


ON ELECTRIC PROPULSION THRUST STANDS: COMPREHENSIVE UNCERTAINTY ANALYSIS AND MODERNIZED CONCEPTUAL DESIGN



David Benjamin Gomez
NSTRF VTE | Summer 2020

Table of Contents

1.	Introduction	2
1.1.	Background	2
1.2.	Governing equations	2
1.3.	Owens thrust stand	7
2.	Uncertainty Analysis	10
2.1.	Classification	10
2.2.	Treatment	12
2.2.1.	Random uncertainties	12
2.2.2.	Systematic uncertainties	15
2.3.	Propagation and combination	16
2.4.	Thrust stand calibration	18
2.4.1.	Response variable and predictor variable	19
2.4.2.	Confidence and prediction intervals for uncertainty analysis	21
2.5.	Thrust measurements	28
2.6.	Random uncertainty	29
2.7.	Systematic uncertainty	29
2.7.1.	Thermal effects	30
2.7.2.	Thrust stand inclination	31
2.7.3.	Misalignment of thrust vector	32
2.7.4.	Other sources	33
2.8.	Total uncertainty	33
3.	Conceptual Design	34
3.1.	Brief configuration trade study	34
3.1.1.	Preliminary design constraints	36
3.1.2.	Torsional pendulum	36
3.1.3.	Hanging pendulum	37
3.1.4.	Inverted pendulum	37
3.1.5.	Displacement-mode vs null-mode	38
3.2.	Preliminary design recommendations	39
4.	Conclusions	40

1. Introduction

1.1. Background

Due to the low thrust-to-weight ratios characteristic of all electric propulsion (EP) devices, obtaining accurate measurements of thrust is a difficult task, despite its importance for performance analysis, thruster qualification, and flight program planning. Most state-of-the-art electric propulsion testing facilities employ pendulum thrust stands—whereby the thrust is measured indirectly, or inferred, from the motion of a pendulum mechanism. These devices have exhibited excellent performance across a wide range of EP thruster sizes and thrust levels. Pendulum thrust stands typically assume one of three forms: torsional pendulums (TP), hanging pendulums (HP), and inverted pendulums (IP). There is no pendulum mechanism that is best suited for all thrusters, as each exhibits their own unique advantages and challenges, and no single pendulum mechanism can be attributed as inherently more accurate than any other, as the uncertainty quantification depends largely on the details of the mechanical design and operational practices.

NASA JPL’s primary electric propulsion testing facility, Owens, which is currently employing an inverted pendulum thrust stand, is leading the development and qualification of the Hall-Effect Rocket with Magnetic-Shielding (HERMES) beginning with the technology demonstration and engineering development units. These thrusters, which are approaching the 100 kg class, are the largest and hottest thrusters that the Owens thrust stand has ever accommodated, and they introduce additional challenges to the thrust measurement process and have invoked unconventional techniques to mitigate unduly large thrust stand uncertainties.

With the ultimate goal of developing an entirely new modernized thrust stand designed specifically for the high-power EP test campaigns currently queued for the Owens facility, in this report, we express critics of the existing thrust stand as well as a proposed conceptual design for the new thrust stand. In an effort to standardize uncertainty analysis techniques within the EP community, we also provided a detailed tutorial on uncertainty analysis techniques for EP thrust stands. This report is organized as follows. In Section 1.2 and Section 1.3, we provide some background information on the dynamics of pendulum thrust stands and an overview of the current thrust stand design for the Owens facility.¹ In Section 2, we provide an uncertainty analysis tutorial for EP thrust stands, such as the one employed by Owens. In Section 2.8, we present our proposed conceptual design for a modernized thrust stand for high-power EP testing with several design recommendations. Finally, in Section 4, we conclude the work.

1.2. Governing equations

The three forms of pendulum thrust stands (torsional, hanging, and inverted) all share a common aspect: they are rotational systems and thus are governed by the rotational spring-mass-damper equation of motion.

$$I\ddot{\theta} + c\dot{\theta} + k_{\theta}\theta = \sum\tau \quad \text{Equation 1}$$

¹These sections may be skipped for the reader already familiar with the dynamics of pendulum thrust and the Owens thrust stand.

$$\left\{ \begin{array}{l} l: \text{moment of inertia} \\ c: \text{damping constant} \\ k_\theta: \text{angular spring constant} \\ \theta: \text{angular displacement} \\ \sum \tau: \text{sum of torques about pivot} \end{array} \right.$$

The fundamental difference between each pendulum type is the effect that gravity has on the dynamics of the system. Shown in Figure 1 are schematics of the dynamics of torsional, hanging, and inverted pendulums.




Figure 1. Schematics of (a) torsional, (b) hanging, and (c) inverted pendulum thrust stands.

As depicted in the torsional pendulum in Figure 1a, the force due to gravity is orthogonal to the XY plane of rotation, and the motion of the pendulum is unaffected by gravity. Thus, the sum of torques about the pivot is simply the applied force times the length of the pendulum².

$$\begin{aligned} \tau_F &= FL \\ \Rightarrow \\ \sum \tau_{TP} &= FL \end{aligned} \quad \text{Equation 2}$$

In contrast, in the hanging pendulum depicted in Figure 1b, we see that the force due to gravity can be decomposed into two components, one parallel to the pendulum arm, and the other orthogonal. The parallel component of the force due to gravity is absorbed via the tension in the pendulum arm, but the orthogonal component creates a torque about the pivot that counteracts the torque due to the applied force. Thus, the sum of torques about the pivot is the *difference* between the torque due to the applied force and the torque due to gravity.

$$\begin{aligned} \tau_F &= FL \\ \tau_g &= mgL \sin(\theta) \approx mgL\theta \end{aligned}$$

² Note that the applied force is assumed to be perpendicular to the pendulum. Hanging and inverted pendulum thrust stands are often constructed using parallel linkages, which keep thrust vector horizontal throughout the range of motion. State-of-the-art thrust stands maintain very small deflections (less than 1 degree), so we can assume the thrust is orthogonal to the pendulum arm throughout the range of motion without significant error—more on this later.

$$\begin{aligned} & \Rightarrow \\ \sum \tau_{HP} &= FL - mgL\theta \end{aligned} \quad \text{Equation 3}$$

Similarly, from the inverted pendulum in Figure 1c, we see that the gravitational force decomposition yields a component parallel to the pendulum arm, which is absorbed via the compression in the arm, and a component perpendicular to the pendulum arm. This time however, the orthogonal component is in the same direction as applied force and amplifies the deflection. Thus, the sum of torques about the pivot in this case is the *sum* of the torques due to the applied force and due to gravity.

$$\begin{aligned} & \tau_F = FL \\ \tau_g &= mgL \sin(\theta) \approx mgL\theta \\ & \Rightarrow \\ \sum \tau_{IP} &= FL + mgL\theta \end{aligned} \quad \text{Equation 4}$$

Using our expressions for the sum of torques for each type of pendulum, we can re-write the general equation of motion for pendulum thrust stands (Equation 1) as follows.

$$I\ddot{\theta} + c\dot{\theta} + k_{\theta}\theta = \begin{cases} \sum \tau_{TP} & \text{for torsional pendulums} \\ \sum \tau_{HP} & \text{for hanging pendulums} \\ \sum \tau_{IP} & \text{for inverted pendulums} \end{cases}$$

\Leftrightarrow

$$I\ddot{\theta} + c\dot{\theta} + k_{\theta}\theta = \begin{cases} FL & \text{for torsional pendulums} \\ FL - mgL\theta & \text{for hanging pendulums} \\ FL + mgL\theta & \text{for inverted pendulums} \end{cases} \quad \text{Equation 5}$$

Notice that for hanging and inverted pendulums, there is a term on the right-hand-side (RHS) that shares a θ with the angular spring constant term on the left-hand-side (LHS). By combining these θ -terms on the LHS, we can see another way of interpreting the effect of gravity on the thrust stand dynamics. Instead of construing the effect of gravity as an additional torque, since we are assuming small angles, the effect of gravity can also be interpreted as either increasing the *effective* angular spring constant for hanging pendulums or decreasing the effective spring constant for inverted pendulums. By defining the *effective angular spring constant* k as the sum of the (original) angular spring constant and the gravitational spring constant, we see that the only difference between the governing equations of the three types of pendulum thrust stands is in the expression for the effective angular spring constant.

$$\begin{aligned} I\ddot{\theta} + c\dot{\theta} + k\theta &= FL \\ & \text{where} \\ k &= \begin{cases} k_{\theta} & \text{for torsional pendulums} \\ k_{\theta} + mgL & \text{for hanging pendulums} \\ k_{\theta} - mgL & \text{for inverted pendulums} \end{cases} \end{aligned} \quad \text{Equation 6}$$

Equation 6 will be useful during our configuration trade study in Section 3.1. For torsional pendulum thrust stands, gravity and the system dynamics are decoupled, and the effective angular spring constant is simply the original angular spring constant. For hanging pendulums, the restoring torque due to gravity increases with—and counteracts—the deflection of the thrust stand, which can be construed as a higher effective angular spring constant. Finally, for inverted pendulums, the torque due to gravity also increases with—but amplifies—the deflection of the thrust stand, which can be construed as a lower effective angular spring constant.

In steady-state, the angular acceleration and angular velocity terms equal zero, and the steady-state deflection θ_{ss} for each pendulum type can be obtained with the following expression.

$$\begin{aligned}
 k\theta_{ss} &= FL \\
 &\Rightarrow \\
 \theta_{ss} &= \frac{FL}{k} = \begin{cases} FL & \text{for torsional pendulums} \\ FL/(k_\theta + mgL) & \text{for hanging pendulums} \\ FL/(k_\theta - mgL) & \text{for inverted pendulums} \end{cases} \quad \text{Equation 7}
 \end{aligned}$$

All else constant, we see from Equation 7 that hanging pendulums exhibit smaller steady-state deflections than inverted pendulums. One important characteristic of inverted pendulums is at once a benefit and a risk factor: inverted pendulums are unstable. The negative sign in the denominator of the expression for the steady-state deflection of inverted pendulums means that as mgL approaches k_θ , the deflection goes to infinity. Careful balancing of the angular spring constant (adjusted with linear or torsional springs) and the mass of the thrust stand (adjusted using mass ballast) can afford relatively small effective angular spring constants and correspondingly large deflections, which is beneficial for increasing the sensitivity and resolution of thrust stands. As the deflections become excessively large however, the small angle assumption breaks down and significant systematic errors might incur. State-of-the-art inverted pendulum thrust stands typically deflect less than 1° , which afford reasonably low errors due to the small angle assumption, i.e., $(\sin\theta - \theta)/\theta \leq 0.13\%$ for $\theta \leq 5^\circ$.

