

Bonding fasteners to wood using statistical data and methods.

Martin Schöön

13th November 2013

This pamphlet highlights the benefits of using a statistical approach in engineering rather than using nominal values, min/max limits or rule-of-thumb margins.

This text is *not* about workmanship or process control. If you are interested in the practical implementation of these hardware-to-wood bonds you should read “The Gougeon Brothers on Boat Construction”.

The reader that can stand my enthusiasm for using science-based methodology will be rewarded by finding out how brewing beer played an important role in developing these methods.

A note on typography and printing: The layout is intended for A5 format paper. It makes for easy on-screen reading and if you have access to Adobe Reader and a printer capable of double-sided printing you can produce a nice little ‘book’ by using the booklet printing option in Reader’s print dialogue.

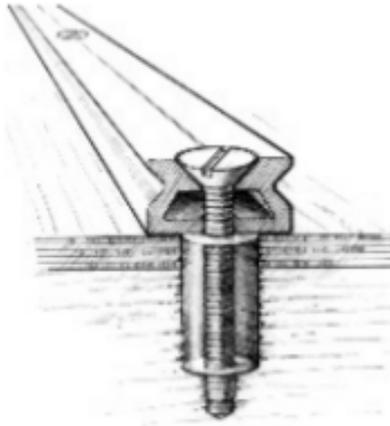


Figure 1: Fastener in oversize hole with epoxy. (From page 365 in the fifth edition of "The Gougeon Brothers on Boat Construction".

1 Bonding fasteners to wood – take one

1.1 Background and data

Attaching hardware to wood is problematic. Metals are hard and wood is soft. One way to improve the interface between fasteners and the wood structure is to use epoxy resin as an interface material. The idea is to increase the wood surface that carries the load without using oversize fasteners. This is done by drilling holes that have larger diameter than the fasteners and filling the extra volume with epoxy resin. The properties of such fastener/epoxy/wood interfaces were published in the 1970s (see page 17). You find my SI units (MPa) translation of the results in Table 1.

Table 1: Measured shear strength in MPa of epoxy/wood interfaces.

	Spruce	Ash	Mahogany
	7.22	9.17	10.14
	7.53	10.73	11.00
	7.76	10.93	11.75
	8.78	11.51	13.93
	11.84	11.54	15.33
	13.17	13.08	15.79
mean	9.38	11.16	12.99
standard deviation	2.51	1.28	2.36

1.2 Design based on the recommended 'safety factor'

Those of you who have looked in the Appendix may have noticed the design recommendation at the bottom of the table on page 17: Use the average strength divided by two for dimensioning holes. In Table 2 I summarize what such a design approach leads to for a 5 kN load.

Table 2: Designs for 2x safety factor

	Spruce	Ash	Mahogany
stress [MPa]	4.96	5.58	6.50
failure rate [%]	3.07	0.0006	0.29
C_{pk}	0.62	1.45	0.92

It is clear that this design approach leads to inconsistent results. 3% failures for spruce is not too impressive. For mahogany the result is ten times better and for ash it is really good at six failures in one million samples.

1.3 Designing using probability distributions

The rule-of-thumb method used above neglects some of the information obtained from measurements. If we use both mean and standard deviation we can do a better job. We can choose a failure rate we

accept and calculate the epoxy-to-wood bond area needed to achieve that failure rate.

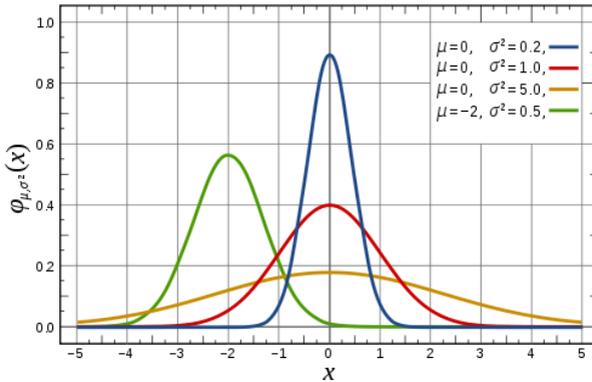


Figure 2: Normal probability density functions.

