Applied Statistics

ProblemSet Solution and Discussion

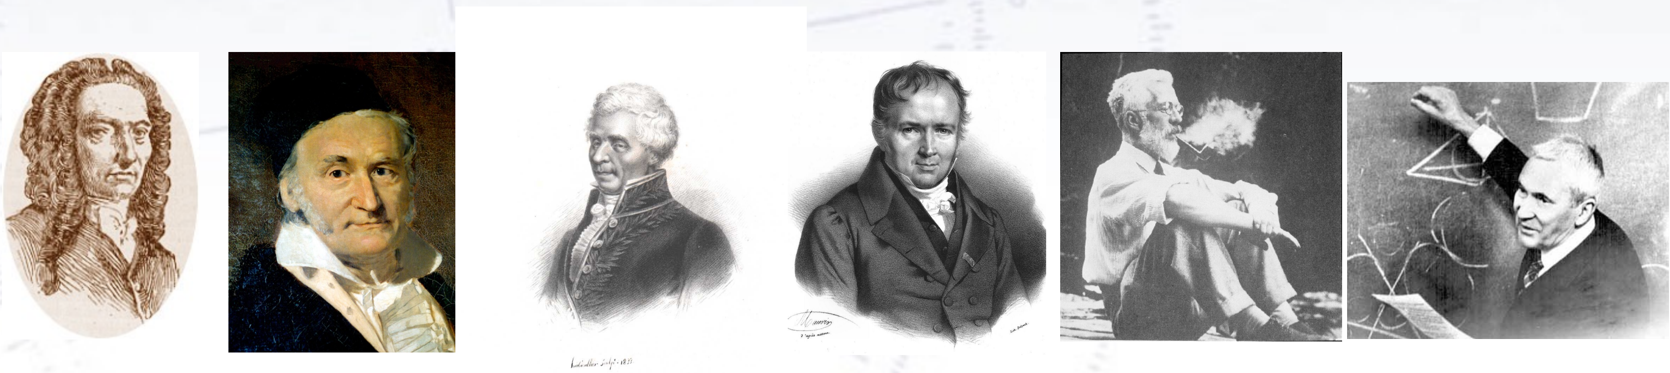

Troels C. Petersen (NBI)

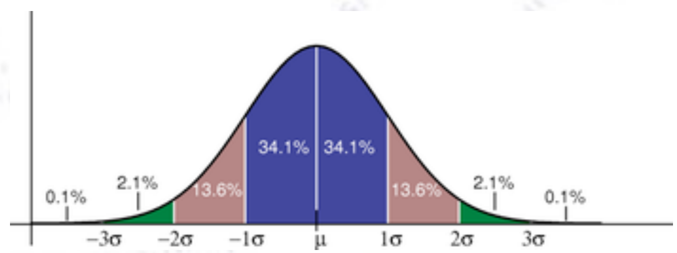

"Statistics is merely a quantisation of common sense"

Problem 1.1

1.1 (6 points) Little Peter goes to the casino and puts money on black ($p = 18/37$).

- In 50 games, what are the chances that he will win exactly 25 times? 26 times or more?
- How many times does he have to play in order to be 95% sure of winning at least 20 times?

$$P(k = 25|50, p) = \frac{50!}{25!^2}(p(1-p))^{25} = 0.110243$$

$$P(k > 25|50, p) = \sum_{k=26}^{50} \frac{50!}{k!(50-k)!} p^k (1-p)^{50-k} = 0.369458$$

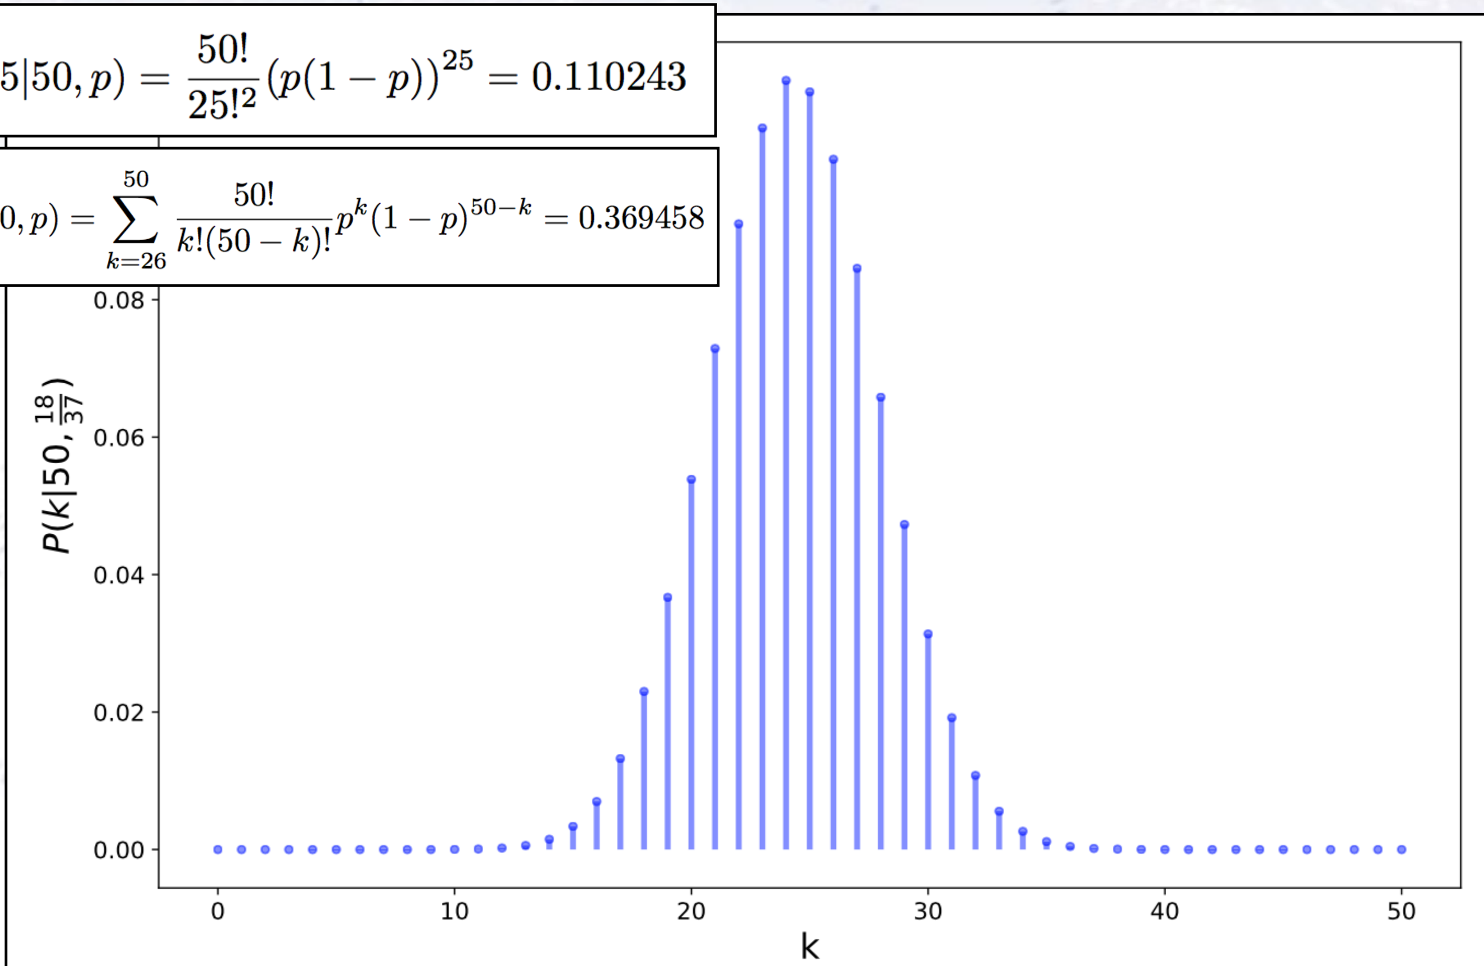

Problem 1.1

1.1 (6 points) Little Peter goes to the casino and puts money on black ($p = 18/37$).

- In 50 games, what are the chances that he will win exactly 25 times? 26 times or more?
- How many times does he have to play in order to be 95% sure of winning at least 20 times?

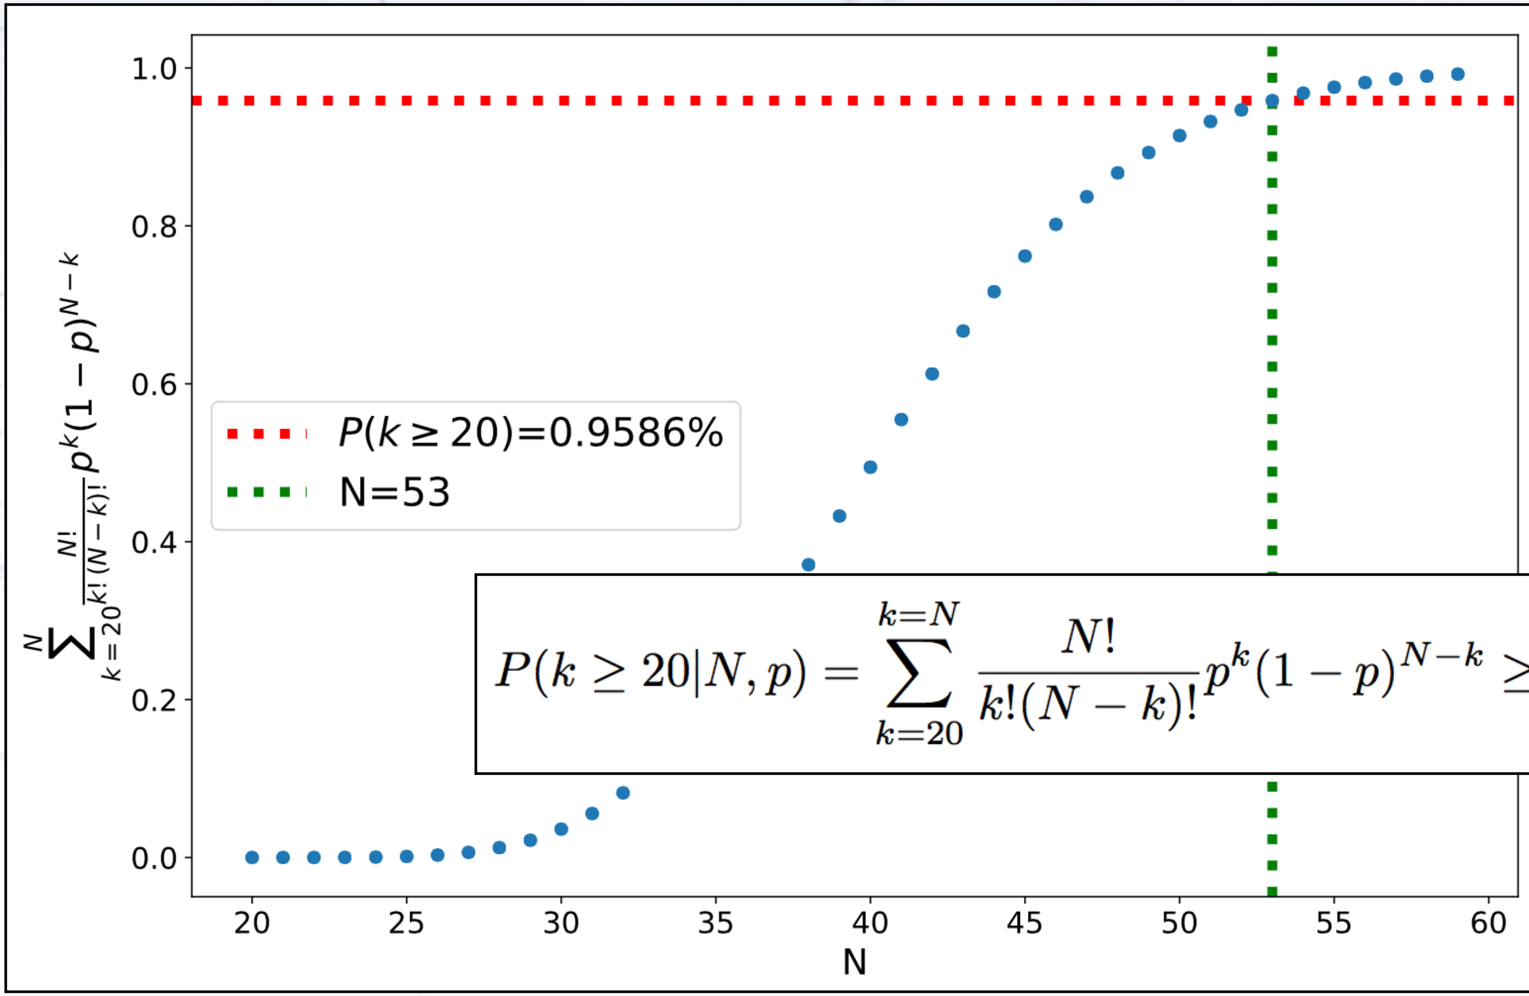

$$P(k \geq 20 | N, p) = \sum_{k=20}^{k=N} \frac{N!}{k!(N-k)!} p^k (1-p)^{N-k} \geq 95\%$$

Problem 1.2

1.2 (4 points) What is the probability of a Gaussian value to lie between 1.25σ and 2.5σ away from the mean?

$$P(1.25\sigma < x < 2.5\sigma) = 2 \cdot (\text{CDF}(2.5) - \text{CDF}(1.25)) = 0.1989$$

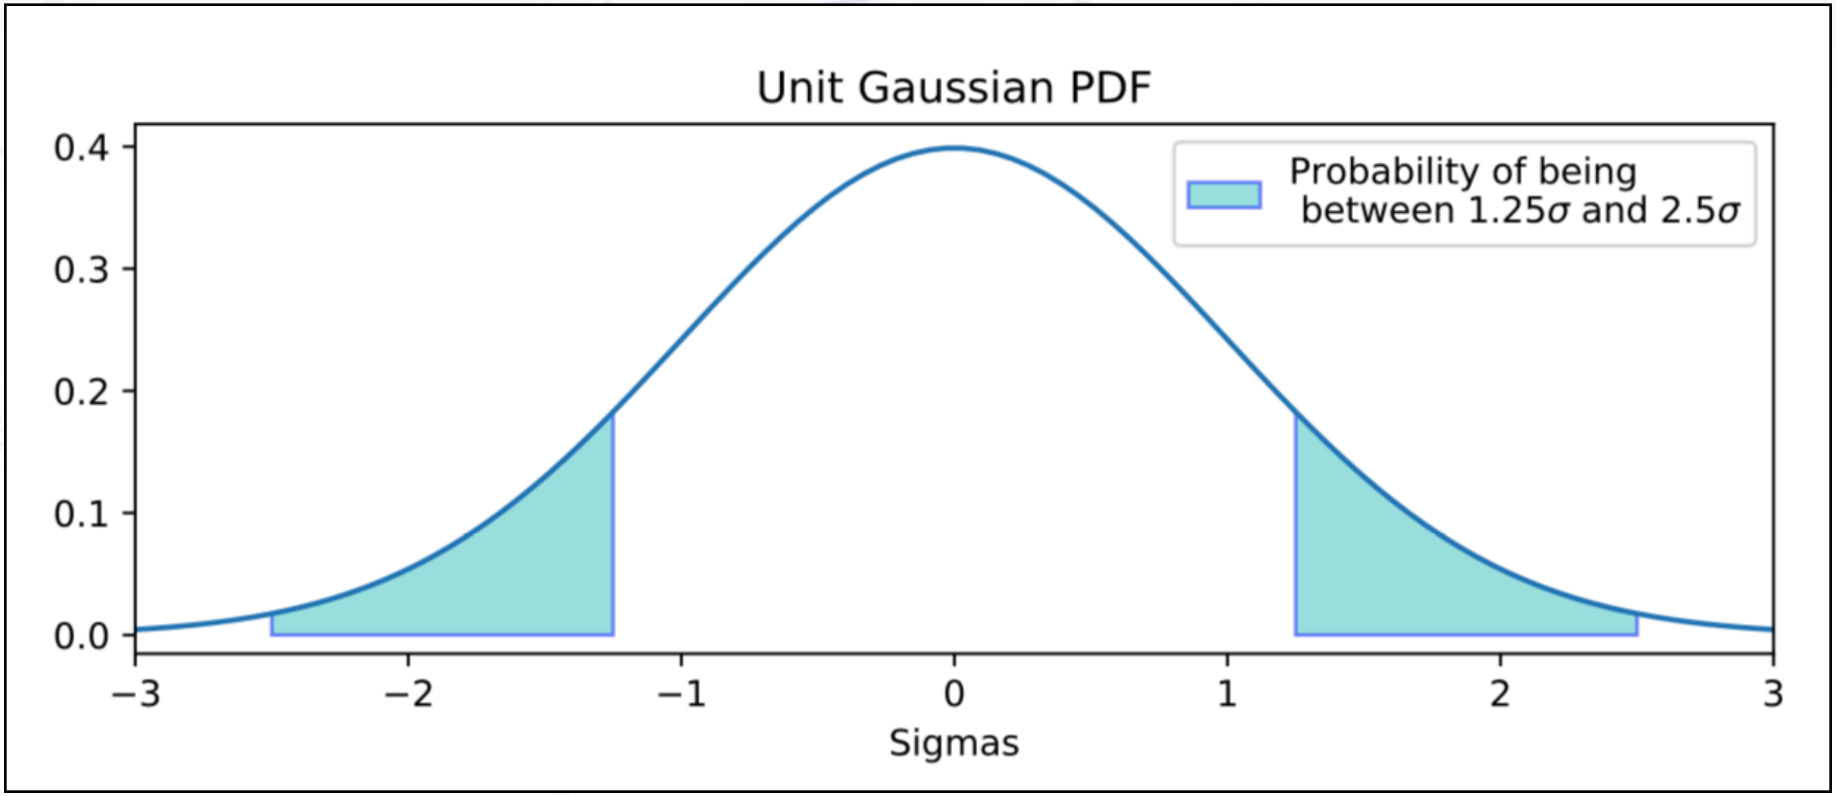

Problem 1.3

1.3 (6 points) The number of S-train delays is counted daily. Assume in the following, that delays are uncorrelated, and that the number of departures is the same every day.

- What distribution should the number of daily delays follow?
- Days with more than 7 delays are considered "delay days". If there were 19 "delay days" in a normal year, what is your estimate for the average number of daily delays?

Independent (?), N large, p small (both possibly varying): **Poisson**

Problem 1.3

1.3 (6 points) The number of S-train delays is counted daily. Assume in the following, that delays are uncorrelated, and that the number of departures is the same every day.

- What distribution should the number of daily delays follow?
- Days with more than 7 delays are considered "delay days". If there were 19 "delay days" in a normal year, what is your estimate for the average number of daily delays?

Independent (?), N large, p small (both possibly varying): **Poisson**

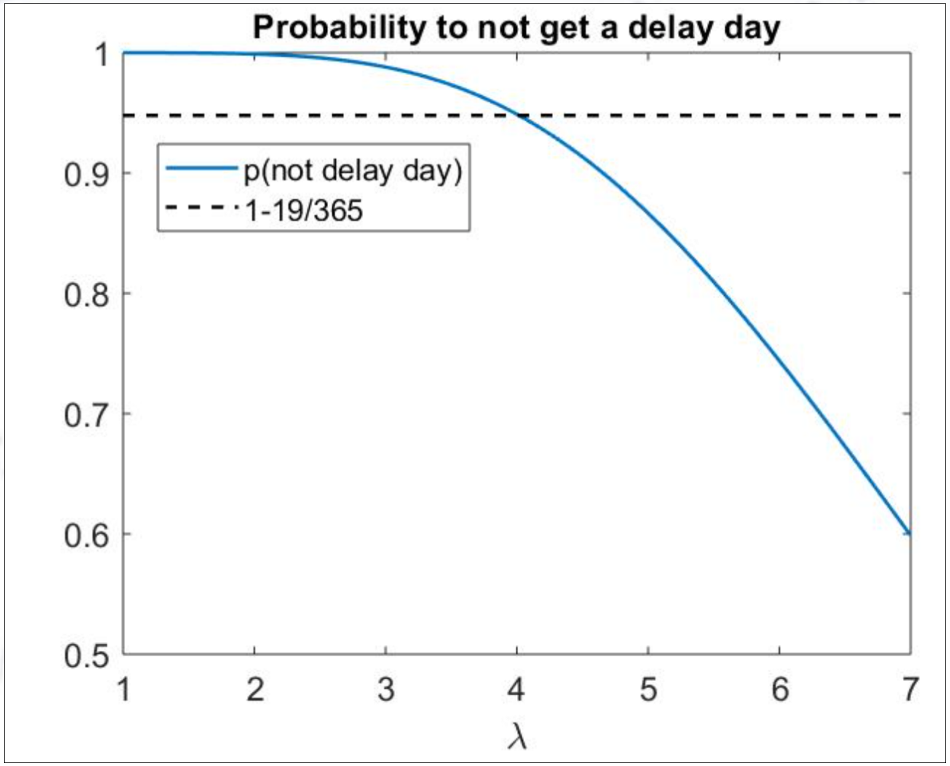

4.02 ± 0.10

Typical mistakes

Problem 1.1:

- Not including 26 in the second problem
- Rounding N to 50, 51, 52, 55 etc

Problem 1.2:

- One-sided gaussian integral
- Misinterpretation of the question, e.g. giving $P(x<2.5)$ and $P(x<1.25)$ without subtracting them.

Problem 1.3:

- Distribution identified as binomial or gaussian, with or without explanation
- Distribution identified as poissonian, but with missing or wrong explanation
- Forgetting to include 8 in calculation
- Fixing the result as lambda=4+/-2 days, using the squared root of lambda as uncertainty on mean.

Problem 2.1

2.1 (13 points) A measurement of a tumor depth (in cm) was done using two methods. The first gave 4 measurements with uncertainty while the second gave 12 without, as shown in the table.

With unc.	2.05 ± 0.11			2.61 ± 0.10			2.46 ± 0.13			2.48 ± 0.12		
Without unc.	2.69	2.71	2.56	2.48	2.34	2.79	2.54	2.68	2.69	2.58	2.66	2.70

- Do the measurements with uncertainty agree with each other? Do those without?
- Which of the two methods provide the most accurate positioning?
- What is your best estimate of the tumors position? And with what uncertainty?

