Harm and Causation: Assessing the Value of Studies of Harm

6

Daniella A. Zipkin and Jeffrey Kushinka

Guide for the Teacher

Harm and causation questions make up a significant portion of the medical literature, in the form of cohort and case-control studies. Teaching harm can work well in the beginning of an evidence-based medicine course, particularly during discussions of study design. The topic flows well from a review of the strengths and weaknesses of cohort and case-control studies. Harm and causation questions come up frequently in various clinical settings. In addition, because results of studies that show negative associations tend to be highlighted often by news outlets and social media, patient exposure to the results generated by these types of studies tends to be high. We recommend covering the following components when teaching harm and causation:

1. Framing a harm or causation question.
2. Selecting the optimal study design. Study selection is covered separately in Chap. 3.
3. Assessing the risk of bias in cohort and case-control studies.
4. Calculating relative risk (when possible) and odds ratios for studies of harm and causation.
5. Describing the appropriate use of odds ratios and their limitations.

D. A. Zipkin (✉)
Department of Medicine, Duke University Health System,
Duke University School of Medicine, Durham, NC, USA
e-mail: daniella.zipkin@duke.edu

J. Kushinka
Department of Internal Medicine, Virginia Commonwealth University School of Medicine,
Richmond, VA, USA
e-mail: Jeffrey.kushinka@vcuhealth.org

D. A. Zipkin (ed.), *Teaching Evidence-Based Medicine*,
https://doi.org/10.1007/978-3-031-11174-7_6

6. App ying the results of harm trials to individual patients.
7. Communicating the results of harm trials to patients.

For each of these sub-topics you will find:

- Core content handout—we recommend learners read ahead of class.
- Samples of articles and accompanying worksheets for exercises to do together during teaching.
- Supplementary material in some cases.
- Links to videos with examples of real time teaching.

While framing the question and selecting the design can be taught in a brief introduction (under 15 min), each of the other topics may require an hour—risk of bias, odds ratios vs. relative risk (risk ratios), applying results, and communicating results.

Study Design for Harm or Causation Questions

Determining harm or causation requires investigating associations between exposures and outcomes. Different types of studies can provide information regarding these associations. *Randomized controlled trials* are the best studies for evidence of causation, because unmeasured variables which may impact the associations will be randomly distributed throughout the subjects. However, it is not often that a randomized trial will detect unexpected harm, and naturally it is unethical to plan an RCT when harm is expected (unless you plan to *reduce* a known harm). In addition, randomized controlled trials may not be designed with a follow up period which is long enough to detect the emergence of relevant harms.

The next best study design would be a *cohort study*, where a group with the exposure or treatment is compared to a group without the exposure, and followed prospectively. However, this study design is weaker, because we are unable to control for factors that influenced who received the exposure [1]. These factors, known as *confounders,* may be driving the apparent associations, with the exposure being investigated actually having little or no impact on the outcome in question. (see a full discussion of confounding on the following page). The manner in which exposed and unexposed subjects are selected is a big determinant of confounding. *Selection bias* results when the study sample does not represent the target population because of the site of recruitment or differences in baseline demographic factors [2]. Cohort studies may suffer from *detection bias*, or the tendency to look more closely for an outcome in one group over another, based on exposure (i.e., if we look more frequently in the exposure group and therefore find an association more frequently, how do we know we are not missing the same association in the control group?). Cohort studies may also be subject to *outcome ascertainment bias*, or the tendency to identify an outcome differently in each group being compared. This can occur if we define the outcome differently in the groups, or if we look for the outcome differently in the groups.

Case-control studies begin with gathering two groups of patients based on outcome status—one group with the outcome of interest, and one without—and then looking retrospectively to determine the degree of the exposure in each group. Because the selection of patients who have had the outcome (cases) and who have not had the outcome (controls) can impact all subsequent investigation about determining the potential exposure, case-control studies are prone to a number of biases. Among these are *selection bias*, *recall bias*, and *interviewer bias*, as all case-control studies require looking back in time to assess the degree of exposure. Because of the high risk of bias in case-control studies, they should be reserved for situations where the outcome is *rare*, making a prospective cohort or randomized trial not feasible. Case-control studies should be considered "hypothesis generating," and should lead to more rigorously designed studies to confirm the findings whenever possible.

