Auntie Spark's Guide to Error Analysis

Many of my fellow students have complained about error analysis and propagation of errors. What are we supposed to do and how do we do it? Well the lectures and the lecture notes seemed clear enough to me, but since they don't seem that way to everyone, I've put together a few examples.

Example 1: A Single Set of Measurements

The first example is measuring air velocity from Pitot tube pressure measurements. We have to start from the theory: The equation for a Pitot tube that relates the pressure difference to the velocity is

$$v = \sqrt{\frac{2\left(P_0 - P\right)}{\rho}} \tag{1.1}$$

where v is the free stream velocity, P_0 is the stagnation pressure (the pressure at the tip of the Pitot tube), P is the free-stream pressure (the pressure at the side tap), and ρ is the air density. As long as I use consistent units, it doesn't matter whether I use SI or American Engineering units. Since the manometer I'm using is a differential one, I can simplify the equation to

$$v = \sqrt{\frac{2\Delta P}{\rho}} \tag{1.2}$$

where ΔP is the pressure difference measured by the manometer. As a starting point I looked up the accuracy of the manometer on the website http://www.dwyer-inst.com/Product/Pressure/Manometers/Digital/Series475#specs. It appears to have an accuracy of ±0.5% of full scale at the temperature range of interest. Since the full-scale differential pressure is 10 inH_2O, the uncertainty of the instrument is ±0.05 inH_2O or ±0.012 kPa. To calculate the uncertainty in the density of air, I needed the formula for the density of air in terms of things I can measure: temperature, atmospheric pressure, and relative humidity. For this exercise let's assume the answer came out to 1.220±0.020 kg/m^3. At one specific speed setting for the wind tunnel I found that the manometer was averaging 1.50 inH_2O and the fluctuations seemed to be ±0.07 inH_2O. Since the manometer doesn't store or transmit data I had to more-or-less eyeball the average and uncertainty from watching the excursions for about a minute. The proper way to add the uncertainty of the measurement with the uncertainty of the instrument depends on whether we assume the errors are random and uncorrelated or not. Let's assume they are. In that case the correct uncertainty is

$$e_{sum} = \sqrt{e_{man}^2 + e_{meas}^2} = \sqrt{0.05^2 + 0.07^2} = \pm 0.086 \text{ inH}_2\text{O} = \pm 0.021 \text{ kPa} \quad . \tag{1.3}$$

Now I can calculate the uncertainty in the velocity. The differential of Equation (1.2) is

$$dv = \frac{1}{\rho}\frac{d\Delta P}{\sqrt{\frac{2\Delta P}{\rho}}} - \frac{\Delta P}{\rho^2}\frac{d\rho}{\sqrt{\frac{2\Delta P}{\rho}}} = \frac{\rho d\Delta P - \Delta P d\rho}{\rho^2 v} \tag{1.4}$$

Again, assuming that the errors in ΔP and ρ are random and uncorrelated, the error in v is:

$$e_v = \frac{1}{\rho^2 v}\sqrt{\left(\rho e_{\Delta P}\right)^2 + \left(\Delta P e_\rho\right)^2} \tag{1.5}$$

For my measurement, then, the velocity is

$$v = \sqrt{\frac{2\Delta P}{\rho}} = \sqrt{\frac{2\left(1.50\text{inH}_2\text{O}\right)}{1.220\text{kg/m}^3}} = 24.75\text{m/s} = 55.37\text{mph} \tag{1.6}$$

and the uncertainty is:

$$\begin{aligned} e_v &= \frac{1}{\rho^2 v}\sqrt{\left(\rho e_{\Delta P}\right)^2 + \left(\Delta P e_\rho\right)^2} \\ &= \frac{1}{\left(1.220 \text{ kg/m}^3\right)^2 24.75 \text{ m/s}}\sqrt{\left(1.220 \text{ kg/m}^3 \cdot 21 \text{ Pa}\right)^2 + \left(374 \text{ Pa} \cdot 0.020 \text{ kg/m}^3\right)^2} \\ &= 0.72 \text{ m/s} \end{aligned}$$

(1.7)

The result is reported as $v = 24.75 \pm 0.72$ m/s. If you knew the confidence interval, e.g., 95% confidence), you'd include that as well. For the curious, the uncertainty in the calculated velocity is about ±3%.

