



Practice of Epidemiology

Bayesian Sample Size Determination for Case-Control Studies When Exposure May be Misclassified

Lawrence Joseph* and Patrick Bélisle

* Correspondence to Prof. Lawrence Joseph, Division of Clinical Epidemiology, McGill University Health Centre, Royal Victoria Hospital, 687 Pine Avenue West, V Building, Room V2.10, Montreal, Quebec, H3A 1A1, Canada (e-mail: lawrence.joseph@mcgill.ca).

Initially submitted February 1, 2013; accepted for publication July 10, 2013.

Odds ratios are frequently used for estimating the effect of an exposure on the probability of disease in case-control studies. In planning such studies, methods for sample size determination are required to ensure sufficient accuracy in estimating odds ratios once the data are collected. Often, the exposure used in epidemiologic studies is not perfectly ascertained. This can arise from recall bias, the use of a proxy exposure measurement, uncertain work exposure history, and laboratory or other errors. The resulting misclassification can have large impacts on the accuracy and precision of estimators, and specialized estimation techniques have been developed to adjust for these biases. However, much less work has been done to account for the anticipated decrease in the precision of estimators at the design stage. Here, we develop methods for sample size determination for odds ratios in the presence of exposure misclassification by using several interval-based Bayesian criteria. By using a series of prototypical examples, we compare sample size requirements after adjustment for misclassification with those required when this problem is ignored. We illustrate the methods by planning a case-control study of the effect of late introduction of peanut to the diet of children to the subsequent development of peanut allergy.

Bayesian methods; case-control study; misclassification error; sample size determination; study design

Abbreviations: ACC, average coverage criterion; ALC, average length criterion; CI, credible interval; HPD, highest posterior density; MWOC, modified worst outcome criterion; OR, odds ratio.

Statistical techniques that adjust for possible biases in observational studies are increasingly common in epidemiology (1–4). Methods for designing studies that will eventually need to be adjusted for such biases, however, have lagged behind. We address this gap by presenting a method for adjusting sample size requirements for case-control studies with possible misclassification bias.

As discussed in the classic text by Schlesselman (5), case-control designs can be used to estimate the effect of an exposure on the probability of a disease or condition. For example, it is hypothesized that the late introduction of peanut to the diet may increase the probability of peanut allergy in children (6). Suppose that a case-control study is to be conducted to estimate the effect of introduction of the food prior to 1 year of age compared with later introduction. Groups of children with and without peanut allergy will be surveyed, with their parents providing information about when their children were first

introduced to peanut or products containing peanut. Exposure misclassification may arise from the inaccurate recall of exposure information, which may also differ in magnitude between cases and controls. Under these circumstances, what sample size is required for accurate estimation of the odds ratio once statistical methods, which adjust for the misclassification due to recall bias, are applied?

Ignoring misclassification, Wickramaratne (7) and Lemeshow et al. (8) reviewed classical sample size methods for hypothesis testing and interval estimation for odds ratios. Frequentist sample size methods depend on accurate point estimates of the required inputs, which here include not only the exposure rates, but also the rates of misclassification within each disease class. As we will show, the estimated sample sizes can be very sensitive to minor changes to these inputs. It is therefore advantageous to consider Bayesian methods, where prior densities not only allow for uncertainty in the inputs, but

Table 1. Two-by-Two Table of Observed Data

	Disease +	Disease –	Total
Exposure +	a	c	$a + c$
Exposure –	b	d	$b + d$
Total	$a + b$	$c + d$	N

incorporate this uncertainty into the sample size requirements. This is especially important in the presence of misclassification, which induces a nonidentified model. As discussed by various authors (9–12), the calculation of sample sizes within nonidentified models is inherently different from regular problems, because the posterior density does not converge to a single point as the sample size increases. Therefore, even infinite sample sizes may not guarantee sufficient accuracy.

From a Bayesian viewpoint, sample size for case-control studies, including examination of the optimal case-to-control ratio, was addressed by De Santis et al. (13) and M'LAN et al. (14). Neither, however, considered the change in sample size resulting from possible exposure misclassification. Devine and Smith (15) addressed the change in sample size requirements induced by misclassification by using frequentist power-based criteria. Gustafson (16) reviewed general Bayesian methods to adjust for misclassified exposure data, and Stamey and Gerlach (17) considered misclassification for case-control studies, but their method requires a validation sample that is not always available.