Fig. 2 shows the probability density functions (PDFs) of four different distributions. The failure rate for one of these distributions for a chosen value of χ is the area under the PDF curve to the left of that χ -value. The result depends on both the location of the curve and how much it is spread out. Pick for instance the value -3 . The area to the left of -3 under the green curve is clearly larger than the corresponding area under the red or the blue curve. It is hard to see from these curves if the failure rate of the yellow distribution is worse than that of the green distribution for $\chi = -3$.

Fig. 3 shows the area under the probability density functions also known as the cumulative density functions (CDF). Now we clearly see how close the green and the yellow distributions are for $\chi = -3$. We also see what the probability of failure is. It is close to 10% and both the red and the blue distributions have much lower failure probabilities.

Both Fig. 2 and Fig. 3 are taken from Wikipedia's article on the normal probability distribution¹. You find all the equations involved there and since spread sheet programs such as Microsoft's Excel and

¹http://en.wikipedia.org/wiki/Normal_distribution

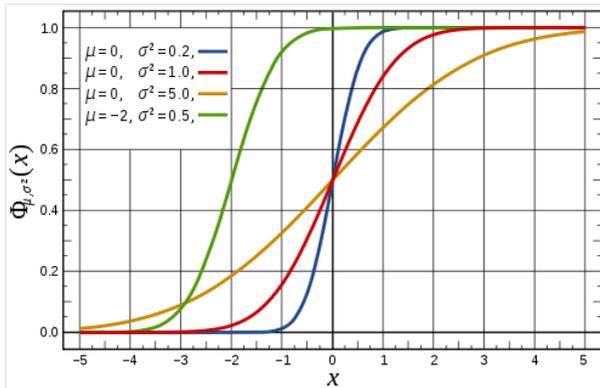


Figure 3: Normal probability cumulative distribution functions.

LibreOffice Calc² come with these functions I don't need to repeat them here. If we have values for mean and standard deviation we can make our favorite program calculate both probability density and cumulative density function values.

1.4 Designing for a 5 kN load and 0.1% failure rate

Using the inverse normal probability function of one of those spread sheet programs I calculated the acceptable shear stress and the corresponding barrel area of the drilled holes for a failure rate of 0.1% and a 5 kN load. The results are shown in Table. 3.

Table 3: Designs for 0.1% failure rate

	Spruce	Ash	Mahogany
stress [MPa]	1.63	7.21	5.70
area [mm ²]	3066	694	877

Spruce is the lightest and softest of the three and it coming out of this with the lowest acceptable stress level is no surprise. Ash is

²<http://www.libreoffice.org/>

actually better than mahogany despite having a lower mean value for its strength. This is due to the much smaller variation in the ash data. Studying the plotted PDFs in Fig. 4 we see what is going on. Below 8 MPa we see there is more area under the yellow mahogany curve than under the red ash curve.

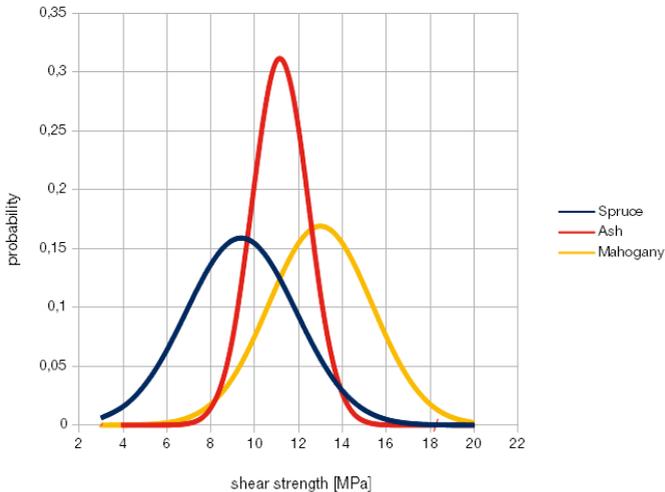


Figure 4: Measured shear strength PDFs.

The failure rate used in this example, 0.1%, was picked at random to produce this example and should not be regarded as some sort of recommended failure rate. Much tougher goals are the norm in manufacturing industry³.

