EXAM PAPERS PRACTICE

Boost your performance and confidence with these topic-based exam questions

Practice questions created by actual examiners and assessment experts

Detailed mark scheme

Suitable for all boards

Designed to test your ability and thoroughly prepare you

4.11 Hypothesis Testing

IB Maths - Revision Notes

4.11.1 Hypothesis Testing

Language of Hypothesis Testing

What is a hypothesis test?

- A hypothesis test uses a **sample of data** in an experiment to test a **statement** made about the **population**
  - The statement is either about a **population parameter** or the distribution of the **population**
- The hypothesis test will look at the probability of observed outcomes happening under set conditions
- The probability found will be compared against a given **significance level** to determine whether there is **evidence to support** the statement being made

What are the key terms used in statistical hypothesis testing?

- Every hypothesis test **must** begin with a clear **null hypothesis** (what we believe to already be true) and **alternative hypothesis** (how we believe the data pattern or probability distribution might have changed)
- A **hypothesis** is an assumption that is made about a particular **population parameter** or the **distribution of the population**
  - A **population parameter** is a numerical characteristic which helps define a population
    - Such as the mean value of the population
  - The **null hypothesis** is denoted H_0 and sets out the assumed population parameter or distribution given that no change has happened
  - The **alternative hypothesis** is denoted H_1 and sets out how we think the population parameter or distribution could have changed
  - When a hypothesis test is carried out, the null hypothesis is **assumed to be true** and this assumption will either be **accepted** or **rejected**
    - When a null hypothesis is accepted or rejected a **statistical inference** is made
- A hypothesis test will always be carried out at an appropriate **significance level**
  - The significance level sets the **smallest probability** that an event could have occurred by chance
    - Any probability smaller than the significance level would suggest that the event is unlikely to have happened by chance
  - The **significance level** must be set **before** the hypothesis test is carried out
  - The **significance level** will usually be 1%, 5% or 10%, however it may vary

One-tailed Tests

What are one-tailed tests?

- A **one-tailed test** is used for testing:
  - Whether a distribution can be used to model the population
  - Whether the population parameter has **increased**
  - Whether the population parameter has **decreased**
- **One-tailed tests** can be used with:
  - Chi-squared test for independence
  - Chi-squared goodness of fit test
  - Test for proportion of a binomial distribution
  - Test for population mean of a Poisson distribution
  - Test for population mean of a normal distribution
  - Test to compare population means of two distributions

Two-tailed Tests

What are two-tailed tests?

- A **two-tailed test** is used for testing:
  - Whether the population parameter has **changed**
- **Two-tailed tests** can be used with:
  - Test for population mean of a normal distribution
  - Test to compare population means of two distributions

Conclusions of Hypothesis Testing

How do I decide whether to reject or accept the null hypothesis?

- A sample of the population is taken and the **test statistic** is calculated **using the observations** from the sample
  - Your GDC can calculate the test statistic for you (if required)
- To decide whether or not to reject the null hypothesis you first need either the ***p*-value** or the **critical region**
- The ***p* - value** is the probability of a value being **at least as extreme** as the test statistic, assuming that the null hypothesis is true
  - Your GDC will give you the *p*-value (if required)
  - If the ***p*-value is less than the significance level** then the **null hypothesis** would be **rejected**
- The **critical region** is the range of values of the test statistic which will lead to the **null hypothesis** being **rejected**
  - If the **test statistic** falls within the **critical region** then the **null hypothesis** would be **rejected**
- The **critical value** is the boundary of the critical region
  - It is the least extreme value that would lead to the rejection of the null hypothesis
  - The **critical value** is determined by the **significance level**

How should a conclusion be written for a hypothesis test?