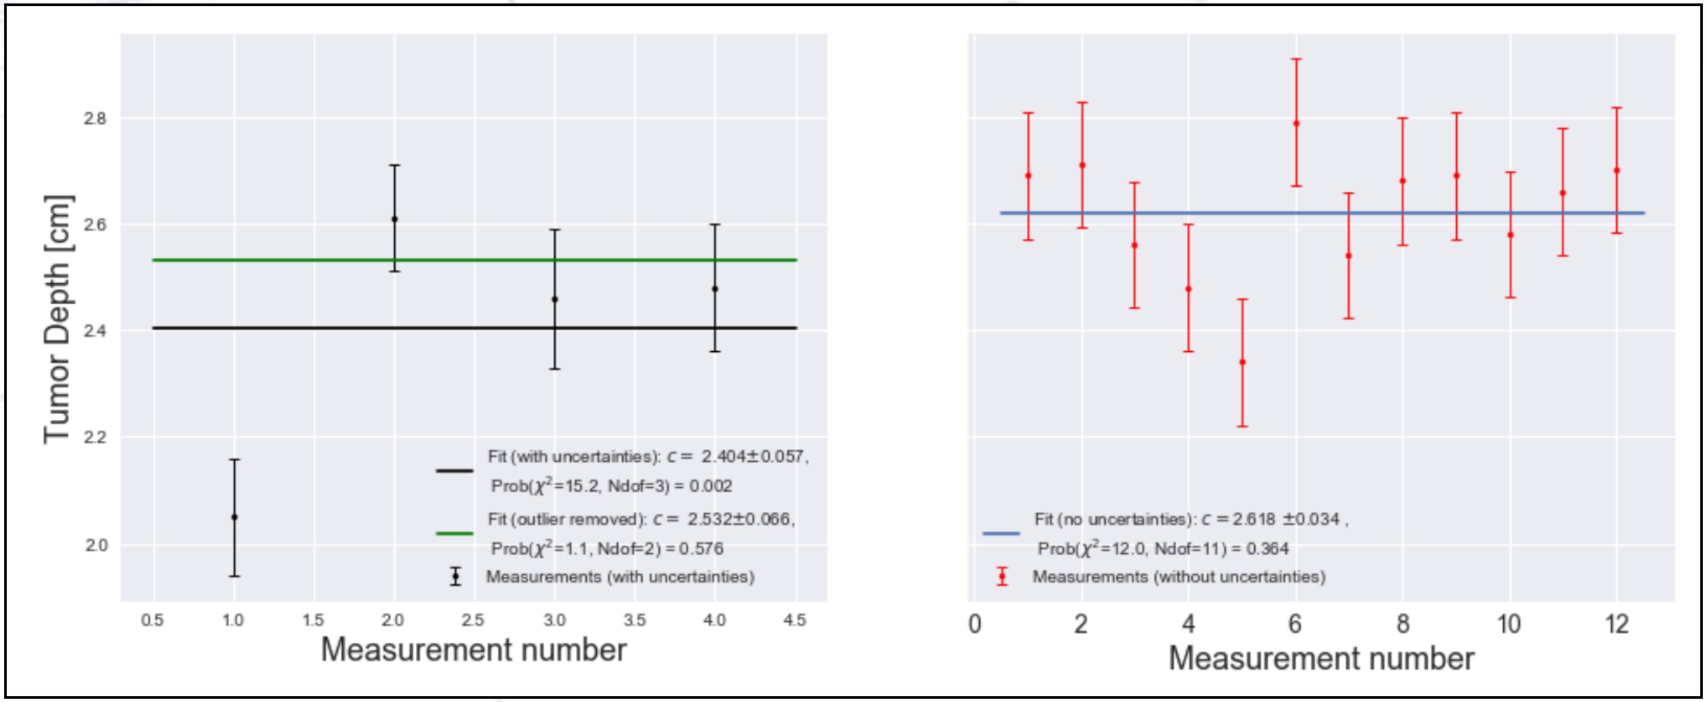

Problem 2.1

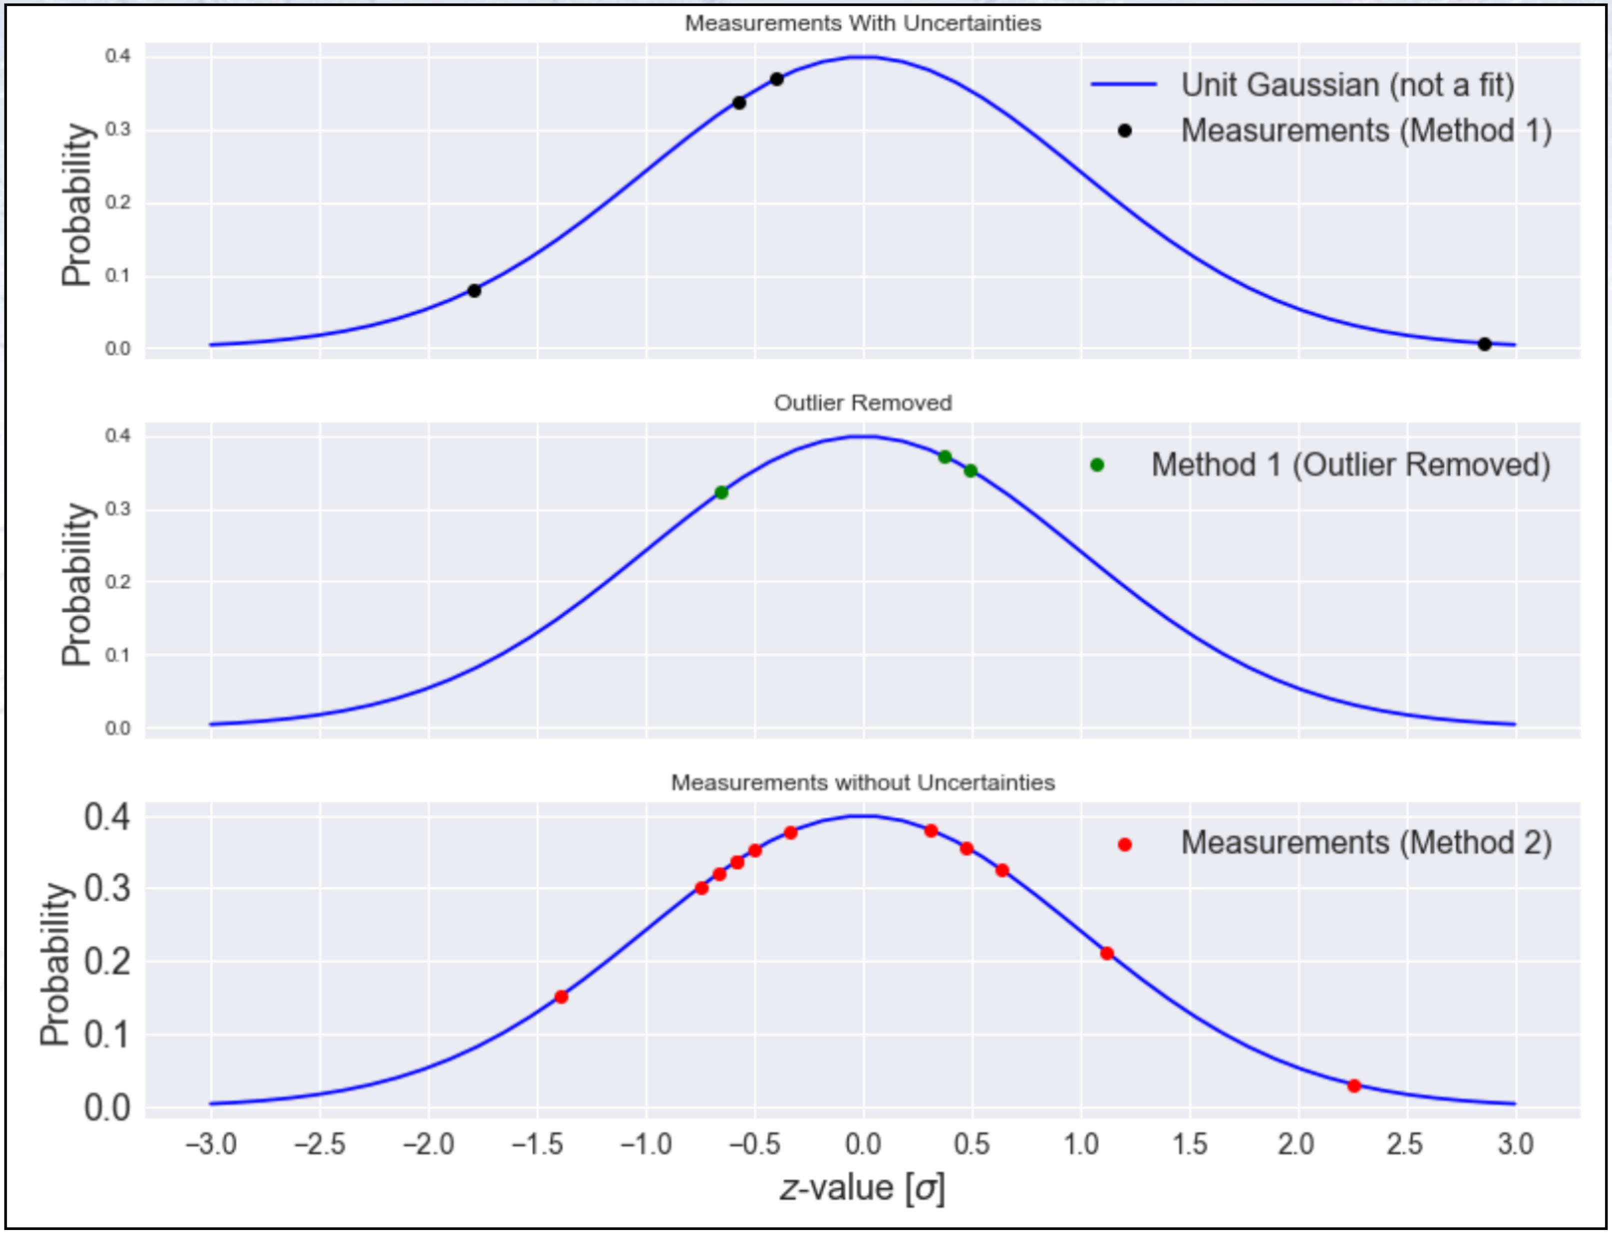

Measurements With Uncertainties
Unit Gaussian (not a fit)
Measurements (Method 1)
Probability
0.4
0.3
0.2
0.1
0.0
Outlier Removed
Method 1 (Outlier Removed)
Probability
0.4
0.3
0.2
0.1
0.0
Measurements without Uncertainties
Measurements (Method 2)
Probability
0.4
0.3
0.2
0.1
0.0
−3.0
−2.5
−2.0
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
2.0
2.5
3.0
z-value [σ]

Problem 2.1

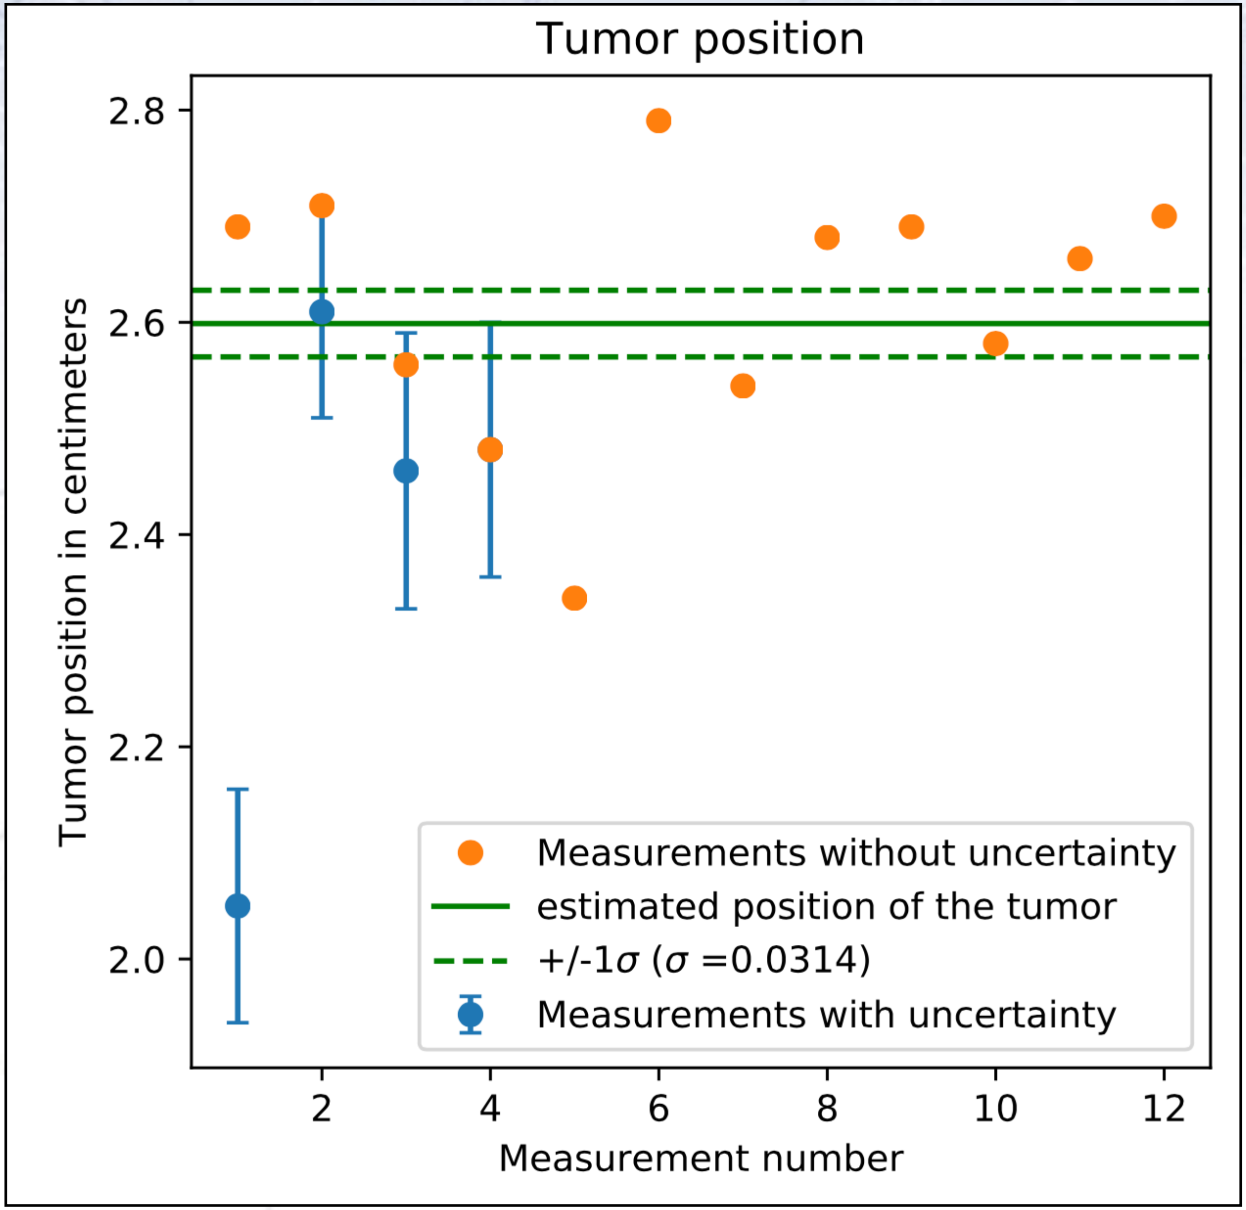

Tumor position
Tumor position in centimeters
2.8
2.6
2.4
2.2
2.0
2
4
6
8
10
12
Measurement number
Measurements without uncertainty
estimated position of the tumor
+/-1σ (σ =0.0314)
Measurements with uncertainty

Problem 2.1

Some rejected points among those without uncertainties!

We also want to look at the measurements without uncertainty. Here we find the mean and standard deviation of the measurements to be: $\mu = 2.6183 \pm 0.0357$, $std = 0.1236$. Again we look to see how far the measurements are from the mean. This time we do it with a one sample test, if z is more than 3σ away, we reject the measurement:

Measurement #	z-value	
Measurement 1	-2.01σ	$\Rightarrow$ accept
Measurement 2	-2.57σ	$\Rightarrow$ accept
Measurement 3	1.64σ	$\Rightarrow$ accept
Measurement 4	3.88σ	$\Rightarrow$ reject
Measurement 5	7.80σ	$\Rightarrow$ reject
Measurement 6	-4.81σ	$\Rightarrow$ reject
Measurement 7	2.20σ	$\Rightarrow$ accept
Measurement 8	-1.73σ	$\Rightarrow$ accept
Measurement 9	-2.01σ	$\Rightarrow$ accept
Measurement 10	1.07σ	$\Rightarrow$ accept
Measurement 11	-1.17σ	$\Rightarrow$ accept
Measurement 12	-2.29σ	$\Rightarrow$ accept

Typical mistakes

Problem 2.1:

– Wrong uncertainty on dataset2: **Forgetting to divide rms by squared(N)**.
– For measurements without uncertainties, using error on mean for distance and rejection.
– Noting that problem is "low statistics" in case with uncertainties!
– Wrong uncertainty on weighted mean and no chi2 test for weighted mean
– Not excluding first point of dataset1, either because chi2 test missing or because they are "not a fan" of this approach
– Forgetting to discuss precision
– Not combining the two datasets, after excluding the first point.
– Combining in weighted average, but forgetting to chi2 test
– Combining without excluding first datapoint, even if the chi2 test had failed
– Not giving a final, unique estimation of the depth

$$x = (2.60 \pm 0.03)\ \text{cm}; \quad \chi^2 = 2.43;\ p = 48.8\%$$

Problem 2.2

2.2 (9 points) The spectral radiance B of a body is given by Planck's Law: $B(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{k_B T}} - 1}$ where ν is the frequency and T is the absolute temperature, while $h = 6.626 \times 10^{-34}\ Js$, $c = 299.7 \times 10^6\ m/s$, and $k_B = 1.381 \times 10^{-23}\ J/K$ are constants of nature.

- Given values of $\nu = (0.566 \pm 0.025) \times 10^{15}$ Hz and $T = (5.50 \pm 0.29) \times 10^3$ K (uncorrelated), what is the expected spectral radiance, B?
- How does the uncertainty change, if there is a correlation of $\rho(\nu, T) = 0.87$?

$$\sigma_B^2 = 4\sigma_\nu^2 \left(\frac{3h\nu^2}{c^2\left(e^{\left(\frac{h\nu}{Tk_B}\right)} - 1\right)} - \frac{h^2\nu^3 e^{\left(\frac{h\nu}{Tk_B}\right)}}{Tc^2k_B\left(e^{\left(\frac{h\nu}{Tk_B}\right)} - 1\right)^2} \right)^2 \Bigg\} I_\nu^2$$

$$+ \sigma_T^2 \frac{4h^4\nu^8 e^{\left(\frac{2h\nu}{Tk_B}\right)}}{T^4c^4k_B^2\left(e^{\left(\frac{h\nu}{Tk_B}\right)} - 1\right)^4} \Bigg\} I_T^2$$

$$\boldsymbol{B(v,T) = (1.933 \pm 0.534) * 10^{-8}\ \frac{J}{m^2}}$$

$$\boldsymbol{B(v,T) = (1.933 \pm 0.370) * 10^{-8}\ \frac{J}{m^2}}$$

Problem 2.2

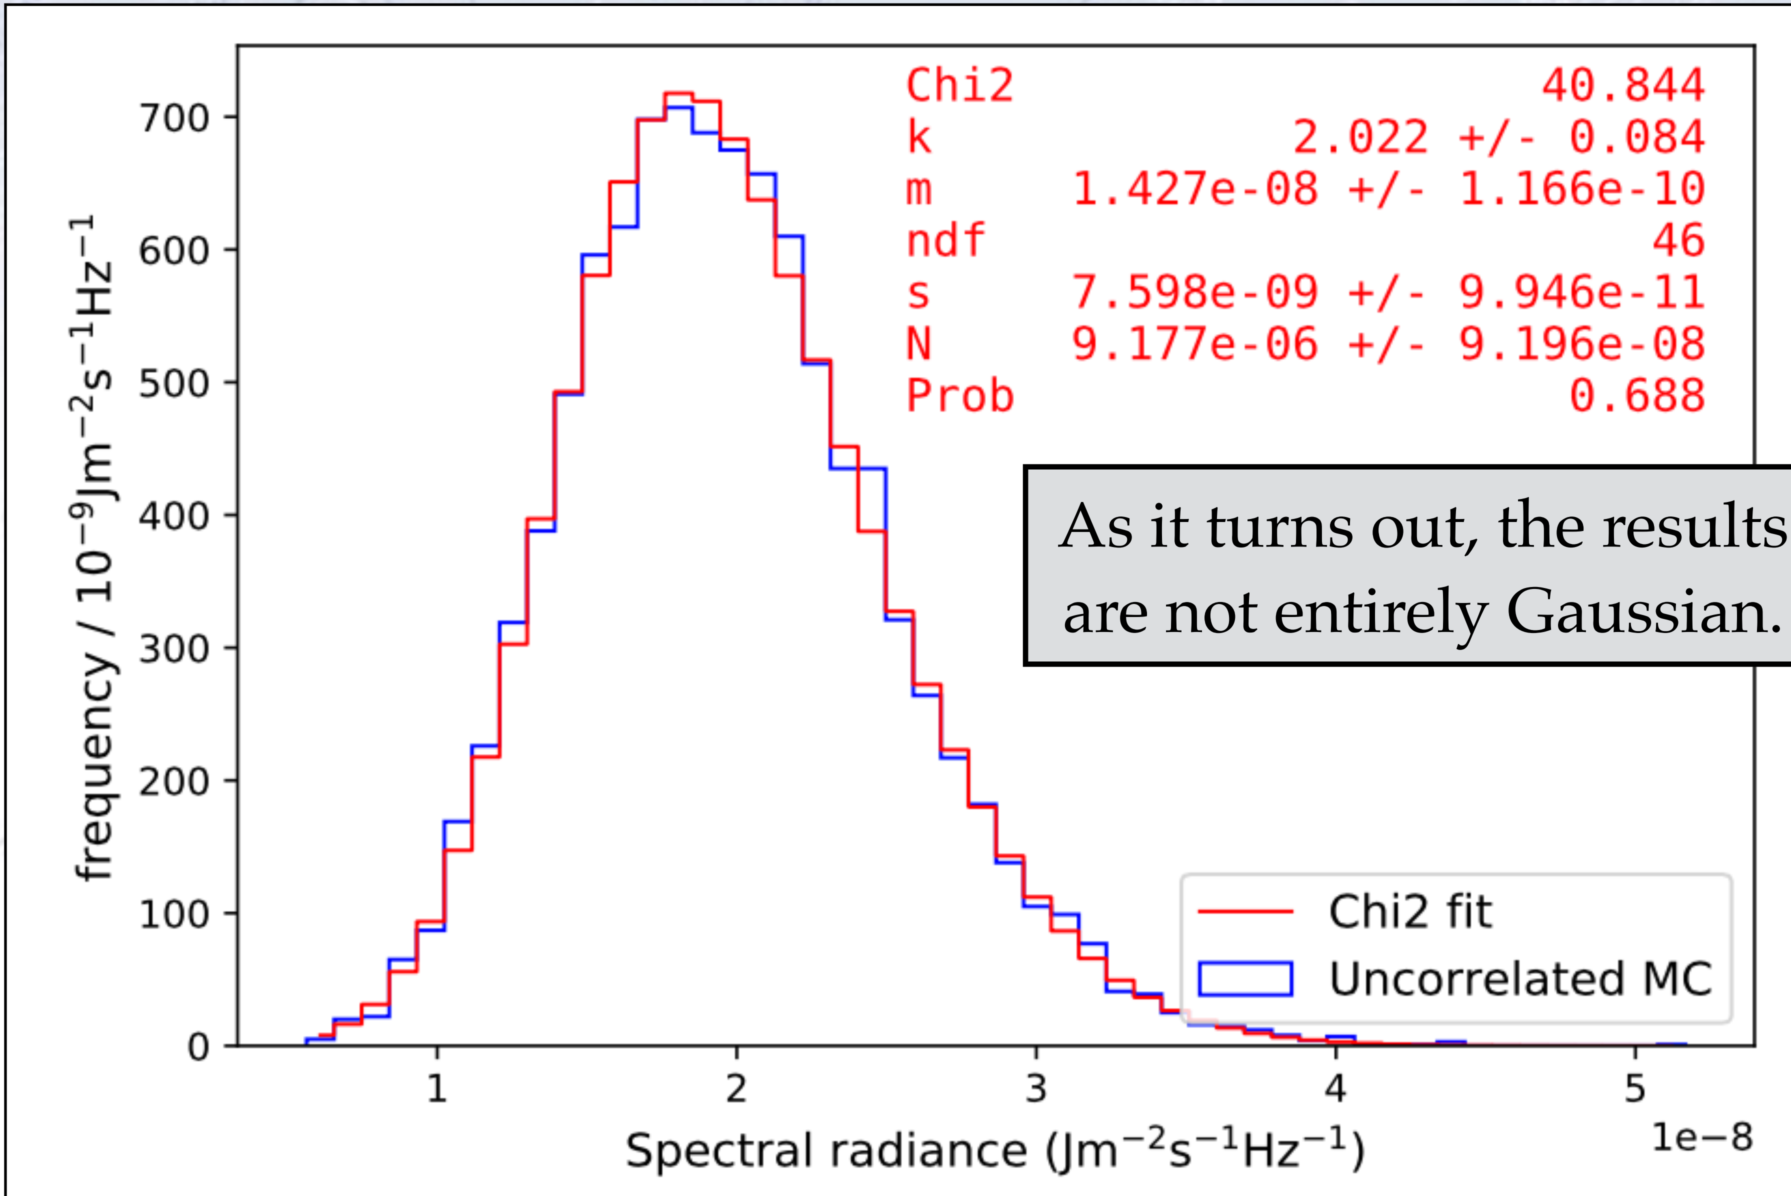

As it turns out, the results are not entirely Gaussian.

Problem 3.1

3.1 (15 points) Let $f(x)$ be a PDF defined as $f(x) = C(1 - e^{-ax})$ for $x \in [0, 2]$ and $a = 2$.

- What is the mean and RMS of $f(x)$? Also, what is the value of C?
- What method(s) can be used to produce random numbers according to $f(x)$? Why?
- Produce 500 random numbers distributed according to $f(x)$ and plot these.
- Fit the numbers you produced above leaving a as a floating parameter.
- Let u be a sum of 5 random values from $f(x)$. Produce 1000 values of u and test if they are consistent with a Gaussian distribution?

