

### Chi-square Distribution

If  $Z$  has a standard normal distribution, then  $Z^2$  has a chi-square distribution, denoted by,  $\chi^2$ , with 1 df.

Thus

$$P(|Z| > 1.96) = 0.05$$

implies that

$$P(\chi^2 > 3.84) = 0.05$$

because  $(1.96)^2 = 3.84$ .

- A Chi-square random variable is nothing more than a sum of squared standard normal RVs.
- If  $Z_1, \dots, Z_n \sim \text{i.i.d. } N(0,1)$  then

$$Z_1^2 + \dots + Z_n^2 = \chi^2 \text{ is chi-square with } n \text{ df.}$$

The general form of the Chi-square statistic is:

$$\chi^2 = \sum_i \frac{(O_i - E_i)^2}{E_i}$$

where  $O_i$ 's are observed counts,  $E_i$ 's are expected counts under the Null hypothesis, and  $i$  denotes the number of 'cells'.

## Contingency Tables

---

The most common use of the Chi-square distribution is with contingency tables. Take the familiar 2x2 table as an example.

Observed table or  $O_i$ 's for  $i=1,\dots,4$

	Sample 1	Sample 2	
Success	a	b	a+b
Failure	c	d	c+d
	a+c	b+d	N

where  $\hat{\theta}_1 = \frac{a}{a+c}$  and  $\hat{\theta}_2 = \frac{b}{b+d}$

(it doesn't matter if we use row or column probabilities)

Under the Null hypothesis that  $H_0: \theta_1 = \theta_2$ , which implies that the rows and columns are independent, we have

Expected table or  $E_i$ 's for  $i=1,\dots,4$

	Sample 1	Sample 2	
Success	$(a+b)(a+c)/N$	$(a+b)(b+d)/N$	a+b
Failure	$(c+d)(a+c)/N$	$(c+d)(b+d)/N$	c+d
	a+c	b+d	N

(Why?)

## Contingency Tables

---

Our test statistic is the standard observed - expected chi-square statistic

$$\begin{aligned}\chi^2 &= \sum_i \frac{(O_i - E_i)^2}{E_i} \\ &= \frac{N(ad - bc)^2}{(a+b)(a+c)(b+d)(c+d)} \\ &= \left[ \frac{\hat{\theta}_1 - \hat{\theta}_2}{\sqrt{\theta_p(1-\theta_p)\left(\frac{1}{a+c} + \frac{1}{b+d}\right)}} \right]^2\end{aligned}$$

$$\hat{\theta}_1 = \frac{a}{a+c} \quad \text{and} \quad \hat{\theta}_2 = \frac{b}{b+d} \quad \theta_p = \frac{a+b}{a+b+c+d}$$

which is simply the Z-statistic for the difference between the two proportions squared. Because the Z-statistic is only approximately normal in large samples, the  $\chi^2$  will only be approximately correct in large samples as well.

Notice that we are testing that conditional probabilities of success given the column ( or sample) are equal.

## Contingency Tables

---

The  $\chi^2$  test statistic is almost always used with a 2x2 table, but the hypotheses being tested changes depending on how we obtained the data.

### Case 1: Prospective or Cross Sectional study

	Exposed	Not Exposed	
Disease	a	b	a+b
No Disease	c	d	c+d
	a+c	b+d	N

where  $\hat{\theta}_1 = \frac{a}{a+c}$  and  $\hat{\theta}_2 = \frac{b}{b+d}$  (a+c fixed or random)

$\theta_1 = P(\text{disease} | \text{exposed})$

$\theta_2 = P(\text{disease} | \text{not exposed})$

$H_0: \theta_1 = \theta_2$  or  $H_0: rr = \theta_1 / \theta_2 = 1$

or

$H_0: or = (\theta_1 / (1 - \theta_1)) / (\theta_2 / (1 - \theta_2)) = 1$

The null hypothesis is often interpreted as “No association between Exposure and Disease” or “The rows and columns are independent”, i.e.  $P(D|E) = P(D|\text{not } E)$  implies that disease status is independent of event E.

