Estimating and contextualizing the attenuation of odds ratios due to non-collapsibility

Stephen Burgess

Department of Public Health & Primary Care, University of Cambridge *

September 26, 2014

Short title: Attenuation of odds ratios due to non-collapsibility.

Keywords: odds ratios, non-collapsibility, confounding, case-control sampling.

*Correspondence to Stephen Burgess. Address: Department of Public Health & Primary Care, Strangeways Research Laboratory, 2 Worts Causeway, Cambridge, CB1 8RN, UK. Telephone: +44 1223 748651. Email: sb452@medschl.cam.ac.uk.

Abstract

The odds ratio is a measure commonly used for expressing the association between an exposure and a binary outcome. A feature of the odds ratio is that its value depends on the choice of the distribution over which the probabilities in the odds ratio are evaluated. In particular, this means that an odds ratio conditional on a covariate may have a different value from an odds ratio marginal on the covariate, even if the covariate is not associated with the exposure (not a confounder). We define the individual and population odds ratios as the ratio of the odds of the outcome for a unit increase in the exposure respectively for an individual in the population, and for the whole population, in which case the odds are averaged across the population. The attenuation of conditional, marginal and population odds ratios from the individual odds ratio is demonstrated in a realistic simulation exercise. The degree of attenuation differs in the whole population and in a case-control sample, and the property of invariance to outcome-dependent sampling is only true for the individual odds ratio. The relevance of the non-collapsibility of odds ratios in a range of methodological areas is discussed. (200 words)

1 Introduction

Non-collapsibility is a phenomenon by which some measures of association, such as odds ratios, vary depending on the choice of covariate adjustment even if the covariate is not associated with the exposure. It has some similarities with confounding, although the two are distinct concepts and differ in a number of ways. Confounding refers to the lack of exchangeability between population groups with different levels of the exposure due to inherent differences in their risk profiles [Greenland and Robins, 1986]. It often results from common causes of the exposure and outcome known as confounders [Miettinen and Cook, 1981]. Bias from confounding can be in any direction, and can lead to an observational association which is in the opposite direction to the true underlying causal association [Davey Smith and Ebrahim, 2002]. In contrast, when there are no confounders, non-collapsibility may still alter the magnitude of an association, although it will not change its direction [Gail et al., 1984]. Non-collapsibility merely requires some source of heterogeneity of risk between individuals or strata in the population, which could be due to the exposure itself.

Non-collapsibility means that, for a population containing strata with heterogeneous levels of risk and in the absence of confounding, the measure of association takes a different value when averaged across the strata of the population as opposed to when considered in an individual strata, even if the measure is identical in each of the strata (unless the probability of the outcome is identically distributed in each strata). A formal definition is that a measure of association is collapsible if, when it is constant across the strata of a covariate (which is distributed independently of the exposure, and is therefore not a confounder), this constant value equals the value obtained from the overall (marginal) analysis. Non-collapsibility is the absence of this property [Greenland et al., 1999].

An (extreme) example of non-collapsibility is given in Table 1. The odds ratio for the exposure is approximately 2 in both strata of the population, but is approximately

1 in a sample consisting of equal numbers of individuals from each stratum. Implicit in this calculation is the assumption that exposure status is not associated with the stratum variable, so that the proportions of exposed and non-exposed individuals are the same in both strata. Stratum membership is therefore not a confounder. If the proportions in each strata were different, then the direction of the association may change on conditioning on stratum membership, a situation known as Simpson’s paradox [Simpson, 1951; Pearl, 2000].

Table 1: Illustrative example of the non-collapsibility of the odds ratio

	Probability of event		Odds ratio
	Non-exposed	Exposed	
Stratum A	0.01	0.02	≈ 2
Stratum B	0.98	0.99	≈ 2
Overall	0.495	0.505	≈ 1

Probabilities of an outcome in two strata with odds ratios conditional on stratum membership, and in an overall group containing equal numbers from both strata with odds ratio marginal on stratum membership

A conditional odds ratio, calculated conditionally on the value of measured covariates (such as the odds ratio conditional on stratum membership in the above example), will typically differ from marginal odds ratio, calculated marginally across the distributions of the covariates (the overall odds ratio). In a logistic regression, this means that the asymptotic value of the regression coefficient for an exposure will change when adjustment is made for a variable associated with the outcome, even if the variable is not a confounder of the exposure–outcome association. (A notable exception is under the null, where all odds ratios coincide at 1.)

Although non-collapsibility is a problem which has been recognized theoretically, little advice is available on the extent to which non-collapsibility may affect odds ratio estimates in practice. As non-collapsibility affects the magnitude of estimates, rather than their direction, it is a less important problem for applied researchers

than confounding. However, in many cases an investigator may want to compare odds ratios for which different levels of adjustment for covariates has been made to assess the compatibility of estimates, or odds ratios calculated using different analysis methods which target differ odds ratio parameters [Janes et al., 2010].

As a motivating example, if a coefficient in a logistic regression analysis increases on adjustment for a covariate from, say, 0.2 to 0.5, can this magnitude of change be accounted for solely by non-collapsibility, or does it reflect that the covariate is an important confounding variable? Alternatively, if an odds ratio of 2.0 is estimated from a logistic regression analysis adjusting for multiple covariates, but an instrumental variable method gives an odds ratio estimate of 1.8, are these estimates compatible?

2 Methods

In this section, we initially define individual, conditional, marginal, and population odds ratios, and provide methods for the estimation of these parameters. We then outline a simulation study to investigate the difference between odds ratio parameters in a range of realistic scenarios.

2.1 Defining odds ratio parameters

We consider the odds ratio of an outcome Y for an exposure X . We assume that the outcome is binary and the probability of the outcome depends only on the exposure X , a known covariate C and an unknown covariate U .

The odds ratio of Y for $X = 1$ versus $X = 0$ is the odds of the outcome at the observed value of the exposure 1 divided by the odds at the value 0. If the exposure is binary, then this represents the ratio of the odds when the exposure is present to the odds when it is absent. We refer to this odds ratio as a marginal odds ratio (MOR), as the probabilities are calculated marginal across strata of the covariate C (and the

unmeasured covariate U).

$$\text{MOR}(1, 0) = \frac{\mathbb{P}(Y = 1|X = 1)}{\mathbb{P}(Y = 0|X = 1)} \bigg/ \frac{\mathbb{P}(Y = 1|X = 0)}{\mathbb{P}(Y = 0|X = 0)} \quad (1)$$

This can be written more concisely as:

$$\text{MOR}(1, 0) = \frac{\text{odds}(Y|X = 1)}{\text{odds}(Y|X = 0)} \quad (2)$$

where $\text{odds}(Y) = \mathbb{P}(Y = 1)/\mathbb{P}(Y = 0)$.

