Malaria Molecular Surveillance Study Design Workshop

Module 4: Statistical power

Remember …

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False		

α sets the **false positive rate** of a test. Using α we can control how often we incorrectly conclude that there is a real effect when there is none.

Power

In power analysis, we also specify an **alternative hypothesis**

H_0: The population prevalence equals p_0

H_1: The population prevalence equals p, which is different from p_0

Power

In power analysis, we also specify an **alternative hypothesis**

H_0: The population prevalence equals p_0

H_1: The population prevalence equals p, which is different from p_0

For example…

I want to test if the prevalence of *pfcrt* K76T mutations is significantly different from 10%. When powering this test, I assume the true prevalence of K76T mutations is 15%.

Power

What is the distribution of my test statistic under the alternative hypothesis?

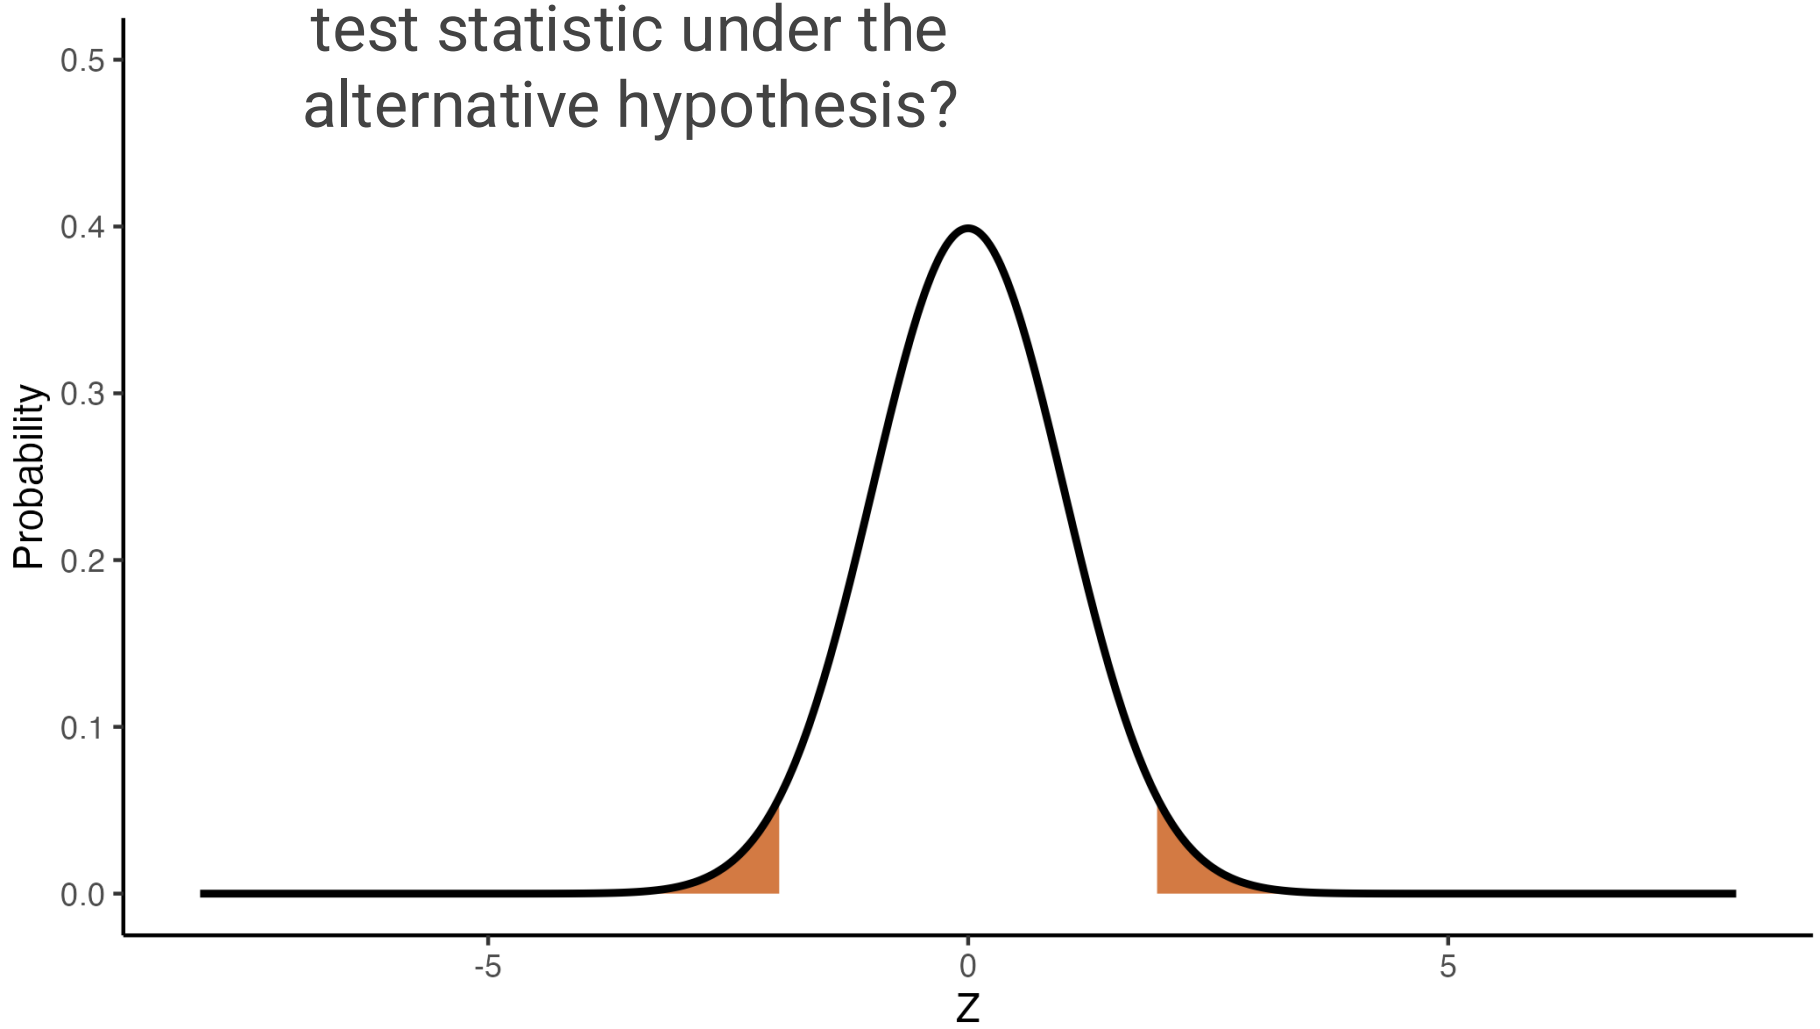

Power

What is the distribution of my test statistic under the alternative hypothesis?

$$Z = \frac{\hat{p} - p_0}{\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}}$$

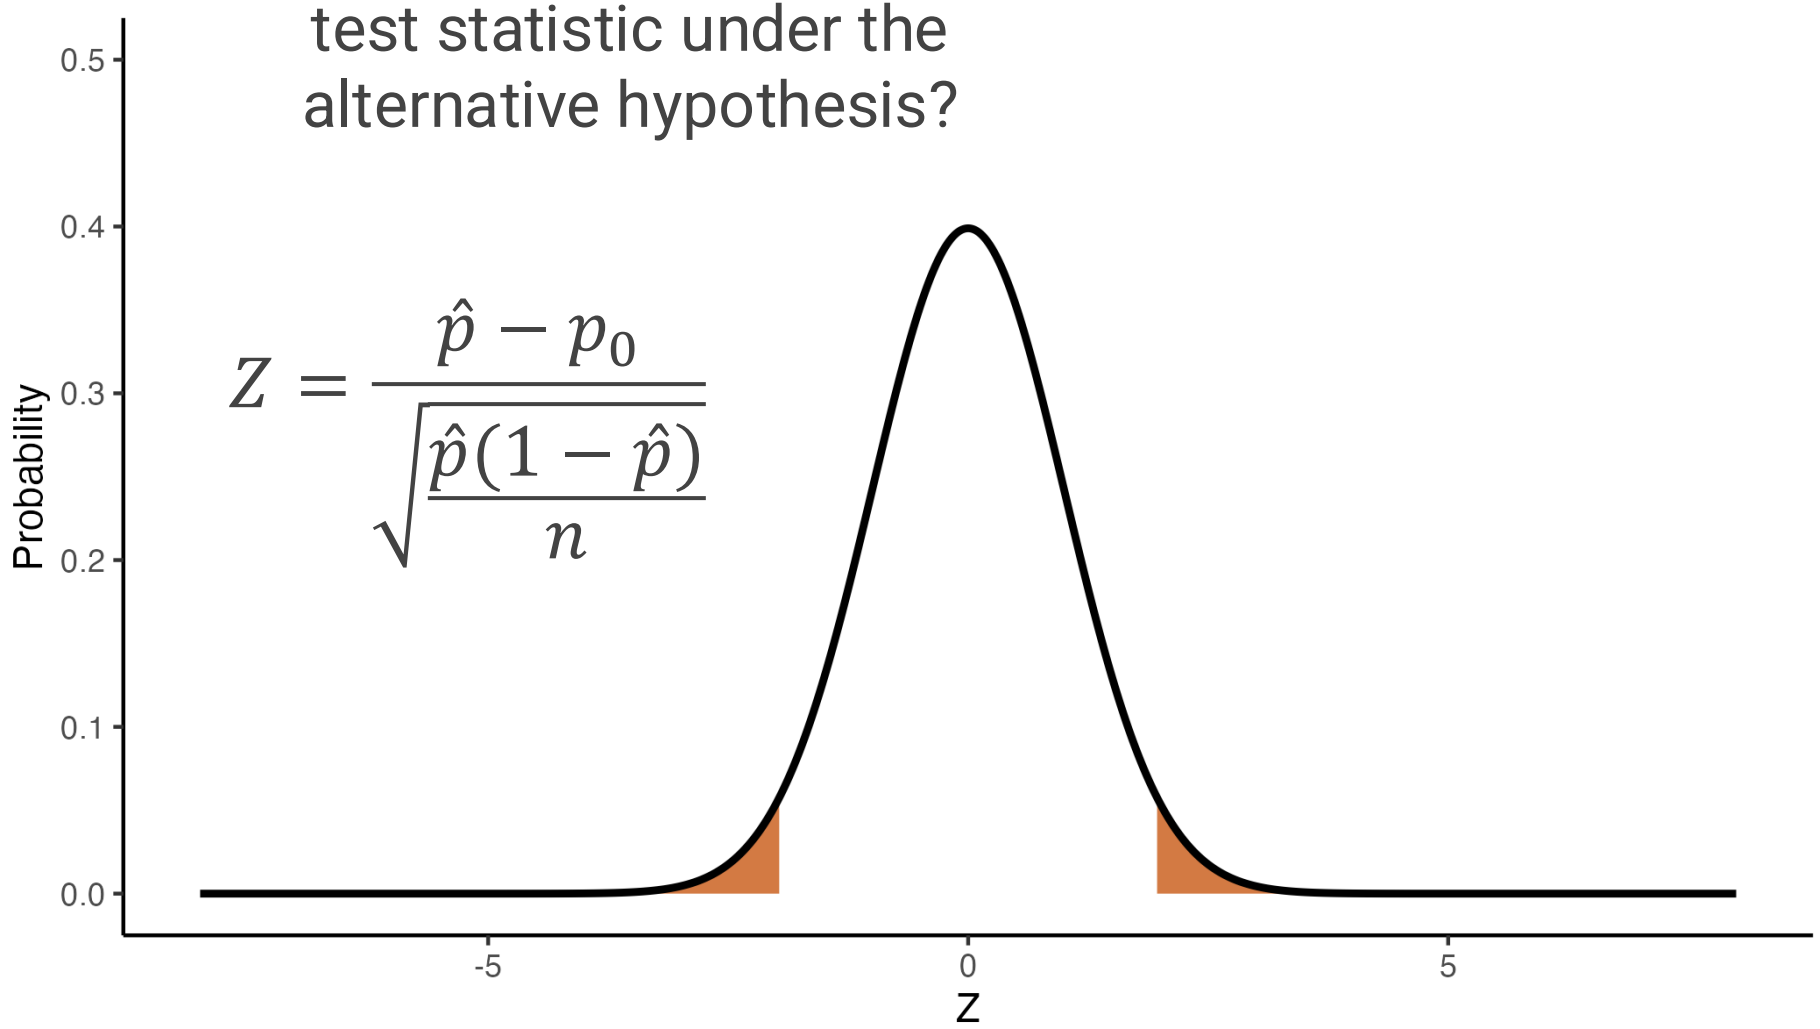

Power

What is the distribution of my test statistic under the alternative hypothesis?