If the manometer had a serial interface, I would have connected it to a computer and used LabVIEW to calculate the mean and uncertainty of the measurement using Student's *t*-test and my desired confidence interval. I could then either have the VI calculate the velocity and uncertainty using the formulas above with my entered instrument and density uncertainties, or I could have transferred the values to a spreadsheet and calculated the velocity and uncertainty there.

For the pedantic, uncertainty, error, and residual have very precise meanings relating to confidence interval, true or population mean, and sample mean. A statistician would state that the residuals are used as estimations of the errors and a statistical analysis of the residuals is used to estimate the confidence interval for a given uncertainty. In common engineering practice, the three are used as synonyms with the term residual being reserved for special occasions.

Example 2: A Linear Fit/Calibration of a Sensor

The second example is calibrating a differential pressure sensor to measure velocity. In this case I used a Freescale MPX53DP to measure the pressure difference between a pressure tap in the tip of a nosecone and the shoulder of a nosecone. This arrangement is a Pitot tube arrangement, but doesn't meet the recommendations for a Pitot tube, so a priori it's unlikely that the proportionality constants will be the same as in the theory, but the functional form should be correct. We have several possibilities for proceeding here. There are four models that need to be accounted for in this analysis. The first model is Bernoulli's law or the Pitot tube equation, Equation (1.2). The second is the assumed relationship between the differential pressure and the output voltage from the sensor,

$$V_{sensor} = K\Delta P + V_{off} \tag{1.8}$$

where K is the sensitivity (nominally 1.2 mV/kPa from the data sheet), and V_{off} is the offset voltage (nominally 20 mV from the data sheet). The third is the relationship between the sensor voltage and the output voltage of the conditioning circuit. For a decently designed circuit the relationship should be linear,

$$V_{out} = GV_{sensor} + V_{off2} \tag{1.9}$$

where G is the circuit gain and V_{off2} is the circuit offset calculated relative to the output. Typical values for G and V_{off2} might be 160 and –3V respectively. The fourth is the relationship between the analog-to-digital converter and the input voltage.

$$N = 2^{bits} \frac{V_{in}}{V_{ref}} \ . \tag{1.10}$$

For my case, I have a 16-bit converter with a range from 0V to 3.3V, so $bits$ = 16 and $V_{ref} = 3.3\text{ V}$.

At this point I have a choice. I can experimentally determine the constants and uncertainties in each of my models and propagate the errors as we've done before, or I can combine the models, get an overall functional form, and I can experimentally determine the constants and uncertainties for the combined model. Since the second method requires ¼ of the experiments of the first method, that's what I'll do.

$$N = m'v^2 + b' \tag{1.11}$$

where

$$m' = \frac{2^{bits-1} G K \rho}{V_{ref}} \tag{1.12}$$

and

$$b' = \frac{2^{bits}\left(G V_{off} + V_{off2}\right)}{V_{ref}} \ . \tag{1.13}$$

It looks like a plot of N versus v^2 should be a straight line and I can get nominal values for the slope and intercept from my models. However, I can measure N much more accurately than I can set v. Linear regression assumes that the uncertainty in x is negligible when compared with the uncertainty in y. In other words, it assumes I can set or measure my independent variable precisely and all of the error is in the dependent variable. What to do?

Well, what am I trying to do anyway? Ultimately, I want to get $N(t)$ during a flight and use it to calculate $v(t)$. Since I'm going to have to do the inverse function eventually anyway, I might as well do it now.

$$v^2 = mN + b \tag{1.14}$$

where $m = N / m'$ and $b = -b' / m'$. I can now take my data, plot v^2 as a function of N, and do the line statistics on m and b.

