Estimate the standard deviation for each of the following rate measurements

Counts	Duration (min)	Rate (/min)	Sigma
210	5000	0.042	
280	6000	0.047	
500	10000	0.050	
1100	20000	0.055	
20	500	0.040	
49	1000	0.049	

Estimate the standard deviation for each of the following flux measurements

		Paddle characteristics				
Counts	Duration (min)	Separation (cm)	Area 1 (cm^2)	Area 2 (cm^2)	Flux (/m^2/s/sr)	Sigma
500	5000	70	400	400	0.510	
250	5000	100	400	400	0.521	
100	5000	150	400	400	0.469	
130	10000	200	400	400	0.542	
160	5000	150	400	650.25	0.461	
180	10000	200	400	650.25	0.461	

Pick several pairs of either rate or flux measurements and work out the

number of standard deviation difference between the measurements

Assume all these trials measure the same true rate.
1) Calculate the average rate over all these trials by summing counts and duration to get a total number of counts and duration and divide the two. (Think of this as combining the data into one big experiment.)
2) Calculate the number of standard deviations (Nsd) each trial departs from this average
3) Highlight any differences among measurements that are "significant"

Counts	Duration (min)	Rate (/min)	Sigma	Nsd
210	5000	0.042		
280	6000	0.047		
500	10000	0.050		
1100	20000	0.055		
20	500	0.040		
49	1000	0.049		

Average	Sigma

Assume all these trials measure the same true flux.
1) Calculate the average flux over all these trials by taking a WEIGHTED AVERAGE
2) Calculate the number of standard deviations (Nsd) each trial departs from this average
3) Highlight any differences among measurements that are "significant"

		Paddle characteristics					
Counts	Duration (min)	Separation (cm)	Area 1 (cm^2)	Area 2 (cm^2)	Flux (/m^2/s/sr)	Sigma	Nsd
500	5000	70	400	400	0.510		
250	5000	100	400	400	0.521		
100	5000	150	400	400	0.469		
130	10000	200	400	400	0.542		
160	5000	150	400	650.25	0.461		
180	10000	200	400	650.25	0.461		

Average	Sigma

WEIGHTED AVERAGE: An average is simply the sum of a series of measurements, X_i, divided by the number of measurements, N. In a WEIGHTED AVERAGE, each measurement carries a weight w_i. The weighted average is just the [sum of (X_i*w_i)]/[sum of (w_i)]. In averaging a series of measurements with known standard deviations, sigma_i, the weight w_i is just 1/(sigma_i)^2. The sigma of a weighted average is 1/SQRT(sum of w_i).

Assume all these trials measure the same true rate.
1) Calculate the average and Nsd deviation as before
2) Calculate the total chi2 and number of degrees of freedom
3) Calculate the probability of this chi2

Counts	Duration (min)	Rate (/min)	Sigma	Nsd
210	5000	0.042		
280	6000	0.047		
500	10000	0.050		
1100	20000	0.055		
20	500	0.040		
49	1000	0.049		

Average	Sigma	Chi2	N d.o.f.	Prob.

Assume all these trials measure the same true flux.
1) Calculate the average flux and Nsd deviation as before
2) Calculate the total chi2 and number of degrees of freedom
3) Calculate the probability of this chi2

		Paddle characteristics					
Counts	Duration (min)	Separation (cm)	Area 1 (cm^2)	Area 2 (cm^2)	Flux (/m^2/s/sr)	Sigma	Nsd
500	5000	70	400	400	0.510		
250	5000	100	400	400	0.521		
100	5000	150	400	400	0.469		
130	10000	200	400	400	0.542		
160	5000	150	400	650.25	0.461		
180	10000	200	400	650.25	0.461		

Average	Sigma	Chi2	N d.o.f.	Prob.

CHI2: The chi2 of a set of measurements X_i to be consistent with an expectation x is sum[(X_i-x)^2/(sigma_i)^2]. Here our "expectation" is the average (weighted or other method) of all of all trials. Note that chi2 is always a positive quantity. The number of "degrees of freedom" in a chi2 test is just the number of data points. However, note that if the data itself is used to determine the expectation of x, then this reduces the number of degrees of freedom by one. The chi2 probability is a number between 0 and 1 that indicates how likely it is that for a given number of degrees of freedom, the chi2 could be greater than the reported chi2. A very small chi2 probability indicates that it is unlikely that the measurements X_i are all consistent with the expectation. Hint: the chi2 probability function in Excel is CHIDIST(chi2,ndof)

Now let's attempt to find a linear correlation between count rates and an external variable (in this case, temperature).

1) Calculate the average and Nsd deviation, against the "prediction"
2) Calculate the total chi2 and number of degrees of freedom
3) Calculate the probability of this chi2
4) Vary the parameters of the function (intercept, slope) and see how probability varies
5) Find the intercept that provides the minimum chi2 with slope=0. What is the probability?
6) Find the intercept that provides the minimum chi2 with slope=0.001. Is the probability significantly higher?
7) Repeat for slopes of 0.002 and 0.003

Counts	Duration (min)	Rate (/min)	Sigma	Temp	Prediction	Nsd
210	5000	0.042		68	0.050	
280	6000	0.047		73	0.050	
500	10000	0.050		71	0.050	
1100	20000	0.055		75	0.050	
20	500	0.040		71	0.050	
49	1000	0.049		70	0.050	

Intercept	Slope	Chi2	N d.o.f.	Prob.
0.05	0			

CHI2 Fitting: This is a simple example of "fitting" using a chi2. (More complex examples are possible!) The minimum chi2 gives the highest probabillity of the prediction explaining the data. Modifying the parameters of the prediction (here, intercept and slope) and watching how the chi2 changes is how a fitter determines the "best values" of parameters meant to describe data. Observation of a much higher probability when introducing a slope parameter as in this example would be evidence for a correlation between the external variable (temperature) and the observed rate.