## Contingency Tables

---

### Case 2: Case-control Study

	Exposed	Not Exposed	
Disease	a	b	a+b
No Disease	c	d	c+d
	a+c	b+d	N

where  $\hat{\theta}_1 = \frac{a}{a+b}$  and  $\hat{\theta}_2 = \frac{c}{c+d}$  (notice denominators)

$\theta_1 = P(\text{exposed} \mid \text{disease})$

$\theta_2 = P(\text{exposed} \mid \text{no disease})$

- Because of our sampling scheme a+b is fixed, we are not able to estimate the risk difference or relative risk. However, we can estimate the OR. See example in "221 Study Design Impact on RD RR OR.xlsx".

So we can still test

$$H_0: OR = (\theta_1/(1-\theta_1)) / (\theta_2/(1-\theta_2)) = 1$$

See the excel sheet Study Design Impact on RD RR OR.xlsx for an example and detailed explanation of why the OR still gives the correct estimate in a case-control study.

## Contingency Tables

---

For contingency tables with  $r$  rows and  $c$  columns, use the same chi-square test statistic:

$$\chi^2 = \sum_i \frac{(O_i - E_i)^2}{E_i}, \text{ with } df = (r-1)(c-1).$$

This is often referred to as an 'test of association'. That is we are testing to see if the rows and columns are associated or if they are 'independent'.

Statistically we are testing to see if **all** of row probabilities, conditional on the rows, are equal. For example:

	1	2	3	4	5
A	$\theta_{a1}$	$\theta_{a2}$	$\theta_{a3}$	$\theta_{a4}$	$\theta_{a5}$
B	$\theta_{b1}$	$\theta_{b2}$	$\theta_{b3}$	$\theta_{b4}$	$\theta_{b5}$
C	$\theta_{c1}$	$\theta_{c2}$	$\theta_{c3}$	$\theta_{c4}$	$\theta_{c5}$
D	$\theta_{d1}$	$\theta_{d2}$	$\theta_{d3}$	$\theta_{d4}$	$\theta_{d5}$
E	$\theta_{e1}$	$\theta_{e2}$	$\theta_{e3}$	$\theta_{e4}$	$\theta_{e5}$

The Chi-square test of no association has the null

$$\begin{aligned} H_0: & \theta_{a1} = \theta_{a2} = \theta_{a3} = \theta_{a4} = \theta_{a5} && \text{and} \\ & \theta_{b1} = \theta_{b2} = \theta_{b3} = \theta_{b4} = \theta_{b5} && \text{and} \\ & \theta_{c1} = \theta_{c2} = \theta_{c3} = \theta_{c4} = \theta_{c5} && \text{and} \\ & \theta_{d1} = \theta_{d2} = \theta_{d3} = \theta_{d4} = \theta_{d5} && \text{and} \\ & \theta_{e1} = \theta_{e2} = \theta_{e3} = \theta_{e4} = \theta_{e5} && . \end{aligned}$$

It's equivalent to condition on the columns instead.

## Contingency Tables

---

Sometimes we wish to combine several contingency tables. When this happens we have to be careful. One way of determining if the data are similar enough to combine is to examine their odds ratio.

	Sample 1	Sample 2	
Success	a	b	a+b
Failure	c	d	c+d
	a+c	b+d	N

where  $\hat{\theta}_1 = \frac{a}{a+c}$  and  $\hat{\theta}_2 = \frac{b}{b+d}$

In this table the estimated *OR* is

$$(\theta_1/(1-\theta_1)) / (\theta_2/(1-\theta_2)) = ad/bc$$

and the confidence interval for the *OR* is:

$$\exp\left(\log(\hat{or}) \pm Z_{\alpha/2} \sqrt{\frac{1}{a} + \frac{1}{b} + \frac{1}{c} + \frac{1}{d}}\right)$$
$$\exp\left(\log(\hat{or}) \pm Z_{\alpha/2} \sqrt{\frac{1}{\hat{\theta}_1(a+c)} + \frac{1}{(1-\hat{\theta}_1)(a+c)} + \frac{1}{\hat{\theta}_2(b+d)} + \frac{1}{(1-\hat{\theta}_2)(b+d)}}\right)$$

⇒ We'll learn the trick behind the CI when we study the delta method for deriving a standard error.

