ODDS RATIOS

PETER M. LANCE

Abstract. One of the most curious differences between the health economics and public health traditions is the heavy reliance of the latter on odds rations, while the focus in the former tends to be marginal effects and relative risks. In this manuscript, we introduce the core concepts motivating the two and discuss interpretational issues.

1. Introduction

In both the health economics and public health traditions, there is a heavy reliance on limited dependent variable models such as logit, in large part due to the fact that many behaviors and outcomes from a health standpoint are binary in nature. Despite this common departure point, researchers in these two traditions have tended to adhere to very different reporting practices. Specifically, health economists tend to prefer to express conclusions regarding the role of key explanatory variables with marginal effects. Public health researchers, on the other hand, have tended to rely heavily on odds-ratios. In this manuscript we describe and then compare the two approaches.

2. The Basics

It is perhaps most useful to begin with a few basic concepts. First, for some potential binary event B (i.e. the event either occurs ($B = 1$) or doesn't ($B = 0$)) let

$$P = Pr(B = 1)$$
$$P_1 = Pr(B = 1|D = 1)$$
$$P_0 = Pr(B = 1|D = 0)$$

where D is some binary characteristic that we think of as playing a role in the conditional distribution of B. D could, for instance, be gender or race (black/white). (We would thus be thinking that one's gender or race has some influence on the probability of the event B happening.)

Then, we have the following potential parameters of interest:

The *Marginal Effect* (of D on the probability of B):

$$ME = P_1 - P_0$$

The marginal effect is the difference in the probability of the event occurring between those for whom $D = 1$ and $D = 0$. The marginal effect is the *difference in risk*.

1991 *Mathematics Subject Classification.* Odds Ratios.
Key words and phrases. Odds Ratios.
Beethoven, who supplied the beat.

The *Relative Risk (Sometimes called the Risk Ratio)* (of B occurring for those for whom $D = 1$ over those for whom $D = 0$):

$$RR = \frac{P_1}{P_0}$$

The relative risk is not the difference in risks. Rather, it is the relative size of the two risks. When you say that the relative risk is 1.7, that means that the risk of something happening for those for whom $D = 1$ is 70 percent larger proportionally than those for whom $D = 0$. It does not mean that the risk for those for whom $D = 1$ is 70 percentage points greater than that for those whom $D = 0$.

The *Odds* (of B Occurring):

$$O = \frac{P}{1 - P}$$

This basically captures the risk of something happening compared to it not happening. If it is 3 it means that the risk of something happening is three times that of it not happening *within some reference group* (be it the whole sample, sub-samples stratified by D, etc.). Since this is a binary event (and hence the outcome where it happens and the outcome where it doesn't is the mutually exclusive and *exhaustive* set of possibilities) we therefore know that if the odds are 3 then $P = .75$. Similarly, if the odds are 4, then $P = .8$. The point is this: the odds tells us the probability or risk of the event occurring. Another way of looking at this is that we have one equation (and a monotonic one at that for $P \in (0, 1)$) and one unknown.

The *Odds Ratio* (for B occurring between those for whom $D = 1$ and those for whom $D = 0$):

$$OR = \frac{\frac{P_1}{1-P_1}}{\frac{P_0}{1-P_0}}$$

The first thing to notice is that simple re-arrangement yields:

$$OR = \frac{P_1}{1 - P_1} \cdot \frac{1 - P_0}{P_0}$$

An immediate implication of this is that, unless D plays no role in the conditional distribution of B (i.e. $P_1 = P_0$), then ***an odds ratio is not the same thing as relative risk!!!*** Generally speaking, you cannot interpret an odds-ratio in relative risk terms. Notice, however, that the closer P_1 and P_0 are the closer the odds ratio will be to 1 and the closer the relative risk and the odds ratio will be.

A relative risk of 1.7 means that the event is 70 percent relatively more likely to occur for those for whom $D = 1$ than those for whom $D = 0$. In other words, the risk of it happening is proportionally 70 percent greater for those for whom $D = 1$ than $D = 0$.[1]

However, an odds ratio of 1.7 cannot be interpreted this way. For instance, when $P_1 = .5$ and $P_0 \cong .37$, the resulting odds-ratio is roughly 1.7 but the relative risk is 1.35 or so, while the marginal effect is .13. Similarly, suppose that the odds ratio is .7. For instance, when $P_1 = .5$ and $P_0 \cong .59$ the odds ratio is around .7. The relative risk is $.5/.59 \cong .85$ and the marginal effect is -.09.