If the exposure is continuous, then we can consider the marginal odds ratio of Y for $X = x_1$ versus $X = x_0$, or the marginal odds ratio of Y for $X = x + 1$ versus $X = x$ (the marginal odds ratio for a unit increase in the exposure), where x_1 , x_0 and x are arbitrarily chosen values of the exposure.

$$\text{MOR}(x_1, x_0) = \frac{\text{odds}(Y|X = x_1)}{\text{odds}(Y|X = x_0)} \quad (3)$$

$$\text{MOR}(x + 1, x) = \frac{\text{odds}(Y|X = x + 1)}{\text{odds}(Y|X = x)} \quad (4)$$

Alternatively, we can consider a conditional odds ratio (COR), where the probabilities are calculated conditional the variable C . For example:

$$\text{COR}(x + 1, x; C = c) = \frac{\mathbb{P}(Y = 1|X = x + 1, C = c)}{\mathbb{P}(Y = 0|X = x + 1, C = c)} \bigg/ \frac{\mathbb{P}(Y = 1|X = x, C = c)}{\mathbb{P}(Y = 0|X = x, C = c)} \quad (5)$$

represents the odds ratio conditional on $C = c$ of Y for $X = x + 1$ versus $X = x$. Although in this paper we refer to the MOR and COR as if they are uniquely defined, in practice there is a spectrum of conditional odds ratios depending on the choice of covariate adjustment.

In a logistic-linear model, where the logit-transformed probability of the outcome is a linear function of X and C (and not a function of U), the COR does not depend on the values of X or C (see Section 2.2). This is the form of model usually assumed in

a logistic regression analysis. In contrast, even in the logistic-linear case, the marginal odds ratio $\text{MOR}(x+1, x)$ depends on the chosen value of x , and additionally on the distribution of C . In order to obtain a single marginal measure of association with a continuous exposure for a given dataset, we additionally define a population odds ratio (POR) as the ratio between the odds for the given population with the observational distribution of the exposure and covariates and the odds for a population identical to the given population except with the exposure increased by one unit for all individuals. We write this as:

$$\text{POR} = \frac{\text{odds}(Y|\sigma_X = 1)}{\text{odds}(Y|\sigma_X = 0)} \quad (6)$$

where the $\sigma_X = \delta$ notation represents a shift in the exposure distribution by δ units. Formally, we have $(X|\sigma_X = \delta) = (X + \delta|\sigma_X = 0)$ and $(X|\sigma_X = 0) = X$. A POR is a population-based measure of association, and as such is of interest to policymakers. Other population-averaged odds ratios could be evaluated; we chose to define the POR in terms of a unit increase in the exposure distribution as it is most similar to the other odds ratio parameters considered. In a similar way, we can think of an individual odds ratio (IOR) as the odds ratio for an individual i , representing the ratio between the odds of the outcome at different levels of the exposure:

$$\text{IOR}(i) = \frac{\text{odds}(Y_i|\sigma_X = 1)}{\text{odds}(Y_i|\sigma_X = 0)} \quad (7)$$

The difference between the POR and IOR is the space over which the probabilities are evaluated. For the POR, this space is the whole population; for the IOR, it is the individual in question. The IOR is of theoretical rather than practical interest, as neither the probability nor the odds of an outcome cannot be estimated for an individual other than by assuming that all explainable variation in the outcome can be expressed as a model using measured covariates. If this model is correct, then the IOR is equal to the COR.

The IOR depends on the values of X , C and U , and so will typically differ for

Table 2: Odds ratio parameters

		Function of ...	Averaged across ...
IOR	Individual odds ratio	X, C, U	None
COR	Conditional odds ratio	X, C	U
MOR	Marginal odds ratio	X	C, U
POR	Population odds ratio	None	X, C, U

List of abbreviations for odds ratios used in this paper together with variables the odds ratios are function of, and variables the odds ratios are averaged across

each individual in the population outside of a logistic-linear model, whereas the POR depends on the distributions of the same variables, but represents an measure averaged across the whole population. In comparison, the COR depends on the values of X and C , but is averaged across the distribution of U , and the MOR depends on the value of X , but is averaged across the distribution of C and U . A list of the abbreviations used for odds ratios in this paper is provided in Table 2.

The POR relates to the MOR (in particular, the $\text{MOR}(\bar{X} + 1, \bar{X})$ where $\bar{X}$ is the mean value of the exposure) in a similar way to the relationship between the average marginal effect and the marginal effect at a representative value (in particular, the mean value of the exposure). The marginal effect (or partial effect) is a measure of association used in the field of econometrics which is calculated as the gradient of change in the outcome for a change in the exposure conditional on other variables staying constant; formally, this is expressed as a partial derivative of the conditional mean of the outcome with respect to the exposure evaluated at a given value of the exposure [Cameron and Trivedi, 2009]. For a non-linear model (such as a logistic model), the marginal effect will depend on the value of the exposure, and hence the marginal effect at the mean (of the exposure) or the average marginal effect are often cited as a single measure of association to summarize the marginal effect of the exposure in the population [Mogstad and Wiswall, 2010]. The average marginal effect

is calculated by averaging the ‘marginal effect at a representative value’ across the distribution of the exposure. The POR is similar, except that the odds functions in the numerator and the denominator of the odds ratio expression are averaged separately, so that the POR is not an average of the $\text{MOR}(X_i + 1, X_i)$ across the distribution of the exposure X .

2.2 Odds ratios in a logistic-linear model

To further demonstrate the differences between odds ratio parameters, we consider expressions of the parameters under a particular mathematical model for the probability of an outcome, a logistic-linear model. We assume that the logit-transformed probability of an outcome π is a linear function of the variables X , C and U :

$$\text{logit}(\pi) = \beta_0 + \beta_1 X + \beta_2 C + \beta_3 U \quad (8)$$

The IOR for a unit increase in the exposure from $X = x$ to $X = x + 1$ for an individual with $C = c$ and $U = u$ is:

$$\begin{aligned} \text{IOR}(x + 1, x; c, u) &= \frac{\frac{\text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u)}{1 - \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u)}}{\frac{\text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u)}{1 - \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u)}} \\ &= \frac{\exp(\beta_0 + \beta_1(x + 1) + \beta_2 c + \beta_3 u)}{\exp(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u)} \\ &= \exp(\beta_1) \end{aligned} \quad (9)$$

This is independent of the values of $X = x$, $C = c$ and $U = u$.

The COR for a unit increase in the exposure from $X = x$ to $X = x + 1$ conditional on $C = c$ is:

$$\text{COR}(x + 1, x; c) = \frac{\frac{\int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(u) du}{1 - \int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(u) du}}{\frac{\int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(u) du}{1 - \int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(u) du}} \quad (10)$$

where $\text{expit}()$ is the inverse of the $\text{logit}()$ function, and the integrals are evaluated over $p(u) du$, the probability distribution of U . If the coefficient β_3 is zero, then the probabilities no longer depend on U , In this case, the COR simplifies to:

$$\text{COR}(x+1, x; c) = \frac{\exp(\beta_0 + \beta_1(x+1) + \beta_2 c)}{\exp(\beta_0 + \beta_1 x + \beta_2 c)} = \exp(\beta_1) \quad (11)$$

which is the same as the corresponding IOR.

The MOR for a unit increase in the exposure from $X = x$ to $X = x + 1$ is:

$$\text{MOR}(x+1, x) = \frac{\frac{\int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(c, u) dc du}{1 - \int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(c, u) dc du}}{\frac{\int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(c, u) dc du}{1 - \int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(c, u) dc du}} \quad (12)$$

where the integrals are evaluated over $p(c, u) dc du$, the joint probability distribution of C and U . Even if $\beta_3 = 0$, the MOR still depends on the value of $X = x$ and the distribution of C .