SAMPLE SIZE DETERMINATION FOR ODDS RATIOS IN THE PRESENCE OF MISCLASSIFICATION

We begin by describing a model for adjusting odds ratios for exposure misclassification, similar to that used by Gustafson et al. (18). We then define several Bayesian sample size criteria and indicate how they can be applied to help design studies with bias-adjusted odds ratios. Technical details and numerical algorithms are given in the Appendix.

Let $i = 1, 2$ index the case and control populations, respectively. For any given sample size N , we observe the 2-by-2 layout given in Table 1. Let the true probability of exposure among cases be given by p_1 and the true probability of exposure among controls be given by p_2 . Let p'_1 and p'_2 be the observed probabilities in the presence of misclassification. The 2 sets of probabilities are related by the equations

$$p'_1 = p_1 \times S_1 + (1 - p_1) \times (1 - C_1) \quad \text{and} \\ p'_2 = p_2 \times S_2 + (1 - p_2) \times (1 - C_2),$$

where S_i and C_i are correct classification rates within case ($i = 1$) or control ($i = 2$) populations, defined by

$$S_i = \Pr\{\text{classified as exposed}|\text{truly exposed}\} \quad \text{and} \\ C_i = \Pr\{\text{classified as not exposed}|\text{truly not exposed}\}.$$

If $S_1 = S_2$ and $C_1 = C_2$, then there are 2 correct classification rate parameters to estimate; otherwise there are 4, giving

a total of $m = 4$ or $m = 6$ unknown parameters when added to p_1 and p_2 . Because all parameters have range $[0, 1]$, β densities may be used as prior distributions. The available data in Table 1 provide only 3 degrees of freedom, meaning that, in practice, the problem is nonidentifiable, and that informative prior distributions need to be placed over a subset of parameters to obtain reasonable inferences. Typically, one supposes some knowledge or limits about the classification rates, after which one can use less informative priors over p_1 and p_2 .

The likelihood function for the observed data is a product of 2 binomial functions, and using the notation from Table 1 is proportional to

$$(p'_1)^a (1 - p'_1)^b (p'_2)^c (1 - p'_2)^d. \quad (2)$$

The posterior density function is proportional to the product of the likelihood function and the prior density over the unknown parameters $p_1, p_2, S_1, S_2, C_1, C_2$. The odds ratio is then estimated by integrating the full posterior density to eliminate the nuisance parameters S_1, S_2, C_1, C_2 and introducing the change of variable odds ratio $= (p_1/(1 - p_1))/(p_2/(1 - p_2))$. Because there is no closed form solution, inference proceeds by Markov chain Monte Carlo methods. We used WinBUGS, version 1.2, software (MRC Biostatistics Unit, Cambridge, United Kingdom) to implement a Gibbs sampler algorithm.

The marginal posterior density of the odds ratio is typically summarized by a highest posterior density (HPD) credible interval. In planning a study, suppose we desire an interval of length l that includes the odds ratio with probability $1 - \alpha$. For example, we may wish to estimate the odds ratio to an accuracy of $l = 0.5$ with a $1 - \alpha = 95\%$ interval. The marginal posterior density of the odds ratio depends on the data, which are, of course, unknown at the planning stage. We can account for this uncertainty in different ways, leading to different sample size criteria. We consider 3 such criteria as follows: the average coverage criterion (ACC), the average length criterion (ALC), and the modified worst outcome criterion (MWOC), defined in detail in M'LAN et al. (14).

Allowing the posterior probability $1 - \alpha$ to vary with each potential data set while holding the credible interval length l fixed leads to a sample size that guarantees the desired posterior probability on average, that is, the ACC sample size. The average is taken over the set of all possible data sets, weighted by the probability that each data set may arise, as determined by the prior densities over all parameters.

Conversely, we can allow the HPD interval length l to vary while fixing the posterior probability at $1 - \alpha$. This ALC sample size averages the lengths of fixed probability HPD intervals over all possible data sets, again weighted by the prior distributions.

