For this reason, it is necessary to repeat the optimization a number of times from different starting points in order to adequately explore the parameter space. For this particular example, we repeated the optimization 10 times (using Matlab's **fmincon** as before) with the locations of the extrema given in Table 3.

Table 3: Location of $(\pi_0, \text{sens}_1, \text{spec}_1, \sigma^2)$ at the maximum and minimum posterior median of θ given the constraints of (15)–(23).

	$(\pi_0, \text{sens}_1, \text{spec}_1, \sigma^2)$	Posterior median θ (exp θ)
max $\hat{\theta}$	(0.01, 0.01, 0.991, 0.7)	4.44 (84.8)
min $\hat{\theta}$	(0.05, 0.5, 0.95, 0.7)	−0.11 (0.90)

Compared to the previous results (10), we see that the lower bound of −0.11 remained unchanged but the upper bound is shrunk towards 0. Thus, the effect of constraints (15)–(22) is to limit the upper extreme of $\hat{\theta}$. The location of min $\hat{\theta}$ is not limited by the additional constraints imposed in (15)–(22), although the location of max $\hat{\theta}$ is. The active constraints at max $\hat{\theta}_M$ and max $\hat{\theta}_U$ are (20), (21) and (22). These can be derived as results from the optimization.

6 The use of the method as a tool for sensitivity analysis

As noted in the introduction and Section 2 of this paper, Bayesian, probabilistic strategies for adjustment of exposure misclassification might suffice in the scenario where misclassification is not too serious, where the data are not too “diluted”. In the situation where exposure misclassification is potentially serious, not only does it become more difficult to elicit prior distributions for the parameters, but the dilution of the data means that it is unlikely that the data can tell us much about the direction of θ . This is when the method of this paper is most useful. For example, we might ask, what values do $sens_1$, $sens_0$, $spec_1$, $spec_0$ have to take in order that we may have evidence of a positive relationship between EMF and childhood leukaemia, or how does departure from the assumption of non-differential misclassification affect inference? These questions can be answered by using the method of this paper.

Before we discuss further, it will be useful to introduce the following terminology. Denoting by $\hat{\theta}_M(\eta)$ the posterior median of θ given η and $\hat{\theta}_L(\eta)$ the posterior 2.5%-ile and $\hat{\theta}_U(\eta)$ the posterior 97.5%-ile, let us define the

Feasible posterior median interval (FPMI) = $[\min \hat{\theta}_M, \max \hat{\theta}_M]$

Feasible posterior credible interval (FPCI) = $[\min \hat{\theta}_L, \max \hat{\theta}_U]$

In Figure 2, we can examine how departures from non-differential misclassification affect the FPMI and FPCI. Here, we see that while departures from non-differential *specificity* increases the FPMI and FPCI,

Figure 2: Sensitivity analyses showing how FPMI and FPCI vary with departures from non-differential misclassification, assuming all other constraints hold.

departures from non-differential *sensitivity* did not affect the intervals very much. In Figure 3, we can look at how restricting the ranges of $sens_1, sens_0, spec_1, spec_0$ affect the intervals. It can be seen that if $spec_1$ and $spec_0$ were greater than 0.97, then there is greater evidence for a positive relationship between EMF and childhood leukaemia. As $sens_0$ and $sens_1$ increase, the width of the FPMI and FPCI decreases. We also see that when $(spec_1, spec_0)$ are > 0.97 , and $(sens_1, sens_0)$ are between 0.2 and 0.4, the FPCI just includes 0. Further increase in $(sens_1, sens_0)$ does not lead to the FPCI being further away from 0.

Finally we may also look at how π_0 affect the FPMI and FPCI. In Figure 4, we can see that as π_0 increases, the upper limit of the FPMI and FPCI become lower, but the lower limit remains nearly the same.

Figure 3: Sensitivity analyses showing how FPMI and FPCI vary with changes in the feasible range of $sens_1, sens_0, spec_1$ and $spec_0$, assuming all other constraints hold.

Figure 4: Sensitivity analyses showing how FPMI and FPCI vary with changes in the feasible range of π_0 , assuming all other constraints hold.

7 Discussion

In this paper, we have introduced a new method for carrying out sensitivity analysis in case-control studies subject to exposure misclassification bias. In traditional Bayesian analysis, in the presence of severe misclassification, results are very sensitive to the prior distributions given, and different investigators may have different prior distributions. A common way to deal with uncertainty in prior distributions is through the use of a hierarchical prior [13]. However, in specifying a hierarchical prior, one is still specifying a *unique* prior distribution. If one does not agree with the prior, one also cannot strictly agree with the posterior. The robust Bayesian approach of this paper offers an alternative. Instead of averaging out results from different prior distributions, we seek out the most extreme inference possible among a specified class of prior distributions. Readers can compare their own beliefs with the feasible region specified for the various parameters. If his/her belief falls within the feasible region, then his/her posterior inference will also fall within the FPFI/FPCI. If the feasible region is wide, however, inevitably the FPFI/FPCI will also be wide. In these situations, perhaps a better use of this method is in exploring what misclassification probabilities are needed in order that the data may confer evidence in support of a positive or a negative association. As a tool for sensitivity analysis, the advantage of this approach is that one does not have to specify unique values for the uncertain parameters, but can instead specify ranges for the parameters.