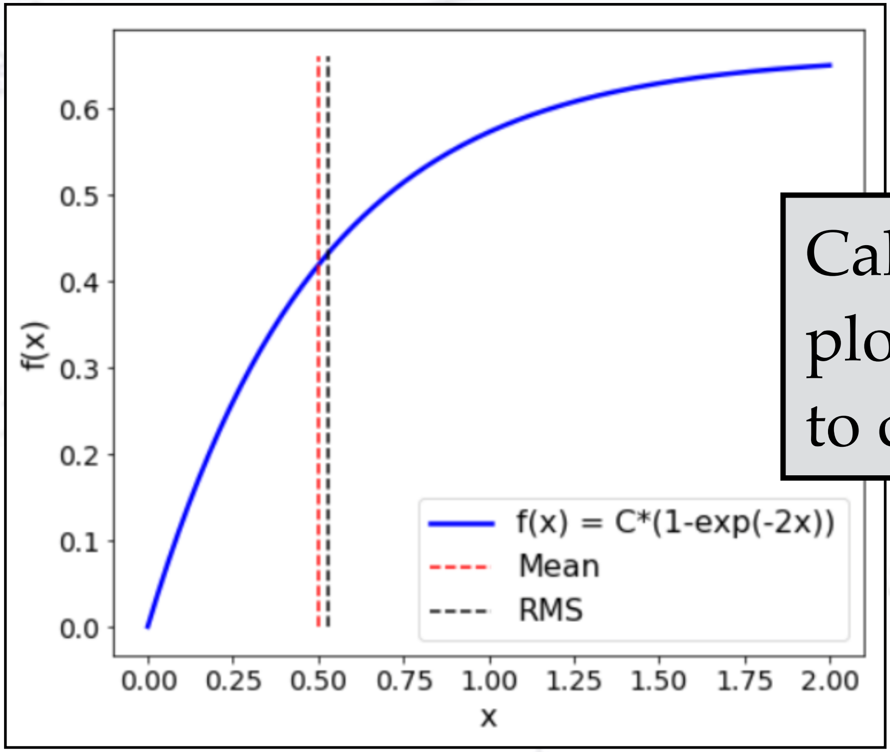

$$\mu = 1.17; \quad \sigma = 0.51$$

Calculating the mean (and even plotting it), it is always healthy to consider, if this is reasonable!

Problem 3.1

Most fitted nicely with a Chi2 fit, and most commented on low statistics.

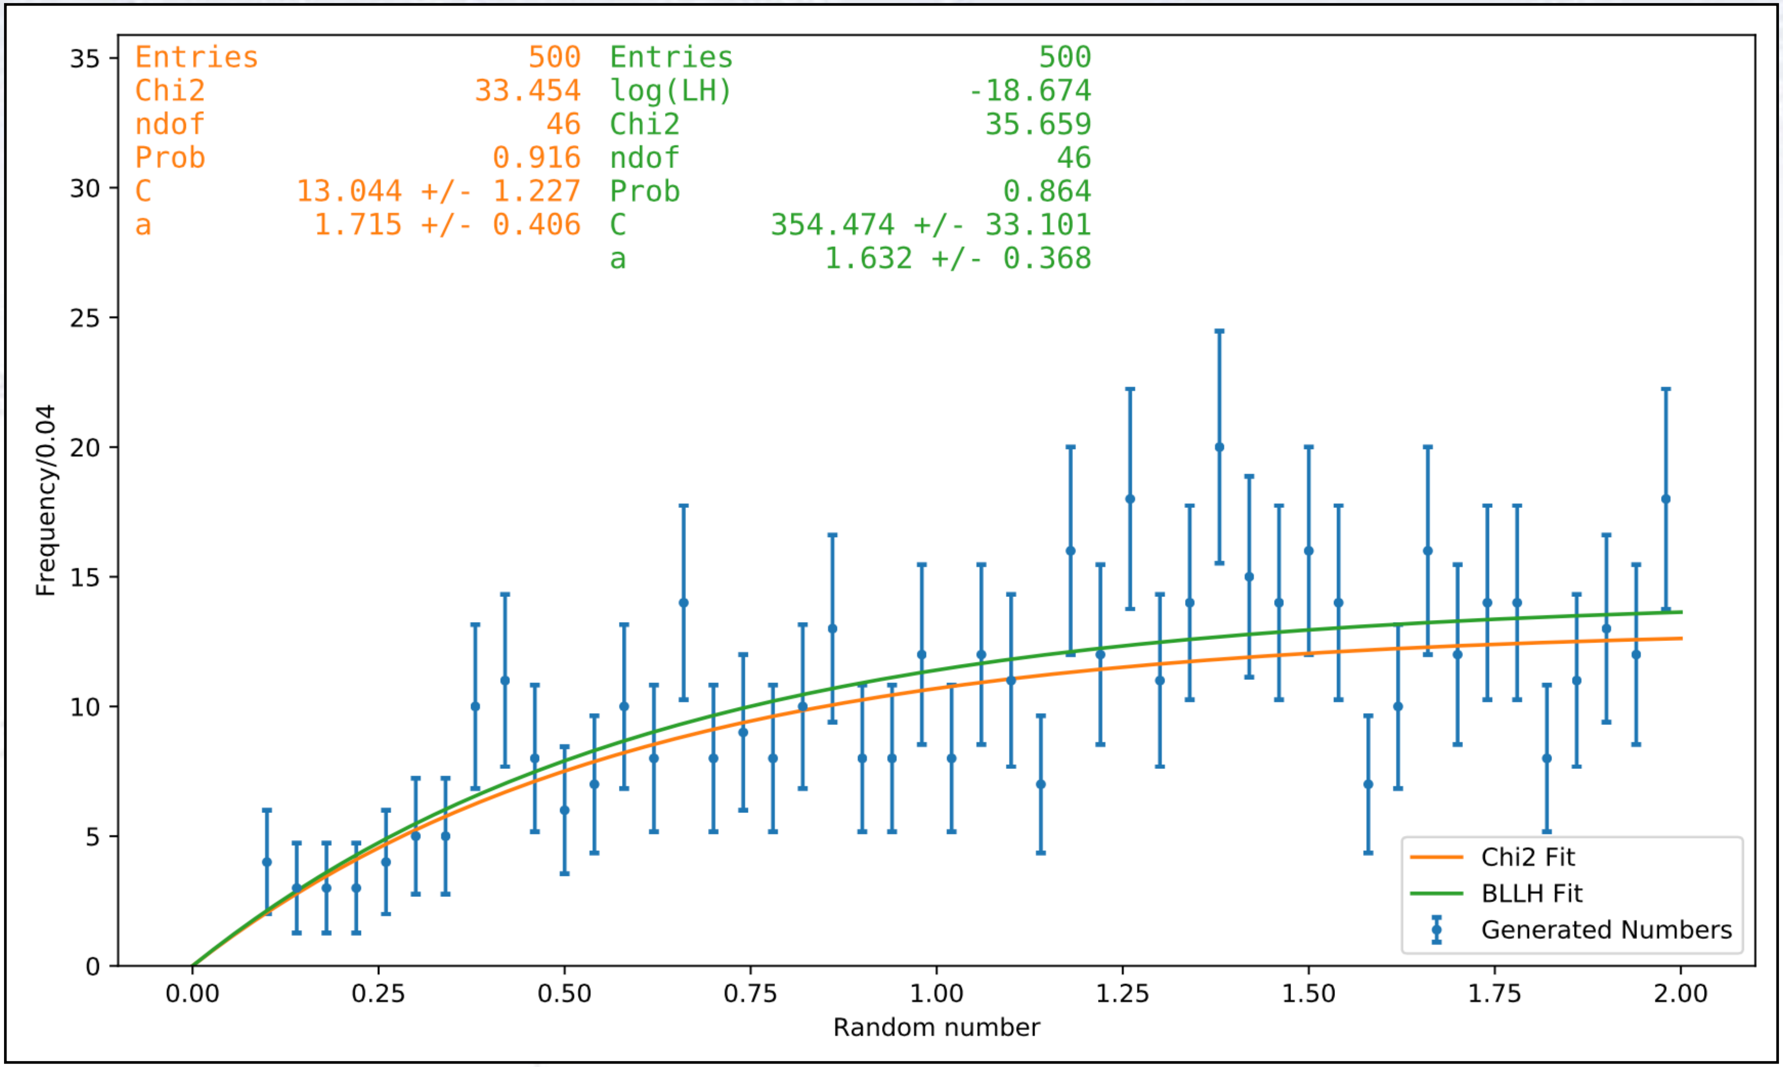

Problem 3.1

Most fitted nicely with a Chi2 fit, and most commented on low statistics.

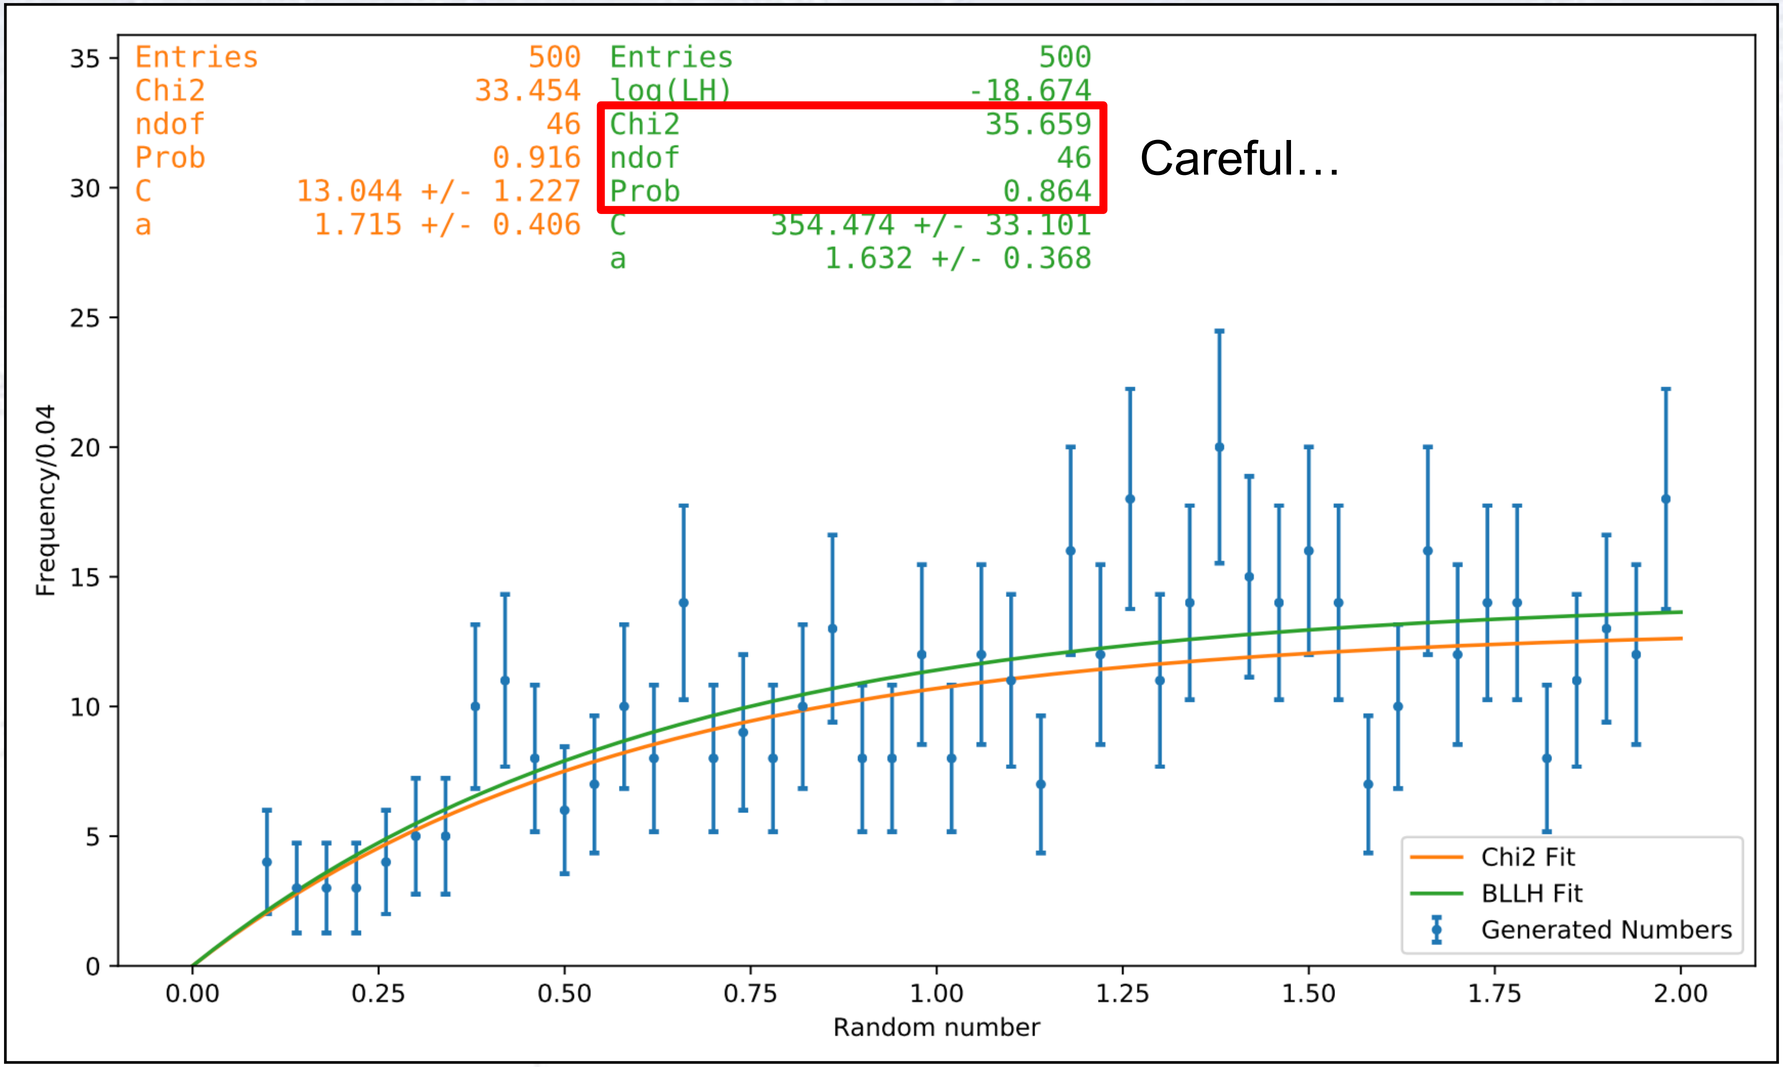

Problem 3.1

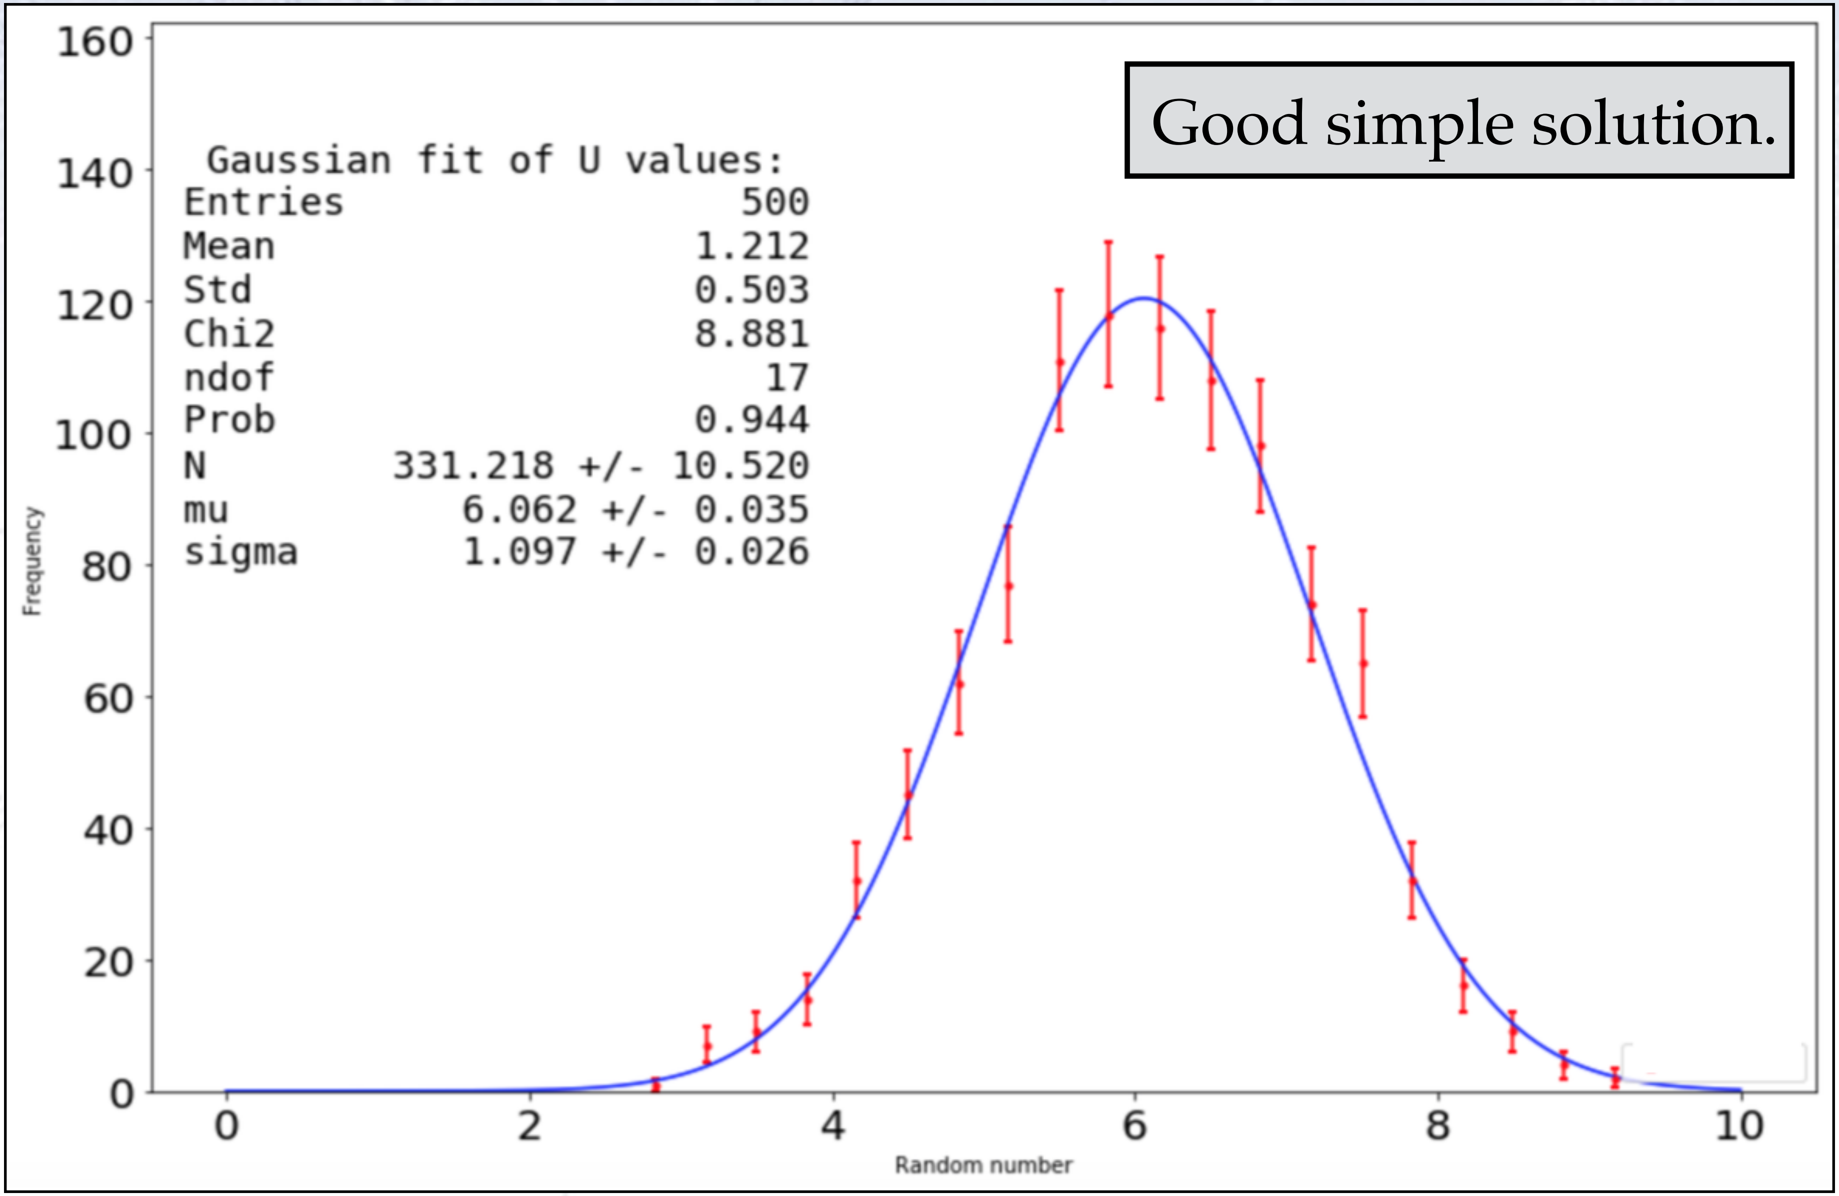

Good simple solution.

Problem 3.1

Great advanced solution.

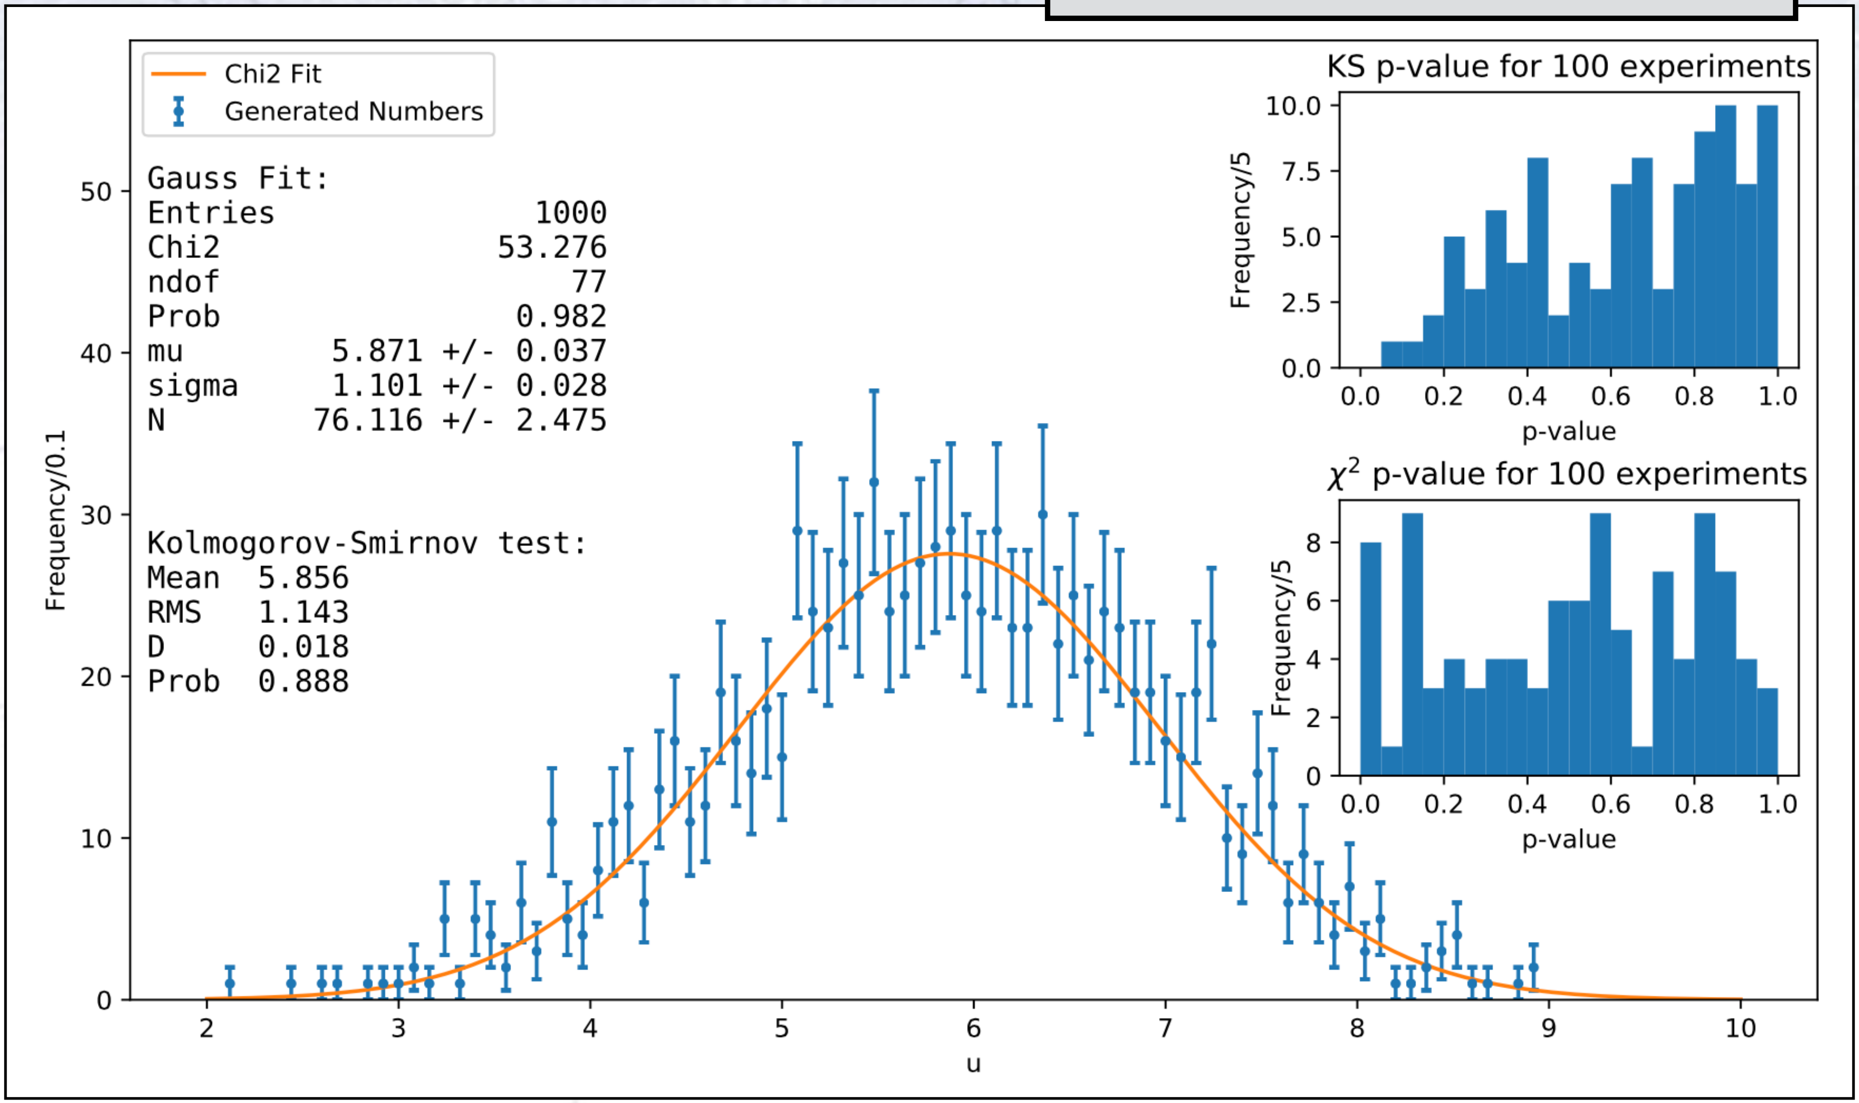

Problem 4.1

4.1 (15 points) The National UFO Reporting Center (NUFORC) has since 1974 catalogued reported UFO sightings. A subset of the data with 64719 entries containing date, time, place, shape, and duration of observation can be found at **www.nbi.dk/~petersen/data_UfoSightings.txt**.

