ANOVA: Analysis of variance

Hypotheses:

H$_0$**:** $\mu_1 = \mu_2 = \ldots = \mu_n$.

H$_A$**:** at least one mean is different.

Assumptions:

- Each observation is independent of all others.
- The errors are normally distributed.
- The error terms have the same variance σ^2.

Distribution

F-square distribution

Test Statistic:

Sources of Variability	d.f.	Sums of Squares	Mean Square	Test Statistic	P-value
Between	$k-1$	SSB	$MSB = \frac{SSB}{k-1}$	$F = \frac{MSB}{MSW}$	$P(F_{k-1,n-k} > F_{obs})$
Within	$n-k$	SSW	$MSW = \frac{SSW}{n-k}$		
Total	$n-1$	SST			

Statistical Conclusion

If p-value $< \alpha$, reject the null hypothesis (H_0).
If p-value $\geq \alpha$, fail to reject the null hypothesis (H_0).

Clinical Conclusion

At least one group mean is different from the others.

Non-parametric alternatives

- Kruskal-Wallis test by ranks

Correlation Analysis:

Correlation $\neq$ Causation!

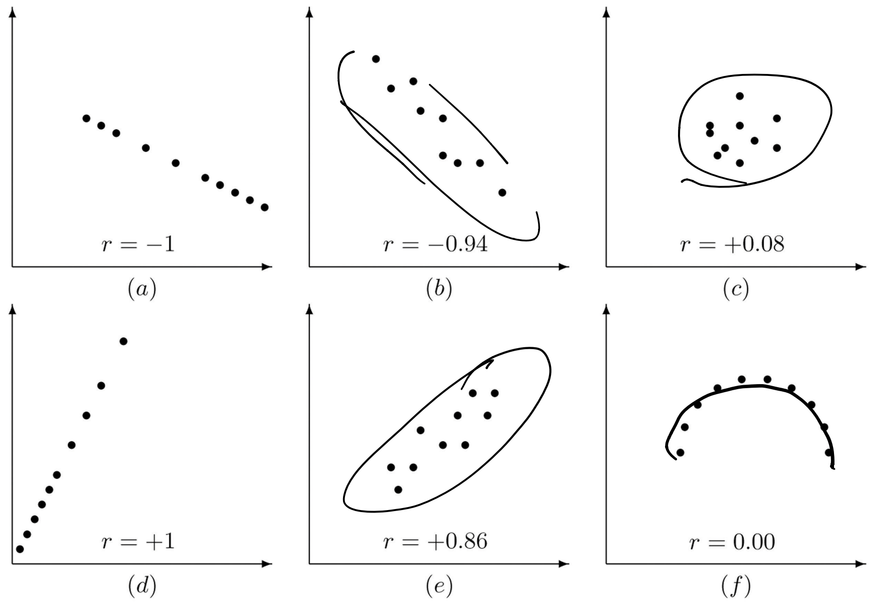

Hypotheses:

H$_0$: $\rho = 0$

H$_A$: $\rho \neq 0$

Assumptions:

- The data must be continuous (interval/ratio).
- There must be a linear relationship between X and Y.
- The data for X and the data for Y must be normally distributed.

Distribution

Student's t-dstribution

$$t = r\sqrt{\frac{n-2}{1-r^2}}$$

Degrees of freedom = $n-2$

Test Statistic:

$$\boxed{r = \frac{cov(X,Y)}{\sigma_x \sigma_Y} = \frac{\sum(x_i-\bar{x})(y_i-\bar{y})}{\sqrt{\sum(x_i-\bar{x})}\sqrt{\sum(y_i-\bar{y})}}}$$

Statistical Conclusion

If p-value $< \alpha$, reject the null hypothesis (H$_0$).

If p-value $\geq \alpha$, fail to reject the null hypothesis (H$_0$).

Clinical Conclusion

There is enough evidence to show X and Y are correlated.

