Odds Ratios and Risk Ratios

Caitlin J. Bakker, PhD, MLIS, AHIP-D (She/Her/Hers)[a]

[a]Discovery Technologies Librarian, Dr. John Archer Library & Archives, University of Regina, Regina, Saskatchewan, Canada, https://orcid.org/0000-0003-4154-8382, caitlin.bakker@uregina.ca

Cite as: Bakker CJ. Odds Ratios and Risk Ratios. Hypothesis. 2026;38(1). doi:10.18060/29536

The previous column introduced probability and odds[1]. This column will build on those pieces to introduce odds ratios and risk ratios, two very useful ways of describing the likelihood of an occurrence between different groups.

When discussing both odds ratios and risk ratios, exposures and events are often mentioned. Often, these can be discussed in terms of disease and negative outcomes, such as exposure to cigarette smoke and developing lung cancer. However, exposure or event can be interpreted more broadly. An event might be considered the outcome of interest, while the exposure might be considered an intervention or some sort of change in condition. For example, if we were flipping a coin and wanted to know whether using the left or right hand had an impact on whether the coin landed on heads or tails, landing on heads or tails would be the event, while the right or left hand would be the exposure.

An odds ratio is a ratio of ratios. Recall from the previous column that "odds" means the "the probability of an event occurring divided by the probability of it not occurring." An odds ratio is calculated by taking the odds of an event in one group (the exposed group), and dividing it by the odds of the event happening in the unexposed group. Using the table, an odds ratio would be calculated as (a / b) / (c / d).

Table 1. 2x2 table to calculate odds ratios and risk ratios

<table>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Event</td></tr>
<tr><td>Yes</td><td>No</td></tr>
<tr><td rowspan="2">Exposure</td><td>Yes</td><td>a</td><td>b</td></tr>
<tr><td>No</td><td>c</td><td>d</td></tr>
</table>

A risk ratio, also called relative risk, is the ratio of two probabilities: the probability (or risk) of an event happening in the exposed group, and the probability (or risk) of it happening in the unexposed group. Remember from the previous column that the definition of probability is "how likely it is that an event would occur." A risk ratio would be calculated as (a / (a + b)) / (c / (c + d)).

A Practical Application

Imagine that a medical library is having an issue with patrons not returning their materials on time, despite a reminder being emailed to them a week before the item is due. The library is wondering whether sending a text message on the day the item is due would improve the number of on-time returns. To test out this idea, they randomly divide their patrons with materials into two groups: the group that will receive both the email and text message (the exposed group), and the group that will receive only the email (the unexposed group). They count how many of the items from each group are returned on time. They test out using this text message reminders for one month.

Table 2. On-time returns and late-returns at one library during a pilot

	On Time Returns	Late Returns
Exposed (Email + Text)	132	175
Unexposed (Email)	117	198

If we were to calculate the odds ratio, it would be as follows:

$$\frac{a/b}{c/d} = \frac{132/175}{117/198} = \frac{0.754}{0.591} = 1.276$$

If we were to calculate the risk ratio, it would be as follows:

$$\frac{(\frac{a}{a+b})}{(\frac{c}{c+d})} = \frac{(\frac{132}{132+175})}{(\frac{117}{117+198})} = \frac{(\frac{132}{307})}{(\frac{117}{315})} = \frac{0.43}{0.371} = 1.16$$

Odds ratios and risk ratios seem to be fairly close, and in our above example, we might draw some fairly similar conclusions: people who receive both an email and a text seem to be more likely to return materials on time than those who receive only email. When looking at the RR, we could conclude that the group who received both email and text were 16% more likely to return materials on time compared to the email alone group. The OR would suggest that the odds of returning materials on time are about 28% higher in the email and text group.

The risk ratio focuses on how likely an event is to happen in one group or another, and it doesn't consider how likely the event is to not happen. In contrast, the odds ratio compares the odds (or chances) of something happening in one group to the odds of it happening in another. The risk ratio can be a bit more immediately intuitive, and the odds ratio is sometimes used in more sophisticated calculations or in cases where it's not possible to calculate the risk (for example, in a case-control study when you purposefully select participants based on outcomes).

Let's imagine that the library that started this text message initiative shared this information with other libraries in their system, who also decided to implement a pilot. There is enthusiasm to see whether implementing this system-wide would improve return rates generally. To help test this initiative, four libraries implement the text message reminder for half of their patrons for one month and track the returns data.

Table 3. On-time returns, late returns, and odds and risk ratios for four libraries participating in a pilot

Library		On Time Returns	Late Returns	Odds Ratio	Risk Ratio
Library 1	Exposed	132	175	1.28	1.16
	Unexposed	117	198		
Library 2	Exposed	14	128	2.03	1.93
	Unexposed	7	130		
Library 3	Exposed	97	5	0.59	0.98
	Unexposed	99	3		
Library 4	Exposed	127	14	8.63	1.76
	Unexposed	82	78		

For Library 2, on-time returns are rare (10% vs. 5%), and the OR (2.03) and RR (1.93) are very similar. In Library 3, on-time returns are very common (95% vs. 97%), and the OR (0.59) exaggerates the difference compared to the RR (0.98).