Finally, the weakest form of evidence about causation is the *case series*. This is simply a series of cases where it was noted that an exposure and an outcome had occurred, and there is no comparison group. Case series are similar to case-control studies in that they are only useful in generating hypotheses that may lead to more rigorous studies.

Confounding Variables

Confounding can be caused by any factor that is associated with both the exposure and the outcome of interest, as depicted in Fig. 6.1. Confounders may be "silently" influencing the outcome more than the exposure being studied. In order to address confounding, one must be able to:

1. Think of all potential confounders,
2. Measure the confounders to the greatest degree of accuracy possible,
3. Input a numeric value for each confounder in a complex equation known as regression analysis (while performing regression analyses is beyond the scope of

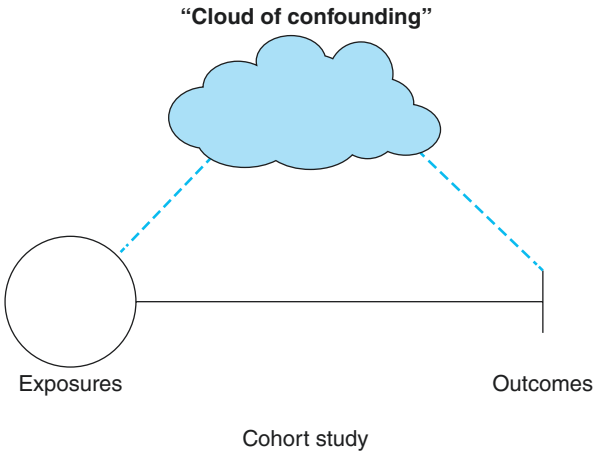

Fig. 6.1 Cohort with cloud of confounding. Confounding can be represented by a "cloud" that hovers over ever cohort. Confounders are variables which correlate with both the exposure and outcome and may be silently driving the observed association

this text, we will review important features of regression analyses to watch for as you interpret studies).
4. Assess whether the relationship between exposure and outcome persists even after adjusting for each confounder, one at a time, in the regression analysis.

Imagine a ridiculous situation with an exposure we know to be harmful: cigarette smoking. Imagine investigators wanted to assess whether cigarette smoking was a contributor to cirrhosis. Imagine these investigators found an association between cigarette smoking and cirrhosis but failed to account for the amount of alcohol consumption in these patients. We would be missing a major variable that is likely to be driving the outcome in question and lead to a potentially spurious association between cigarette smoking and cirrhosis.

Many confounders in psychosocial domains are important drivers of associations and cannot be reliably measured. Behavioral confounders, such as dietary habits, exercise, optimism, self-care, and utilization of support services are good examples of this. As a result, even studies which adjust for multiple confounding variables can never eliminate all confounding. The only way to eliminate the impact of confounding is to conduct a randomized trial, where randomization evenly distributes all confounders across the groups, including the confounders we cannot think of or cannot measure.

Assessing Bias in Studies of Harm or Causation

We take all of the factors above into account when assessing the extent of bias in a study of harm or causation (Box 6.1).

Box 6.1 Assessing the Risk of Bias in a Cohort Study

- Assessing the risk of bias in a cohort study:
  - Was the study population representative of the target population?
  - Were patients similar with respect to risk factors for the outcome, aside from the exposure of interest?
  - Was statistical adjustment for confounding variables described clearly and include all important variables?
  - Were outcomes explored in a similar fashion in exposed and non-exposed subjects?
  - Was follow-up time long enough for important outcomes to have emerged?
- Assessing the risk of bias in a case-control study.
  - Were cases and controls gathered from the same population?
  - Were cases and controls matched with regard to socio-demographic variables and clinical variables known to impact the likelihood of exposure?
  - Was the detection of exposures reliable and carried out in the same way in both groups?