With our compact expression for the pendulum dynamics for each pendulum thrust stand (Equation 6), we can write the equation of the motion in standard form to obtain expressions for the damping coefficient and the natural frequency.

$$\begin{aligned}
 I\ddot{\theta} + c\dot{\theta} + k\theta &= FL \\
 &\Leftrightarrow \\
 \ddot{\theta} + \frac{c}{I}\dot{\theta} + \frac{k}{I}\theta &= \frac{FL}{I} \quad \text{Equation 8}
 \end{aligned}$$

The undamped natural frequency and damping coefficient are defined as follows.

$$\omega_n \equiv \sqrt{\frac{k}{I}} \quad \text{Equation 9}$$

$$\zeta \equiv \frac{c}{2} \sqrt{\frac{1}{Ik}} \quad \text{Equation 10}$$

⇒

$$\ddot{\theta} + 2\zeta\omega_n\dot{\theta} + \omega_n^2\theta = \frac{FL}{I} \quad \text{Equation 11}$$

If we model the thrust stands as a rod-mass system, then $I = mL^2$. The solution to is $\theta = \theta(t)$, which is temporal response of the thrust stand subject to some $F = F(t)$. Since the determination of the effective spring constant, moment of inertia, and damping constant is not trivial, end-to-end calibrations are conducted that statistically “learn” the relationship between the steady-state deflection of the thrust stand and the applied force giving $F = F(\theta)$, or more practically (as we will see), $F = F(x)$.

Going back to our expression for the steady-state deflection (Equation 7), the *angular sensitivity* is defined as the steady-state angular displacement over the applied force.

$$\theta_{ss} = \frac{FL}{k}$$

⇒

$$S_\theta \equiv \frac{\theta_{ss}}{F} = \frac{L}{k} \quad \text{Equation 12}$$

However, it often much more convenient to measure the linear deflection as opposed to the angular deflection. Assuming the deflection is made at a distance L from the pivot point, the linear steady-state deflection would be as follows³. Note that we again assume small angles.

$$x_{ss} = L \sin(\theta_{ss}) \approx L\theta_{ss} = \frac{FL^2}{k} \quad \text{Equation 13}$$

The *linear sensitivity*, which is the most commonly reported form of thrust stand sensitivity, is defined as the steady-state linear displacement over the applied force.

³ Note that the depictions in Figure 1 show deflections in the different directions, but for simplicity’s sake, we will henceforth assume that all “linear deflections” are in the “x”-direction.

$$S_x \equiv \frac{x_{ss}}{F} = \frac{L^2}{k} \quad \text{Equation 14}$$

From Equation 14, we see that in general we can achieve higher sensitivities by designing larger thrust stands or by using lower effective angular spring constants. This is why inverted pendulums are appealing: because the effective angular spring constant of inverted pendulums is the difference between two terms ($k_{IP} = k_\theta - mgL$), we can design k_θ , m , and L , such that the denominator of Equation 14 is arbitrarily small, and the linear sensitivity is correspondingly large. Care must be taken to ensure that $k_\theta > mgL$, lest the thrust stand become unstable.

1.3. Owens thrust stand

The Owens thrust stand currently employs a displacement-type, parallel-linkage, inverted pendulum thrust stand. The thrust stand is shown in Figure 2, and a schematic is shown in Figure 3.



Figure 2. NASA JPL's Owens thrust stand (with the magnetically shielded H9 onboard).




Figure 3. Schematic of NASA JPL’s Owens thrust stand (with notional Hall thruster onboard).

There are plenty of excellent qualities about this thrust stand that make it a “state-of-the-art” inverted pendulum thrust stand. For example, Owens’ thrust stand features an in-situ calibration system composed of a pulley-weight system; it features closed-loop inclination, damping, and thermal controllers, and uses a high-precision Linear Variable Differential Transducer (LVDT) to monitor the thrust stand displacement. However, there are also several unsavory design features that deserve attention; these are listed and briefly described next.

1. **Thruster mounting subsystem.** The thruster mounting system consists of several interacting components: the thruster, the power line tower, the thruster mounting plate, the thruster cooling column. Sitting on top of a single column, the thruster mounting plate is a thin (4 mm) aluminum plate that supports the weight of the thruster and power line tower and all the associated harnessing. The power line tower (upper right of Figure 3) is ~0.7 m tall 80/20 assembly and weighs enough such that the thruster mounted plate tilts, sometimes to the point where it contacts the stationary copper cooling shroud. This tilting can cause several issues: it can add significant friction to the motion of the thrust stand, reducing the deflection, thus leading to erroneously low reported thrusts; additionally, it can cause a misalignment between the thrust vector and the intended direction of motion, again, leading to erroneously low reported thrusts. Furthermore, the weight of the power line tower pulls the center of mass towards the right (in Figure 3). Were the center of mass to drift outside the parallel pendulum arms, the thrust stand would topple over. This failure mode is made all the more probable given the narrowness of the parallel linkage pendulum, spanning roughly the axial “depth” of the HERMeS, as indicated in Figure 3. While the weight of the HERMeS has thus far kept the center of mass above the pendulum system, this may not be the case for smaller thrusters. Also related to the power line tower are the power harnesses themselves: they are in-line with the direction of thrust and as they heat up and cool down (due to resistive or Ohmic heating) they expand and contract, pushing and pulling on the thrust stand.

2. **Mechanical subsystems.** The majority of the Owens thrust stand is fabricated from thin Aluminum structures. Given the inevitable temperature variability inherent in all thrust stands, Aluminum may not be the best material choice given its rather large coefficient of thermal expansion. There are also mechanical structures that are inherently weak in their design. For example, the inclination pivot point is a 1/4"-20 threaded rod that has been carved on either side to create a literal weak point in the rod about which the entire thrust stand pivots. There is simply no easy way to determine the reliability of this mechanical feature. What is more, the location of the inclination pivot is more or less in the center of the thrust stand. Given that the thrust stand center of mass is likely to drift left or right, it is not certain which way the thrust stand will tilt if subject to inclination drifts. Furthermore, in order to provide higher sensitivity, the steel strip flexures are very thin, which raises concern for column buckling as the weight of the thrusters under test continues to approach the 100-kg class. There is also the aluminum wave spring, which was bent manually and is poorly characterized. The stiffness of this wave spring and flexures has not been determined, neither theoretically nor experimentally, and perhaps more importantly, the linear regimes of these flexible components have not been established. Per Equation 7, as the mass (and thrust) of the thruster increases, the steady-state deflection also increases, which makes the knowledge of the linear regime of the flexible components all the more necessary.
3. **Electronic subsystems.** The thrust stand is an electro-mechanical device that houses 4 primary electronic components: the displacement sensor, the active damping control, the calibration spool motor, and the inclination motor. The displacement sensor is a Linear Variable Differential Transducer (LVDT) that is positioned at the end of a ~0.5 m long rod attached to the upper plate of the inverted pendulum. Since the (linear) thermal expansion is proportional to the length of the material, temperature variations inside the copper shroud can cause expansions and contractions of the LVDT arms which are magnified due to its length. Since the Owens thrust stand operates in displacement mode, artificial deflections are interpreted as artificial thrust, and the reported thrust measurements can be erroneously high or low due to thermal expansions/contractions of the LVDT arm. The calibration spool motor also deserves some attention. User stories claim that the deployment and retraction of the calibration weights cause unnecessary vibrations from the stepper motor, which cause the calibration weights to swing erratically. Also, the inclination motor warrants consideration. Unlike the calibration pulley motor, the inclination control system is a lead screw assembly driven by a conventional DC motor. User stories claim that the control authority of the DC motor is sometimes not sufficient to raise/lower the lead screw or tilt the thrust stand.

Many of the previously described design weak points are amplified by hotter and heavier thrusters with higher thrust levels. Given the test campaigns in the queue for the Owens facility, a re-design of the thrust stand is warranted. It is clear that many of these design choices can have an effect on thrust measurement accuracy, but how shall this be quantified? An important aspect of assessing the quality of a thrust stand design is the quantification of its accuracy via an uncertainty analysis. In an effort to standardize practices for conducting uncertainty analysis, the next section constitutes a tutorial on the subject.

2. Uncertainty Analysis

The performance of thrust stands for electric propulsion are described by several metrics, as depicted in Figure 4.

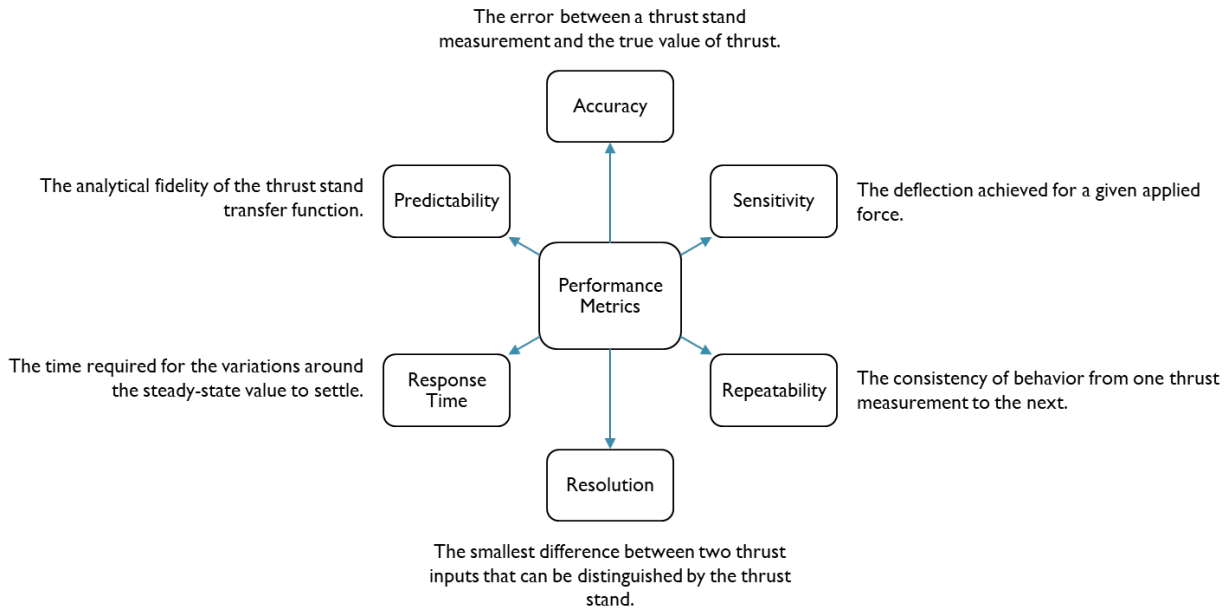


Figure 4. Performance metrics of thrust stands for electric propulsion testing.

While all are important considerations in the design and analysis of thrust stands, in this section we will focus on quantifying the accuracy of the Owens thrust stand, where the accuracy is expressed in terms of uncertainty in the thrust measurements. The goals of this section are to (1) provide a tutorial on uncertainty analysis for electric propulsion thrust stands in general and (2) apply the uncertainty analysis framework to the displacement-mode inverted pendulum thrust stand of the Owens chamber. Section 2 is organized as follows. We will first consider the classification and treatment of common experimental uncertainties, then their propagation and combination. We will then examine the important role of the calibration process and see how it simplifies the uncertainty analysis. Finally, we will apply this framework to an EP thrust stand.

