Biostatistics – Spring 2026
Lecture 03: Fisher's Exact Test

Dr. Zaid T. Al-Khaledi
Department of Statistics and Informatics
University of Mosul

February 25, 2026

Introduction

In Lecture 2 we learned the Chi-square test of independence for contingency tables. That test is very popular and very useful, but it is not always reliable.

This lecture answers a practical question:

> When should I **avoid** Chi-square, and what should I use instead?

We focus on 2×2 tables and learn Fisher's Exact Test, which is designed for small samples.

By the end of this lecture, you should be able to:

- Why Chi-square Can Be Wrong for Small Samples.
- Why We Use Fisher's Exact Test.
- Fisher's Exact Test: the 2x2 Table and Probability.
- Example 1 (Manual Fisher).
- Example 2 (Biomedical Example).
- Example 3 (When Chi-square and Fisher Both Work).

1. Why Chi-square Can Be Wrong for Small Samples

1.1 The key idea behind Chi-square

Chi-square compares observed frequencies O_{ij} with expected frequencies E_{ij} under independence:

$$E_{ij} = \frac{(\text{row total}_i)(\text{column total}_j)}{n}.$$

Then it uses the statistic

$$\chi^2 = \sum_{i=1}^{r} \sum_{j=1}^{c} \frac{(O_{ij} - E_{ij})^2}{E_{ij}}.$$

The important point is this:

> The Chi-square p-value comes from an **approximation** to the true sampling distribution.

1.2 When the approximation becomes poor

The approximation becomes poor when the table is **sparse** (many small counts).

A common practical rule for a 2×2 table:

- If any expected count $E_{ij} < 5$, Chi-square may be inaccurate.
- If several expected counts are below 5, Chi-square is usually not recommended.

Why is this a problem? Because with small counts, the true distribution of the statistic is not close enough to a Chi-square curve.

1.3 What do we do instead?

When the sample is small, we use an **exact** method:

> Fisher's Exact Test: it calculates the exact probability of the observed pattern of counts.

2. Why We Use Fisher's Exact Test (and not other tests)

There are several tests for association in 2×2 tables:

- Chi-square test (approximate; good for large samples)
- Likelihood-ratio Chi-square (also approximate; good for large samples)
- Fisher's Exact Test (exact; best for small samples)

So why Fisher, specifically?

2.1 Fisher is exact (not approximate)

Fisher's test uses the **Hypergeometric distribution**. It does not rely on "large sample" assumptions. That is the main reason we use it in biomedical small studies.

2.2 Fisher is designed for 2×2 tables with fixed margins

In many designs, the row totals and/or column totals are fixed by design. For example:

- Case-control studies: number of cases and controls can be fixed.
- Small clinical comparisons: total treated and untreated may be fixed.

Fisher conditions on these fixed totals and evaluates how unusual the observed table is under H_0.

2.3 Fisher stays valid when expected counts are tiny

Even if one cell has expected count < 1, Fisher remains valid. Chi-square does not.

2.4 The price we pay

Fisher can be conservative (sometimes p-values are not as small as Chi-square for the same table), and for very large samples it can be slower. But for the course level: the rule is simple:

> Small sample / small expected counts $\Rightarrow$ Fisher.
>
> Large sample $\Rightarrow$ Chi-square is fine.

3. Fisher's Exact Test: the 2×2 Table and Probability

Consider the 2×2 table:

	B_1	B_2	Row total
A_1	a	b	r_1
A_2	c	d	r_2
Col total	c_1	c_2	n

where

$$r_1 = a + b, \quad r_2 = c + d, \quad c_1 = a + c, \quad c_2 = b + d, \quad n = r_1 + r_2.$$

3.1 Hypotheses

$$H_0: \text{ A and B are independent.} \qquad H_1: \text{ A and B are associated.}$$

3.2 Exact probability of one table (hypergeometric form)

When row totals and column totals are fixed, the only free value is a. Once a is known, the rest are determined:

$$b = r_1 - a, \quad c = c_1 - a, \quad d = r_2 - c.$$

Under H_0, the probability of observing a specific value a is:

$$P(a) = \frac{\binom{r_1}{a}\binom{r_2}{c_1 - a}}{\binom{n}{c_1}}.$$

This form is easier than factorials and is the standard Fisher formula.