- Plot the distribution of duration of observation, and calculate both mean and median.
- Do these durations follow the same distribution on the East and West coast?
- What is the correlation between day in the year and time of the day of observation?
- Considering only the West Coast, is the distribution of number of observations uniform over the seven week days? How about when considering only Monday to Thursday?

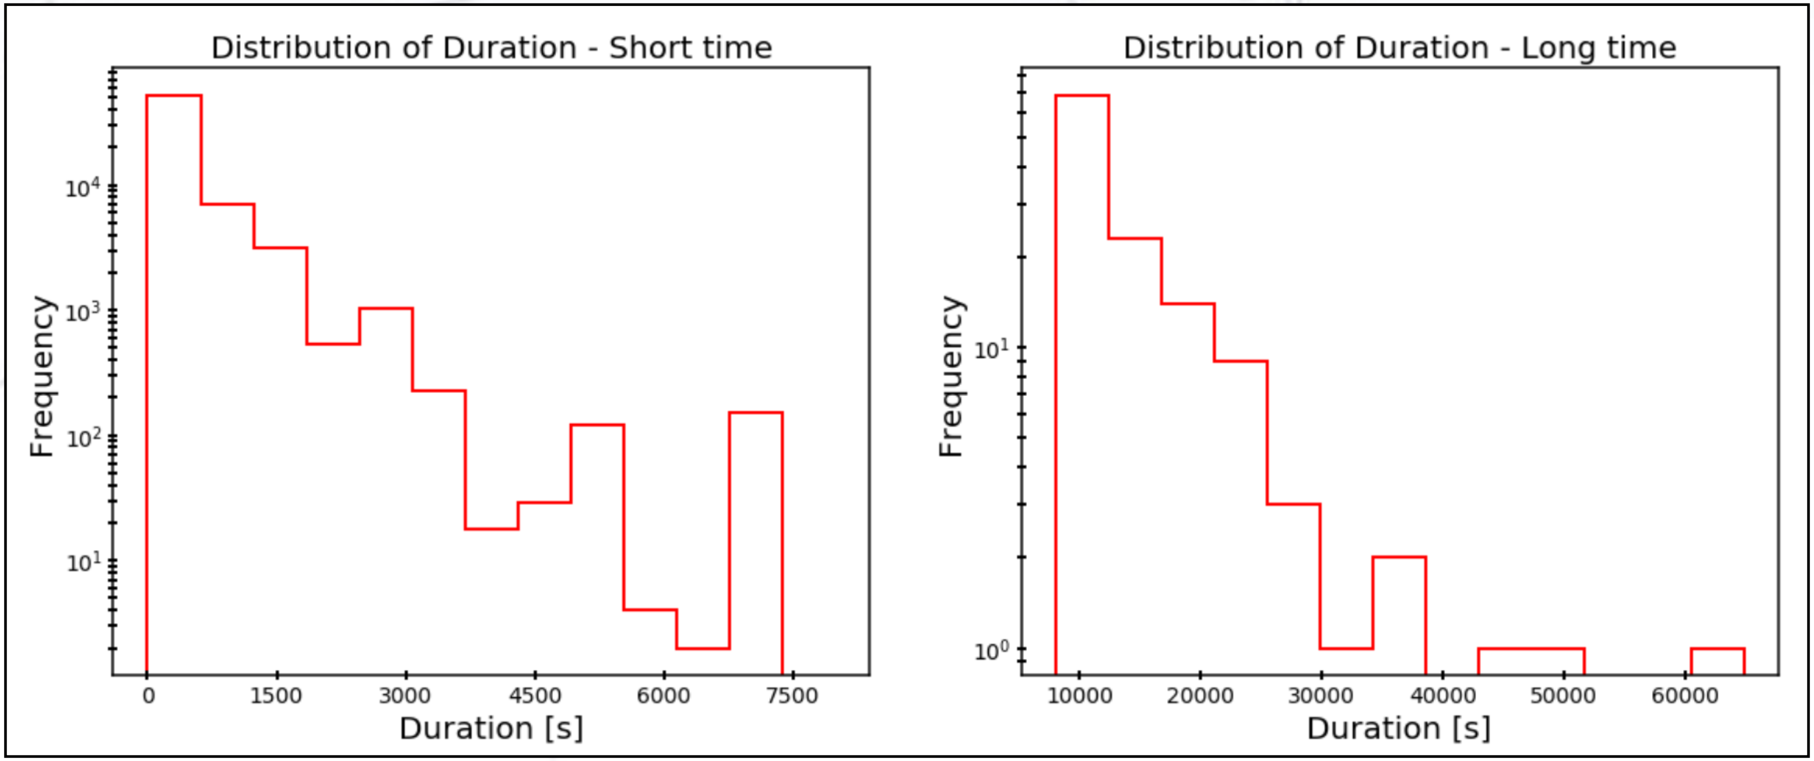

Problem 4.1

Great solution, but figure hard to read

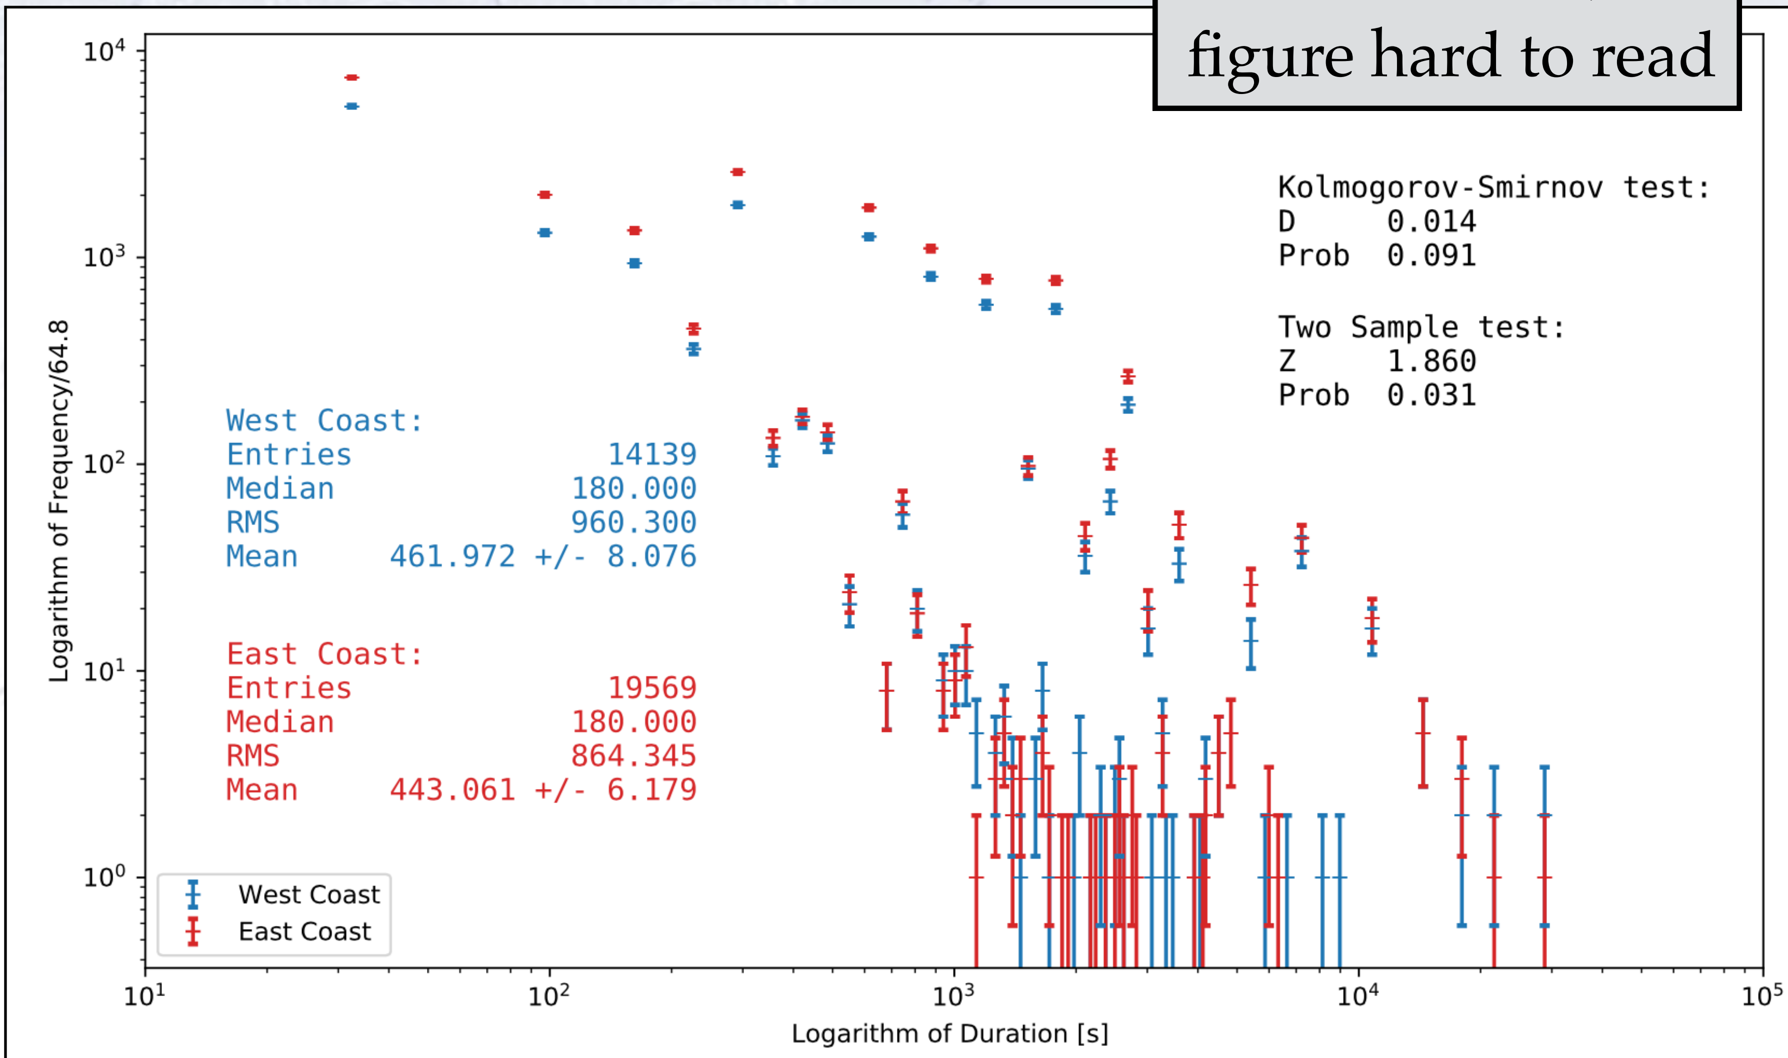

Problem 4.1

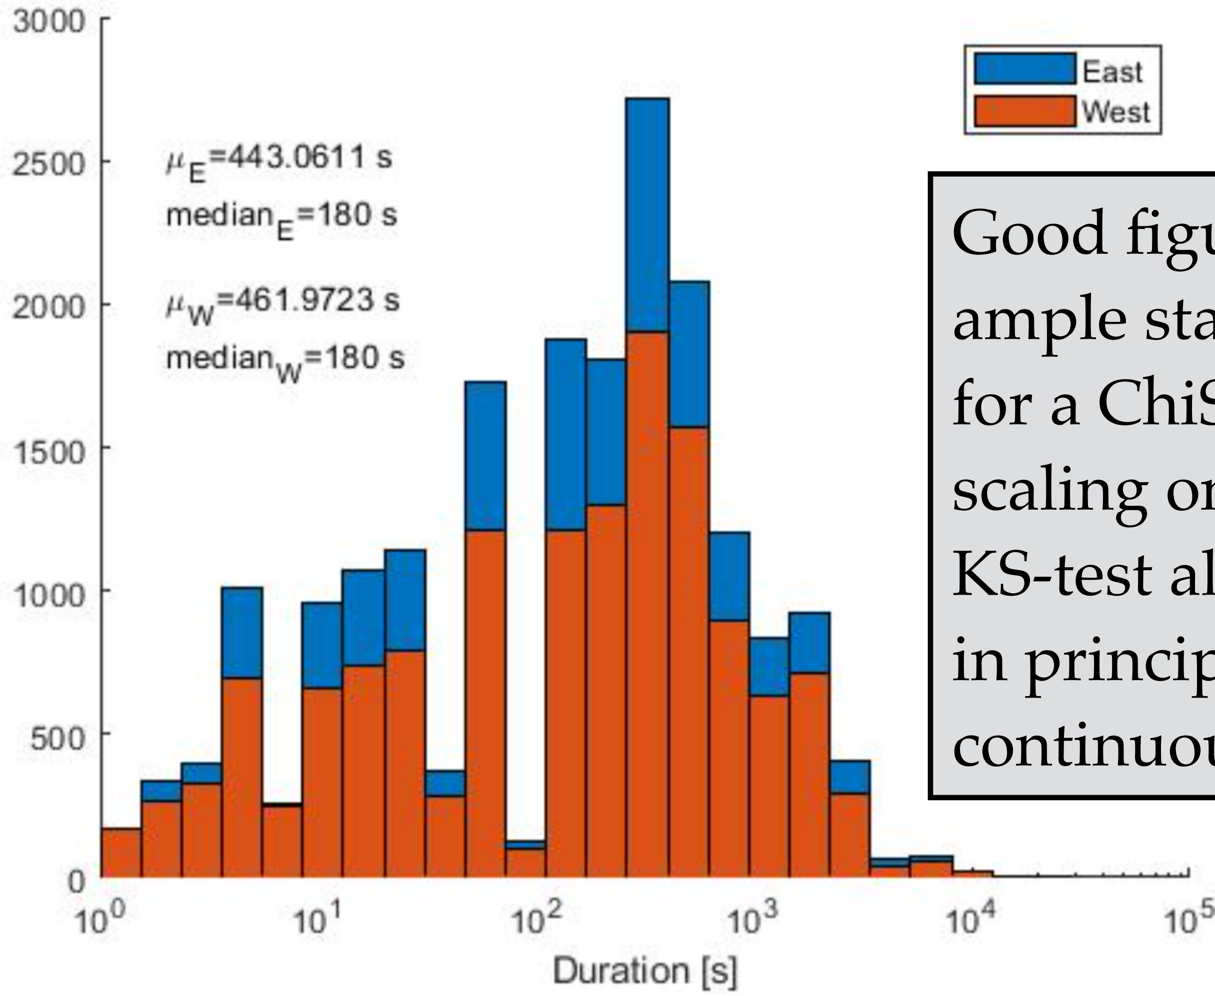

Figure 7: The two different distributions look the same to the eye.

Good figure, and there is ample statistics in each bin for a ChiSquare test (after scaling one histogram). KS-test also OK, though data in principle needs to be continuous.

Problem 4.1

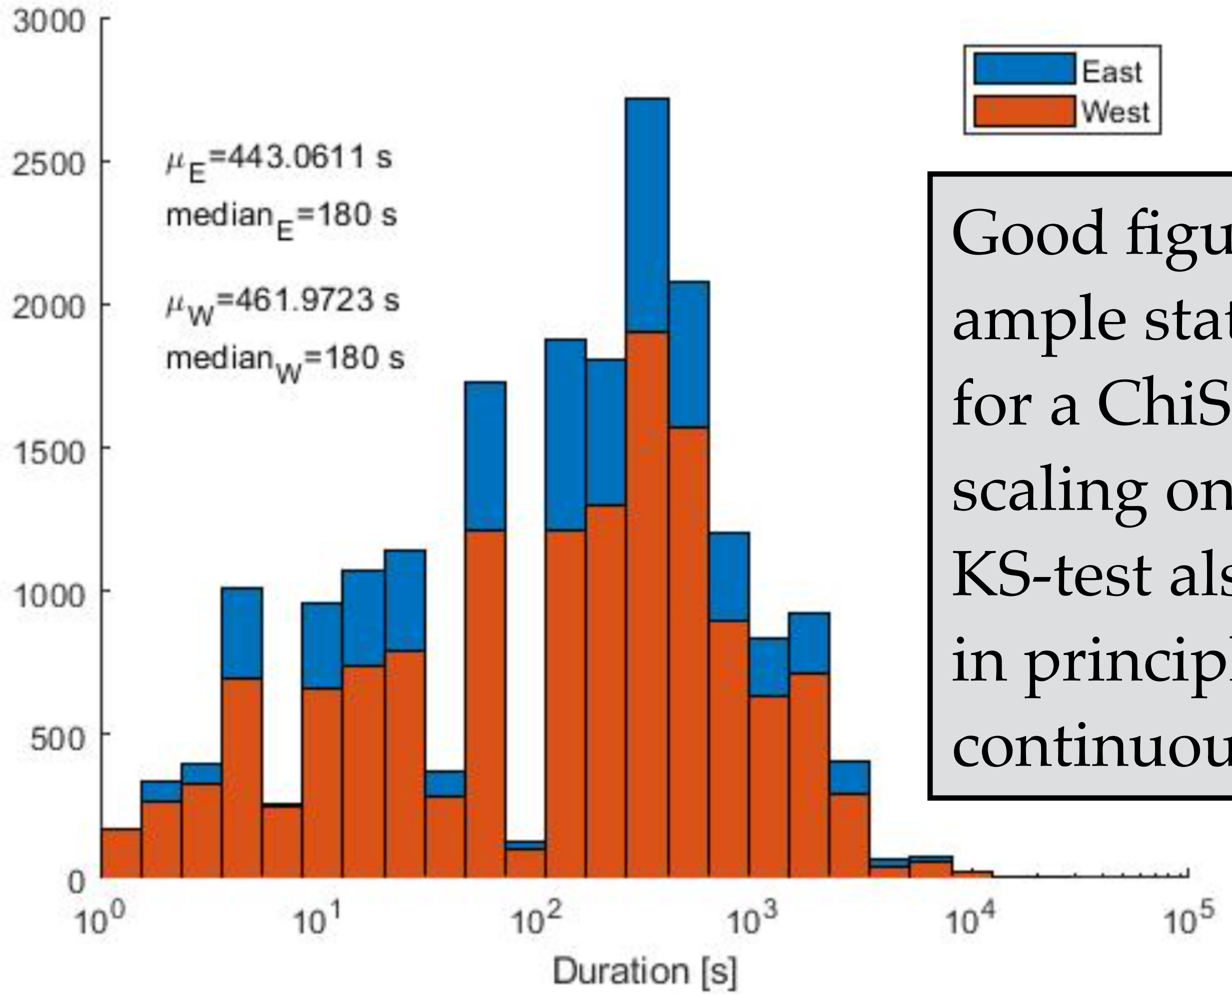

Figure 7: The two different distributions look the same to the eye.

Good figure, and there is ample statistics in each bin for a ChiSquare test (after scaling one histogram). KS-test also OK, though data in principle needs to be continuous.

Quantify, please…

Useful plot?

