PHS2000B Lab 3

Interaction

02/18/2020

Contents

1 What is interaction and why do we study it?

Up to this point in the semester, our primary focus has been on average causal effects in a population —how we define them, the assumptions necessary to estimate them (positivity, exchangeability, and consistency), and what we can still learn about them when some of those assumptions aren't met (e.g. via sensitivity analyses); however in population health we often have questions that go far beyond just the average causal effect. We also want to know things like the subsets of the population for whom the effect occurs (or doesn't occur) and how the effect changes in response to other exposures or interventions. That is we want to be able to answer questions like:

- Does the effect of statins on cardiovascular risk differ among men and women?

- Is there a combined effect of air pollution and smoking on the risk of lung canccer such that double exposure is higher than the effect of each individually?
- Is the risk of hospitalization lower for children infected by the novel coronavirus (2019-nCoV) as compared to adults?
- Is there a social gradient in the effects of breast screening on breast cancer mortality?

In order to answer these sorts of questions we need a theory that goes far beyond simple causal contrasts between dichotomous treatments and incorporates more of the heterogeneity we experience in the real world. In this section, we'll explore the concept of *interaction* and see how it relates to the causal framework we've defined previously. We'll also discover that the traditional notion of *interaction* covered in most introductory statistics courses actually includes quite distinct causal phenomena.

2 Measures of interaction

In most introductory statistics/epidemiology courses, there is some discussion of the possibility that measures of association between an outcome (Y) and exposure (A) may differ across levels or strata of another variable (V). This is generally presented as *an interaction between A and V*. As a simple example, consider the relationship between smoking, asbestos exposure, and lung cancer from Hilt et al. (1986) below:

Table 1: Risk of lung cancer by smoking and asbestos status.

	Asbestos	No Asbestos
Smoker	0.0450	0.0095
Non-smoker	0.0067	0.0011

Clearly there are differences in lung cancer risk across strata of smoking and asbestos exposure with the absolue risk being the highest in the doubly exposed group and lowest in the doubly unexposed group.

2.1 Additive measures of interaction

We could summarize this relationship between smoking and asbestos on risk of lung cancer using several metrics. One way might be to calculate the extent to which the effect in the doubly exposed exceeds the effect among either of the singly exposed categories, e.g. whats the difference in risk between those exposed to smoking and asbestos versus those exposed to only asbestos or only smoking. Using the numbers above this would be given by:

$$(0.0450 - 0.0011) - [(0.0095 - 0.0011) + (0.0067 - 0.0011)] = 0.0299$$

where the values suggests there is an excess risk of from being exposed to both smoking and asbestos of about 3 percentage points over what we would expect if we just added the individual effects of asbestos and smoking. This is an example of an *additive* interaction because it's measured on the absolute difference scale. In general for two binary exposures and a binary outcome the calculation above can be simplified to

$$p_{11} - p_{10} - p_{01} + p_{00}.$$

A positive result implies a positive or "super-additive" interaction (meaning that double exposure contributes excess risk over single exposure) while a negative result implies a negative or "sub-additive" interaction (meaning that double exposure reduces risk over single exposure).

2.2 Multiplicative measures of interaction

Alternatively, we could examine whether the effect in the doubly exposed exceeds the effect among either of the singly exposed categories on the relative scale. For example using smoking and asbestos data we could compare the relative risk in the doubly exposed to the product of the relative risks in the singly exposed, e.g.

$$\frac{\frac{0.0450}{0.0011}}{\frac{0.0095}{0.0011} \cdot \frac{0.0067}{0.0011}} = 0.7776$$

This time the value implies that there is about 22% lower risk from being exposed to both smoking and asbestos over what we would expect from the product of the individual effects of asbestos and smoking. This is an example of a *multiplicative* interaction because it's measured on the relative scale (rather than on the absolute difference scale). Using risk ratios, it's given in general by

$$\frac{RR_{11}}{RR_{10}RR_{01}}$$

for two binary exposures and a binary outcome. A value greater than one implies a positive multiplicative interaction between the exposures while a value less than 1 implies a negative interaction. Beyond the risk ratio, there are several other ratio effect measures common in epidemiology which can also be used to assess for interaction on the multiplicative scale like the odds ratio, the hazard ratio, and the incidence rate ratio.

$$\frac{OR_{11}}{OR_{10}OR_{01}}$$
$$\frac{HR_{11}}{HR_{10}HR_{01}}$$
$$\frac{IRR_{11}}{IRR_{10}IRR_{01}}$$

Be wary that each could produce slightly different results as they each assess interaction on their own scale and each measure has it's quirks[1].

2.3 Scale dependency

The example above nicely demonstrates the general rule that interaction is *scale dependent*. On the one hand, on the additive scale we have evidence that exposure to smoking and asbestos is associated with an increased risk over exposure to either alone, while on the other hand, on the multiplicative scale we have evidence that exposure to both is associated with decreased risk over exposure to either alone.