## Contingency Tables

---

There is a test called the Mantel-Haenszel test, which tests if the odds ratio is equal in different  $r \times c$  contingency tables. This test is often called the 'test of Homogeneity'.

We will discuss it in detail later, but now we will examine a case when two  $2 \times 2$  tables are combined.

### **Example -- Simpson's Paradox**

This example is from an investigation of the relationship between smoking and cancer (data are fictitious). Gender is associated with both of these variables.

Here we will see that the odds of having cancer is greater for smokers, for both females and males.

However when we combine the data we see that smoking is actually protective for cancer. (really?)

## Contingency Tables

---

<b>Males</b> cancer	Smoking		
	Yes	No	
Yes	<b>9</b>	<b>51</b>	60
No	<b>6</b>	<b>43</b>	49
	15	94	109

$$\text{or(males)} = (9 \times 43) / (6 \times 51) = 1.26$$

$$P(C|S) = 9/15 = 0.60$$

$$P(C|nS) = 51/94 = 0.543$$

$$RR = 1.11. \text{ Risk Diff} = 0.057.$$

<b>Females</b> cancer	Smoking		
	Yes	No	
Yes	<b>14</b>	<b>7</b>	21
No	<b>19</b>	<b>12</b>	31
	33	19	52

$$\text{or(females)} = (14 \times 12) / (19 \times 7) = 1.26$$

$$P(C|S) = 14/33 = 0.42$$

$$P(C|nS) = 7/19 = 0.368$$

$$RR = 1.14. \text{ Risk Diff} = 0.052.$$

Both the male and the female have the same trend: the odds of developing cancer are higher among smokers than non-smokers.

Since the estimated odds ratios are similar why not simply combine the two tables?

## Contingency Tables

---

<b>combined</b>	Smoking		
cancer	Yes	No	
Yes	<b>23</b>	<b>58</b>	81
No	<b>25</b>	<b>55</b>	80
	48	113	161

$$\text{or}(\text{combined}) = (23 \times 55) / (58 \times 25) = 0.872$$

$$P(C|S) = 23/48 = 0.479$$

$$P(C|nS) = 58/113 = 0.513$$

$$RR = 0.934. \text{ Risk Diff} = -0.034.$$

And we see that smoking is now protective! Why did does the trend reverse?

- What does this imply about simply testing for homogeneity and then combining over tables?

⇒ Be very, very careful.

Sometimes, it is better to present stratified data when the effect appears the same, than it is to present the combined data.

## Contingency Tables

---

Another example:

*Ask Marilyn (Parade Magazine, April 28, 1996)*

A reader writes in to ask:

"A company decided to expand, so it opened a factory generating 455 jobs. For the 70 white collar positions, 200 males and 200 females applied. Of the females who applied, 20% were hired, while only 15% of the males were hired. Of the 400 males applying for the blue collar positions, 75% were hired, while 85% of the females were hired.

A federal Equal Employment enforcement official noted that many more males were hired than females, and decided to investigate. Responding to charges of irregularities in hiring, the company president denied any discrimination, pointing out that in both the white collar and blue collar fields, the percentage of female applicants hired was greater than it was for males.

But the government official produced his own statistics, which showed that a female applying for a job had a 58% chance of being denied employment while male applicants had only a 45% denial rate. As the current law is written, this constituted a violation....

Can you explain how two opposing statistical outcomes are reached from the same raw data?"

## Contingency Tables

---

Here is the data:

<b>White</b>	Hired		
	Yes	No	
Male	<b>30</b>	<b>170</b>	200
Female	<b>40</b>	<b>160</b>	200
	70	330	400

$$P(H|M) = 30/200 = 0.15$$

$$P(H|F) = 40/200 = 0.20$$

<b>Blue</b>	Hired		
	Yes	No	
Male	<b>300</b>	<b>100</b>	400
Female	<b>85</b>	<b>15</b>	100
	385	115	500

$$P(H|M) = 300/400 = 0.75$$

$$P(H|F) = 85/100 = 0.85$$

<b>All</b>	Hired		
	Yes	No	
Male	<b>330</b>	<b>270</b>	600
Female	<b>125</b>	<b>175</b>	300
	455	445	900

$$P(H|M) = 330/600 = 0.55$$

$$P(H|F) = 125/300 = 0.42$$

## Contingency Tables

---

Marilyn, correctly notes that, even though all the figures presented are correct, the two outcomes are not opposing:

"Say a company tests two treatments for an illness. In trial No. 1, treatment A cures 20% of its cases (40 out of 200) and treatment B cures 15% of its cases (30 out of 200). In trial No. 2, treatment A cures 85% of its cases (85 out of 100) and treatment B cures 75% of its cases (300 out of 400)....