Rather than averaging over potential data sets, a conservative approach would be to ensure a maximum length of l and a minimum probability of $1 - \alpha$, regardless of the data set that occurs, termed the worst outcome criterion sample size. In practice, there is often at least 1 data set that leads to very poor accuracy, so that the worst outcome criterion sample size is infinite. For example, this is always the case when sampling from the posterior density of an odds ratio in which the rate

in the denominator might be close to 0, causing instability in the odds ratio estimate. Therefore, in this article we use the MWOC, where the desired length and posterior probability are guaranteed over a subset of all data sets with a given probability. For example, we might choose the sample size such that the desired l and $1 - \alpha$ are guaranteed over 95% of all data sets. We denote this by MWOC[95] or, more generally, MWOC[$100 \times (1 - \gamma)$], where γ represents the probability that a randomly selected data set will not satisfy the length and/or posterior probability requirements. The use of the MWOC avoids the situation of having to select an unnecessarily large sample size to guard against highly improbable data.

Because there are no closed form solutions, we used the numerical algorithm detailed in the Appendix to estimate the optimal sample size. A user-friendly program called SSCOR (Sample Size Calculations for Odds Ratios), which implements all of the above methods, is available at www.medicine.mcgill.ca/epidemiology/Joseph/. We next use this software to determine sample sizes for various scenarios that may occur in the planning of case-control studies, comparing situations with and without misclassification.

SAMPLE SIZES FOR PROTOTYPICAL SCENARIOS

The misclassification of exposure and the subsequent need to adjust the odds ratio to account for these errors has important implications for the design of case-control studies. In general, the larger the misclassification error and the more uncertain one is of the magnitude of this error, the more uncertainty there will be in the final odds ratio estimate, and consequently, the larger the sample size requirements will be. In this section, we present some prototypical scenarios that illustrate the degree to which sample sizes must increase with increasing amounts of misclassification.

Throughout, we will assume that a 95% HPD interval is desired. We will consider 3 values for the true odds ratio, approximately centered on odds ratio values of 0.7, 1, and 1.5, with desired total HPD interval lengths of 0.4, 0.2, and 0.8, respectively. The latter were chosen to be sufficiently small, such that definitive inferences can be made. For example, if the true odds ratio is 0.7 and the length of the HPD interval is 0.4, then a 95% credible interval close to (0.5, 0.9) can be expected, which is sufficiently far from the null value to convincingly demonstrate a protective effect. On the other hand, if the odds ratio is 1, then a total length of 0.2 will result in an interval similar to (0.9, 1.1), which will often be close enough to the null value to conclude no clinically important effect. Similarly, when the odds ratio is 1.5, the interval will be approximately (1.1, 1.9), which we assume is far enough from the null value to conclude a positive effect.

We assumed a $\beta(10, 90)$ prior for the exposure rate within the case group. This density provides a mean rate of 10%, which is typical of many exposures, and the parameters sum to 100, providing knowledge equivalent to 100 prior observations. To obtain odds ratios of 0.7, 1, and 1.5, we modeled the exposure rate among controls by $\beta(13.7, 86.3)$, $\beta(10, 90)$ and $\beta(6.9, 93.1)$ densities, respectively. These provide median odds ratios of 0.694 (95% prior credible interval (CI): 0.281,

1.65), 1 (95% CI: 0.387, 2.58), and 1.52 (95% CI: 0.550, 4.47) centered close to the target values.

For each of these scenarios, we calculated sample sizes assuming no misclassification, as well as low, moderate, and high degrees of misclassification. For moderate and high degrees of misclassification, we also considered narrower and wider prior densities around the central value, because knowledge about the misclassification rate can have as much of an effect on the sample size as the misclassification rate itself. Low misclassification was defined by a $\beta(681.5, 13)$ prior probability of correct classification and a density with 95% range from 0.97 to 0.99, implying a misclassification rate between 1% and 3%. We used $\beta(116, 12)$ and $\beta(214.35, 70.8)$ densities to represent the moderate and high misclassification rates with wider prior ranges, respectively, implying error rates between 5% and 15% and between 20% to 30%, respectively. For narrower prior ranges, we used misclassification rates between 9% and 11% for the moderate error rate and between 24% and 26% for high rates, corresponding to $\beta(3103.9, 344)$ and $\beta(5400.3, 1799.5)$ densities for correct classification, respectively.

For each of the above $3 \times 6 = 18$ scenarios, we considered the following 4 criteria: the ACC, ALC, MWOC[50], and MWOC[90]. We calculated 2 sample sizes for each scenario (except no misclassification), depending on whether one assumes the same or allows for different misclassification rates within the diseased and nondiseased populations. This determines whether there are 4 or 6 unknown parameters to estimate. Thus, we considered a total of 132 different scenarios. An upper limit of 100,000 subjects per group was set, because larger studies would usually not be practical.