2.1. Classification

In the context of experimental measurements, the words “uncertainty” and “error” are often used interchangeably, but there is a useful distinction between the two. *Uncertainty* refers to a *range* of possible values within which one believes the true value of the measurement is likely to reside. Consider the following statement: the measured thrust is 700 ± 10 mN. This translates to “the best estimate of the true (but unknown) thrust is 700 mN, but it is likely to be as high as 710 or as low as 690.” Note that uncertainties are usually symmetric about our best estimate of the true value, but they need not be—i.e., the positive uncertainty can be different from the negative uncertainty. A measurement *error* refers to the *difference* between the best estimate of the true value and the true value itself. Notice that, by definition, an error requires knowledge of the true value of the underlying measurement. A notional depiction of the distinction between errors and uncertainties is provided in Figure 5.

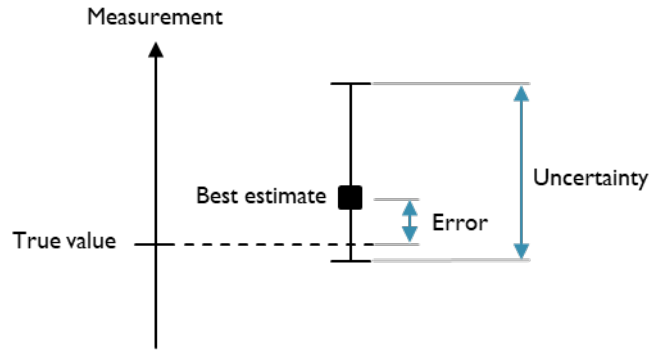


Figure 5. Depiction of experimental uncertainties and errors.

Uncertainties can be classified into two groups: *random* and *systematic*. Random uncertainties are normally distributed with a mean of zero, whereas systematic uncertainties tend to bias the measurement in one direction or the other. The familiar depiction of random and systematic uncertainties is shown in the left panel of Figure 6; to highlight the difficulty in handling systematic uncertainties, a more realistic version is shown on the right.

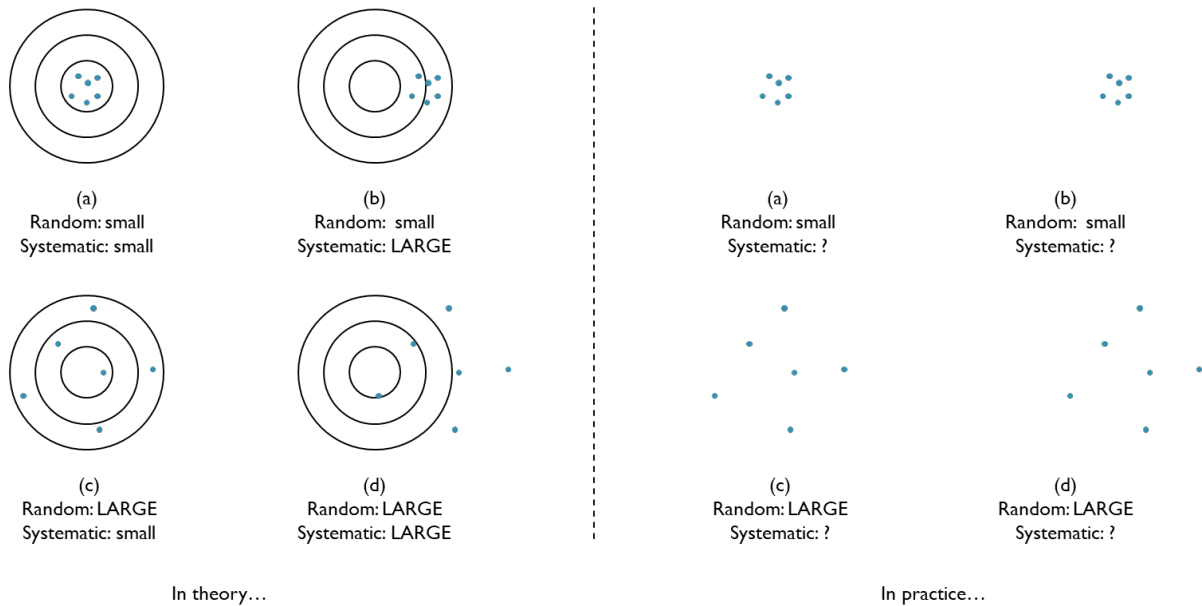


Figure 6. Random and systematic uncertainties in theory and in practice, depicted via bullet holes at target practice.

In the vast majority of experimental measurements, including the measurement of thrust in EP testing, one does not know the true value of the underlying measurement. The treatment of random uncertainties is thus fundamentally different from that of systematic uncertainties. **Random uncertainties are usually quantified via statistical methods and lend themselves to minimization through repeated sampling and averaging; however, for the reasons just discussed, systematic uncertainties are harder to identify and evaluate, so their treatment is not as straightforward, as will be discussed next.**

How one handles measurement infidelity is determined by its classification. The following classification chart shown in Figure 7 will be useful for categorizing such sources of measurement infidelity.

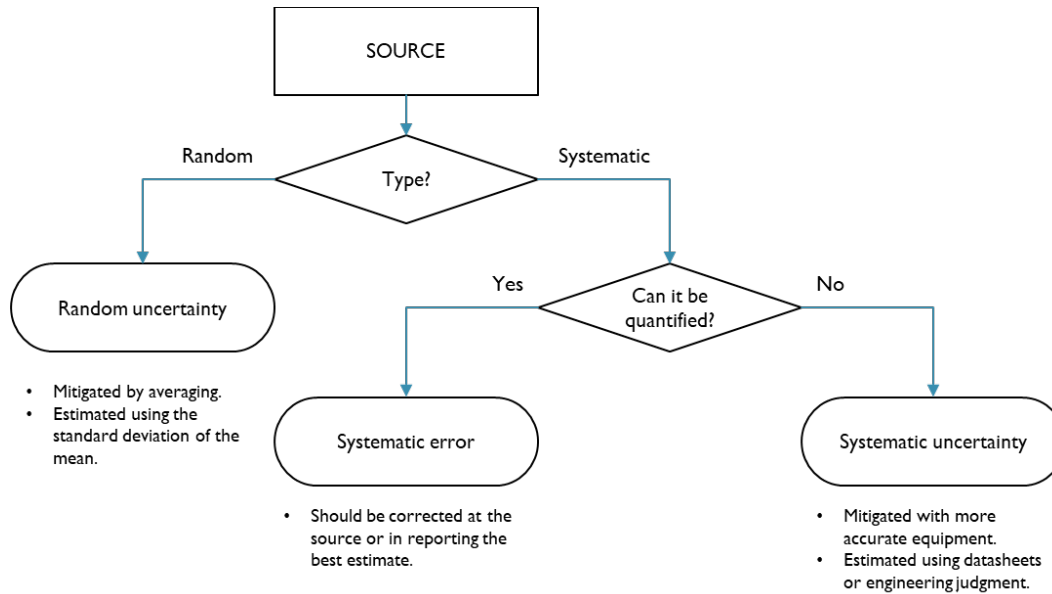


Figure 7. Classification tree of potential sources of measurement infidelity.

When presented with a potential source of measurement infidelity, first determine whether it is random or systematic. If it is random, then it can only be *random uncertainty*—there is no “*random error*” in this context. On the other hand, if it is systematic, can it be quantified? If so, then it is a *systematic error*. If it cannot be quantified, it should be rendered as a *systematic uncertainty*. As an illustration of how to use the classification tree, consider the following potential sources of measurement infidelity with a DM-IP thrust stand.

- **Mechanical vibrations.** Mechanical vibrations manifest as noise in the thrust stand response. Noise is a type of *random uncertainty*. It can be mitigated with sufficient averaging and estimated using the standard deviation of the mean (discussed further in the next sub-section).
- **Thrust vector misalignment.** If the thrust vector were misaligned, the thrust stand would only measure the projection of the thrust vector parallel to the line of motion. A misaligned thrust vector would cause consistent and systematic underestimation. Since we do not know the exact misalignment of the thrust vector, it cannot be quantified, only estimated, and will be rendered as a *systematic uncertainty*. Its treatment will be discussed in the next sub-section.
- **A calibration weight was left on the thruster pulley during thruster operation.** The erroneous thrust measurement would be systematically high, but since we know the weight of the calibration loads, the measurement infidelity can be quantified; thus, it is a *systematic error*. It should be corrected at the source (e.g., remove the weight) or corrected for in the reporting of the best estimate (e.g., subtract the weight of the calibration load from the best estimate of the reported thrust).

2.2. Treatment

Following the classification of sources of measurement infidelity, the next question is how to treat them.

2.2.1. Random uncertainties

Consider first *random uncertainties*. Random uncertainties are treated using statistical methods, so a brief review of relevant statistics is provided next.

Suppose we take a sample of N measurements of some quantity x . The best estimate \hat{x} of the true value is the *mean of the sample*, \bar{x} , computed as follows.

$$\hat{x} = \bar{x} \equiv \frac{1}{N} \sum_{i=1}^N x_i \quad \text{Equation 15}$$

The *standard deviation of the sample*⁴ is determined as follows.

$$s_x \equiv \sqrt{\frac{\sum_{i=1}^N (x_i - \bar{x})^2}{N - 1}} \quad \text{Equation 16}$$

Notice the $N - 1$ in the denominator of the standard deviation of the sample. Why not just N ? Any sum of squares has associated with it a certain number of degrees of freedom ν . This is the number of independent pieces of information involving the number used to compile the sum of squares.

$$dof = \nu \equiv \# \text{ observations} - \# \text{ estimated parameters} \quad \text{Equation 17}$$

How many independent observations go in the calculation of the standard deviation of the sample? Well, we have N observations: x_1, x_2, \dots, x_N . But we also know one parameter: \bar{x} . If we know \bar{x} and were only given $N - 1$ observations, then we can immediately solve for the last observation. It is no longer independent. So, we have $N - 1$ degrees of freedom (or $N - 1$ independent observations) since one of the observations can be computed using the sample mean \bar{x} and the other $N - 1$ observations.