3.3 What is the p-value in Fisher?

Conceptually:

> The Fisher p-value is the sum of probabilities of the observed table and all tables that are "as extreme as" the observed table (in the direction of association).

In many software packages, two-sided Fisher p-value is computed by summing all tables with probability less than or equal to the observed table probability.

We will use that practical idea in our examples.

4. Example 1 (Manual Fisher with Small Numbers)

Suppose a small study tests whether a new disinfectant is related to bacterial growth (Yes/No).

	Growth Yes	Growth No	Total
New Disinfectant	1	7	8
Old Disinfectant	6	2	8
Total	7	9	16

This is a small table, and some expected counts are small, so Fisher is suitable.

Step 1: Identify margins

$$r_1 = 8, \quad r_2 = 8, \quad c_1 = 7, \quad c_2 = 9, \quad n = 16.$$

Let a be the count in cell (New disinfectant, Growth Yes). Observed $a_{\text{obs}} = 1$.

Step 2: Find all possible values of a

Because counts cannot be negative,

$$\max(0, c_1 - r_2) \le a \le \min(r_1, c_1).$$

Here:

$$\max(0, 7 - 8) = 0 \le a \le \min(8, 7) = 7.$$

So $a = 0, 1, 2, 3, 4, 5, 6, 7$ are possible.

Step 3: Compute $P(a)$ using combinations

$$P(a) = \frac{\binom{8}{a}\binom{8}{7-a}}{\binom{16}{7}}.$$

Compute the observed probability:

$$P(1) = \frac{\binom{8}{1}\binom{8}{6}}{\binom{16}{7}}.$$

Now,

$$\binom{8}{1} = 8, \quad \binom{8}{6} = 28, \quad \binom{16}{7} = 11440.$$

So

$$P(1) = \frac{8 \times 28}{11440} = \frac{224}{11440} \approx 0.0196.$$

Step 4: Complete Computation of Fisher p-value

Since the denominator is fixed and equals $\binom{16}{7} = 11440$, it is enough to compute the numerator and choose the values of a that satisfies $\binom{8}{a}\binom{8}{7-a} \le 224$. Hence, we compute $\binom{8}{a}\binom{8}{7-a}$ for all possible values of a:

a	$\binom{8}{a}\binom{8}{7-a}$
0	$1 \cdot 8 = 8$
1	$8 \cdot 28 = 224$
2	$28 \cdot 56 = 1568$
3	$56 \cdot 70 = 3920$
4	$70 \cdot 56 = 3920$
5	$56 \cdot 28 = 1568$
6	$28 \cdot 8 = 224$
7	$8 \cdot 1 = 8$

Two-Sided Fisher p-value

For the two-sided Fisher p-value (practical definition), we sum all probabilities less than or equal to $P(1)$.

So we take the values where $\binom{8}{a}\binom{8}{7-a} \leq 224$. Here we have:

$$P(0),\ P(1),\ P(6),\ P(7) \leq 0.0196,$$

Therefore, the two-sided p-value is:

$$p - value = P(0) + P(1) + P(6) + P(7)$$

$$= \frac{8 + 224 + 224 + 8}{11440} = \frac{464}{11440} = \frac{29}{715} \approx 0.04056.$$

Final Interpretation

Since the two-sided p-value is

$$p - value \approx 0.0406 < 0.05,$$

we conclude:

> There is statistically significant evidence at the 5% level that disinfectant type and bacterial growth are associated.

5. Example 2 (Fisher in a Biomedical Context: Rare Event)

A small clinical observation checks whether a medicine is related to a rare side effect.