Although any joint density over $p_1, p_2, S_1, S_2, C_1, C_2$ can be used, aside from the 4-parameter model that sets $S_1 = S_2$ and $C_1 = C_2$, we have chosen independent priors for each parameter. Although equation 1 exposes the relationship between p'_1 and p'_2 and the parameters from which they are derived, there is no particular reason to suspect that the correct classification probabilities are dependent on p_1 and p_2 . Nevertheless, should this be the case, our methods can be easily modified by using a different joint prior over the parameter space.

Table 2 presents the resulting sample sizes assuming $g = 1$. When there is no misclassification, the sample sizes range from a low of 813 for the MWOC[50] criterion when the odds ratio is 0.7 to a high of 32,072 for the most strict criterion, the MWOC[90], with an odds ratio of 1. The ACC and ALC sample sizes are intermediate to these extremes. In general, the sample sizes were largest for the narrowest interval around an odds ratio of 1 and smallest for an odds ratio of 0.7 with an interval length of 0.4.

Under the $m = 4$ parameter model, even low rates of misclassification greatly increase the desired sample sizes. For example, there was an approximately 20% increase in sample sizes when the odds ratio was 0.7 but often more than a doubling of the sample size when the odds ratio was 1, including 2 cases in which the 100,000 ceiling was reached, under the MWOC[90]. As expected, moderate and high levels of misclassification require even larger sample sizes, and having a better knowledge of the misclassification rate decreases the

Table 2. Sample Sizes per Group for the Prototypical Scenarios, Assuming a Case-to-Control Ratio of $g = 1$

Criterion	OR	Interval Length	No Misclassification	Low Misclassification	Moderate Misclassification and Narrow Prior	Moderate Misclassification and Wide Prior	High Misclassification and Narrow Prior	High Misclassification and Wide Prior
Identical misclassification parameters (4-parameter model)								
ACC	0.7	0.4	1,341	1,665	3,422	3,698	12,326	14,223
ALC	0.7	0.4	1,016	1,254	2,588	3,053	9,567	12,519
MWOC[50] ^a	0.7	0.4	813	1,004	2,055	2,344	7,495	9,210
MWOC[90] ^b	0.7	0.4	2,696	3,288	6,688	7,288	25,018	28,721
ACC	1	0.2	15,992	40,760	>100,000	>100,000	>100,000	>100,000
ALC	1	0.2	11,427	20,690	53,802	>100,000	>100,000	>100,000
MWOC[50]	1	0.2	8,595	10,745	23,454	57,437	92,838	>100,000
MWOC[90]	1	0.2	32,072	>100,000	>100,000	>100,000	>100,000	>100,000
ACC	1.5	0.8	3,514	8,936	41,237	>100,000	>100,000	>100,000
ALC	1.5	0.8	2,176	3,920	10,862	>100,000	72,435	>100,000
MWOC[50]	1.5	0.8	1,434	1,962	4,751	9,273	20,604	47,454
MWOC[90]	1.5	0.8	7,353	>100,000	>100,000	>100,000	>100,000	>100,000
Distinct misclassification parameters (6-parameter model)								
ACC	0.7	0.4	1,341	2,391	6,226	>100,000	>100,000	>100,000
ALC	0.7	0.4	1,016	1,653	3,957	>100,000	>100,000	>100,000
MWOC[50]	0.7	0.4	813	1,234	2,797	>100,000	24,810	>100,000
MWOC[90]	0.7	0.4	2,696	9,793	>100,000	>100,000	>100,000	>100,000
ACC	1	0.2	15,992	>100,000	>100,000	>100,000	>100,000	>100,000
ALC	1	0.2	11,427	>100,000	>100,000	>100,000	>100,000	>100,000
WMOC[50]	1	0.2	8,595	>100,000	>100,000	>100,000	>100,000	>100,000
MWOC[90]	1	0.2	32,072	>100,000	>100,000	>100,000	>100,000	>100,000
ACC	1.5	0.8	3,514	>100,000	>100,000	>100,000	>100,000	>100,000
ALC	1.5	0.8	2,176	>100,000	>100,000	>100,000	>100,000	>100,000
MWOC[50]	1.5	0.8	1,434	4,095	17,569	>100,000	>100,000	>100,000
MWOC[90]	1.5	0.8	7,353	>100,000	>100,000	>100,000	>100,000	>100,000

Abbreviations: ACC, average coverage criterion; ALC, average length criterion; g , case-to-control ratio; MWOC, modified worst outcome criterion; OR, odds ratio.

