



JMP016: Salk Vaccine

Relative Risk

Produced by:

Dr. DeWayne Derryberry,
Idaho State University Department of Mathematics

Salk Vaccine

Relative Risk

Key ideas

Relative risk, observational study, cohort study, randomized experiments.

Background

In the 1950s polio was responsible for 6% of the deaths in 5- to 9-year-olds and left many more crippled. In 1954 a study was undertaken to determine the effectiveness of the Salk polio vaccine. The study was unusual in that almost 2 million first-, second- and third-graders participated in the study.

The study can be thought of in two parts:

- **Randomized Experiment:** In some areas of the country parents consented to have their second-grade children participate in a double-blind randomized trial. These children were randomly assigned to receive the vaccine or a salt solution (placebo) by injection.
- **Cohort Study:** In other areas of the country parents were less receptive to the notion of random assignment, and second-graders received the vaccine, while first- and third-graders were given nothing and treated as a control group.

When comparable groups are followed forward in time, the study is often called a cohort study. Both the randomized experiment and the cohort study began before the 1954 polio season and ended after the 1954 polio season.

See *The Biggest Public Health Experiment Ever: The 1954 Field Trial of the Salk Poliomyelitis Vaccine*, Paul Meier, <http://www.stat.luc.edu/StatisticsfortheSciences/MeierPolio.htm>.

The Task

Determine whether the vaccine was effective in the cohort study. And, if the vaccine was effective, quantify the degree of the effectiveness.

We will leave a similar analysis of the randomized experiment as an exercise.

The Data [Salk.jmp](#)

The data set contains the results of the experiment.

Polio Yes (got polio) or no (did not get polio)
Treatment Vaccine (in the vaccine group) or control (in the control group)
Count The number of children in each classification. Count has been assigned the "Freq" role in the JMP data table.

Exhibit 1 The Data

Polio	Treatment	Count
yes	vaccine	57
no	vaccine	200,688
yes	placebo	142
no	placebo	201,087

Design Issues

We prefer a randomized experiment over an observational study. Why? In an experiment we have created two groups (by randomization) that we expect to be very similar in all sorts of ways. Given that we can never anticipate all possible confounding variables, randomization, especially when big groups are involved, tends to balance out both known and unknown confounding variables across the two groups (treatment and control).

This cohort study, which involved children from all the same schools, with second-graders in the treatment group (vaccinated) and first- and third-graders in the control group, tends to balance out possible confounding variables in much the same way as random assignment. Both groups, for example, have the same average age (probably the reason for including first- and third-graders), but more importantly, the treatments and controls should have about the same racial and gender composition, about the same socioeconomic profile, etc.

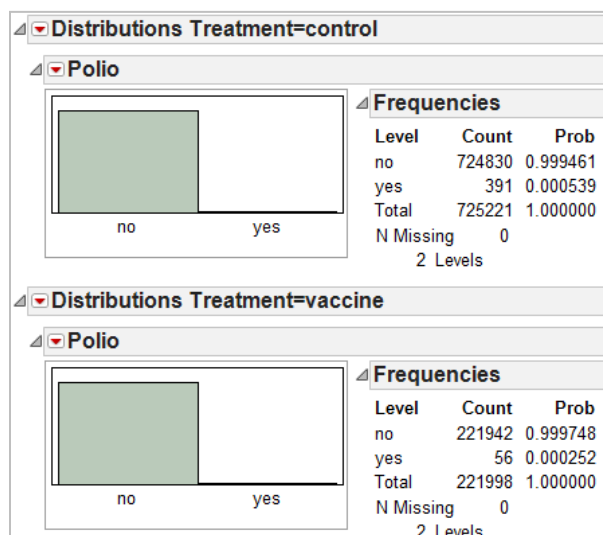
Given all of this, why is a cohort study not as good as an experiment? There are two reasons: lack of blinding and the placebo effect. Because a cohort study is not double-blind, certain biases threaten the validity of the study. Because the subjects know whether or not they got the vaccine, this could influence their behavior. In fact, those with the vaccine may take more chances and those without the vaccine may be more cautious. This would have the effect of making the vaccine appear less effective than it actually is.