So, in two trials, treatment A scored 20% and 85%. Also in two trials, treatment B scored only 15% and 75%. No matter how many people were in those trials, treatment A (at 20% and 85%) is surely better than treatment B (at 15% and 75%), right?

*Wrong! Treatment B performed better. It cured 330 (300+30) out of the 600 cases (200+400) in which it was tried --a success rate of 55%...By contrast, treatment A cured 125 (40+85) out of the 300 cases (200+100) in which it was tried, a success rate of only about 42%."*

She notes that this is exactly what happened to the employer. Because so many more men applied for the blue collar positions, even if the employer hired all the women who had applied for blue collar positions, it couldn't satisfy the government regulations.

### Exact Tests for Probabilities

To estimate the expected value of a random variable on the basis of  $n$  independent observations, we use the sample mean. To set a confidence interval, we can use the fact that

$$\frac{\bar{X}_n - \mu}{\hat{SE}(\bar{X}_n)}$$

has an approximate standard normal distribution. (The bigger  $n$  is, the better the approximation.) From this we find, when  $n$  is large,

$$95\% \text{ Confidence Interval: } \bar{X}_n \pm 1.96 \hat{SE}(\bar{X}_n)$$

Level  $\alpha = 0.05$  Test: Reject  $H_0: \mu = \mu_0$  vs  $H_A: \mu > \mu_0$  if

$$\frac{\bar{X}_n - \mu_0}{\hat{SE}(\bar{X}_n)} > 1.645.$$

If when the hypothesis  $H_0: \mu = \mu_0$  is true, the value of  $SE(\bar{X}_n)$  is known, then that value often replaces the estimate,  $\hat{SE}(\bar{X}_n)$ , in the denominator of the test statistic.

## Contingency Tables

---

For example, this is the case when the  $X$ 's have a Bernoulli( $\theta$ ) distribution, because  $H_0: \theta = \theta_0$  implies that  $SE(\bar{X}_n) = \sqrt{\theta_0(1-\theta_0)/n}$ . The usual test statistic is  $(\bar{X}_n - \theta_0) / \sqrt{\theta_0(1-\theta_0)/n}$ .

To test the hypothesis that the probability of heads with my 40¢ piece is 0.50 (just like most other coins) we toss it 80 times, and observe the number of heads. Suppose we see 50 heads in our 80 tosses.

Under the null hypothesis,  $H_0$ , we know the variance. It is  $\theta_0(1-\theta_0)$ , which in the present example, is 1/4. So instead of estimating the variance, we can use the test statistic,

$$\frac{\sqrt{n}(\hat{\theta} - 0.5)}{\sqrt{0.25}},$$

rejecting the hypothesis  $H_0: \theta = 0.5$  vs  $H_A: \theta > 0.5$  if the test statistic exceeds the critical value found in the z-table.

## Contingency Tables

---

For our observations (50 heads in 80 tosses),

$$\hat{\theta} = 50/80 = 0.625,$$

so the observed value of the approximate normal test statistic is

$$\frac{\sqrt{80}(0.625-0.5)}{\sqrt{0.25}} = 2.236$$

and the approximate p-value is  $P(Z \geq 2.236) = 0.013$ . This evidence is strong enough to justify rejecting  $H_0$  at the 5% level, but not at the 1% level ( $0.01 < p < 0.05$ ).

The general procedure is this: Reject  $H_0: \theta = \theta_0$  if the approximate p-value is  $\leq \alpha$  :

$$P\left(Z > \frac{\sqrt{n}(\hat{\theta}_{obs} - \theta_0)}{\sqrt{\theta_0(1 - \theta_0)}}\right) < \alpha$$

This test is based on the fact that if  $H_0$  is true,

$$E(\hat{\theta}) = \theta_0 \quad \text{Var}(\hat{\theta}) = \theta_0(1 - \theta_0)/n$$

and if  $n$  is large, the probability distribution of  $\hat{\theta}$  can be approximated by the normal distribution with this mean and variance.