- Your conclusion **must** be written in the **context** of the question
- Use the **wording in the question** to help you write your conclusion
  - If **rejecting the null** hypothesis your conclusion should state that there is **sufficient evidence** to suggest that the null hypothesis is unlikely to be true
  - If **accepting the null** hypothesis your conclusion should state that there is **not enough evidence** to suggest that the null hypothesis is unlikely to be true
- Your conclusion **must not** be definitive
  - There is a chance that the test has led to an **incorrect conclusion**
  - The outcome is **dependent on the sample**
    - a **different sample** might lead to a **different outcome**
- The conclusion of a **two-tailed test** can state if there is evidence of a change
  - You should not state whether this change is an increase or decrease
  - If you are testing the difference between the means of two populations then you can **only conclude that the means are not equal**
    - You can not say which population mean is bigger
    - You'd need to use a **one-tailed** test for this

Exam Tip

- Accepting the null hypothesis does **not** mean that you are saying it is true
  - You are simply saying there is not enough evidence to reject it

4.11.2 Chi-squared Test for Independence

Chi-Squared Test for Independence

What is a chi-squared test for independence?

- A chi-squared (χ^2) **test for independence** is a hypothesis test used to test whether **two variables are independent** of each other
  - This is sometimes called a χ^2 **two-way test**
- This is an example of a **goodness of fit** test
  - We are testing whether the data fits the model that the variables are independent
- The chi-squared (χ^2) distribution is used for this test
- You will use a **contingency table**
  - This is a two-way table that shows the **observed frequencies** for the different combinations of the two variables
    - For example: if the two variables are hair colour and eye colour then the contingency table will show the frequencies of the different combinations

Why might I have to combine rows or columns?

- The **observed** values are used to calculate **expected** values
  - These are the expected frequencies for each combination **assuming that the variables are independent**
    - Your GDC can calculate these for you after you input the observed frequencies
- The **expected values** must all be **bigger than 5**
- If one of the expected values is less than 5 then you will have to **combine the corresponding row or column** in the matrix of **observed values** with the **adjacent** row or column
  - The decision between row or column will be based on which seems the **most appropriate**
    - For example: if the two variables are age and favourite TV genre then it is more appropriate to combine age groups than types of genre

What are the degrees of freedom?

- There will be a **minimum number of expected values** you would need to know in order to be able to calculate all the expected values
- This minimum number is called the **degrees of freedom** and is often denoted by ν
- For a **test for independence** with an $m \times n$ contingency table
  - $\nu = (m-1)\times(n-1)$
  - For example: If there are 5 rows and 3 columns then you only need to know **2 of the values** in **4 of the rows** as the rest can be calculated using the totals

What are the steps for a chi-squared test for independence?

- **STEP 1**: Write the **hypotheses**
  - H_0 : Variable *X* is independent of variable *Y*
  - H_1 : Variable *X* is not independent of variable *Y*
    - Make sure you clearly write what the variables are and don't just call them *X* and *Y*
- **STEP 2**: Calculate the **degrees of freedom** for the test
  - For an $m \times n$ contingency table
  - Degrees of freedom is $\nu = (m-1)\times(n-1)$
- **STEP 3**: Enter your **observed frequencies** into your GDC using the option for a 2-way test
  - Enter these as a matrix
  - Your GDC will give you a matrix of the **expected values** (assuming the variables are independent)
    - If any values are 5 or less then combine rows/columns and **repeat step 2**
  - Your GDC will also give you the χ^2 statistic and its *p*-value
  - The χ^2 statistic is denoted as χ^2_{calc}
- **STEP 4**: Decide whether there is evidence to reject the null hypothesis
  - EITHER compare the **χ^2 statistic** with the given **critical value**
    - If χ^2 statistic **>** critical value then **reject H_0**
    - If χ^2 statistic **<** critical value then **accept H_0**
  - OR compare the ***p*-value** with the given **significance level**
    - If *p*-value **<** significance level then **reject H_0**
    - If *p*-value **>** significance level then **accept H_0**
- **STEP 5**: Write your **conclusion**
  - If you **reject H_0**
    - There is sufficient evidence to suggest that variable *X* is not independent of variable *Y*
    - Therefore this suggests they are **associated**
  - If you **accept H_0**
    - There is insufficient evidence to suggest that variable *X* is not independent of variable *Y*
    - Therefore this suggests they are **independent**

How do I calculate the chi-squared statistic?