1.5 C_{pk}

In the last row of Table 2 I introduced a new parameter called C_{pk} or process capability index. C_{pk} is the distance between the mean of the distribution and the nearest requirement limit. The distance is

³http://en.wikipedia.org/wiki/Six_Sigma

measured in the 'unit' three standard deviations of the distribution. C_{pk} is also described in a Wikipedia article⁴ but since it is not a standard function in office suites I repeat its equations here.

$$C_{pk} = \min \left[\frac{UL - mean}{3sdev}, \frac{mean - LL}{3sdev} \right] \quad (1)$$

$$C_{pk} = \left| \frac{limit - mean}{3sdev} \right| \quad (2)$$

Eq. 1 is used for double-sided requirement intervals where UL is the upper limit and LL the lower limit. Eq. 2 applies to single-sided situations.

C_{pk} is just another way of expressing failure rate but it is in very wide-spread use in manufacturing which warrants its mentioning here.

⁴http://en.wikipedia.org/wiki/Cpk_Index

2 Digging deeper

Now I will pull the rug from under your feet. In Section 1.4 I 'forgot' to mention two problems with the design approach I used. What I have in mind is the fact that the means and standard deviations I used are based on a very small sample base – only six measured values for each wood species – and I just assumed without checking that the variation in the shear strength of the wood-epoxy bond has a normal distribution.

2.1 Estimates of mean and standard deviation

Based on a sample of a population we can only calculate *estimates* of the true mean and standard deviation. The confidence intervals of these estimates depend on the sample size and with just a few samples like in Section 1 we get large confidence intervals.

The relationship between estimated and real mean and standard deviation has been studied since the late 1800s. The most famous contribution was made by William Sealy Gosset little over 100 years ago. He worked at the Guinness Brewery in Dublin and was forbidden to publish scientific work. He did anyway and used a pseudonym: Student. R.A. Fisher later dubbed the distribution that approximates the normal distribution "Student's distribution"⁵. We owe it to the beer brewing industry that we have ways to pin down how good our estimates of normal distributions are!

Here is how to calculate the confidence intervals.

$$mean_{est} - \frac{t \cdot sdev_{est}}{\sqrt{n}} < mean_{true} < mean_{est} + \frac{t \cdot sdev_{est}}{\sqrt{n}} \quad (3)$$

Where n is the sample size and t is the argument to the cumulative t-distribution function with $(n - 1)$ degrees of freedom resulting in the output p . The level of confidence is $(1 - p)$.

$$sdev_{est} \sqrt{\frac{n-1}{x_2}} < sdev_{true} < sdev_{est} \sqrt{\frac{n-1}{x_1}} \quad (4)$$

⁵http://en.wikipedia.org/wiki/Student_t_distribution.

Where x_1, x_2 are the inputs to the cumulative χ^2 distribution (Chi2 distribution) with $(n - 1)$ degrees of freedom resulting in the output values p_1, p_2 . The level of confidence is $p_2 - p_1$.

This math may look scary but both Microsoft Excel⁶ and LibreOffice Calc⁷ come supplied with the necessary functions for calculating t, x_1 and x_2 .

2.2 Confidence intervals for the wood/bolt bonding example

Using the equations of Section 2.1 I have calculated 90% confidence intervals for the means and standard deviations of Section 1.1. The result is found in Table 4.

Table 4: 90% confidence intervals.

	Spruce	Ash	Mahogany
mean in interval [MPa]	7.32–11,45	10.11–12.21	11.05–14.93
sdev in interval [MPa]	1.69–5,24	0.86–2.67	1.59–4.93

With uncertainties like this our designs in Section 1.4 are based on a rather shaky foundation. The small size of the sample is the problem. If we pretend we have 30 samples, still a small sample size, for each wood species and re-calculate we find the confidence intervals shrink noticeably, see Table 5.

Table 5: 90% confidence intervals based on 30 samples.