The associated spreadsheet, FittingAPitotNose.xlsx, is an implementation for one data set. I attempted to set the velocity, v_{set}, in the wind tunnel in 2 m/s increments. The attempted settings are Column A in the spreadsheet. I then averaged the ADC counts (N_{meas}) at each setting for Column B. I could have done statistics on each measurement and done a weighted least-squares fit, but my line fitting should give me sufficient statistics. I then calculated the square of each set velocity, v_{set}^2 (Column C), and used Excel's built-in line fitting functions to calculate the slope, m (Cell B35), and the intercept, b (Cell B36), of my v^2 vs. N fit. I then calculated the v^2 from the fit, v_{calc}^2 (Column D), for each N_{meas}. I need these for my residuals, which are calculated in Column E and then used to calculate the root mean squared residual, S_e, in Cell E35. For a given N_{meas}, call it N_j, the sample standard error in v^2, S_y, is calculated by

$$S_y = S_e \sqrt{\frac{1}{N} + \frac{(x-\bar{x})^2}{\sum_{i=1}^{N}(x_i-\bar{x})^2}} = S_e \sqrt{\frac{1}{count} + \frac{(N_j-\bar{N})^2}{\sum_{i=1}^{count}(N_i-\bar{N})^2}} \quad , \tag{1.15}$$

(given in the Spring 2015 Data Analysis Lecture Notes) which is Column G. Note that the values are largest at the two ends and smallest in the middle. The degrees of freedom are 29 = count–2 in this case (Cell B38) and the significance level, α, is chosen as 0.05 (95% confidence interval) in Cell B39. The Students *t*-test value is calculated in Cell B40. The resultant plus and minus confidence intervals are calculated in Columns H and I and plotted along with the best-fit line and the data points in Figure 1.