- You are **expected** to be able to use your **GDC** to calculate the χ^2 statistic by inputting the matrix of the observed frequencies
- Seeing how it is done by hand might deepen your understanding but you are **not expected** to use this method
- **STEP 1**: For each **observed frequency** O_i calculate its **expected frequency** E_i
  - Assuming the variables are independent
    - $E_i = P(X = x) \times P(Y = y) \times \text{Total}$
    - Which simplifies to $E_i = \dfrac{\text{Row Total} \times \text{Column Total}}{\text{Overall Total}}$

STEP 2 Calc late the χ^2 statistic sing the form la

- $$\chi^2_{calc} = \sum \frac{(O_i - E_i)^2}{E_i}$$
  - You do not need to learn this formula as your GDC calculates it for you
- To calculate the *p*-value you would find the probability of a value being bigger than your χ^2 statistic using a χ^2 distribution with ν degrees of freedom

Exam Tip

Note for Internal Assessments (IA)

- If you use a χ^2 test in your IA then beware that the outcome may not be accurate if there is only 1 degree of freedom
  - This means it is a 2×2 contingency table

Worked example

At a school in Paris, it is believed that favourite film genre is related to favourite subject. 500 students were asked to indicate their favourite film genre and favourite subject from a selection and the results are indicated in the table below.

	Comedy	Action	Romance	Thriller
Maths	51	52	37	55
Sports	59	63	41	33
Geography	35	31	28	15

It is decided to test this hypothesis by using a χ^2 test for independence at the 1% significance level.

The critical value is 16.812.

a) State the null and alternative hypotheses for this test.

H_0 : Favourite subject is independent of favourite film genre
H_1 : Favourite subject is not independent of favourite film genre

b) Write down the number of degrees of freedom for this table.

$\nu = (\text{rows} - 1) \times (\text{columns} - 1) = (3-1) \times (4-1)$

$\boxed{\nu = 6}$

c) Calculate the χ^2 test statistic for this data.

Type matrix into GDC

χ^2 statistic $= 12.817...$

$\boxed{\chi^2_{calc} = 12.8 \ (3 \text{ sf})}$

d) Write down the conclusion to the test. Give a reason for your answer.

$12.8 < 16.812$

Accept H_0 as χ^2 statistic < critical value. There is insufficient evidence to suggest that favourite subject is not independent of favourite film genre. Therefore this suggests they are independent.

4.11.3 Goodness of Fit Test

Chi-Squared GOF: Uniform

What is a chi-squared goodness of fit test for a given distribution?

- A chi-squared (χ^2) **goodness of fit test** is used to test data from a sample which suggests that the population has a given distribution
- This could be that:
  - the proportions of the population for different categories follows a **given ratio**
  - the population follows a **uniform distribution**
    - This means all outcomes are **equally likely**

What are the steps for a chi-squared goodness of fit test for a given distribution?

- **STEP 1**: Write the **hypotheses**
  - H_0 : Variable X can be modelled by the given distribution
  - H_1 : Variable X cannot be modelled by the given distribution
    - Make sure you clearly write what the variable is and don't just call it X
- **STEP 2**: Calculate the **expected frequencies**
  - Split the total frequency using the given ratio
  - For a uniform distribution: divide the total frequency N by the number of possible outcomes k
- **STEP 3**: Calculate the **degrees of freedom** for the test
  - For k possible outcomes
  - Degrees of freedom is $v = k - 1$
- **STEP 4**: Enter the **frequencies** and the **degrees of freedom** into your GDC
  - Enter the observed and expected frequencies as two separate lists
  - Your GDC will then give you the χ^2 statistic and its p-value
  - The χ^2 statistic is denoted as χ^2_{calc}
- **STEP 5**: Decide **whether** there is evidence to **reject the null** hypothesis
  - EITHER compare the **χ^2 statistic** with the given **critical value**
    - If χ^2 statistic **>** critical value then **reject H_0**
    - If χ^2 statistic **<** critical value then **accept H_0**
  - OR compare the **p-value** with the given **significance level**
    - If p-value **<** significance level then **reject H_0**
    - If p-value **>** significance level then **accept H_0**
- **STEP 6**: Write your **conclusion**
  - If you **reject H_0**
    - There is sufficient evidence to suggest that variable X does not follow the given distribution
    - Therefore this suggests that the data is **not distributed as claimed**
  - If you **accept H_0**
    - There is insufficient evidence to suggest that variable X does not follow the given distribution
    - Therefore this suggests that the data is **distributed as claimed**

Worked example

A car salesman is interested in how his sales are distributed and records his sales results over a period of six weeks. The data is shown in the table.