	Spruce	Ash	Mahogany
mean in interval [MPa]	8.46–10.31	10.69–11.63	12.12–13.86
sdev in interval [MPa]	2.07–3.21	1.06–1.64	1.95–3.02

Does this mean a sample of six is useless? Yes and no. It depends on our needs and the consequences of a failure. Using the methods demonstrated above we can calculate how uncertain our estimates are.

⁶T.INV or T.INV.2T, CHISQ.INV and CHISQ.INV.RT

⁷TINV and CHIINV

If we need a really low failure rate there are two ways to get there: Either add a substantial 'margin' in the design guided by the calculations or if that isn't possible, reduce the level of uncertainty by additional measurements. We can even estimate how big that experiment need to be to reach a given level of confidence.

2.3 Is this data really from a normally distributed population?

All we have done so far do only work if the sampled population has a variation that has a normal distribution. If the true distribution is something else all estimates based on the normal distribution are bogus. We clearly need a way to test if the data is normally, or whatever, distributed.

Constructing a histogram based on measured data may be the first thing to attempt. Fig. 5 demonstrates the drawback of this approach: it does not work well with small sample sizes.

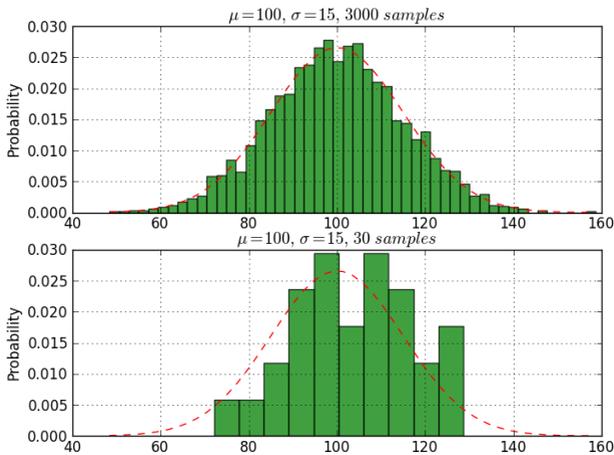


Figure 5: Examples of histograms constructed from samples of different sizes.

There is clearly room for improvement and there are indeed a number of alternative ways of testing if sampled data fits a certain distribution. Here I only cover something called the normal probability plot⁸. In Fig. 6 we have such plots for the high and low sample size data of Fig. 5. The basic function of the plot is to indicate if sampled data has a normal distribution. If the dots form a straight line this is the case and we see that even the small sample is fairly easy to judge.

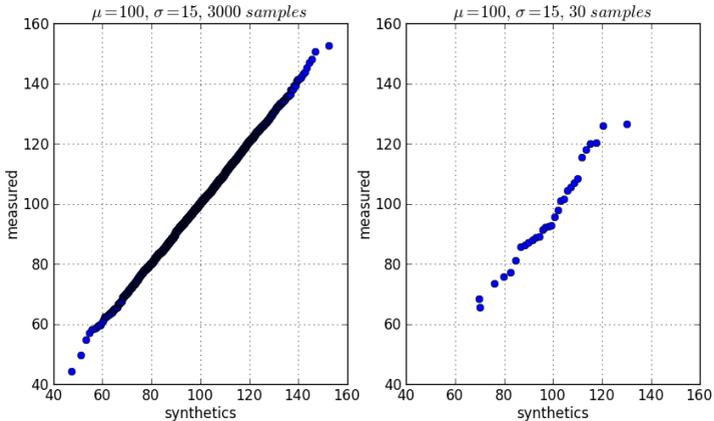


Figure 6: Normal probability plot examples.

The article on normal probability plots in Wikipedia is reasonably clear about how to construct these plots but I include a short procedure description here for completeness. It is fairly easy to do this with spreadsheet programs such as Microsoft's Excel and LibreOffice Calc. I used Python^{9,10} to make Fig. 6.