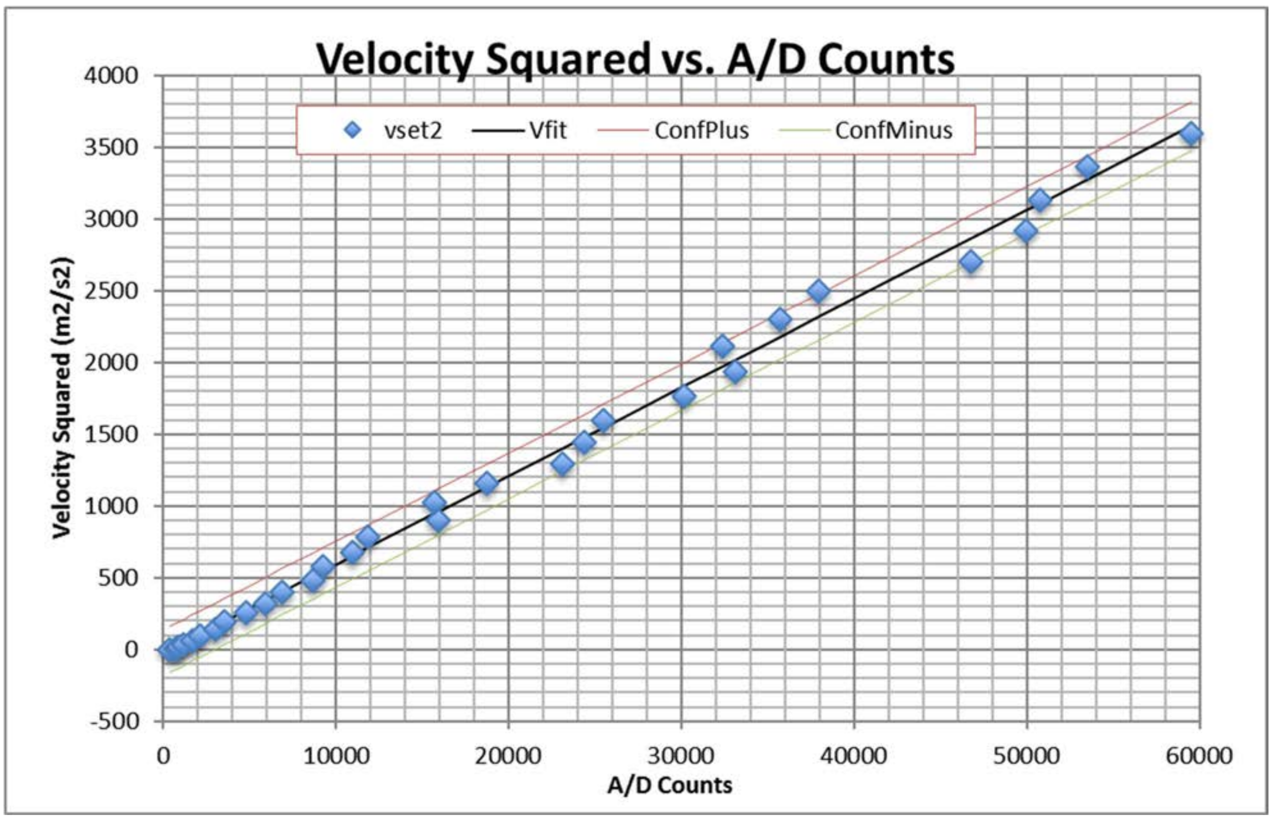

Figure 1: The confidence interval for the velocity squared as a function of the ADC counts

Even though the confidence interval boundaries look like straight lines, they are actually curved with the minimum in the middle and the maxima at the two ends. For this case it's hard to see the curvature.

Finally, the confidence interval in v_{calc} is calculated from the confidence interval in v_{calc}^2 with the standard propagation of errors technique where

$$\varepsilon_v = \frac{\partial\sqrt{v^2}}{\partial v^2}\varepsilon_{v^2} = \frac{1}{2v}\varepsilon_{v^2} \ , \qquad (1.16)$$

and $\varepsilon_{v^2} = tS_y$ (Column K). A plot of the confidence interval as a function of the calculated velocity is shown in Figure 2. Note that the uncertainty in velocity increases greatly as the measured velocity decreases, whereas the uncertainty in the square of the velocity doesn't vary much with velocity.