[1] While we're on the subject, once again it is important not to confuse the relative risk with the marginal effect or, to put marginal effect slightly differently, the *difference* in risk.

The second thing to notice is that you cannot unequivocally recover underlying risks from the odds ratio. To see this, first consider again the odds. The odds are a ratio fundamentally involving only one risk. Therefore, algebraically this is a situation of one equation and one unknown, which is in principal solvable. The odds ratio, by contrast, is a ratio of ratios, each of which involves a different unknown (P_1 and P_0). Therefore, if you have an odds ratio of 2, that could mean that the odds for $D = 1$ is 2 while the odds for $D = 0$ is 1, or that the odds for $D = 1$ is 4 while the odds for $D = 0$ is 2, or that that odds for $D = 1$ is 8 while $D = 0$ is 4, etc. In terms of the numerical examples in the preceding paragraph, lots of other $\{P_1, P_0\}$ combinations will yield the odds ratios 1.7 and .7.

This is in principal an important conceptual drawback to odds ratios. The outcome of behavioral interest is the risk or probability of the event B occurring. What we wish to know about is how various traits or characteristics of the individual, his or her household or community, his of her workplace environment, etc. influence the probability of the event B occurring. In other words, we are after the *conditional distribution* of B. The metric of ultimate interest for the outcome is thus the underlying probability of B occurring. Certainly marginal effects are in this metric. From the odds, the outcome in this metric can easily be recovered. However, the question one has to ask, given that our scientific goal is to understand the role of D in the conditional distribution of B, is why the odds are particularly interesting: they tell us **nothing** about the role of D in the conditional distribution of B. Maybe the odds of an outcome occurring for African Americans is 3. By itself, so what?

You cannot recover the outcome in this metric from the relative risk but, like the odds, the relative risk at least has the advantage of possessing a clear behavioral interpretation. But it can be a deceiving one. Suppose, for instance, we were considering the urgency of some intervention to create greater equity in the outcome B with respect to D (for instance, we might wish to know whether it is worth the public investment, perhaps in gender-specific preventive interventions that will require a great deal of input of research resources to develop, to equalize the risk of some discrete health outcome occurring across men and women). We estimate the relative risk

$$RR = \frac{P_1}{P_0}$$

to be 3. The would be true if $P_1 = .9$ and $P_0 = .3$. Were we to express the marginal effect of D in this case it would be .6, which is pretty huge: the two groups face wildly different risks of a pretty common health outcome occurring. Surely we would have to consider allocating scarce resources to close this gap. The problem, however, is that you get the same relative risk if $P_1 = .0009$ and $P_0 = .0003$. Clearly this health outcome is a big deal to those who experience it, but from the standpoint of setting policy priorities for the population as a whole this would seem to involve little yield toward equity in terms of the absolute resources committed.

Like the relative risk, odds ratios are also not directly interpretable in terms of the metric of the outcome of interest. Indeed, they tell us nothing about the marginal distribution of B with respect to D in the behavioral metric of interest. Worse still, the odds ratio is not even readily behaviorally interpretable in principal in the fashion that the relative risk is. The odds ratio is a ratio of ratios, but what does that *mean* from a behavioral standpoint? I have rarely heard a convincing answer to this question, and even those that are are pretty strained. For instance, in his class notes Prof. Edward Norton has the following to say:

> "Here is one interpretation that will only make sense after studying duration models. Think of the outcome as being time dependent (in the above example, the time it takes for someone to start smoking). The odds ratio is the ratio of the *time to outcome* for one type of observation relative to another. For instance, if the odds ratio was 1.2 in the example above, then on average women take 20 percent longer to start smoking."

Of course, this is a fairly strained interpretation (and I'm not even sure I fully buy it as such), and in any case we aren't approaching this in duration model context.