1. Calculate mean and standard deviation for the measured data.
2. Construct a vector containing normal cumulative probability data

⁸http://en.wikipedia.org/wiki/Normal_probability_plot

⁹<http://www.python.org/>

¹⁰<https://code.google.com/p/pythonxy/>

points x (n is the sample size)

$$x_i = \begin{cases} 1 - \sqrt[n]{0.5} & \text{for } i = 1 \\ \sqrt[n]{0.5} & \text{for } i = n \\ \frac{i-0.3175}{n+0.365} & \text{otherwise} \end{cases} \quad (5)$$

3. Calculate the inverse normal cumulative probability function with the data from step 2 as input. Put another way: Find the input to the normal CDF that result in the data calculated in step 2.
4. Plot points (x,y) in a rectangular grid where x is the data from step 3 and y is the ordered, measured data.

What you are plotting is measured data against 'synthetic' data with the same mean and standard deviation. If the points do not fall on a straight line it indicates that sampled data may have some other distribution.

Let us test our wood data like this (I used OpenOffice Calc to create Figs. 7, 8 and 9).

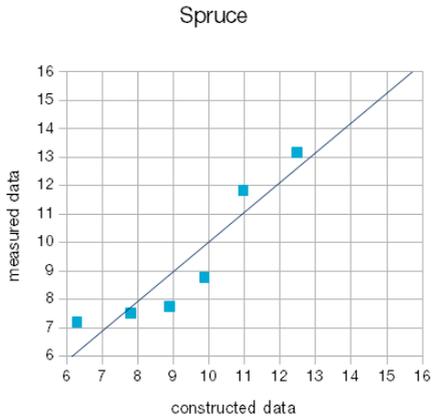


Figure 7: Spruce probability plot.

Both ash and mahogany looks decent but if I were to do something serious using spruce I would want to do some additional testing.

Ash

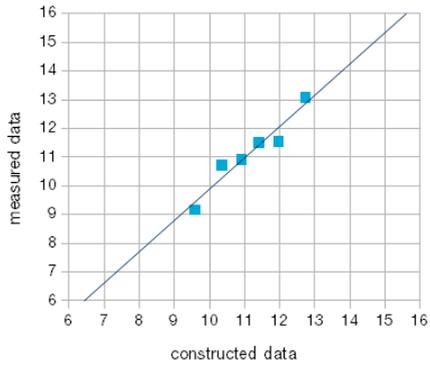


Figure 8: Ash probability plot.

Mahogany

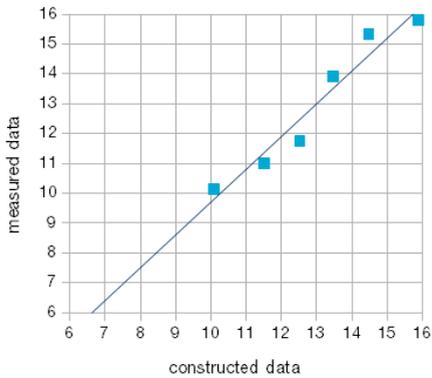


Figure 9: Mahogany probability plot.

2.4 Are my processes stable?

You may have noticed the order of the data in Table 1. If this is the case for a real measurement series it should be reason for concern. Your data may pass all statistical tests mentioned so far and still be corrupt if the mean and/or variation are not stable over time!

Fig. 10 is borrowed from W. Edwards Deming¹¹. Data looking like this indicates that either the manufacturing or the measurement process or both are unstable.

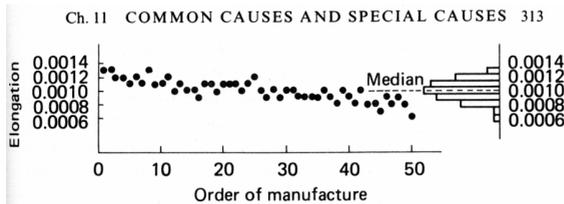


Figure 10: Spring stiffness variation over time.

2.5 How to handle aging products and variable load conditions

In a commercial context engineers are faced with the problem of designing for an acceptable field return rate given a reality where delivered products see a wide variation in load cases and where their capability to withstand those loads degrade over time. This is a rather big subject and I will only touch it.

I have created a basic example loosely based on the wood bonding example of Section 1. The load now has a normal distribution designed such that 5 kN is one standard deviation away from the mean value 4.6 kN.