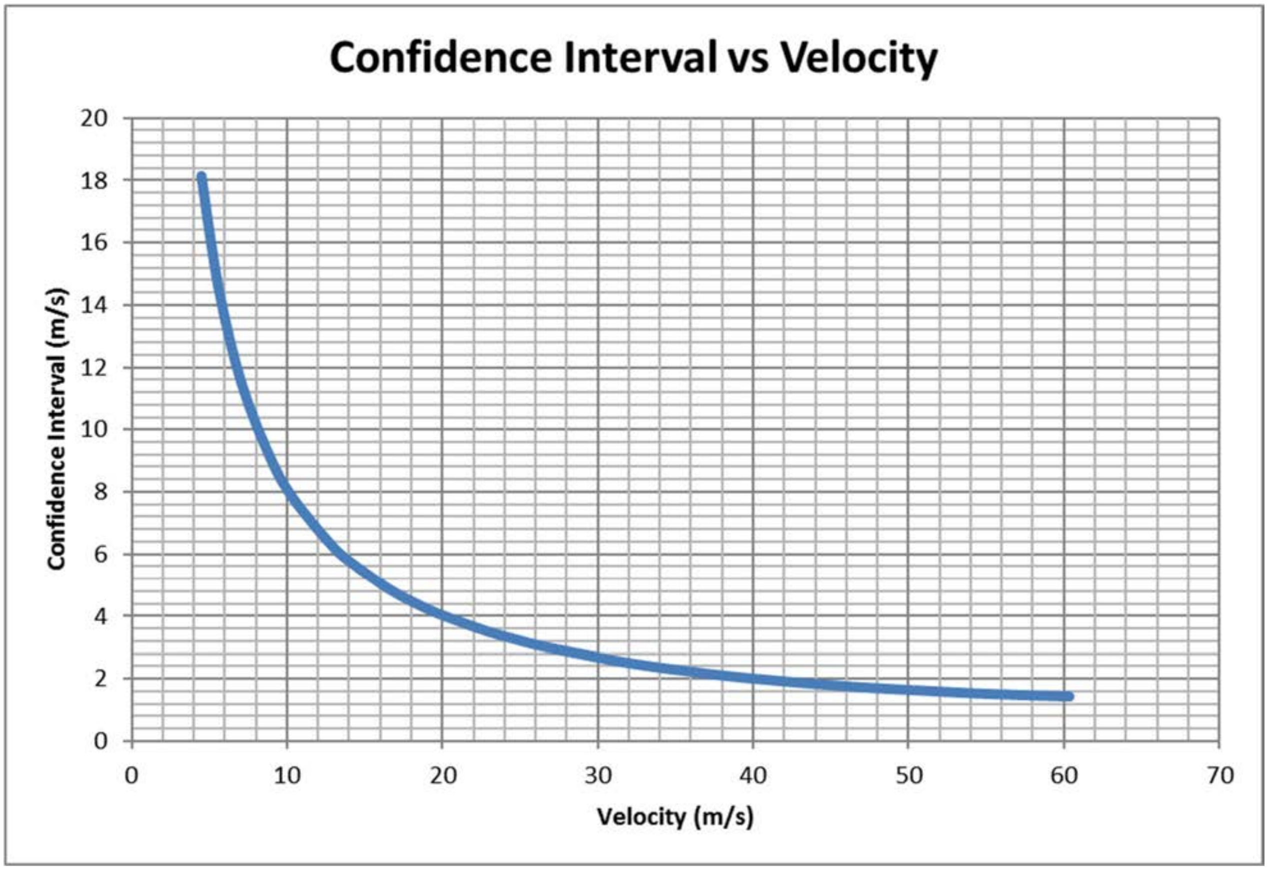

Figure 2: The Confidence Interval of the Velocity as a Function of the Velocity

If you had generated these data while calibrating your Pitot tube, how would you use your results to analyze actual flight data from your Pitot tube? Your data from the flight would be a sequence of numbers from your ADC, in other words a set of *N*'s. For each N_i you would calculate the v_i^2 from your calibration (linear fit) *m* and *b* and calculate the confidence interval for that v_i^2 using Equation (1.15). Remember that all of your values in the equations except for N_i (the N_j in the equation) and v_i^2 come from your original calibration data set. You then calculate your $v_i = \sqrt{v_i^2}$, and your confidence interval from Equation (1.16). You can then plot your velocity versus time with your error bars or error band.

Example 3: Integrating and Differentiating Time Series Data

The third example is calculating acceleration and position from a set of velocity measurements (say from a speedometer). If you had a perfect noiseless function for the velocity, $v(t)$, then the standard method for calculating the acceleration would be to differentiate,

$$a(t) = \frac{dv(t)}{dt} \tag{1.17}$$

and the standard method for calculating the position would be to integrate,

$$r(t) = \int_0^t v(t')dt' \tag{1.18}$$

Since we're differentiating and integrating with respect to time, we need models for the errors or noise that are functions of time. There are three common functions that are used to model the errors or noise: The first is a constant bias, $\varepsilon_b(t) =$ constant $= \varepsilon_b$. This sort of error usually results from an inaccuracy or offset in the calibration. The second is to assume that there is a random noise component, $\varepsilon_n(t)$, that is Gaussian white noise with standard deviation, σ. Note that this model doesn't work for the Pitot tube in Example 2. There the uncertainty varies with velocity. Here it doesn't: The noise is constant regardless of velocity. The third is to assume that the nearly constant bias exhibits flicker noise. On the time scale of a flight we shouldn't have to worry about the flicker noise, so we'll ignore it. Our resulting model is

$$v_m(t) = v(t) + \varepsilon_b + \varepsilon_n(t) \ , \tag{1.19}$$

where $v_m(t)$ is the measured velocity and $v(t)$ is the true velocity. The integration of the model will give:

$$r_m(t) = \int_0^t \left[v(t')dt' + \varepsilon_b dt' + \varepsilon_n(t')dt' \right] = r(t) + \varepsilon_b t + \sigma\sqrt{\Delta t}\sqrt{t} \ , \tag{1.20}$$

where Δt is the sample time or the time between measurements (the reciprocal of the sample rate). The error in the position is then the measured position minus the true position or

$$e_m(t) = r_m(t) - r(t) = \varepsilon_b t + \sigma\sqrt{\Delta t}\sqrt{t} \tag{1.21}$$

The error term that grows linearly with time can be eliminated by carefully measuring the DC bias and subtracting it from the measured velocity. The term that grows as the square root of time is a classic random walk and can't be easily eliminated. The standard deviation of the noise can usually be decreased by buying a more expensive sensor (military-grade accelerometers and rate gyros are very expensive). Also, the sampling rate can be increased to decrease Δt .