In general, interaction measured on one scale does not necessarily imply interaction on another. As further demonstration consider the table below which show results from two hypothetical studies.

Table 2: Risk of outcome Y by cross-classified exposure to A and V.

Study 1			Study 2		
	$A=1$	$A=0$		$A=1$	$A=0$
$V=1$	0.10	0.07	$V=1$	0.10	0.04
$V=0$	0.05	0.02	$V=0$	0.05	0.02

In the first study, there's no evidence of interaction on additive scale as $0.10 - 0.07 - 0.05 + 0.02 = 0$ but there is evidence for negative interaction on the multiplicative scale as $(0.10/0.02)/[(0.07/0.02)(0.05/0.02) =$

[1]For instance, if you recall from PHS2000A the odds ratio is not always collapsible across strata of other covariates which presents particular challenges when interest lies in comparing effects across strata (as in interaction analyses).

0.57. In the second study, this time there's evidence of positive interaction on the additive scale as $0.10 - 0.05 - 0.04 + 0.02 = 0.03$ but there is no evidence of interaction on the multiplicative scale as $(0.10/0.02)/[(0.05/0.02)(0.04/0.02) = 1$. Indeed, Greenland et al. (2008) show that, in general, if two exposures have effects on the outcome, then there must be interaction between them on some scale, i.e. if there's no interaction on the additive scale there must be an interaction on the multiplicative scale and vice versa.

How are we to reconcile findings from different scales when estimates diverge? Should we prefer the additive measure or the multiplicative? Like so many things in life, the answer depends on the question being asked. In lecture, we mentioned that multiplicative interaction was more consistent with how certain biological systems interact, but that the additive scale was the most natural scale when answering important public health questions like where to direct resources for intervention. However, even these rules of thumb are somewhat crude and fail to capture all possibilities. In general you're better off thinking hard about which scale is more relevant to your research question. When in doubt, you can and should present evidence on multiple scales.

2.4 Other measures of interaction

2.4.1 RERI

Sometimes we may also be interested in assessing additive interaction when only relative measures are available or reported (or as we'll see later when we want to transform the output from certain statistical models to the additive scale). Notice that if we take our original additive measure $p_{11} - p_{10} - p_{01} + p_{00}$ and divide through by p_{00} we get

$$\text{RERI} = RR_{11} - RR_{10} - RR_{01} + 1$$

This quantity is known as the *relative excess risk due to interaction* (RERI) and, like the former measure, implies positive additive interaction when RERI > 0 and negative additive interaction when RERI < 0. Under the rare outcome assumption, we could use odds ratios in place of the risk ratios above.

2.4.2 Attributable proportion

Another measure of additive interaction is the attributable proportion defined by

$$\text{AP} = \frac{RR_{11} - RR_{10} - RR_{01} + 1}{RR_{11}}.$$

It measures the proportion of risk in the doubly exposed group that is due to the interaction. You may notice that it is just the RERI scaled by RR_{11}. Like the RERI, the AP also implies positive additive interaction when AP > 0 and negative additive interaction when AP < 0.

3 Causal structures of interaction

In the previous section, we focused on different ways that we could quantify interaction (i.e. changes in measures of association across levels of another variable), but thus far we've said nothing about what these quantities imply about the world. What does it mean that an association varies across levels of another variable? Does that mean there is some synergy or joint effect of the two variables or does it mean there are simply differences in mechanism or response across subgroups due to some other source. Under what circumstances can I interpret any of these causally? In this section, we'll attempt to add structure to these

sorts of questions by using counterfactual notation to more rigorously define different sources of observed interactions.

3.1 Effect heterogeneity

In some settings, we may be interested in the effect of a treament or exposure A on an outcome Y but may be curious if the effect of A on Y is different in different subgroups of the population. For example, going back to one of our original questions, maybe we want to understand if the effect of statins on risk of cardiovascular mortality differs between males and females (perhaps because we have some biological theory for why it might). This might be a useful enterprise if we want to target treatment to women or if we have concerns about whom the pharmaceutical industry prioritizes. More formally, we want to take the average treatment effect of A on Y in the whole population, which we defined previously as:

$$\mathbb{E}[Y_{a=1} - Y_{a=0}]$$

and we want to parition into subgroups defined by some variable V and then examine whether the effect is the same across different levels of V,[2] e.g.

$$\underbrace{\mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 1]}_{\text{effect of } A \text{ in } V=1 \text{ group}} \stackrel{?}{=} \underbrace{\mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 0]}_{\text{effect of } A \text{ in } V=0 \text{ group}}$$

or

$$(\mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 1]) - (\mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 0]) \stackrel{?}{=} 0$$

We define this as effect heterogeneity although the same concept is variously called effect modification and moderation in different disciplines. In words, the expression above is asking whether the effect of A comparing the universe in which everyone in subgroup $V = 1$ was given $a = 1$ to the universe in which everyone in subgroup $V = 1$ was not given $a = 1$ is the same as that which would be observed in subgroup $V = 0$ comparing the universe in which everyone in that subgroup was given treatment versus the universe in which they were not. Notice this is on the additive scale, but we could ask a similar question on the multiplicative scale, i.e.