$$Z = \frac{\hat{p} - p_0}{\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}}$$

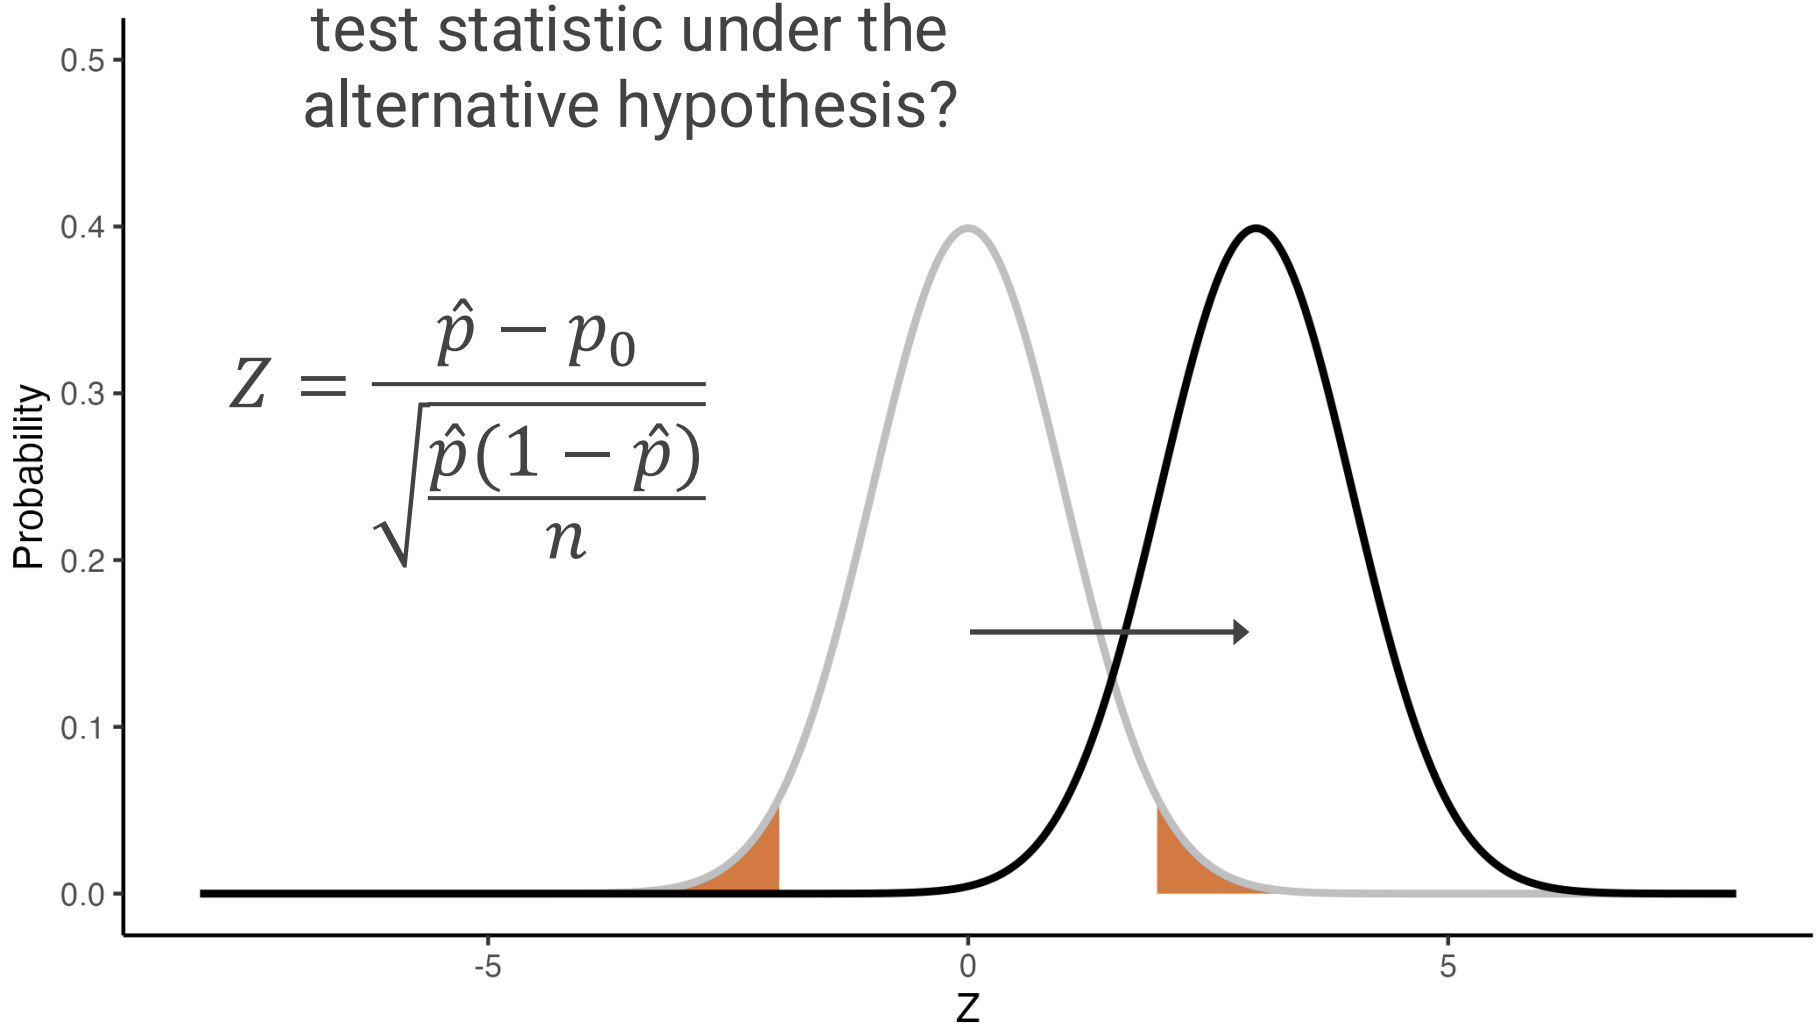

Power

Area under curve gives the probability of rejecting H_0 if H_1 is true

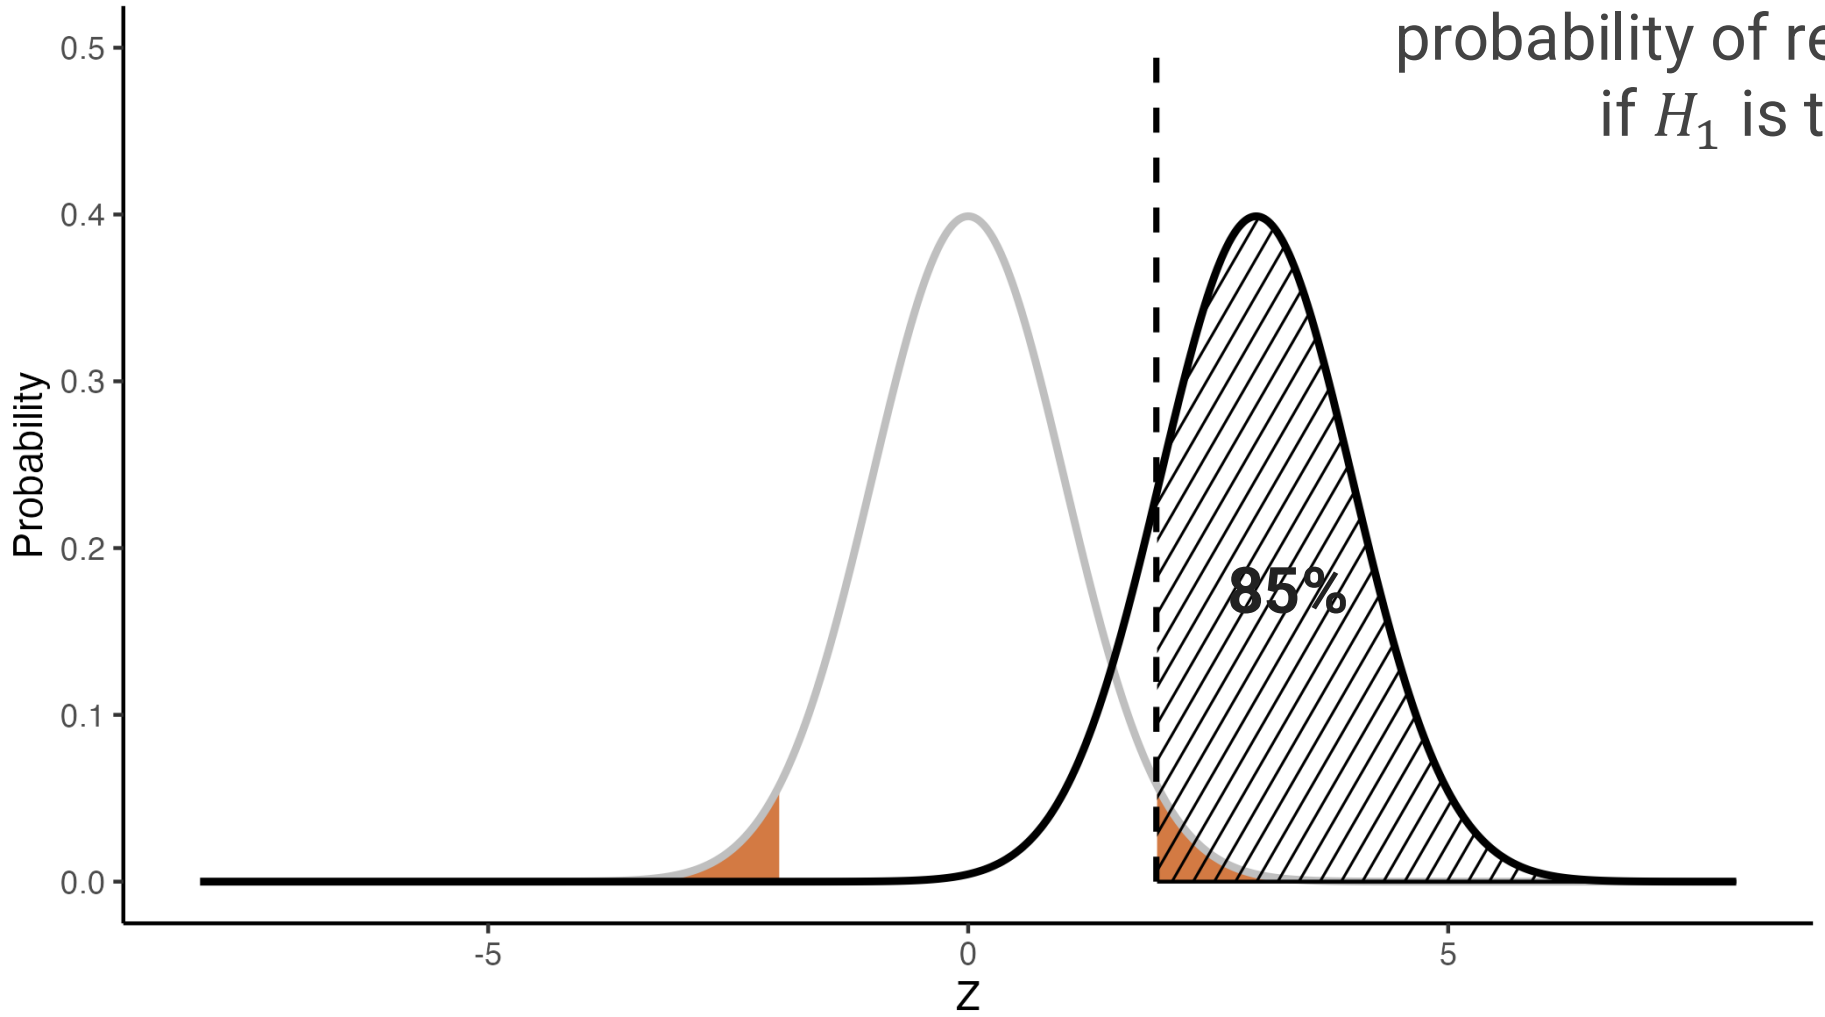

Power

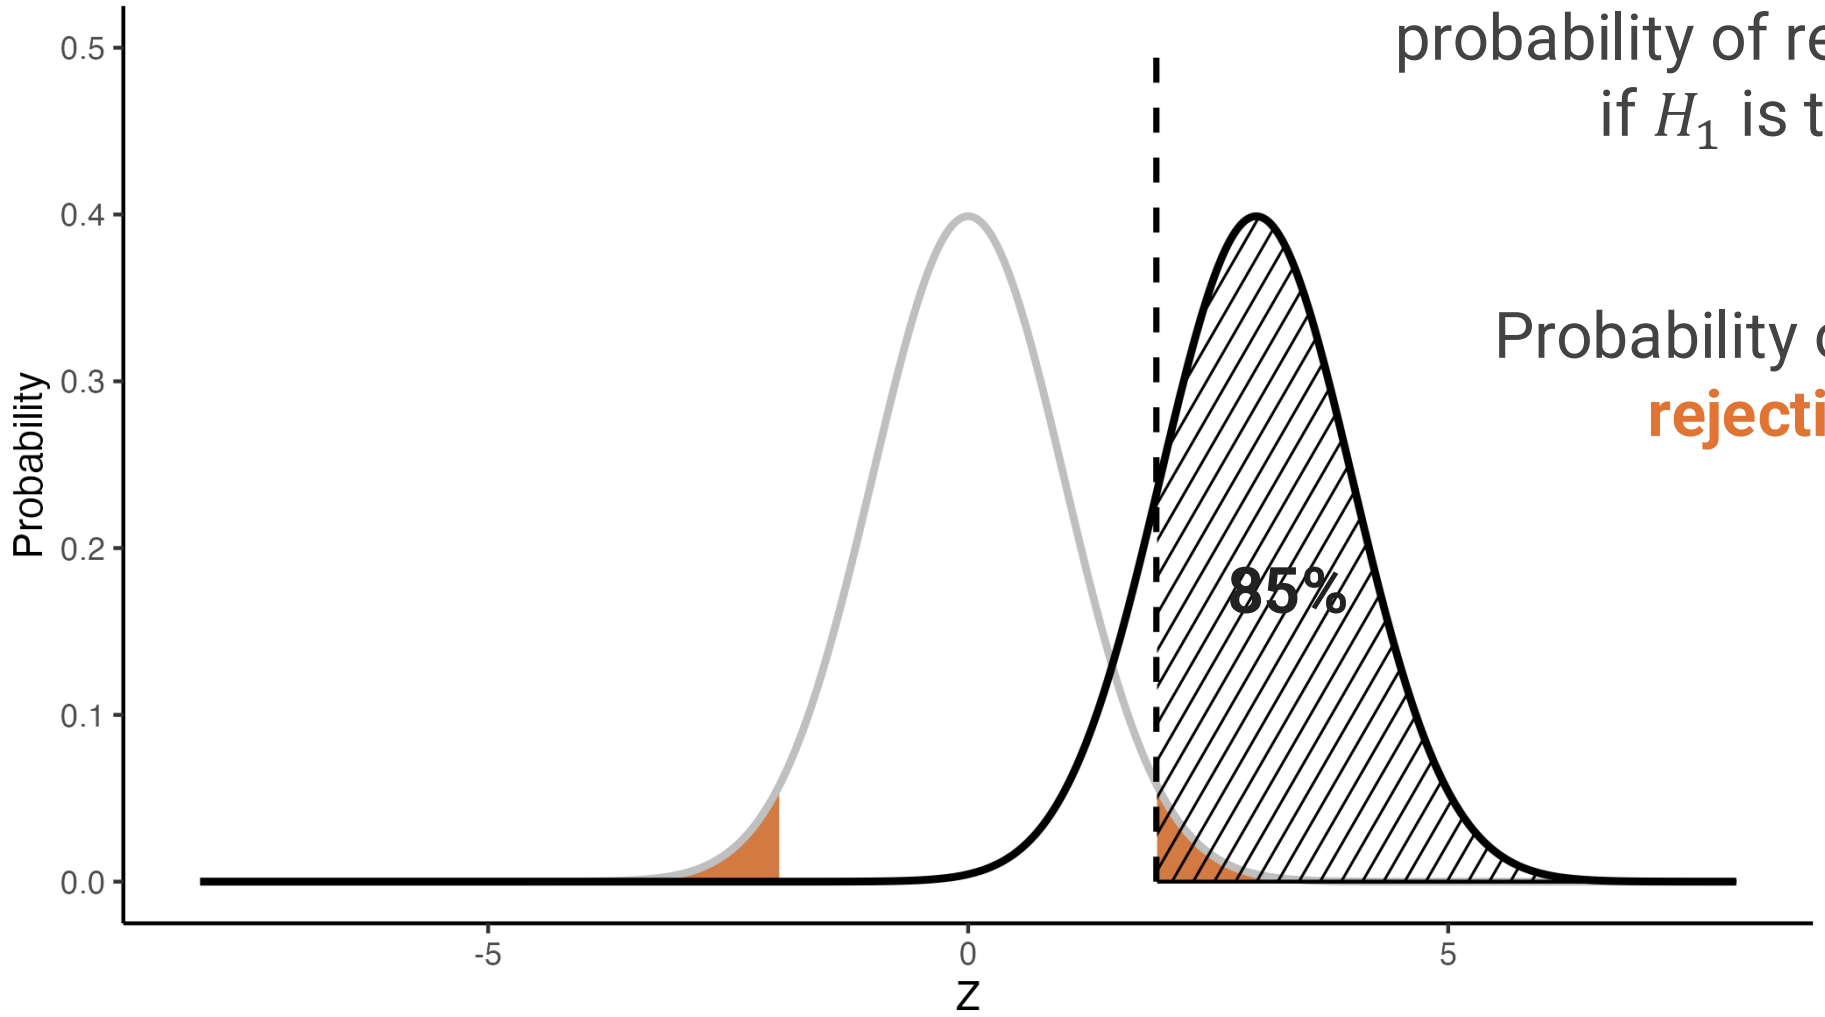

Area under curve gives the probability of rejecting H_0 if H_1 is true

Probability of **correctly rejecting** H_0

Power

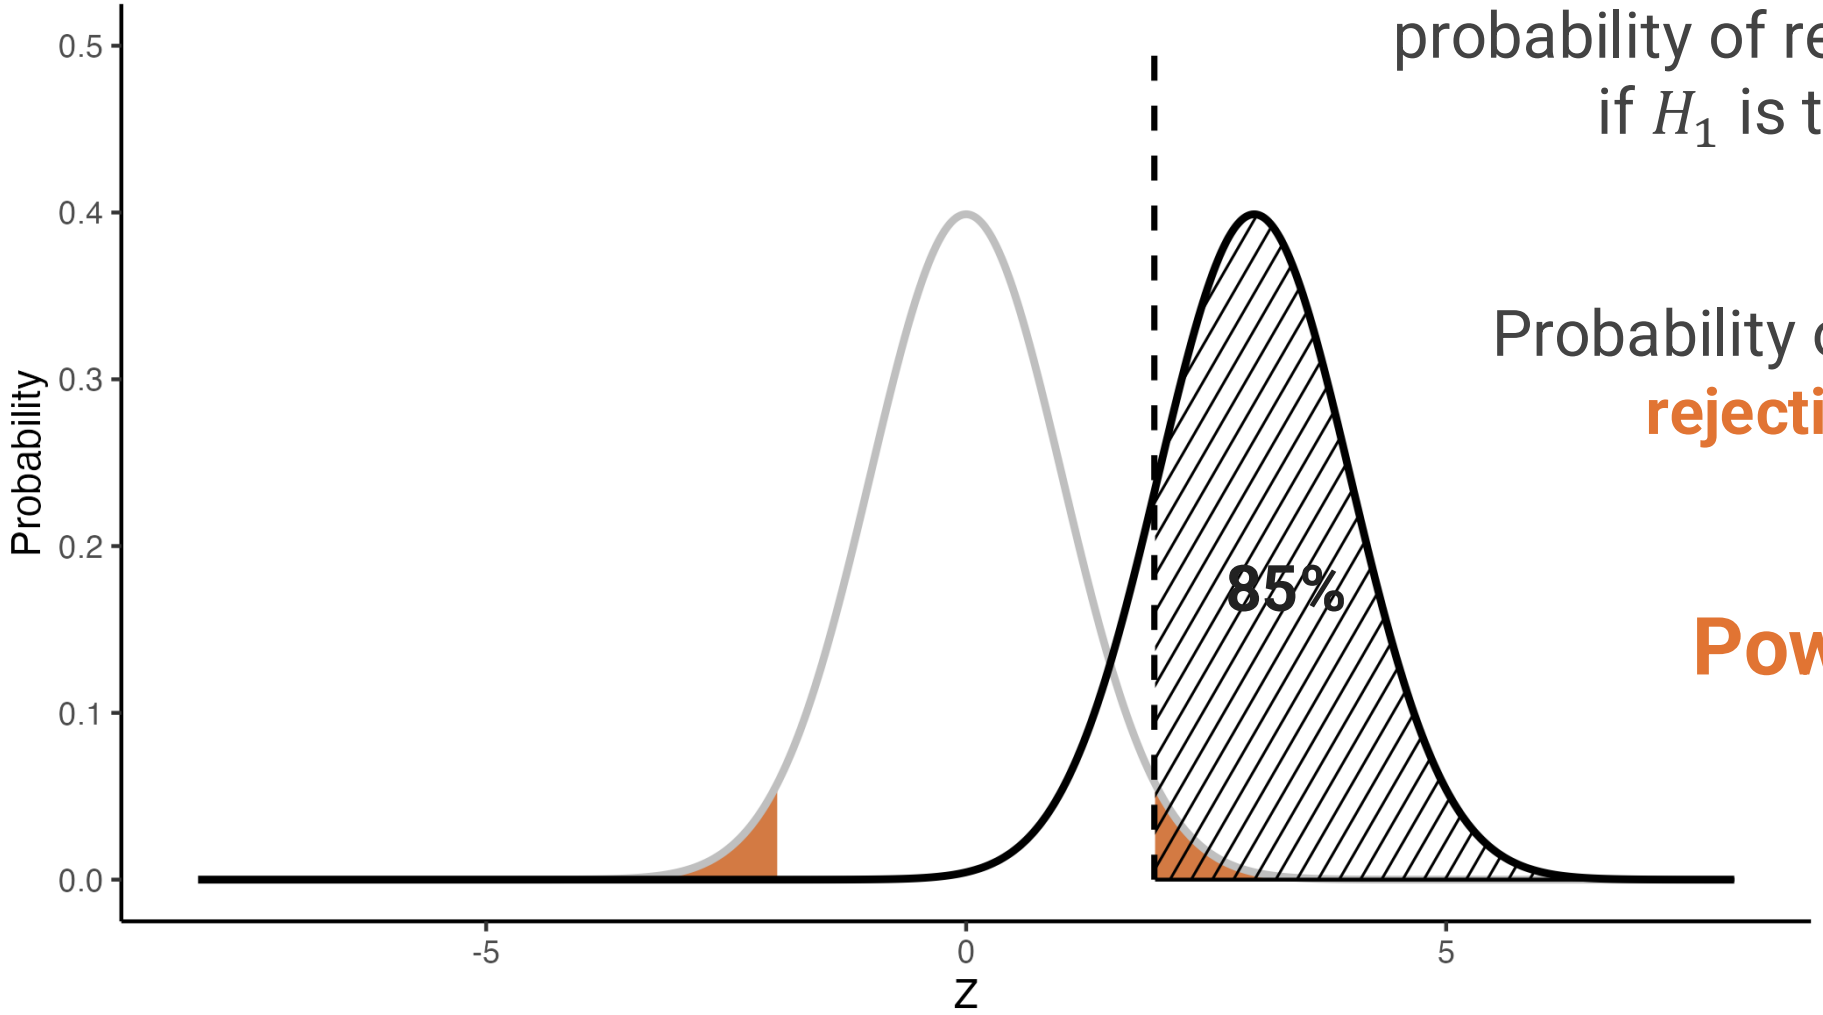

Area under curve gives the probability of rejecting H_0 if H_1 is true