To do a proper job of propagating the error in velocity to the error in acceleration requires extensive information in stochastic signals. The Wikipedia articles are a good place to start for the interested reader. However, there is a fairly straightforward method to get an approximate confidence interval or error bar estimate. If you have a multisample measurement of the signal with a known input (for example 100 samples at a known velocity), you can take the numerical derivative and then measure the standard deviation of the resultant samples and use that value in your calculations.

An example of these techniques is shown in the spreadsheet: SampleTimeSeriesErrors.xlsx. As a reminder, in this example I have a vertical velocity sensor (on a rocket) that has a fixed offset and random errors that don't vary with time or velocity. I figure you can do the work to propagate the errors from a device like in Example 2.

To calculate my offset and standard deviation I took two seconds of data before launch when I knew the velocity was zero. On the spreadsheet these values correspond to $t = -2$ to $t = 0$ in Columns A and B. I then tracked the vertical velocity from $t = 0$ to apogee. For the two seconds of pre-flight data with known velocity, I calculated the average (Cell C2) and the standard deviation (Cell C1). The average was used as the best guess for the constant offset, and all of the measured velocities had the average subtracted from them to get the corrected velocity. I was too lazy to bother with Students *t*-test, so I used ± the standard deviation for my 68% confidence interval and for plotting my error bars (Column D). The results are shown in Figure 3. The error bars are the horizontal points that are close to the round data points.