$$\frac{\mathbb{E}[Y_{a=1} \mid V = 1]}{\mathbb{E}[Y_{a=0} \mid V = 1]} \Big/ \frac{\mathbb{E}[Y_{a=1} \mid V = 0]}{\mathbb{E}[Y_{a=0} \mid V = 0]} \stackrel{?}{=} 1$$

Given that we are only concerned about the causal effect of intervening on A, we can estimate these effects under the same set of assumptions we defined previously with respect to A, i.e. exchangeability, consistency, and positivity. Under these conditions we can evaluate whether there is effect heterogeneity by

$$\mathbb{E}[Y \mid A = 1, V = 1] - \mathbb{E}[Y \mid A = 0, V = 1] - \mathbb{E}[Y \mid A = 1, V = 0] + \mathbb{E}[Y \mid A = 0, V = 0] \stackrel{?}{=} 0$$

which you may notice bears a striking resemblance to $p_{11} - p_{10} - p_{01} + p_{00}$ when Y is binary, and that's because under the causal assumptions above we can evaluate whether there is effect heterogeneity by checking whether there is an interaction between A and V.

For a more concrete example, consider the following data from a hypothetical randomized trial in which we've randomized participants to A and measured their outcomes Y as well as some data on baseline covariate V. In this case, randomization implies that we have marginal exchangeability with respect to A (which you can verify using the counterfactual outcomes) and therefore we can identify the effects of A on Y from the observed data by subtracting the means in the $A = 1$ and $A = 0$ groups.

[2] Note that here we should also specify that V here is generally a pre-treatment variable, i.e. it temporally proceeds A, because otherwise it's less clear if it's an effect modifier or a mediator, a distinction which we'll cover in more detail later.

Table 3: Example effect heterogeneity by V

ID	A	V	Y	$Y_{a=0}$	$Y_{a=1}$
1	0	0	0	1	0
2	0	0	0	1	0
3	0	0	0	1	0
4	0	0	0	1	0
5	0	1	1	0	1
6	0	1	1	0	1
7	0	1	1	0	1
8	0	1	1	0	1
9	1	0	1	1	0
10	1	0	1	1	0
11	1	0	1	1	0
12	1	0	1	1	0
13	1	1	0	0	1
14	1	1	0	0	1
15	1	1	0	0	1
16	1	1	0	0	1

Using the observed data $\mathbb{E}[Y \mid A = 1] - \mathbb{E}[Y \mid A = 0] = \frac{4}{8} - \frac{4}{8} = 0$ which is the same as the difference in potential outcomes $\mathbb{E}[Y_{a=1} - Y_{a=0}] = \frac{8}{16} - \frac{8}{16} = 0$ and suggests there is no average effect of A on Y. However, if we look within strata defined by V we find that among those with $V = 1$ the effect is $\mathbb{E}[Y \mid A = 1, V = 1] - \mathbb{E}[Y \mid A = 0, V = 1] = \frac{4}{8} - \frac{0}{8} = 0.5$ and among those with $V = 0$ the effect is $\mathbb{E}[Y \mid A = 1, V = 0] - \mathbb{E}[Y \mid A = 0, V = 0] = 0 - \frac{0}{8} = -0.5$ implying we find equal and opposite effects of A on Y in the subgroups. Comparing this again to the potential outcomes we see that we would observe the same phenomena among those with $V = 1$ and $V = 0$ if we could observe both their counterfactual outcomes.

In the observational setting, marginal exchangeability of exposure/treatment is often untenable and so instead we often consider conditional exchangeability within levels of a set of covariates C representing likely confounders. Thus when evaluating effect heterogeneity in the observational setting, we generally have to first condition[3] on a set of variables C to make the conditional exchangeability assumption plausible prior to evaluating whether the the effect of A varies across levels of V.

$$\mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 1, C = c] \stackrel{?}{=} \mathbb{E}[Y_{a=1} - Y_{a=0} \mid V = 0, C = c]$$

Notice that up to this point we have made no mention of why the effect of A on Y varies across levels of V. That's because the theory for heterogeneous effects doesn't require us to know precisely why it varies in order to verify that it varies. In fact, there are many possibile causal mechanisms that could lead to effect heterogeneity/modification by V. In their paper Vanderweele and Robins (2007) attempted to classify these into categories as shown below using DAGs. In some cases, the effect of A on Y differs by V because V has either a direct (a) or indirect (b) on Y, in other cases V is simply a proxy for another variable that has an effect on Y (e), or V could simply be related to Y through some unknown common mechanism U (c) or even a feature of some selection mechanism (d). The larger point is that the assumptions and process for estimating the effects will be the same in all these cases, but you as the researcher may have to think or need a strong theory to explain why the effect of A differs by V.

