Design and analysis of experiments
Lecture 1

Jakob G. Rasmussen
Department of Mathematics
Aalborg University
Denmark

Practical information

- Teacher: Jakob G. Rasmussen
- Textbook:
  Douglas C. Montgomery
  Design and Analysis of Experiments, International Student Version, 8th edition, 2012
- Evaluation:
  - At least 80% attendance (i.e. 8 lectures)
  - Complete 2 handin exercises
- Participant-list
- Computers & chords
- R & RStudio

How not to do statistics - example 1

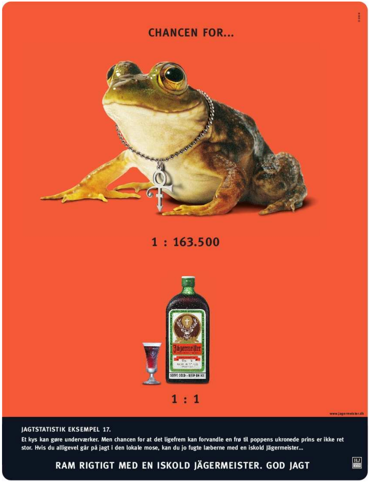

- This commercial claims that when you kiss a frog it will only very rarely turn into a prince, but you will always have succes with a Jägermeister
- Or does it... ?

How not to do statistics - example 2

Figure 2. Winning times for the Boston Marathon for men (1897 through 1992) and women (1972 through 1992) augmented with the best linear fit. Data from The 1993 Information Please Sports Almanac (p. 619–620).

How not to do statistics - example 3

Do storks really bring babies?

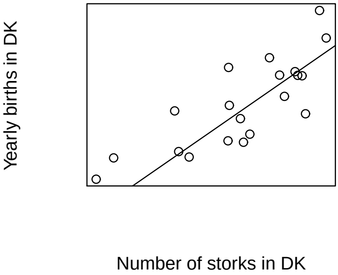

How not to do statistics - example 3

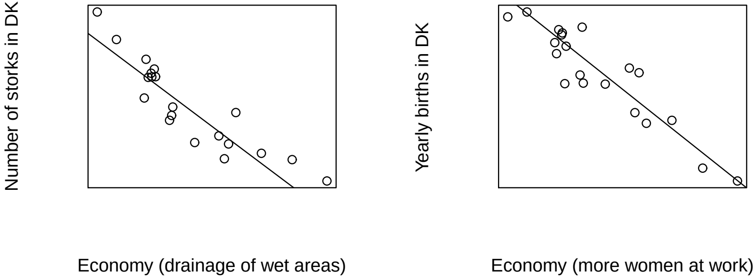

As the economy has grown, both the number of storks and the number of births have decreased - the economy is a confounding variable.

Examples - a few conclusions

- Example 1: Make sure you express what you think you express
- Example 2: Do not overinterpret your model, and derive conclusions that it does not support
- Example 3: Tendencies in a model can be something completely different

But of course there are many other things that can go wrong - a solid understanding of statistics is important to analyse experiments as is a solid understanding of the problem at hand.

Introduction to R

- R demonstration, part 1 - no statistics so far
- Exercises 1 + 2

Experiments

What is an experiment?

- Designed: Designing the experiment from scratch
- Observation, prospective: Making new observations without interfering
- Observation, retrospective: Using previously recorded observations

The structure of experiments

1. Recognizing the problem
2. Selecting the response variable
3. Choice of factors, levels, ranges
4. Experimental design
5. Performing the experiment
6. Statistical analysis
7. Conclusions

We will focus on 6. in this course, but this has many implications for the other items.

Experiments - a bit of good advice

- Use your knowlegde of the problem
- Keep it simple
- Practical and statistical significance are not the same
- Iterative experiments

Chapter 1 contains lots of other useful information, but now we turn to random variables.

Discrete random variables

- Discrete random variables take values in a countable set
- Examples of experiments leading to discrete random variables
  - Roll a die or flip a coin
  - Number of visits on a webpage
  - Number of failures in a production
- Examples of discrete distributions
  - Binomial distribution
  - Poisson distribution
- Discrete distributions are important, but not the focus of this course!