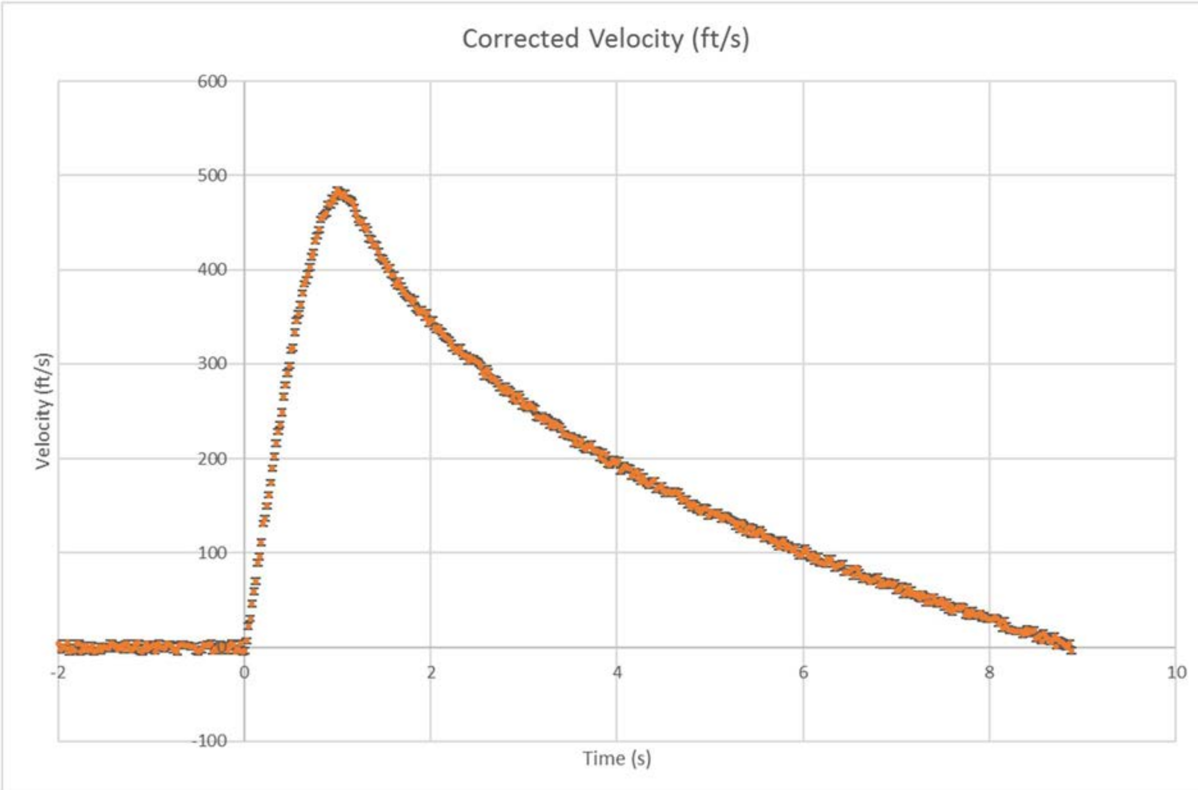

Figure 3: The corrected Velocity with Error Bars

Next, I numerically integrated the corrected velocities from $t = 0$ onward to get the altitude (Column E). I assumed that my subtracting took care of the offset term in Equation (1.20), and I used my calculated standard deviation and the sample rate of 50 SPS to calculate the random-walk term in the same equation (Column F). I restarted the time to 0 at both $t = -2$ and $t = 0$ to show how the term grows with time. See Figure 4 for results. The error bars are plotted on the graph but the growth is so slow that they are hard to see. A simple plot of Column F versus Column A would show the growth in uncertainty better.

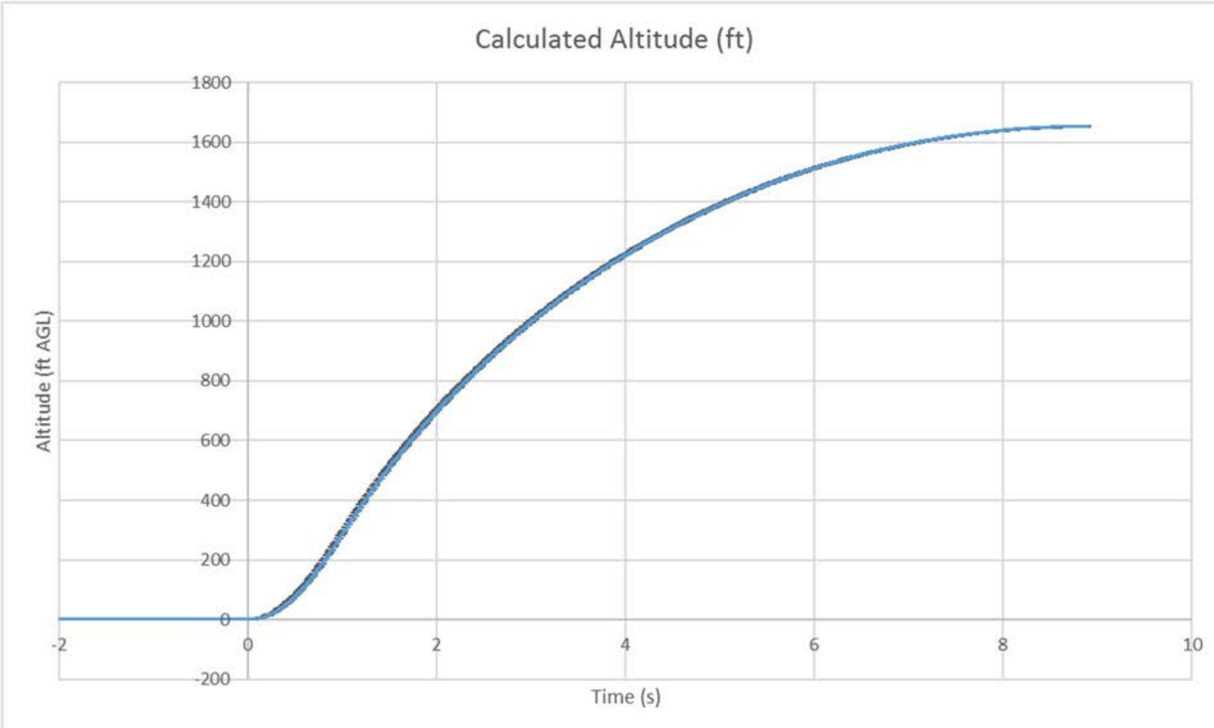

Figure 4: The Numerically Integrated Altitude with Error Bars from the Corrected Velocity

Finally, I took the numerical derivative of Column C to get the acceleration in Column G. I calculated the standard deviation for the first two seconds when I know the acceleration was 0 to get my standard deviation (Cell G1) and used it for my confidence interval and error bars for the acceleration (Column H). The noise-magnifying effect of taking numerical derivatives is very obvious in Figure 5.