The POR for a unit increase in the exposure is:

$$\text{POR}(x+1, x) = \frac{\frac{\int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(x, c, u) dx dc du}{1 - \int \text{expit}(\beta_0 + \beta_1(x+1) + \beta_2 c + \beta_3 u) p(x, c, u) dx dc du}}{\frac{\int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(x, c, u) dx dc du}{1 - \int \text{expit}(\beta_0 + \beta_1 x + \beta_2 c + \beta_3 u) p(x, c, u) dx dc du}} \quad (13)$$

where the integrals are evaluated over $p(x, c, u) dx dc du$, the joint probability distribution of X , C and U . Again, even if $\beta_2 = \beta_3 = 0$ (that is, the probability of an outcome is a function of X and no other variables), the POR still depends on the distribution of X .

2.3 Simulation study

In order to investigate the numerical differences between individual, conditional, marginal and population odds ratios in practice, we simulate 1000 datasets on 100 000

individuals from the following data-generating model:

$$\begin{aligned}
X &\sim \mathcal{N}(0, 1) \\
C &\sim \mathcal{N}(0, \sigma_C^2) \\
U &\sim \mathcal{N}(0, \sigma_U^2) \\
Y &\sim \text{Bernoulli}(\pi) \\
\text{logit}(\pi) &= \beta_0 + \beta_1 X + C + U
\end{aligned} \tag{14}$$

To ensure that the differences between odds ratios are due to non-collapsibility not due to confounding, we simulate data with no confounding. We vary parameters corresponding to the prevalence of the outcome ($\beta_0 = -3, -4, -5$), the magnitude of the effect of the exposure on the outcome ($\beta_1 = 0.2, 0.4, 0.6, 0.8$), and the variances of the distributions of the measured and unmeasured confounders ($\sigma_C^2, \sigma_U^2 = 0, 0.5, 1$). Logistic regression analysis adjusting for a number of covariates in a real-data example gave a variance of the estimated linear predictor $\text{logit}(\hat{\pi})$ of 1.90 in a case-control study and 0.79 in a cohort study, although the variance of the true linear predictor may be greater due to the presence of unmeasured covariates [Burgess, 2012]. If $\beta_0 = -3$ and the variance of the linear predictor is 1, then the 2.5th and 97.5th percentiles of the probability of an outcome are 0.007 and 0.26 respectively; if the variance is 3, then the percentiles are 0.002 and 0.60. If $\beta_0 = -5$, the corresponding values are 0.001 and 0.05 (variance 1), and 0.0002 and 0.17 (variance 3).

The conditional odds ratio is estimated by logistic regression of Y on X and C . Marginal and population odds ratios are obtained by taking averaging over the probabilities for the original population and for a simulated population identical to the original but with different values for the exposure. Odds ratios are calculated

using these probabilities averaged across the empirical distributions of X , U and C :

$$\widehat{\text{MOR}}(1, 0) = \frac{\sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 \times 1 + c_i + u_i)}{1 - \sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 \times 1 + c_i + u_i)} \bigg/ \frac{\sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 \times 0 + c_i + u_i)}{1 - \sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 \times 0 + c_i + u_i)} \quad (15)$$

$$\widehat{\text{POR}} = \frac{\sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1(x_i + 1) + c_i + u_i)}{1 - \sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1(x_i + 1) + c_i + u_i)} \bigg/ \frac{\sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 x_i + c_i + u_i)}{1 - \sum_i \frac{1}{n} \text{expit}(\beta_0 + \beta_1 x_i + c_i + u_i)} \quad (16)$$

This is a Monte Carlo approach to calculating the probabilities in the marginal and population odds ratios, and the expressions $\widehat{\text{MOR}}(1, 0)$ and $\widehat{\text{POR}}$ will tend towards the parameters $\text{MOR}(1,0)$ and POR as the sample size increases. Although the exposure X is continuous, the evaluation of $\text{MOR}(1,0)$ does not make use of the distribution of X , and so is valid for an arbitrarily-distributed exposure, including a binary exposure. If the exposure were binary, $\text{MOR}(1,0)$ would equal POR , and so this simulation analysis also addresses the attenuation of the POR with a binary exposure.

In addition to odds ratios estimated for the whole population, we also calculate odds ratios for a case-control sample. The case-control sample contains all the cases in the population (those individuals with $Y = 1$) and includes controls (those individuals with $Y = 0$) with probability determined by the proportion of case participants in the simulated dataset, such that the overall sample contains on average equal proportions of cases and controls.

2.4 Interaction between exposure and covariate

The above data-generating model does not allow for an interaction between the exposure and a covariate. If there is an interaction, then the individual and conditional odds ratios are non-trivial functions of the covariate, and there is no single individual or conditional odds ratio. The marginal and population odds ratios are still well-defined quantities. A conditional odds ratio could be estimated within strata of a

categorical confounder, or else an interaction term can be incorporated into a logistic regression model.

To illustrate the behaviour of the marginal and population odds ratios when there is an interaction between the exposure and the measured covariate, as well as the coefficients from logistic regression analyses with and without an interaction term, we perform a further simulation study with the following data-generating model:

$$\begin{aligned}
X &\sim \mathcal{N}(0, 1) \\
C &\sim \mathcal{N}(0, \sigma_C^2) \\
U &\sim \mathcal{N}(0, \sigma_U^2) \\
Y &\sim \text{Bernoulli}(\pi) \\
\text{logit}(\pi) &= \beta_0 + \beta_1 X + C + \beta_2 X C + U
\end{aligned} \tag{17}$$

All parameters are taken as in the previous simulation study, except that $\beta_0 = -3$ throughout, a more limited set of values for β_1 is considered ($\beta_1 = 0.4, 0.8$) and the interaction parameter β_2 is set to three values: 0, 0.2 and 0.4.

3 Results

Results from the simulation study with no interaction term are presented in Tables 3 and 4, and Figures 1–3. We give the conditional log odds ratio (CLOR), the marginal log odds ratio for $X = 1$ versus $X = 0$ (MLOR(1,0)), and the population log odds ratio (PLOR). The coefficient β_1 in each case is the individual log odds ratio (ILOR). The results presented are the mean estimates across simulations. Monte Carlo error in the results due to the limited number of simulations is around 0.001 for the CLOR, and is negligible for the MLOR(1,0) and PLOR.

Considering estimates for the whole population, we see that the PLOR is attenuated from the CLOR even when there is no variation in the risk of the outcome

from covariates, whether measured or unmeasured. In general, the attenuation of the PLOR from the ILOR is greater than that of the CLOR or MLOR(1,0). Attenuation of all three parameters is greater when the outcome is more common, when the effect of the exposure is greater, and when the variances of the covariates are greater. The magnitude of attenuation of the PLOR from the ILOR when the total variance of the covariates (measured and unmeasured) is one ranges from 2% for $\beta_0 = -5, \beta_1 = 0.2$, to 15% for $\beta_0 = -3, \beta_1 = 0.8$. The CLOR is approximately equal to the ILOR when there are no unmeasured covariates ($\sigma_U^2 = 0$). Increasing attenuation of the CLOR can be observed as more of the variation in the covariates is unmeasured rather than measured, by comparing scenarios where the total variance of the covariates is fixed, but the division between measured and unmeasured covariates changes. For example, in the first row of Table 3, for $\beta_0 = -3, \beta_1 = 0.2$ the CLOR is 0.201 when $\sigma_C^2 = 1, \sigma_U^2 = 0$, 0.197 when $\sigma_C^2 = \sigma_U^2 = 0.5$, and 0.191 when $\sigma_C^2 = 0, \sigma_U^2 = 1$. The MLOR(1,0) and PLOR are similar for $\beta_1 = 0.2$, but diverge as β_1 increases.