In summary, marginal effects are directly behaviorally interpretable, and in the metric of interest: the probability of the event B occurring. Relative risks also carry a clear behavioral interpretation. However, they are not in the metric of the probability of interest, and thus the scale of the behavioral inference they provide is unclear, a serious drawback from the standpoint of, for instance, finding the most cost effective policies toward which to commit scarce resources. Odds ratios are neither readily behaviorally interpretable in any fashion obvious to me (beyond the vacuous statement "well, the odds ratio is the ratio of the odds") or in the metric of immediate interest.

One interesting question is whether there is any kind of mapping from relative risk to the odds ratio, so that one could at least unequivocally make the case that a given odds ratio implies some level of relative risk. The answer, unfortunately, is no. Consider the case of $P_1 = .6$ and $P_0 = .4$. The relative risk in this instance is 1.5, while the odds ratio is 2.25. Now, consider the case of $P_1 = .06$ and $P_0 = .04$. The relative risk is still 1.5, but the odds ratio is now around 1.53. The point is that there is not a one-to-one mapping between odds ratios and relative risks.

If there is any general rule one can cite, it has been my casual observation that odds ratios tend to exaggerate relative risk. Odds ratios greater than 1 tend to exaggerate the increase in risk (OR $>$ RR) while those less than 1 exaggerate the *decrease* in risk (OR $<$ RR). The difference seems to get worse the more common is the event in question (notice the example in the last paragraph: when $P_1 = .6$ and $P_0 = .4$ the odds ratio and the relative risk are very different, but when $P_1 = .06$ and $P_0 = .04$ they are very close, and the change has all been to the odds ratio). When events are common, odds ratios and relative risks will tend to be pretty close in cases only when the odds ratio is close to 1.

The only convincing substantive argument I have heard for odds ratios appeals to case-control studies. In other words, it appeals to studies that have a group of cases (e.g. patients with some kind of cancer) and then try to recruit controls who do not have the cancer but are otherwise similar to them. This has some really obvious problems. For one thing, confounding is a problem. For another, it can represent a species of choice-based sampling.

The claim, however, is that you can compute odds ratios with case control studies, but not risks, and hence relative risks. I found the following example online[2]. Suppose you want to understand the association between esophageal cancer and alcohol consumption. Suppose as well that we gathered interviews with 200 cases and 775 controls, with the following results.

[2] At astor.som.jhmi.edu/ esg/GRENADA/RRandOR.ppt.

	Alcohol (mg/day)		
	> 80	≤ 80	
Case	96	104	200
Control	109	666	775
	205	770	

One question for present purposes is whether, with this information, one can calculate the probability of esophageal cancer if you drink more than 80 grams of alcohol per day. The problem is that, from this case-control data, you cannot know the rate of esophageal cancer for the population overall (or, rather, for the population for which you wish to make inferences with this sample). The reason is that case-control studies are basically like a species of choice-based sampling (where the "choice" is having esophageal cancer) that are thus not representative of some well-defined population.[3] Typically, they over-sample "cases."[4] That means that we cannot establish true risk of esophageal cancer in the population. That means that we cannot estimate the relative risks either.[5]

The odds ratio is a bit different story. It is simply:

$$\frac{\frac{\frac{96}{205}}{\frac{109}{205}}}{\frac{\frac{104}{770}}{\frac{666}{770}}} = \frac{96}{109} \cdot \frac{666}{104} = 5.64$$

The reason the odds ratio can be computed but risks cannot be computed is sometimes not immediately clear. However, if you look carefully at these fractions one element of the explanation is evident. The reference group for the odds are the columns, not the rows: those who fall in the different points support for D, not B. There is another interesting way of looking at this in the logit framework that I will discuss in the next section.

3. Odds Ratios, Marginal Effects and Relative Risks in the Logit Model

3.1. **The Logit Model.** At this point it is perhaps useful to conduct a basic review of the logit model. I do this for two reasons. First, it is always good to clear the cobwebs out of my own mind. Second, it is very important to remember that the logit approach is derived from a behavioral model.