Now consider Equation 16 more carefully. Each measurement x_i has some residual $(x_i - \bar{x})$ associated with it. We cannot simply take the average of the residuals, as many would cancel out, so we take the sum of the squares of the residuals to estimate the *average deviation* of all the measurements from the mean of our sample. However, and this is crucial, **our answer for the best estimate $\hat{x} = \bar{x}$, represents a judicious combination of all N measurements in our sample, and we have every right to believe that it will be more reliable than any one of the measurements taken alone in estimating the true value of the underlying measurement.** If we repeat this procedure many times: (1) collecting N measurements of the same quantity x and (2) computing the *sample mean* \bar{x} , we will have a *distribution of sample means*. Statistical theory, more specifically, the central limit theorem, tells us that this distribution is centered on the true value of the underlying measurement x_t and has a standard deviation, called the *standard deviation of the mean*, computed as follows.

$$\sigma_{\bar{x}} \equiv \frac{s_x}{\sqrt{N}} \quad \text{Equation 18}$$

⁴ The reader will no doubt find other names for the statistical metrics discussed in this paper. We have chosen terminology that is most common and that accurately describes the underlying metric.

A depiction of the distinction between the *standard deviation of the sample* s_x and the *standard deviation of the mean* $\sigma_{\bar{x}}$ is provided in Figure 8.




Figure 8. Distinction between the *standard deviation of the sample* s_x (left panel) and the *standard deviation of the mean* $\sigma_{\bar{x}}$ (right panel).

Notice how on the left panel of Figure 8 the distribution is centered on the mean of the j^{th} sample \bar{x}_j . Now if we take many samples and compute many sample means, we will get a distribution of sample means, and that is depicted on the right panel of Figure 8. Observe that the distribution of the sample means is centered on the true value of the underlying measurement x_t .

Of course, the true value of the underlying measurement x_t is unknown, so when we compute a *sample mean* \bar{x} , we have no idea how close it is to x_t . However, we can use the *standard deviation of the mean* $\sigma_{\bar{x}}$ to construct confidence intervals about our *sample mean* \bar{x} , which will tell us the probability that the true value x_t lies within our intervals, and is computed as follows.

$$CI = \bar{x} \pm t_{v,\alpha} \sigma_{\bar{x}} \quad \text{Equation 19}$$

Here, t is a value that depends on the number of degrees of freedom v in the sample and desired confidence level $(1 - \alpha)\%$ and must be looked up in a *Two-Tailed Student t -Distribution table*.⁵ Notice from Equation 18 that the standard deviation of the mean decreases with the square root of N . So, the more observations we include in our sample, the tighter our confidence intervals will be.

The 95% CIs for two hypothetical samples are shown below. Notice that for one sample, the confidence interval does not contain the true value of the underlying measurement and for the other, it does.

⁵ The t -distribution is like Z -distribution (that is, the standard normal distribution) but shorter and fatter. As the number of degrees of freedom increases, the t -distribution approaches the Z -distribution (for $N > 30$ they are essentially the same).




Figure 9. Two hypothetical confidence intervals for two sample means.

The 95% confidence interval translates to following: “Upon repeated sampling, 95% of the intervals constructed in this way will contain the true but unknown value of the underlying measurement.” In other words, if we were to take 100 samples of N measurements and compute 100 sample means and 100 confidence intervals, roughly 95 of them would contain the true value of the underlying measurement. And the best part (or the worst part) is that we will never know. We can express this as a random uncertainty U_{rand} in the following form.

$$x_t = \bar{x} \pm U_{rand}$$

with

$$U_{rand} = t_{v,\alpha} \sigma_{\bar{x}}$$

Random uncertainties should only be ascertained in this statistical way. A single data point on its own does not lend itself to random uncertainty analysis.

2.2.2. Systematic uncertainties

Now consider the pernicious treatment of systematic uncertainties. There is no simple theory for how to handle systematic errors. The fundamental problem is to decide how to estimate the systematic uncertainty and then how to combine it with random uncertainty. On estimating systematic uncertainty, recall the Uncertainty Classification Chart in Figure 7. If we have identified a source of systematic measurement infidelity and we can quantify it, then we have identified a *systematic error* and it should be corrected either at the source or in the reporting of the best estimate. However, if it cannot be quantified and is thus a *systematic uncertainty*, how should we handle it? Experimental convention is to model the systematic uncertainty as a type of random uncertainty, i.e., normally distributed with a mean of zero. For example, most (commercial) experimental instrumentation come with a datasheet that reports the accuracy of the instrument, e.g., the thermocouple is accurate to ± 5 °C. What this is telling us is that the probe is likely biased, but the extent of the biased cannot be resolved beyond ± 5 °C, i.e., the systematic error is uncertain \rightarrow systematic uncertainty. Recall that systematic errors have been quantified, e.g., the thermocouple has a bias of 3 °C. This is a luxurious scenario. If quantified in this way, systematic errors should be corrected immediately, either at the source (e.g., repeating the measurement with an un-biased thermocouple) or in the reporting of the best estimate (e.g., subtracting 3 °C from the reported temperature). In contrast, systematic uncertainties are just that: they are uncertain. The instrumentation may be biased, or it may not be. In this scenario, we assume that the systematic uncertainty is normally distributed (much like random

uncertainties). We address how to propagate and combine systematic and random uncertainties in the next section.

2.3. Propagation and combination

The most common approach to the propagation of uncertainties is as follows. Suppose that we have a function φ of several variables $\omega_1, \dots, \omega_n$ —i.e., $\varphi = \varphi(\omega_1, \dots, \omega_n)$. Each variable ω_i has associated with it a total uncertainty U_{tot,ω_i} , which itself may be comprised of systematic and random uncertainties. If the uncertainties in $\omega_1, \dots, \omega_n$ are, and this is crucial, *independent and normally distributed*, then the total uncertainty in φ is as follows.

$$U_{tot,\varphi} = \sqrt{(U_{\varphi,\omega_1})^2 + \dots + (U_{\varphi,\omega_n})^2} \quad \text{Equation 20}$$

Equation 20 says that the total uncertainty in φ is the Sum in Quadrature (SIQ) of *the partial uncertainties in φ due to the uncertainties in $\omega_1, \dots, \omega_n$* . The partial uncertainty in φ due to the uncertainty in ω_i is given as follows.

$$U_{\varphi,\omega_i} = \left. \frac{\partial \varphi}{\partial \omega_i} \right|_{\hat{\omega}_1, \dots, \hat{\omega}_n} U_{tot,\omega_i} \quad \text{Equation 21}$$

Equation 21 says that the partial uncertainty in φ due to the total uncertainty in ω_i is the partial derivative φ with respect to ω_i (this is a measure of the influence that ω_i has on φ) multiplied by the total uncertainty in ω_i . Note that the derivative term is evaluated at the best estimates of $\omega_1, \dots, \omega_n$. Finally, the total uncertainty in ω_i , as mentioned previously, may be comprised of both systematic and random uncertainties. If the systematic uncertainty U_{sys} and random uncertainty U_{rand} are *independent and normally distributed*, then the total uncertainty in the variable ω_i is as follows.

$$U_{tot,\omega_i} = \sqrt{(U_{sys,\omega_i})^2 + (U_{rand,\omega_i})^2} \quad \text{Equation 22}$$

An illustration of this nested application of uncertainty propagation and combination is shown below.

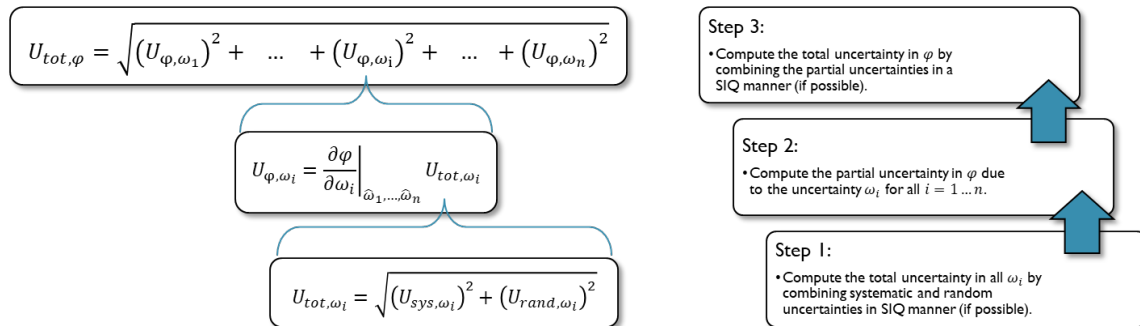


Figure 10. Propagation and combination of experimental uncertainties.

In general, a set of independent and normally distributed uncertainties can be combined in this Summation in Quadrature (SIQ) manner⁶. However, if the set of uncertainties cannot be deemed both *independent and normally distributed* then the uncertainties cannot be combined via a SIQ; they can only be combined via their direct sum. This is because the SIQ accounts for the possibility that some of the high measurement uncertainty would cancel with some of the low measurement uncertainty. For example, suppose we wish to compute the uncertainty in $\varphi = \omega_1 + \omega_2$ where $\omega_1 = 5 \pm 2$ and $\omega_2 = 5 \pm 3$. If there were no uncertainty, then the value of φ would always be 10, but what is the worst-case scenario? Suppose both ω_1 and ω_2 were measured at their lower uncertainty bounds, i.e., $\omega_1 = 3$ and $\omega_2 = 2$, in this case, $\varphi = 5$; or suppose they were measured at their upper uncertainty bounds, i.e., $\omega_1 = 7$ and $\omega_2 = 8$, in this case, $\varphi = 15$. In either worst-case scenario, the uncertainty in φ is the direct sum of the uncertainties in ω_1 and ω_2 .⁷

$$U_{tot,\varphi} = U_{tot,\omega_1} + U_{tot,\omega_2} = 2 + 3 = 5$$

These worst-case scenarios are not very likely however. If the uncertainties in ω_1 and ω_2 are *independent* (i.e., the value of one uncertainty does not affect the value of the other) and *normally distributed* (i.e., they have a normal distribution with a mean of zero), then there is a good chance that the uncertainty in both measurements will “cancel out.” The SIQ method of combining uncertainties accounts for this possibility and the resulting SIQ uncertainty is always smaller than the direct sum.

$$\begin{aligned} U_{tot,\varphi} &= \sqrt{(U_{tot,\omega_1})^2 + (U_{tot,\omega_2})^2} \\ &= \sqrt{2^2 + 3^2} \cong 3.6 \end{aligned}$$

In general, the direct sum can be considered an upper-bound on a combination of uncertainties. Expressed mathematically,

$$\begin{aligned} U_{tot,\varphi} &= \sqrt{(U_{\varphi,\omega_1})^2 + \dots + (U_{\varphi,\omega_n})^2} \\ &\quad \text{with} \\ U_{tot,\varphi} &\leq U_{\varphi,\omega_1} + \dots + U_{\varphi,\omega_n}. \end{aligned}$$

To summarize, one should use the SIQ method to combine sets of uncertainty only when the sets can be deemed *independent* and are *normally distributed*. When a source of uncertainty does not meet these criteria, it should be added to the total uncertainty via a direct sum.

⁶ This is also called Root Sum of Squares (RSS), but we will avoid this acronym so as not to confuse with Residual Sum of Squares, discussed shortly.

⁷ Note that the partial derivatives of F with respect to ω_1 and ω_2 are both equal to one.