## Contingency Tables

---

STATA does this test for me at the command **prtest** which instructs it to perform "One- and two-sample tests of proportions":

```
. prtesti 80 50 0.5
```

One-sample test of proportion x: Number of obs = 80

Variable	Mean	Std. Err.	z	P> z	[95% Conf. Interval]
x	.625	.0541266	11.547	0.0000	.5189138 .7310862

Ho: proportion(x) = .5

Ha: x < .5	Ha: x ~ = .5	Ha: x > .5
z = 2.236	z = 2.236	z = 2.236
P < z = 0.9873	P >  z  = 0.0253	P > z = 0.0127

But if  $H_0$  is true, we know not only the mean and the variance; we know the entire probability distribution of  $\hat{\theta}$ . Since  $\hat{\theta} = X/n$ , where  $X \sim \text{binomial}(n, \theta_0)$ , we know that the probability that  $\hat{\theta}$  will equal  $k/n$  is precisely

$$\binom{n}{k} \theta_0^k (1 - \theta_0)^{n-k} \quad \text{for } k = 0, 1, \dots, n$$

Thus we can use the exact probability distribution—we don't have to use an approximation.

## Contingency Tables

---

For the 80 tosses of my 40¢ piece, with  
 $H_0: \theta = 0.5$  and  $H_A: \theta > 0.5$  the exact p-value is

$$p\text{-value} = P(\hat{\theta} \geq 0.625) = P(X \geq 50 \mid \theta = 0.5)$$

$$= \sum_{x=50}^{80} \binom{80}{x} (0.5)^x (1 - 0.5)^{80-x}$$

$$= 0.016496$$

STATA does this test for me at the command **bitest** which instructs it to perform a "Binomial probability test":

```
. bitesti 80 50 0.5
```

N	Observed k	Expected k	Assumed p	Observed p
80	50	40	0.50000	0.62500
<hr/>				
Pr(k >= 50)		= 0.016496	(one-sided test)	
Pr(k <= 50)		= 0.990842	(one-sided test)	
Pr(k <= 30 or k >= 50)		= 0.032993	(two-sided test)	

## Contingency Tables

---

The real usefulness of the exact test procedure appears when the sample size is small. We can perform precise hypothesis tests, even when we have only a small sample.

*In general, to test  $H_0: \theta = \theta_0$  vs  $H_A: \theta > \theta_0$  at level  $\alpha = 0.05$  we will reject  $H_0$  if  $\hat{\theta} = X/n$  is too large, i.e., if  $X$  is too large. How large? Instead of rejecting if the approximate p-value is  $\leq 0.05$ , calculate the exact p-value and reject  $H_0$  if it is  $\leq 0.05$ :*

$$p\text{-value} = P( X \geq x_{obs} ) = \sum_{x_{obs}}^n \binom{n}{k} \theta_0^k (1 - \theta_0)^{n-k}$$

## Contingency Tables

---

Example: Test  $H_0: \theta = 1/3$  vs  $H_A: \theta > 1/3$ . If  $n = 4$  and we observe 3 successes, what is the p-value?

$$P(X \geq 3 \mid \theta = 1/3) = \binom{4}{3} (1/3)^3 (2/3) + (1/3)^4 = \frac{1}{9} = 0.111111.$$

If  $n=4$  and we observe 4 successes, what is the p-value?

$$P(X \geq 4 \mid \theta = 1/3) = (1/3)^4 = \frac{1}{81} = 0.012346.$$

Suppose we want a two-sided test of

$$H_0: \theta = 1/3 \text{ vs } H_A: \theta \neq 1/3$$

when  $n = 4$  and we observe 4 successes. What is the p-value?