In a cohort study we cannot measure the placebo effect, so we cannot adjust for it. If we only have a cohort study, even if the vaccine is effective, we cannot know to what extent this is a placebo effect. Nevertheless, cohort studies may be all that is available when studying humans. Certainly among observational studies, well-planned cohort studies are some of the most reliable.

Analysis

Was the vaccine effective in the cohort study? Exhibit 2 seems to indicate the vaccine is effective. Those in the control group had a polio rate at about 0.0539%, while those in the vaccine group had a polio rate of 0.0252%.

Exhibit 2 Preliminary Results for the Cohort Study



(Analyze > Distribution; select **Polio** as Y, Columns and **Treatment** as By. **Count**, which was previously assigned the "Freq" role, appears automatically in the Freq field. Click OK.

For a horizontal layout select Stack under the top red triangle.

To assign the "Freq" role before analysis, right-click on the variable in the data table and select Preselect Role > Freq.

To increase the number of decimal places displayed in the Frequencies table, double-click on a value in the **Prob** column and increase values in the Width and Dec fields.)

In medical studies where rare events are discussed, it is usual to express these numbers as cases per 100,000. Using this terminology, we would say there were 53.9 polio cases per 100,000 in the control group and 25.2 cases per 100,000 in the vaccine group.

Certainly the vaccine looks effective, but the sample sizes are small for count data. We need to be aware that, to a large extent, the study is only as good as the sample size of the cell with the lowest counts. While some cells have counts of over 200,000, others have counts of less than 500. For count data, cell counts of less than 500, and especially less than 100, are not necessarily large. For example, political opinion polls, working with count data, usually require a sample size of about 1,200. So, this study may not have a lot of power to detect a vaccine effect.

Let's formally test whether the vaccine was effective. We are interested in the following hypotheses:

H_0 : polio rates are the same for the vaccine group and the control group.

H_a : polio rates are lower in the vaccine group than the control group.

If the vaccine does appear to be effective, then we will want an estimate of the effect size, with the appropriate confidence interval. Since we're interested in a one-sided test, we'll conduct a two-sample test for proportions.

Exhibit 3 provides convincing evidence (p -value $\ll 0.01$) that the vaccine is effective. However, the reduction in the polio rate is estimated at only between 0.0198% and 0.0369%.

Exhibit 3 A Two-Sample Test of Proportions for the Cohort Study

Two Sample Test for Proportions			
Description	Proportion Difference	Lower 95%	Upper 95%
P(yes control)-P(yes vaccine)	0.000287	0.000198	0.000369
Adjusted Wald Test		Prob	
P(yes control)-P(yes vaccine) >= 0	<.0001*		
P(yes control)-P(yes vaccine) <= 0	1.0000		
P(yes control)-P(yes vaccine) = 0	<.0001*		
Response Polio category of interest			
<input type="radio"/> no			
<input checked="" type="radio"/> yes			

(Analyze > Fit Y by X, select **Polio** as Y, Response and **Treatment** as X, Factor. **Count**, which was previously assigned the "Freq" role, appears automatically in the Freq field. Click OK

Click on the top red triangle and select Two Sample Test for Proportions. This option is only available for 2 x 2 tables.

The default category of interest, in this case, is **no**. Change the category of interest to **yes** (got polio) using the radio button at the bottom of the output.)

With such a small difference, this does not seem to be very informative. One issue is that polio is just rare – almost no one got polio in either the treatment or the control group.

For rare events, the relative risk is a better measure for comparing probabilities. Relative risk is the ratio of two probabilities. In this case, it's a ratio of the probability of getting polio in the control group to the probability of getting polio with the vaccine.

The relative risk for the cohort study is 2.137 (Exhibit 4).

Exhibit 4 Relative Risk of Polio in the Cohort Study

Relative Risk			
Description	Relative Risk	Lower 95%	Upper 95%
P(yes control)/P(yes vaccine)	2.13731	1.615345	2.827936

(Select Relative Risk from the top red triangle. Change the response polio category from **no** to **yes**, and click OK.)

Since relative risk is a ratio, a value greater than 1.0 indicates that the vaccine was effective, and a value of less than 1.0 indicates that the vaccine was ineffective.