[3]Note that this means we have different heterogenous effects for each level of C, if we want to take these conditional effects and combine them to get single estimates in each level of V like we did above we need to use either standardization or IP weighting.

Figure 1: Different structures of effect modification/heterogeneity by V.

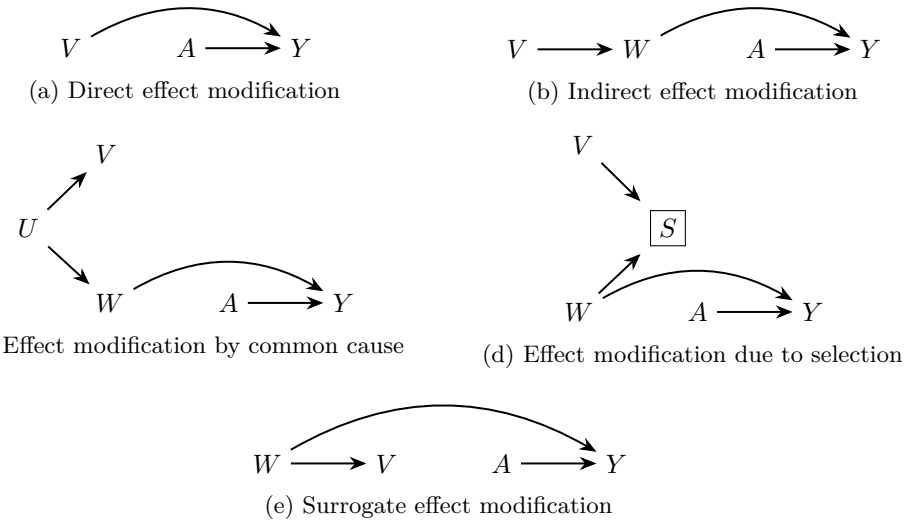

(a) Direct effect modification

(b) Indirect effect modification

(c) Effect modification by common cause

(d) Effect modification due to selection

(e) Surrogate effect modification

3.2 Causal interaction

In other settings, we may be interested in the joint effects of intervening on two treaments or exposures A and V as compared to intervening on one or the other individually. For instance, in the asbestos and smoking example we may wonder whether a joint intervention to both reduce asbestos exposure and encourage smoking cessation would reduce risk of lung cancer among miners more than either intervention would individually. Using potential outcome notation we want to estimate something like:

$$\underbrace{(\mathbb{E}[Y_{a=1,v=1}] - \mathbb{E}[Y_{a=0,v=0}])}_{\text{effect of joint intervention}} - \left(\underbrace{(\mathbb{E}[Y_{a=1,v=0}] - \mathbb{E}[Y_{a=0,v=0}])}_{\text{effect of intervening on } A \text{ alone}} + \underbrace{(\mathbb{E}[Y_{a=0,v=1}] - \mathbb{E}[Y_{a=0,v=0}])}_{\text{effect of intervening on } V \text{ alone}} \right) \stackrel{?}{=} 0$$

on the additive scale or likewise something like

$$\frac{\mathbb{E}[Y_{a=1,v=1}]}{\mathbb{E}[Y_{a=0,v=0}]} \Big/ \left(\frac{\mathbb{E}[Y_{a=1,v=0}]}{\mathbb{E}[Y_{a=0,v=0}]} \cdot \frac{\mathbb{E}[Y_{a=0,v=1}]}{\mathbb{E}[Y_{a=0,v=0}]} \right) \stackrel{?}{=} 1$$

on the multiplicative scale. In this case where both A and V are binary, we need to define four counterfactual outcomes:

1. $Y_{a=1,v=1}$ the outcome when both A and V are set to 1.
2. $Y_{a=1,v=0}$ the outcome when A is set to 1 and V is set to 0.
3. $Y_{a=0,v=1}$ the outcome when A is set to 0 and V is set to 1.
4. $Y_{a=0,v=0}$ the outcome when both A and V are set to 0.

And to identify effects from data we now need additional exchangeability assumptions about **both** A and V. That is we need both of the following conditions to hold:

$$Y_{a,v} \perp\!\!\!\perp A$$

$$Y_{a,v} \perp\!\!\!\perp V$$

Graphically this corresponds with d-separation requirements for both A and V, i.e. there should be no backdoor paths between A or V and the outcome Y.