^a MWOC[50] denotes the MWOC at the 50% probability level, that is, guaranteeing the desired credible interval length and coverage over 50% of all data sets.

^b MWOC[90] denotes the MWOC at the 90% probability level, that is, guaranteeing the desired credible interval length and coverage over 90% of all data sets.

Table 3. Sample Sizes for Both Groups Combined for Planning a Study Concerning the Early Introduction of Peanut to the Diet

Criterion by Case-to-Control Ratio	No Misclassification	Low Misclassification	Moderate Misclassification
ACC			
1	2,908	3,242	4,918
2	3,078	3,441	5,160
3	3,468	3,876	5,848
ALC			
1	2,758	3,056	4,600
2	2,925	3,249	4,914
3	3,280	3,664	5,516
MWOC[50]^a			
1	2,654	2,958	4,440
2	2,843	3,141	4,734
3	3,196	3,564	5,336
MWOC[95]^b			
1	5,014	5,784	8,912
2	5,427	6,159	9,468
3	6,200	7,048	10,868

Abbreviations: ACC, average coverage criterion; ALC, average length criterion; MWOC, modified worst outcome criterion.

^a MWOC[50] denotes the MWOC at the 50% probability level, that is, guaranteeing the desired credible interval length and coverage over 50% of all data sets.

^b MWOC[95] denotes the MWOC at the 95% probability level, that is, guaranteeing the desired credible interval length and coverage over 95% of all data sets.

sample size compared with when this rate is less well known, a priori.

The situation is considerably worse when the misclassification rates are not assumed to be identical across groups, that is, when $m = 6$. Indeed, even low degrees of misclassification often create sample sizes above 100,000, except for the scenarios with an odds ratio of 0.7. High degrees of misclassification almost always lead to very large sample sizes.

Gustafson et al. (18) derived the posterior density of the odds ratio under misclassification by assuming an infinite sample size. Typically, the posterior density of the odds ratio narrows substantially up to a certain sample size, past which there are diminishing returns. When combined with the Monte Carlo methods given in the appendix of Dendukuri et al. (10), it is possible to determine whether an infinite sample size will satisfy a given criterion. For example, Table 2 indicates that the ACC sample size for an odds ratio of 1 for moderate misclassification with a narrow prior is greater than 100,000, but a sample size of just under 200,000 per group (400,000 in total) is sufficient to satisfy the ACC, reaching an average posterior probability of 0.952 at 200,000. On the other hand, even an infinite sample size is not sufficient for this same situation under high misclassification with a wide prior.

SAMPLE SIZE REQUIRED TO ACCURATELY ESTIMATE THE EFFECT OF LATE INTRODUCTION TO PEANUT ON PEANUT ALLERGY

The overall prevalence of peanut allergy among children in Montreal, Canada, is approximately 1.5% (19). Suppose we anticipate that 60% of cases have late exposure (95% prior interval from 55% to 65%, represented by a $\beta(229.8, 153.2)$ density) compared with 30% of controls with late exposure (95% prior interval from 25% to 35%, represented by a $\beta(100.5, 234.5)$ density), giving an odds ratio close to 3.5 (95% CI: 2.58, 4.80). If a case-control study is being planned, what should the sample size be so that the odds ratio is estimated to within a total HPD interval length of 1?

The sample size will depend on many factors, including which sample size criterion will be used, how many controls will be selected for each case, whether one allows for possible misclassification errors in parental information about the timing of first introduction to peanut in the diet, and the degree to which this misclassification is assumed to be known.

Table 3 provides the required sample sizes under a wide variety of possible design choices. Somewhat unrealistically, if no misclassification is assumed, for a case-to-control ratio of $g = 1$, the sample sizes range from less than 3,000 to more than 6,000, depending on the criterion selected. If one is content to ensure a posterior HPD length of 1 only on average (or median), sample sizes close to 3,000 are needed. If one wants to be 95% certain of obtaining an HPD interval of 1 or less, then sample sizes in the 5,000 to 6,000 range are required. The final choice can be based on the trade-off between the certainty of obtaining an HPD interval of length 1 versus the costs associated with the larger sample sizes. The effect of the case-to-control ratio can also be gleaned from Table 3, in which higher values of g raise the total sample size requirements by a few percentage points. Although a ratio of $g = 1$ is optimal, difficulties in finding cases, which are relatively rare, may lead to other choices.