	Side Effect Yes	Side Effect No	Total
Medicine A	0	10	10
Medicine B	3	7	10
Total	3	17	20

Here we clearly have a very small cell (0). Chi-square is risky. Fisher is appropriate.

Step 1: Identify margins

$$r_1 = 10, \quad r_2 = 10, \quad c_1 = 3, \quad c_2 = 17, \quad n = 20.$$

Let a be the count in cell (Medicine A, Side Effect Yes).
Observed:

$$a_{\text{obs}} = 0.$$

Step 2: Possible values of a

$$\max(0, c_1 - r_2) = \max(0, 3 - 10) = 0$$
$$\min(r_1, c_1) = \min(10, 3) = 3.$$

So

$$a = 0, 1, 2, 3.$$

Step 3: Compute $P(a)$

$$P(a) = \frac{\binom{10}{a}\binom{10}{3-a}}{\binom{20}{3}}.$$

First compute the denominator:

$$\binom{20}{3} = 1140.$$

Now compute the numerator for all possible values of a:

a	$\binom{10}{a}\binom{10}{3-a}$
0	$1 \cdot 120 = 120$
1	$10 \cdot 45 = 450$
2	$45 \cdot 10 = 450$
3	$120 \cdot 1 = 120$

Therefore,

$$P(0) = \frac{120}{1140} \approx 0.1053$$
$$P(1) = \frac{450}{1140} \approx 0.3947$$
$$P(2) = \frac{450}{1140} \approx 0.3947$$
$$P(3) = \frac{120}{1140} \approx 0.1053$$

Two-Sided Fisher p-value

The observed probability is

$$P(0) = 0.1053.$$

For the practical two-sided Fisher definition, we sum all probabilities less than or equal to $P(0)$.

From the table:

$$P(0) = P(3) = 0.1053$$

while

$$P(1), P(2) = 0.3947 > 0.1053.$$

Thus we include only $a = 0$ and $a = 3$.

$$p\text{-value} = P(0) + P(3)$$

$$= \frac{120 + 120}{1140} = \frac{240}{1140} = \frac{4}{19} \approx 0.2105.$$

Final Interpretation

Since

$$p\text{-value} \approx 0.2105 > 0.05,$$

we conclude:

> We do not have sufficient statistical evidence at the 5% level to conclude that medicine type and side effect occurrence are associated.

This does not mean that no association exists; it means that the data do not provide strong enough evidence to detect one.

6. Example 3 (When Chi-square and Fisher Both Work)

Consider a larger 2×2 table:

	Positive Test	Negative Test	Total
Exposed Group	30	70	100
Not Exposed	20	80	100
Total	50	150	200

Expected counts are:

$$E_{11} = \frac{(100)(50)}{200} = 25, \quad E_{12} = 75, \quad E_{21} = 25, \quad E_{22} = 75,$$

all are ≥ 5. So Chi-square is safe. Fisher is also possible, but not necessary.

So in practice:

- Large sample $\Rightarrow$ Chi-square is fine and fast.
- Small sample / tiny expected $\Rightarrow$ Fisher is safer.

Homework (HW)

HW1 (Fisher Exact Test)

Use Fisher's Exact Test at $\alpha = 0.05$:

	Cancer	No Cancer	Total
Coffee Drinkers	4	1	5
Non-Drinkers	1	8	9
Total	5	9	14

Required:

1. Write H_0 and H_1.
2. Identify r_1, r_2, c_1, c_2, n.
3. List all possible a values.
4. Compute $P(a_{\text{obs}})$.
5. Compute the two-sided Fisher p-value.
6. Make a decision about H_0.

HW2 (Decide the Test)

For each table below, answer:

- Chi-square or Fisher?
- Why?

(A)

	+	−
Treated	2	18
Control	9	11

(B)

	+	−
*Group*1	40	60
*Group*2	35	65