Week	1	2	3	4	5	6
Number of sales	15	17	11	21	14	12

A χ^2 goodness of fit test is to be performed on the data at the 5% significance level to find out whether the data fits a uniform distribution.

a) Find the expected frequency of sales for each week if the data were uniformly distributed.

If uniformly distributed all expected frequencies are equal

$$\text{Expected frequency} = \frac{15+17+11+21+14+12}{6}$$

$$\boxed{\text{Expected frequency} = 15}$$

b) Write down the null and alternative hypotheses.

H_0: Number of sales can be modelled by a uniform distribution

H_1: Number of sales can not be modelled by a uniform distribution

c) Write down the number of degrees of freedom for this test.

$\nu = 6 - 1$

$\boxed{\nu = 5}$

d) Calculate the *p*-value.

Type two lists into GDC

Observed	15	17	11	21	14	12
Expected	15	15	15	15	15	15

$p = 0.4933...$

$\boxed{p = 0.493 \ (3sf)}$

e) State the conclusion of the test. Give a reason for your answer.

$0.493 > 0.05$

Accept H_0 as p-value > significance level
There is insufficient evidence to suggest that number of sales can not be modelled by a uniform distribution. Therefore this suggests it is uniformly distributed.

Chi-Squared GOF: Binomial

What is a chi-squared goodness of fit test for a binomial distribution?

- A chi-squared (χ^2) **goodness of fit test** is used to test data from a sample suggesting that the population has a **binomial distribution**
  - You will either be **given a precise binomial distribution** to test $\mathrm{B}(n, p)$ with an assumed value for *p*
  - Or you will be asked to test whether a binomial distribution is **suitable without being given an assumed value** for *p*
    - In this case you will have to calculate an **estimate** for the value of *p* for the binomial distribution
    - To calculate it divide the mean by the value of *n*
    - $p = \frac{\bar{x}}{n} = \frac{1}{n} \times \frac{\sum fx}{\sum f}$

What are the steps for a chi-squared goodness of fit test for a binomial distribution?

- **STEP 1**: Write the **hypotheses**
  - H_0 : Variable *X* can be modelled by a binomial distribution
  - H_1: Variable *X* cannot be modelled by a binomial distribution
    - Make sure you clearly write what the variable is and don't just call it *X*
    - If you are given the assumed value of *p* then state the precise distribution $\mathrm{B}(n, p)$
- **STEP 2**: Calculate the **expected frequencies**
  - If you were not given the assumed value of *p* then you will first have to **estimate it** using the **observed data**
  - Find the probability of the outcome using the binomial distribution $\mathrm{P}(X = x)$
  - Multiply the probability by the total frequency $\mathrm{P}(X = x) \times N$
  - You will have to combine rows/columns if any expected values are 5 or less
- **STEP 3**: Calculate the **degrees of freedom** for the test
  - For *k* outcomes (after combining expected values if needed)
  - Degrees of freedom is
    - $\nu = k - 1$ if you were **given** the assumed value of *p*
    - $\nu = k - 2$ if you had to **estimate** the value of *p*
- **STEP 4**: Enter the **frequencies** and the **degrees of freedom** into your GDC
  - Enter the observed and expected frequencies as two separate lists
  - Your GDC will then give you the χ^2 statistic and its *p*-value
  - The χ^2 statistic is denoted as χ^2_{calc}
- **STEP 5**: Decide whether there is **evidence** to **reject the null** hypothesis

- EITHER compare the χ^2 **statistic** with the given **critical value**
  - If χ^2 statistic **>** critical value then **reject H_0**
  - If χ^2 statistic **<** critical value then **accept H_0**
- OR compare the p-value with the given significance level
  - If p-value **<** significance level then **reject H_0**
  - If p-value **>** significance level then **accept H_0**

- **STEP 6**: Write your **conclusion**
  - If you **reject H_0**
    - There is sufficient evidence to suggest that variable X does not follow the binomial distribution $\mathrm{B}(n, p)$
    - Therefore this suggests that the data **does not follow** $\mathrm{B}(n, p)$
  - If you **accept H_0**
    - There is insufficient evidence to suggest that variable X does not follow the binomial distribution $\mathrm{B}(n, p)$
    - Therefore this suggests that the data **follows** $\mathrm{B}(n, p)$

Worked example

A stage in a video game has three boss battles. 1000 people try this stage of the video game and the number of bosses defeated by each player is recorded.