Probability of **correctly rejecting** H_0

Power!

Power

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False		

Power

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False	False negative	True positive

Power

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False	False negative $1 - \text{Power}$	True positive Power

Power

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False	False negative $1 - \text{Power}$	True positive Power

Power is the probability of **correctly rejecting** the null hypothesis. It is the chance that we find something interesting, given that it is there.

Power

		Conclusion about H_0	
		Fail to reject	**Reject**
Truth about H_0	**True**	True negative $1 - \alpha$	False positive α
	False	False negative β	True positive $1 - \beta$

Power is the probability of **correctly rejecting** the null hypothesis. It is the chance that we find something interesting, given that it is there.

How do we calculate power?

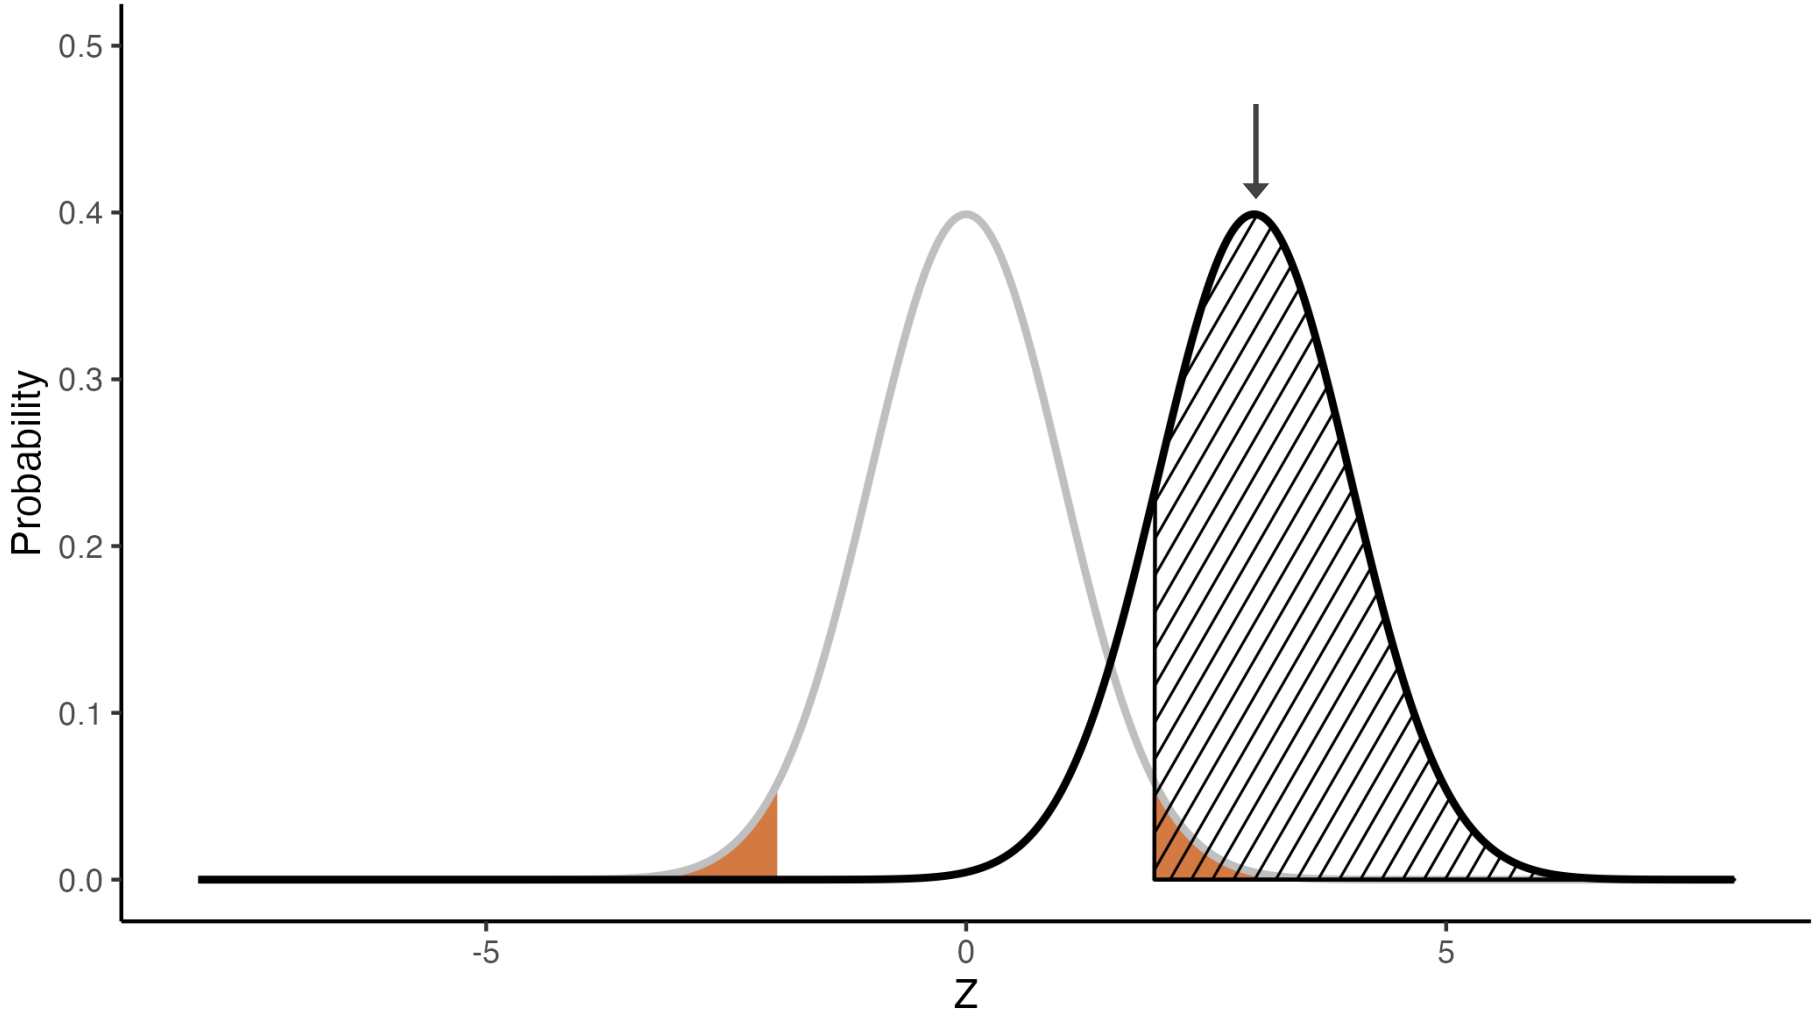

How do we calculate power?