2.4. Thrust stand calibration

Next, we will examine the important role of the calibration process and the effect it has on the uncertainty analysis. Theoretically, one could compute the effective angular stiffness (Equation 6) of a thrust stand by carefully evaluating all the flexible components and by knowing the pendulum mass and length. This process is difficult however because the laboratory testing of an electric thruster often requires complex interfaces between the thrust stand and the test facility (e.g., the propellant/coolant lines and electronic harnessing), which contribute parasitic stiffness to the system. It is much more practical to perform an end-to-end *calibration* of the entire thrust stand assembly. The primary purpose of the thrust stand calibration is to learn the relationship between the deflection (or sensor output) and the applied force. This relation is most commonly learned via linear regression, and as we will see, there are aspects of linear regression that greatly simplify the uncertainty analysis.

DM-IP thrust stands are calibrated by assuming the steady-state linear deflection of the thrust stand is proportional to the applied force, e.g., $F = k_x x_{ss}$. From Equation 14, we see that the constant of proportionality k_x is the inverse of the linear sensitivity S_x .

$$S_x \equiv \frac{x_{ss}}{F} \Rightarrow F = \frac{1}{S_x} x_{ss}$$

and knowing ...

$$F = k_x x_{ss}$$

\Rightarrow

$$k_x = \frac{1}{S_x}$$

Thrust stand calibration involves applying known forces F parallel to the line of motion and monitoring the deflection sensor output x of the thrust stand. The application of known forces is commonly performed with a pulley-weight system. The loading profile typically assumes one of the following forms: the applied calibration weights are increased incrementally (Figure 11a), or the weights are sequentially loaded and unloaded (Figure 11b). The latter has the advantage of correcting for (to some extent) the effect of zero drift, to be discussed shortly. Both methods yield the calibration curve shown in Figure 11c. Note that displacement sensors typically output a voltage that is converted to a displacement via the sensor sensitivity G (m/V). Also note that the profiles are often mirrored, i.e., the loads are stepped up then stepped down, which provides twice as many data points and is recommended practice—this is not shown in the figure to avoid excessive clutter.

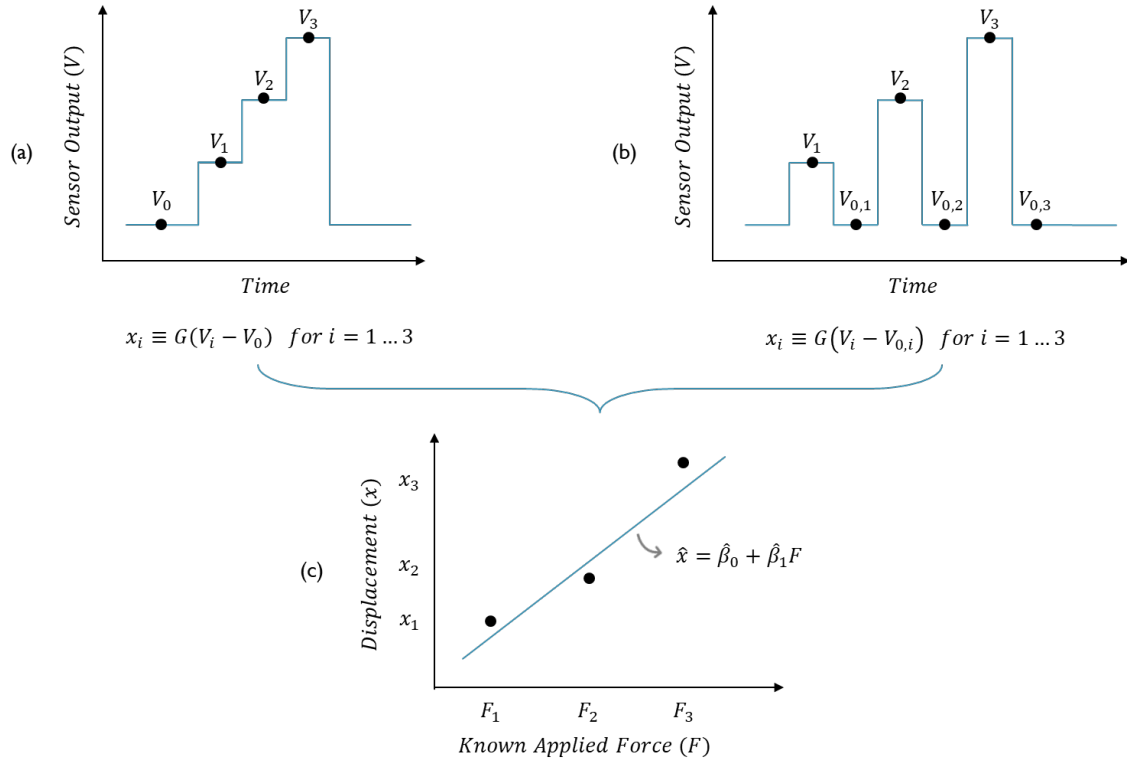


Figure 11. Typical calibration profiles of electric propulsion pendulum thrust stands. Not shown is the fact that the loading profiles are often mirrored, i.e., the loads are stepped up then stepped down, which provides twice as many data points and is recommended practice

2.4.1. Response variable and predictor variable

There is the delicate subject of which variable to use as the response variable and which to use as the predictor. There are two methods.

1. The first method puts the applied force on the y-axis (response) and the displacement on the x-axis (predictor)—i.e., it regresses F on x as follows.

$$\hat{F} = \hat{\beta}'_0 + \hat{\beta}'_1 x \quad \text{Equation 23}$$

At first glance, this method seems appropriate since our end goal is to predict the force given a value of displacement. However, least-squares regression assumes (or requires) that the predictor is fixed while the response is random. Equation 23 is the inverse of this statement since the predictor is the random variable (the displacement) while the response is the fixed variable (the known calibration loads)⁸. This method thus is referred to as *inverse regression*.

2. The second method has the displacement as the response and the applied force as the predictor—i.e., it regresses x on F , as follows.

⁸ There will usually be much more relative uncertainty in the displacement sensor output than the known loads due to, for example, random uncertainties from mechanical vibrations or electrical noise, as well as systematic uncertainties from thermal effects or distortions in the thrust stand.

$$\hat{x} = \hat{\beta}_0 + \hat{\beta}_1 F \quad \text{Equation 24}$$

We mentioned that least-squares regression assumes (or requires) that the predictor is fixed while the response is random. Is the predictor truly fixed? We cannot know the calibration weights with absolute certainty, but we can ensure that their relative uncertainty is much smaller than the relative uncertainty in the displacement sensor, and thus satisfy this requirement. As discussed in Section 2.2.1, this can be accomplished by repeated sampling of the calibration weights. Once we have the model trained, we need to invert Equation 24 to predict the force using the deflection as shown below. Hence, this method is known as *inverse prediction*.

$$F = -\frac{\hat{\beta}_0}{\hat{\beta}_1} + \frac{1}{\hat{\beta}_1} \hat{x} \quad \text{Equation 25}$$

Which should we use? The former method (inverse regression), while not supported by theory, has been found to work well in practice for highly linear datasets. The latter method (inverse prediction) is slightly unintuitive but has the advantage of having theoretical support. Being supported by statistical theory will allow us to leverage other aspects of linear regression, such as confidence and prediction intervals in quantifying the thrust stand random uncertainty (to be discussed in the next section). For these reasons, **it is recommended that the predictor be the known applied force while the response variable be the change in displacement sensor output (inverse prediction)**.

It is tempting to believe that the “slope” terms in Equation 23 and Equation 25 are inverses of each other, and that the best-fit line will be the same regardless of which method is used, but that is usually not the case.⁹ Consider the process of estimating the coefficients by the recommended *inverse prediction* method (Equation 24). If using least-squares regression, the best estimates of the coefficients are those that minimize the *Residual Sum of Squares* (RSS), with the residuals e_i being defined as the difference between the actual response x_i and the predicted response \hat{x}_i .

$$e_i \equiv x_i - \hat{x}_i \quad \text{Equation 26}$$

$$RSS \equiv \sum_{i=1}^n (e_i)^2 = \sum_{i=1}^n (x_i - \hat{x}_i)^2 = \sum_{i=1}^n (x_i - \hat{\beta}_0 - \hat{\beta}_1 F_i)^2 \quad \text{Equation 27}$$

The RSS is a quadratic function with two unknowns and admits a global minimum, which can be solved for the following closed-form expressions of the estimated coefficients.

$$\hat{\beta}_0 = \bar{x} - \hat{\beta}_1 \bar{F} \quad \text{Equation 28}$$

⁹ Due to the manner in which we measure thrust, the intercept coefficient is of little importance as it cancels out in the computation of thrust.

$$\hat{\beta}_1 = \frac{\sum_{i=1}^n (F_i - \bar{F})(x_i - \bar{x})}{\sum_{i=1}^n (F_i - \bar{F})^2} = \frac{\widehat{cov}(F, x)}{\widehat{var}(F)} \quad \text{Equation 29}$$

If we were to swap the regressor and predictor and perform *inverse regression* (Equation 23), we would get the following closed-form expressions for the estimated coefficients.

$$\hat{\beta}'_0 = \bar{x} - \hat{\beta}'_1 \bar{F} \quad \text{Equation 30}$$

$$\hat{\beta}'_1 = \frac{\sum_{i=1}^n (F_i - \bar{F})(x_i - \bar{x})}{\sum_{i=1}^n (x_i - \bar{x})^2} = \frac{\widehat{cov}(F, x)}{\widehat{var}(x)} \quad \text{Equation 31}$$

Observe that inverse of $\hat{\beta}'_1$ is not the same as $\hat{\beta}_1$ unless $\widehat{var}(F) = \widehat{var}(x)$, which is usually not the case.

2.4.2. Confidence and prediction intervals for uncertainty analysis

If we proceed with the theory-supported *inverse prediction* method and regress x on F , we can leverage certain aspects of linear regression to conduct our uncertainty analysis—namely, confidence and prediction intervals, as will be discussed next.