One solution: Calculate the one-sided test p-value and double it. This gives  $p = 0.02469$ . But this procedure can lead to p-values  $> 1$ . For example, for testing  $H_0: \theta = 1/3$  vs  $H_A: \theta > 1/3$  when  $n = 5$ , and  $x = 2$  is observed, the (one-sided) p-value is

$$p\text{-value} = P(\hat{\theta} \geq 0.40) = P(X \geq 2 \mid \theta = 1/3) = 0.539094$$

Doubling this gives a two-sided "p-value" of 1.078

Other suggestions have been made, and the previous release of STATA reported one of them, as well as the simple "double the one-sided p-value" procedure. But they also have their unsatisfactory aspects. Statistics has no definitive solutions for these very simple problems (What is the "right way" to do exact two-sided tests of hypotheses about binomial probabilities, and what is the "right way" to set exact two-sided confidence intervals?).

An important advantage of the "approximate normal" testing and confidence interval machinery is that in the case of the nice symmetric normal distribution it is clear what to do about two-sided tests and confidence intervals.

What to do when you face this problem?

- (a) Follow whatever conventions you find in your department/discipline.
- (b) Report one-sided p-value, clearly identified.

Most readers will simply double it if they "want to be conservative."

## Contingency Tables

---

Exercise: (For education and entertainment only)  
Suppose you and your friend pass the time by playing a game (instead of studying Biostatistics). Your friend insists that it is "all luck," but you think the game requires some skill, of which you are confident that you have more than your friend. Suppose you keep score and find that you have won 8 games out of ten.

(a) What probability model might you use to test the hypothesis

*$H_0$ : All luck vs  $H_A$ : You are better player ?*

(b) What is the p-value?

### Fisher's Exact Test of Equality of Two Probabilities

We have seen three ways to test the hypothesis that two binomial probabilities are equal:

(i) calculate

$$\frac{|\hat{\theta}_1 - \hat{\theta}_2|}{\sqrt{\hat{\theta}_p(1 - \hat{\theta}_p)\left(\frac{1}{n_1} + \frac{1}{n_2}\right)}},$$

(approximate)  $p\text{-value} = P(|Z| \geq |z_{obs}|)$ , and

(ii) calculate

$$\chi^2 = \sum \frac{(\text{observed} - \text{expected})^2}{\text{expected}}$$

(approximate)  $p\text{-value} = P(\chi^2 \geq \chi_{obs}^2)$ .

These two tests are actually equivalent. When the numbers of observations in the four cells are not very large, a "continuity correction" is often made in an effort to improve the accuracy of the approximate p-value or Type I error rate.

## Contingency Tables

---

The correction reduces the value of the test statistic, thereby increasing the p-value, and making it more difficult to reject the null hypothesis:

(iii) calculate

$$\chi_c^2 = \sum \frac{(|\text{observed} - \text{expected}| - 0.5)^2}{\text{expected}}$$

$$(\text{approximate}) \quad p\text{-value} = P(\chi^2 \geq \chi_{c\text{ obs}}^2) .$$

Exercise: If I perform a chi-square test on a 2-by-2 table without the continuity correction, then later do the test with the correction, how will the second p-value (the corrected one) compare to the first?

- ☐ It will definitely be smaller.
- ☐ It will be the same.
- ☐ It might be larger or smaller, depending on the numbers in the particular table that I am analyzing.
- ☐ It will definitely be larger.

## Contingency Tables

---

Another solution was invented by the greatest of the statisticians, Sir Ronald A. Fisher. It is called "Fisher's Exact Test," and provides exact p-values no matter how small the numbers in the table.

	Success	Failure	
Sample 1	a	b	$a+b=n_1$
Sample 2	c	d	$c+d=n_2$
	$a+c$	$b+d$	n

This test is based on recognizing an important fact about the probabilities of tables of this sort: if the success probabilities in the two rows are equal ( $H_0$ ), then given the total number of successes,  $a+c$ , the conditional probability that  $a$  of the successes will occur in the first row (and  $c$  in the second) is

$$\frac{\binom{a+c}{a} \binom{b+d}{b}}{\binom{n}{a+b}}$$

## Contingency Tables

---

It is as if we simply had  $n$  subjects of which  $a+c$  are "successes" and chose  $n_1 = a+b$  of them at random to be "sample 1," letting the rest be "sample 2." The above expression gives the probability that, if we divide the  $n$  subjects in this way, we will find that  $a$  of the  $a+c$  successes are in "sample 1."