$$E[Z] = \frac{|p - p_0|}{\sqrt{\frac{p\,(1-p)}{n}}}$$

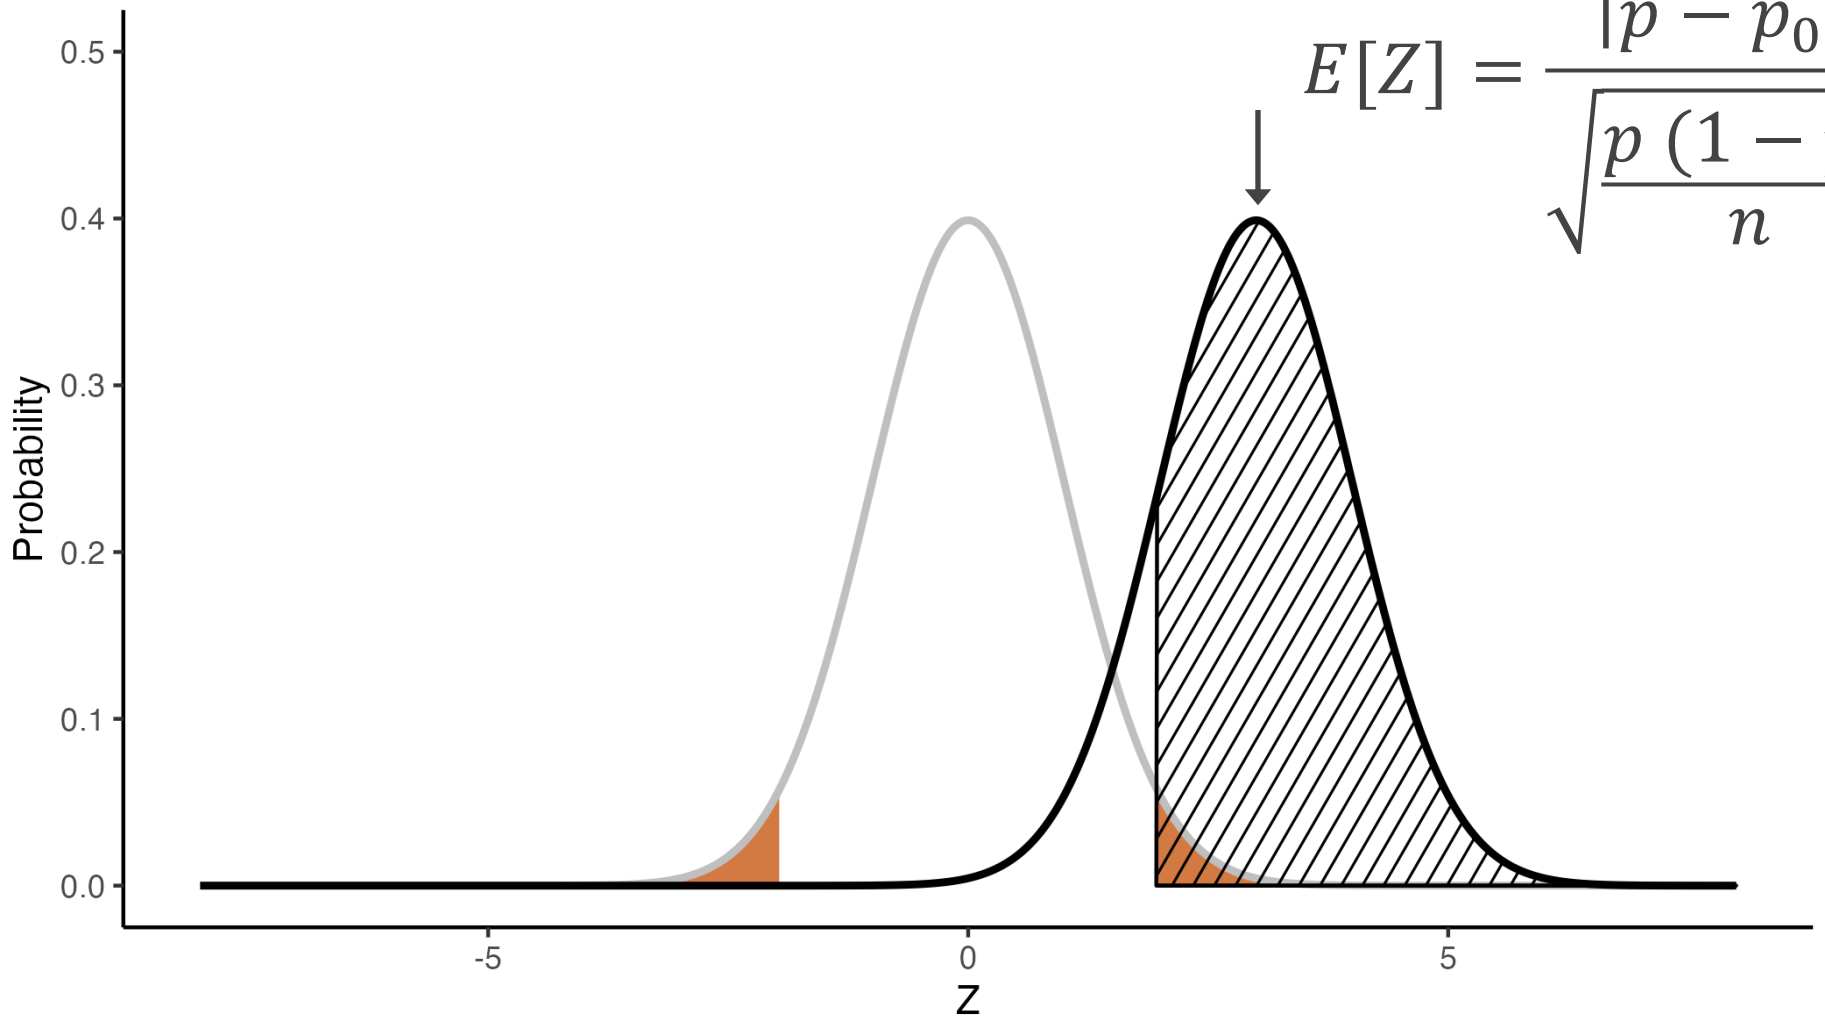

How do we calculate power?

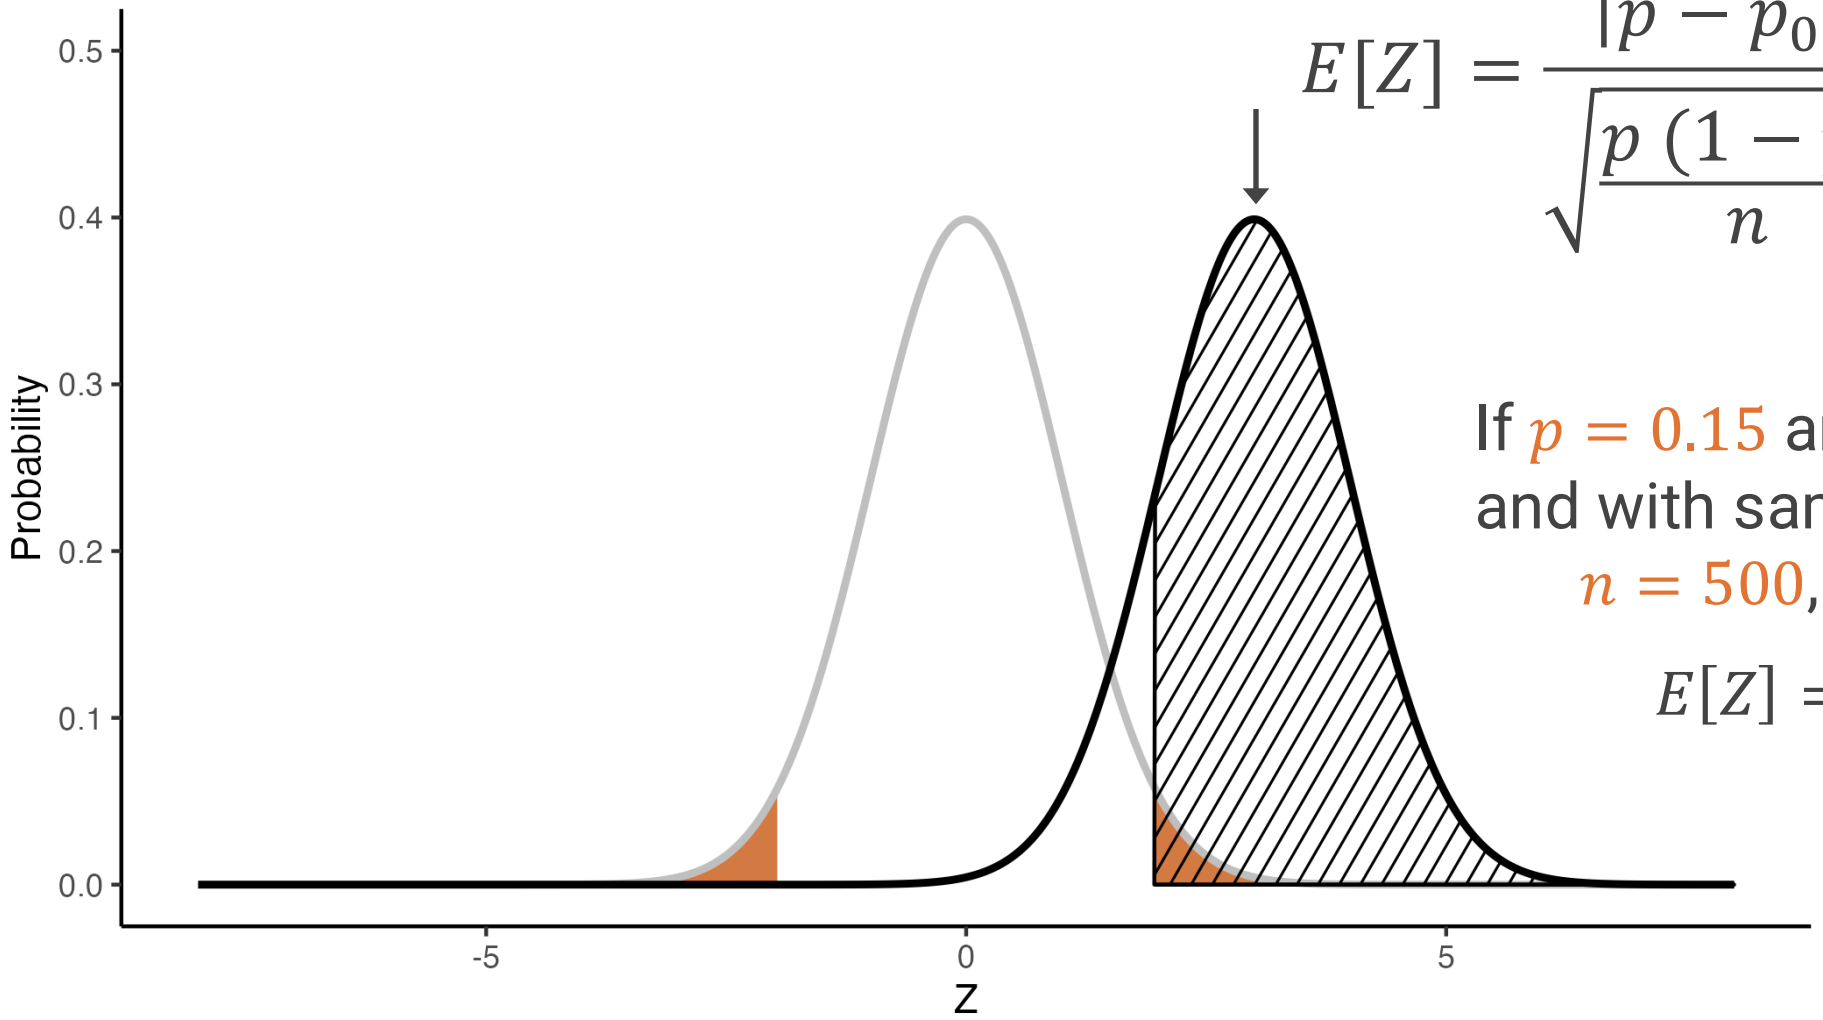

$$E[Z] = \frac{|p - p_0|}{\sqrt{\frac{p\,(1-p)}{n}}}$$

If $p = 0.15$ and $p_0 = 0.1$, and with sample size of $n = 500$, we find

$$E[Z] = 3.13$$

How do we calculate power?