Non-parametric alternatives

- Spearman Correlation

Simple Linear Regression:

Hypotheses:

intercep

slope

$$y_i = \beta_0 + \beta_1 x_i + \varepsilon_i$$

H$_0$**:** $\beta_1 = 0$

H$_A$**:** $\beta_1 \neq 0$

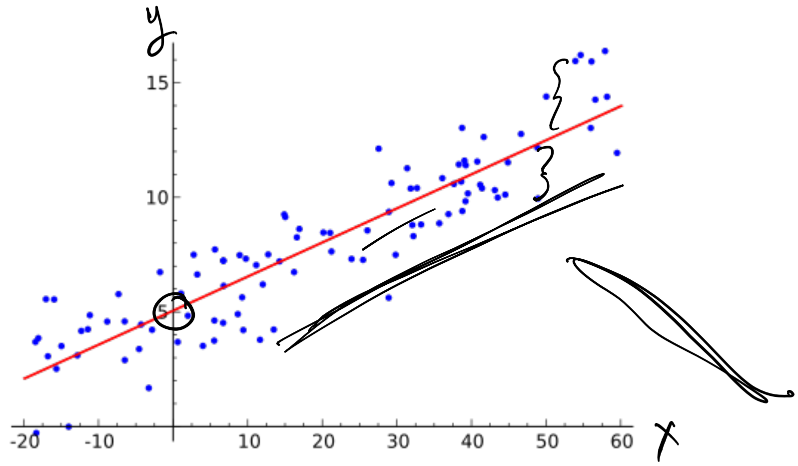

Assumptions:

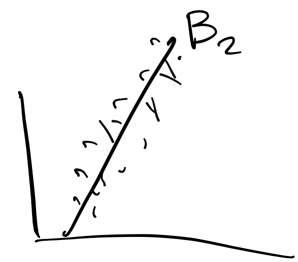

- Predictor variables (x) are fixed variables.
- There is a linear relationship between independent (x) and dependent (y) variables.
- The response data has constant variance (homoscedasticity).
- All errors are independent and uncorrelated.

Distribution

Student's t-dstribution

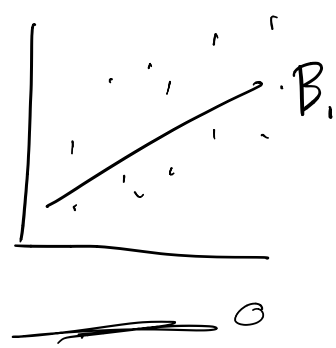

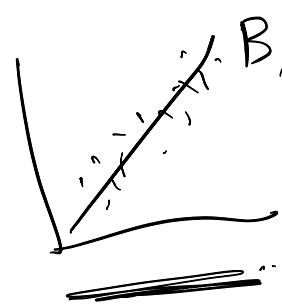

Test Statistic:

$$t = \frac{\hat{\beta}_1}{se_{\beta_1}}$$

Degrees of freedom $= n - 2$

Statistical Conclusion

If p-value $< \alpha$, reject the null hypothesis (H_0).
If p-value $\geq \alpha$, fail to reject the null hypothesis (H_0).

Clinical Conclusion

There is enough evidence to show a relationship between X and Y.

Non-parametric alternatives

- Nonparametric regression

Logistic Regression:

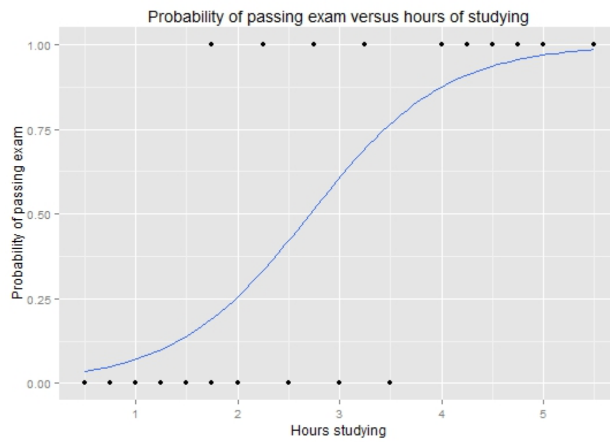

Hypotheses:

$$logit(p_i) = log(\frac{p_i}{1-p_i}) = \boxed{\beta_0 + \beta_1 x_i + \varepsilon_i}$$

$\mathbf{H}_0$: $\beta_1 = 0$

$\mathbf{H}_A$: $\beta_1 \neq 0$

Assumptions:

- Predictor variables (x) are fixed variables.
- Dependent variable (y) is binary.
- The independent variables and log odds are linear.