To measure the quality of the regression model, common practice is the use the *Coefficient of Determination* (R^2) which can be computed with the *Total Sum of Squares* (TSS), *Explained Sum of Squares* (E_xSS), and *Residual Sum of Squares* (RSS), as defined in below and depicted in Figure 12.

$$TSS \equiv \sum_{i=1}^n (x_i - \bar{x})^2 \quad \text{Equation 32}$$

$$E_xSS \equiv \sum_{i=1}^n (\hat{x} - \bar{x})^2 \quad \text{Equation 33}$$

$$RSS \equiv \sum_{i=1}^n (x_i - \hat{x})^2 \quad \text{Equation 34}$$

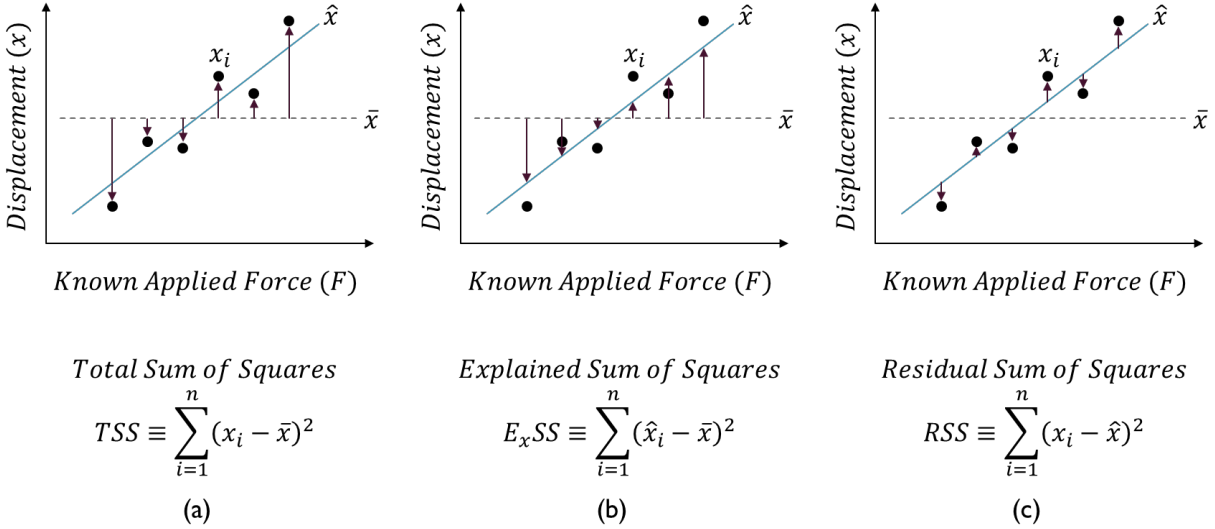


Figure 12. Depiction of (a) Total Sum of Squares, (b) Explained Sum of Squares, (c) and Residual Sum of Squares.

The TSS is a measure of the total variability in the response variable. The E_xSS is the amount of variability *explained* by the regression model. The RSS is the residual variability left unexplained by the regression model. A fundamental equation in regression analysis says that the E_xSS plus the RSS is always the TSS.

$$TSS = E_xSS + RSS \quad \text{Equation 35}$$

The *Coefficient of Determination* R^2 is defined as the ratio of the *Explained Sum of Squares* to the *Total Sum of Squares* and is bounded by 0-1.

$$R^2 \equiv \frac{E_xSS}{TSS} = 1 - \frac{RSS}{TSS} \quad \text{Equation 36}$$

Furthermore, the response is defined as the mean (expected) value of the response, given a very large number of predictors of the same value. Thus, when we predict an \hat{x}_i (the points on the blue line in Figure 12), what we are really predicting is the average deflection for many calibration-loads of the same weight F_i .

$$\hat{x}_i = E[x|F_i] \quad \text{Equation 37}$$

The distribution of x for each value of F_i is assumed to be equivalent (homoscedasticity), and this distribution has a variance called the common variance of the residuals. The *population common variance of the residuals* σ_{RSS}^2 is unknown but can be estimated with the *sampling common variance of the residuals* s_{RSS}^2 , the latter of which is the RSS normalized by its number of degrees of freedom—hence, s_{RSS}^2 is also

called the *Mean Residual Sum of Squares (MRSS)*.¹⁰ Recall that any sum of squares has associated with it a certain number of degrees of freedom. From Equation 17, recall that the number of degrees of freedom is the number of observations n minus the number of estimated parameters. For a simple linear model, there are two estimated parameters, $\hat{\beta}_0$ and $\hat{\beta}_1$, so the number of degrees of freedom is $n - 2$.

$$\sigma_{RSS}^2 \approx s_{RSS}^2 = MRSS = \frac{RSS}{dof} = \frac{\sum_{i=1}^n (x_i - \hat{x})^2}{n - 2} \quad \text{Equation 38}$$

The square root of the sample common variance of the residuals s_{RSS} is called *Residual Standard Error*,¹¹ but it is best thought of as the *standard deviation of the residuals*. For example, compare Equation 39 below with Equation 16. The only difference is that the data x_i in Equation 39 is compared to the model \hat{x} and not the sample mean of the data \bar{x} , and that the sum of squares in Equation 39 has 2 degrees of freedom (two unknowns, $\hat{\beta}_0$ and $\hat{\beta}_1$), while the sum of squares in Equation 16 has only 1.

$$s_{RSS} = \sqrt{\frac{\sum_{i=1}^n (x_i - \hat{x})^2}{n - 2}} \quad \text{Equation 39}$$

The concepts of homoscedasticity and the standard deviation of the residuals are depicted in Figure 13.

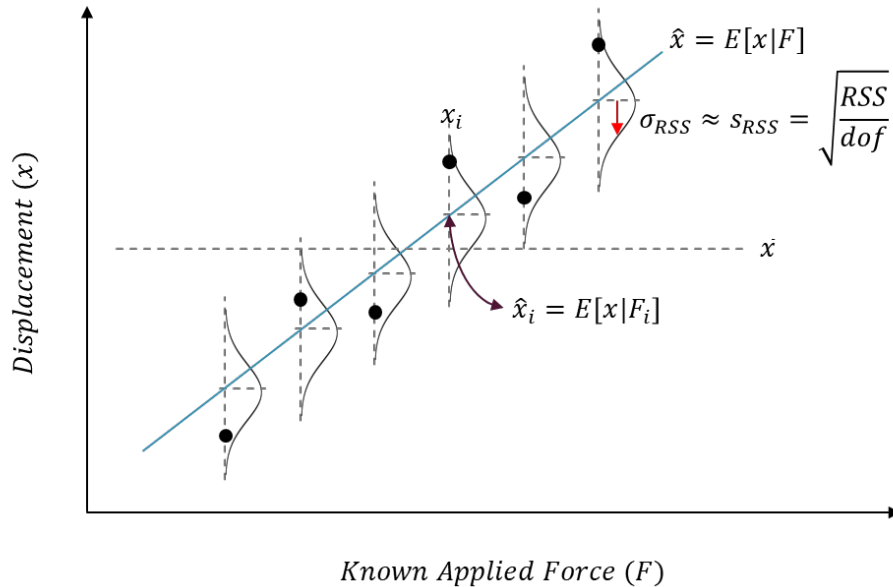


Figure 13. Depiction of homoscedasticity and the standard deviation of the residuals.

The standard deviation of the residuals in Equation 39 plays a key role in calculating the confidence intervals about both the estimated coefficients and the predicted response.

¹⁰ The sampling common variance goes by many names; it is also called the *variance of the residuals*, the *Mean Standard Error* or the *Mean Squared Error (both MSE)*.

¹¹ Other names include the *Root Mean Squared Error (RMSE)* or the *Standard Error of the Estimates (SEE)*.

Upon repeated calibrations, we will obtain slightly different estimates of coefficients $\hat{\beta}_0$ and $\hat{\beta}_1$. Much like the distribution of sampling means, after many calibrations, we would obtain a sampling distribution of each coefficient. Statistical theory tells us that these distributions are normally distributed, and their standard deviations, called the *standard error of the estimated coefficients*, can be computed from (by taking the square root of) the following equations.

$$SE^2(\hat{\beta}_1) = \sigma_{RSS}^2 \left[\frac{1}{\sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 40}$$

$$SE^2(\hat{\beta}_0) = \sigma_{RSS}^2 \left[\frac{1}{n} + \frac{\bar{F}^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 41}$$

In the above equations, the *population common variance of the residuals* σ_{RSS}^2 is unknown, and is thus estimated with the *sample common variance of the residuals* s_{RSS}^2 (A.K.A. the MRSS), in which case, Equation 40 and Equation 41 would be *estimates* of the standard errors of the estimated coefficients. Note that in the EP community, it is common practice to perform many repeated calibrations. If this is the case, and all the estimated coefficients are recorded, then the standard error of the estimated coefficients need not be estimated—it can be computed directly by taking the standard deviation of the mean of the estimated coefficients using Equation 18.

Once we have the regression model, it is straightforward to apply Equation 24 for prediction purposes. Given a value of the predictor F_0 , we can predict the value of the response \hat{x}_{F_0} . There are two sorts of uncertainty associated with this prediction; their sources follow.

1. The coefficient estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ are estimates for β_0 and β_1 —that is, the *best-fit line* \hat{x} is only an estimate for the *true population line* x_t . We use *confidence intervals* (CI) to determine how much our *best-fit line* \hat{x} will vary from *true (but unknown) population line* x_t .
2. Even if we knew the true population line (i.e., the true values of the coefficients β_0 and β_1), the response value cannot be predicted perfectly because of random uncertainty. We use *prediction intervals* (PI) to determine how much our predicted response \hat{x}_i will vary from actual data x_i .

Confidence intervals are computed using the *standard error of the regression model* $SE(\hat{x})$, which can be computed from (by taking the square root of) the following expression.

$$SE^2(\hat{x}) = \sigma_{RSS}^2 \left[\frac{1}{n} + \frac{(F_0 - \bar{F})^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 42}$$

⇒

$$CI = \hat{x} \pm t_{v,\alpha} SE(\hat{x})$$

As discussed previously, t is a value determined from a t-distribution table using the degrees of freedom v and confidence level $1 - \alpha$. Suppose we compute a 95% confidence interval about the regression model. The 95% CI should be interpreted as follows, “95% of the confidence intervals constructed in this way will contain the true but unknown population line.”

Prediction intervals are computed using the *standard error of the predicted response* $SE(\hat{x}_i)$, which can be computed from (by taking the square root of) the following expression.

$$SE^2(\hat{x}_i) = \sigma_{RSS}^2 \left[1 + \frac{1}{n} + \frac{(F_0 - \bar{F})^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 43}$$

\Rightarrow

$$PI = \hat{x}_i \pm t_{v,\alpha} SE(\hat{x}_i)$$

Suppose we compute a 95% prediction interval. The 95% PI should be interpreted as follows, “95% of the prediction intervals constructed in this way will contain the true but unknown value of the response variable.”

In both Equation 42 and Equation 43, σ_{RSS}^2 is estimated with s_{RSS}^2 . A visual comparison of the confidence and prediction intervals is presented in Figure 14.