The function $\phi()$ gives us the area under the curve

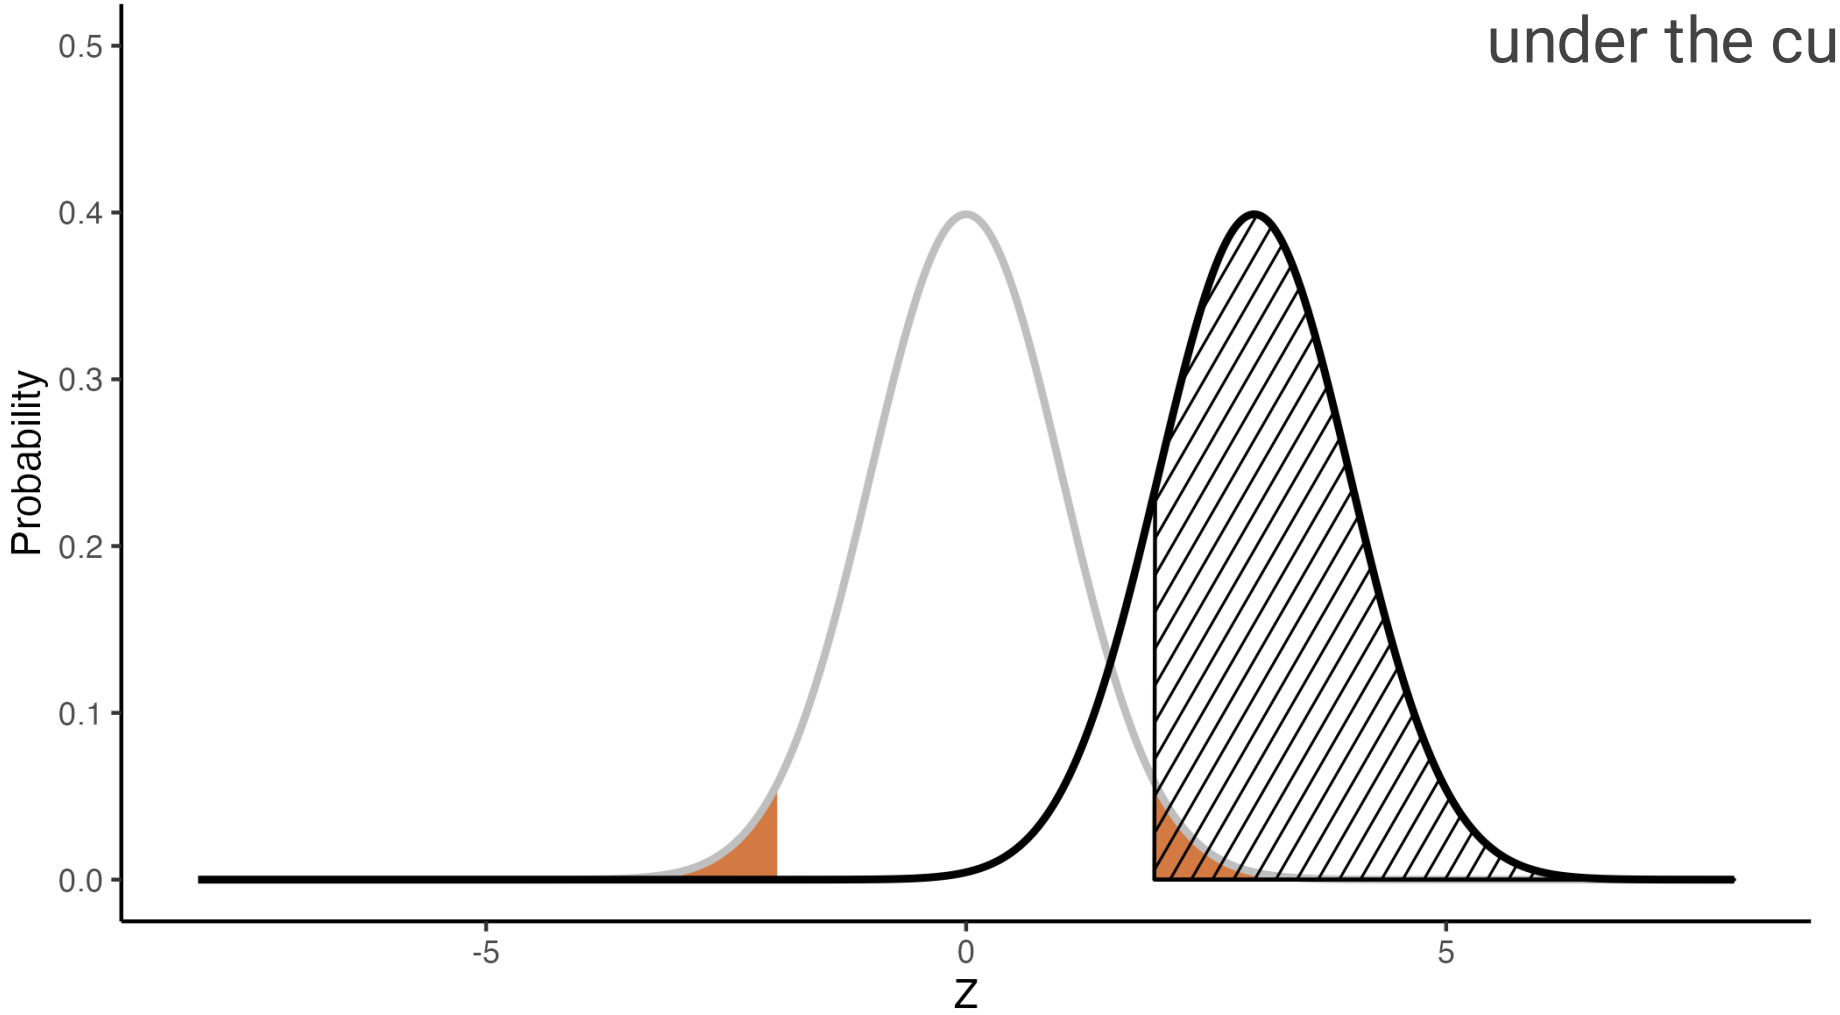

How do we calculate power?

The function $\phi()$ gives us the area under the curve

$$\text{Power} = 1 - \phi(z_{1-\alpha/2} - E[Z])$$

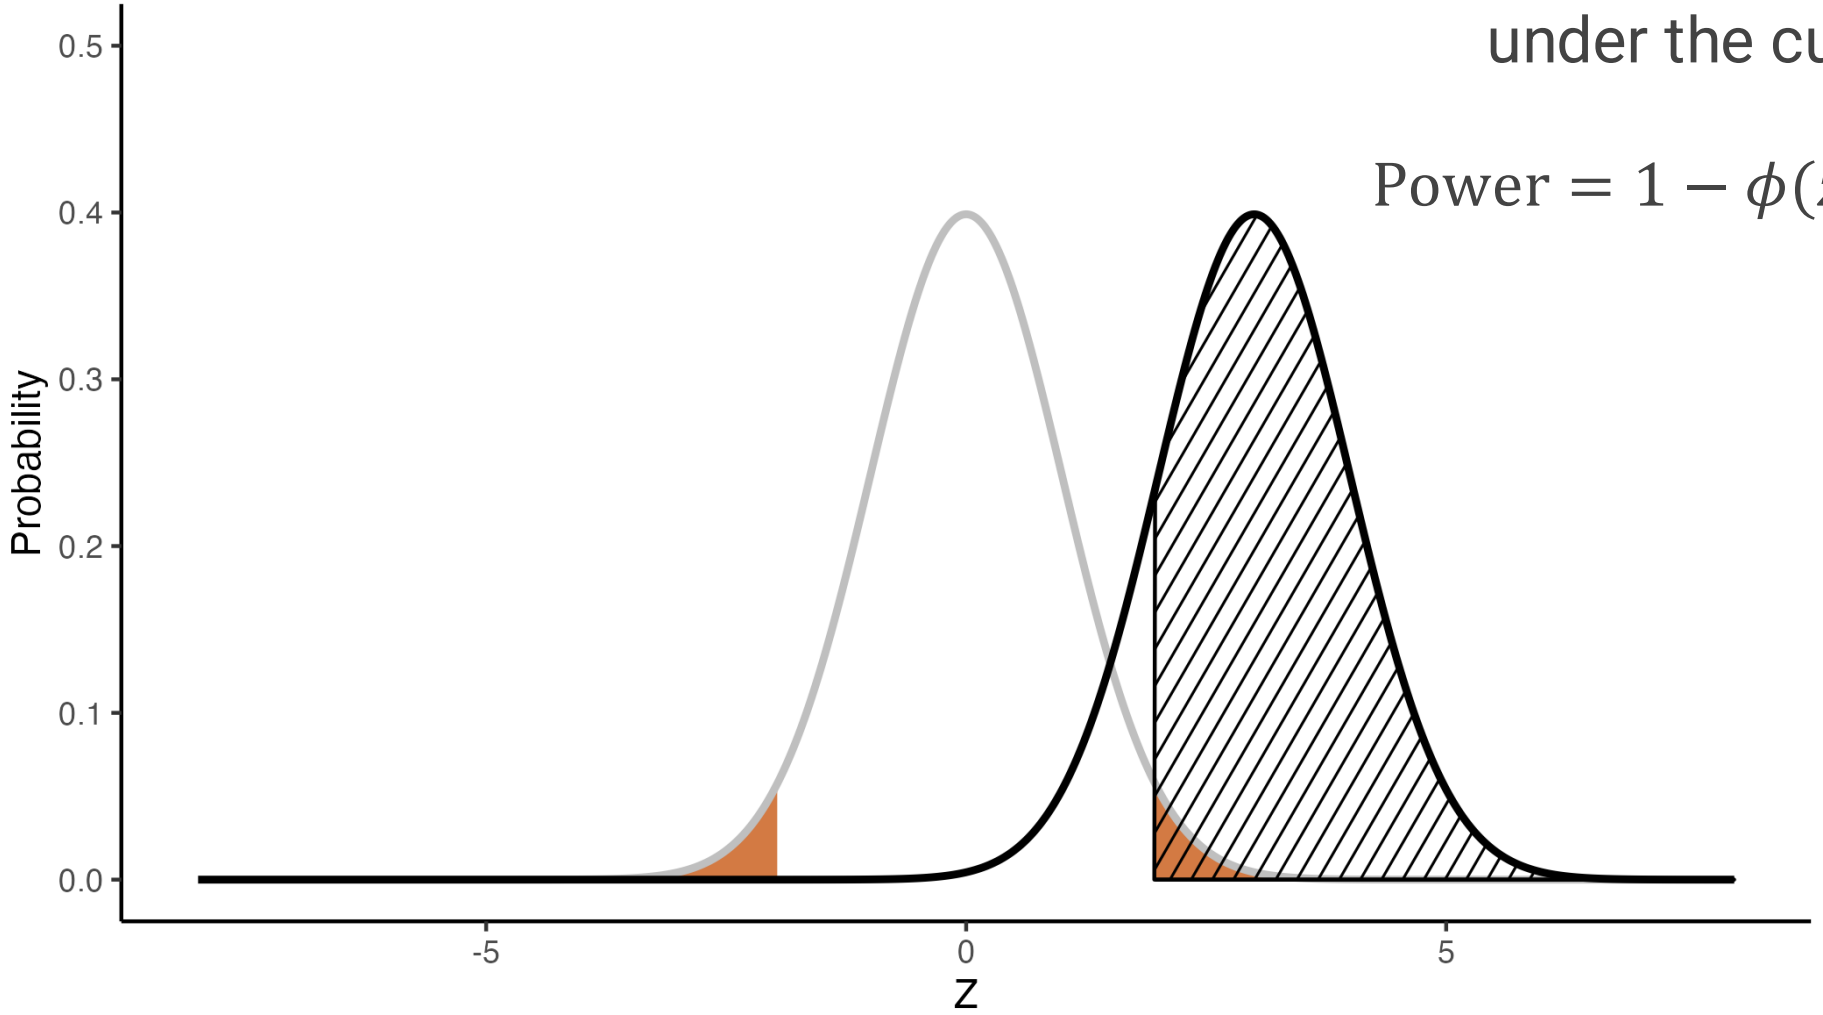

How do we calculate power?

