Comparing interventions using inverse probability weighted risk functions

EPID 722

Spring 2021

jessedwards@unc.edu

Where are we?

Readings: Cole SR, Hernán MA. Adjusted survival curves with inverse probability weights. Comput. Methods Programs Biomed. 2004;75(1):45–9

Today: Lecture and discussion

Thursday: Exercise and discussion

Where have we been?

First 6 weeks of the course: Methods to estimate risk

- In settings with censoring and truncation using nonparametric models (week 2)
- Using parametric models (week 3)
- Generalized to external target populations (week 4)
- In settings with informative censoring (week 5)
- In settings with competing events (week 6)

Hazard ratio interlude

Midterm

Where are we going?

Now, we will discuss methods to compare risks.

Why compare risks?

1. To describe outcomes for subsets of the target population (usually to assess whether there is some type of disparity in risk)
2. To compare outcomes under interventions/policies/plans
3. Others?

"Descriptive" comparison of risk: We describe differences in risk even in cases where we may not wish to intervene to change someone's group membership

Comparison of outcomes to estimate a causal effect. Typically of an intervention we could hypothetically implement that would move people from one group to another.

What do I mean by “intervention?”

Anything you might think about changing in the world.

Synonyms: treatment, exposure plan, policy, strategy

Risks in the world as it could be

Epidemiologists are often called upon to estimate the distribution of outcomes that would occur under a specific exposure, treatment, or intervention.

If everyone in the target population receives that intervention, this is easy.

If some people in the target population receive that intervention, this is harder.

If no one in the target population receives that intervention, this is difficult.

Potential outcomes

Let T_i^a represent the time from origin to event that *would have occurred* for person i under intervention a

The risk under intervention a is $F_{T^a}(t)$ where

$$F_{T^a}(t) = P(T_i^a \leq t)$$

Formalization with potential outcomes

We can compare risks under multiple (possibly) counterfactual exposure plans using risk ratios and differences:

Risk difference:

$$F_{T^1}(t) - F_{T^0}(t) = P\left(T_i^1 \leq t\right) - P\left(T_i^0 \leq t\right)$$

Risk ratio:

$$F_{T^1}(t)/F_{T^0}(t) = P\left(T_i^1 \leq t\right)/P\left(T_i^0 \leq t\right)$$

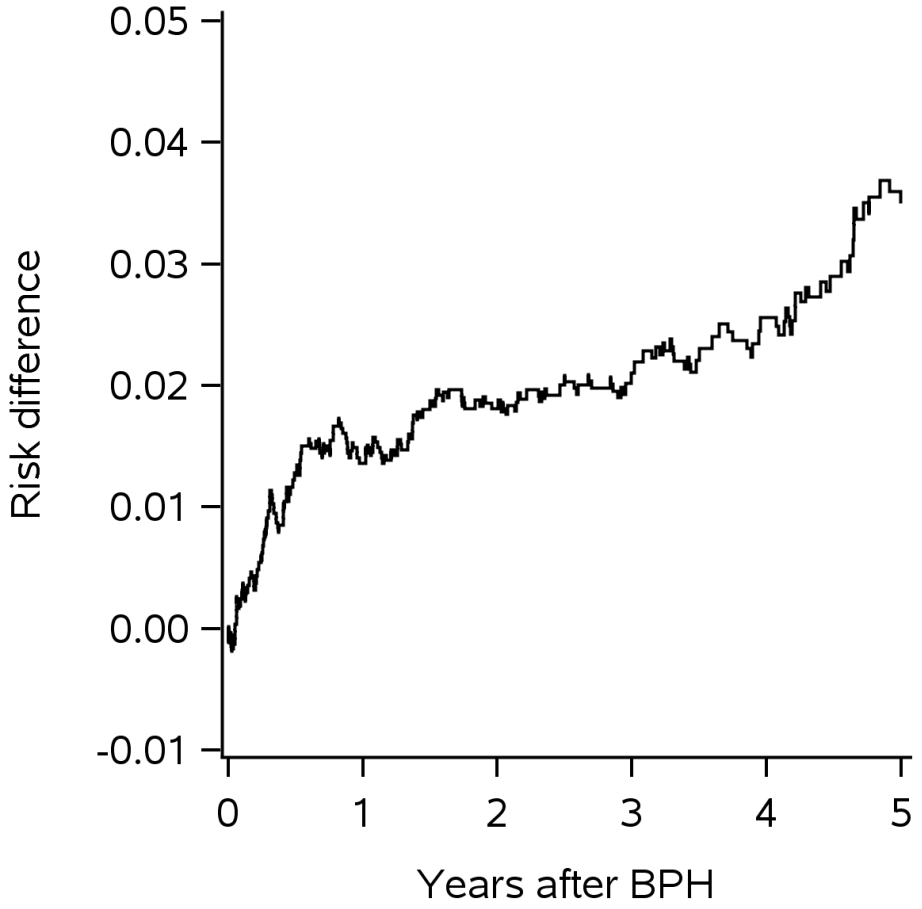

Why formalize our approach with potential outcomes?

To clarify the parameter of interest:

e.g., $F_{T^1}(t)$ rather than $F_T(t|A = 1)$

(What is the difference?)

Other formalizations are possible:

e.g., the “do” operator (Pearl, *Causality*)

Potential outcomes are (useful) fiction

Potential outcomes cannot be measured.

But we can use assumptions to make inferences about them using the observed data.

For details, see Edwards JK, Cole SR, Westreich D. All your data are always missing: incorporating bias due to measurement error into the potential outcomes framework. International journal of epidemiology. 2015 Aug 1;44(4):1452-9.

Linking potential outcomes to data

We let the observed outcome T_i stand in for at most 1 potential outcome for each person in our study.

$T_i^a = T_i$ when $A_i = a$
$T_i^a = .$ when $A_i \neq a$

"Counterfactual consistency"

Missingness of T^a is determined by the intervention assignment* in the observed data.

We usually "observe" T^a for at most one value of a.

To identify $F_{T^a}(t)$ in the presence of missing T^a, we let the experience of those with $A = a$ stand in for those with $A \neq a$.

* This could be a true assignment, as in a trial, or assignment by some other mechanism (e.g., participant choice)