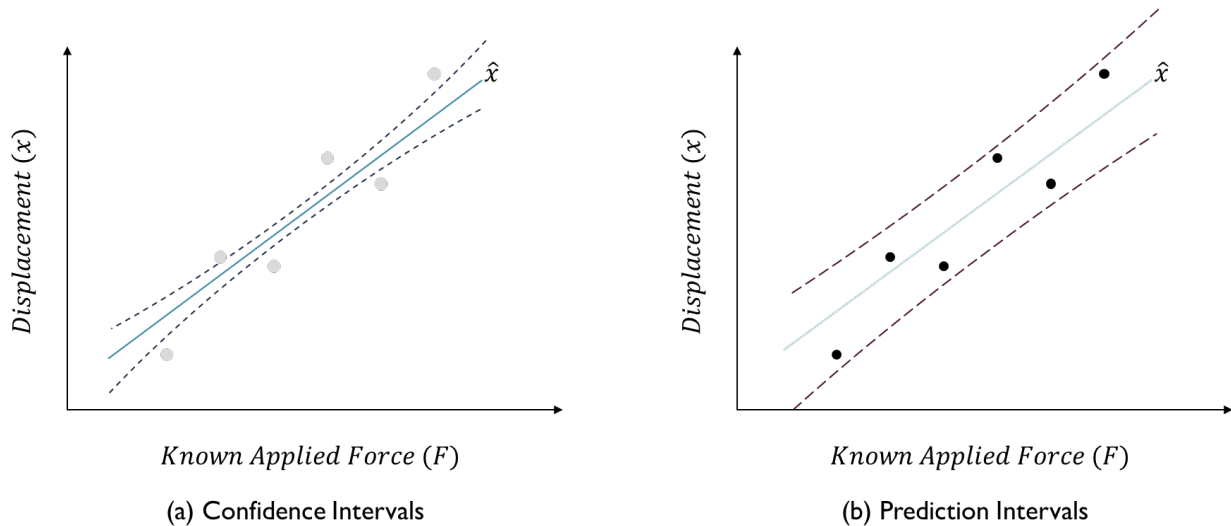


Figure 14. Comparison between (a) confidence intervals of the regression model and (b) prediction intervals of the predicted response.

Consider the left panel of Figure 14. The curved dashed lines represent the confidence intervals. They are bounds that apply to the regression line itself which is just an estimate of the true but unknown population line. What the confidence interval says is that if we took 100 samples and made 100 calibration curves and computed 100 confidence intervals in this way, roughly 95 of them would contain the true population line. Now consider the right panel of Figure 14. The dashed lines here represent the prediction intervals. What the prediction interval says is that if we took a single value of the applied force F_i and predicted the displacement and computed the prediction interval, there is a 95% probability that this interval contains the true but unknown displacement x_i for that calibration load F_i . In other words, the prediction interval is used to quantify the uncertainty surrounding the predicted displacement for single value of the applied load.

Furthermore, there are three important features which can be gleaned from Equation 42 and Equation 43 as depicted in Figure 14. First, there is a subtle flaring in these intervals with the smaller intervals being near the mean of the data. This is due to the $(F_0 - \bar{F})^2$ terms in both Equation 42 and Equation 43. In other words, the models are more accurate near the mean of the data. This should be taken into consideration when selecting the calibration weight sets. Second, the prediction intervals are always larger than confidence intervals. This is due to the addition of the “1” in Equation 43. This makes sense since the prediction intervals is the uncertainty surrounding a prediction from a single value of the predictor, whereas the confidence interval is the uncertainty surrounding the average (expected) deflection of many predictors of the same value (see Equation 37). Third, note that both the confidence and prediction intervals are sets of *vertical error bars* on the predicted response. This is a key point that will be revisited shortly.

We can use the prediction intervals to quantify the random uncertainty in our predictions. Consider the recommended *x on F* regression model given by Equation 24. We can use the model two ways. We *could* use the model as trained, but what we really need is to invert the model and perform *inverse prediction*, as discussed previously, and predict the applied force given a value of the displacement. The uncertainty in the former *as trained* method can be characterized by Equation 44 (which is the same as Equation 43, except now it is the standard error of the predicted displacement from F_0); whereas the uncertainty in the *inverted prediction* method can be characterized by Equation 45. In the latter case, to minimize random uncertainty, the value taken for the deflection is often the average of many measurements, say N , and is denoted \bar{x}_0 .

$$SE^2(\hat{x}_{F_0}) = \sigma_{RSS}^2 \left[1 + \frac{1}{n} + \frac{(F_0 - \bar{F})^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 44}$$

⇒

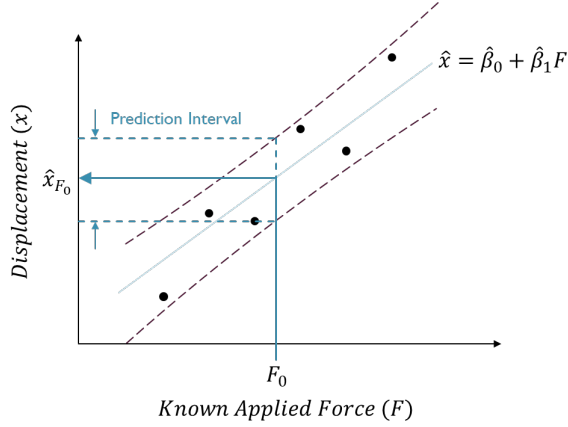
$$PI(\text{as trained}) = \hat{x}_i \pm t_{v,\alpha} SE(\hat{x}_{F_0})$$

$$SE^2(\hat{F}_{\bar{x}_0}) = \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right] \quad \text{Equation 45}$$

⇒

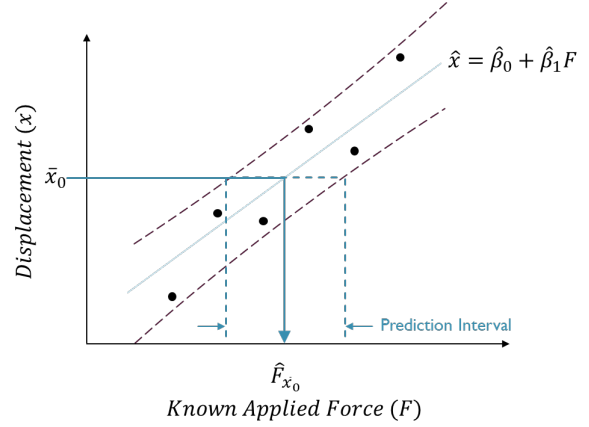
$$PI(\text{inverse prediction}) = \hat{F}_{\bar{x}_0} \pm t_{v,\alpha} SE(\hat{F}_{\bar{x}_0})$$

These two prediction methods and their associated prediction intervals are depicted in Figure 15.



$$SE^2(F_0 \rightarrow \hat{x}_{F_0}) = \sigma_{RSS}^2 \left[1 + \frac{1}{n} + \frac{(F_0 - \bar{F})^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \right]$$

(a) as trained method



$$SE^2(\bar{x}_0 \rightarrow \hat{F}_{\bar{x}_0}) = \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right]$$

(b) inverse prediction method

Figure 15. Prediction intervals (a) using the regression model as trained vs. (b) using the inverted prediction method.

One important item to consider is the fact that whichever method we chose, the variable that is used as the predictor is assumed known with absolute certainty. For example, in Figure 15a, when the regression model is used *as trained*, the predictor F_0 is assumed known with no uncertainty and the displacement \hat{x}_{F_0} incurs the uncertainty. In contrast, in Figure 15b where the regression model is inverted, the predictor, in this case the displacement \bar{x}_0 , is assumed known with no uncertainty and the estimated force $\hat{F}_{\bar{x}_0}$ incurs the uncertainty. In either method, there is some uncertainty that is ignored. This requirement, although never perfectly met, can be adequately satisfied by ensuring sufficiently small random uncertainty in the predictor, whatever it may be, via sufficient sampling as discussed in Section 2.2.

The inverse prediction interval given by Equation 45 (and depicted in Figure 15b) is at the heart of conducting an uncertainty analysis with a thrust stand calibration curve via the *inverse prediction* method. There are several important features that may be gleaned from this expression.

1. The standard error of the prediction interval decreases as the estimated slope coefficient $\hat{\beta}_1$ increases. We will see shortly that the inverse of the estimated coefficient ($1/\hat{\beta}_1$) is in fact the linear sensitivity of the thrust stand S_x . The point here is that the prediction interval decreases as the linear sensitivity increases, so this is an added advantage of designing higher sensitivity thrust stands.
2. The standard error of the predicted response $SE^2(\hat{F}_{\bar{x}_0})$ includes the *standard error of the estimated slope coefficient* $SE^2(\hat{\beta}_1)$. To see this, simply distribute the (square of the) *standard deviation of the residuals* σ_{RSS}^2 in Equation 45 and compare the result to Equation 40.

$$\begin{aligned} SE^2(\hat{F}_{\bar{x}_0}) &= \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right] \\ &= \frac{1}{\hat{\beta}_1^2} \left[\frac{\sigma_{RSS}^2}{N} + \frac{\sigma_{RSS}^2}{n} + \frac{\sigma_{RSS}^2}{\sum_{i=1}^n (F_i - \bar{F})^2} \cdot \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2} \right] \end{aligned}$$

$$= \frac{1}{\hat{\beta}_1^2} \left[\frac{\sigma_{RSS}^2}{N} + \frac{\sigma_{RSS}^2}{n} + SE^2(\hat{\beta}_1) \cdot \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2} \right]$$

In other words, the prediction intervals constructed using Equation 45 includes a term that accounts for the random uncertainty in the slope coefficient. Recall that standard error of the estimated coefficient need not be estimated—it can be computed directly by taking the standard deviation of the mean of the estimated coefficients using Equation 18.

3. The prediction interval is directly proportional to *standard deviation of the residuals* σ_{RSS}^2 , which is computed by Equation 38. Recall that square root of the common variance is called the *residual standard error* s_{RSS} , given by Equation 39, and is a measure of the random variability in the response variable, which is assumed constant for all values of the predictor. Thus, the prediction interval also includes a term that accounts for the random uncertainty in the displacement response.
4. The prediction interval decreases as both the number of calibration points n increases and as the number of data points taken to compute the mean of the displacement (\bar{x}_0) N increases.
5. The prediction interval decreases as $\bar{x}_0 \rightarrow \bar{x}$. The measurements will be more accurate if the displacement is near the middle of the range experienced under the range of calibration loads.

In short, the prediction interval given by the *inverse prediction* method (Equation 45) is used when the regression model is trained *x on F*, then inverted to predict force given a value of the deflection. The standard error of the predicted response includes the random uncertainty in the estimated slope coefficient, the random uncertainty in the displacement variable, the number of calibration points, and the number of data point used to compute the mean displacement \bar{x}_0 . **We strongly recommend the inverse prediction method coupled with the inverse prediction intervals to be used to perform the random uncertainty analysis for a thrust measurement.**

With the aforementioned techniques for conducting uncertainty analysis using the tools of linear regression, we can now apply this framework to electric propulsion thrust stands. To keep the discussion grounded, we will focus on inverted pendulum thrust stands. First, we will review best practices for taking a thrust measurement, then we will examine the total random uncertainty in a thrust stand measurement. Finally, we will tackle estimating systematic uncertainties.