Our best estimate is that the vaccine cuts the chances of getting polio by 113.7% ($2.137 - 1$). That is, around 114% more second-graders would have gotten polio if no vaccine was given, compared to a population given the vaccine.

The confidence interval provides an interval estimate for the relative risk (see Exhibit 4). We want to know whether the entire interval is greater than 1.0, indicating the vaccine was effective, or whether 1.0 is contained in the interval, suggesting that it is unclear if the vaccine was effective. The 95% CI for relative risk is (1.615, 2.828). So, the estimated reduction in the polio rate, due to the vaccine, was between 61.5% ($1.615 - 1$) and 182.8% ($2.828 - 1$).

For two-by-two tables the χ^2 test (Exhibit 5) and the two-sample test of proportions (Exhibit 3) produce very similar results.

Exhibit 5 A Chi Square Analysis of the Cohort Study

Contingency Table				Tests			
		Polio		N	DF	-LogLike	RSquare (U)
		no	yes	947219	1	17.013996	0.0044
Treatment	Count						
	Total %						
	Col %						
	Row %						
	control	724830	391				
		76.52	0.04				
		76.56	87.47				
		99.95	0.05				
	vaccine	221942	56				
		23.43	0.01				
		23.44	12.53				
		99.97	0.03				
		946772	447				
		99.95	0.05				
				947219			

Test	ChiSquare	Prob>ChiSq
Likelihood Ratio	34.028	<.0001*
Pearson	29.659	<.0001*

Fisher's Exact Test	Prob	Alternative Hypothesis
Left	<.0001*	Prob(Polio=yes) is greater for Treatment=control than vaccine
Right	1.0000	Prob(Polio=yes) is greater for Treatment=vaccine than control
2-Tail	<.0001*	Prob(Polio=yes) is different across Treatment

(Displayed by default in Fit Y by X. Hint: If desired, right-click on the contingency table to add or remove summary data.)

There are several p -values associated with the χ^2 analysis in Exhibit 5. Fisher's exact test produces both one- and two-sided p -values, while the Pearson test (the test found in most introductory textbooks) and the likelihood ratio test produce only two-sided results. Unless sample sizes are very small, all three tests will produce similar results (we omit the details). We will look more closely at some of the computations associated with the tests in the exercises.

Alternatively, we might be interested in knowing the strength of the association between our two variables (whether students get polio and whether they were vaccinated). Several measures of association are available. However, we will only consider one that is particularly useful and easy to compute for two-by-two tables, Gamma (also called Yule's-Q in the case of a two-by-two table). Gamma (γ), which takes on values from -1 to 1, is often used to measure the association between two ordinal variables.

With a two-by-two table, Gamma is calculated as: $\gamma = \frac{ad-bc}{ad+bc}$

Where the table is of the form:

a	b
c	d

$$\gamma = \frac{(724830 \times 56) - (391 \times 221942)}{(724830 \times 56) + (391 \times 221942)} \sim -0.3626$$

This ratio will take on its highest positive value (1.0) when $b = 0$ or $c = 0$. In our table (displayed in Exhibit 5) this would occur in the unlikely event that none of the control group got polio, but *all* of the vaccine group got polio! On the other hand, this ratio takes on its lowest value (-1.0) when $a = 0$ or $d = 0$. This would occur if everyone in the control group got polio, but no one in the vaccine group did. In our problem,

This is the first entry in Exhibit 6.

Exhibit 6 Gamma, a Measure of the Degree of Association

Measures of Association				
Measure	Value	Std Error	Lower 95%	Upper 95%
Gamma	-0.3626	0.0621	-0.4843	-0.2410
Kendall's Tau-b	-0.0056	0.0008	-0.0072	-0.0040
Stuart's Tau-c	-0.0002	0.0000	-0.0003	-0.0001
Somer's D C R	-0.0003	0.0000	-0.0004	-0.0002
Somer's D R C	-0.1091	0.0157	-0.1398	-0.0784
Lambda Asymmetric C R	0.0000	0.0000	0.0000	0.0000
Lambda Asymmetric R C	0.0000	0.0000	0.0000	0.0000
Lambda Symmetric	0.0000	0.0000	0.0000	0.0000
Uncertainty Coef C R	0.0044	0.0014	0.0017	0.0071
Uncertainty Coef R C	0.0000	0.0000	0.0000	0.0001
Uncertainty Coef Symmetric	0.0001	0.0000	0.0000	0.0001

(From the Fit Y by X output, select Measures of Association from the top red triangle.)