A major contribution of this paper is its computational aspect. It is increasingly realized that data collected from observational studies cannot give unbiased estimates of epidemiological quantities of interest (e.g. [17, 20]), and that standard confidence intervals often underestimate the true uncertainty associated with estimates as they ignore bias. Therefore it has often been suggested that realistic models of epidemiologic data should take into account of uncertainty of bias by using models that integrate data with subjective “expert” knowledge [17, 20, 25]. Some have proposed that we seek out feasible region of inference given subjectively specified constraints for the unidentifiable parameters [31]. Application of these approaches, however, has been limited to simple scenarios, where such bounds can be computed analytically. In this paper, this is extended through the use of an optimization algorithm to situation where no such analytical solution exists. The method of this paper also paves the way for more general robust Bayesian inference, which has yet to become popular despite its philosophical integrity [2, 3]. The present paper shows that this type of inference is computationally feasible in a situation involving three unidentifiable parameters, although the existence of multiple local minima/maxima remains a potential difficulty in the general adoption of this method. Because the computation is suitably fast in the example of this paper, optimization can be repeated at a large number of starting points to ensure that the failure to locate the global optima is unlikely. As the number of unidentifiable parameters increases, this may become more and more infeasible, and further research is needed to help us locate the global optima in those situations.

Appendix A: robust Bayes interpretation of FPMIs and FPCIs

In robust Bayes analysis, we seek to summarize the many possible posterior inferences arising from a *class* of prior distributions. Assuming θ and η are independent and supposing $\hat{\theta}$ is a posterior percentile of θ , seeking the minimum and maximum of $\hat{\theta}$ can be thought of as seeking the minimum and maximum posterior percentile among the class of prior distributions which have zero density outside the feasible region of η . To see this, note that the posterior density of θ can be written as:

$$p(\theta|X) = \int_{\eta} p(\theta, \eta|X) d\eta$$

Denoting the cumulative distribution of θ given X by $F_{\theta|X} : \theta \rightarrow p$, which maps θ onto the percentiles p , we have:

$$\begin{aligned}
F_{\theta|X}(\theta) &= \int_{-\infty}^{\theta} \int_{\eta} p(\theta, \eta|X) d\eta d\theta \\
&= \int_{-\infty}^{\theta} \int_{\eta} \frac{p(X|\theta, \eta)p(\theta)p(\eta)p(X|\eta)}{p(X|\eta)p(X)} d\eta d\theta \quad (\text{by Bayes' theorem}) \\
&= \int_{\eta=-\infty}^{\theta} p(\theta|X, \eta) d\theta p(\eta|X) d\eta \\
&= \int_{\eta} F_{\theta|X, \eta}(\theta) p(\eta|X) d\eta
\end{aligned}$$

Hence, we see that the cumulative distribution of θ given X is a weighted average of the cumulative distribution of θ given X and η . Now, if our prior distribution of η belongs to a class that has zero mass for values outside the feasible region of η , denoted $\mathcal{E}$, i.e.:

$$p(\eta) = 0 \quad \forall \eta \notin \mathcal{E}$$

then

$$p(\eta|X) = 0 \quad \forall \eta \notin \mathcal{E}$$

and because averages cannot be greater than the maximum or less than the minimum,

$$\min_{\eta \in \mathcal{E}} F_{\theta|X, \eta}(\theta) \leq F_{\theta|X}(\theta) \leq \max_{\eta \in \mathcal{E}} F_{\theta|X, \eta}(\theta)$$

Now because our Bayesian estimates $\hat{\theta}$ are percentile functions $F^{-1} : p \rightarrow \theta$, which is the inverse of the cumulative distribution function. Since the cumulative distribution function $F(\theta)$ is necessarily a monotonically increasing function, we have:

$$\min_{\eta \in \mathcal{E}} \hat{\theta} = \min_{\eta \in \mathcal{E}} F_{\theta|X, \eta}^{-1}(p) \leq F_{\theta|X}^{-1}(p) \leq \max_{\eta \in \mathcal{E}} F_{\theta|X, \eta}^{-1}(p) = \max_{\eta \in \mathcal{E}} \hat{\theta}$$

Thus, by finding $\min_{\eta} \hat{\theta}$ and $\max_{\eta} \hat{\theta}$, we give bounds to $F_{\theta|X}^{-1}(p)$. Note that when we give bounds to $F_{\theta|X}^{-1}(p)$, we are assuming that the prior distribution of θ is the same as the prior distribution we use to calculate the bounds (i.e. $F_{\theta|X}^{-1}(p)$ and $F_{\theta|X, \eta}^{-1}(p)$ share the same prior distribution for θ). For this to be possible, the prior distribution of θ must not depend on η .