The design of the epoxy bond is done in such a way that its failure rate with respect to a 5 kN limit is 0.1% – the same as in the design

¹¹W. Edwards Deming: "Out of the Crisis", The MIT Press, ed. 2000, ISBN-13 978-0-262-54115-2

example of Section 1.4. As time goes by the wood-epoxy bond ages resulting in a lower mean value and increased standard deviation for its strength. All this is illustrated in Fig. 11. NB, this aging data is not based on measurements. It is pure fiction!

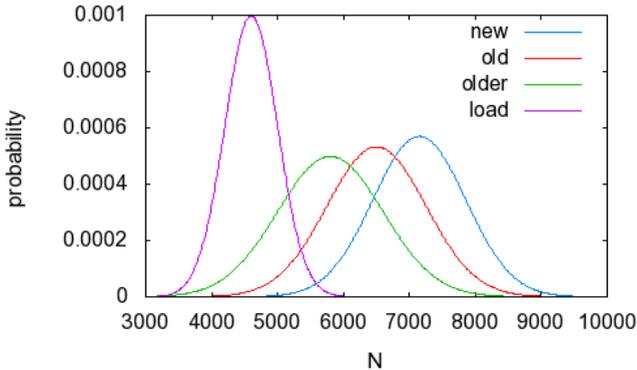


Figure 11: Illustrating aging equipment and a variable load.

In this illustration the probability of failure is *not* the area under the strength probability curve to the left of the 5 kN limit as it depends on the probabilities of *both* the load and the strength.

A brute force Monte Carlo simulation, where randomly generated loads and strength numbers are compared, is the most straight-forward way to estimate the failure rate in this case. The number of failures – when load sample exceeds strength sample – are added and divided by the total number of Monte Carlo Runs. The results of such a simulation are listed in Table 6.

Even though the distribution of the load exceeds the old 5 kN limit and overlaps the ‘new’ strength distribution I get a lower probability of failure: 0.075% vs 0.1%. This 25% reduction hints at the potential for saving money/weight by better understanding what our designs will face once released into the wild.

I have used a more mathematically-minded process to generate Fig. 12 to show how the failure probability changes as the overlap increases. I had to multiply the ‘new’ curve by ten to make it visible. The results of Table 6 are the areas under the curves in Fig. 12.

Table 6: Calculated failure probabilities.

new	0.075%
old	1.3%
older	9%

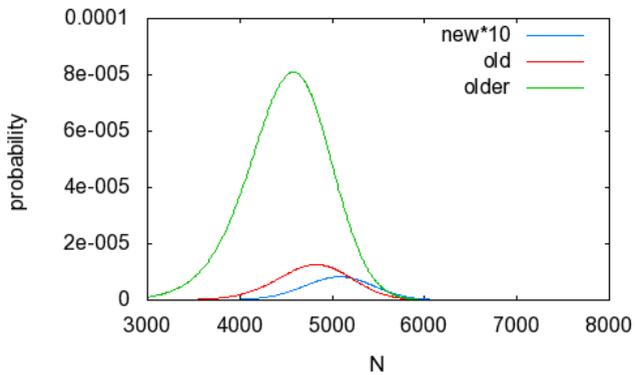


Figure 12: Failure probability curves.

Appendix

Experimental data on bolts bonded to wood

The following is copied from page 301 in the first edition of "The Gougeon Brothers on Boat Construction".

B-1—Wood/Epoxy and Threaded Rod Bonding

The following was excerpted from a research paper entitled "Wood as an Engineering Material with Applications of Concentrated Loading." This paper is the product of research performed by Kurt Keidel at Ohio State University in 1977.

In the experiment, mild steel threaded rod was bonded to wood using WEST System resin. They used samples of three types of wood: Honduras mahogany, white ash and Sitka spruce. The diameter of the bolt hole was 1/16 inch larger than the rod diameter. This allowed full encapsulation of the rod by the epoxy. The hole and threaded rod were precoated before assembly.

The testing apparatus was designed so that one end of the threaded rod was secured to the stationary head, while a movable head pressed evenly against the sample block (see Figure B-1). Stress and strain were monitored during the test.