Comparing estimates from the case-control sample to those from the whole population, we see that the CLOR is slightly less attenuated from the ILOR, but the MLOR(1,0) and the PLOR are more attenuated. This is because of the greater heterogeneity of risk between individuals in the case-control sample, leading to a greater discriminatory role of the measured covariate into high and low risk groups, which is captured by the CLOR, but not by other measures. The mean estimates of the CLOR are similar under case-control sampling and in the population; this is a major motivation of the use of the odds ratio as a measure of association. The MLOR(1,0) and PLOR are substantially different in the case-control sample and in the population. Although the PLOR for the case-control sample does not have a natural interpretation as a parameter of interest, it may be relevant to some methods for the estimation of odds ratios in case-control samples. The maximum attenuation of the CLOR from the ILOR is 10%; while the maximum attenuation of the MLOR(1,0) and PLOR from the ILOR are 35% and 44% respectively.

Table 3: Simulation results using entire population

			$\sigma_U^2 = 0$			$\sigma_U^2 = 0.5$			$\sigma_U^2 = 1$		
$\sigma_C^2 =$			0	0.5	1	0	0.5	1	0	0.5	1
$\beta_1 = 0.2$	$\beta_0 = -5$	CLOR	0.201	0.200	0.201	0.198	0.197	0.197	0.195	0.195	0.191
		MLOR(1,0)	0.200	0.199	0.196	0.199	0.196	0.192	0.196	0.192	0.187
		PLOR	0.200	0.199	0.196	0.199	0.196	0.192	0.196	0.192	0.187
	$\beta_0 = -4$	CLOR	0.199	0.200	0.200	0.197	0.195	0.195	0.192	0.189	0.188
		MLOR(1,0)	0.200	0.197	0.192	0.197	0.192	0.185	0.192	0.185	0.177
		PLOR	0.200	0.197	0.191	0.197	0.191	0.185	0.191	0.185	0.177
	$\beta_0 = -3$	CLOR	0.199	0.200	0.200	0.193	0.192	0.190	0.185	0.182	0.180
		MLOR(1,0)	0.200	0.193	0.185	0.193	0.185	0.175	0.185	0.175	0.166
		PLOR	0.200	0.193	0.184	0.193	0.184	0.174	0.184	0.174	0.165
$\beta_1 = 0.4$	$\beta_0 = -5$	CLOR	0.399	0.399	0.402	0.396	0.396	0.392	0.392	0.387	0.382
		MLOR(1,0)	0.400	0.397	0.392	0.397	0.392	0.384	0.392	0.384	0.372
		PLOR	0.399	0.396	0.390	0.396	0.390	0.380	0.390	0.380	0.368
	$\beta_0 = -4$	CLOR	0.401	0.400	0.400	0.394	0.391	0.387	0.383	0.377	0.371
		MLOR(1,0)	0.400	0.394	0.383	0.394	0.383	0.369	0.383	0.368	0.353
		PLOR	0.398	0.391	0.378	0.391	0.378	0.364	0.378	0.364	0.348
	$\beta_0 = -3$	CLOR	0.400	0.400	0.400	0.386	0.382	0.379	0.369	0.362	0.359
		MLOR(1,0)	0.400	0.386	0.367	0.386	0.367	0.348	0.367	0.348	0.330
		PLOR	0.396	0.380	0.361	0.380	0.361	0.342	0.361	0.342	0.324
$\beta_1 = 0.6$	$\beta_0 = -5$	CLOR	0.599	0.599	0.601	0.596	0.590	0.587	0.585	0.578	0.571
		MLOR(1,0)	0.600	0.596	0.587	0.596	0.587	0.574	0.587	0.574	0.556
		PLOR	0.597	0.590	0.578	0.590	0.578	0.561	0.578	0.561	0.542
	$\beta_0 = -4$	CLOR	0.600	0.600	0.600	0.588	0.582	0.578	0.571	0.563	0.555
		MLOR(1,0)	0.600	0.589	0.572	0.589	0.572	0.550	0.572	0.550	0.526
		PLOR	0.593	0.578	0.556	0.578	0.556	0.532	0.556	0.533	0.508
	$\beta_0 = -3$	CLOR	0.601	0.600	0.600	0.576	0.571	0.568	0.549	0.541	0.536
		MLOR(1,0)	0.600	0.577	0.548	0.577	0.548	0.519	0.548	0.519	0.491
		PLOR	0.584	0.557	0.527	0.557	0.527	0.499	0.527	0.499	0.473
$\beta_1 = 0.8$	$\beta_0 = -5$	CLOR	0.800	0.800	0.801	0.792	0.787	0.780	0.776	0.765	0.757
		MLOR(1,0)	0.800	0.794	0.781	0.794	0.781	0.762	0.781	0.762	0.738
		PLOR	0.791	0.777	0.756	0.776	0.756	0.730	0.756	0.730	0.702
	$\beta_0 = -4$	CLOR	0.800	0.800	0.799	0.781	0.773	0.768	0.754	0.743	0.734
		MLOR(1,0)	0.800	0.785	0.760	0.785	0.760	0.729	0.760	0.729	0.696
		PLOR	0.779	0.752	0.720	0.752	0.720	0.687	0.720	0.687	0.655
	$\beta_0 = -3$	CLOR	0.801	0.801	0.800	0.764	0.758	0.754	0.726	0.717	0.712
		MLOR(1,0)	0.800	0.767	0.727	0.767	0.727	0.687	0.727	0.687	0.650
		PLOR	0.757	0.716	0.677	0.716	0.677	0.640	0.677	0.640	0.608

Mean estimates of the conditional (CLOR), marginal (MLOR(1,0)) and population (PLOR) log odds ratios from simulation analysis when individual log odds ratio is β_1 using whole dataset