True

T^1	T^0	1	A	T
10	5	1	0	
10	5	1	0	
10	5	1	0	
10	5	1	0	
4	2	1	1	
4	2	1	1	

$E(T^1) =$

$E(T^0) =$

$E(T|A=1) =$

$E(T|A=0) =$

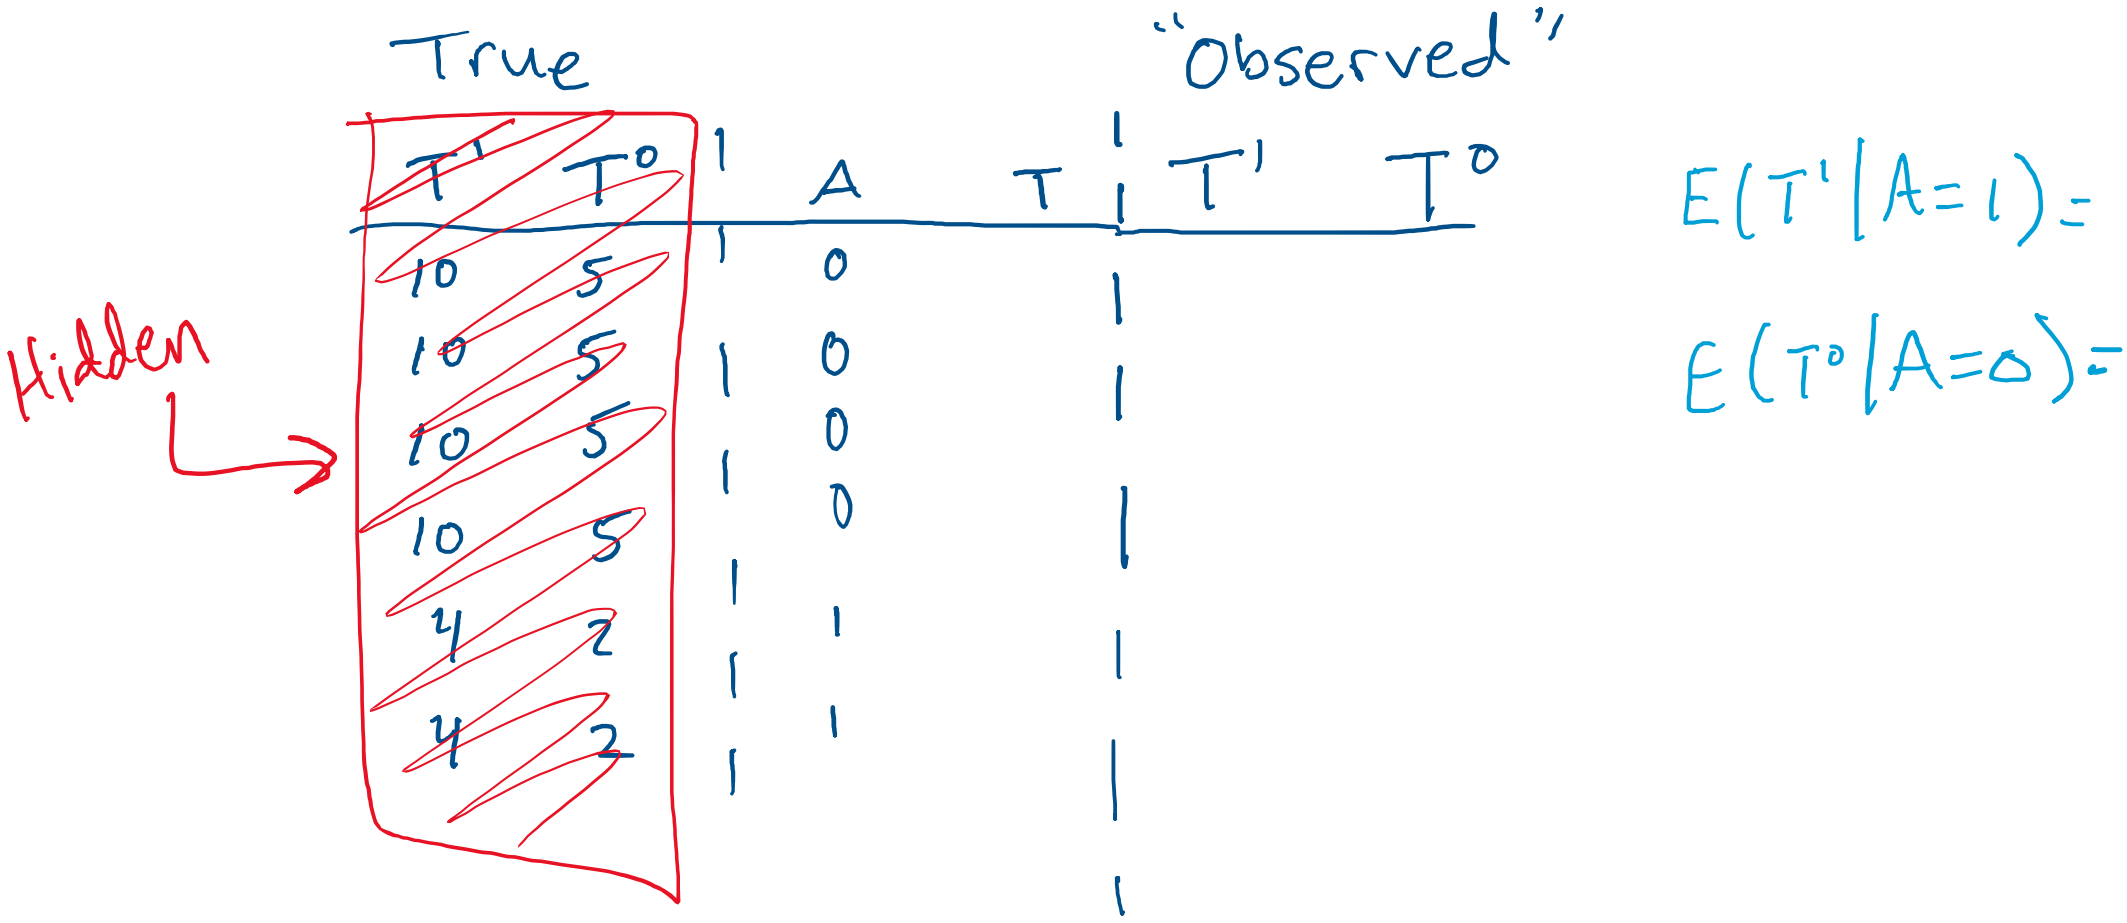
True
"Observed"
Hidden
T^1
T^0
1
A
T
T^1
T^0
10
5
0
4
2
$E(T^1|A=1)=$
$E(T^0|A=0)=$

In randomized trials, …

In expectation, randomization balances covariate history (including the unseen potential outcomes) between people assigned to each arm.

In other words,

When treatments/exposures are randomized, T^1 and T^0 are missing completely at random.

In this setting, observed exposure A is independent of T^a, $A \coprod T^a$.

We will refer to this concept as “exchangeability”.

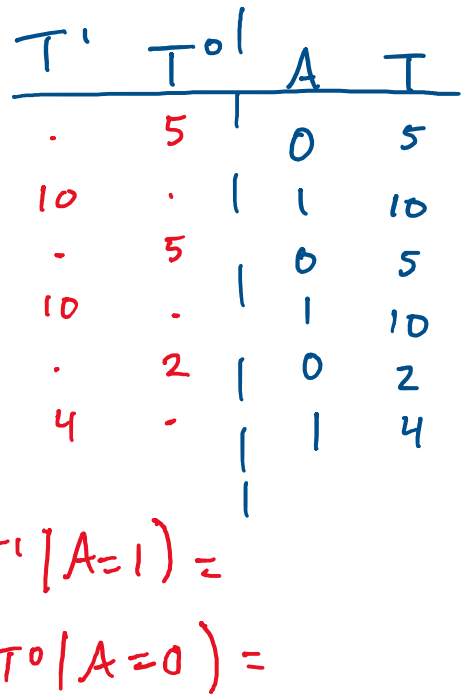

Risk in observational studies

Because A is not randomized, it is unlikely that $A \coprod T^a$, in most settings.

(Why?)

However, within strata of covariates L, we may believe that A is independent of T^a, or $T^a \coprod A|L$

We call this assumption "conditional exchangeability."

When we assume conditional exchangeability, we must also assume positivity, or $P(A = a|L = l) > 0$ for all l found in the target population.

In observational studies, under the conditional exchangeability assumption, we believe T^a is missing at random within strata of L.

- "Complete case" analysis will fail
- But we can reweight "complete cases" to have the same distribution of L as the full data.

IP-weighted risk functions

Under the assumptions of conditional exchangeability and positivity, we may estimate $F_{T^a}(t)$ as the inverse probability weighted risk function for group $A = a$ in the observed data (details above).

The inverse probability weights are:

$$\pi_i^a = \frac{1}{P(A_i = a|L_i)}$$

IP-weighted risk functions

Under the assumptions of conditional exchangeability and positivity, we may estimate $F_{T^a}(t)$ as the inverse probability weighted risk function for group $A = a$ in the observed data (details above).

The (stabilized) inverse probability weights are:

$$\pi_i = \frac{P(A_i = a)}{P(A_i = a|L_i)}$$

Why stabilize?

Improves the performance of the weights (reduces extreme values) and weights sum to original sample size.

Stabilized IP-weights

The stabilized IP weights should have a mean of 1.

To estimate the denominator of the weights $P(A_i = a|L_i)$, we could use

- Tabular methods (in low dimensional settings)
- A parametric regression model
- A machine learning approach

In this course, we will primarily use logistic regression.

Estimating stabilized IP-weights

If we use a parametric model to estimate denominator of weights, we must carefully consider how to model covariates L.

Ideally, we would fit a saturated model.

If fitting a saturated model is not possible (e.g., when some covariates are continuous), it is best to be as flexible as possible.

Assuming by default that continuous covariates have a linear (or log-linear) relationship with the outcome is dangerous and could result in uncontrolled confounding.