Given that Gamma can take on values from -1 to 1, this does not indicate a particularly strong relationship between whether a child got the vaccine and also got polio. This is not surprising. A large value for gamma can only occur if one group has a high polio rate and the other has a very low polio rate. But, polio was rare in both groups. This reflects a subtle fact: if a disease is extremely rare, even a vaccine that is highly effective (reducing the chance of getting the disease by half!) can have a seemingly small impact.

Summary

Statistical Insights

The evidence strongly suggests the vaccine is effective in fighting polio. Our best estimate is that there would be a 61.5% to 182.8 % increase in polio cases if we went from a situation where everyone received the vaccine to a situation where no one received the vaccine.

This is definitely not a random sample. It appears to include about all of the second-graders in the US in 1954. Whether the vaccine would be more or less effective in other geographic regions, with different overall health and different ethnic makeup, is unclear. Medical and epidemiological experts would have some sense of the generalizability of the results. It is also unclear how the results might apply to older or younger children, and to what extent the virus will develop an adaptation over time, reducing the effectiveness of the vaccine.

Implications

Because this was an observational study, there is always some concern that the results may not hold up in future applications of the vaccine, or the vaccine may not be as effective as first thought. However, as observational studies go, this one was very good.

JMP® Features and Hints

This case first used the Distribution platform to explore polio rates for the control and vaccine groups. Count was assigned the Freq role in the data table, which allowed it to automatically populate the Freq field in the analysis dialog windows.

Then, we tested the hypothesis that the polio rates were the same for the control and vaccine groups using a one-sided, two-sample proportions test. The relative risk was used to compare polio rates and provide an estimate of the reduction in polio cases with the vaccine. The two-sided Pearson and likelihood ratio tests and one-sided Fisher's exact test, which display by default, were also used to compare polio rates.

Finally, a measure of association, Gamma, was introduced. Although Gamma is primarily used for ordinal variables, it's ideal for two-by-two tables.

Exercises

Below are the results for the randomized experiment:

Polio	Treatment	Count
yes	vaccine	57
no	vaccine	200,688
yes	placebo	142
no	placebo	201,087

1. Enter the data in a JMP data table – remember to preset the role of Counts to “Freq”. Perform a similar statistical analysis for the experimental data. How do the p -values for the tests, confidence interval for relative risk, and gamma compare to the cohort study.
2. The likelihood ratio test and Person’s test both produce values that are χ^2 with one degree of freedom, if the null hypothesis is true.

Let C_i be the observed counts for the four cells in the two-by-two table and let E_i be the expected counts in the two-by-two table, then

- a. The likelihood ratio is computed as: $2 \times \sum_{k=1}^4 C_k \times \ln \left(\frac{C_k}{E_k} \right)$

verify the value of the likelihood ratio in Exhibit 5.

- b. Pearson’s test statistic is the familiar:

$$\sum_{k=1}^4 \frac{(C_k - E_k)^2}{E_k}$$

verify the value of Pearson’s test statistic given in Exhibit 5.

3. For math geeks. The confidence interval for relative risk is as follows:

Let \hat{p}_v and n_v be the sample proportion and sample size for the vaccine group, and \hat{p}_c and n_c be the sample proportion and sample size for the control group. The confidence interval for the relative risk is:

$$\left(\frac{\hat{p}_c}{\hat{p}_v} e^{-K}, \frac{\hat{p}_c}{\hat{p}_v} e^K \right)$$

where $K = 1.96 \times \sqrt{\frac{1-\hat{p}_v}{n_v \hat{p}_v} + \frac{1-\hat{p}_c}{n_c \hat{p}_c}}$

Verify the relative risk confidence interval found in Exhibit 4. (This formula, found in Alan Agresti’s book on Categorical Data Analysis, is derived using the δ –method).