Table 4: Simulation results using case-control sample

			$\sigma_U^2 = 0$			$\sigma_U^2 = 0.5$			$\sigma_U^2 = 1$		
$\sigma_C^2 =$			0	0.5	1	0	0.5	1	0	0.5	1
$\beta_1 = 0.2$	$\beta_0 = -5$	CLOR	0.197	0.199	0.199	0.200	0.200	0.197	0.196	0.195	0.195
		MLOR(1,0)	0.199	0.196	0.189	0.196	0.189	0.178	0.189	0.178	0.164
		PLOR	0.198	0.196	0.188	0.196	0.188	0.177	0.188	0.177	0.163
	$\beta_0 = -4$	CLOR	0.200	0.201	0.201	0.198	0.196	0.195	0.192	0.192	0.188
		MLOR(1,0)	0.196	0.190	0.178	0.190	0.178	0.163	0.178	0.163	0.149
		PLOR	0.196	0.189	0.177	0.189	0.177	0.162	0.177	0.162	0.147
	$\beta_0 = -3$	CLOR	0.201	0.199	0.200	0.193	0.192	0.192	0.186	0.184	0.182
		MLOR(1,0)	0.191	0.179	0.162	0.179	0.162	0.147	0.162	0.147	0.134
		PLOR	0.190	0.177	0.161	0.177	0.161	0.145	0.161	0.145	0.132
$\beta_1 = 0.4$	$\beta_0 = -5$	CLOR	0.402	0.399	0.403	0.399	0.396	0.398	0.390	0.389	0.388
		MLOR(1,0)	0.397	0.391	0.377	0.391	0.377	0.353	0.377	0.353	0.326
		PLOR	0.396	0.387	0.369	0.387	0.369	0.343	0.369	0.343	0.315
	$\beta_0 = -4$	CLOR	0.400	0.399	0.401	0.395	0.393	0.390	0.384	0.382	0.379
		MLOR(1,0)	0.392	0.378	0.354	0.378	0.354	0.324	0.354	0.324	0.296
		PLOR	0.388	0.371	0.343	0.371	0.343	0.312	0.343	0.312	0.283
	$\beta_0 = -3$	CLOR	0.401	0.400	0.399	0.387	0.384	0.383	0.369	0.366	0.363
		MLOR(1,0)	0.379	0.354	0.322	0.354	0.322	0.291	0.322	0.291	0.266
		PLOR	0.372	0.342	0.308	0.342	0.308	0.278	0.308	0.278	0.254
$\beta_1 = 0.6$	$\beta_0 = -5$	CLOR	0.600	0.603	0.604	0.596	0.596	0.595	0.589	0.586	0.580
		MLOR(1,0)	0.595	0.585	0.562	0.585	0.562	0.526	0.562	0.526	0.486
		PLOR	0.589	0.569	0.535	0.570	0.535	0.492	0.535	0.492	0.449
	$\beta_0 = -4$	CLOR	0.600	0.602	0.600	0.591	0.588	0.591	0.575	0.569	0.575
		MLOR(1,0)	0.586	0.564	0.526	0.564	0.526	0.482	0.526	0.482	0.439
		PLOR	0.572	0.536	0.489	0.536	0.489	0.443	0.489	0.443	0.403
	$\beta_0 = -3$	CLOR	0.601	0.601	0.599	0.579	0.574	0.574	0.553	0.547	0.543
		MLOR(1,0)	0.566	0.527	0.478	0.527	0.478	0.433	0.478	0.433	0.396
		PLOR	0.537	0.486	0.436	0.486	0.436	0.394	0.436	0.394	0.360
$\beta_1 = 0.8$	$\beta_0 = -5$	CLOR	0.802	0.802	0.801	0.794	0.790	0.787	0.782	0.778	0.771
		MLOR(1,0)	0.792	0.778	0.746	0.778	0.746	0.697	0.746	0.697	0.642
		PLOR	0.770	0.730	0.674	0.730	0.674	0.614	0.674	0.615	0.560
	$\beta_0 = -4$	CLOR	0.801	0.800	0.801	0.783	0.781	0.776	0.761	0.754	0.749
		MLOR(1,0)	0.779	0.748	0.696	0.748	0.696	0.636	0.696	0.636	0.581
		PLOR	0.733	0.672	0.607	0.672	0.607	0.549	0.607	0.549	0.500
	$\beta_0 = -3$	CLOR	0.801	0.800	0.801	0.770	0.765	0.762	0.733	0.728	0.722
		MLOR(1,0)	0.751	0.696	0.631	0.696	0.631	0.572	0.631	0.572	0.524
		PLOR	0.669	0.599	0.538	0.599	0.538	0.488	0.538	0.488	0.449

Mean estimates of the conditional (CLOR), marginal (MLOR(1,0)) and population (PLOR) log odds ratios from simulation analysis when individual log odds ratio is β_1 using case-control sample

Figure 1: Mean estimates of the conditional (CLOR, circle), marginal (MLOR(1,0), triangle) and population (PLOR, cross) log odds ratios from simulation analysis when individual log odds ratio is β_1 (dashed line) using whole dataset (black) and case-control sample (grey) for $\beta_0 = -3$

Figure 2: Mean estimates of the conditional (CLOR, circle), marginal (MLOR(1,0), triangle) and population (PLOR, cross) log odds ratios from simulation analysis when individual log odds ratio is β_1 (dashed line) using whole dataset (black) and case-control sample (grey) for $\beta_0 = -4$

Figure 3: Mean estimates of the conditional (CLOR, circle), marginal (MLOR(1,0), triangle) and population (PLOR, cross) log odds ratios from simulation analysis when individual log odds ratio is β_1 (dashed line) using whole dataset (black) and case-control sample (grey) for $\beta_0 = -5$

3.1 Interaction between exposure and covariate

Results from the simulation study with an interaction term between the exposure and the measured covariate are presented in Table 5. We give the beta-coefficient for the exposure ($\hat{\beta}_X$) from logistic regression without an interaction term (logistic-1), the beta-coefficients for the exposure ($\hat{\beta}_X$) and the interaction between the exposure and measured covariate ($\hat{\beta}_{XC}$) from logistic regression with an interaction term (logistic-2), the marginal log odds ratio for $X = 1$ versus $X = 0$ (MLOR(1,0)), and the population log odds ratio (PLOR). All of these results are obtained in the whole population rather than the case-control sample. When $\sigma_C^2 = 0$, the covariate C is identically zero and there is no interaction term; this scenario is presented for consistency with the previous tables of results and for comparison.

The MLOR(1,0) and the PLOR are attenuated compared to the coefficient from a logistic regression without an interaction term. The degree of attenuation increases as the interaction term increases. However, the logistic coefficient also increases, so some degree of increased attenuation would be expected from Table 3. Comparing the attenuation at $\beta_1 = 0.4, \beta_2 = 0.4, \sigma_C^2 = 1, \sigma_U^2 = 1$ from 0.551 to 0.513 (MLOR(1,0)) or to 0.466 (PLOR) with the attenuation in the original simulation at $\beta_0 = -3, \beta_1 = 0.6, \sigma_C^2 = 1, \sigma_U^2 = 1$ from a CLOR of 0.536 to 0.491 (MLOR(1,0)) or to 0.473 (PLOR) suggests a similar degree of attenuation for the MLOR(1,0), but additional attenuation for the PLOR. The MLOR(1,0) and PLOR are in some cases greater than the coefficient from logistic regression with an interaction term when $\beta_2 > 0$; however, this is not surprising as the CLOR for the exposure is greater than β_1 when $C > 0$.

4 Relevance to applied research

In this section, we investigate a range of areas of applied research in which the non-collapsibility of the odds ratio has direct relevance. The list of topics is not exhaustive,