Number of bosses defeated	0	1	2	3
Frequency	490	384	111	15

A χ^2 goodness of fit test at the 5% significance level is used to decide whether the number of bosses defeated can be modelled by a binomial distribution with a 20% probability of success.

a) State the null and alternative hypotheses.

H_0: Number of bosses defeated can be modelled by the binomial distribution $B(3, 0.2)$

H_1: Number of bosses defeated can not be modelled by the binomial distribution $B(3, 0.2)$

b) Assuming the binomial distribution holds, find the expected number of people that would defeat exactly one boss.

Let $X \sim B(3, 0.2)$

Using GDC $P(X = 1) = 0.384$

Expected $1000 \times 0.384 = 384$

$\boxed{\text{Expected frequency of } 1 = 384}$

c) Calculate the p-value for the test.

Find the other expected frequencies

For 0: $1000 \times P(X=0) = 1000 \times 0.512 = 512$

For 2: $1000 \times P(X=2) = 1000 \times 0.096 = 96$

For 3: $1000 \times P(X=3) = 1000 \times 0.008 = 8$

Type two lists into GDC

Observed	490	384	111	15
Expected	512	384	96	8

$\nu = 4 - 1 = 3$

$p = 0.02426\ldots$

$\boxed{p = 0.0243 \text{ (3sf)}}$

d) State the conclusion of the test. Give a reason for your answer.

$0.0243 < 0.05$

Reject H_0 as p-value < significance level
There is sufficient evidence to suggest that the number of bosses defeated can not be modelled by the binomial distribution $B(3, 0.2)$.

Chi-Squared GOF: Normal

What is a chi-squared goodness of fit test for a normal distribution?

- A chi-squared (χ^2) **goodness of fit test** is used to test data from a sample suggesting that the population has a **normal distribution**
  - You will either be **given a precise normal distribution** to test $\mathrm{N}(\mu, \sigma^2)$ with assumed values for μ and σ
  - Or you will be asked to test whether a normal distribution is **suitable without being given assumed values** for μ and/or σ
    - In this case you will have to calculate an **estimate** for the value of μ and/or σ for the normal distribution
    - Either use your GDC or use the formulae
    - $\bar{x} = \dfrac{\sum fx}{\sum f}$ and $s_{n-1}^2 = \dfrac{n}{n-1} s_n^2$

What are the steps for a chi-squared goodness of fit test for a normal distribution?

- **STEP 1**: Write the **hypotheses**
  - H_0 : Variable X can be modelled by a normal distribution
  - H_1 : Variable X cannot be modelled by a normal distribution
    - Make sure you clearly write what the variable is and don't just call it X
    - If you are given the assumed values of μ and σ then state the precise distribution $\mathrm{N}(\mu, \sigma^2)$
- **STEP 2:** Calculate the **expected frequencies**
  - If you were not given the assumed values of μ or σ then you will first have to estimate them
  - Find the probability of the outcome using the normal distribution $\mathrm{P}(a < X < b)$
    - Beware of unbounded inequalities $\mathrm{P}(X < b)$ or $\mathrm{P}(X > a)$ for the class intervals on the 'ends'
  - Multiply the probability by the total frequency $\mathrm{P}(a < X < b) \times N$
  - You will have to combine rows/columns if any expected values are 5 or less
- **STEP 3**: Calculate the **degrees of freedom** for the test
  - For k class intervals (after combining expected values if needed)
  - Degrees of freedom is
    - $\nu = k - 1$ if you were **given** the assumed values for **both** μ and σ
    - $\nu = k - 2$ if you had to **estimate either** μ or σ but **not both**
    - $\nu = k - 3$ if you had to **estimate both** μ and σ
- **STEP 4**: Enter the **frequencies** and the **degrees of freedom** into your GDC