If one more realistically assumes a low misclassification rate of 1%–3%, input as a $\beta(681.5, 13)$ density for the correct classification probability, and assumes the rate to be equal in the 2 groups ($m = 4$), then sample sizes rise by roughly 10% across all criteria. Under a moderate rate of misclassification of 9%–11%, represented by a $\beta(3103.852, 344.0302)$ density for the correct classification probability, the sample sizes rise by about 60% compared with the no misclassification scenario, although they may remain feasible.

It is interesting to consider what may happen if one plans a sample size for a study ignoring measurement error, but later analyzes the data considering measurement error. For example, if there are equal numbers of cases and controls ($g = 1$), the MWOC[50] from Table 3 suggests a total sample size of 2,654, assuming no measurement error. If this size is used for the study, which is subsequently analyzed by using a moderate rate of misclassification of between 9% and 11%, then the length of an HPD interval of probability 0.95 will be 1.2, which is about 20% wider than the original planned length.

We can also evaluate the effect of the $m = 6$ parameter model. If we use a moderate rate of misclassification of 9%–11% in both groups but allow distinct parameters for these rates, the ALC sample size under $g = 1$ and an HPD length of 1 is 5,248. However, it is also reasonable to assume that

misclassification rates are lower in the peanut allergic group (cases) compared with the nonallergic group (controls), because cases may make more effort to recall or may remember the history more accurately, given the likely reaction that would have occurred in the child upon early ingestion. For example, we might assume a moderate misclassification rate of 9%–11% for cases and a larger misclassification rate of 15%–25% for controls. Although the latter interval is wide, it remains plausible that many controls might not accurately remember this history, and the exact recall rate would typically not be accurately known. With $g = 1$ and using the ALC criterion with a total HPD interval length of 1, the total sample size is greater than 200,000 (more than 100,000 per group). Under these conditions, obtaining HPD lengths of 1 may not be feasible, even though the inputs are entirely plausible. Because an odds ratio of 3.5 is far from the null value of 1, a study with lower accuracy is still informative. Doubling the width from 1 to 2 reduces the sample size considerably, to a manageable 362.

All of the above sample sizes are fully Bayesian, in the sense that relatively strong prior information is assumed for each parameter and is used not only for the purposes of planning the study, but also within the eventual analysis. One cannot use noninformative (e.g., $\beta(1,1)$ or uniform prior densities) across all parameters, because the problem is nonidentifiable. However, it is possible to use relatively weak prior information for the exposure prevalence values within the case and control groups, provided good prior information is available for the misclassification parameters. This strategy can be used by researchers who prefer to “let the data speak for themselves” at the analysis stage, while still planning their studies to accommodate possible misclassification errors.

Forexample, we can consider the prior densities used above, but divide each parameter by 10, reducing the prior effective sample size. Thus, the β density for the probability of exposure in cases changes from $\beta(229.8, 153.2)$ to $\beta(22.98, 15.32)$, and the β density for the controls changes from $\beta(100.5, 234.5)$ to $\beta(10.05, 23.45)$. Summing across β parameters gives a prior effective sample size of 71.8 compared with the previous size of 718. With $g = 1$, an ALC sample size of 5,202 is required to attain a total HPD length of 1, even in the absence of misclassification. Under moderate misclassification of 9%–11%, the sample size roughly doubles to 9,626. Reduction of the effective prior sample size by another factor of 10 to 7.18 returns sample sizes that are larger than 200,000, even with no misclassification.

Another strategy is to separate the “design priors” used to generate the data from the “analysis priors” used at the analysis stage (20). Although the nonidentified nature of our problem requires at least some informative priors at both design and analysis stages, one can place uniform priors on p_1 and p_2 at the analysis stage. This will provide sample sizes for researchers wanting to use minimal prior information for the odds ratio at the analysis stage. Under this strategy, the ALC sample sizes under moderate misclassification are 5,908, 6,483, and 7,564 for $g = 1, 2,$ and $3,$ respectively.