TEACH IT!!

Bias in Studies of Harm or Causation

15 min:

Discussion of bias in harm studies centers around confounding. Start with a diagram on the board of a schematic of a cohort study, a circle on the left, with a horizontal line extending forward in time to the right, moving from exposures to outcomes. Above the cohort, consider drawing a "cloud" to represent confounding. The cloud should connect to both the exposures side and the outcomes side via dotted lines (Fig. 6.1).

Give the group a simple example to examine: for instance, there are cohort studies which have found associations between diet soda consumption and diabetes [3]. This isn't immediately intuitive, given the lack of sugar in diet soda, and it's not clear if there is a compound in diet soda causing the association. This example is on an accessible and familiar topic, allows the group to think about confounding, and lends itself well to a review of different sorts of confounders. Have the group brainstorm confounders in this situation—invariably, some version of the following will emerge: pre-existing obesity, socioeconomic status, diet/fast food intake, social groupings, personality factors not otherwise specified, etc.

With some confounders listed in your "cloud", remind the group that in order to adjust for confounding we must think of the confounders, measure them, and plug them into a mathematical formula called regression analysis. How easily can we measure obesity? Quite easily, use the body mass index! How easily can we measure dietary intake? This turns out to be much harder. Estimates about dietary intake are notoriously fraught with inaccuracies due to self reporting and the social expectations about dietary intake. How about social groupings? Even tougher—we can't measure that well at all.

Conclude with the point that we can NEVER eliminate confounding altogether. Studies do their best in identifying the most important confounders and adjusting for them, but it's not a perfect process. One should assume every cohort has residual confounding.

Explore other sources of bias through discussion:

What exposures would lead to more interactions with the health care system, for instance, and therefore a greater likelihood that the outcome of interest will be discovered? This is an example of detection bias.

Did outcomes have enough time to develop? If time was insufficient, and fewer outcomes are found, this will lead to imprecise estimates due to greater random error with smaller numbers.

30 min:

Add examples to the above discussion, and pre-select one or two studies for which the group can do a full assessment of bias.

If time is short, one option is to "pre-digest" the paper, highlighting key paragraphs where the answers on bias can be found. If you have more time, allow learners to read the article on the spot and then discuss.

Touch on the following concepts, through discussion as a group, and have learners take turns speaking:

Who were the patients? Do they represent a group in which this questions is important?

Were confounders fully assessed and adjusted for?

Was the outcome equally likely to be detected in those with and without the exposure?

Was follow-up complete, and was the time frame sufficient to see important outcomes?

You can use Worksheet 6.0, in the Appendix, as a guide for critical appraisal for learners.

For additional techniques on teaching adjustment for confounding, we recommend the Teaching Tips article entitled "Tips for Teachers of Evidence-Based Medicine: Adjusting for Prognostic Imbalances (Confounding Variables) in Studies on Therapy or Harm" [4].

Harm Math and the Magnitude of Association

In this section, we assess the strength of the association between exposure and outcome, ask whether or not a dose-response relationship exists, and look at the precision of the estimate. A dose-response relationship means that the magnitude of the association increases with increasing "dose," or amount, of the exposure. Precision refers to the confidence interval around the point estimate—the larger the confidence interval, the greater the variability and uncertainty of the estimate, and the lower the precision.

Definitions

$$\textbf{Odds} = \frac{(N \text{ with event})}{(N \text{ without event})}$$

$$\textbf{Risk} = \frac{(N \text{ with event})}{(\text{total } N)}$$

Table 6.1 The 2 × 2 table

	Outcome present	Outcome absent
Exposed	a	b
Not exposed	c	d

Results will usually be presented as a relative risk (RR, same as "risk ratio"), hazard ratio (HR, a more sophisticated risk ratio which accounts for changes in event accrual in the studies over time), or odds ratio (OR). We will tackle hazard ratios a little later. For now, it is important to review the differences between odds and risk, because odds ratios are always at least a little bit inflated compared to risk ratios. We intuitively think in terms of risk, so this inflation could prove deceptive when we interpret study results.