Continuous random variables

- Continuous random variables take values in (an interval of) the real numbers
- Examples of experiments leading to continuous random variables
  - Height of a person in this room
  - The amount of sugar in a patients blood
  - The time until the next accident in an airport
- Examples of continuous distributions
  - Normal (or Gaussian) distribution
  - t-distribution
  - χ^2 (Chi-squared) distribution
  - F-distribution
- We will have a brief look at these four continuous distributions

Normal distribution

- Density function for the normal distribution:

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

- μ is the mean, and σ^2 is the variance

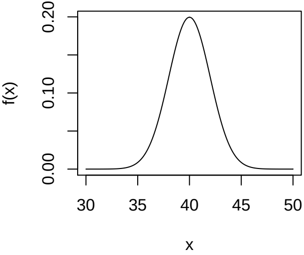

Standard normal distribution

- $N(0, 1)$ is called the standard normal distribution
- Standardization: if $X \sim N(\mu, \sigma^2)$, then

$$Z = \frac{X - \mu}{\sigma} \sim N(0, 1)$$

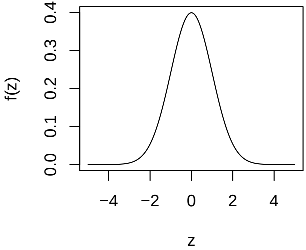

Central limit theorem

- Data can often be assumed to be normally distributed - however, when this does not hold, the central limit theorem is useful
- Central limit theorem:
  If $X_1, \ldots, X_n$ are independent and identically distributed with (finite) mean μ and (finite) standard deviation σ, then for $S_n = X_1 + \cdots + X_n$ we get

$$\lim_{n \to \infty} \frac{S_n - n\mu}{\sqrt{n}\sigma} \sim N(0, 1)$$

- As as rule of thumb if $n > 30$ then the sum is approximately normally distributed

χ^2-distribution

- Let $Z_1, \ldots, Z_n \sim N(0,1)$ be independent.
- Definition:

$$U = Z_1^2 + \cdots + Z_n^2 \sim \chi_n^2$$

  where n is the number of degrees of freedom (d.f.)
- Mean:

$$\mathbb{E}[Z_i^2] = \mathbb{E}[(Z_i - 0)^2] = \text{Var}(Z_i) = 1$$
$$\mathbb{E}[U] = \mathbb{E}[Z_1^2] + \cdots + \mathbb{E}[Z_n^2] = n$$

- Density function for χ^2-distribution with 10, 30 and 50 d.f.:

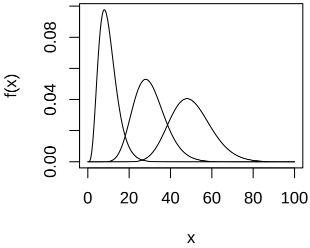

t-distribution

- Let $Z \sim N(0,1)$ and $U \sim \chi^2_n$ be independent.
- Definition:

$$t = \frac{z}{\sqrt{U/n}} \sim t_n = \frac{N(0,1)}{\sqrt{\chi^2_n/n}} \qquad (n = \text{d.f.})$$

- The mean is 0 for $n > 1$ (undefined for $n = 1$)
- Density function for t-distribution with 1, 3 and ∞ d.f.:

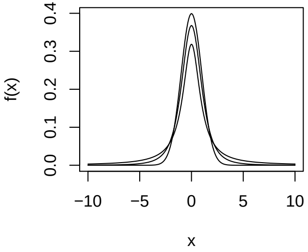

F-distribution

- Let $U \sim \chi^2_n$ and $V \sim \chi^2_m$ be independent.
- Definition:

$$F = \frac{U/n}{v/m} \sim F_{n,m} = \frac{\sqrt{\chi^2_n/n}}{\sqrt{\chi^2_m/m}} \qquad ((n,m) = \text{d.f.})$$

- Density function for F-distribution with various d.f.:

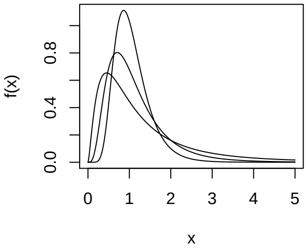

R

- R demonstration, part 2
- Exercise 3