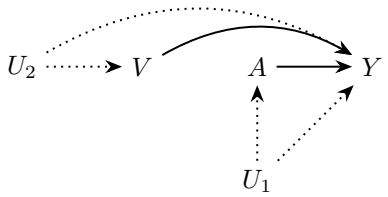

If these conditions hold and the normal positivity and consistency assumptions hold then the comparison of the joint effect to the effect of each exposure alone is given by

$$\mathbb{E}[Y \mid A = 1, V = 1] - \mathbb{E}[Y \mid A = 0, V = 1] - \mathbb{E}[Y \mid A = 1, V = 0] + \mathbb{E}[Y \mid A = 0, V = 0] \stackrel{?}{=} 0.$$

Once again, notice that this expression bears a striking resemblance to the one in the effect heterogeneity section and to our additive interaction expression $p_{11} - p_{10} - p_{01} + p_{00}$ when Y is binary. **They are all identified by the same expression in the data!** This is a really fundamental point that demonstrates the utility of counterfactual notation, the data alone can't tell us whether an interaction is an example of effect heterogeneity, causal interaction, or just a spurious association. Only our knowledge about the context of how the data were generated and the set of assumptions we are willing to make can differentiate between them.

Satisfying two sets of exchangeability assumptions may seem like a tall order (and it is). A canonical example of when this would be justified is a factorial RCT in which participants were randomized to received A and V, A alone, V alone, or neither. In the observational setting, we can again relax the assumptions to conditional exchangeability for A and V within some sufficient set of covariates C sufficient to control for confounding for both exposures.

3.3 Summary

The big picture of doing interaction analyses is that we are choosing a **statistical measure** (additive or multiplicative) and adding to it some **assumptions**, from which we can draw various **possible conclusions** about the presence of effect heterogeneity, causal interaction, or neither. Let "NUCA $\times$1" denote the assumption of no unmeasured confounding on only one of the two exposures, and let "NUCA $\times$2" denote the assumption of no unmeasured confounding on both exposures.

Type of statistical measure	**Assumptions**	**Possible conclusions**
Additive	NUCA $\times$1	Effect heterogeneity
	NUCA $\times$2	Causal interaction
Multiplicative	NUCA $\times$1	Effect heterogeneity
	NUCA $\times$2	Causal interaction

4 Interaction via statistical models

We saw in the first section that one way to estimate measures of interaction is to stratify by your two (or more) variables of interest and calculate relevant measures within each strata. While this may be sufficient for binary exposure variables often when we have more complex exposures or in the observational setting when we need to additionally stratify by another set of variables to ensure conditional exchangeability this quickly becomes impractical. Fortunately, we can use statistical models in these circumstances to estimate exact or approximate corrollaries to the interaction measures discussed previously. In general these measures will map onto product terms in our regression models.

4.1 Linear regression

For a continous outcome we can estimate the following model for A and V

$$\mathbb{E}[Y \mid A, V] = \beta_0 + \beta_1 A + \beta_2 V + \beta_3 AV$$

Using some basic manipulation, it can be shown that the coefficient on the product term β_3 corresponds to

$$\beta_3 = \mathbb{E}[Y \mid A = 1, V = 1] - \mathbb{E}[Y \mid A = 1, V = 0] - \mathbb{E}[Y \mid A = 0, V = 1] + \mathbb{E}[Y \mid A = 0, V = 0]$$

Which is just the same expression for additive interaction seen previously. Therefore an estimated β_3 that is greater than 0 suggests positive interaction and an estimated β_3 that is less than 0 suggests negative interaction.

4.2 Binomial regression

When the outcome is binary and we want to model the risk differences we might use a binomial model like

$$\Pr[Y = 1 \mid A, V] = \beta_0 + \beta_1 A + \beta_2 V + \beta_3 AV$$

Again the coefficient on the product term β_3 corresponds to

$$\begin{aligned}\beta_3 &= \Pr[Y = 1 \mid A = 1, V = 1] - \Pr[Y = 1 \mid A = 1, V = 0] - \Pr[Y = 1 \mid A = 0, V = 1] + \Pr[Y = 1 \mid A = 0, V = 0] \\ &= p_{11} - p_{10} - p_{01} + p_{00}\end{aligned}$$

which is the same expression for additive interaction.

4.3 Log-binomial regression

Alternatively, when the outcome is binary and we want to model the risk ratio we might use a log-binomial model like

$$\log \Pr[Y = 1 \mid A, V] = \beta_0 + \beta_1 A + \beta_2 V + \beta_3 AV$$

This time the **exponentiated** coefficient on the product term corresponds to

$$e^{\beta_3} = \frac{\frac{\Pr[Y=1|A=1,V=1]}{\Pr[Y=1|A=0,V=0]}}{\frac{\Pr[Y=1|A=1,V=0]}{\Pr[Y=1|A=0,V=0]}\frac{\Pr[Y=1|A=0,V=1]}{\Pr[Y=1|A=0,V=0]}} = \frac{RR_{11}}{RR_{10}RR_{01}}$$

which is now the measure of multiplicative interaction on the risk ratio scale.