References

1. Bashir S, Duffy S. The correction of risk estimates for measurement error. *Ann Epidemiol* 1997;7:154–64.
2. Berger J. The robust Bayesian viewpoint (with discussion). In: Kadane, J, editor. *Robustness of Bayesian analysis*. Amsterdam: North-Holland, 1984:63–144.
3. Berger J. Robust Bayesian analysis: sensitivity to the prior. *J Stat Plann Infer* 1990;25:303–28.
4. Bross I. Misclassification in 2 by 2 tables. *Biometrics* 1954;10:478–86.
5. Chu H, Wang Z, Cole S, Greenland S. Sensitivity analysis of misclassification: a graphical and a Bayesian approach. *Ann Epidemiol* 2006;16:834–41.
6. Chu R, Gustafson P, Le N. Bayesian adjustment for exposure misclassification in case-control studies. *Stat Med* 2010;29:994–1003.

7. Copeland K, Checkoway H, McMichael A, Holbrook R. Bias due to misclassification in the estimation of relative risk. *Am J Epidemiol* 1977;105:488–95.
8. Diamond E, Lilienfeld A. Effects of errors in classification and diagnosis in various types of epidemiological studies. *Am J Public Health Nations Health* 1962;52:1137–44.
9. Fox M, Lash T, Greenland S. A method to automate probabilistic sensitivity analyses of misclassified binary variables. *Int J Epidemiol* 2005;34:1370–6.
10. Friedman D, Hatch E, Tarone R, Kaune W, Kleinerman R, Wacholder S, et al. Childhood exposure to magnetic fields: residential area measurements compared to personal dosimetry. *Epidemiology* 1996;7:151–5.
11. Garthwaite P, Kadane J, O'Hagan A. Statistical methods for eliciting probability distributions. *J Am Stat Assoc* 2005;100: 680–700.
12. Gelman A. The boxer, the wrestler, and the coin flip: a paradox of robust Bayesian inference and belief functions. *Am Stat* 2006;60:146–50.
13. Gelman A, Carlin J, Stern H, Rubin D. Bayesian data analysis. Boca Raton, FL: Chapman and Hall/CRC, 2004.
14. Greenland S. Variance-estimation for epidemiologic effect estimates under misclassification. *Stat Med* 1988;7:745–57.
15. Greenland S. Sensitivity analysis, Monte Carlo risk analysis, and Bayesian uncertainty assessment. *Risk Anal* 2001;21:579–83.
16. Greenland S. Bayesian perspectives for epidemiologic research III. Bias analysis via missing-data methods. *Int J Epidemiol* 2009;38:1662–73.
17. Greenland S. Overthrowing the tyranny of null hypotheses hidden in causal diagrams. In: Dechter R, Geffner H, and Halpern J, editors. *Heuristics, probabilities, and causality: a tribute to Judea Pearl*. London: College Press, 2010:365–82.
18. Greenland S, Kheifets L. Leukemia attributable to residential magnetic fields: results from analyses allowing for study biases. *Risk Anal* 2006;26:471–82.
19. Gustafson P. Measurement error and misclassification in statistics and epidemiology: impacts and Bayesian adjustments. Boca Raton, FL: Chapman and Hall, 2004.
20. Gustafson P. On model expansion, model contraction, identifiability and prior information: two illustrative scenarios involving mismeasured variables. *Stat Sci* 2005;20:111–29.
21. Gustafson P, Greenland S. Curious phenomena in Bayesian adjustment for exposure misclassification. *Stat Med* 2006;25:87–103.
22. Gustafson P, Le N, Saskin R. Case-control analysis with partial knowledge of exposure misclassification probabilities. *Biometrics* 2001;57:598–609.
23. Jurek A, Greenland S, Maldonado G. How far from non-differential does exposure or disease misclassification have to be to bias measures of association away from the null? *Int J Epidemiol* 2008;37:382–5.
24. Kheifets L, Shimkhada R. Childhood leukemia and EMF: review of the epidemiologic evidence. *Bioelectromagnetics* 2005;7: S51–9.
25. Lash T, Fox M, Fink A. Applying quantitative bias analysis to epidemiologic data. New York, NY: Springer, 2009.
26. Linet M, Hatch E, Kleinerman R, Robison L, Kaune W, Friedman D, et al. Residential exposure to magnetic fields and acute lymphoblastic leukemia in children. *N Engl J Med* 1997;337:1–7.
27. Lunn D, Thomas A, Best N, Spiegelhalter D. WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Stat Comput* 2000;10:325–37.
28. Lyles R. A note on estimating crude odds ratios in case-control studies with differentially misclassified exposure. *Biometrics* 2002;58:1034–6.
29. MacLehose R, Gustafson P. Is probabilistic bias analysis approximately Bayesian? *Epidemiology* 2012;23:151–8.
30. Morrissey M, Spiegelman D. Matrix methods for estimating odds ratios with misclassified exposure data: extensions and comparisons. *Biometrics* 1999;55:338–44.
31. Vansteelandt S, Goetghebeur E, Kenward M, Molenberghs G. Ignorance and uncertainty regions as inferential tools in a sensitivity analysis. *Stat Sinica* 2006;16:953–79.
32. Zelen M, Parker R. Case-control studies and Bayesian inference. *Stat Med* 1986;5:261–9.