Understanding why this happens requires examining how odds ratios work. Odds ratios compare the odds of an event happening and the odds of it not happening between two groups. In Library 3, the exposed group had 97 on time returns per 5 late returns (odds = 19.4), while the unexposed group had 99 on time returns per 3 late returns (odds = 33). Both groups had very high on-time return rates (95% vs. 97%), but the odds are very different because we are dividing by such a small number of late returns.

In contrast, in Library 2, most items are returned late in both groups. Because of this, the denominators (128 late returns in the exposed group, 130 in the unexposed group) are large and similar to the total sample size. Remember that risk ratio uses the total sample size for each group, while an odds ratio is the ratio of events to non-events. When the number of non-events is similar to the total sample size, the odds ratio and risk ratio are also similar. Generally, for common events, the risk ratio is preferred, while either the odds ratio or the risk ratio can be used when events are rare.

If we look more closely at Library 1, we also see that the differences between the two groups are quite small. Among the patrons who received a text message, 43% of the items were returned on time. Among those who didn't receive a text message, 37% of the items were returned on time. The OR (1.28) and the RR (1.16) are relatively close together. If we compare that to Library 4, where 90% of materials in the exposed group were returned compared to 51% in the unexposed group, the OR (8.63) and RR (1.76) are quite different. Generally, the larger the difference between groups, the larger the difference between OR and RR, and the odds ratio will appear to exaggerate the difference. For that reason, risk ratio may be preferable in cases where there are large differences in event rates between groups, while either odds ratios or risk ratios can be used when differences are small.

Interpreting Odds Ratios and Risk Ratios

Odds ratios and risk ratios can be interpreted in the same way. Where the OR or RR = 1, it means that there is no association between the exposure and the event. Where the OR or RR is greater than 1, it means that the odds or risks are higher for the exposed group, while where the OR or RR is less than 1, it means that the odds or risks are lower for the exposed group[2-4].

We calculate OR and RR in part so that we can generalize to a larger population. In our example, we are interested in learning whether or not adding the text message reminder improves on-time returns for all libraries in our system over time. With OR and RR, we can assess this using confidence intervals, which are most frequently calculated as 95% confidence intervals, but could also be at other levels, such as 90% or 99%.

A confidence interval (or CI) represents uncertainty about where the true value lies, where the true value refers to the OR or RR for the broader population we're trying to understand. A 95% confidence interval tells us that we can be 95% confident that the true value falls somewhere within this range.

For the first library (shown in Table 2), the 95% confidence interval for the odds ratio (1.28) is

0.93-1.76. This range includes 1, which means that the result is not statistically significant. Recall that an OR or RR equal to 1 means that there is no association between the exposure and the event. If our CI includes 1, this means it is possible that, for our broader population (in this case, all system libraries over time), the true value is 1 and therefore there is no association between exposure and outcome. In general, when a 95% CI includes 1, the p-value is greater than 0.05; when it does not include 1, the p-value is less than 0.05.

The width of the CI can also provide us with some insight[5]. Where the CI is narrow (meaning that the range between the lower CI and upper CI is small), this indicates greater precision, while a wider CI would indicate lower precision. In this context, precision refers to how much uncertainty there is in the estimate. High precision indicates that there is a smaller range of plausible values and that we are confident that the true value lies within this smaller range. Higher precision typically comes from larger sample sizes or less variability in the data.

Beyond statistical significance and precision, it's also important to remember that a result might be statistically significant but may not be important from a practical standpoint. For example, an OR of 1.05 with a CI of 1.01-1.09 is significant. However, knowing that the odds are 5% higher (and that the increase in odds is likely between 1% and 9%) may not be large enough to warrant the investment of time, energy and resources to implement a particular initiative.

Conclusion

Understanding the distinctions between odds ratios and risk ratios, including when they should be used and how to interpret them, can equip librarians as information consumers, helping us to critically evaluate research findings, and can strengthen our practice by enabling us to make evidence-informed decisions based on data. However, while statistics can provide valuable information, they should always be considered alongside the practical significance and local context.

References

1. Bakker CJ. Statistics for librarians: Probability and odds. *Hypothesis*. 2025;37(1). doi:10.18060/28737

2. Tenny S, Hoffman MR. *Odds Ratio*. In: StatPearls. StatPearls Publishing; 2023. http://www.ncbi.nlm.nih.gov/books/NBK431098/

3. Tenny S, Hoffman MR. *Relative Risk*. In: StatPearls. StatPearls Publishing; 2023. http://www.ncbi.nlm.nih.gov/books/NBK430824/

4. Hancock M, Kent P. Interpretation of dichotomous outcomes: Risk, odds, risk ratios, odds ratios and number needed to treat. *Journal of Physiotherapy*. 2016;62(3):172-174. doi:10.1016/j.jphys.2016.02.016

5. Hazra A. Using the confidence interval confidently. *Journal of Thoracic Disease*. 2017;9(10):4125-4130. doi:10.21037/jtd.2017.09.14