- Enter the observed and expected frequencies as two separate lists
- Your GDC will then give you the χ^2 statistic and its *p*-value
- The χ^2 statistic is denoted as χ^2_{calc}

- **STEP 5**: Decide whether there is **evidence** to **reject the null** hypothesis
  - EITHER compare the **χ^2 statistic** with the given **critical value**
    - If χ^2 statistic **>** critical value then **reject H_0**
    - If χ^2 statistic **<** critical value then **accept H_0**
  - OR compare the ***p*-value** with the given **significance level**
    - If *p*-value **<** significance level then **reject H_0**
    - If *p*-value **>** significance level then **accept H_0**
- **STEP 6**: Write your **conclusion**
  - If you **reject H_0**
    - There is sufficient evidence to suggest that variable *X* does not follow the normal distribution $\mathrm{N}(\mu, \sigma^2)$
    - Therefore this suggests that the data **does not follow** $\mathrm{N}(\mu, \sigma^2)$
  - If you **accept H_0**
    - There is insufficient evidence to suggest that variable *X* does not follow the normal distribution $\mathrm{N}(\mu, \sigma^2)$
    - Therefore this suggests that the data **follows** $\mathrm{N}(\mu, \sigma^2)$

Worked example

300 marbled ducks in Quacktown are weighed and the results are shown in the table below.

Mass (g)	Frequency
$m < 450$	1
$450 \leq m < 470$	9
$470 \leq m < 520$	158
$520 \leq m < 570$	123
$m \geq 570$	9

A χ^2 goodness of fit test at the 10% significance level is used to decide whether the mass of a marbled duck can be modelled by a normal distribution with mean 520 g and standard deviation 30 g.

a) Explain why it is necessary to combine the groups $m < 450$ and $450 \leq m < 470$ to create the group $m < 470$ with frequency 10.

Combine categories if expected frequencies are 5 or less

$300 \times P(X < 450 \mid X \sim N(520, 30^2)) = 300 \times 0.00981... = 2.944...$

The expected frequency is less than 5 so combine with the next category.

b) Calculate the expected frequencies, giving your answers correct to 2 decimal places.

Let $X \sim N(520, 30^2)$

300 × probability

Mass (g)	Probability	Expected frequency
$m < 470$	0.047790...	14.34
$470 \leq m < 520$	0.452209...	135.66
$520 \leq m < 570$	0.452209...	135.66
$m \geq 570$	0.047790...	14.34

c) Write down the null and alternative hypotheses.

H_0: Mass of the marbled ducks can be modelled by the normal distribution $N(520, 30^2)$

H_1: Mass of the marbled ducks can not be modelled by the normal distribution $N(520, 30^2)$

d) Calculate the χ^2 statistic.

Enter the observed and expected frequencies into GDC $\nu = 4 - 1 = 3$

χ^2 statistic $= 8.162...$

$\boxed{\chi^2_{calc} = 8.16 \text{ (3sf)}}$

e) Given that the critical value is 6.251, state the conclusion of the test. Give a reason for your answer.

$8.16 > 6.251$

Reject H_0 as χ^2 statistic > critical value. There is sufficient evidence to suggest that the mass of the marbled ducks can not be modelled by the normal distribution $N(520, 30^2)$.

Chi-squared GOF: Poisson

What is a chi-squared goodness of fit test for a Poisson distribution?

- A chi-squared (χ^2) goodness of fit test is used to test data from a sample suggesting that the population has a Poisson distribution
  - You will either be **given a precise Poisson distribution** to test $\mathrm{Po}(m)$ with an assumed value for *m*
  - Or you will be asked to test whether a Poisson distribution is **suitable without being given an assumed value** for *m*
    - In this case you will have to calculate an **estimate** for the value of *m* for the Poisson distribution
    - To calculate it just calculate the mean
    - $m = \dfrac{\sum fx}{\sum f}$

What are the steps for a chi-squared goodness of fit test for a Poisson distribution?