Continuing some of the notation from above, suppose that all of the individuals in some population of interest can be characterized (in terms of their background characteristics) by D and another characteristic x (for present purposes, it does not really matter if x is continuous or not in nature). Our goal is to learn about the conditional distribution of B:

$$Pr(B = 1|D, x)$$

[3]Remember, our goal with them is typically to be able say something about the role of some environmental or behavioral factor in shaping the risk of becoming a 'case'.

[4]We can think of them as having pure choice-based sampling as a departure point, and then try to water down the bias by recruiting some controls.

[5]Typically, case-control studies effectively over-sample cases. If you simply gathered a random sample from the population of interest, it would in many cases have to be huge to collect enough cases. Case-control studies are, perhaps unsurprisingly, a particularly popular approach with relatively rare conditions.

Individuals derive some level of latent utility or satisfaction whether they experience the outcome $B = 1$ or the outcome $B = 0$. For instance, the utility or satisfaction that the individual experiences if $B = 1$ could be given by

$$B_1^* = \delta_{10} + \delta_{11} \cdot D + \delta_{12} \cdot x + \xi_1$$

Similarly, utility or satisfaction when $B = 0$ is given by

$$B_0^* = \delta_{00} + \delta_{01} \cdot D + \delta_{02} \cdot x + \xi_0$$

These can be regarded as indirect utility functions, and as such it is easy in principal to incorporate into them any variables that influence the costs of the two decisions.[6] Notice that these equations relate utility or satisfaction under the respective values of B to some observed characteristics or circumstances of the individual

With this framework in hand, an easy decision rule to determine B is readily apparent. Specifically, $B = 1$ if $B_1^* \geq B_0^*$ and $B = 0$ if $B_0^* > B_1^*$. Focusing on the former case, we have:

$$B_1^* \geq B_0^*$$

or

$$\delta_{10} + \delta_{11} \cdot D + \delta_{12} \cdot x + \xi_1 \geq \delta_{00} + \delta_{01} \cdot D + \delta_{02} \cdot x + \xi_0$$

or, re-arranging,

$$\xi_1 - \xi_0 \geq (\delta_{00} - \delta_{10}) + (\delta_{01} - \delta_{11}) \cdot D + (\delta_{02} - \delta_{12}) \cdot x$$

or

$$\xi_0 - \xi_1 \leq (\delta_{10} - \delta_{00}) + (\delta_{11} - \delta_{01}) \cdot D + (\delta_{12} - \delta_{02}) \cdot x$$

Now, remember that we only observe the binary outcome B and not the latent satisfaction levels B^*, and so our modeling must be done in terms of what we can observe (B), or rather the probability that $B = 1$ (we could also model explicitly in terms of the probability $B = 0$ since $Pr(B = 1) + Pr(B = 0) = 1$). The point is, we need some way to develop a tractable, and ideally closed-form, expression for the cumulative probability

$$Pr(B = 1|D, x) = Pr\left(\xi_0 - \xi_1 \leq (\delta_{10} - \delta_{00}) + (\delta_{11} - \delta_{01}) \cdot D + (\delta_{12} - \delta_{02}) \cdot x\right)$$

or, combining terms for parsimony,

$$Pr(B = 1|D, x) = Pr\left(\xi_0 - \xi_1 \leq \beta_0 + \beta_1 \cdot D + \beta_2 \cdot x\right)$$

The specific motivation for the logit model is to assume that the errors ξ follow a Type-1 Extreme value distribution, because cumulative density of the difference between two Type-I Extreme Value distributed random variables is given by the familiar logistic function. Specifically, suppose that ϑ is a random variable created by

[6] Actually, given that these are indirect utility functions, the harder thing to explain is how background characteristics like race, gender, age, education, etc. enter into their right-hand sides. The easiest way to motivate this is to assume that these factors play a role in the health production function. Interactions between these sorts of factors and economic variables like the price of care can be motivated by taking an intertemporal value function approach and simply assuming a sufficiently complex expectations formation process, implying that one cannot rule out any particular functional form from a Taylor-series approximation of the expected value function. Perhaps more behaviorally convincingly, interactions could be justified simply by assuming that one's per-period utility function is not separable in consumption and health.