The RR is the risk in the exposed group divided by the risk in the unexposed group. The OR is the odds in the exposed group divided by the odds in the unexposed group. These are both ratios, so a value greater than 1 represents an increase in risk or odds, less than 1 a decrease in risk or odds. Remember the basic definitions of risk and odds as we move forward, beginning with Table 6.1.

Calculations
RR = [a/(a + b)]/[c/(c + d)]
OR for cohort studies (prospective: odds of outcome given certain exposure) = [a/b]/[c/d] or ad/cb
OR for case-control studies (retrospective: odds of exposure given certain outcome) = [a/c]/[b/d] or ad/cb
[Notice that mathematically, these OR calculations start in different places, but come out to be the same!]

Odds Are Always Larger Than Risk!
Odds ratios will most closely approximate risk ratios when:

- The event or outcome is rare
- The risk difference is small
- The study is large.

Compromising on these factors will cause the odds ratio to start deviating from the risk ratio, often by an unacceptably large gap.
Odds ratios are also appropriate for:

Case-control studies—When outcomes and exposures are dichotomous (i.e., they are either present or absent), they lend themselves well to calculation of odds ratios. In addition, with case-control studies, the concept of "total N,"

the denominator in a risk ratio, is not applicable, because we recruited an arbitrary number of study participants to make up that population.

Regression analyses—the statistical process of evaluating predictors of an outcome works best with odds ratios, because odds ratios can be multiplied and inserted into complex mathematical formulas easily. The output of a regression analysis will be an odds ratio, but authors can then choose to convert it to a risk ratio for publication. While the details of conducting regression analysis are beyond the scope of this text, knowing that the process occurs with odds ratios helps to explain why some prospective cohorts will present odds ratios for their main outcomes. One must ask why they did not convert back to risk ratios—is it possible that the inflated number suited their aims more?

How do we interpret an odds ratio or risk ratio once it is calculated? Think of it as a relative increase or decrease in odds. The math here is similar to what we reviewed in the Therapy chapter. Keep in mind that any ratio (risk ratio or odds ratio) of 1 means there is no difference between the groups being compared. Therefore, for any ratio not equal to 1, the distance from 1 tells us the relative odds or relative risk. For instance, if the odds ratio is 1.3, that represents a 30% relative increase in odds. We can't make sense of this number without knowing the baseline risk for the condition. Imagine the baseline risk is 2%. A 30% increase in that risk would move the risk from 2% to 2.6%. Thus, it is important to bring the relative change back to absolute terms. Please see Chap. 4 for an explanation of these concepts. The *number needed to harm* can be calculated in the same way as the number needed to treat—is it simply the reciprocal of the absolute risk increase. It should be noted that we cannot calculate a number needed to harm for case-control studies because they are retrospective and reflect an arbitrary number of subjects.

The *precision* of these estimates can be assessed by examining the confidence interval around the estimate. In a study which demonstrates an association between an exposure and an adverse outcome, the lower limit of the CI provides a minimal estimate of the strength of the association. In a study which has failed to demonstrate an association, the upper boundary of the CI tells you how big an adverse effect may still be present, despite the failure to show a statistically significant association.

Factors that influence *clinical decision-making* regarding harm include the strength of the association, the magnitude of the risk, the available alternatives, and the possible adverse consequences of minimizing exposure. If there is significant bias in the study design and the association is weak (OR of less than 2.0), then it is probably best to wait for other data to confirm and strengthen the finding. Nonetheless, once even a small possibility of harm exists, the ethical, legal, and societal impact may trump the evidence. Health systems may need to act on the potential harm even if "truth" has not been confirmed.

TEACH IT!!