4.4 Logistic regression

Finally, when the outcome is binary and we want to model the odds ratio we might use a log-binomial model like

$$\operatorname{logit} \Pr[Y = 1 \mid A, V] = \beta_0 + \beta_1 A + \beta_2 V + \beta_3 AV$$

This time the **exponentiated** coefficient on the product term corresponds to

$$e^{\beta_3} = \frac{\frac{\text{Odds}[Y=1|A=1,V=1]}{\text{Odds}[Y=1|A=0,V=0]}}{\frac{\text{Odds}[Y=1|A=1,V=0]}{\text{Odds}[Y=1|A=0,V=0]}\frac{\text{Odds}[Y=1|A=0,V=1]}{\text{Odds}[Y=1|A=0,V=0]}} = \frac{OR_{11}}{OR_{10}OR_{01}}$$

which is now the measure of multiplicative interaction on the odds ratio scale.

4.5 Additional notes on statistical models

Some other thoughts:

- once you start using models note that it becomes harder to spot positivity violations as models will tend to extrapolate over sparse strata by assuming some sort of (linear) relationship, this can be useful but note that it's adding some additional assumptions (i.e. that the form is linear).
- you can get additive measures interaction from the log-binomial and logistic regression models by manipulating the output. see the examples at the end and our the included R code for reference on how to do this.

5 Additional topics

5.1 Qualitative interaction

In some settings, it's possible that exposure may have positive effects in one subgroup but negative effects on another subgroup (e.g. the hypothetical RCT data in the effect heterogeneity section above). This phenomena is often given a special name: *qualitative* interaction. Interestingly, unlike general statistical interactions, qualitative interactions are **not** scale dependent. If there is a qualitative interaction on the additive scale there will be a qualitative interaction on the multiplicative scale as well, and vice versa (this is just a pure accounting identity).

Qualitative interactions are often important because they imply that there exists subgroups for whom exposure is harmful (when the hypothesis is that effect should be protective for everyone like with a drug) or alternatively that there exists subgroups for whom exposure is beneficial (when the hypothesis is that the exposure is harmful for everyone). Thus, the existence of qualitative interaction may be important for determining whether everyone should be treated (or alternatively whether exposure should in fact be removed for everyone).

5.2 Mechanistic/sufficient cause interaction

We have largely omitted discussion of mechanism, i.e. how and why two factors combine to produce an interaction, because often the presence of a statistical interaction tells us little about them (and in some circumstances can be outright misleading) and our theory of the estimation of causal effect also does not require knowledge of them. However, there are certain, very proscribed, circumstances in which interaction might tell us something about mechanism.

To understand when this is the case we have to first introduce the notion of a *sufficient cause*. As defined by Rothman, a sufficent cause is the minimal set of events, conditions or characteristics that inevitably produce some outcome. Sufficient causes can be represented graphically via the famous causal pies. Each pie represents a different mechanism whereby an individual may develop the outcome. Furthermore, a *sufficient cause interaction* exists between two exposures A and V occur together in a sufficient cause and this would imply that, at least for some portion of the population, exposure to A and V is a route to developing the outcome.

It turns out that under some stringent conditions, a statistical interaction between A and V in the data can imply the presence of a sufficient cause interaction between A and V and therefore that A and V are part of a mechanism for at least some individuals in the population. In order to conclude this, first all the assumptions for causal interactions (i.e. exchangeability with respect to both exposures, consistency, positivity) still apply. In addition, either of the following conditions[4] must be met:

[4]Note that these only apply to interactions on additive scale, in Tyler's book he also develops similar conditions for sufficient cause interaction on the multiplicative scale but these are beyond the scope for our course.

Option 1 A and V both have positive *monotonic*[5] effects on the outcome and $p_{11} - p_{10} - p_{01} + p_{00} > 0$ (or likewise $RERI > 0$).

Option 2 alternatively, if $p_{11} - p_{10} - p_{01} > 0$ (or likewise $RERI > 1$) this would still imply a sufficient cause interaction even if effects are not *monotonic*.

5.3 Case-only estimator (genetics)

In some circumstances, particularly in the field of genetics, we may want to assess interaction between two exposures on a binary outcome using only data obtained among the cases (think a full genome sequencing study of people with a rare genetic disorder). It turns out that if the distributions of the the two exposures A and V are statistically independent in the population then the odds ratio reelating A and V among the cases is equal to the interaction measure $\frac{RR_{11}}{RR_{10} \cdot RR_{01}}$ on the multiplicative scale:

$$\frac{\Pr[A = 1 \mid V = 1, Y = 1] / \Pr[A = 0 \mid V = 1, Y = 1]}{\Pr[A = 1 \mid V = 0, Y = 1] / \Pr[A = 0 \mid V = 0, Y = 1]} = \frac{RR_{11}}{RR_{10} \cdot RR_{01}}$$

This presents huge advantages when the independence assumption is true as we only need to collect data on the cases; however it can be severely biased when this assumption does not hold. In the literature it is common for people to argue this is the case for interactions bewtween genetic and environmental exposures but this may not be true.

The assumption can be relaxed to a conditional independence after conditioning on some set of covariates C. In practice the case-only estimator is usually conducted using a logistic regession on the case-only data of the form:

$$\operatorname{logit} \Pr[A = 1 \mid V, C] = \alpha_0 + \alpha_1 V + \alpha_2^t C$$

Here the exponentiated α_1 would estimate the multiplicative measure of interaction $\frac{RR_{11}}{RR_{10} \cdot RR_{01}}$ on the risk ratio scale.