Table 5: Simulation results with interaction between exposure and covariate

			$\sigma_U^2 = 0$			$\sigma_U^2 = 0.5$			$\sigma_U^2 = 1$		
$\sigma_C^2 =$			0	0.5	1	0	0.5	1	0	0.5	1
$\beta_1 = 0.4$	$\beta_2 = 0$	$\hat{\beta}_X$ (logistic-1)	0.400	0.400	0.400	0.387	0.382	0.379	0.368	0.363	0.359
		$\hat{\beta}_X$ (logistic-2)	0.400	0.400	0.400	0.387	0.385	0.384	0.368	0.368	0.367
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.000	0.000	-	-0.009	-0.008	-	-0.014	-0.013
		MLOR(1,0)	0.400	0.386	0.367	0.386	0.367	0.348	0.367	0.348	0.330
		PLOR	0.396	0.380	0.361	0.380	0.361	0.342	0.361	0.342	0.324
	$\beta_2 = 0.2$	$\hat{\beta}_X$ (logistic-1)	0.401	0.480	0.538	0.386	0.450	0.497	0.369	0.420	0.460
		$\hat{\beta}_X$ (logistic-2)	0.401	0.400	0.400	0.386	0.384	0.384	0.369	0.365	0.364
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.200	0.200	-	0.178	0.177	-	0.161	0.161
		MLOR(1,0)	0.400	0.464	0.494	0.386	0.433	0.455	0.367	0.403	0.420
		PLOR	0.396	0.446	0.463	0.380	0.416	0.429	0.361	0.389	0.400
	$\beta_2 = 0.4$	$\hat{\beta}_X$ (logistic-1)	0.400	0.557	0.660	0.386	0.514	0.601	0.368	0.475	0.551
		$\hat{\beta}_X$ (logistic-2)	0.400	0.401	0.401	0.386	0.383	0.381	0.368	0.363	0.362
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.400	0.399	-	0.366	0.363	-	0.338	0.336
		MLOR(1,0)	0.400	0.549	0.621	0.386	0.504	0.563	0.367	0.464	0.513
		PLOR	0.396	0.509	0.545	0.380	0.470	0.503	0.361	0.435	0.466
$\beta_1 = 0.8$	$\beta_2 = 0$	$\hat{\beta}_X$ (logistic-1)	0.800	0.800	0.800	0.765	0.759	0.754	0.724	0.717	0.712
		$\hat{\beta}_X$ (logistic-2)	0.800	0.800	0.800	0.765	0.764	0.762	0.724	0.725	0.723
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.001	0.000	-	-0.017	-0.013	-	-0.025	-0.021
		MLOR(1,0)	0.800	0.767	0.727	0.767	0.727	0.687	0.727	0.687	0.650
		PLOR	0.757	0.716	0.677	0.716	0.677	0.640	0.677	0.640	0.608
	$\beta_2 = 0.2$	$\hat{\beta}_X$ (logistic-1)	0.800	0.868	0.919	0.764	0.815	0.856	0.726	0.767	0.800
		$\hat{\beta}_X$ (logistic-2)	0.800	0.800	0.801	0.764	0.762	0.761	0.726	0.724	0.722
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.200	0.199	-	0.170	0.172	-	0.148	0.152
		MLOR(1,0)	0.800	0.836	0.836	0.767	0.784	0.777	0.727	0.734	0.726
		PLOR	0.757	0.747	0.726	0.716	0.704	0.685	0.677	0.664	0.648
	$\beta_2 = 0.4$	$\hat{\beta}_X$ (logistic-1)	0.800	0.935	1.025	0.765	0.871	0.948	0.725	0.814	0.880
		$\hat{\beta}_X$ (logistic-2)	0.800	0.801	0.800	0.765	0.760	0.760	0.725	0.722	0.720
		$\hat{\beta}_{XC}$ (logistic-2)	-	0.400	0.401	-	0.356	0.358	-	0.324	0.326
		MLOR(1,0)	0.800	0.910	0.943	0.767	0.845	0.868	0.727	0.785	0.804
		PLOR	0.757	0.775	0.764	0.716	0.729	0.721	0.677	0.687	0.682

Mean estimates of the beta-coefficient for the exposure ($\hat{\beta}_X$) from logistic regression without an interaction term (logistic-1), the beta-coefficients for the exposure ($\hat{\beta}_X$) and the interaction between the exposure and measured covariate ($\hat{\beta}_{XC}$) from logistic regression with an interaction term (logistic-2), and the marginal (MLOR(1,0)) and population (PLOR) log odds ratios from simulation analysis using whole dataset

but gives a sense of the wide variety of subject areas for which the question of the interpretation of odds ratio estimates is relevant.

4.1 Randomized trials and observational studies

In a randomized controlled trial (RCT), adjustment for covariates is not generally necessary, as the distribution of all covariates would be thought to be equal in expectation in each arm of the trial [Hauck et al., 1998; Steyerberg et al., 2000]. Additionally, unlike in linear regression, adjustment for covariates in logistic regression does not uniformly translate into estimators having reduced standard errors [Ford et al., 1995]. Choice of covariate adjustment should therefore be made based on the particular odds ratio desired to be estimated [Ford and Norrie, 2002].

In an observational study, adjustment for covariates is necessary, as several of the covariates may be important confounders of the exposure–outcome association. If the assumption that there is no unmeasured confounding is satisfied and most of the explainable variation in the outcome model is accounted for, then the conditional odds ratio from a logistic analysis may be close to an individual odds ratio. However, if there is substantial heterogeneity between individuals beyond that explained by the measured covariates, then the conditional odds ratio will be attenuated for the individual odds ratio. In both cases, the conditional odds ratio will be an overestimate of the average effect in the population of a change in the exposure, which is the population odds ratio. A marginal odds ratio is estimated in an analysis accounting for covariates using inverse probability weighting, although stabilization of the weights may change the interpretation of the estimate [Kaufman, 2010]. A marginal odds ratio can be obtained from a conditional odds ratio by integration across the distribution of the measured covariates [Zhang, 2008]. If the exposure is continuous, then a population odds ratio can be obtained by further integrating across the distribution of the exposure.

4.2 Longitudinal studies

Longitudinal studies comprise data on participants at multiple timepoints. Two approaches to the analysis of longitudinal data are random-effects and marginal models [Diggle et al., 2002]. In a random-effects model, regression coefficients are assumed to differ between individuals. A random effect is a unmeasured variable assumed to capture the specific characteristics of an individual. It is estimated as part of the analysis for each individual using the repeated measurements. The coefficients in a random-effects model represent subject-specific estimates, which are conditional on the value of the random effect variable. In a marginal model, the marginal expectation of the outcome is modelled as a function of covariates. Within-individual variation is modelled separately. The marginal expectation of the outcome is the same as that estimated in a cross-sectional study, and ignores individual-level heterogeneity. The coefficients in a marginal model represent population-averaged parameters, which are marginal across the distribution of the random effect variable. If the parameter of interest is an odds ratio, then the subject-specific (conditional) and population-averaged (marginal) estimates will typically differ [Zeger et al., 1988].

4.3 Matched and unmatched analyses

In a matched study design, each participant is paired to another who has similar characteristics according to matching variables chosen by the investigator. A matched analysis compared individuals within these pairs, whereas an unmatched analysis compared individuals ignoring the pairing structure of the data [Duffy et al., 1989]. The matched analysis estimates a parameter conditional on the matching variables, while the unmatched analysis estimates a parameter marginal on the same variables. If a relative risk parameter is estimated, then these parameters may be equal, but if an odds ratio parameter is estimated, then they will only be equal at the null [Zou, 2007].

4.4 Mediation

The difference between a regression coefficient with and without adjustment for a covariate has been interpreted as part of the evidence for the mediating role of the covariate in the association between the exposure and outcome [Judd and Kenny, 1981]. In a logistic model of association, such differences may reflect non-collapsibility rather than mediation; although non-collapsibility would generally increase the magnitude of the coefficient, whereas mediation would generally decrease it. It may be that partial mediation in a logistic model is masked by non-collapsibility, as the decrease in the regression coefficient due to mediation and the increase due to non-collapsibility may numerically cancel out. Methods from the causal inference approach to mediation analysis for odds ratios have been discussed [VanderWeele and Vansteelandt, 2010]; however, these invoke the ‘rare disease assumption’ in their derivation [Greenland and Thomas, 1982] and so sidestep the issue of non-collapsibility.