Harm Math and Applying Results to Patients

15–20 min:

Have learners fill out the Worksheet 6.1, available in the Appendix. This compares odds and risk for different shaded portions of the pie chart. Discuss with them, and be sure they notice that as the proportion of the shaded area gets larger, odds and risk diverge more and more. The answer key is provided in Worksheet 6.2.

Move on to Worksheet 6.3 attached at the end of this chapter, or provide a similar example utilizing simple numbers. A humorous scenario never hurts! As learners move through the calculations, make the following observations as a group. We provide answers to this imaginary scenario in Worksheet 6.4 for reference.

Odds ratios, like Odds, differ from Risk Ratios when event rates are large.

Odds ratios are more inflated compared to Risk Ratios when the risk difference is larger.

For case-control studies, you cannot calculate risk ratios.

The mathematical result for a cohort study vs. a case-control study for the same dataset will be numerically the same. What differs is how you say it. Have learners practice putting the odds ratio into a sentence for both a cohort study and a case-control study. For instance, say the odds ratio is 2.5. In a cohort study, you might say "the exposed group had 2.5 times greater odds of having the outcome than the non-exposed group". In a case-control study, you might say "those with the outcome had a 2.5 times greater odds of having been exposed than those without the outcome".

10–15 min:

Follow the exercise above by looking at real world examples and interpreting the magnitude of the results. For this portion, it is ok to utilize abstracts only, rather than the full studies, because you'd like the group to look at the results and imagine how to communicate them to patients. For this exercise, you can assume the risk of bias in the selected papers was low and move straight to results.

Discuss the odds ratio or hazard ratio presented in the abstract and put it into words. This is a relative number—i.e., a "relative increase in odds" or a "relative increase in risk."

Provide a patient case around the study of interest, and estimate that patient's baseline risk for the condition. Utilize medical databases, or clinical judgment, with the scenario you have set up.

Remember that odds ratios and hazard ratios are relative increases or decreases in risk. This means that in order to assess the real magnitude, you need to determine the baseline odds or risk of the outcome, and then multiply by the relative change reflected in the ratio.

Example: a 42-year-old man with no medical history aside from persistent gastroesophageal reflux has been stable on a proton-pump inhibitor (PPI) for several years. He recently learned that the PPI was associated with kidney failure based on a news report of a new study [5] *and stopped taking it. His acid reflux symptoms are severe again. How can you counsel him?*

> *A quick read of the abstract tells you that this well-done cohort study found an association between PPI use and incident chronic kidney disease in adults aged 63 on average. Adjusted analyses found a HR of 1.50* [5]*. How do you apply this to your patient?*

Discuss what you've learned: if we trust this study, what is the magnitude of the impact for our patient? Many studies of harm report small to moderate odds ratios. Relative increases to harms with low baseline risks will result in very small changes. These changes may or may not impact how we counsel patients about these harms! This is particularly true when all cohort studies and case-control studies struggle with bias, and these results may be subject to error. This may be an appropriate place to remind the learners about several key points: the discussion of harms in the lay media is often rather alarmist and overstates the impact of relatively small odds ratios and risk ratios, and these studies identify associations only and are NOT proof of cause and effect.

Appendix

Worksheet 6.0—Critical appraisal for studies of harm or causation

Harm and causationcriticalappraisal worksheet (cohort and case control)

Assessing the risk of bias in a cohort study	
Was the study population representative of the target population?	
Were patients similar with respect to risk factors for the outcome, aside from the exposure of interest?	
Was statistical adjustment for confounding variables described clearly with all important variables included?	
Were outcomes explored in a similar fashion in exposed and non-exposed subjects?	
Was follow-up time long enough for important outcomes to have emerged?	
Assessing the risk of bias in a case control study	
Were cases and controls gathered from the same population?	
Were cases and controls matched with regard to socio-demographic variables and clinical variables known to impact the likelihood of exposure?	
Was the detection of exposures reliable and carried out the same way in both groups?	
Assessing the results in a cohort or case control study	
What is the magnitude and precision of the results?	
Can the results be applied to my patient population?	