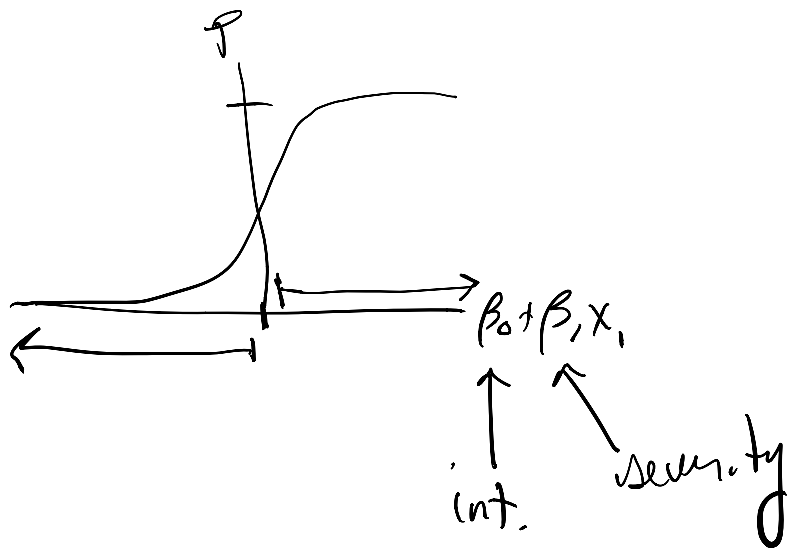

Distribution

Student's t-dstribution

$log(\frac{p}{1-p})$

Test Statistic:

$$t = \frac{\hat{\beta}_1}{se_{\beta_1}}$$

Degrees of freedom = $n - 2$

Statistical Conclusion

If p-value $< \alpha$, reject the null hypothesis (H_0).
If p-value $\geq \alpha$, fail to reject the null hypothesis (H_0).

Clinical Conclusion

There is enough evidence to show that x is predictive of y with an odds ratio of e^{β_1}.

Non-parametric alternatives

- Nonparametric logistic regression

Power and Sample Size Calculations:

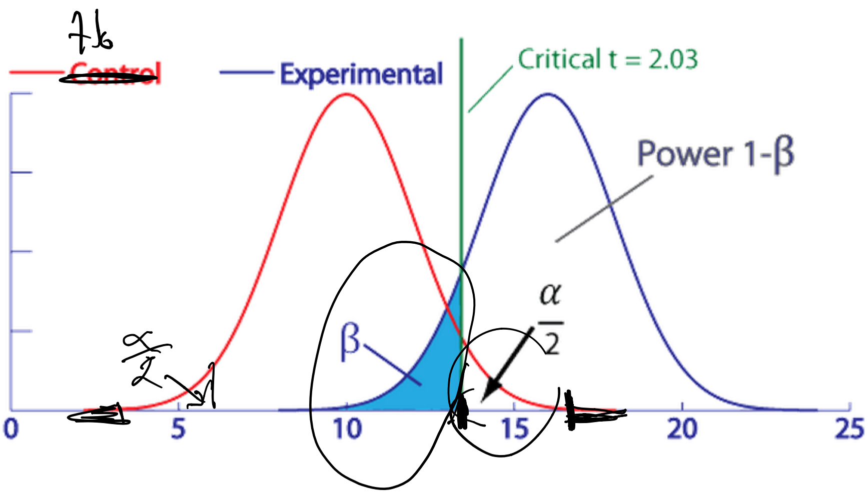

Motivations:

- Balances the Type I error rates with Type II error rates; minimizes the likelihood of false conclusions.
- Allows you to balance resources by requiring the necessary and sufficient number of participants.
- Most granting and regulatory agencies require us to compute sample sizes.

Working Parts:

n The total number of participants required to complete the experiment as designed.

β The probability of committing a Type II error, or the probability of failing to reject the null hypothesis when the alternative hypothesis is true. Power is represented by $1-\beta$, and is usually defaulted to 80%. For more conservative experiments (e.g., clinical trials), a power of 90% may seem reasonable. In contrast, exploratory experiments may use a lower power (70%).

α The probability of committing a Type I error, or the probability of rejecting the null hypothesis when the null hypothesis is true. The default value of α is 0.05.

d The size of the effect you wish to uncover. Effect sizes represent the magnitude of the difference between the two groups in the experiment. Effect sizes typically come from preliminary data or previously reported levels in the literature. In the event an evidence-based effect size cannot be identified, standard conventions of "small", "medium", and "large" exist.

Statistical test Sample sizes and power calculations depend on knowing the statistical test to be performed, for example, a t-test or ANOVA. Since the calculations depend on this knowledge, these working part is *required*.

Types of calculations:

- Computing required sample size (n).
- Computing observed effect sizes (d).
- Computing achieved power ($1-\beta$).