- **STEP 1**: Write the **hypotheses**
  - H_0 : Variable *X* can be modelled by a Poisson distribution
  - H_1 : Variable *X* cannot be modelled by a Poisson distribution
    - Make sure you clearly write what the variable is and don't just call it *X*
    - If you are given the assumed value of *m* then state the precise distribution $\mathrm{Po}(m)$
- **STEP 2**: Calculate the **expected frequencies**
  - If you were not given the assumed value of *m* then you will first have to **estimate it** using the **observed data**
  - Find the probability of the outcome using the Poisson distribution $\mathrm{P}(X = x)$
  - Multiply the probability by the total frequency $\mathrm{P}(X = x) \times N$
    - If *a* is the smallest observed value then calculate $\mathrm{P}(X \le a)$
    - If *b* is the largest observed value then calculate $\mathrm{P}(X \ge b)$
  - You will have to combine rows/columns if any expected values are 5 or less
- **STEP 3**: Calculate the **degrees of freedom** for the test
  - For *k* outcomes (after combining expected values if needed)
  - Degree of freedom is
    - $v = k - 1$ if you were **given** the assumed value of *m*
    - $v = k - 2$ if you had to **estimate** the value of *m*
- **STEP 4**: Enter the **frequencies** and the **degree of freedom** into your GDC
  - Enter the observed and expected frequencies as two separate lists
  - Your GDC will then give you the χ^2 statistic and its *p*-value

  - The χ^2 statistic is denoted as χ^2_{calc}
- **STEP 5**: Decide whether there is **evidence** to **reject the null** hypothesis
  - EITHER compare the χ^2 **statistic** with the given **critical value**
    - If χ^2 statistic **>** critical value then **reject H_0**
    - If χ^2 statistic **<** critical value then **accept H_0**
  - OR compare the p-value with the given significance level
    - If p-value **<** significance level then **reject H_0**
    - If p-value **>** significance level then **accept H_0**
- **STEP 6**: Write your **conclusion**
  - If you **reject H_0**
    - There is sufficient evidence to suggest that variable X does not follow the Poisson distribution $\text{Po}(m)$
    - Therefore this suggests that the data **does not follow** $\text{Po}(m)$
  - If you **accept H_0**
    - There is insufficient evidence to suggest that variable X does not follow the Poisson distribution $\text{Po}(m)$
    - Therefore this suggests that the data **follows** $\text{Po}(m)$

Worked example

A parent claims the number of messages they receive from their teenage child within an hour can be modelled by a Poisson distribution. The parent collects data from 100 one hour periods and records the observed frequencies of the messages received from the child. The parent calculates the mean number of messages received from the sample and uses this to calculate the expected frequencies if a Poisson model is used.

Number of messages	Observed frequency	Expected frequency
0	9	7.28
1	16	a
2	23	24.99
3	22	21.82
4	16	14.29
5	14	7.49
6 or more	0	b

A χ^2 goodness of fit test at the 10% significance level is used to test the parent's claim.

a) Write down null and alternative hypotheses to test the parent's claim.

We are not given a specific Poisson distribution

H_0: Number of messages received can be modelled by a Poisson distribution

H_1: Number of messages received can not be modelled by a Poisson distribution

b) Show that the mean number of messages received per hour for the sample is 2.62.

$$\boxed{m = \frac{\Sigma fx}{\Sigma f} = \frac{0\times9 + 1\times16 + 2\times23 + 3\times22 + 4\times16 + 5\times14}{9 + 16 + 23 + 22 + 16 + 14} = 2.62}$$

c) Calculate the values of a and b, giving your answers to 2 decimal places.

Let $X \sim Po(2.62)$

$a = 100 \times P(X=1) = 100 \times 0.19074... = 19.074$ $\boxed{a = 19.07 \ (2dp)}$

$b = 100 \times P(X \geq 6) = 100 \times 0.05052... = 5.05$ $\boxed{b = 5.05 \ (2dp)}$

d) Perform the hypothesis test.

Calculate degree of freedom $\nu = k-2$ m was estimated

$\nu = 7-2 = 5$

Enter observed and expected frequencies in GDC

$p = 0.03515... < 0.1$

Reject H_0 as p-value < significance level.
There is sufficient evidence to suggest that a Poisson distribution can not model the number of messages received.