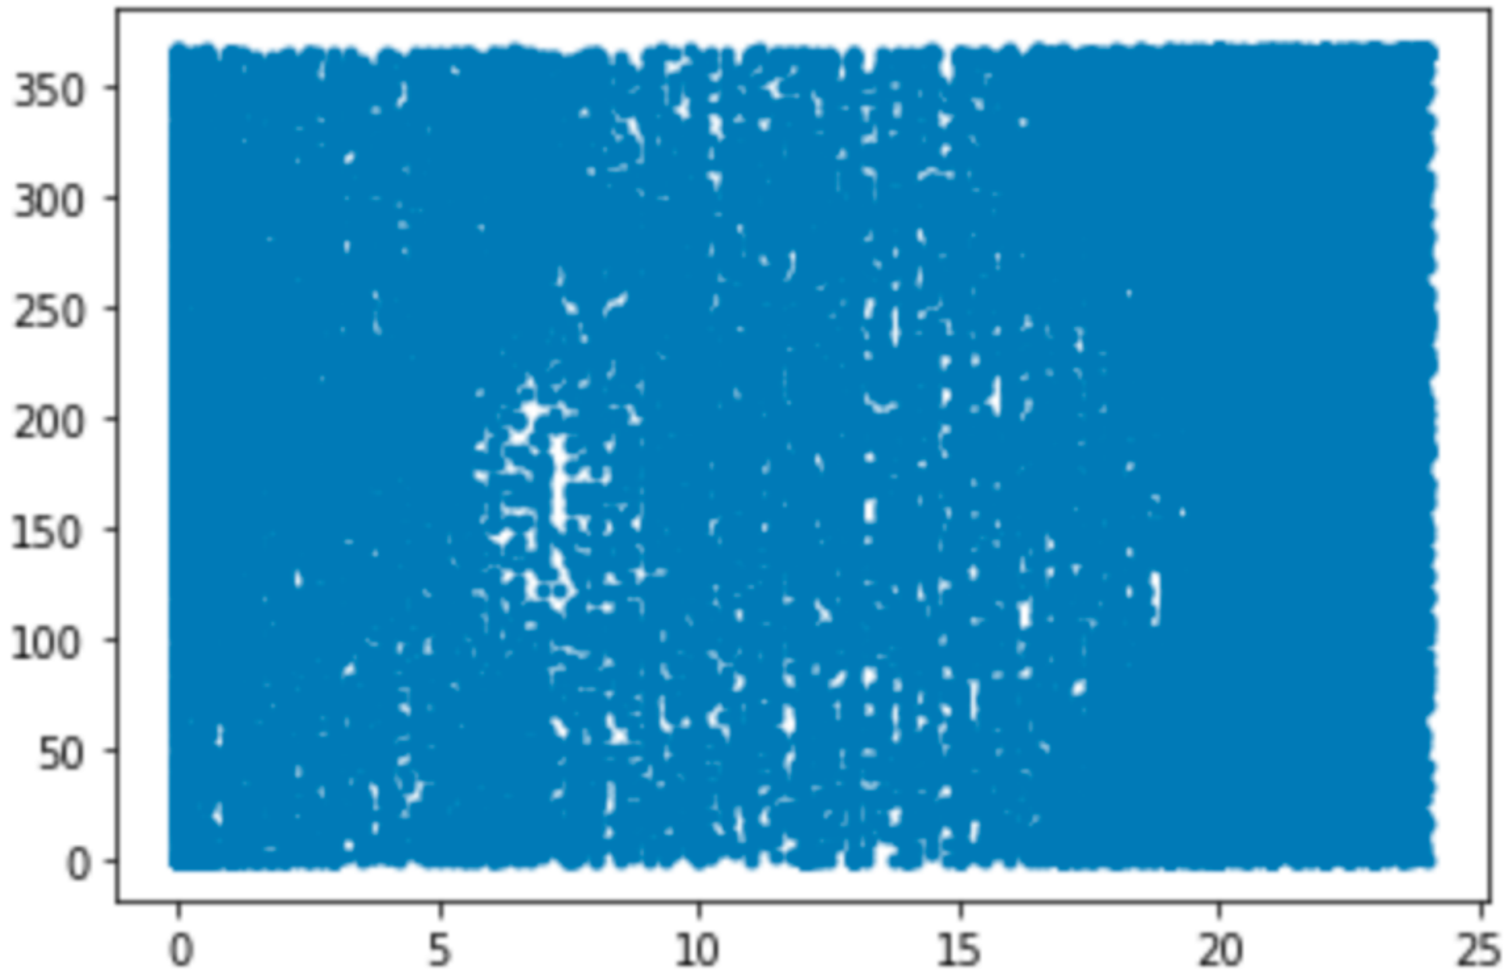

Figure 10: on y-axis we have day in year, on x-axis we have time of the day

Problem 4.1

The linear correlation is 0.024, which is small. But that does not exclude (co)relations…

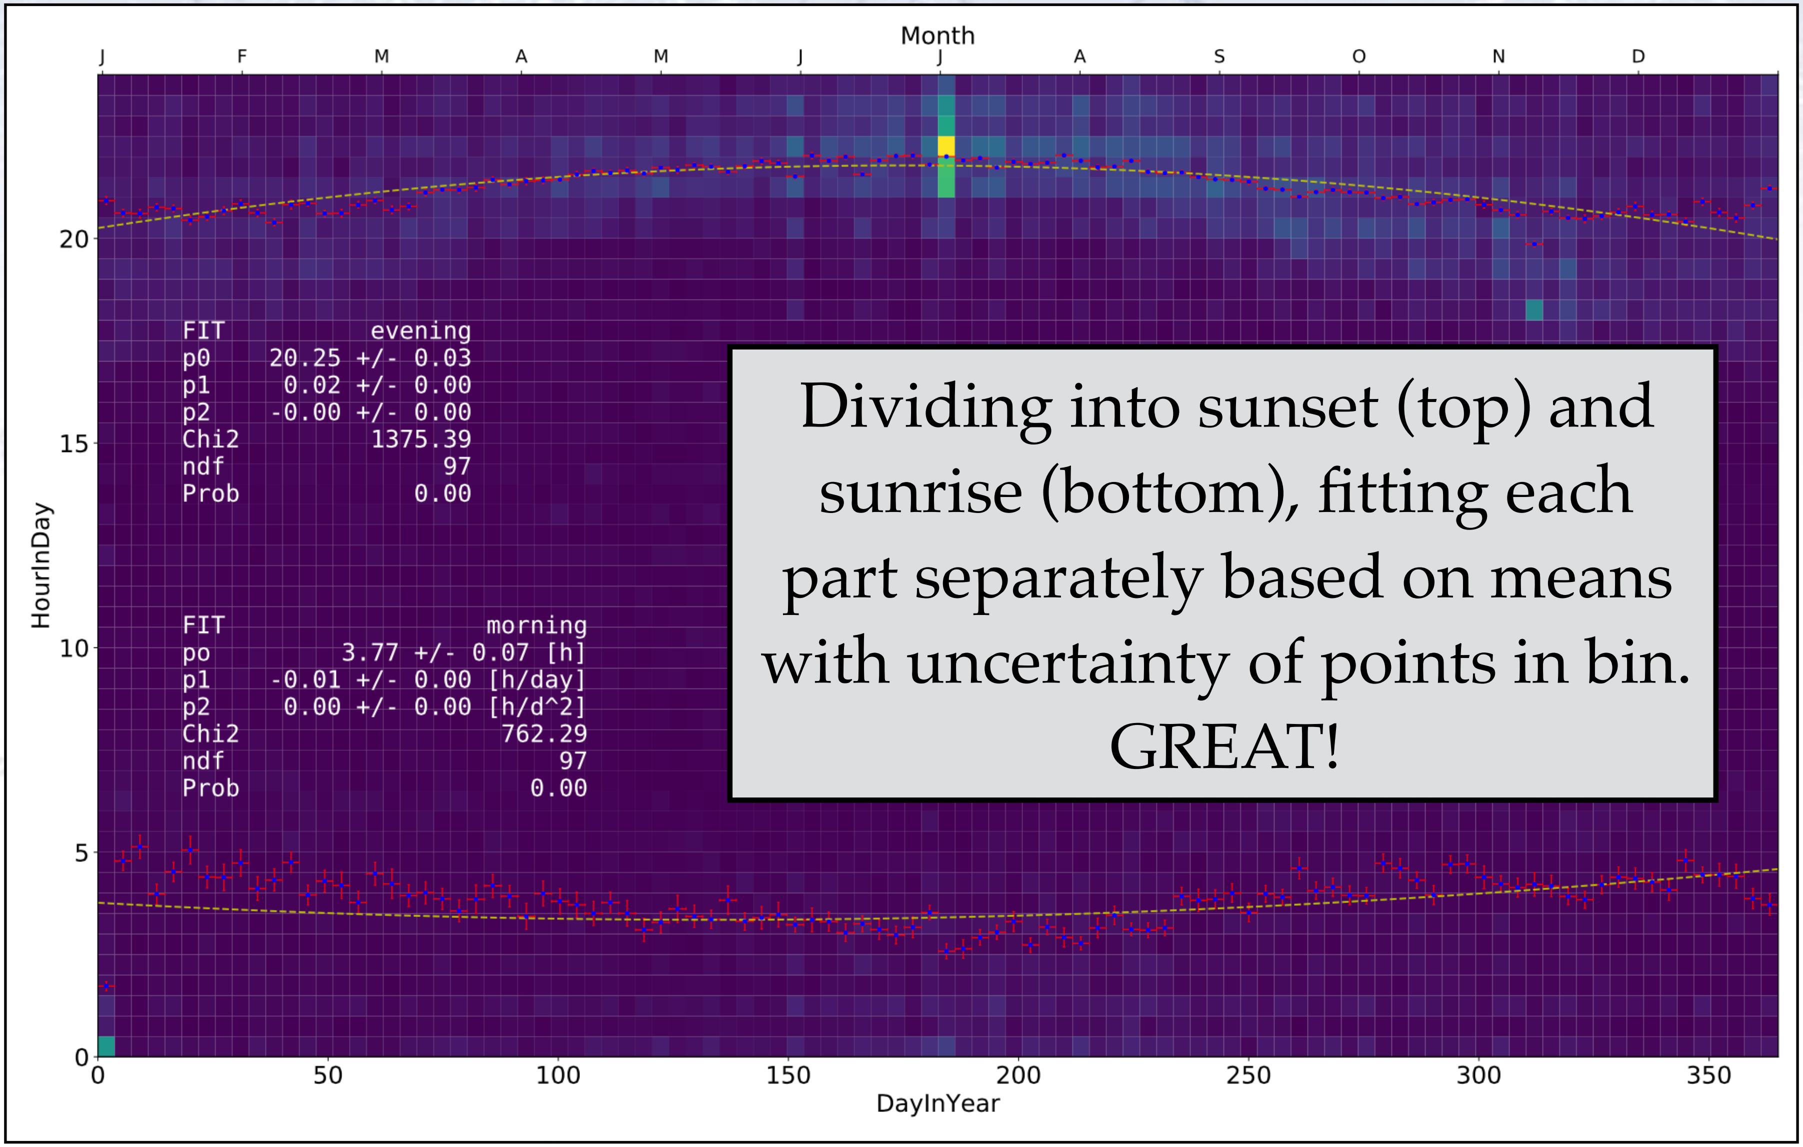

Problem 4.1

The linear correlation is 0.024, which is small. But that does not exclude (co)relations…

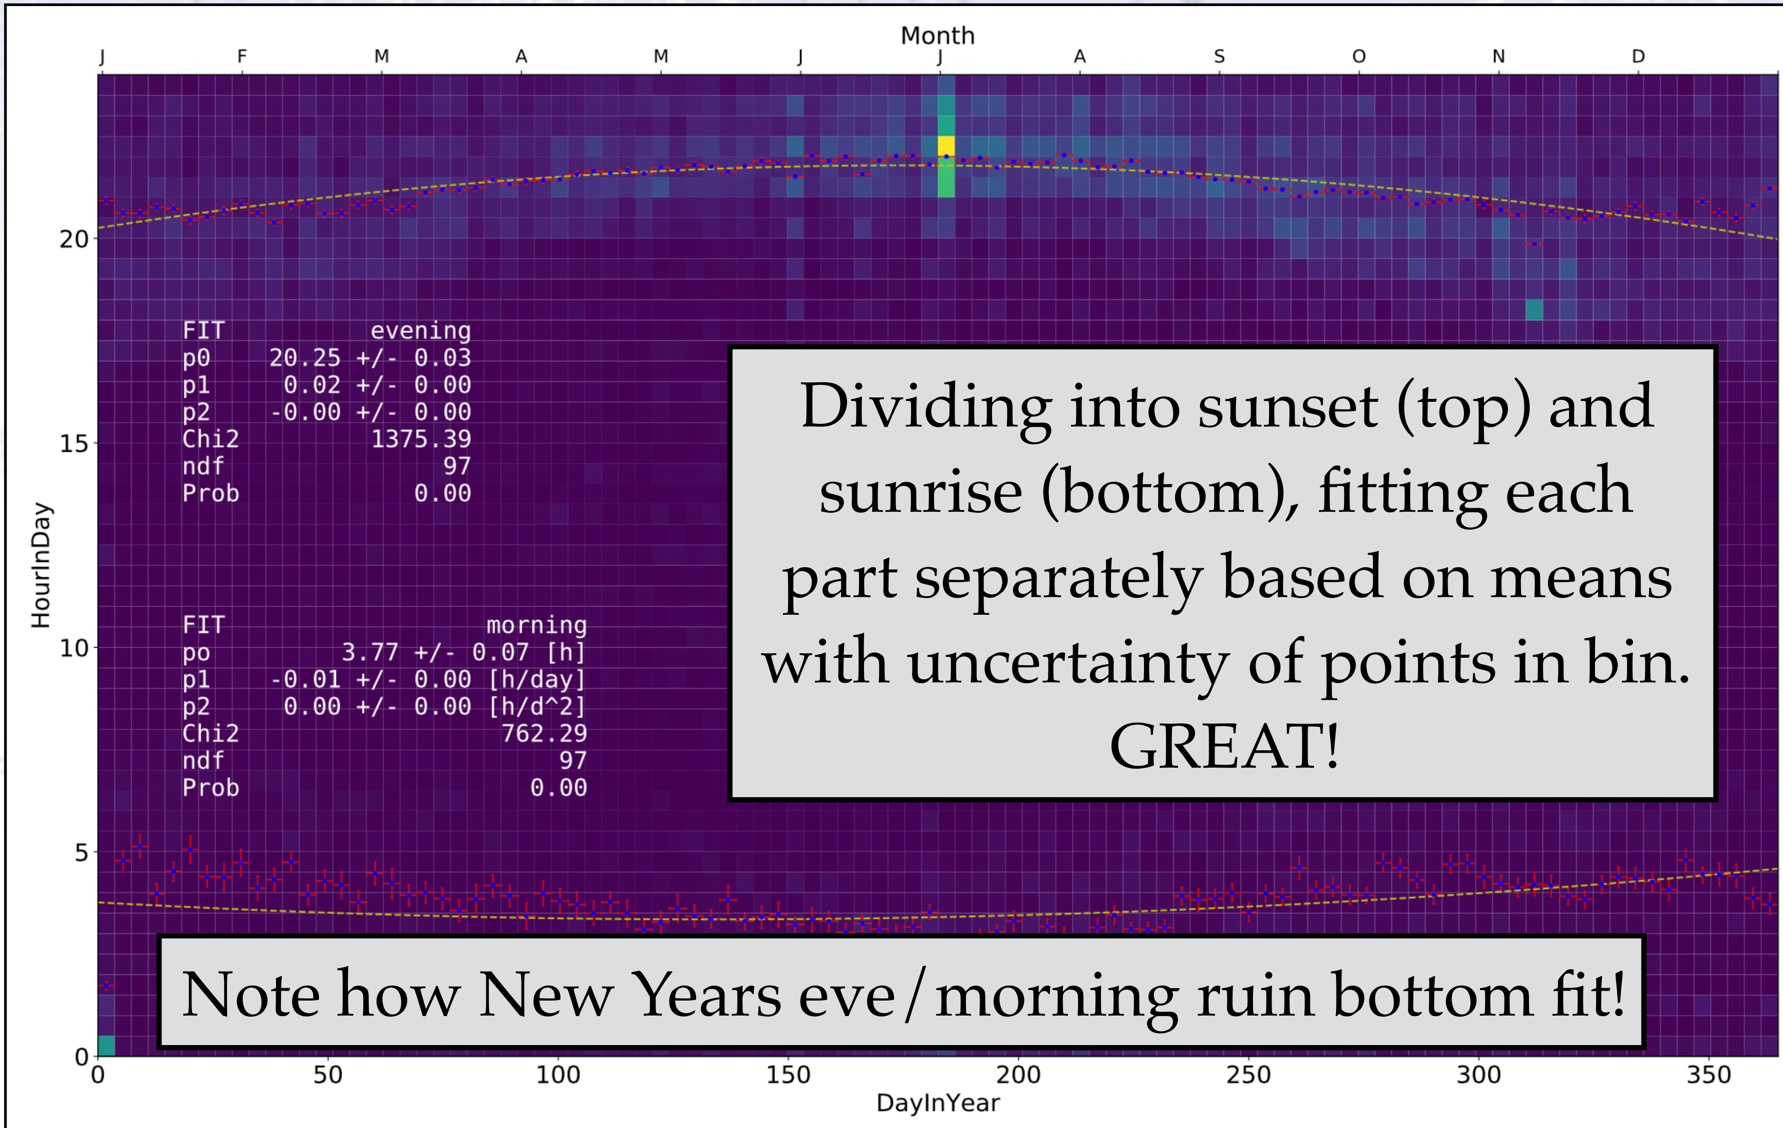

Problem 4.1

The weekly distribution is clearly not flat. Considering Monday-Thursday, it is on the verge of being it. A simple Chi2 fit is the solution… p-value = 0.022

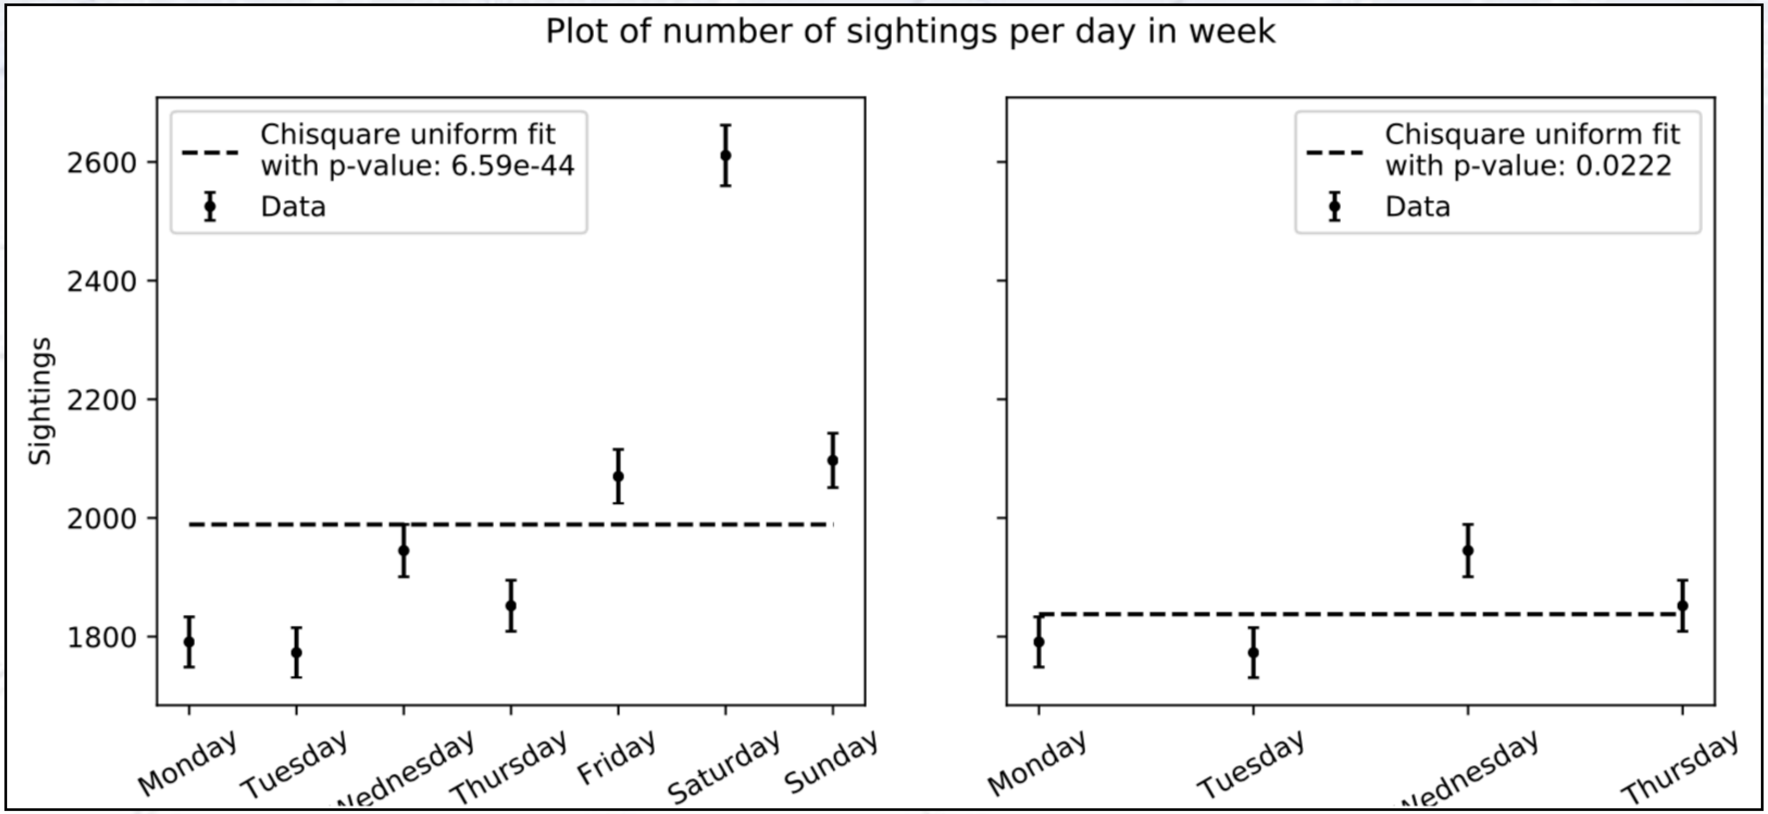

Typical problems

Problem 4.1:

1. **Not enough error bars on fitted data**. Really tough to judge ChiSquare without the whole picture!
2. Some very advanced (although not necessarily fruitful) fitting to the 2d histograms - nice!
3. Lots of weird p-values
4. Lots of folks made one draw from a uniform distribution and then argued with a single Pearson ChiSquare for the constancy of the observations without commentary on the method. Very few did the fit for a constant value.
5. Uninformative plots
6. Some folks accepted hypothesis outright, instead of "rejecting the null" — didn't penalise this.
7. Not enough plots to argue from, in general.
8. If they plotted the 2d histogram (or even scatterplot), then they usually got the pcorr ok.