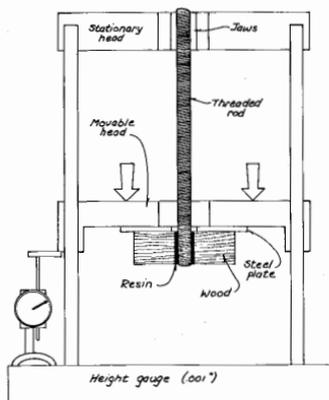


Figure B-1.

Ultimate shear failure (in psi)

	Honduras mahogany	White ash	Sitka spruce
Sample failures	2290.83	1897.28	1273.34
	1471.30	1584.61	1092.25
	2020.21	1329.94	1126.20
	1595.93	1556.31	1910.38
	1703.97	1673.74	1046.97
	2224.11	1669.50	1716.75
Mean shear	1884	1618.66	1360.00
Standard deviation	342.42	185.30	363.87
Allowable shear failure (with safety factor of 2)	942.2	809.28	603.83

Does my data make sense – the Anscombe quartet

In Section 2 we studied the fact that there might be some problem with our raw data and here I do some more of that but throw away statistical tools and bring in good old common sense. Francis Anscombe has created four sets of data to demonstrate the value of inspecting data and applying common sense. All four data sets fit the straight line $y = 3.00 + 0.500 * x$ equally well.

Table 7: Anscombe's quartet

I		II		III		IV	
x	y	x	y	x	y	x	y
10.0	8.04	10.0	9.14	10.0	7.46	8.0	6.58
8.0	6.95	8.0	8.14	8.0	6.77	8.0	5.76
13.0	7.58	13.0	8.74	13.0	12.74	8.0	7.71
9.0	8.81	9.0	8.77	9.0	7.11	8.0	8.84
11.0	8.33	11.0	9.26	11.0	7.81	8.0	8.47
14.0	9.96	14.0	8.10	14.0	8.84	8.0	7.04
6.0	7.24	6.0	6.13	6.0	6.08	8.0	5.25
4.0	4.26	4.0	3.10	4.0	5.39	19.0	12.50
12.0	10.84	12.0	9.13	12.0	8.15	8.0	5.56
7.0	4.82	7.0	7.26	7.0	6.42	8.0	7.91
5.0	5.68	5.0	4.74	5.0	5.73	8.0	6.89

Inspecting Table 7 we see right away that there is something special about the fourth data set. Plotting Anscombe's data-sets shows us convincingly how dangerous a 'blind' numerical analysis can be. The message of Fig. 13 is clear enough – there is no need for further comments.

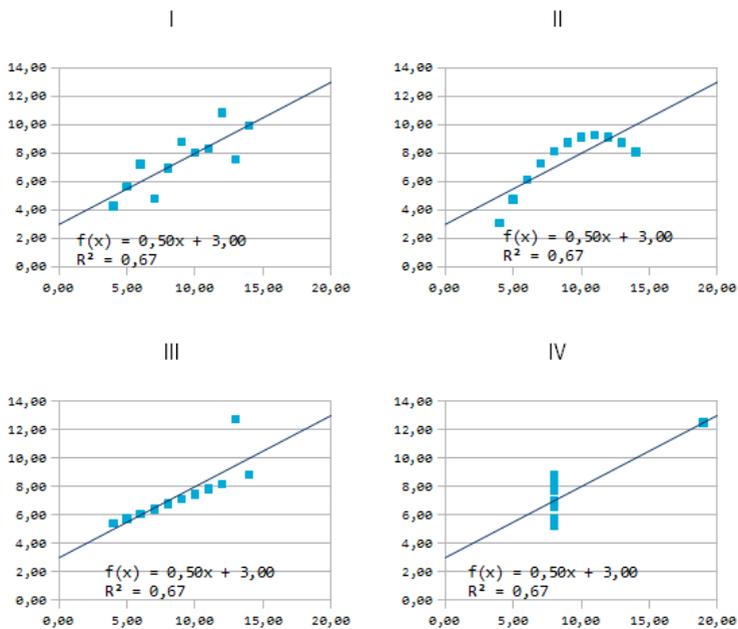


Figure 13: Anscombe's data-sets.