5.4 Statistical considerations

Throughout this lab we really have not paid much attention to questions of sampling variability and how to distinguish true interaction from sampling noise. This was purposeful to make sure we were set on the theory, but know that in practice sampling variability adds another layer of complication. In general, estimating interactions requires much larger sample sizes than estimating main effects of two exposures alone. This is because the power to detect even moderately sized interactions quickly gets vanishingly small when we consider the ideas that interactions are often smaller in magnitude than main effects and that generally a small proportion of individuals are in each of the four or more conditions that define the interaction effect. The homework has a nice simulation that will hopefully help illustrate this, but in general just know that if you're primary interest is in assessing interactions you may want to plan to collect A LOT of data.

[5] Monotonic hear means that A and V both always increase risk or alternatively in potential outcomes notation this implies that individuals of type $(Y_{a=1,v=1} = 0, Y_{a=0,e=1} = 1)$, $(Y_{a=1,v=1} = 0, Y_{a=1,e=0} = 1)$, $(Y_{a=1,v=0} = 0, Y_{a=0,e=0} = 1)$, and $(Y_{a=0,v=1} = 0, Y_{a=0,e=0} = 1)$ can't exist

6 Practice problems

6.1 Problem #1

You conducted a study to assess whether asbestos exposure (E) interacts with a genetic exposure, G, in their associations with respiratory dysfunction. Assume the outcome is rare. All variables are binary. You fit a logistic regression model with the following output:

```
##
## Call:
## glm(formula = Y ~ G * E, family = binomial(link = "logit"))
##
## Deviance Residuals:
##     Min       1Q   Median       3Q      Max
## -1.8930  -1.3799   0.6039   0.9876   1.2691
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept)  -0.2131     0.1503  -1.417    0.156
## G             0.5008     0.3091   1.620    0.105
## E             0.6774     0.1713   3.955 7.64e-05 ***
## G:E           0.6444     0.3930   1.640    0.101
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##     Null deviance: 1335.7  on 999  degrees of freedom
## Residual deviance: 1283.3  on 996  degrees of freedom
## AIC: 1291.3
##
## Number of Fisher Scoring iterations: 4
```

(a) Estimate RR_{11}, i.e., the risk ratio comparing the doubly-exposed to the doubly-unexposed. (You can leave the math unsimplified throughout.)

Answer:

$$RR_{11} = e^{0.501+0.677+0.644} = 6.18$$

(b) Estimate the RERI for this study.

Answer:

$$RERI = e^{0.501+0.677+0.644} - e^{0.501} - e^{0.677} + 1 = 3.57$$

(c) Estimate the proportion of risk in the doubly-exposed that is due to interaction.

Answer:

$$AP = \frac{RERI}{RR_{11}} = \frac{3.57}{6.18} = 0.58$$

(d) Is there evidence for sufficient-cause interaction between G and E, if so what assumptions are necessary for this to be causal?

Answer: There is evidence for sufficient cause interaction because $RERI > 1$ which meets

stronger criteria that doesn't require monotonicity assumption. However, we still need to assume marginal exchangeability of G and E as well as consistency and positivity.

6.2 Problem #2

You are reading a paper reporting on the interaction between a drug (E) and counseling (X) on recovery from depression (Y). The drug was randomly assigned, but counseling was not. The paper reports an RERI of 2.5 (95% CI 2.1 - 2.9). For each of the following statements made in the manuscript, **state whether you believe or don't believe it** on the basis of the reported RERI. If you don't believe it, describe what other evidence or assumptions would be needed to convince you.

(a) "More people would recover if the drug were given preferentially to individuals who received counseling."

SOLUTION: Yes. This is a statement about effect heterogeneity. Since the drug is randomized and because the RERI provides evidence for additive statistical interaction, this statement is reasonable.

(b) "If physicians prescribed both the drug and counseling to the same patients, more patients overall would recover than if physicians invested the same resources to prescribing the drug and counseling to different groups of patients."

SOLUTION: No. This is a statement about causal interaction. We need to assume that both E and X are unconfounded, but we only know that E is randomized.

(c) "There are individuals who would recover if they both took the drug and received counseling, but who would not recover if they just received counseling or just took the drug."

SOLUTION: No. This is a statement about mechanistic interaction. Even without monotonicity assumptions, this RERI is large enough to conclude sufficient cause interaction; however, we would again need to assume that both E and X are unconfounded.