Problem 4.2

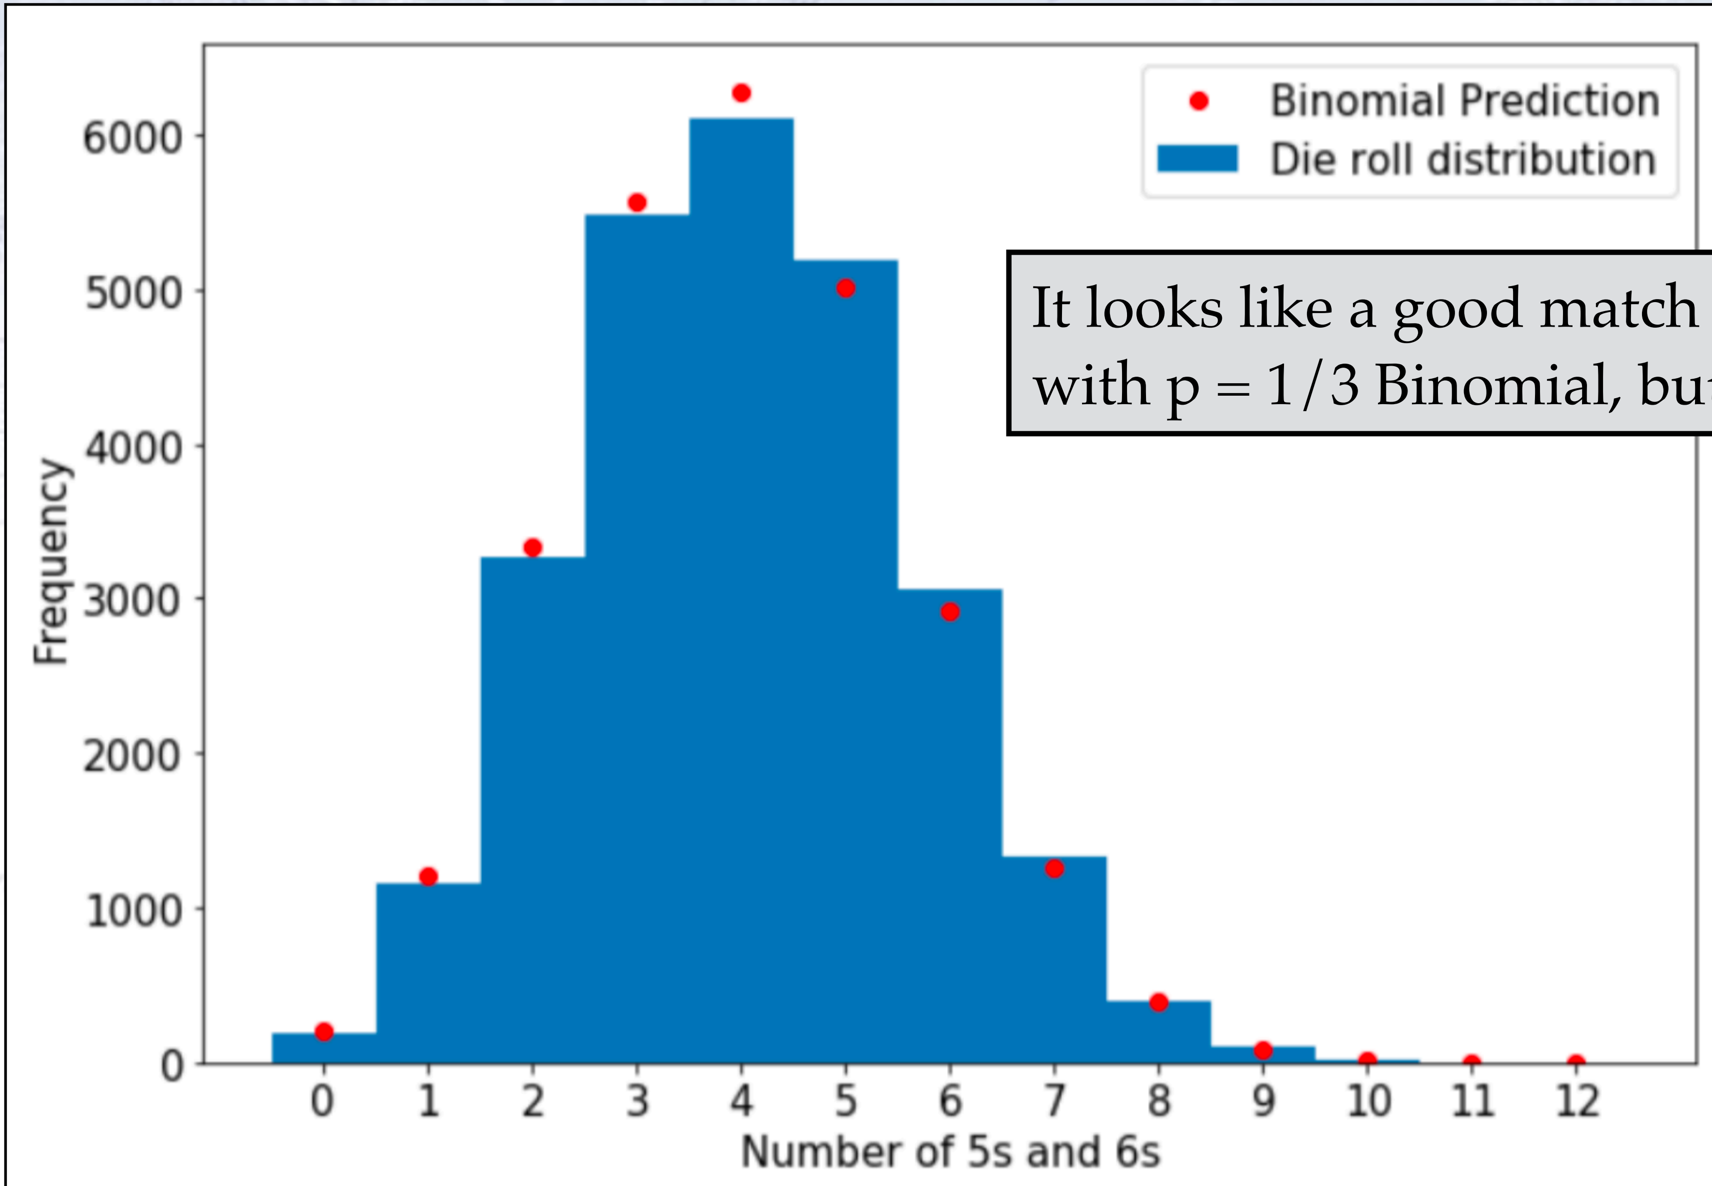

It looks like a good match with p = 1/3 Binomial, but…

Problem 4.2

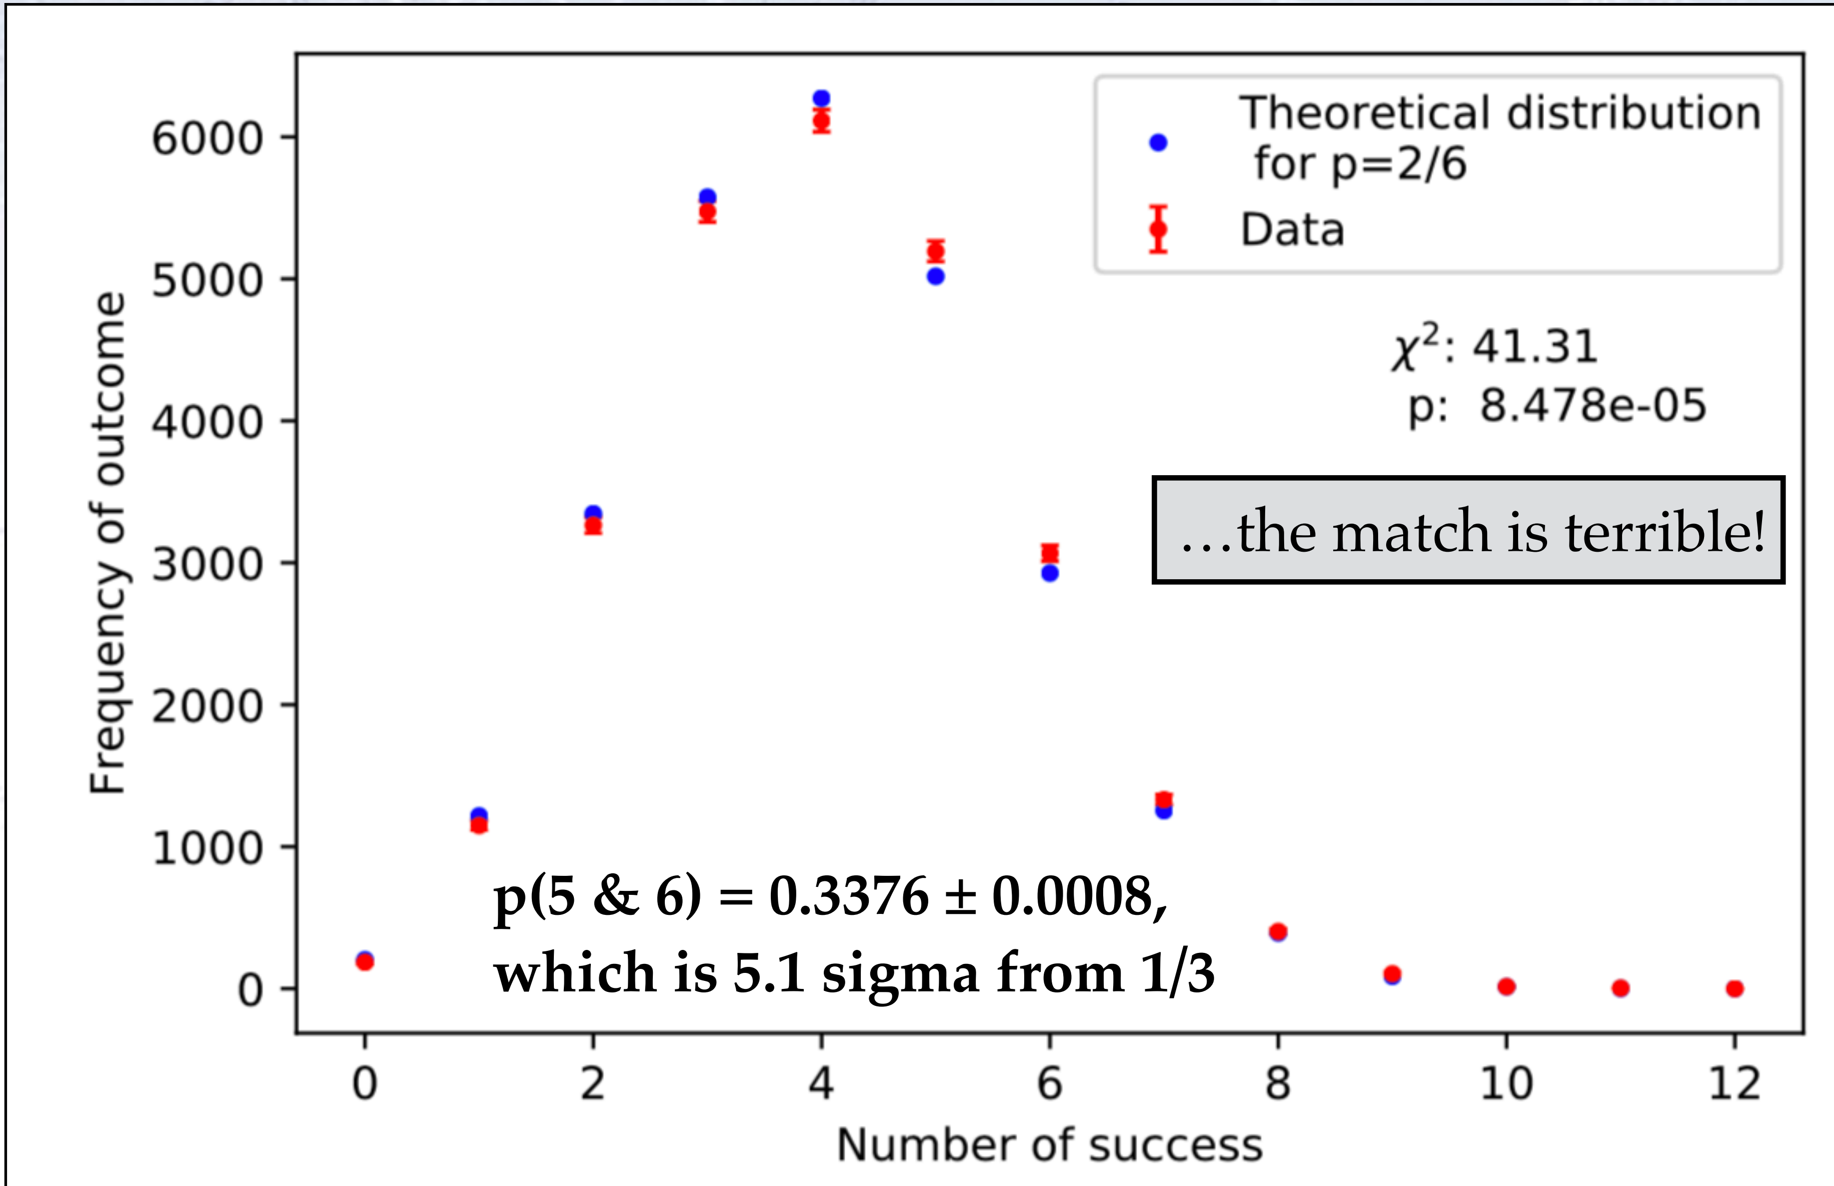
Theoretical distribution for p=2/6
Data
χ^2: 41.31
p: 8.478e-05
…the match is terrible!
p(5 & 6) = 0.3376 ± 0.0008,
which is 5.1 sigma from 1/3
Frequency of outcome
Number of success
0
1000
2000
3000
4000
5000
6000
0
2
4
6
8
10
12

Problem 4.2

4.2 (13 points) To test the fairness of dice, you roll 12 dices and count the number 5s and 6s (N_{56}). Repeating this many times yielded the following result:

Number of 5s & 6s	0	1	2	3	4	5	6	7	8	9	10	11	12
Observed frequency	185	1149	3265	5475	6114	5194	3067	1331	403	105	14	4	0

- What distribution should the number of 5s and 6s follow?
- Compare the data with the expected distribution. Does this hypothesis match the data well?
- Fit the data and test if alternative hypotheses match the data better. Also, determine the probability for a 5 or a 6, and decide if the dice are consistent with being fair.

Subpart 1

The chances of landing a 5 or a 6 on a die should be $p = \frac{2}{6}$. As we can assume they are independent, and we have 12 trials, the should be binomially distributed with siza parameter $N = 12$ and probability parameter $p = \frac{2}{6}$.

Subpart 2

As we knew the uncertainties on the counts, and the expected values $f(k) = N \cdot Binomial(k,12,2/6)$, where N is the total number of trails, we could just do a simple χ^2 test, which returned a probability of 0.036%, leting us reject that the data follows this pattern.

Subpart 3

Here, we did a simple fit, where we allowed adjustment of the probability, as seen on figure 14. It found the probability to be 0.337626 ± 0.000845, with a fit probability of 36.88%. Notice that the resulting probability is many sigmas away from 0.333333...., which leads us to conclude that the dice are not completely fair, with a -slightly- higher chance of a 5 or a 6.

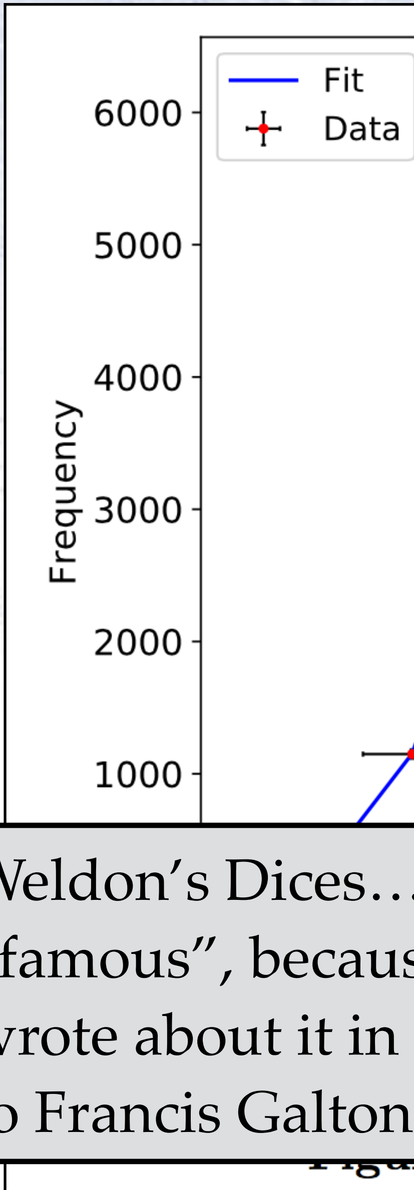

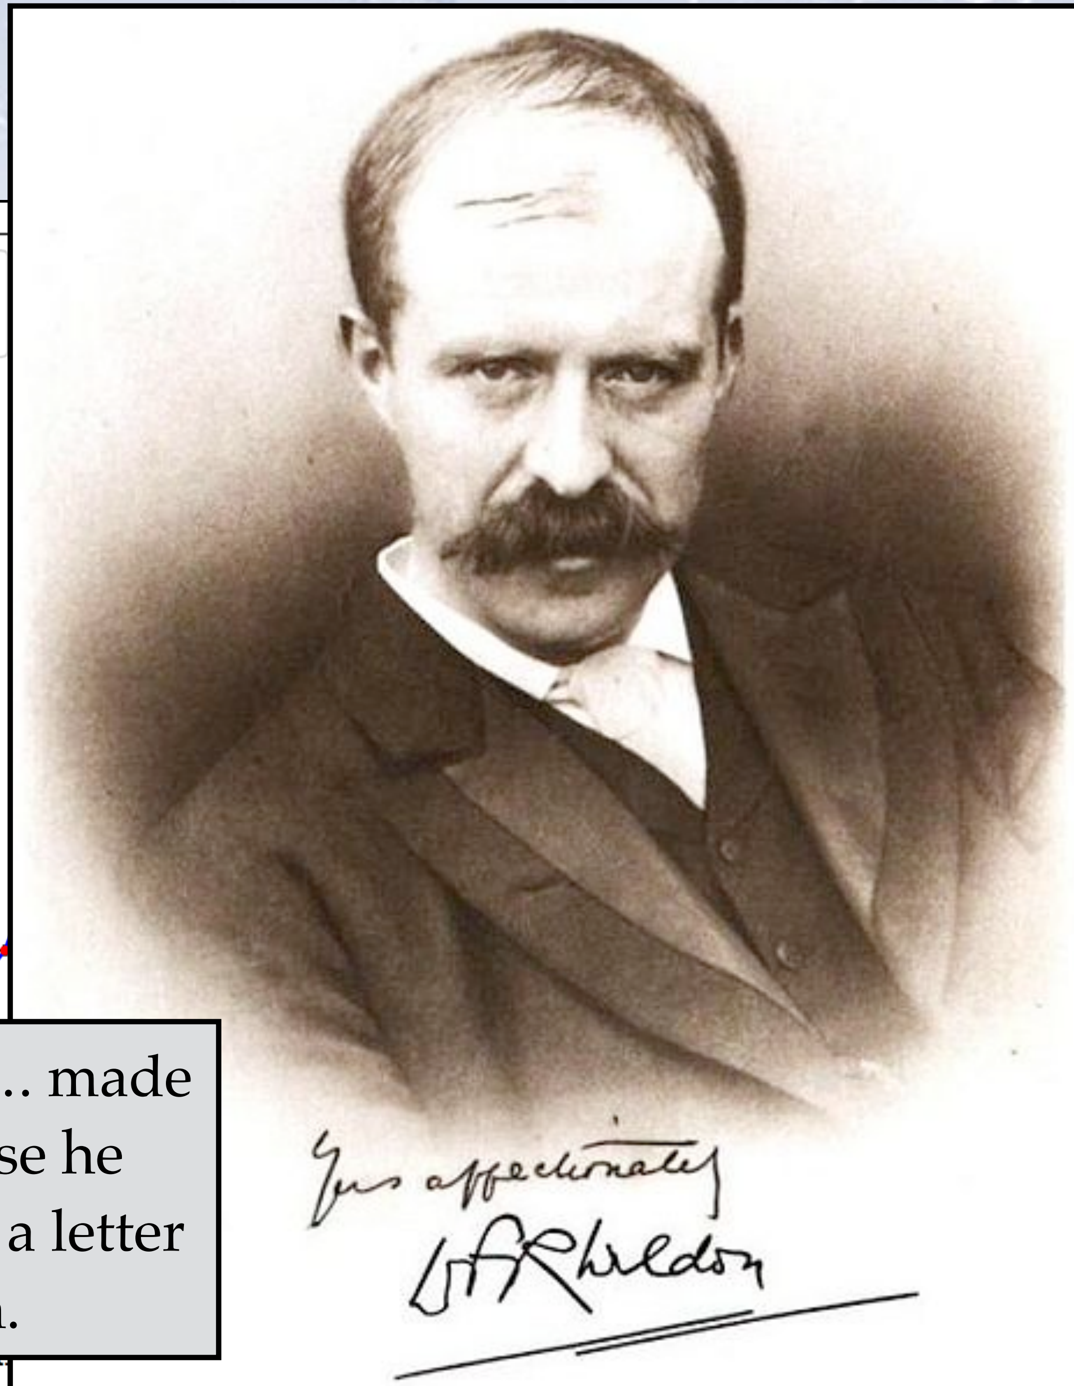

l" / weird
ution
tiful) fit?

Weldon's Dices…. made "famous", because he wrote about it in a letter to Francis Galton.

Typical problems

Problem 4.2:

1. Many people had n and N as free parameters next to p in the binomial fit, while these should be fixed because the number of dices and throws is known. This gave them a different value for p (0.335 +/- 0.005), which is less in tension with 1/3.
2. Many didn't quote errors on p, and just concluded that it was similar to 1/3.
3. A lot of people that used the KS test got a wrong value. Many got a p-value of 1. Not sure what went wrong there.
4. I also found that a lot of people interpreted the question slightly differently. For the second point, they just calculated whether a binomial in general was a good fit, not one specifically with p=1/3. And for the third point, they tested a Gaussian and Poisson as alternative hypotheses. Some people also didn't understand what was meant by "determine the probability for a 5 or a 6".
5. Thinking Poisson because of 26306 repetitions!

Problem 5.1

5.1 (19 points) Gamma ray spectra are used to identify radioactive isotopes in material from the very sharp peaks they produce. The file **www.nbi.dk/~petersen/data_GammaSpectrum.txt** contains 47173 measurements obtained from uranium ore, where each number refers to a *channel number* in the detector, which can be translated into an *energy* (roughly in the 0-1000 keV range).

- The lead isotope ^{214}Pb produces three known, low energy gamma ray peaks at energies E: 242 keV (7.4 %), 295 keV (19.3 %), and 352 keV (37.6 %). Fit these three ^{214}Pb peaks.
- For these three peaks, compare the *relative* distance $r = (E_3 - E_2)/(E_2 - E_1)$ (in channel number) with the corresponding tabular value (in energy). Does the relative distance match?
- The more energetic bismuth ^{214}Bi gamma rays produce peaks at 609 keV (46.1 %) and 1120 keV (15.1 %). From peaks of your choice, determine the energy scale (i.e. how to determine energy from channel number) and test if it is linear.
- Is the energy resolution (i.e. peak width) constant or does it change with energy?
- A theoretical calculation predicts a (small) peak in the spectrum in the range 700-800 keV. Does the data support that prediction? And if so, at what energy?
- Do you find any other peaks or features in the spectrum? If so, quantify your findings.

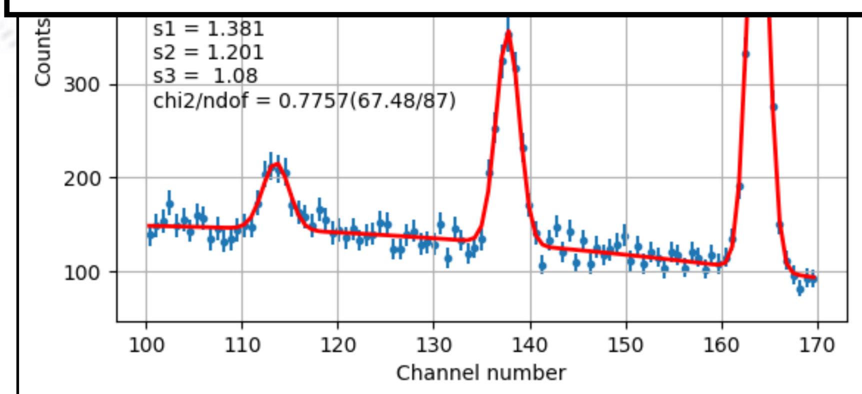

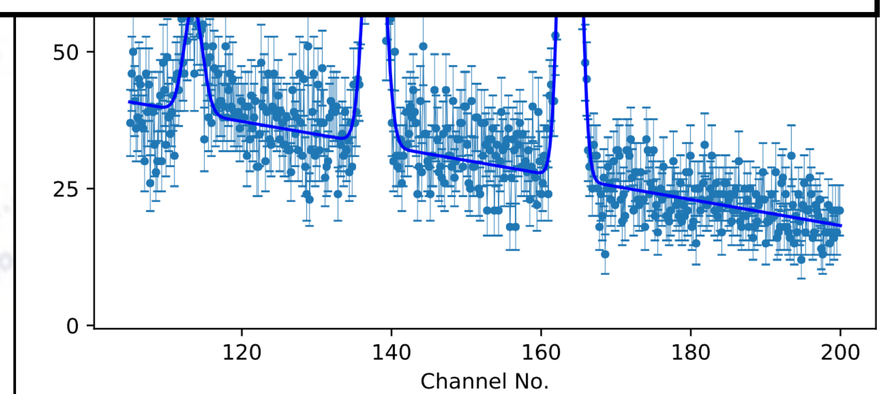

Problem 5.1

While one can try to subtract the background, it is best simply to fit for it.
Advanced fits like these are nice, but it can be done with three separate fits, as the peaks a very independent.

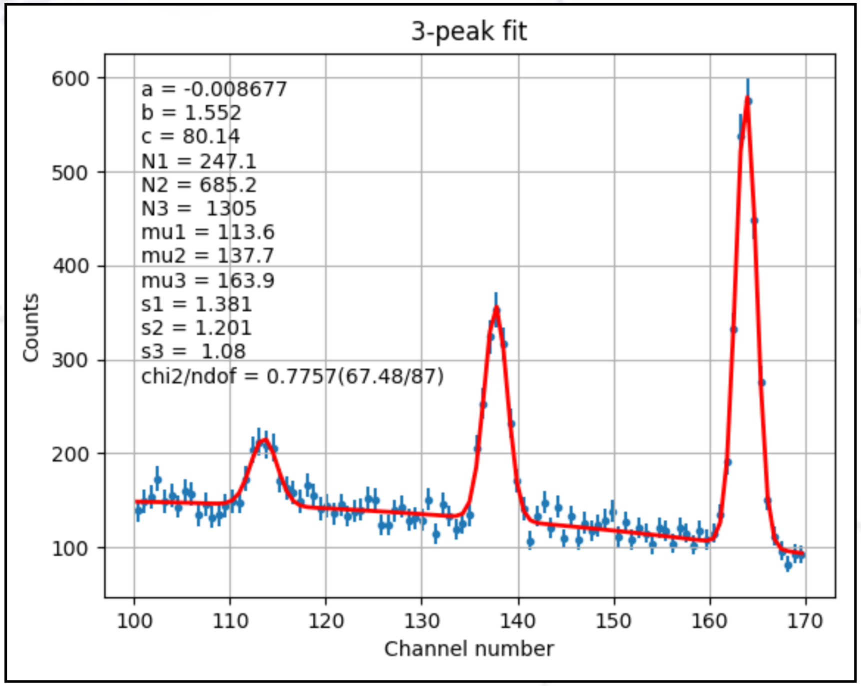

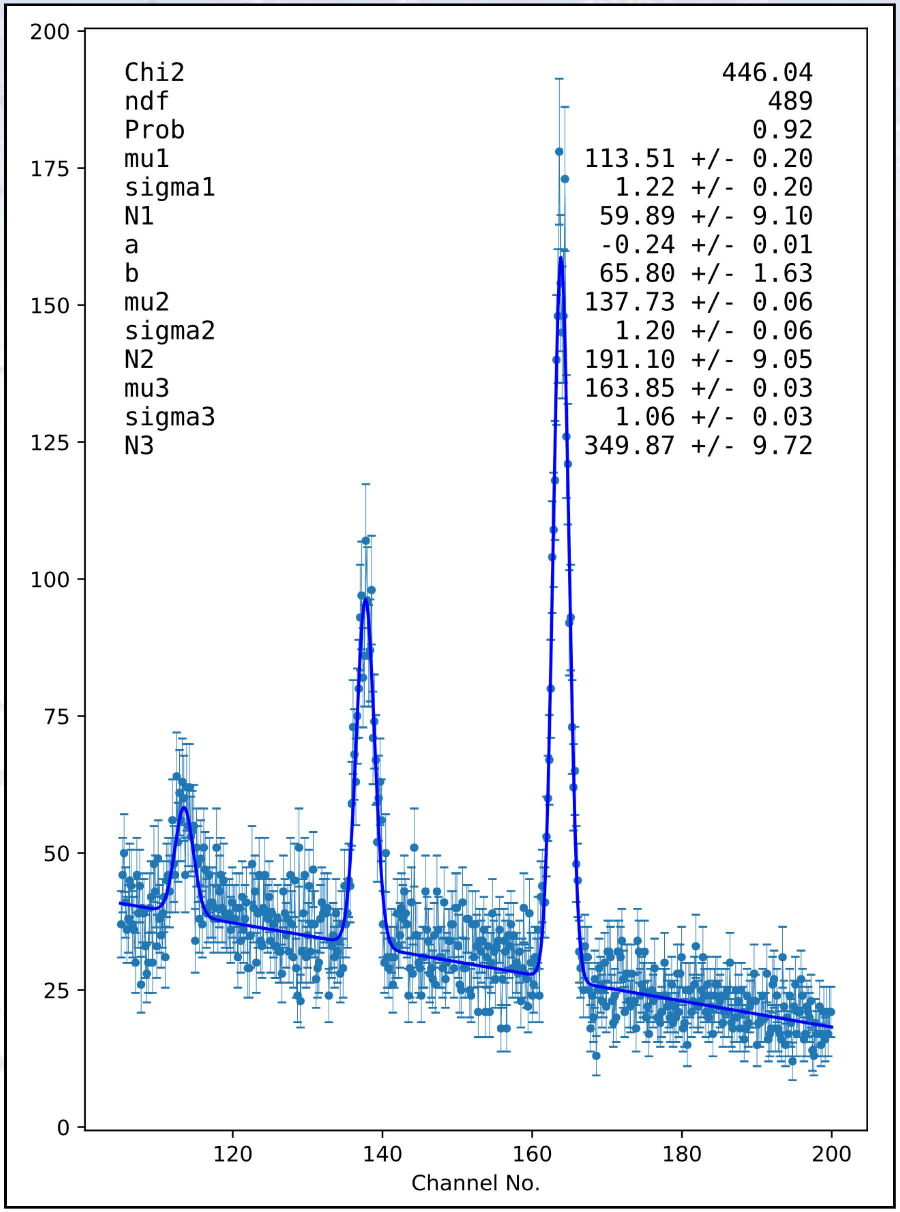

Problem 5.1 - 3 Gaussians…

Why will a fit using 3 Gaussians not fit the third peak?