4.5 Meta-analysis

The conditional odds ratio in different studies will vary when there are unmeasured covariates even if the individual odds ratio in each study is the same [Mood, 2010]. This is because the distribution of the covariates will vary from study to study. This means that a fixed-effects meta-analysis model, where the same parameter is assumed to be estimated in each study, requires the assumption not only that the individual odds ratio is the same in each study, but also either that there are no unmeasured covariates, or that all covariates, measured and unmeasured, are identically distributed in each study. However, unless the distributions of covariates in studies substantially vary, the difference between conditional odds ratios due to non-collapsibility is likely to be small. Although a random-effects meta-analysis model, where the estimated parameter is allowed to vary between studies, seems more reasonable, it is unclear what the precise interpretation of the pooled odds ratio may be, other than a weighted

average of conditional odds ratios from each study.

4.6 Instrumental variable analysis

The method of instrumental variables was developed to enable asymptotically unbiased estimation of regression parameters in the absence of data on confounders. When the outcome is continuous and the relationship between the exposure and outcome is linear, instrumental variable methods consistently estimate the association between the exposure and outcome [Didelez et al., 2010b]. For a non-collapsible measure of association (such as, but not limited to, the odds ratio) or more generally for a non-linear model, instrumental variable estimates have been labelled “forbidden regressions” [Hausman, 1983] as they do not consistently estimate the parameter for the association between the exposure and outcome in the data-generating model. In the case of an odds ratio and a linear-logistic model, this is the individual odds ratio [Palmer et al., 2008]. However, inconsistency of the parameter estimated by instrumental variable methods is not a manifestation of bias, but rather a symptom of non-collapsibility [Vansteelandt et al., 2011]. If no covariates are adjusted for, then the parameter estimated by a two-stage instrumental variable analysis is an odds ratio which is similar to a population odds ratio, except that it is conditional on the instrumental variable. The parameter will be numerically similar to the population odds ratio if the instrumental variable does not explain a large proportion of the variance in the exposure [Burgess and CHD CRP Genetics Collaboration, 2013].

The difference between odds ratio estimates due to non-collapsibility also presents difficulties when comparing odds ratios estimated using different analysis methods [Kaufman, 2010]. For example, odds ratios estimated from an instrumental variable analysis and from a logistic regression analysis are estimating different target parameters, and so will differ even in the absence of confounding [Pang et al., 2013].

5 Discussion

In this paper, we have considered different odds ratios depending on the choice of variables the odds ratio is a function of, and those it is averaged over. Odds ratios will therefore differ depending on the choice of covariate adjustment, levels of the exposure compared, and population over which the comparison is made. We have defined individual, conditional, marginal and population odds ratios, and seen how these differ in a simulation study. We have discussed the relevance of these differences to applied and theoretical researchers in a selection of methodological areas. Although this paper has focused on odds ratios, similar arguments could be made about other non-collapsible measures of association.

Although the numerical results of the simulation study relate only to the limited range of scenarios considered (binary and normally distributed continuous exposure, no confounding, normally distributed covariates), the intention of the paper is not to give a table of values for a researcher to know precisely the magnitude of attenuation of an odds ratio, but rather to know some of the factors influencing the magnitude of attenuation, and to have a sense of the potential relevance of non-collapsibility compared to other such phenomena affecting parameter estimates (for example, confounding, regression dilution bias, external validity) in interpreting odds ratio estimates.

5.1 Why estimate an odds ratio?

One natural reaction to the problems of the estimation of the odds ratio is to wonder why odds ratios are estimated at all. While the problems of non-collapsibility are not generally encountered with risk difference or risk ratio (relative risk) measures, they are not unique to the estimation of odds ratios and also occur for other non-collapsible measures of association (such as the coefficients from a probit regression model). The odds ratio is commonly used in epidemiological research in case-control studies, as it is taught that the odds ratio estimated in the case-control sample equals

the odds ratio in the general population, which in turn approximates the risk ratio [Breslow, 1996] (under the rare disease assumption [Greenland and Thomas, 1982]). While the equality in the whole population and in the case-control sample is only formally true for the individual odds ratio, the simulations in this paper suggest that this will be approximately true for the conditional odds ratio from a multivariable adjusted logistic regression model provided that the covariates explain a reasonable proportion of the true variability in the outcome risk. However, this equality does not hold even approximately for a population odds ratio, particularly when the odds ratio is far from 1.

Despite these criticisms, there is no obvious alternative measure of association in a case-control setting [Mood, 2010]. The probability of an event in the ascertained population is precisely determined by the sampling ratio of cases to controls, and so measures based on risk differences are likely to be highly sensitive to this sampling ratio. Log-linear regression may be employed in some cases, but the estimated probability of an event in the log-linear model may exceed 1, particularly if the sample contains a high ratio of cases to controls. Hence, it is likely that odds ratios will continue to be used in case-control settings in spite of their technical deficiencies.

5.2 Relevance of odds ratio parameters

It is important to note that none of the odds ratios defined and explored in this paper is the “correct” odds ratio. Each of the odds ratios simply represents a different quantity. Although the odds ratios are generally different, each is equal to unity under the null hypothesis, and therefore valid statistical inference can be made to test the null hypothesis based on any of the odds ratios [Vansteelandt et al., 2011]. Adjustment for a covariate which is not a confounder in a logistic regression analysis results in coefficients with less precision, although the efficiency of the coefficients is improved if the covariate is a predictor of the outcome [Robinson and Jewell, 1991].

Insofar as odds ratio parameters are of intrinsic interest (in many cases, other

measures such as absolute risk differences will be more informative), we contend that the odds ratio parameters of fundamental interest are the individual odds ratio, representing the ratio between odds at distinct levels of the exposure for an individual in the population, and the population odds ratio, representing the ratio between the average odds at distinct distributions of the exposure for the population as a whole. The individual odds ratio is generally of interest to patients or practitioners wanting to make decisions for a single individual; the population odds ratio is generally of interest to policymakers wanting to make decisions for the whole population [Stock, 1989]. If the exposure is binary rather than continuous, then there is no distinction between the marginal and population odds ratios. As previously stated, the individual odds ratio is a conceptual quantity rather than a practical one; in practice, a conditional odds ratio, representing the change in risk for an average individual with given levels of the covariates, will be used to convey a measure of individual risk.

Other odds ratio parameters are of interest if they approximate either of these parameters. For instance, estimates from a logistic regression analysis are of interest if most of the explainable variation in the outcome is adjusted for, in which case the estimate is close to the individual odds ratio; and estimates from an instrumental variable analysis are of interest if the instrument does not explain a large proportion of the variation in the exposure, in which case the estimate is close to the population odds ratio.

5.3 Case-control setting

One of the main justifications for the use of odds ratios is that the odds ratio is invariant to outcome-dependent sampling. This means that the same odds ratio is targeted in a case-control sample and in the full population [Didelez et al., 2010a]. While this is true in a logistic-linear model for the individual odds ratio, and therefore for the conditional odds ratio where there are no unmeasured covariates, it is not true in general for marginal or population odds ratios, nor for conditional odds ratios

with unmeasured covariates. This weakens the justification for the use of case-control data to estimate population-level parameters; although valid statistical inference using logistic regression can still be made.