DISCUSSION

The vast majority of case-control studies ignore misclassification, and even those that might consider this possibility

ignore the uncertainty in a priori knowledge of these misclassification rates. This is true both at the design and analysis phases, leading to sample sizes that are typically much too small and final estimates with credible intervals that are much too narrow. The methods presented here are important to the planning of such studies and serve as a warning that ignoring misclassification in study planning and analysis may lead to wildly optimistic interval estimates.

We have discussed different sample size criteria, which lead to different sample sizes for any given problem. A natural question, therefore, is which one to use. Clearly, the MWOC for low values of γ is more conservative than either the ACC or ALC, which guarantee the target values for posterior probability and length only on average. As we have done in Tables 2 and 3 here, we have found it useful to calculate the sample sizes that result from all criteria, including the MWOC $[1 - \gamma]$ for various values of γ , to develop a fuller understanding of the inherent tradeoffs between sample size and the risk of not meeting target values for interval length and posterior probability. On the basis of this information, a final sample size may be selected that balances statistical rigor with practical concerns. It is especially important for study designers to appreciate that, in many cases, the desired estimation accuracy cannot be attained even with an infinite sample size. Clearly, designing studies to have the lowest possible misclassification rates is important, if possible. In some cases, it must be admitted that no study design will result in misclassification rates low enough to derive sufficiently accurate estimates.

Although we have applied our methods to the design of case-control studies with exposure misclassification, similar methods can be developed for other study designs and different sources of bias. As bias adjustment methods gain in popularity at the analysis stage (1–4), methods for designing such studies will also increase in importance.

ACKNOWLEDGMENTS

Author affiliations: Division of Clinical Epidemiology, McGill University Health Centre, Montreal, Canada (Lawrence Joseph, Patrick Bélisle); and Department of Epidemiology and Biostatistics, McGill University, Montreal, Canada (Lawrence Joseph).

Conflict of interest: none declared.

REFERENCES

1. Greenland S. Multiple-bias modeling for analysis of observational data. *J R Stat Soc Ser A*. 2005;168(2):267–306.
2. Greenland S, Lash TL. Bias analysis. In: Rothman KJ, Greenland S, Lash TL, eds. *Modern Epidemiology*, 3rd ed. Philadelphia, PA: Lippincott-Wolters-Kluwer; 2008:345–380.
3. Hoggatt KJ, Greenland S, Ritz B. Adjustment for response bias via two phase analysis: an application. *Epidemiology*. 2009;20(6):872–879.
4. Greene N, Greenland S, Olsen J, et al. Estimating bias from loss to follow-up in the Danish National Birth Cohort. *Epidemiology*. 2011;22(6):815–822.