Worksheet 6.1—Odds exercise, blank

Risk vs. Odds comparison

$$\text{Risk} = \frac{\text{N\_event}}{\text{Total N}}$$

$$\text{Odds} = \frac{\text{N\_event}}{\text{N\_without\_event}}$$

As the frequency of an event increases, what do you notice about the risk and the odds of that event?

	Risk	Odds
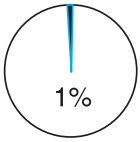		
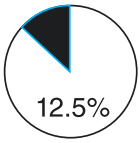		
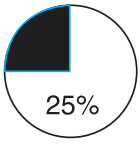		
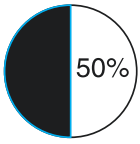		
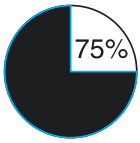		

Worksheet 6.2—Odds exercise, answers

Risk vs. Odds comparison- Answers

$$\text{Risk} = \frac{\text{N\_event}}{\text{Total N}}$$

$$\text{Odds} = \frac{\text{N\_event}}{\text{N\_without\_event}}$$

As the frequency of an event increases, odds and risk diverge more and more, with odds becoming unacceptably inflated. We should only rely on odds for rare events.

	Risk	Odds
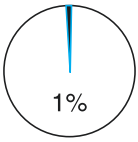	$\frac{1}{100} = 0.01$	1:99 = 0.010101…
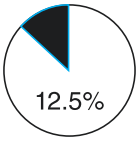	$\frac{12.5}{100} = 0.125$	12.5:87.5 = 0.143
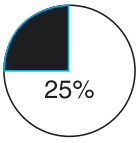	$\frac{25}{100} = 0.25$	25:75 = 0.33
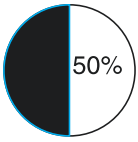	$\frac{50}{100} = 0.5$	50:50 = 1
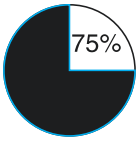	$\frac{75}{100} = 0.75$	75:25 = 3

Worksheet 6.3—Odds ratio exercise, blank

Odds Ratios can be inflated too!

The Case: You are at a large family picnic when your second cousin once removed, Jack, comes up to you and asks your help settling a medical question. He thinks that the best way to avoid a major belly ache after the picnic is to eat only hot dogs, which are a new offering this year, but your great aunt Millie thinks it's best to eat only hamburgers like they did in the old days. Because you are an ambitious student of EBM, you immediately design a study. You enroll everyone at the picnic who chose either hot dogs or hamburgers, and you arrange to call everyone the next morning to gather data…

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dogs	40	60
Control: Hamburgers	20	80

Relative Risk (RR) of Hot Dog consumption leading to belly ache:

RR = Risk with Hot Dog/Risk with Hamburger =

Odds Ratio (OR) of Hot Dog consumption leading to belly ache:

OR = Odds with Hot Dog/Odds with Hamburger =

What do you notice?

You decide to look at this important question again at next year's picnic, because your EBM teacher said all studies should be replicated! Here is the data gathered this time:

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dog	30	70
Control: Hamburger	25	75

Relative Risk (RR):

RR =

Odds ratio (OR):

OR =

What do you notice, now? What happened as we narrowed the risk difference?

In year three, you get a bit of a late start planning, so you decide to perform this research in a different way. You get the full family phone list the day after the picnic, and then call each person to ask them if they have a belly ache, and what they ate.

What is the name of this study design? _____

Shockingly, your numbers come out looking very similar to the last set. Recalculate the Odds Ratio only, now from the vantage point of this new study design, moving backwards from outcomes to exposures. **(Why are we only calculating the Odds Ratio, and not a Risk Ratio this time?)**

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dogs	30	70
Control: Hamburgers	25	75

Odds ratio (OR):

OR = Odds in belly ache group/Odds in feeling great group =

What do you notice?