5.4 Summary

In conclusion, although confounding is a more serious problem to investigators, non-collapsibility is a theoretical issue with important consequences. In answer to the motivating questions, ignoring statistical uncertainty, a coefficient in a logistic regression analysis increasing from 0.2 to 0.5 on adjustment for a covariate would seem to be greater than expected due to non-collapsibility alone, and a more likely conclusion is that the covariate is a confounder. However, based on the simulation results in the paper, an increase from 0.7 to 0.8 could plausibly be solely due to non-collapsibility, especially if the outcome is more common and the covariate explains a considerable proportion of the variation in the outcome. The simulations also suggest that the difference between odds ratio estimates of 2.0 from a logistic regression analysis and 1.8 from an instrumental variable analysis could be entirely due to non-collapsibility, particularly if the outcome is common. Understanding of non-collapsibility and the potential magnitude of its impact on estimates is important for applied researchers in differentiating between odds ratio estimates from different analysis methods or with adjustment for different covariates, and for theoretical researchers in understanding the ‘bias’ of odds ratio estimates from different methods [Austin et al., 2007; Stampf et al., 2010].

Acknowledgment

Stephen Burgess is supported by the Wellcome Trust (grant number 100114). No specific funding was received for the writing of this manuscript.

References

- Austin, P., Grootendorst, P., and Anderson, G. 2007. Conditioning on the propensity score can result in biased estimation of common measures of treatment effect: a Monte Carlo study. *Statistics in Medicine*, 26(4):734–753.
- Breslow, N. E. 1996. Statistics in epidemiology: the case-control study. *Journal of the American Statistical Association*, 91(433):14–28.
- Burgess, S. 2012. *Statistical issues in Mendelian randomization: use of genetic instrumental variables for assessing causal associations. Chapter 4: Collapsibility for IV analyses of binary outcomes*. PhD thesis, University of Cambridge. Available at <http://www.dspace.cam.ac.uk/handle/1810/242184>.
- Burgess, S. and CHD CRP Genetics Collaboration 2013. Identifying the odds ratio estimated by a two-stage instrumental variable analysis with a logistic regression model. *Statistics in Medicine*, 32(27):4726–4747.
- Cameron, A. C. and Trivedi, P. K. 2009. *Microeconometrics using Stata. Section 10.6: Marginal effects*. Stata Press College Station, TX.
- Davey Smith, G. and Ebrahim, S. 2002. Data dredging, bias, or confounding. *British Medical Journal*, 325(7378):1437.
- Didelez, V., Kreiner, S., and Keiding, N. 2010a. Graphical models for inference under outcome-dependent sampling. *Statistical Science*, 25(3):368–387.
- Didelez, V., Meng, S., and Sheehan, N. 2010b. Assumptions of IV methods for observational epidemiology. *Statistical Science*, 25(1):22–40.
- Diggle, P., Heagerty, P., Liang, K.-Y., and Zeger, S. 2002. *Analysis of longitudinal data. Chapter 7: Generalized linear models for longitudinal data*, volume 25. Oxford University Press, USA.

- Duffy, S. W., Rohan, T. E., and Altman, D. G. 1989. A method for combining matched and unmatched binary data application to randomized, controlled trials of photocoagulation in the treatment of diabetic retinopathy. *American Journal of Epidemiology*, 130(2):371–378.
- Ford, I. and Norrie, J. 2002. The role of covariates in estimating treatment effects and risk in long-term clinical trials. *Statistics in Medicine*, 21(19):2899–2908.
- Ford, I., Norrie, J., and Ahmadi, S. 1995. Model inconsistency, illustrated by the Cox proportional hazards model. *Statistics in Medicine*, 14(8):735–746.
- Gail, M., Wieand, S., and Piantadosi, S. 1984. Biased estimates of treatment effect in randomized experiments with nonlinear regressions and omitted covariates. *Biometrika*, 71(3):431–444.
- Greenland, S. and Robins, J. 1986. Identifiability, exchangeability, and epidemiological confounding. *International Journal of Epidemiology*, 15(3):413–419.
- Greenland, S., Robins, J., and Pearl, J. 1999. Confounding and collapsibility in causal inference. *Statistical Science*, 14(1):29–46.
- Greenland, S. and Thomas, D. 1982. On the need for the rare disease assumption in case-control studies. *American Journal of Epidemiology*, 116(3):547–553.
- Hauck, W., Anderson, S., and Marcus, S. 1998. Should we adjust for covariates in nonlinear regression analyses of randomized trials? *Controlled Clinical Trials*, 19(3):249–256.
- Hausman, J. 1983. *Handbook of econometrics*, volume 1, chapter Specification and estimation of simultaneous equation models, pages 391–448. Elsevier. see footnote 60.
- Janes, H., Dominici, F., and Zeger, S. 2010. On quantifying the magnitude of confounding. *Biostatistics*, 11(3):572–582.

- Judd, C. M. and Kenny, D. A. 1981. Process analysis estimating mediation in treatment evaluations. *Evaluation Review*, 5(5):602–619.
- Kaufman, J. S. 2010. Marginalia: comparing adjusted effect measures. *Epidemiology*, 21(4):490–493.
- Miettinen, O. and Cook, E. 1981. Confounding: essence and detection. *American Journal of Epidemiology*, 114(4):593–603.
- Mogstad, M. and Wiswall, M. 2010. Linearity in instrumental variables estimation: Problems and solutions. Technical report, Forschungsinstitut zur Zukunft der Arbeit. Bonn, Germany.
- Mood, C. 2010. Logistic regression: Why we cannot do what we think we can do, and what we can do about it. *European Sociological Review*, 26(1):67–82.
- Palmer, T., Thompson, J., Tobin, M., Sheehan, N., and Burton, P. 2008. Adjusting for bias and unmeasured confounding in Mendelian randomization studies with binary responses. *International Journal of Epidemiology*, 37(5):1161–1168.
- Pang, M., Kaufman, J. S., and Platt, R. W. 2013. Studying noncollapsibility of the odds ratio with marginal structural and logistic regression models. *Statistical Methods in Medical Research*.
- Pearl, J. 2000. *Causality: models, reasoning, and inference. Chapter 6: Simpson’s paradox, confounding and collapsibility*. Cambridge University Press.
- Robinson, L. and Jewell, N. 1991. Some surprising results about covariate adjustment in logistic regression models. *International Statistical Review*, 59(2):227–240.
- Simpson, E. H. 1951. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241.

- Stampf, S., Graf, E., Schmoor, C., and Schumacher, M. 2010. Estimators and confidence intervals for the marginal odds ratio using logistic regression and propensity score stratification. *Statistics in Medicine*, 29(7–8):760–769.
- Steyerberg, E., Bossuyt, P., and Lee, K. 2000. Clinical trials in acute myocardial infarction: Should we adjust for baseline characteristics? *American Heart Journal*, 139(5):745–751.
- Stock, J. 1989. Nonparametric policy analysis. *Journal of the American Statistical Association*, 84(406):567–575.
- VanderWeele, T. and Vansteelandt, S. 2010. Odds ratios for mediation analysis for a dichotomous outcome. *American Journal of Epidemiology*, 172(12):1339–1348.
- Vansteelandt, S., Bowden, J., Babanezhad, M., and Goetghebeur, E. 2011. On instrumental variables estimation of causal odds ratios. *Statistical Science*, 26(3):403–422.
- Zeger, S., Liang, K., and Albert, P. 1988. Models for longitudinal data: a generalized estimating equation approach. *Biometrics*, 44(4):1049–1060.
- Zhang, Z. 2008. Estimating a marginal causal odds ratio subject to confounding. *Communications in Statistics – Theory and Methods*, 38(3):309–321.
- Zou, G. Y. 2007. One relative risk versus two odds ratios: implications for meta-analyses involving paired and unpaired binary data. *Clinical Trials*, 4(1):25–31.