5. Schlesselman JJ. *Case-Control Studies: Design, Conduct, Analysis*. New York, NY: Oxford University Press; 1982.
6. Du Toit G, Katz Y, Sasieni P, et al. Early consumption of peanuts in infancy is associated with a low prevalence of peanut allergy. *J Allergy Clin Immunol*. 2008;122(5):984–991.
7. Wickramaratne PJ. Sample size determination in epidemiologic studies. *Stat Methods Med Res*. 1995;4(4):311–337.
8. Lemeshow S, Hosmer D, Klar J, et al. *Adequacy of Sample Size in Health Studies*. Chichester, England: John Wiley & Sons Ltd; 1990.
9. Rahme E, Joseph L, Gyorkos T. Bayesian sample size determination for estimating binomial parameters from data subject to misclassification. *J R Stat Soc Ser C Appl Stat*. 2000;49(1):119–128.
10. Dendukuri N, Rahme E, Blisle P, et al. Bayesian sample size determination for prevalence and diagnostic studies in the absence of a gold standard test. *Biometrics*. 2004;60(2):388–397.
11. Gustafson P. Sample size implications when biases are modelled rather than ignored. *J R Stat Soc Ser A*. 2006;169(4):865–881.
12. Stamey J, Seaman J, Young D. Bayesian sample size determination for inference on two binomial populations with no gold standard classifier. *Stat Med*. 2005;24(19):2963–2976.
13. De Santis F, Perone Pacifico M, Sambucini V. Optimal predictive sample size for case-control studies. *J R Stat Soc Ser C Appl Stat*. 2004;53(3):427–441.
14. M'LAN CE, Joseph L, Wolfson DB. Bayesian sample size determination for case-control studies. *J Am Stat Assoc*. 2006;101(474):760–772.
15. Devine OJ, Smith JM. Estimating sample size for epidemiologic studies: the impact of ignoring exposure measurement uncertainty. *Stat Med*. 1998;17(12):1375–1389.
16. Gustafson P. *Measurement Error and Misclassification in Statistics and Epidemiology*. New York, NY: Chapman and Hall; 2004.
17. Stamey J, Gerlach R. Bayesian sample size determination for case-control studies with misclassification. *Comput Stat Data Anal*. 2007;51(6):2982–2992.
18. Gustafson P, Le ND, Saskin R. Case-control analysis with partial knowledge of exposure misclassification probabilities. *Biometrics*. 2001;57(2):598–609.
19. Ben-Shoshan M, Kagan R, Alizadehfah R, et al. Is the prevalence of peanut allergy increasing? A five-year follow-up study on the prevalence of peanut allergy in primary school children in Montreal. *J Allergy Clin Immunol*. 2009;123(4):783–788.
20. Gelfand AE, Wang F. A simulation-based approach to Bayesian sample size determination for performance under a given model and for separating models. *Stat Sci*. 2002;17(2):193–208.
21. Chen M-H, Shao Q-M. Monte Carlo estimation of Bayesian credible and HPD intervals. *J Comput Graph Stat*. 1999;8(1):69–92.
1. Sample M_1 random values from the joint prior density of $p_1, p_2, S_1, S_2, C_1, C_2$. This involves selecting values from a β density for each parameter.
2. For each of the M_1 sets of parameters sampled in step 1, use equations (1) and (2) to calculate the probabilities of falling into each of the 4 cells defined by Table 1.
3. Select a tentative value for the sample size N , keeping in mind that a case-to-control ratio other than $g = 1$ may be selected. For each of the M_1 random situations, draw M_2 random 2-by-2 tables a, b, c, and d, by using the probabilities calculated in step 2. This is equivalent to sampling data from the marginal distribution of the data. In practice, $M_2 = 1$ is sufficient.
4. For each of the $M_1 \times M_2$ data sets, run the Gibbs sampler algorithm via WinBUGS (15) to derive samples from the posterior densities for p_1 and p_2 . By using these values, calculate odds ratio = $(p_1/(1 - p_1))/(p_2/(1 - p_2))$, a sample from the posterior density of the odds ratio adjusted for misclassification error.
5. Use the method of Chen and Shao (21) to calculate an HPD interval from each posterior sample from step 4, and hence calculate the length or posterior probability of each sample, as required by the chosen criterion. This method assumes unimodality of the posterior density, which, in our experience, is satisfied. If not, symmetric intervals can be substituted.
6. To implement the ACC criterion, compare the average coverage of HPD intervals of length l to the predetermined value of the coverage $1 - \alpha$. If the average coverage is greater (smaller) than the desired $1 - \alpha$, return to step 1 using a smaller (greater) sample size N . Continue until the criterion is met, using, for example, a bisectional search or model-based strategy to select the next N . A model-based strategy can use the pairs of N and average coverage values to create a fitted curve to predict the most likely value of N required, refining the model after each step. In practice, we have found a model of the form

$$\text{avg cov} = \alpha + \beta \Phi \left(\frac{\log(N) - \mu}{\sigma} \right)$$

fits the data well, where α and β are regression parameters to be estimated, μ and σ are measures of central tendency and spread of the logarithms of the sample sizes selected, respectively, and $\Phi(\cdot)$ is the cumulative normal density. To implement the ALC, compare the average length of the HPD intervals with fixed coverage $1 - \alpha$ by using a similar search strategy as for the ACC until the desired average length is attained. Finally, to implement the MWOC, for each N we must compare the proportion of samples that satisfy both the desired length and coverage, stopping when the proportion matches the desired $1 - \gamma$.

The ratio of controls to cases is another important design choice. Sample sizes can be calculated across a range of values for g , selecting the value that leads to the smallest sample size that is feasible in a given study, considering the availability of cases and controls.

APPENDIX

We now present the details of our numerical algorithm that determines the required sample sizes for estimating odds ratios in the presence of exposure misclassification. The algorithm is illustrated for the most general case in which $S_1 \neq S_2$ and $C_1 \neq C_2$, but similar steps can be followed when the exposure classification probabilities are equal within case and control populations.