The function $\phi()$ gives us the area under the curve

$$\text{Power} = 1 - \phi(z_{1-\alpha/2} - E[Z])$$

$$\text{Power} = 0.85$$

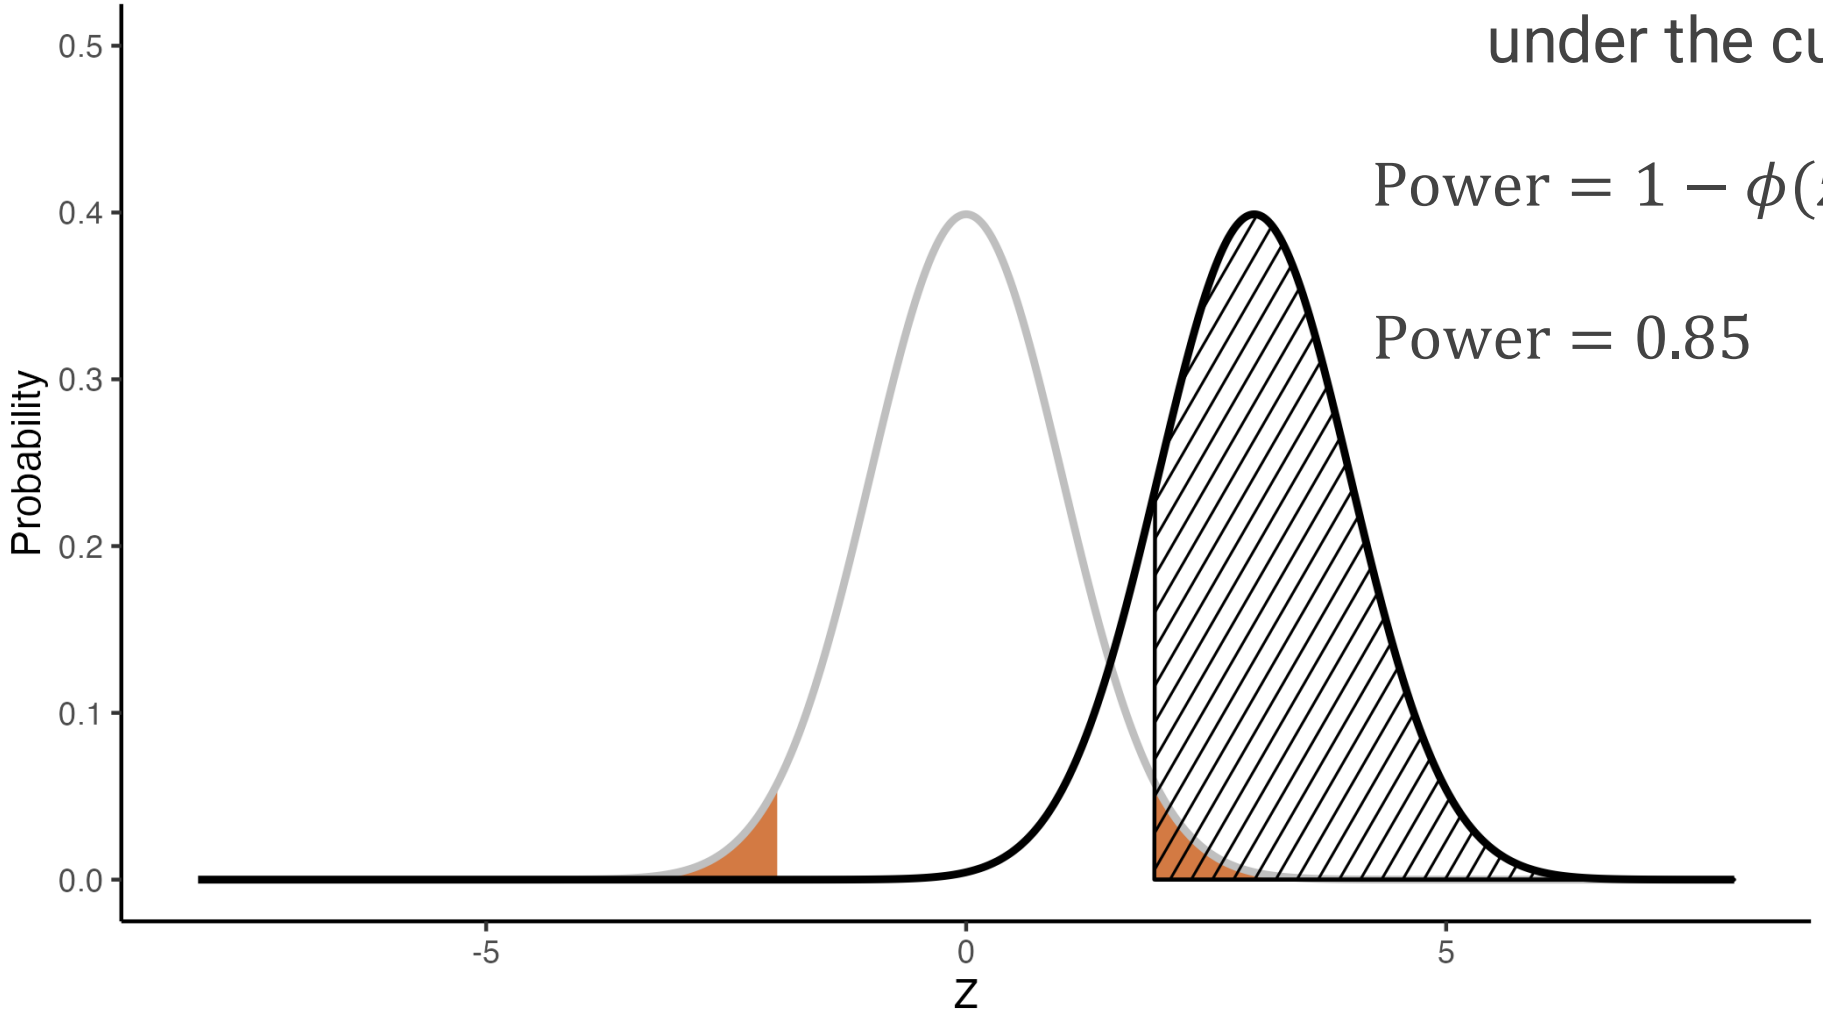

Power as a function of sample size

$$\text{Power} = 1 - \phi(z_{1-\alpha/2} - E[Z])$$

Power as a function of sample size

$$\text{Power} = 1 - \phi\left(z_{1-\alpha/2} - \frac{|p - p_0|}{\sqrt{\frac{p(1-p)}{n}}}\right)$$

Power as a function of sample size

$$\text{Power} = 1 - \phi\left(z_{1-\alpha/2} - \frac{|p - p_0|}{\sqrt{\frac{p(1-p)}{n}}}\right)$$

Power varies as a function of sample size

Power as a function of sample size

$$\text{Power} = 1 - \phi\left(z_{1-\alpha/2} - \frac{|p - p_0|}{\sqrt{\frac{p(1-p)}{n}}}\right)$$