Estimating IP-weighted risk functions

$$F_{T^a}(t) = 1 - S_{T^a}(t),$$

$$= 1 - \prod_{k:R_k \leq t} \left\{ 1 - \frac{y_k^w(a)}{N_k^w(a)} \right\},$$

Where $y_k^w(a)$ is the weighted number of events at time k in group $A = a$, and $N_k^w(a)$ is the weighted number of subjects in the risk set at time k in group $A = a$.

Visualizing IPW and regression adjustment using DAGs

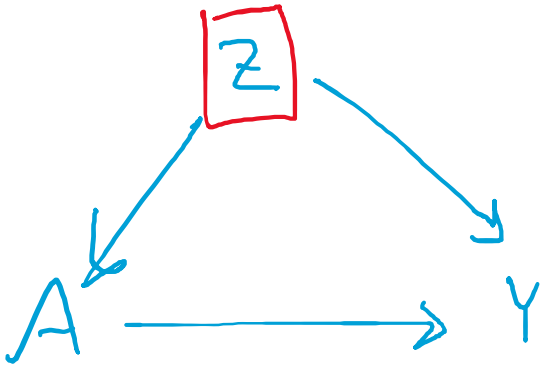

Stratification

(includes restriction + regression adjustment)

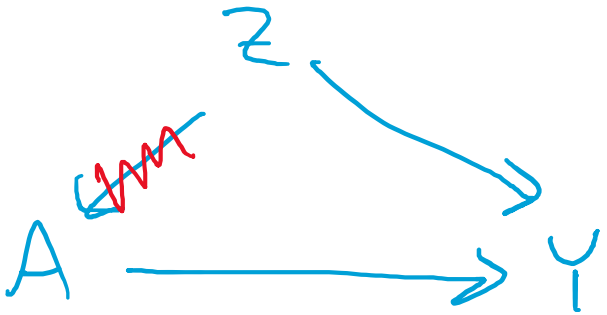

Standardization

(IP weighting or g computation)

Important caution

Including variables in inverse probability of treatment weights does NOT account for informative censoring based on those variables.

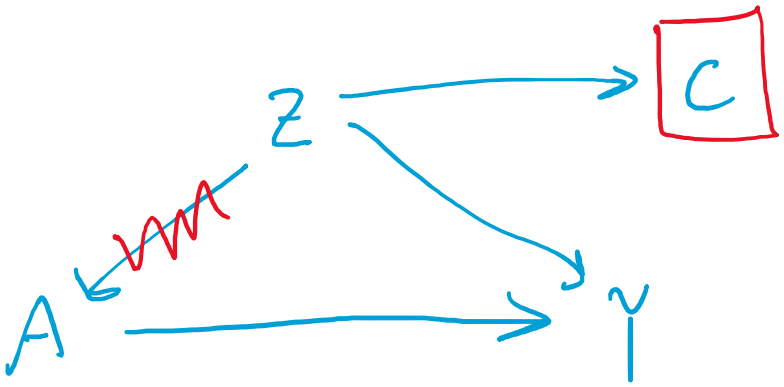

Combining IPW and IPCW

Treatment weights and censoring weights can be multiplied together:

$$\pi_i(t) = \pi_i^a \times \pi_i^c(t)$$

$$\pi_i(t) = \frac{P(A_i = a)}{P(A_i = a|L_i)} \times \frac{P(C_i > t|A_i)}{P(C_i > t|A_i, Z_i)}$$

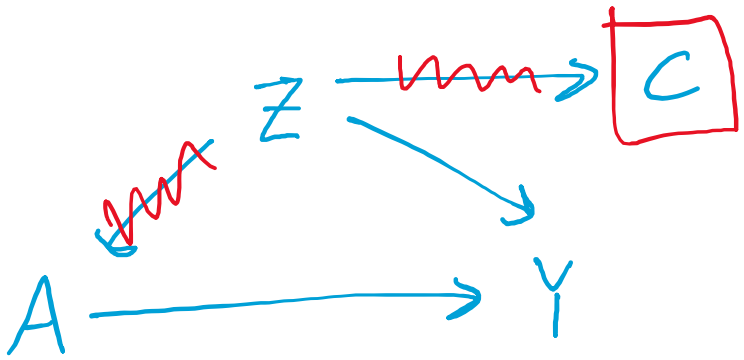

Applying weights

This week, we focus on applying weights to the Kaplan-Meier estimator to estimate counterfactual risks.

However, we could also apply weights to the Aalen-Johansen estimator to estimate counterfactual risks in settings with competing events.

Or, we could apply weights to regression models to estimate weighted RRs, RDs, ORs, or HRs.

Variance

When using weights to *standardize the covariate distribution in each treatment arm to the covariate distribution in the entire study sample*, the **robust variance** provides a conservative estimate of the variance for parameters of interest.

However, when alternative standardization techniques are used

e.g.:

- using weights to standardize the covariate distribution in each treatment arm to the covariate distribution in the treated, or
- using inverse odds weights for transportability

the robust variance is not guaranteed to be conservative, and bootstrapping is recommended.

Variance

In these settings, we will use the nonparametric bootstrap to obtain estimates of the standard error of the risk at each timepoint of interest.

We first saw the bootstrap in Exercise 2.

Recall, algorithm:

1. Draw B samples of size N from the observed data
2. For each sample, compute the desired risk (e.g., 10-year risk)
3. The standard deviation of estimates across the B samples is the standard error of the estimate

VARIATIONS

Alternative IP weights

The stabilized weights presented in this lecture are the weights we would use to estimate the effect of intervention a in the entire study sample.

These weights standardize the covariate distribution in each group (exposed and unexposed) to the covariate distribution in the entire study sample.

Following Sato [1], we can adapt these weights if we instead want to estimate the effect of a *among those exposed to* a.

[1] Sato and Matsuyama, Epidemiology 2003;14: 680–686

Effect of treatment in the treated

To estimate effect of a in those observed to be exposed to a, $\mathrm{E}(F_{T^1}(t) - F_{T^0}(t)|A = a)$:

1. Estimate the risk among people with $A = a$, had they been assigned to intervention a
    1. This is simply $P(T^a \leq t|A = a) = P(T \leq t|A = a)$, or the observed risk of the outcome in the exposed.
2. Estimate the risk among people with $A = a$, had they been assigned to intervention a'.
    1. Let risk among those with $A = a'$ stand in for risk under a'
    2. Reweight participants with $A = a'$ to have the same distribution of covariates Z as those with $A = a$

Effect of treatment in the treated

To estimate effect of a in those observed to be exposed to a, $\mathrm{E}(F_{T^1}(t) - F_{T^0}(t)|A = 1)$:

1. Estimate the risk among people with $A = 1$, had they been assigned to intervention 1
    1. Assign people observed to have $A = 1$ a weight of 1

2. Estimate the risk among people with $A = 1$, had they been assigned to intervention 0.
    1. Assign people observed to have $A = 0$ a weight of $\pi_i = P(A = 1|Z_i)/P(A = 0|Z_i)$

How would we adapt this algorithm to estimate the effect of treatment in the untreated?

Alternative IP weights (ii)

The stabilized weights presented in this lecture are the weights we would use to estimate the effect of assigning everyone to exposure a compared to assigning *no one* to exposure a.

We can adapt our application of these weights if we instead want to compare the risk we would see under exposure a to the risk we would expect to see under the status quo (also known as the "natural course").

What is the risk under the status quo?

The risk we observe under the observed exposure distribution.

(We can estimate this using the Kaplan-Meier estimator in the entire study sample).

Steps to compare risk under intervention to risk under the status quo

1. In one copy of the dataset (the reference copy, or the copy representing the status quo), give everyone a weight of $\pi_i = 1$. Assign everyone a new variable $X_i = 0$.
2. In a second copy of the dataset, limit to those who followed the plan of interest (e.g., treated individuals). Then give these individuals a weight of $\pi_i = 1/P(A = 1|Z_i)$. Assign everyone in this dataset a new variable $X = 1$.
3. Concatenate these datasets.
4. Continue with standard analytic techniques (weighted Cox model or weighted KM estimators) using X as the exposure variable and π as the weight.

EXAMPLE

Context

Say we have a study of 1000 patients receiving routine physical exams at UNC on October 1, 2021.

Your goal is to follow them all for 1 year to assess flu incidence. We will use various techniques to estimate flu incidence under various strategies that involve providing flu vaccines on the date of the routine physical exam.