subtracting one Type-I Extreme Value distributed random variable from another. Then

$$Pr(\vartheta \leq a) = \frac{exp(a)}{1 + exp(a)}$$

Thus, under this assumption about the distribution of the errors ξ,

$$Pr(B = 1|D, x) = \frac{exp\,(\beta_0 + \beta_1 \cdot D + \beta_2 \cdot x)}{1 + exp\,(\beta_0 + \beta_1 \cdot D + \beta_2 \cdot x)}$$

Since $Pr(B = 0|D, x) = 1 - Pr(B = 1|D, x)$ it is trivial to establish that

$$Pr(B = 0|D, x) = \frac{1}{1 + exp\,(\beta_0 + \beta_1 \cdot D + \beta_2 \cdot x)}$$

From this basic logic the multinomial logit model can also be derived quite easily.[7]

From this derivation, two things should be clear. First, the logit model is, first and foremost, a *behavioral* model: it assumes an underlying decision making process behind the observed outcome B. If B^* were observable, we would likely simply regress $B_1^* - B_0^*$ on D and x. Because it is not observable, we are instead forced to rely on modeling the probability that $B = 1$. Why we would wish to get even further from the heart of the matter is not clear. The objective in estimating this behavioral model is to learn how D and x influence the actual distribution of the behavioral variable of interest. Second, and on a closely related note, logit is, as posed, a *regression* model. Once again, the goal is and should be to understand how D and x affect the distribution of the dependent variable, $Pr(B = 1|D, x)$, in the same fashion that the real objective in any generic regression model $E(y|x)$ is the marginal effect of x on y. Any reporting convention that moves one away from reporting the direct influence of D and x on the $Pr(B = 1)$ is of unclear value.

3.2. **Marginal Effects and Relative Risks in the Logit.** Within the logit model framework, marginal effects and relative risks are fairly straightforward. They are, respectively,

$$ME = Pr(B = 1|D = 1, x) - Pr(B = 1|D = 0, x)$$
$$= \frac{exp\,(\beta_0 + \beta_1 + \beta_2 \cdot x)}{1 + exp\,(\beta_0 + \beta_1 + \beta_2 \cdot x)} - \frac{exp\,(\beta_0 + \beta_2 \cdot x)}{1 + exp\,(\beta_0 + \beta_2 \cdot x)}$$

---

[7]Many derivations of the logit model that one encounters suffer a small but important mistake. Specifically, their logic is as follows:

> "Assume that an individual has the single indirect utility function
>
> $$B^* = \lambda_0 + \lambda_1 \cot D + \lambda_2 \cdot x + \psi$$
>
> The individual then chooses $B = 1$ if
>
> $$B^* = \lambda_0 + \lambda_1 \cdot D + \lambda_2 \cdot x + \psi \geq 0$$
>
> or, re-arranging,
>
> $$\psi \leq -\lambda_0 - \lambda_1 \cdot D - \lambda_2 \cdot x$$
>
> If one assumes that ψ follows a Type-I Extreme Value distribution this gives rise to the choice probability
>
> $$Pr(B = 1|D, x) = \frac{exp(-\lambda_0 - \lambda_1 \cdot D - \lambda_2 \cdot x)}{1 + exp(-\lambda_0 - \lambda_1 \cdot D - \lambda_2 \cdot x)}"$$

The problem with this is that the cumulative density of a Type-I Extreme Value distributed random variable is *not* given by the logistic function. It is only the cumulative density of the *difference* of two such random variables that is given by the logistic function.

and

$$
\begin{aligned}
RR &= \frac{Pr(B=1|D=1,x)}{Pr(B=1|D=0,x)} \\
&= \frac{\frac{exp(\beta_0+\beta_1+\beta_2\cdot x)}{1+exp(\beta_0+\beta_1+\beta_2\cdot x)}}{\frac{exp(\beta_0+\beta_2\cdot x)}{1+exp(\beta_0+\beta_2\cdot x)}} \\
&= \frac{exp\left(\beta_0+\beta_1+\beta_2\cdot x\right)}{1+exp\left(\beta_0+\beta_1+\beta_2\cdot x\right)}\cdot\frac{1+exp\left(\beta_0+\beta_2\cdot x\right)}{exp\left(\beta_0+\beta_2\cdot x\right)} \\
&= exp\left(\beta_1\right)\cdot\frac{1+exp\left(\beta_0+\beta_2\cdot x\right)}{1+exp\left(\beta_0+\beta_1+\beta_2\cdot x\right)}
\end{aligned}
$$

Notice that both of these are a function of the other explanatory variables in the logit model.