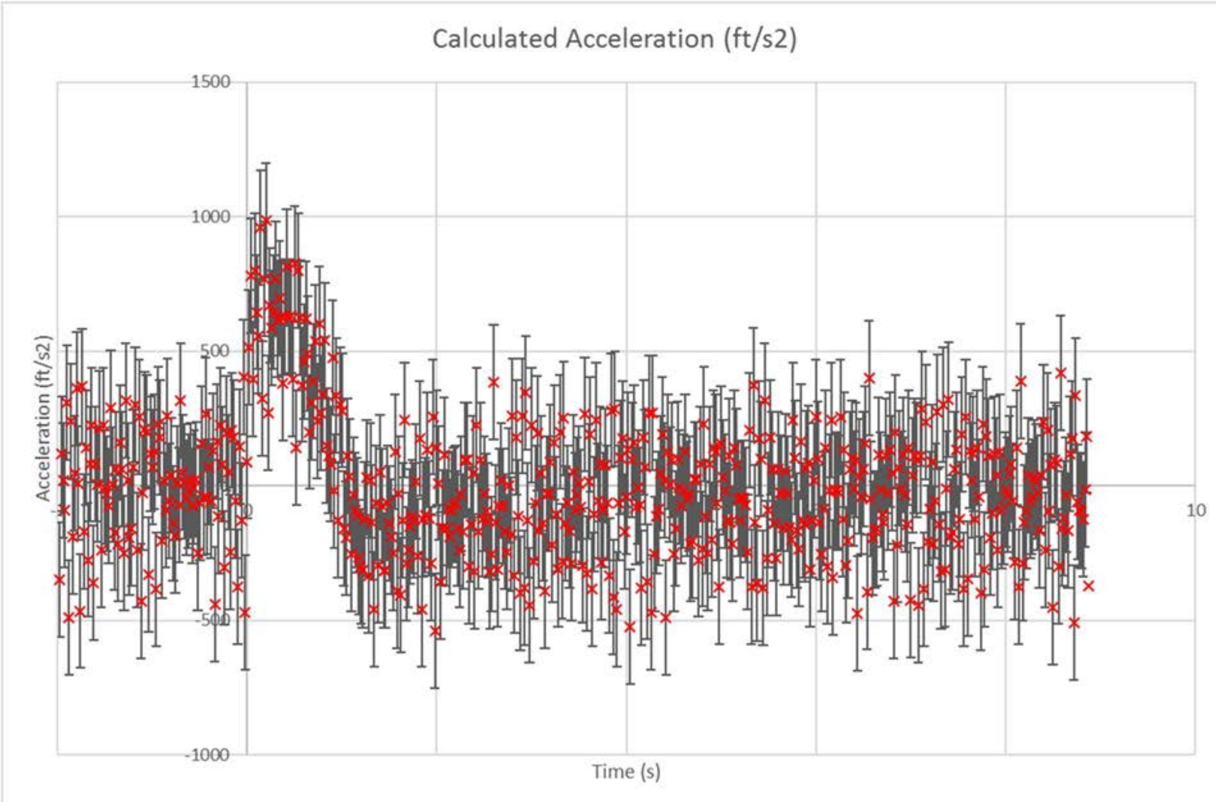

Figure 5: The Numerically Differentiated Acceleration with error bars from the Corrected Velocity

It's clear that the errors in taking numerical derivatives are huge. If you use any of the techniques for reducing the noise in a derivative, such as lowpass filtering or fitting a spline, you can use the same method of calculating the standard deviation on the processed signal for a stretch where you know the signal isn't changing to get error estimates. Just be aware that those techniques can remove some of the actual signal as well as the noise, and the calculated error bars will not reflect the missing signal. In fact, most of the techniques we have described are for errors or noise that is random with a zero mean. They cannot detect systematic errors, and you have to be aware that they can't.