All patients have agreed to contact the study team if they get the flu during the risk period.

Assume no patients had received a flu shot prior to their exam date and that patient outcomes are independent.

Discuss

Say we have a study of 1000 patients receiving routine physical exams at UNC on October 1, 2021.

All patients have agreed to contact the study team if they get the flu during the risk period.
Assume no patients had received a flu shot prior to their exam date.

Your goal is to estimate the 1-year incidence of flu among this group of individuals if they were to all receive the flu shot on that date.

How would you conduct this study if you:

1. All patients wanted a flu shot and you could give out 1000 flu shots?
2. All patients wanted a flu shot, but you could only give out 500 flu shots?
3. Only 500 patients wanted a flu shot (and you had >500 flu shots)?

Context

Say we are in the setting where only 500 patients wanted a flu shot.

Desire for the flu shot is affected only by binary covariate Z, which was associated with a 3-fold increase in desire for the vaccine. Z also affected risk of flu.

We wish to estimate the risk of flu in the **natural course** and counterfactual risk under several scenarios.

Natural course

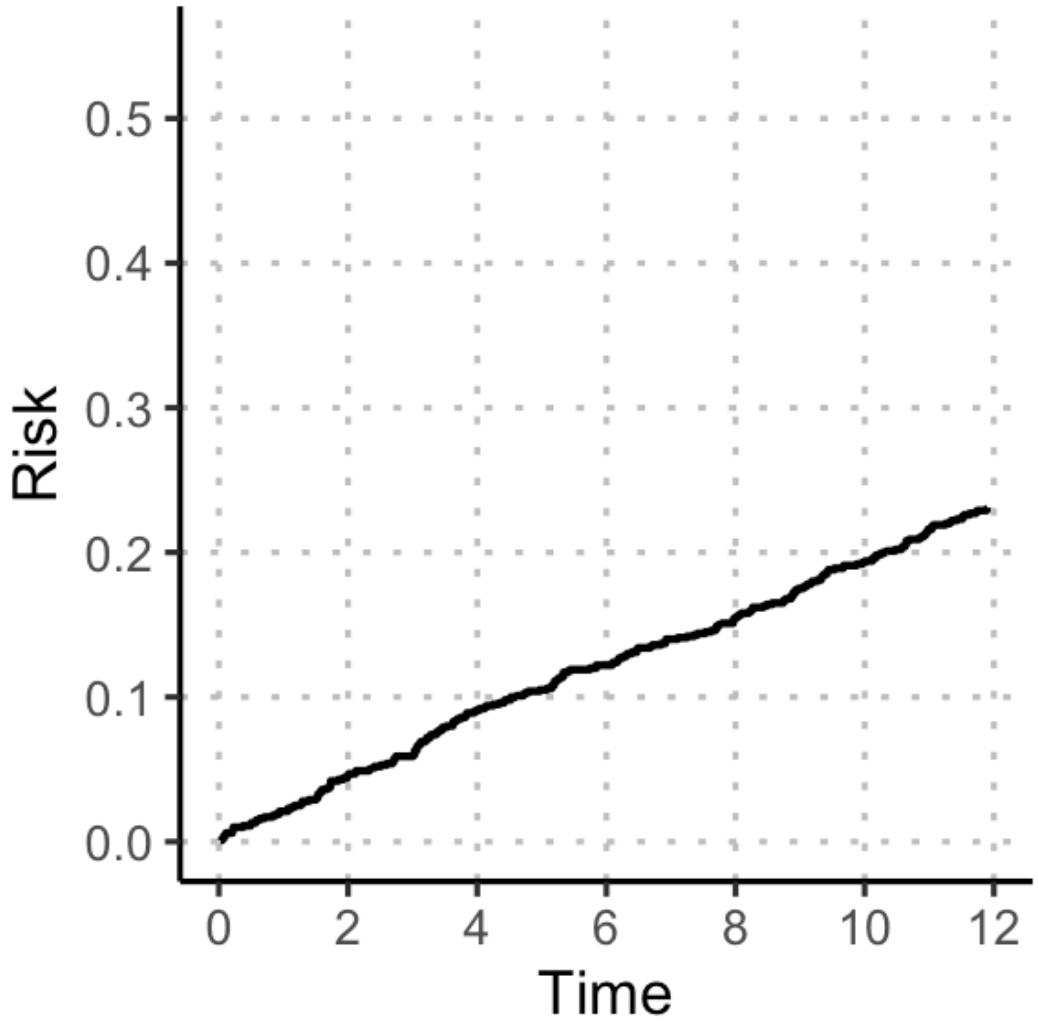

Crude

How does risk differ between those **vaccinated** and those **not vaccinated**?

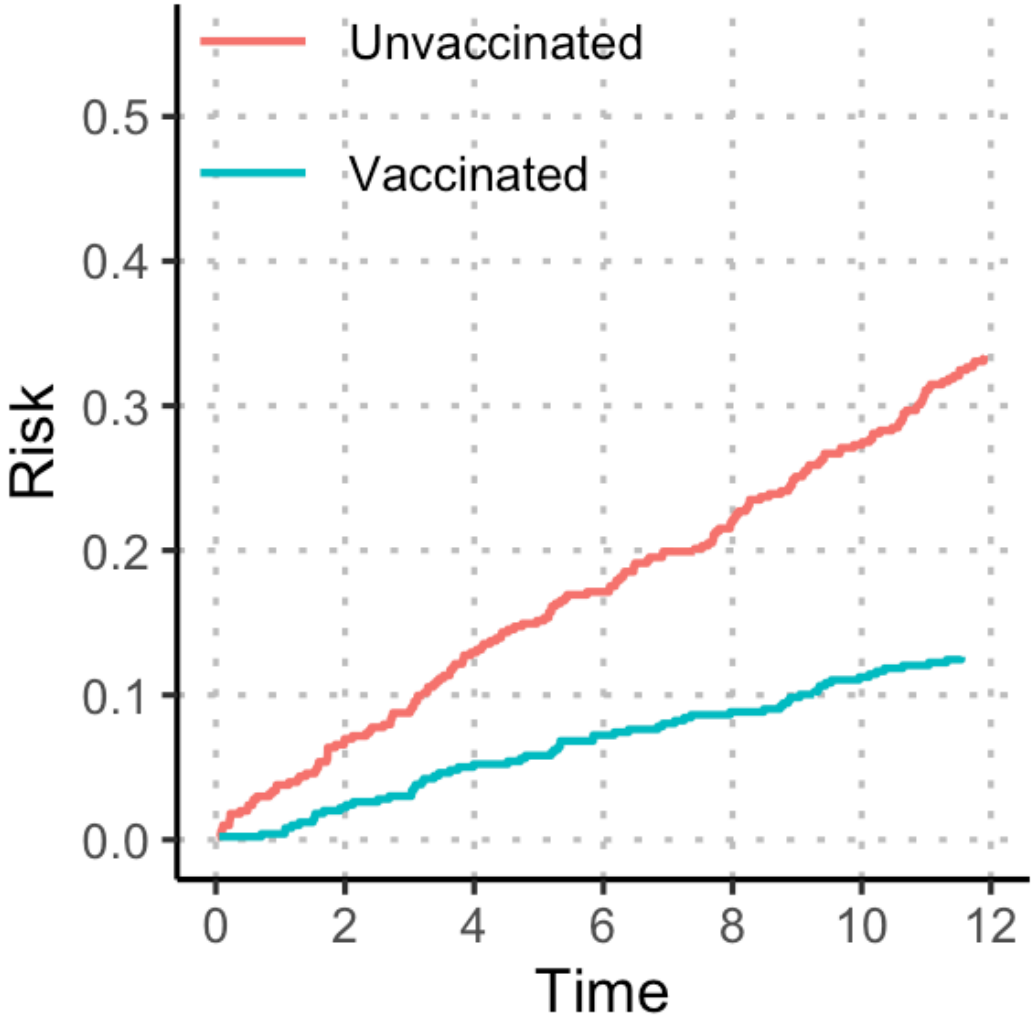

Crude

Steps:

1) Stratify KM estimator by vaccine status
2) Fit crude cox model
3) Plot

HR = 0.33

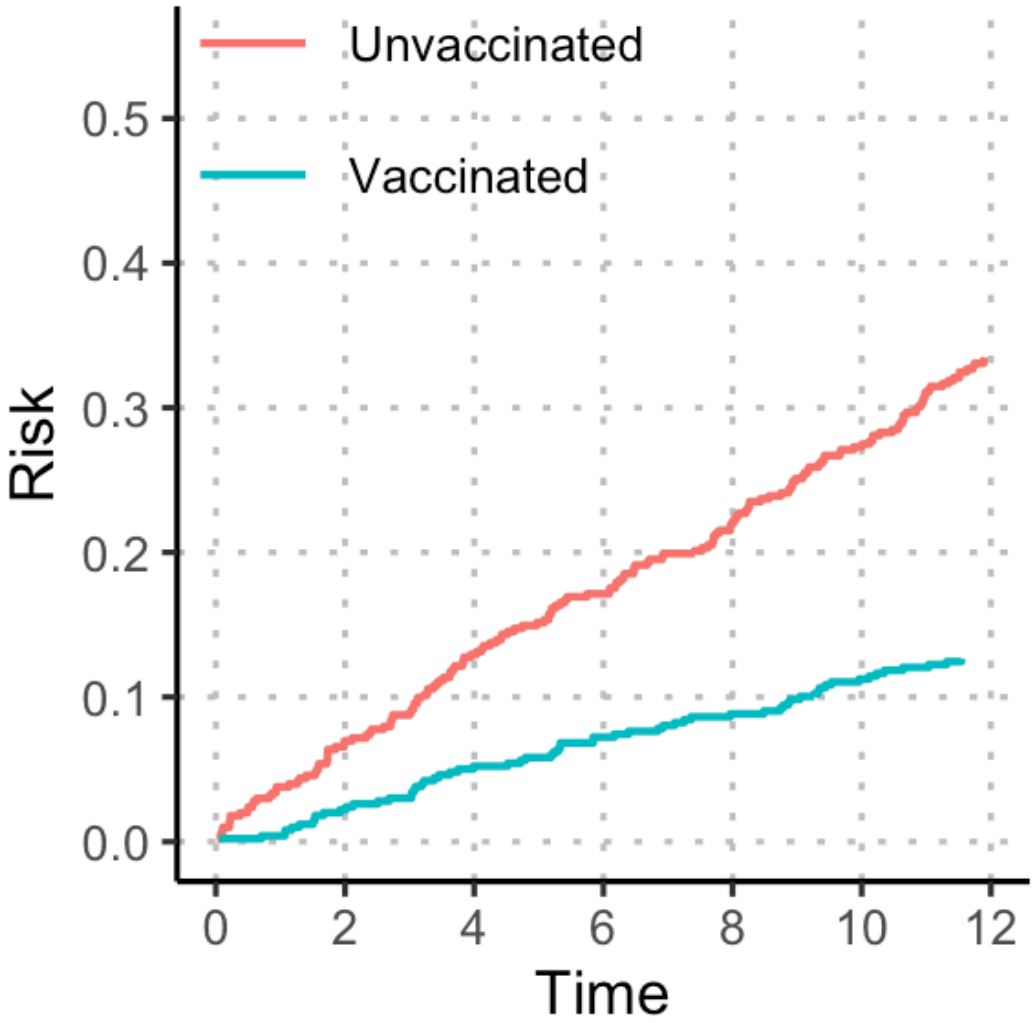

Average treatment effect

How would the risk in the entire study sample **had everyone been vaccinated** compare to the risk **had no one been vaccinated**?

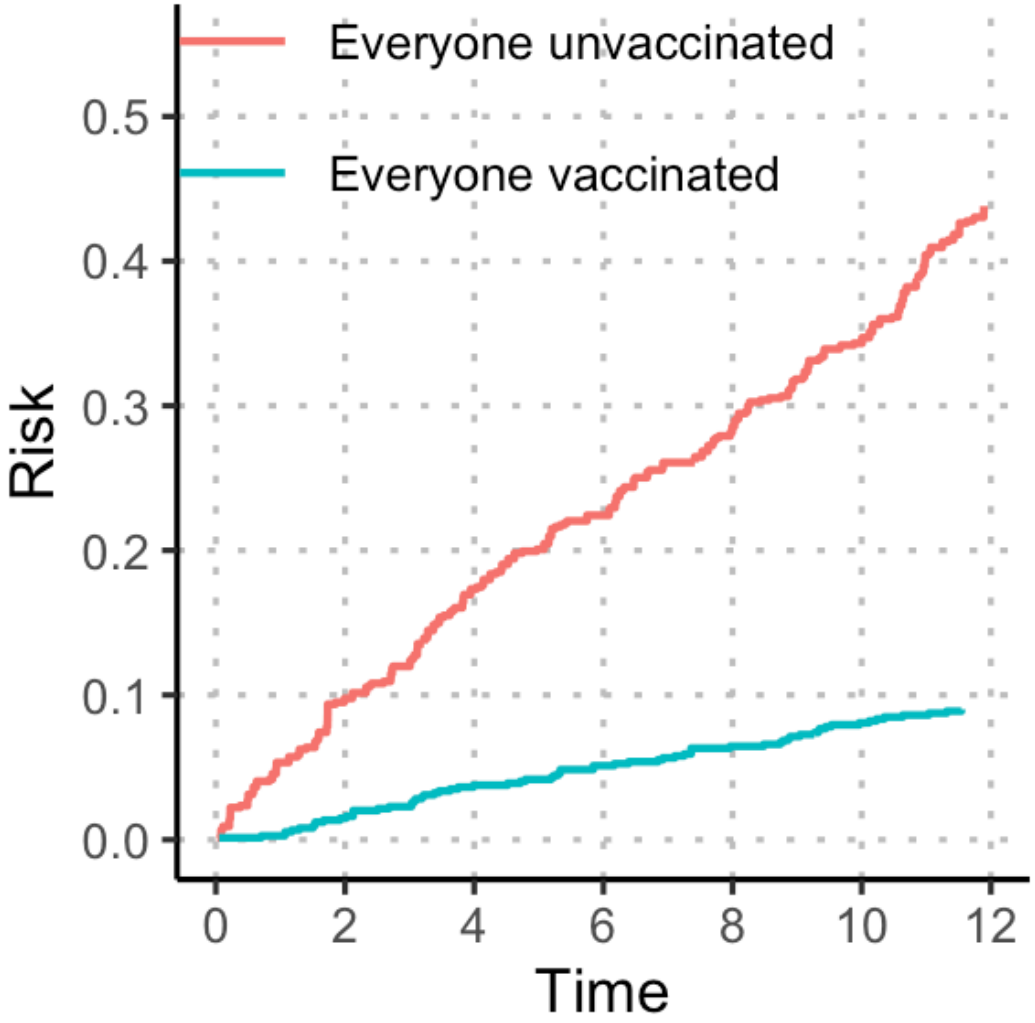

Average treatment effect

Steps:

1) Estimate IPW as

$$\frac{P(A = a)}{P(A = a|Z)}$$

2) Fit weighted KM estimator stratified by A and Cox model

3) Plot

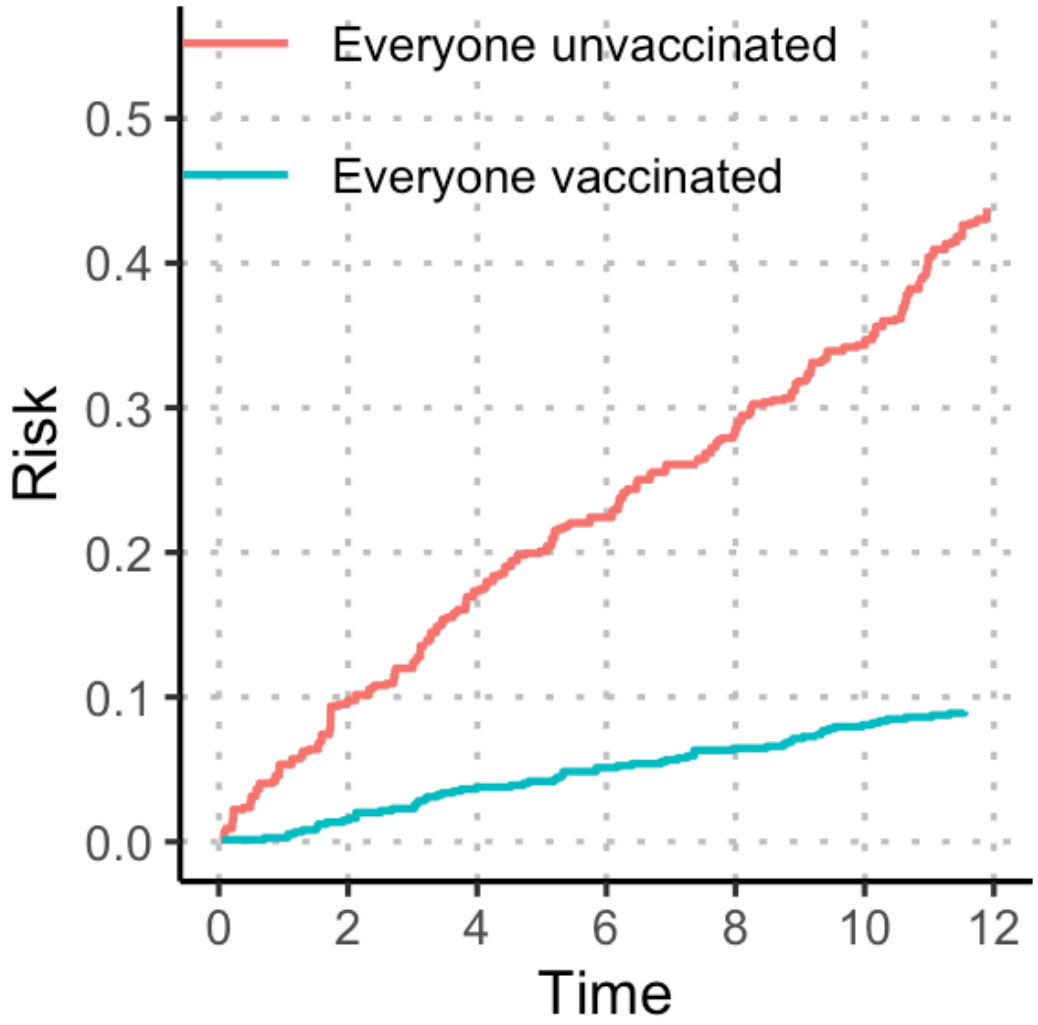

HR = 0.16

Alternate target population: Effect of treatment in the treated

How does **risk among vaccinated patients** compare to **the risk we would have seen in that group had they not been vaccinated**?

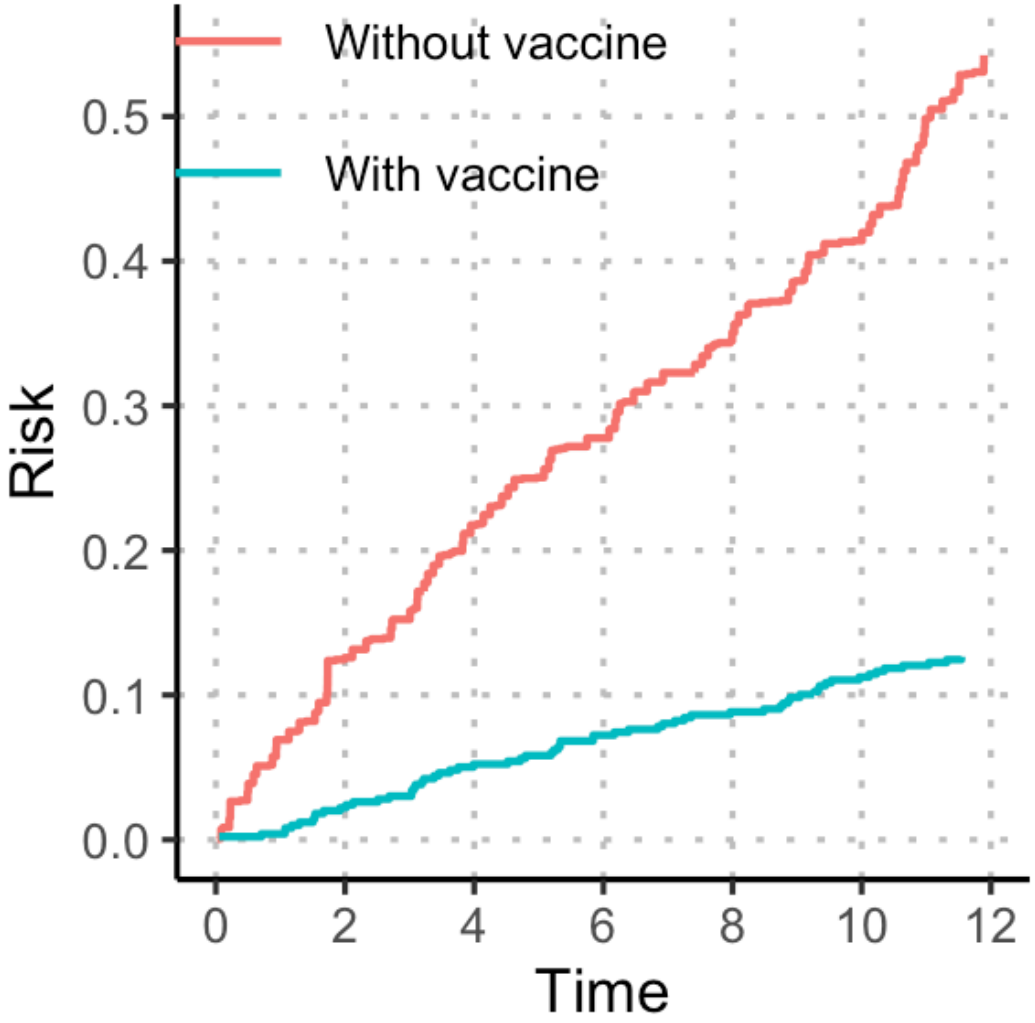

Alternate target population: Effect of treatment in the treated

Steps:

1) If $A = 1$, weights are set to 1
2) If $A = 0$, weights are

$$\frac{P(A = 1|Z)}{P(A = 0|Z)}$$

2) Fit weighted KM estimator stratified by A and Cox model

3) Plot

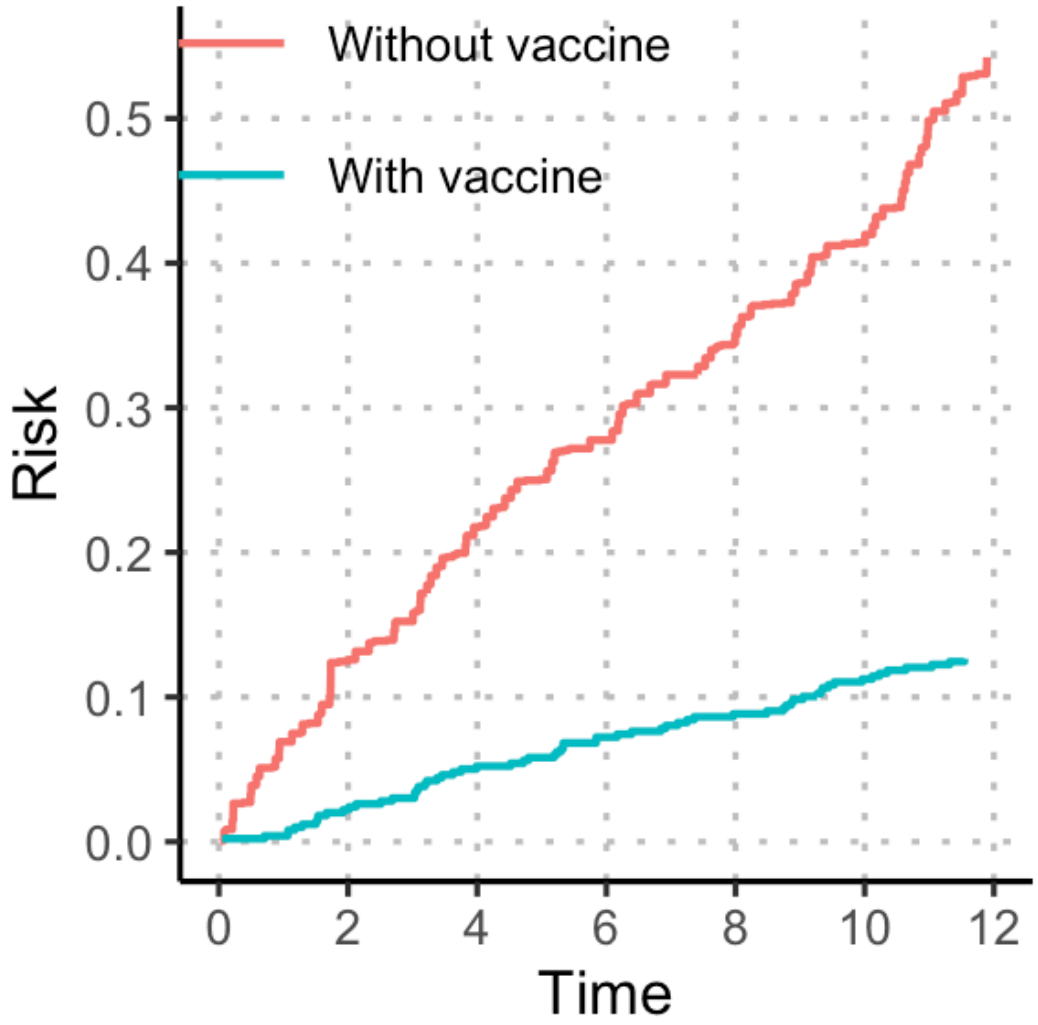

HR = 0.17

Alternate target population: Effect of treatment in the UNtreated

How does **risk among unvaccinated patients** compare to **the risk we would have seen in that group had they been vaccinated**?

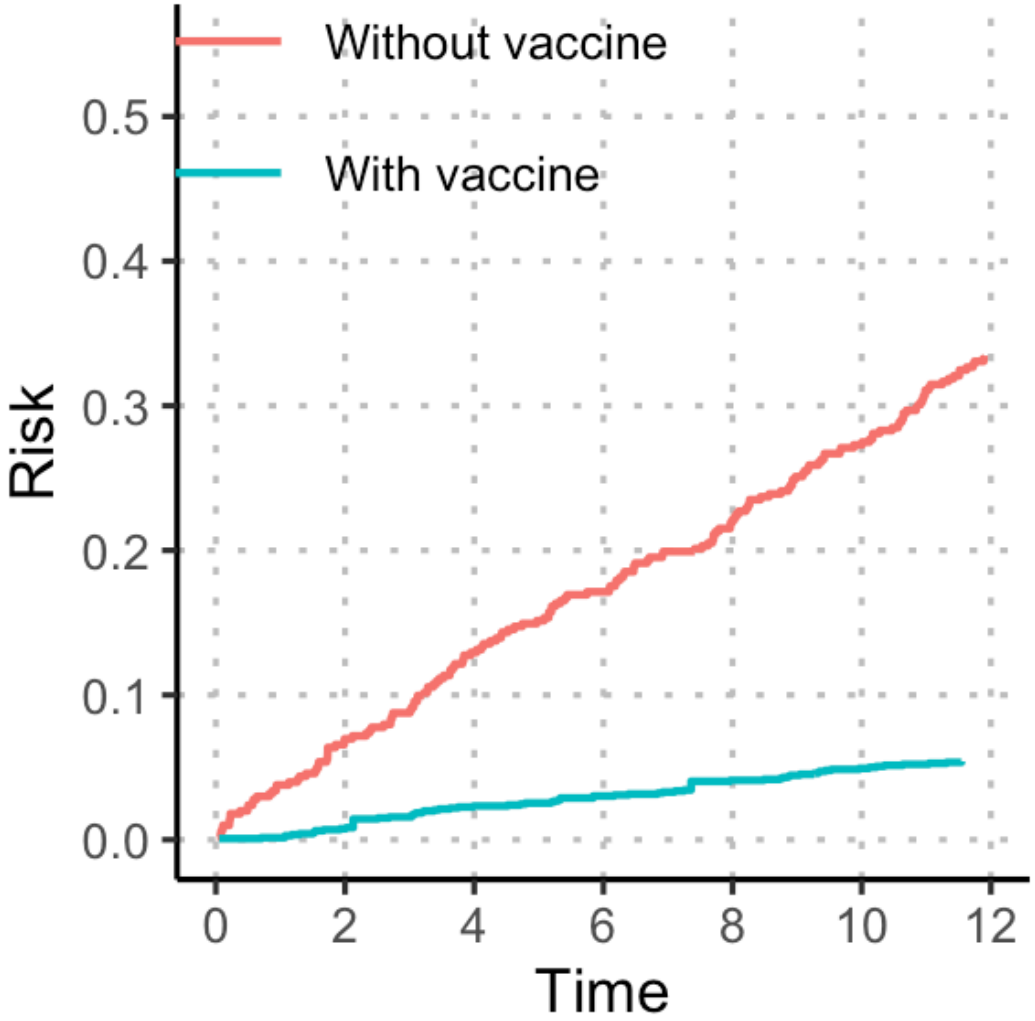

Alternate target population: Effect of treatment in the UNtreated

Steps:

1) If $A = 0$, weights are set to 1
2) If $A = 1$, weights are

$$\frac{P(A = 0|Z)}{P(A = 1|Z)}$$

2) Fit weighted KM estimator stratified by A and Cox model

3) Plot

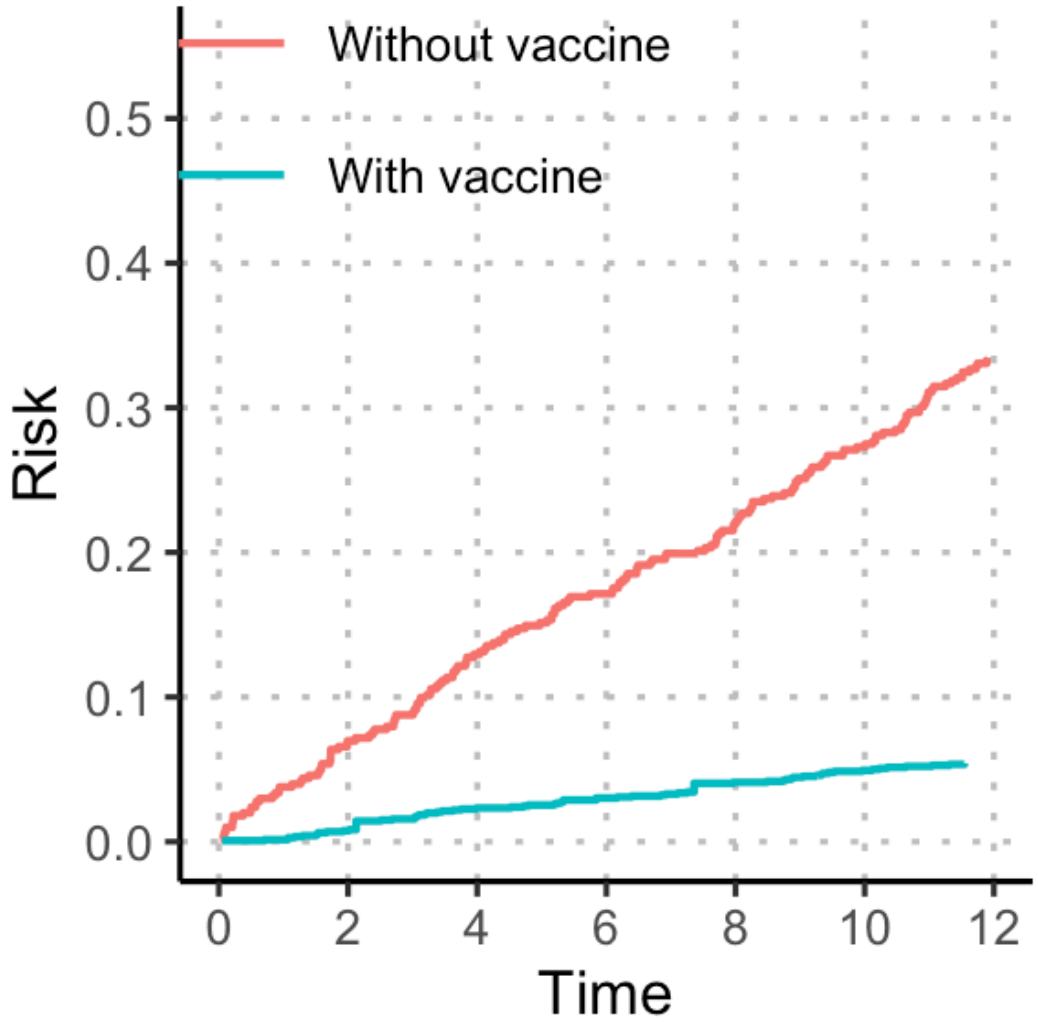

HR = 0.14

Alternate contrast: Comparing to natural course

By how much would we have reduced risk if we had **vaccinated everyone** (compared to the risk we saw under the **observed distribution of vaccine**)?