A marginal effect can (and should) be computed for every observation in the sample. The idea here is that we want to evaluate everyone's probability of choosing $B = 1$ in two states of the world: when $B = 1$ and when $B = 0$. One should thus compute $Pr(B = 1|D = 1, x)$ and $Pr(B = 1|D = 0, x)$ for each observation in the sample. In other words, for their given empirical values of the other explanatory x, compute their choice probability under the two states of the world. To get some overall sense of the marginal effect or risk ratio with respect to D for the sample, simply average these across every observation in the sample. The alternative approach is to compute the marginal effect and risk ratio for a "representative agent": one for whom $x = \bar{x}$, the mean of x across the sample. In other words, compute $Pr(B = 1|D = 1, \bar{x})$ and $Pr(B = 1|D = 0, \bar{x})$ and the resulting marginal effect and risk ratio.

3.3. **Odds Ratios.** Under the logit model odds ratios are a slightly different story. First, note that the odds when $D = 1$ are:

$$
\begin{aligned}
&\frac{Pr(B=1|D=1,x)}{1-Pr(B=1|D=1,x)} \\
&= \frac{\frac{exp(\beta_0+\beta_1+\beta_2\cdot x)}{1+exp(\beta_0+\beta_1+\beta_2\cdot x)}}{1-\frac{exp(\beta_0+\beta_1+\beta_2\cdot x)}{1+exp(\beta_0+\beta_1+\beta_2\cdot x)}} \\
&= \frac{\frac{exp(\beta_0+\beta_1+\beta_2\cdot x)}{1+exp(\beta_0+\beta_1+\beta_2\cdot x)}}{\frac{1}{1+exp(\beta_0+\beta_1+\beta_2\cdot x)}} \\
&= exp\left(\beta_0+\beta_1+\beta_2\cdot x\right)
\end{aligned}
$$

By similar logic,

$$
\begin{aligned}
&\frac{Pr(B=1|D=0,x)}{1-Pr(B=1|D=0,x)} \\
&= exp\left(\beta_0+\beta_2\cdot x\right)
\end{aligned}
$$

The odds ratio is thus

$$
\begin{aligned}
&\frac{\frac{Pr(B=1|D=1,x)}{1-Pr(B=1|D=1,x)}}{\frac{Pr(B=1|D=0,x)}{1-Pr(B=1|D=0,x)}} \\
&= \frac{exp\left(\beta_0+\beta_1+\beta_2\cdot x\right)}{exp\left(\beta_0+\beta_2\cdot x\right)} = exp\left(\beta_1\right)
\end{aligned}
$$

As we can see, under the logit model the odds ratio does not depend on the level of other covariates x. $exp(\beta_1)$ is, I believe, what STATA reports when it presents odds ratios. I think that, when confronted with a continuous variable, STATA is computing the odds ratio as the ratio of the estimated odds of the outcome when the continuous variable increases by one unit. In other words, for a continuous variable, compute the odds. Then increase the variable by one unit and re-compute the odds. Finally, divide the odds following the one unit increase by the odds before the one unit increase.[8]

For simplicity, now ignore the other covariates x (in other words, assume they were not in our model and did not belong there, so that the logit model simply related the probability of $B = 1$ to D: $Pr(B = 1|D)$). Notice that the marginal effect and relative risk are functions of β_0:

$$\begin{aligned} ME &= Pr(B = 1|D = 1, x) - Pr(B = 1|D = 0, x) \\ &= \frac{exp\,(\beta_0 + \beta_1)}{1 + exp\,(\beta_0 + \beta_1)} - \frac{exp\,(\beta_0)}{1 + exp\,(\beta_0)} \end{aligned}$$

and

$$\begin{aligned} RR &= \frac{Pr(B = 1|D = 1)}{Pr(B = 1|D = 0)} \\ &= exp\,(\beta_1) \cdot \frac{1 + exp\,(\beta_0)}{1 + exp\,(\beta_0 + \beta_1)} \end{aligned}$$

However, even with this adjustment the odds ratio is still

$$OR = exp\,(\beta_1)$$

The really striking difference is that the odds ratio is not a function of β_0, whereas the marginal effect and relative risk are.