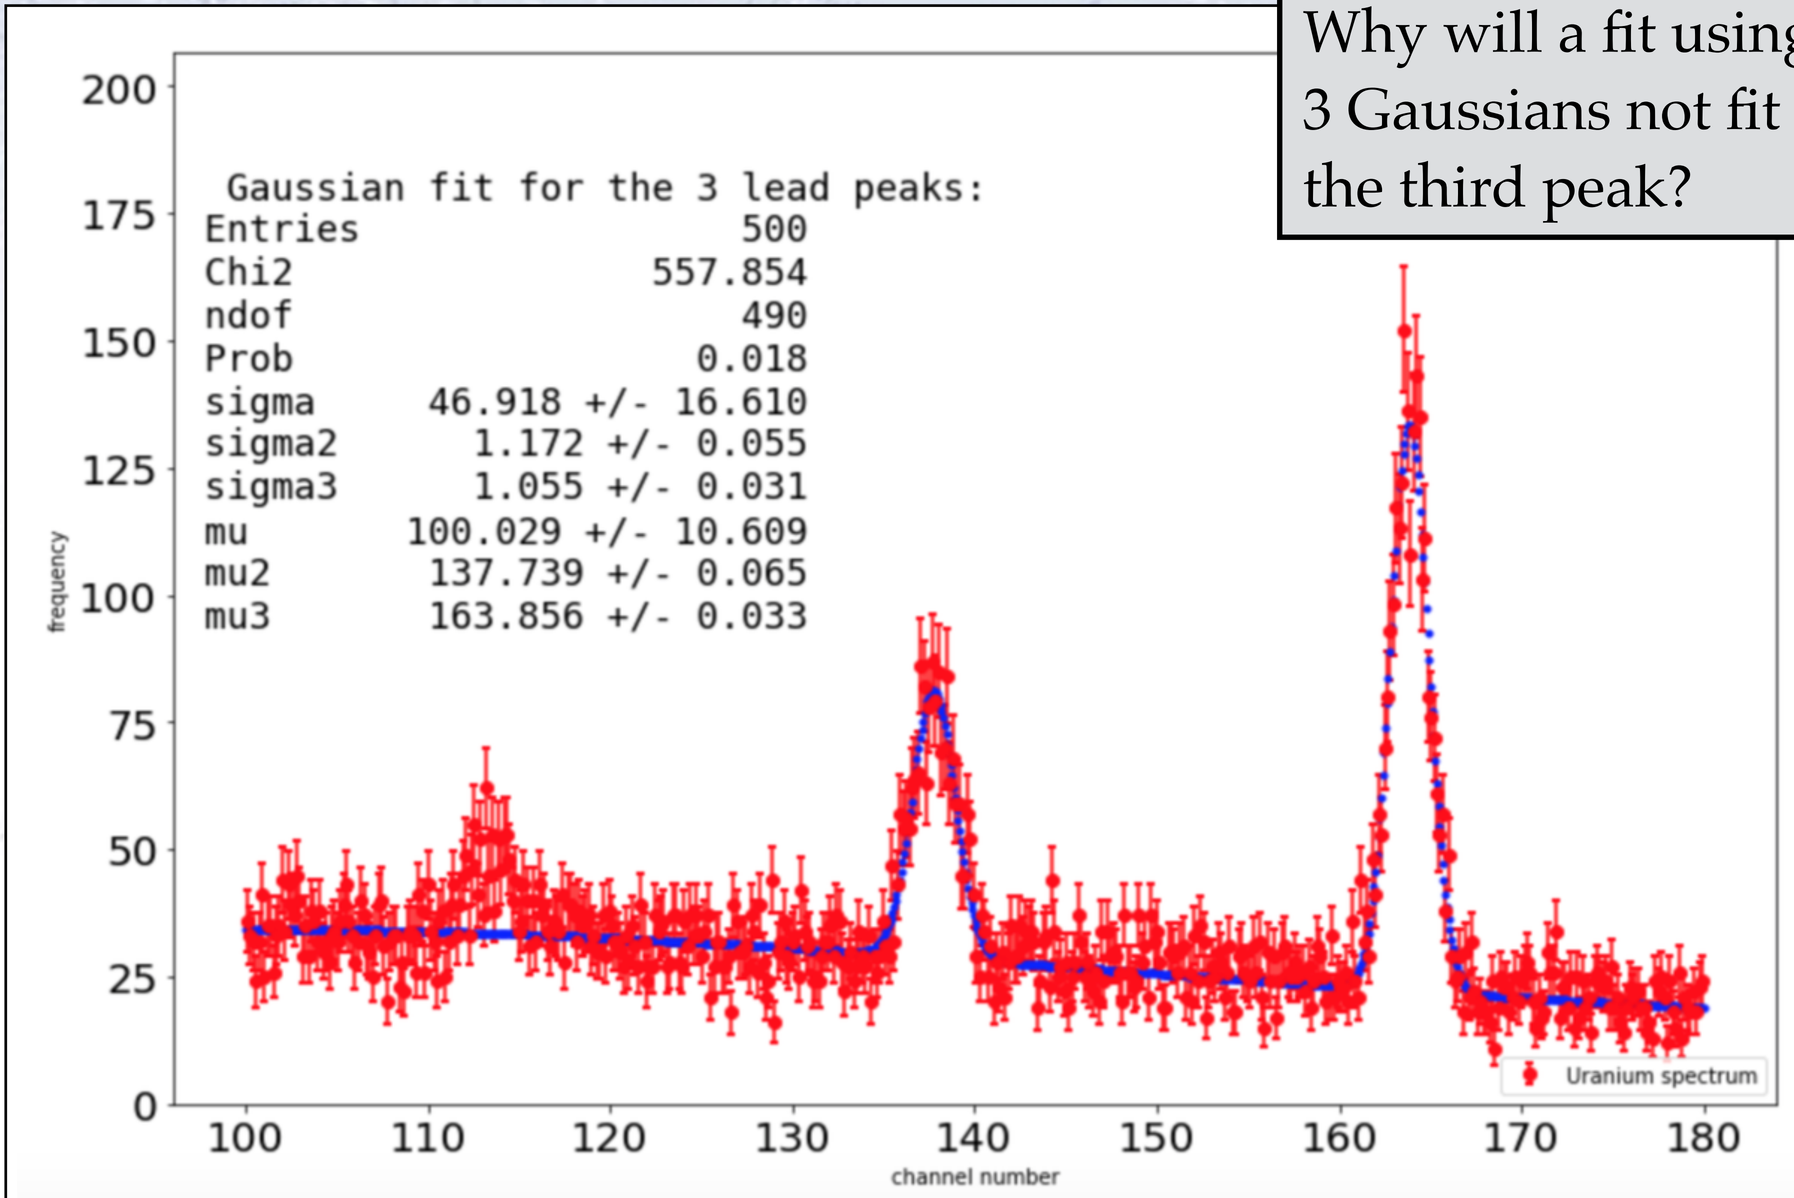

Problem 5.1 - No fit solution!

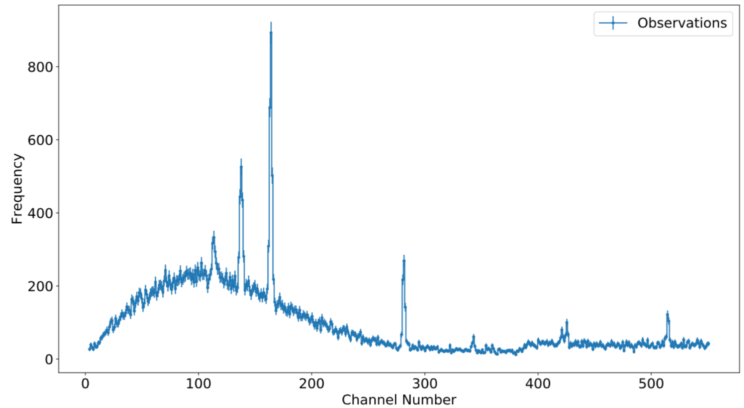

Figure 8: Binned data from uranium core, with a bin-width of 1.10.

From Figure 8 the three peaks corresponding to the peaks at energies 242 keV, 295 keV and 352 keV are visible, within the first 200 channels. In order to fit them I would fit a sum of three Gaussian distributions with mean and standard deviation according to the three peaks, using Chi2Regression with Minuit.

5.1.2 Had the fitting in the previous question worked I would have used the mean from each peak, to calculate a relative distance of the peaks r_{peaks} and compared it with the relative distance of the energies $r_{energies}$:

$$r_{peaks} = \frac{\mu_3 - \mu_2}{\mu_2 - \mu_1}$$

$$r_{energies} = \frac{E_3 - E_2}{E_2 - E_1} = \frac{(352-295)keV}{(295-242)keV} = \frac{57}{53} \approx 1.0755$$

5.1.3 Depending on how well r_{peaks} and $r_{energies}$ match I would choose to determine the energy scale from the three first peaks. Maybe the energy scaling would allow to determine the peaks of 214Bi gamma rays at 609 keV and 1120 keV, and from those five peaks I would test if it was linear.

5.1.4 From the fits I would determine the peak width, using mean and standard deviation. Using the scaling from **5.1.3** I would determine if the energy resolution changes with energy.

5.1.5 I would assume the small peak in the spectrum range 700-800 keV can be found in data around channel number 420, as I assume the peak at channel number 280 and the peak at channel 350 are from 214Bi gamma rays. Something I would test had I known the energy scale.

5.1.6 From the data it appears that at channel number 400 and on the frequency rises. This may be due to background measurements. Furthermore there are a total of 7 visible peaks in data, and only six peaks have been mentioned and could possibly be paired with actual gamma ray peaks in the above. Had I known the energy I might have determined that one of the peaks was just noise, or a value within the range of an actual gamma ray peak.

Problem 5.1

With $P(\chi^2 = 198, ndof = 195) = 0.424$ **we seem to have found a peak at** (743.9 ± 0.4) **keV**

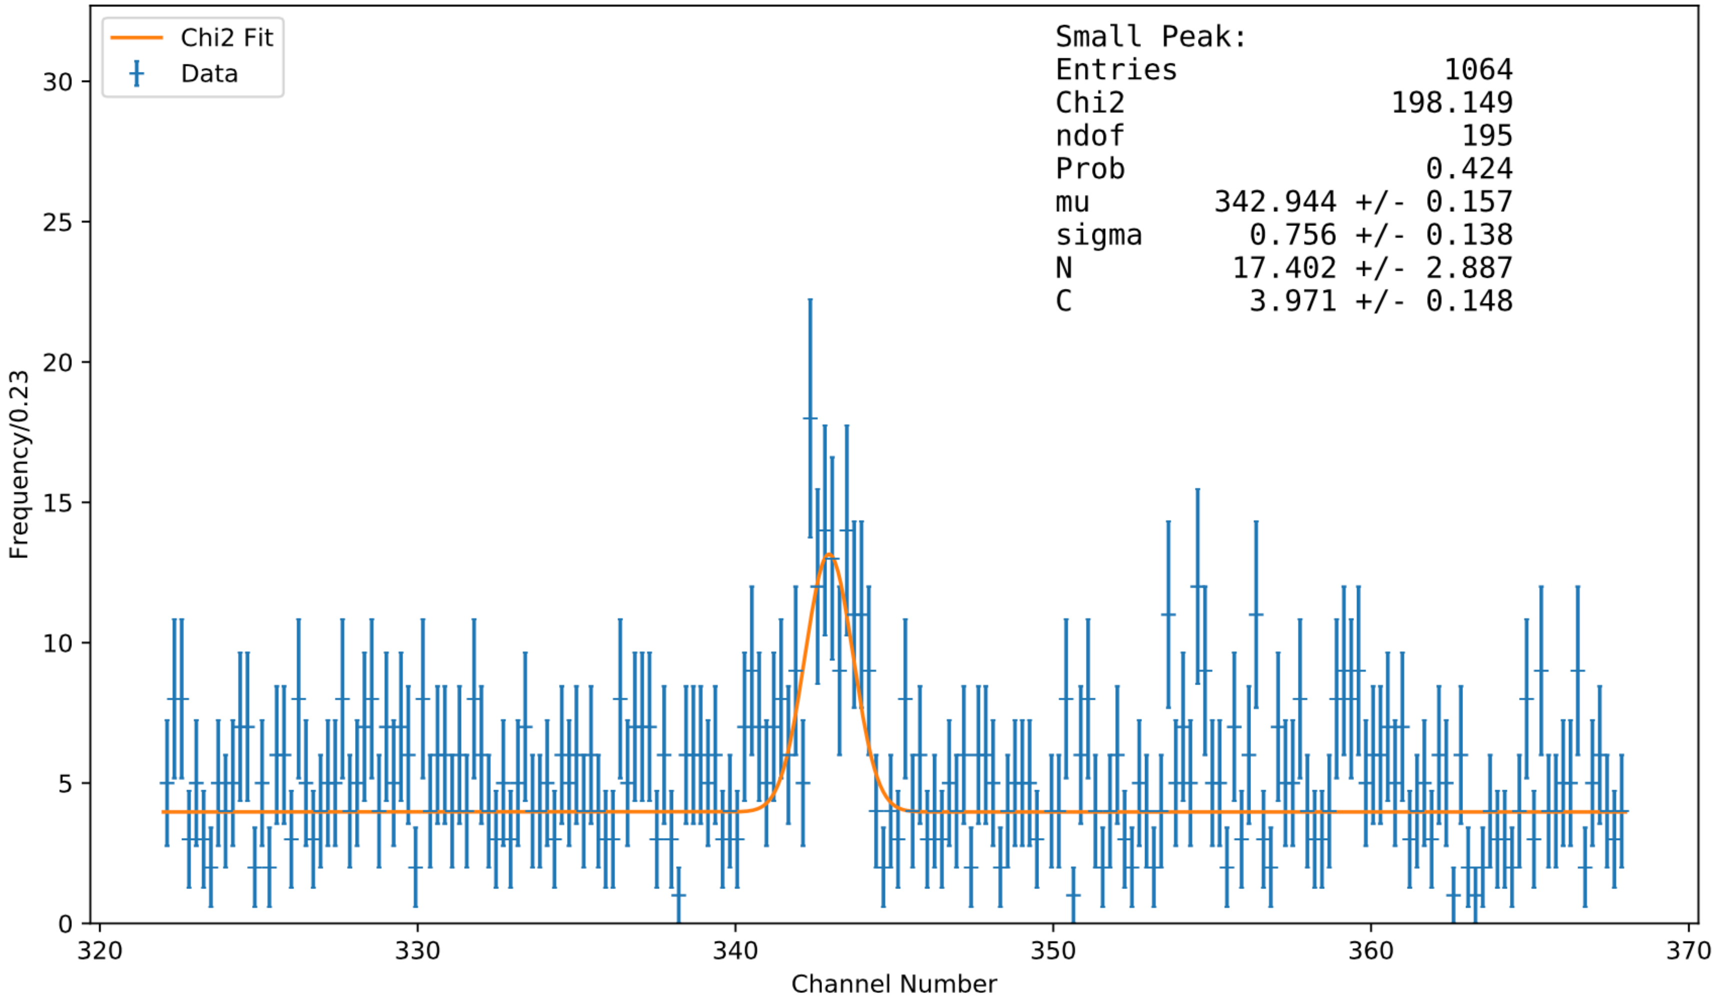

Problem 5.1

With $P(\chi^2 = 198, ndof = 195) = 0.424$ **we seem to have found a peak at** (743.9 ± 0.4) **keV**

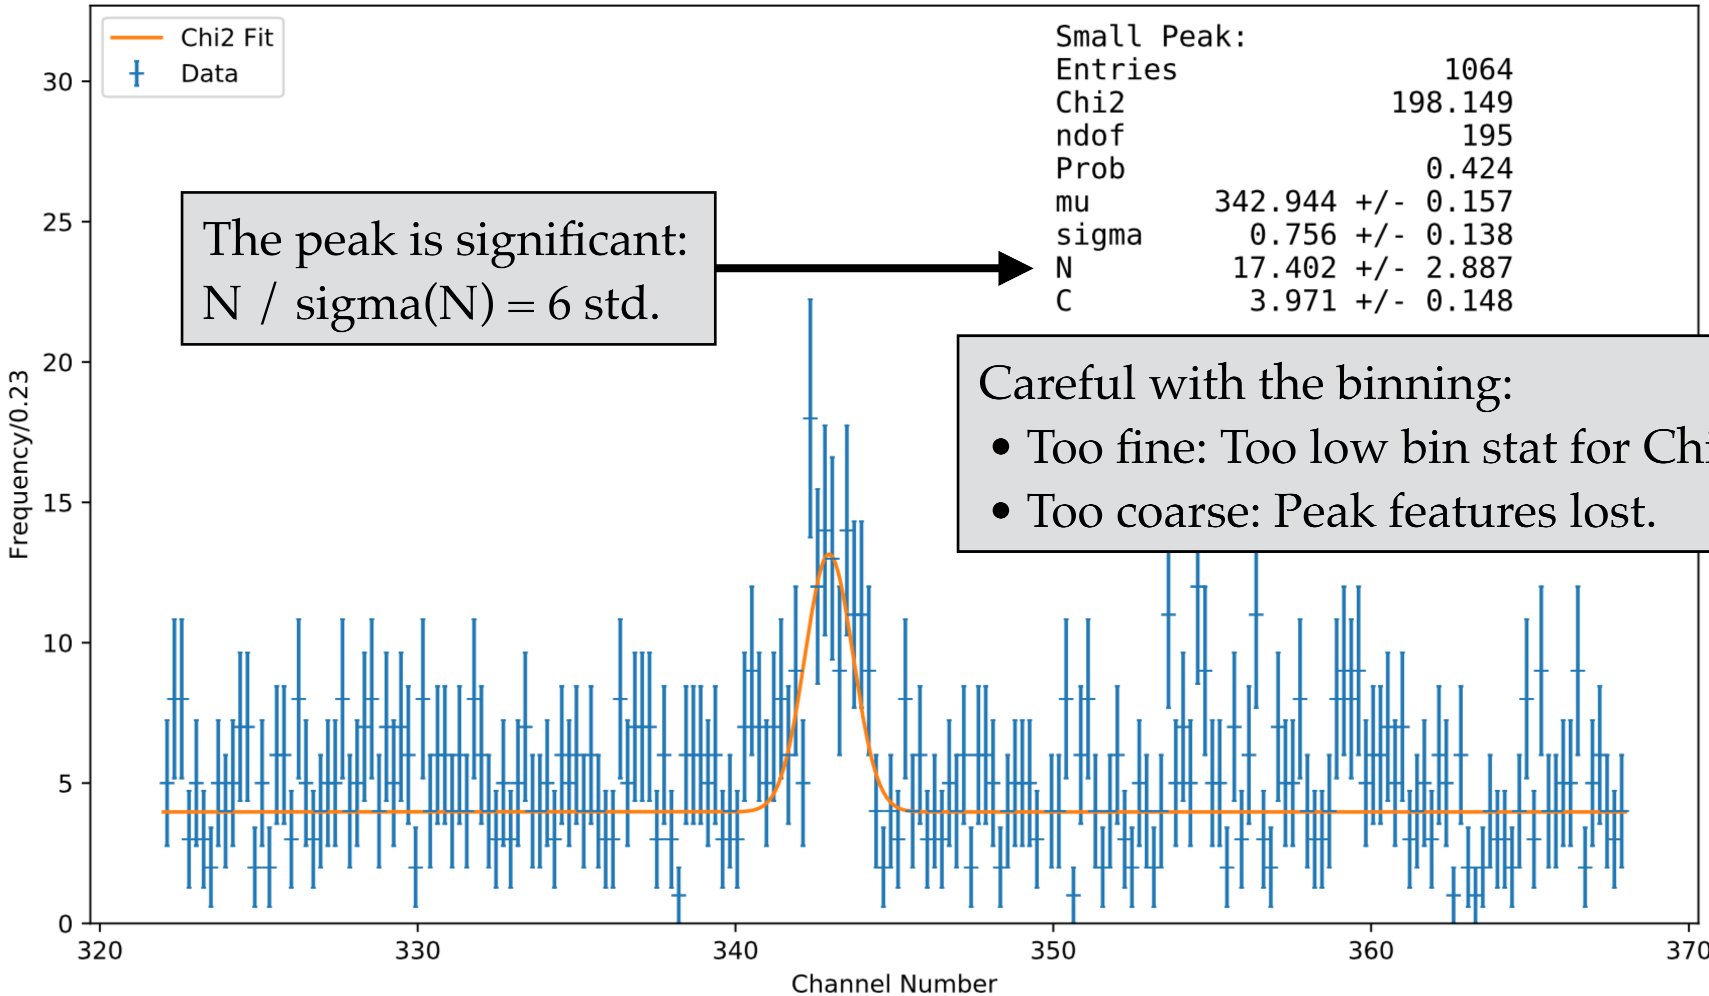

The peak is significant:
N / sigma(N) = 6 std.

Careful with the binning:

- Too fine: Too low bin stat for Chi2.
- Too coarse: Peak features lost.