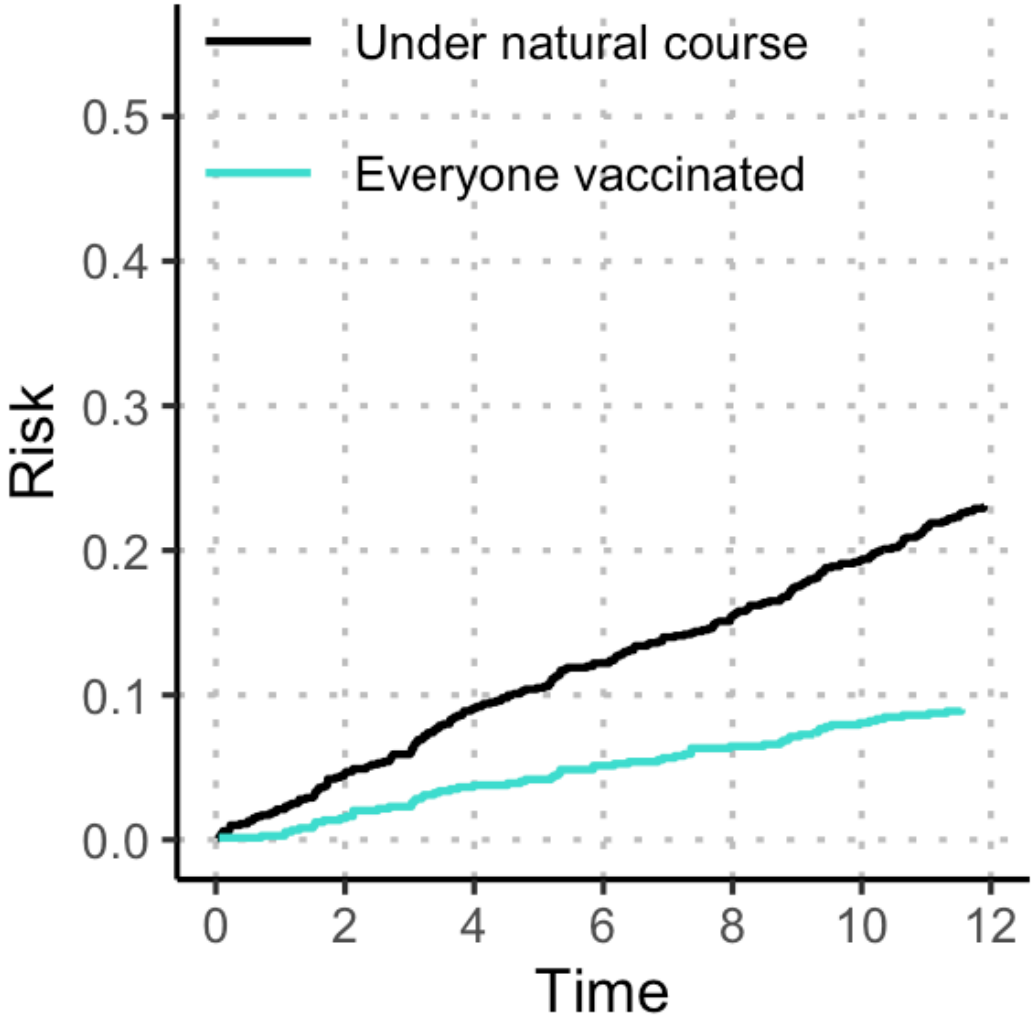

Alternate contrast: Comparing to natural course

Steps:

1) In the observed data, set $X = 0$ and weight = 1
2) In a second copy of the observed data, compute weights as
   1) If $A = 0$, weight = 0
   2) If $A = 1$, weight = $\frac{P(A = 1)}{P(A = 1|Z)}$
3) Concatenate data from steps 1 and 2
4) Fit weighted KM estimator stratified by X and Cox model (with X as exposure)
5) Plot

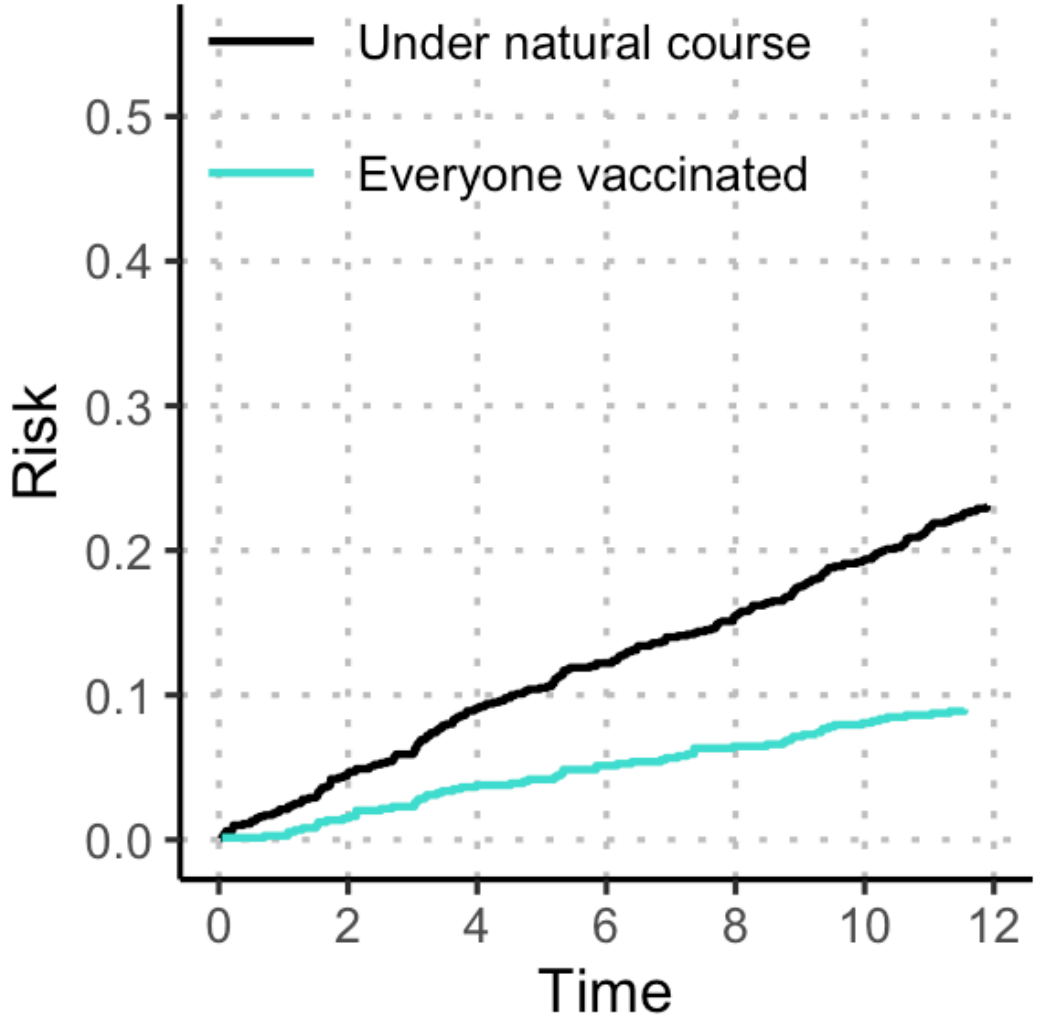

HR = 0.36

Comparing interventions using inverse probability weighted risk functions

EPID 722

Spring 2021

jessedwards@unc.edu