Power varies as a function of sample size

Small sample

$n = 50$

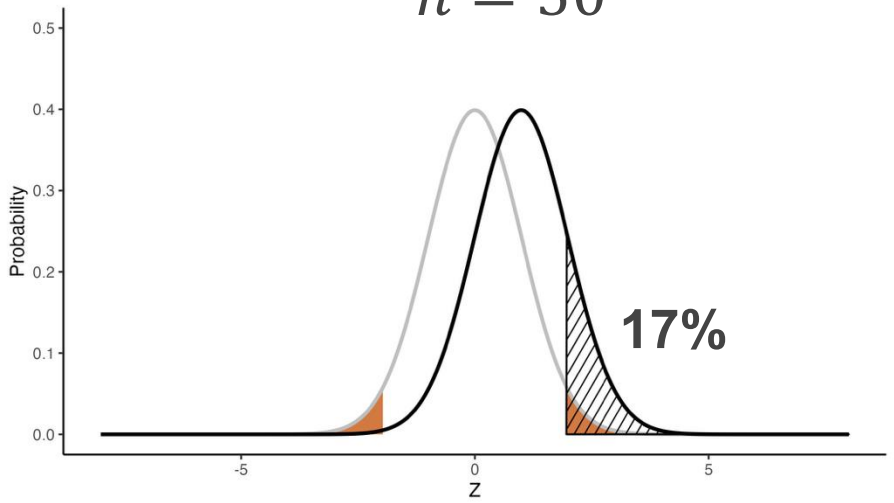

Large sample

$n = 500$

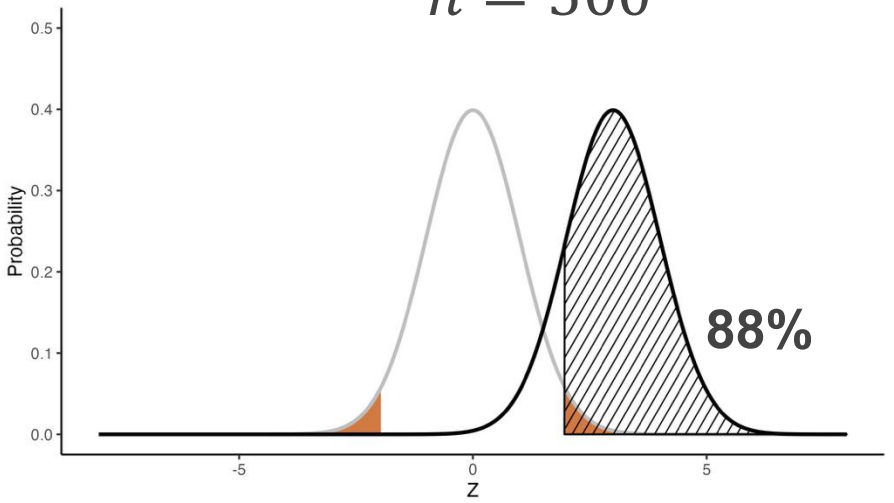

Power curves

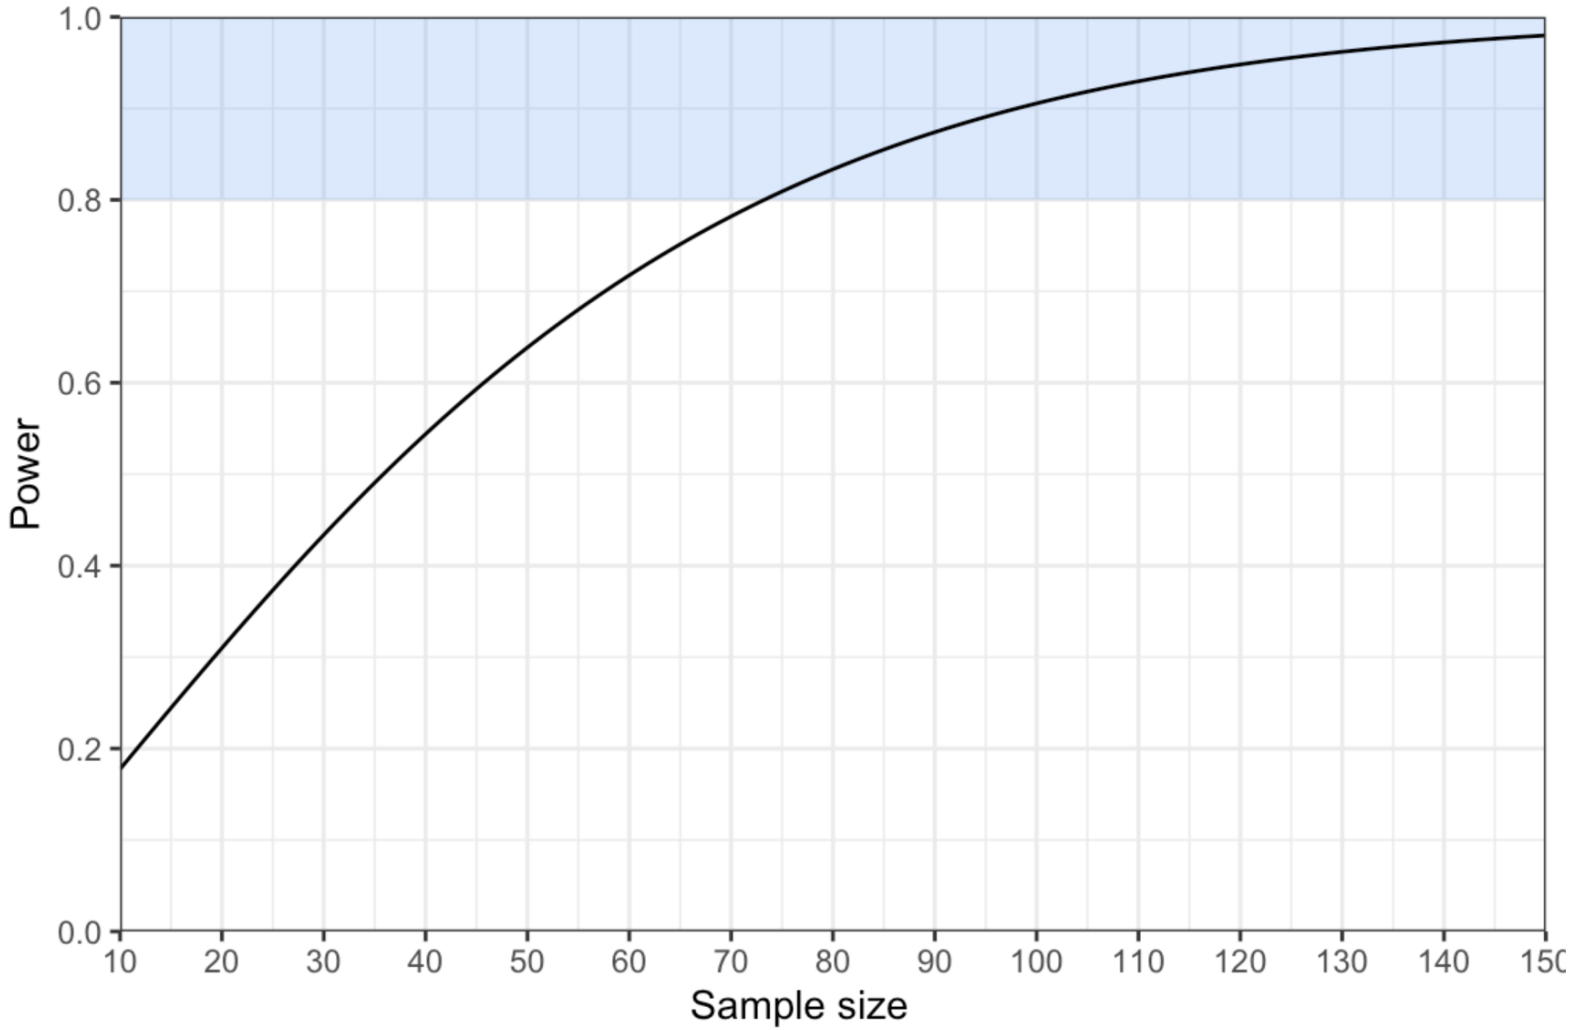

80% power

Power curves

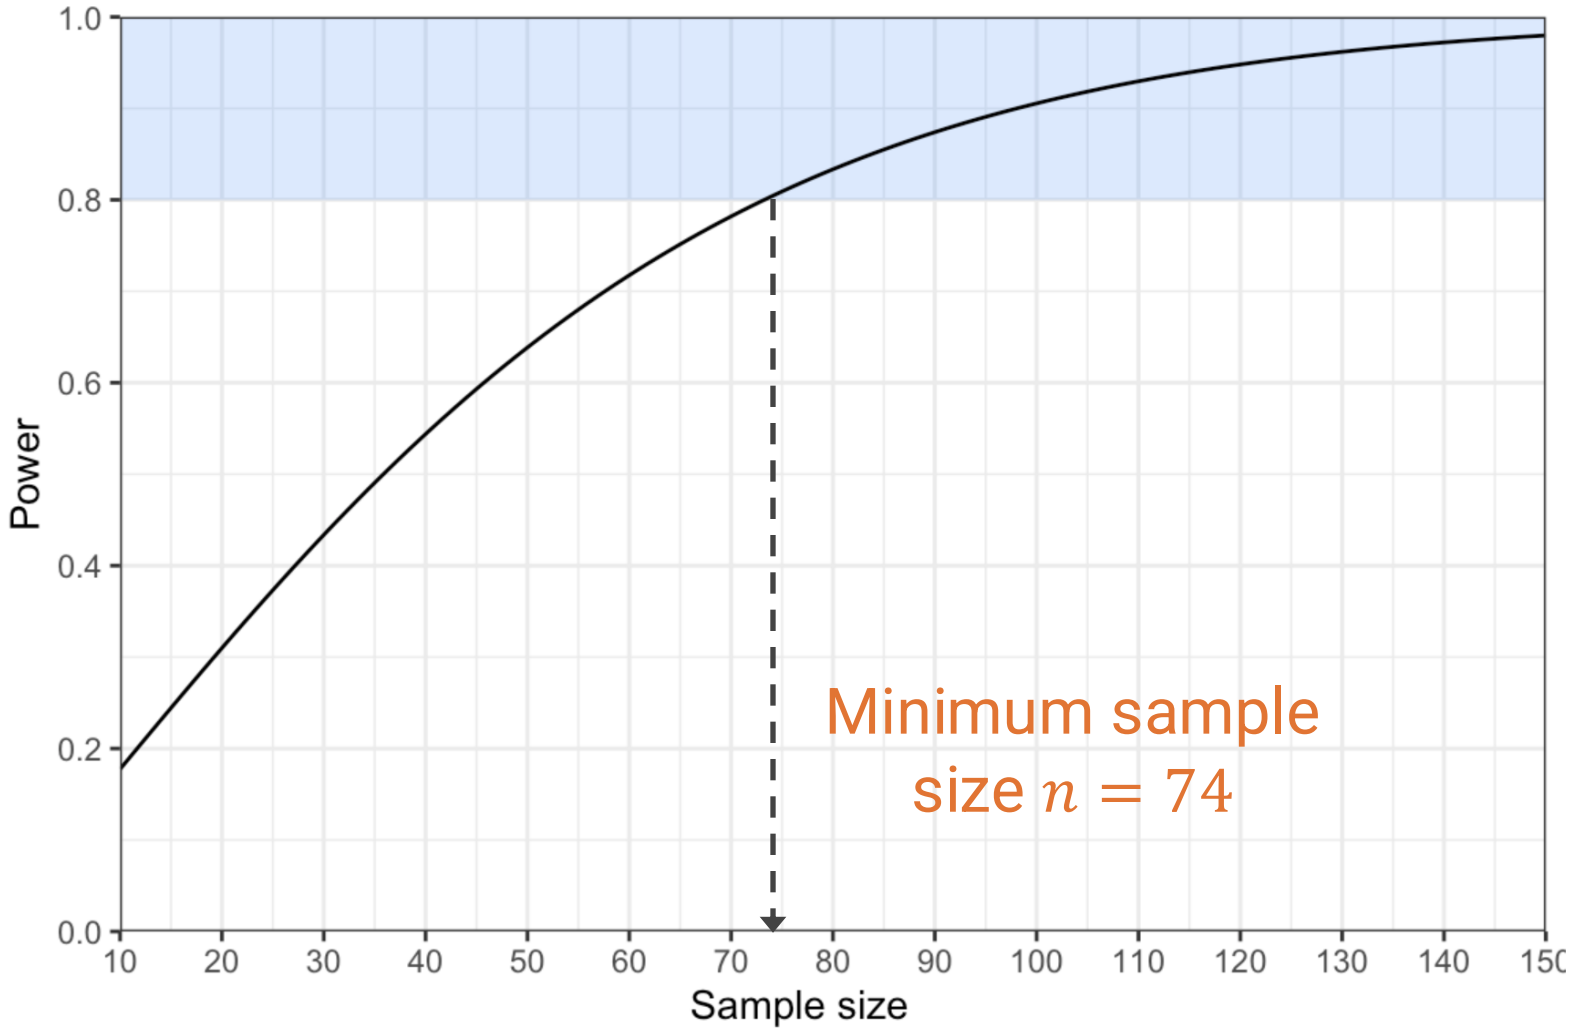

80% power

Sample size formulae

$$\text{Power} = 1 - \phi\left(z_{1-\alpha/2} - \frac{|p - p_0|}{\sqrt{\frac{p(1-p)}{n}}}\right)$$

Can we reverse-engineer this to find the value of n that achieves a target power?

Sample size formulae

$$\text{Power} = 1 - \phi\left(z_{1-\alpha/2} - \frac{|p - p_0|}{\sqrt{\frac{p(1-p)}{n}}}\right)$$

Can we reverse-engineer this to find the value of n that achieves a target power?

$$n = \left(z_{1-\beta} + z_{1-\frac{\alpha}{2}}\right)^2 \frac{p(1-p)}{(p-p_0)^2}$$

For 80% power, we find $z_{1-\beta} = 0.84$

Summary

- **Power** is the true positive rate. It is the chance of **correctly rejecting the null hypothesis**.
- Power increases with **sample size**. We can use power curves or sample size formulae to choose a value of n

Module 4 activity

Format: Interactive R code, accessed through the web

- Test for change in prevalence
- Using power curves and sample size tables
- Test for detection of rare *pfk13* variant

Workshop materials

https://mrc-ide.github.io/MMS-SD_workshop/