	Success	Failure	
Sample 1	8	2	$10=n_1$
Sample 2	1	4	$5=n_2$
	9	6	$15=n$

If sample 1 and sample 2 really represent, in effect, a random division of the 15 subjects into groups of 10 and 5, what is the probability that 8 of the 9 successes will go into sample 1?

$$\frac{\binom{9}{8}\binom{6}{2}}{\binom{15}{10}} = \frac{135}{3003} = 0.045$$

## Contingency Tables

---

Fisher's exact test calculates as the p-value the probability of the observed table, plus the probabilities of all tables that are "as extreme or more so" than the one observed.

For a one-sided test "More extreme" tables are those with the same margins as the observed table and either

- (i) even larger values of " $a$ " than the one observed (when the observed " $a$ " is greater than the value expected under  $H_0$ ,  $(a+b)(a+c)/n$ , as it is in our example), or
- (ii) even smaller values of " $a$ " than the one observed (when the observed " $a$ " is smaller than the value expected under  $H_0$ ).

For a two-sided test "More extreme" tables are those with the same margins as the observed table, and which are "more improbable" under  $H_0$  ( regardless of whether " $a$ " is in the same side of the expected value as the one observed or not).

## Contingency Tables

---

In our example, the observed  $a=8$  is larger than the value expected under  $H_0$ ,  $10(9/15)=6$ , so the table

	Success	Failure	
Sample 1	9	1	$10=n_1$
Sample 2	0	5	$5=n_2$
	9	6	$15=n$

is more extreme in the same direction, with probability

$$\frac{\binom{9}{9}\binom{6}{1}}{\binom{15}{10}} = \frac{6}{3003} = 0.002$$

Since there are no more extreme tables in this direction (no tables with the same margins and even larger values of " $a$ " than 9), the one-sided p-value is  $0.045 + 0.002 = 0.047$ .

## Contingency Tables

---

To calculate the two-sided p-value, we must find the tables (with the same margins) that are as improbable as the one observed (or more so) in the other direction (few successes in sample 1, i.e., small values of " $a$ "). The first candidate is

	Success	Failure	
Sample 1	4	6	$10=n_1$
Sample 2	5	0	$5=n_2$
	9	6	$15=n$

$$\frac{\binom{9}{4}\binom{6}{6}}{\binom{15}{10}} = \frac{126}{3003} = 0.042$$

	Success	Failure	
Sample 1	5	5	$10=n_1$
Sample 2	4	1	$5=n_2$
	9	6	$15=n$

The probability of this table under  $H_0$  (conditional probability, given a total of 9 successes) is greater than that of the observed table:

$$\frac{\binom{9}{5}\binom{6}{5}}{\binom{15}{10}} = \frac{756}{3003} = 0.252$$

## Contingency Tables

---

so this last table's probability is not included in the two-sided p-value, which is the probability of the observed table (with  $a = 8$ ) plus the probabilities of the two more extreme tables (with  $a = 9$  and  $a = 4$ ):  $0.045 + 0.002 + 0.042 = 0.089$ . (Note that  $a = 3$  is even further from the expected value than  $a = 4$ , but it is not possible, given the marginal totals of the table.)

STATA does the calculation for you:

```
. tabi 8 2 \ 1 4
```

row	col	1	2	Total
1		8	2	10
2		1	4	5
Total		9	6	15

  

Fisher's exact =	0.089
1-sided Fisher's exact =	0.047

## Contingency Tables

---

Fisher's test does not generalize to  $r \times c$  tables as easily as the Chi-squared test does; however, there are acceptable methods used in practice. R implements two of them with the command `fisher.test`.

Extend our last example by adding sample 3.

	Success	Failure	
Sample 1	8	2	$10=n_1$
Sample 2	1	4	$5=n_2$
Sample 3	0	4	$4=n_3$
	9	10	$19=n$

Notice that knowing the value of the margins and one cell no longer gives you all the cell values.

```
x <- matrix(c(8, 1, 0, 2, 4, 4), ncol = 2)
x
chisq.test(x)
# notice the chi-squared test gives a warning
fisher.test(x)
fisher.test(x, simulate.p.value=TRUE)
```