Try to speak the results of this study in a sentence! Compare that to how you would speak the results of the previous study.

Worksheet 6.4—Odds ratio exercise, answers

Odds Ratios can be inflated too!

The Case: You are at a large family picnic when your second cousin once removed, Jack, comes up to you and asks your help settling a medical question. He thinks that the best way to avoid a major belly ache after the picnic is to eat only hot dogs, which are a new offering this year, but your great aunt Millie thinks it's best to eat only hamburgers like they did in the old days. Because you are an ambitious student of EBM, you immediately design a study. You enroll everyone at the picnic who chose either hot dogs or hamburgers, and you arrange to call everyone the next morning to gather data…

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dogs	40	60
Control: Hamburgers	20	80

Relative Risk (RR) of Hot Dog consumption leading to belly ache:

RR = Risk with Hot Dog/Risk with Hamburger = (40/100)/(20/100) = 2.0

Odds Ratio (OR) of Hot Dog consumption leading to belly ache:

OR = Odds with Hot Dog/Odds with Hamburger = (40/60)/(20/80) = 2.67

What do you notice? The estimate is inflated when using odds! Why?
Because it's not a rare event, and it's not a large population.

You decide to look at this important question again at next year's picnic, because your EBM teacher said all studies should be replicated! Here is the data gathered this time:

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dog	30	70
Control: Hamburger	25	75

Relative Risk (RR):

RR = (30/100)/(25/100) = 1.2

Odds ratio (OR):

OR = (30/70)/(25/75) = 1.29

What do you notice, now? What happened as we narrowed the risk difference?
Still inflated, but less so. Why? Even with a relatively common event, risk difference will also play into how inflated the odds ratio can get.

In year three, you get a bit of a late start planning, so you decide to perform this research in a different way. You get the full family phone list the day after the picnic, and then call each person to ask them if they have a belly ache, and what they ate.

What is the name of this study design? Case-control!

Shockingly, your numbers come out looking very similar to the last set. Recalculate the Odds Ratio only, now from the vantage point of this new study design, moving backwards from outcomes to exposures. **(Why are we only calculating the Odds Ratio, and not a Risk Ratio this time?** Because we cannot calculate Risk in a case control – there is no "total N" for the denominator, we selected the population ourselves…**)**

	Outcome: Belly Ache	No Outcome: Feeling great!
Exposure: Hot Dogs	30	70
Control: Hamburgers	25	75

Odds ratio (OR):

OR = Odds in belly ache group/Odds in feeling great group =
(30/25)/(70/75) = 1.29

What do you notice? 1.29 is exactly the same as the last time was calculated this!

It turns out, the math doesn't change if you are a cohort study or a retrospective case control study. It is the same either way – all that changes is how you communicate those results.

Try to speak the results of this study in a sentence! Compare that to how you would speak the results of the previous study.

In the case control version, we say "those with belly ache had a 1.29 greater odds of having eaten hot dogs."

In the cohort version on the previous page, we say "those who ate hot dogs had a 1.29 greater odds of having a belly ache."

References

1. Grimes DA, Shulz KF. Bias and causal associations in observational research. Lancet. 2002;359:248–52.
2. Ellenberg JH. Cohort studies: selection bias in observational and experimental studies. Stat Med. 1994;13:557–67.
3. Nettleton JA, Lutsey PL, et al. Diet soda intake and risk of incident metabolic syndrome and type 2 diabetes in the Multi-Ethnic Study of Atherosclerosis (MESA). Diabetes Care. 2009;32(4):688–94.
4. Kennedy CC, Jaeschke R, et al. Tips for teachers of evidence-based medicine: adjusting for prognostic imbalances (confounding variables) in studies on therapy or harm. J Gen Intern Med. 2008;23(3):337–43.
5. Lazarus B, Chen Y, et al. Proton pump inhibitor use and the risk of chronic kidney disease. JAMA Intern Med. 2016;176(2):238–46.