6.3 Problem #3

An `R` function [available online][https://osf.io/748dp/] estimates additive interactions and uses them to test for mechanistic interactions [2]. You can just copy/paste the function definition into your `R` console or use the below code to automatically download and read in the code from an online public repository:

```
# downloads the code into your current working directory
if(!file.exists("interaction_code.R")){
  library(rvest)
  source_code <- read_html("https://mfr.osf.io/render?url=https://osf.io/748dp/?action=download%26mode=
  html_nodes("pre") %>% html_text()
  sink("interaction_code.R")
  cat(source_code)
  sink()
}
source("interaction_code.R")

# now you'll just need to call the function and interpret the results
# see the OSF repository for documentation and vignettes: https://osf.io/7ccpp/
```

Now we'll simulate data with two exposures (`E1` and `E2`) and confounder `C`:

```r
n=5000
set.seed(451)
E1 = rbinom(n, size=1, p=0.2)  # generate a binary genetic exposure
set.seed(1984)
E2 = rbinom(n, size=1, p=.75)

# binary confounder that is correlated with both exposures
# linear predictor from which to generate C
linpred = 1 + 1.5*E1 + 2.2*E2
# expit to back-transform logits into probabilities
probs = exp(linpred) / (1 + exp(linpred))
# generate C
set.seed(451)
C = rbinom(n, size=1, p=probs)

# true betas for Y
beta0 = log(1)
betaE1 = log(1.17)
betaE2 = log(1.15)
betaE1E2 = log(1.7)

# for generation of Y
linpred = beta0 + betaE1*E1 + betaE2*E2 + betaE1E2*E1*E2
probs = exp(linpred) / (1 + exp(linpred))
set.seed(22)
Y = rbinom(n, size=1, p=probs)
```

Fit the logistic regression model including firstly the main effects and interaction of the exposures of interest, followed by any covariates:

```r
m = glm(Y ~ E1*E2 + C, family=binomial(link="logit"))
summary(m)
```

```
##
## Call:
## glm(formula = Y ~ E1 * E2 + C, family = binomial(link = "logit"))
##
## Deviance Residuals:
##     Min       1Q   Median       3Q      Max
## -1.6143  -1.2324   0.7963   1.1234   1.1635
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept)  0.03282    0.16088   0.204 0.838365
## E1           0.23828    0.14837   1.606 0.108279
## E2           0.07668    0.07428   1.032 0.301918
## C            0.01890    0.16113   0.117 0.906634
## E1:E2        0.61931    0.17272   3.586 0.000336 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
```

```
##
##     Null deviance: 6863.1  on 4999  degrees of freedom
## Residual deviance: 6757.5  on 4995  degrees of freedom
## AIC: 6767.5
##
## Number of Fisher Scoring iterations: 4
```

Then pass it to `additive_interactions`:

```
additive_interactions(m)
```

```
##    Stat        Est       CI.lo     CI.hi      p.val.0 p.val.epi
## 1  RERI 1.19661040  0.65430493 1.7389159 7.636672e-06 0.9981552
## 2    AP 0.47011133  0.30355991 0.6366628 1.580897e-08        NA
## 3    E1 0.17411092 -0.05795168 0.4061735 7.071139e-02        NA
## 4    E2 0.05157272 -0.04212606 0.1452715 1.403422e-01        NA
## 5 E1:E2 0.77431636  0.50702237 1.0416103 6.823338e-09        NA
##   p.val.suff.cause
## 1        0.2386735
## 2               NA
## 3               NA
## 4               NA
## 5               NA
```

What happens if we assume monotonicity for at least one of the exposures?

```
additive_interactions(m, monotone=1)
```

```
##
## Assuming AT LEAST ONE of exposures has positive monotonic effect
##
##    Stat        Est       CI.lo     CI.hi      p.val.0 p.val.epi
## 1  RERI 1.19661040  0.65430493 1.7389159 7.636672e-06 0.2386735
## 2    AP 0.47011133  0.30355991 0.6366628 1.580897e-08        NA
## 3    E1 0.17411092 -0.05795168 0.4061735 7.071139e-02        NA
## 4    E2 0.05157272 -0.04212606 0.1452715 1.403422e-01        NA
## 5 E1:E2 0.77431636  0.50702237 1.0416103 6.823338e-09        NA
##   p.val.suff.cause
## 1        0.2386735
## 2               NA
## 3               NA
## 4               NA
## 5               NA
```

6.4 Problem #4: (Very optional)

This problem is very optional. You will NOT have to do anything like this on exams, homework, or in real life.

Prove that the case-only estimator on slide 18 is indeed equal to the multiplicative interaction when $A \perp\!\!\!\perp V$. Hint: Start with the right hand side (the multiplicative interaction). Use Bayes' Rule to flip the direction of conditioning for each of the 3 terms. Under independence, $P(A = a, V = v) = P(A = a)P(V = v)$; use this in each of the resulting denominators. Then various terms should cancel out.

7 References

1. VanderWeele, TJ. "On the distinction between interaction and effect modification." *Epidemiology* 20.6 (2009): 863-871.
2. Mathur, MB, & VanderWeele, TJ. (2018). R Function for Additive Interaction Measures. *Epidemiology*, 29(1), e5-e6.