From this observation, we can see through the lens of the logit model why the odds ratio can be computed for case-control studies but relative risks (and marginal effects) cannot. The main problem with the case-control approach is that, by effectively over-sampling cases relative to what a true probability sample might have yielded, it undermines our ability to estimate the average risk level for the outcome of interest, which is driven by the constant parameter β_0. Odds ratios are still reasonable because they do not involve β_0.

There are a few rules of thumb about odds ratios and logit models, which I pilfer from Edward Norton's class notes. If β_1 is close to zero, then the odds ratio will be approximately $1 + \beta_1$. This approximation is an underestimate if $\beta_1 > 0$ and an overestimate if $\beta_1 < 0$. If $\beta_1 = 0.7$, then the odds ratio is 2. If $\beta_1 = 3$ then the odds ratio is 20.

4. Odds Ratios: Deadly Intellectual Weapons

The discussion above has mainly been about the properties of odds ratios, relative risks and marginal effects. Perhaps the greatest drawback to odds ratios, one which does not admittedly reflect a substantive defect to them, is the way that they have been interpreted in the media. Specifically, journalists often present results based on odds ratios as if they reflected relative risk. The *Language Log*, a blog concerned primarily with the abuse and torture of language, has a good post on this (see http://itre.cis.upenn.edu/ myl/languagelog/archives/004767.html). It uses as an

[8] Though there might be other methods. A quick web search yielded some results on the subject (none of which I have read), such as Moser, B. and L. Coombs. (2004) "Odds ratios for a continuous outcome variable without dichotomizing" *Statistics in Medicine* 23(13): 1843-60.

example an article in the *New England Journal of Medicine* that found disparities in treatment standards by race and gender, findings which prompted the following in the popular press:

> "Physicians said they would refer blacks and women to heart specialists for cardiac catheterization tests only 60 percent as often as they would prescribe the procedure for white male patients."-The Washington Post
>
> "[Doctors] refer blacks and women to heart specialists 60% as often as they would white male patients."-LA Times
>
> "Doctors are only 60% as likely to order cardiac catheterization for women and blacks as for men and whites."-NY Times

In actuality, the study reported an odds ratio of around .6 (specifically, of .574) for women and blacks. The average reader of these newspapers might, however, have assumed that blacks and women were only 60 percent as likely, proportionately, to receive cardiac catheterization. Since the referral rate for cardiac catheterization for white men was 90.6 percent, the casual reader of there papers would assume that that for blacks and women was .6·90.6=54.4 percent, a huge gap. However, since the journal article reported an odds ratio, we can readily compute the odds for men as

$$\frac{.906}{1-.906} = 9.64$$

The odds ($odds_{bw}$) for blacks and women thus must be solvable from

$$\frac{odds_{bw}}{9.64} = .6$$

or

$$odds_{bw} = 9.64 \cdot .6 = 5.784$$

From this odds, it is easy to compute the referral rate for blacks and women as 84.7 percent. That still represents a gap, but nowhere near as dramatic as that for a referral rate of 54.4 percent.

This sort of sloppiness has real consequences. It exaggerates our sense of risk from various behaviors and environmental factors, leading to counterproductively strong adjustments in lifestyles, undue fear and potential mis-allocation of resources as, for instance, undue fears translate into political pressure to change funding priorities at institutions like the NIH. Unfortunately, it is also widespread: more often than not, when I hear someone like Katie Couric report the results of a new health study in terms suggestive of relative risk, further inspection of the article cited reveals that the media was reporting an odds ratio as if it were a relative risk.

MEASURE Evaluation, Carolina Population Center, University of North Carolina at Chapel Hill. (CB#8120, University Square, 123 West Franklin St., Chapel Hill, NC. 27516-2524

E-mail address: pmlance@email.unc.edu