Problem 5.1

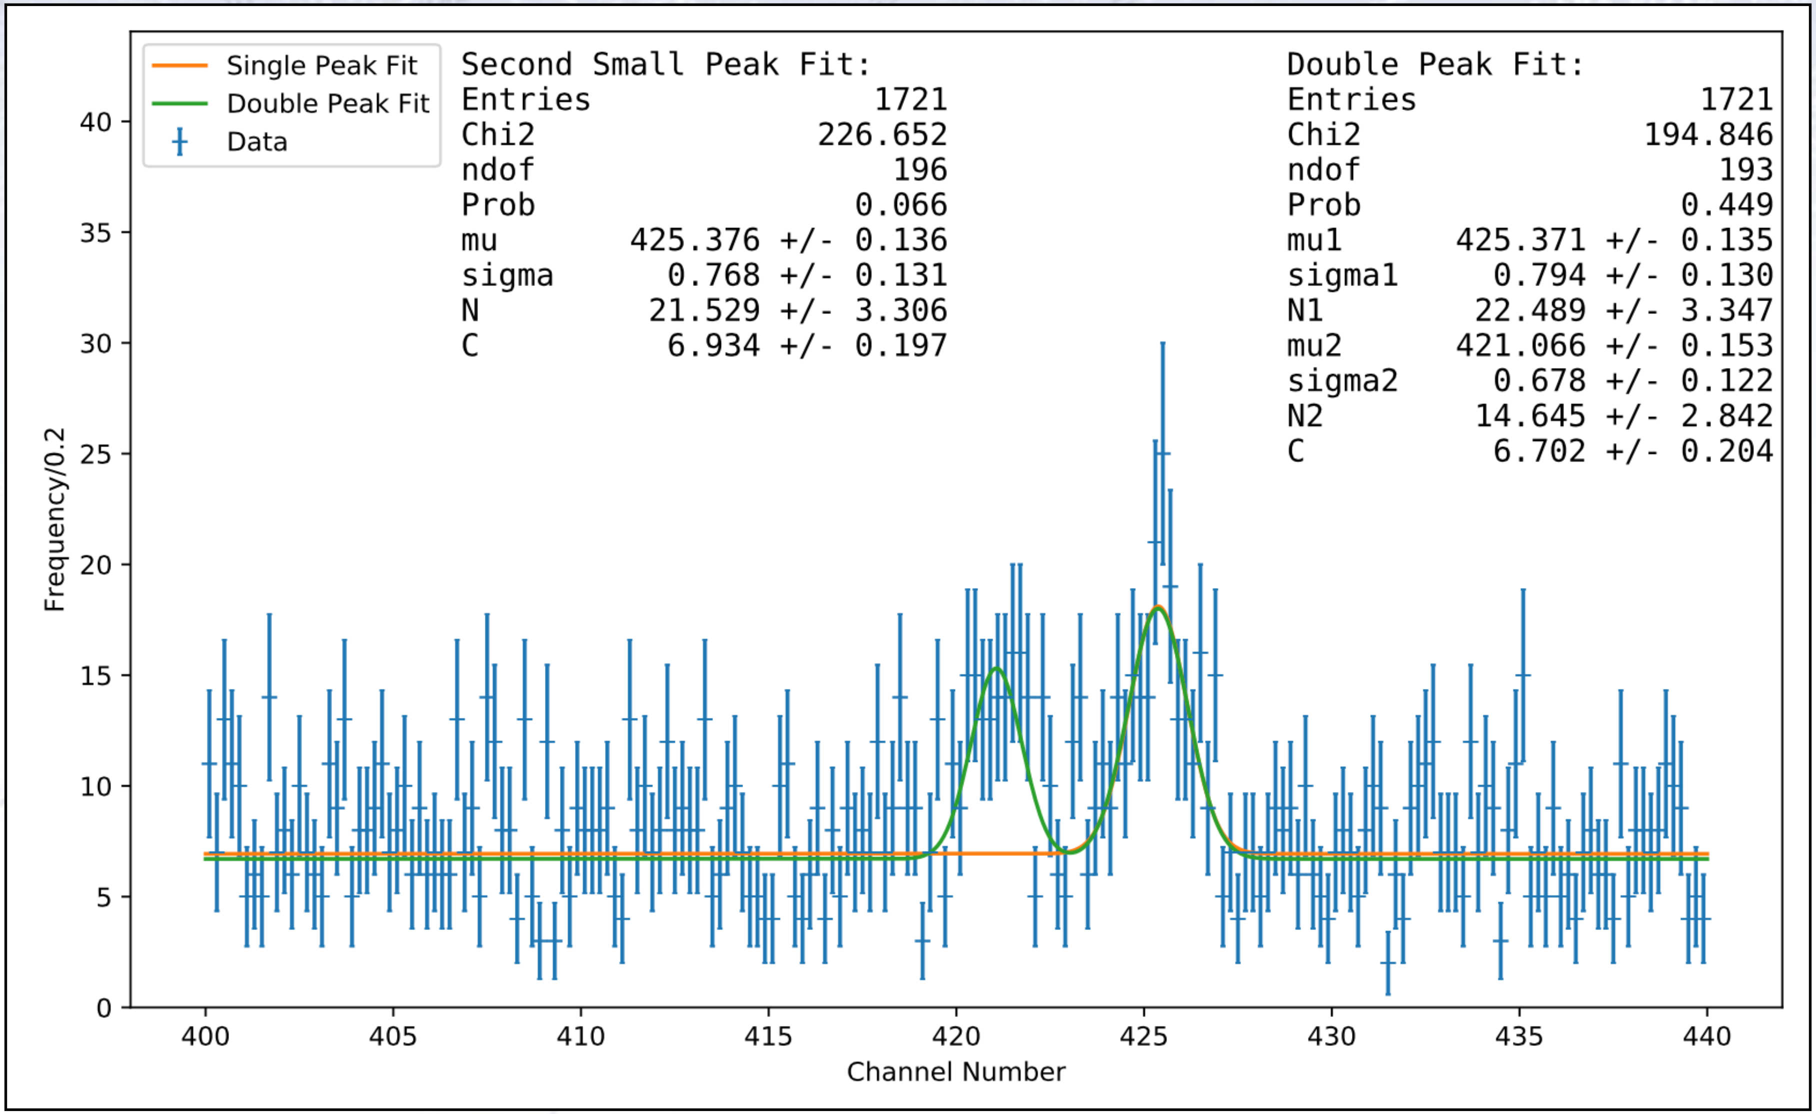

Single Peak Fit
Double Peak Fit
Data
Second Small Peak Fit:
Entries 1721
Chi2 226.652
ndof 196
Prob 0.066
mu 425.376 +/- 0.136
sigma 0.768 +/- 0.131
N 21.529 +/- 3.306
C 6.934 +/- 0.197
Double Peak Fit:
Entries 1721
Chi2 194.846
ndof 193
Prob 0.449
mu1 425.371 +/- 0.135
sigma1 0.794 +/- 0.130
N1 22.489 +/- 3.347
mu2 421.066 +/- 0.153
sigma2 0.678 +/- 0.122
N2 14.645 +/- 2.842
C 6.702 +/- 0.204
Frequency/0.2
Channel Number
0
5
10
15
20
25
30
35
40
400
405
410
415
420
425
430
435
440

Problem 5.1

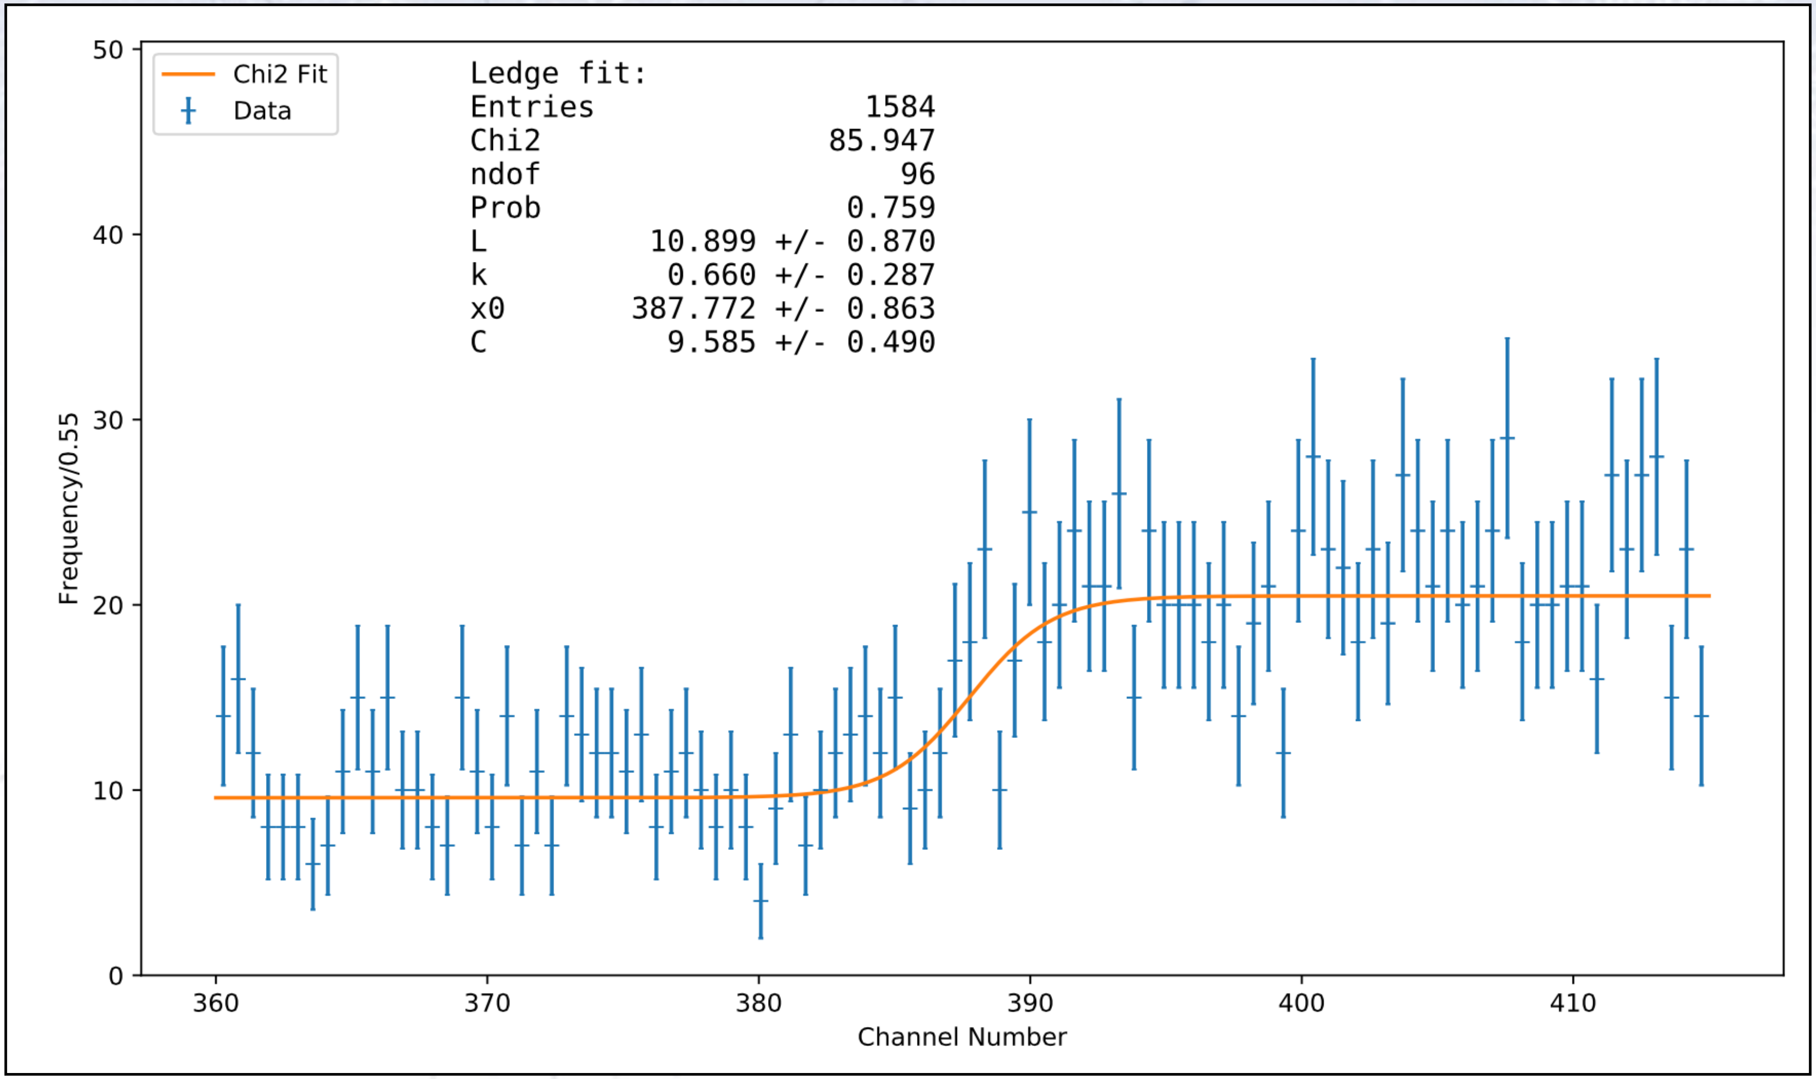
Chi2 Fit
Data
Ledge fit:
Entries 1584
Chi2 85.947
ndof 96
Prob 0.759
L 10.899 +/- 0.870
k 0.660 +/- 0.287
x0 387.772 +/- 0.863
C 9.585 +/- 0.490
Frequency/0.55
Channel Number
0
10
20
30
40
50
360
370
380
390
400
410

Problem 5.1 - inspiration

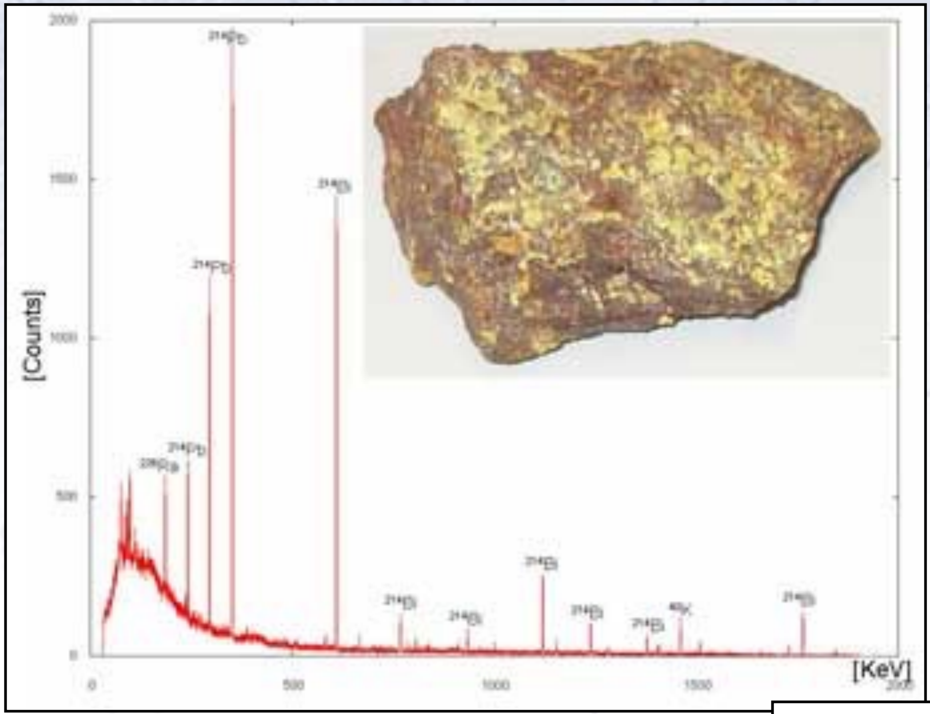

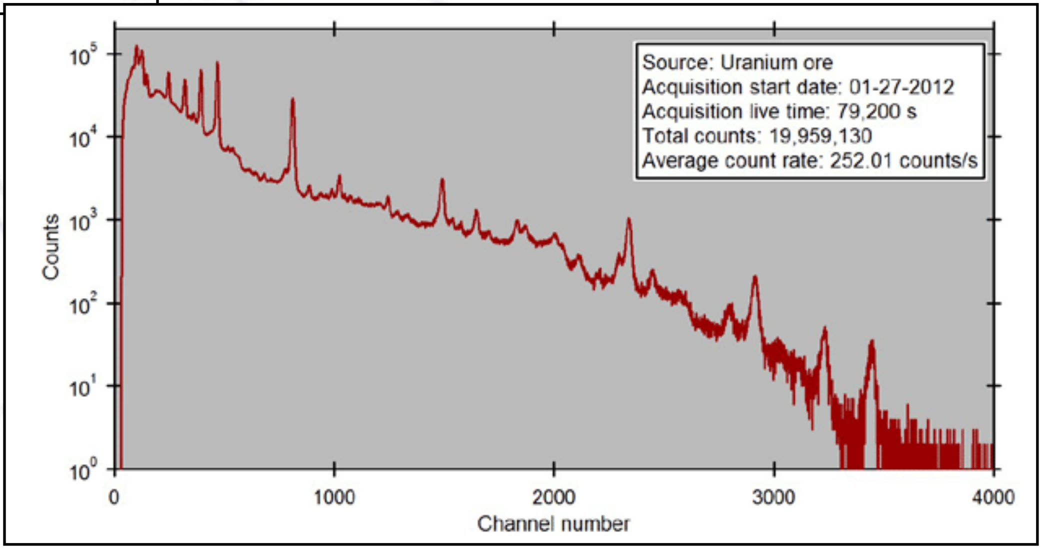

“Doing statistic problems during the entire night is like partying.
You can't stop yourself, but the next day you feel miserably…”

[End of an exam solution]

Comment on code sharing!

Out of interest, we ran Moss (Measure Of Software Similarity) on your code, which is an automatic system for determining the similarity of programs (e.g. detecting plagiarism in programs).

Don't worry - nothing suspicious was found. Thank you!

Moss Results

Tue Jan 22 04:54:30 PST 2019

Options -l python -d -m 4

[How to Read the Results | Tips | FAQ | Contact | Submission Scripts | Credits]

File 1	File 2	Lines Matched
py/[redacted]/ (32%)	py/[redacted]/ (43%)	411
py/[redacted]/ (7%)	py/[redacted]/ (6%)	68
py/[redacted](8%)	py/[redacted]/ (6%)	82
py/[redacted](7%)	py/[redacted]/ (13%)	52
py/[redacted]/ (4%)	py/[redacted]/ (4%)	95
py/[redacted]/ (20%)	py/[redacted]/ (14%)	84
py/[redacted] (3%)	py/[redacted]/ (0%)	68
py/[redacted]/ (7%)	py/[redacted]/ (7%)	54
py/[redacted]/ (5%)	py/[redacted]/ (5%)	69
py/[redacted]/ (11%)	py/[redacted]/ (12%)	62
py/[redacted](3%)	py/[redacted](2%)	51
py/[redacted] (4%)	py/[redacted]/ (4%)	43
py/[redacted]/ (3%)	py/[redacted](4%)	30
py/[redacted]/ (2%)	py/[redacted]/ (2%)	55
py/[redacted]/ (2%)	py/[redacted]/ (2%)	77

Comment on code sharing!

Out of interest, we ran Moss (Measure Of Software Similarity) on your code, which is an automatic system for determining the similarity of programs (e.g. detecting plagiarism in programs).

Don't worry - nothing suspicious was found. Thank you!

Moss Results

Tue Jan 22 04:54:30 PST 2019

Options -l python -d -m 4

File 1	File 2	Lines Matched
py/[redacted]/ (7%)	py/[redacted]/ (7%)	54
py/[redacted]/ (5%)	py/[redacted]/ (5%)	69
py/[redacted]/ (11%)	py/[redacted]/ (12%)	62
py/[redacted](3%)	py/[redacted](2%)	51
py/[redacted] (4%)	py/[redacted]/ (4%)	43
py/[redacted]/ (3%)	py/[redacted](4%)	30
py/[redacted]3/ (2%)	py/[redacted]/ (2%)	55
py/[redacted]/ (2%)	py/[redacted]3/ (2%)	77

Moss Results

Tue Jan 22 05:11:04 PST 2019

Options -l python -d -m 1000000

[How to Read the Results | Tips | FAQ | Contact | Submission Scripts | Credits]

File 1	File 2	Lines Matched
py[redacted](85%)	py/[redacted]0/ (23%)	936
py[redacted]36%)	py/[redacted]48%)	458
py[redacted] (31%)	py/[redacted]%)	424

Your results….

Total score	KU ID
100	
69	djc619
77.5	cpv752
98	nhs907
68.5	ckz831
88.5	wkr446
0	gsl528
0	jmk397
81	xgt955
68.5	kwp513
76.5	jrz619
75.5	qkg982
77	szm331
63.5	qzh229
79	zrf802
71.5	ptj469
30.5	rbn433
46.5	ksj465
67	xjl924
57.5	xgj708
51.5	vqx956
89	hwl460
75.5	mnk942
47	
66	kgt154
87	qzs982
100	lgb543
68	pfq906
72.5	bdl889
87.5	jhm221
52	mgx632
0	mnx938
72	hzn955
9	lxq472
71.5	fcz936
80	qxk428
69	dxp300

Total score	KU ID
100	
94	pfl888
51.5	hvw680
31	ncs601
58	zts164
98	pcm615
74.5	nwl935
64	zql906
57	pdg115
87	bkq651
80	mfs627
94.5	gsr903
0	hwl460
74.5	zkv499
87.5	tnh658
0	gtv101
51	gtl238
0	ksw803
72	htd809
33	jlw351
62.5	rhf801
56.5	jnr144
68.5	nsl738
69	pml305
98	gsc967
0	mhf548
71.5	cxq235
0	klq995
80.5	vgj803
76.5	zws783
71	mvw615
81	qvb164
50.5	rjt488
72	zfg663
68.5	mzw962
89	hlg223
81.5	pmx895

Total score	KU ID
100	
88.5	rch246
37.5	fxq291
0	hgj942
56	hfb792
86	dpf150
72.5	sgw622
76.5	bvf365
95.5	dwz764
69.5	sqv821
88	kqc695
88.5	qzk800
64.5	qrd689
68	svk776
92	nwb154
70	bwr366
57.5	dmc472
75	lbc622
84.5	wjv651
79.5	smj783
102	rqt552
37	str224
76	ngh299
77.5	rng399
0	qvn822
81.5	znl919
63.5	zdl473
93.5	pwn274
98	msk377
86.5	zfj803
76.5	pzj861
94.5	mfv505
90.5	bmc999
95	hgb619
78	jrh351
61	lwq229
75	mds274

Total score	KU ID
100	
74	hzq483
81	nbf686
34.5	cfg939
95.5	mpb982
35.25	zwx233
47.5	nqj779
98	dlk339
85	dvj919
76.5	zqp405
79.5	mbc442
66.434	Average
0.664	
72.586	Average - 0s
0.726	

1.1.1	1.1.2	1.2	1.3.1	1.3.2	2.1.1	2.1.2	2.1.3	2.2.1	2.2.2	3.1.1	3.1.2	3.1.3	3.1.4	3.1.5	4.1.1	4.1.2	4.1.3	4.1.4	4.2.1	4.2.2	4.2.3	5.1.1	5.1.2	5.1.3	5.1.4	5.1.5	5.1.6	Total score	KU ID
3	**3**	**4**	**3**	**3**	**5**	**4**	**4**	**4**	**5**	**3**	**3**	**3**	**3**	**3**	**3**	**4**	**4**	**4**	**4**	**4**	**5**	**3**	**3**	**3**	**3**	**3**	**4**	**100**	
2.82	2.65	3.5	2.519	2.257	3.434	3.16	2.368	3.37	4.319	2.227	2.556	2.87	1.602	2.287	2.556	2.866	3.204	2.523	3.454	2.546	2.667	2.343	1.551	1.917	1.505	2.083	1.727	66.434	Average
0.94	0.883	0.875	0.84	0.752	0.687	0.79	0.592	0.843	0.864	0.742	0.852	0.957	0.534	0.762	0.852	0.717	0.801	0.631	0.864	0.637	0.533	0.781	0.517	0.639	0.502	0.694	0.432	0.664	
2.874	2.779	3.533	2.642	2.439	3.434	3.454	2.51	3.37	4.401	2.29	2.629	2.925	1.73	2.47	2.654	3.005	3.426	2.753	3.587	2.75	3.236	2.75	1.861	2.379	2.083	2.778	2.784	72.586	Average - 0s
0.958	0.926	0.883	0.881	0.813	0.687	0.864	0.628	0.843	0.88	0.763	0.876	0.975	0.577	0.823	0.885	0.751	0.857	0.688	0.897	0.688	0.647	0.917	0.62	0.793	0.694	0.926	0.696	0.726	

<u>Problem 2.1 (tumor depth):</u>

This is essentially the TableMeasurement problem (with less statistics). Especially the measurements without uncertainties gave rise to problems.

<u>Problem 3.1.4 (generating numbers):</u>

Fitting the 500 numbers gave rise to bins with low statistics (unless you binned coarsely!).

<u>Problem 4.2 (Weldon's dice):</u>

It was clear to most, that this was Binomial, but realising that the dice are not exactly fair was harder!
Simply plotting (even without errors) was NOT enough. p(5+6) = 0.3378 +- 0.0008, which is >5σ from 1/3.

<u>Problem 5.1 (Gamma spectrum):</u>

The second problem on the ratios required proper error propagation, which few did.
The fourth problem on the peak resolution (i.e. fitted widths) was a matter of comparing them (ChiSquare).
The last problem is harder, as a "general open search" is less textbook and more reality!

- There was a double peak barely significant.
- There was a change in the background level.