2.5. Thrust measurements

The recommended practice for taking a thrust measurement is as follows.

1. Start the thruster and allow it to reach thermal equilibrium and steady-state operation.
2. Obtain an adequate number of measurements of the thrust stand displacement to calculate a mean value $\bar{x}_{ON} = \frac{1}{N} \sum_{i=1}^N x_{ON,i}$ with a sufficiently small standard deviation of the mean.
3. Turn off thrust power and flow rate and obtain measurements of the thrust stand displacement to characterize the zero (no load) value $\bar{x}_{OFF} = \frac{1}{N} \sum_{i=1}^N x_{OFF,i}$.

The difference between the thruster ON and OFF states will be considered the thrust response for that operating point. This method has the advantage of correcting for (to some extent) long-term drifts in the thrust stand zero.

$$F_T = \bar{F}_{ON} - \bar{F}_{OFF} \quad \text{Equation 46}$$

Thus, our thrust measurement F_T is the difference between two regression predictions, \bar{F}_{ON} and \bar{F}_{OFF} , where \bar{F}_{ON} and \bar{F}_{OFF} , as per Equation 25, are given by the following.

$$\bar{F}_{ON} = \frac{\bar{x}_{ON}}{\hat{\beta}_1} - \frac{\hat{\beta}_0}{\hat{\beta}_1} \quad \text{Equation 47}$$

$$\bar{F}_{OFF} = \frac{\bar{x}_{OFF}}{\hat{\beta}_1} - \frac{\hat{\beta}_0}{\hat{\beta}_1} \quad \text{Equation 48}$$

Recall our discussion in Section 2.3 regarding the combination and propagation of uncertainties. Here we have a function F_T of two variables, \bar{F}_{ON} and \bar{F}_{OFF} . Our objective in the next two subsections is to determine the random and systematic uncertainties in both \bar{F}_{ON} and \bar{F}_{OFF} . Give those figures, we will then combine the random and systematic uncertainties in \bar{F}_{ON} to obtain a *total* uncertainty in \bar{F}_{ON} , then we combine the random and systematic uncertainties in \bar{F}_{OFF} to obtain a *total* uncertainty in \bar{F}_{OFF} .

2.6. Random uncertainty

As discussed, the most direct method for conducting random uncertainty analysis is to use the inverse prediction intervals for both \bar{F}_{ON} and \bar{F}_{OFF} , which can be computed from Equation 45. In other words, when reporting \bar{F}_{ON} and \bar{F}_{OFF} , we can compute the 95% prediction intervals using Equation 45 and write the following.

$$F_{ON} = \bar{F}_{ON} \pm t_{v,\alpha} SE(\bar{F}_{ON}) \quad \text{Equation 49}$$

$$F_{OFF} = \bar{F}_{OFF} \pm t_{v,\alpha} SE(\bar{F}_{OFF}) \quad \text{Equation 50}$$

The terms that follow after the \pm in Equation 49 and Equation 50 are the random uncertainties in the predicted responses.

$$U_{rand,\bar{F}_{ON}} = t_{v,\alpha} SE(\bar{F}_{ON}) \quad \text{Equation 51}$$

$$U_{rand,\bar{F}_{OFF}} = t_{v,\alpha} SE(\bar{F}_{OFF}) \quad \text{Equation 52}$$

2.7. Systematic uncertainty

The process for estimating the systematic uncertainty in a thrust stand measurement is not as straightforward as random uncertainties and is more an exercise in engineering judgement. Recall from Section 2.1. that systematic errors which cannot be quantified are called *systematic uncertainties*, and they are usually modeled like random uncertainties, i.e., with normal distributions. Sources of systematic errors include thermal drift, mechanical friction, external electromagnetic forces, non-uniformities in vacuum facilities, thermal expansion, and drift in plumbing/cabling. However, and this is crucial, many systematic errors are

nulled via the calibration process. **Because the calibration process is a mapping of the thrust stand response to a known applied force, whatever systematic imperfections that alter the response from what it would have been (were those imperfections not present) are of little significance so long as those systematic imperfections are present during both thrust stand calibration (when the “mapping” is statistically learned) and thruster operation (when the “mapping” is applied).** This feature of the calibration process reduces the sources of systematic errors to only those that differ between thruster operation and calibration, with thermal effects being the dominant culprit. Simply put, the thermal environment of the thrust stand may be different between calibration and thruster operation due to the heat generated by the thruster. For this reason, modern thrust stands employ thermal control via water cooling systems. These systems are used to maintain consistent temperatures between calibration and thruster operation; however, temperature variation can still exist, and it is worth examining how thermal effects can alter a thrust measurement.

2.7.1. Thermal effects

Thermal effects can creep into the thrust stand in many ways, but generally, they cause either *gain shifts*, that is, changes to the stiffness of the thrust stand—or *zero shifts*, that is, changes to the thrust stand “zero” state. For example, springs generally become less stiff with increasing temperature, so for the same applied force, the spring would deflect more if it were hotter (gain shift). Similarly, a spring can undergo thermal expansion (or contraction) under temperature variation, which could change the neutral length of the spring and alter the zero-position of the thrust stand (zero shift). This could happen, for example, with the “waterfall” power lines. There is no standard method for estimating the gain and zero shifts of the thrust stand due to thermal effects. We provide three options.

1. Experimental examination: one could simply perform a series of calibrations under varying degrees of thermal load and examine the relationship (if any) between the thermal load and calibration constants. This could involve the creation of a mass and thermal simulator of the thruster of interest.
2. Thermo-mechanical simulation: using a high-fidelity CAD model of the thrust stand, software exists that can predict thermal expansion and contraction under various thermal maps. However, a challenge here is that this method requires an accurate thermal mapping of the thrust stand and all its components.
3. Analytical study: one could use simple beam theory and heat transfer analysis to estimate the variation in the calibration constants. Similar to the former, this method requires an accurate thermal mapping of the thrust stand and all its components.

Recall that any *systematic error* that is identified should be corrected at the source or in the reporting of the final measurement. Otherwise, one should report that the systematic bias cannot be resolved beyond some threshold and treat it as a *systematic uncertainty*. Once the systematic uncertainty has been identified, common practice is to combine it with the random uncertainty via the *Summation in Quadrature* (SIQ) method. For example, a possible gain shift should be treated as a systematic uncertainty in \bar{F}_{ON} ($U_{sys, \bar{F}_{ON}}$) and should be combined with the corresponding random uncertainty in \bar{F}_{ON} ($U_{rand, \bar{F}_{ON}}$) via the SIQ method to obtain the total uncertainty in \bar{F}_{ON} ($U_{tot, \bar{F}_{ON}}$). Likewise, a possible zero shift should be treated as a systematic uncertainty in \bar{F}_{OFF} ($U_{sys, \bar{F}_{OFF}}$) and should be combined with the corresponding random uncertainty in \bar{F}_{OFF} ($U_{rand, \bar{F}_{OFF}}$) via the SIQ method to obtain the total uncertainty in \bar{F}_{OFF} ($U_{tot, \bar{F}_{OFF}}$). This procedure is shown below.

$$U_{tot, \bar{F}_{ON}} = \sqrt{(U_{sys, \bar{F}_{ON}})^2 + (U_{rand, \bar{F}_{ON}})^2} \quad \text{Equation 53}$$

$$U_{tot, \bar{F}_{OFF}} = \sqrt{(U_{sys, \bar{F}_{OFF}})^2 + (U_{rand, \bar{F}_{OFF}})^2} \quad \text{Equation 54}$$

2.7.2. Thrust stand inclination

A relatively large source of systematic uncertainty that is not nulled via the calibration due to the uncertainty in the inclination of the thrust stand. For this reason, modern thrust stands employ active inclination control (as depicted in Figure 3). Were the entire thrust stand tilted by some inclination i about the inclination pivot, then a component of the weight of the thruster would exist in the line of motion causing a deflection in the thrust stand not generated by thrust, as shown below. We refer to this erroneous force as the “lean force.”




Figure 16. Notional inclination error due to a component of the weight force in the direction of thrust stand deflection.

Notice that the lean force can be positive (artificially increases the measured thrust) or negative (artificially decreases the measured thrust) depending on the direction of the inclination. If we knew i , then we could correct our measurement as follows.

$$\begin{aligned} \hat{F} &= F + F_{lean} \\ &\Rightarrow \\ F &= \hat{F} - F_{lean} = \hat{F} - mg \sin i \cos i \end{aligned}$$

But since we do not know i , the uncertainty in our measurement can be estimated as the difference between the true thrust and the estimated thrust.

$$U_{F,i} = F - \hat{F} = \hat{F} - mg \sin i \cos i - \hat{F} = -mg \sin i \cos i$$

Here we have an expression for the uncertainty in the force as a function of the inclination. Note the negative sign. If i were positive (or leaning the right as depicted in Figure 16), then the measured thrust \hat{F} would be larger than the true thrust generated by the thruster F , since the measured thrust would include the lean force, and the error would come out negative. If i were negative (leaning to the left), then the measured thrust would be less than the true thrust, and the error would come out positive, since $\sin -i = -\sin i$. Also note that the uncertainty in F due to i is symmetric (normally distributed with a mean of zero), meaning that if quantified in this way it can be combined with other systematic uncertainties in F via the SIQ method.

2.7.3. Misalignment of thrust vector

Another source of systematic uncertainty that is not nulled via the calibration process is the potential misalignment of the thrust vector due to improper mounting. (A similar analysis can be conducted for a misalignment of the calibration pulley mechanism, but we will only discuss the thrust vector here.) If the thrust vector is misaligned, then the thrust stand will only measure the projection of the thrust vector onto the line of motion of the thrust stand. Suppose we cannot resolve alignment of the thruster beyond some imaginary cone with a half-angle of α degrees. That is, the thruster vector is within $\pm\alpha$ of the line of motion, as shown below.




Figure 17. Notional misalignment of thrust vector.

How shall we treat this? First, observe that no matter the orientation of the misaligned thrust vector (e.g., whether above, below, or to the side of the line of motion), the projection onto the line of motion will always be less than the true thrust. The force causing the horizontal deflection is reduced to $\vec{F} \cos(\alpha)$. Since our displacement sensor is only sensitive to displacements in the x-direction, our estimated thrust \hat{F} measures the projection of the true thrust on the x-axis. If we knew α , then we could correct our measurement as follows.

$$\begin{aligned}\hat{F} &= F \cos(\alpha) \\ \Rightarrow \\ F &= \frac{\hat{F}}{\cos(\alpha)}\end{aligned}$$

But since we do not know α , the uncertainty in our measurement can be estimated as the difference between the true thrust and the estimated thrust.

$$U_{F,\alpha} \cong F - \hat{F} = \frac{\hat{F}}{\cos(\alpha)} - \hat{F} = \hat{F}(\sec(\alpha) - 1)$$

Here we have an expression for the uncertainty in the force as a function of the misalignment angle α . Notice that $\sec(\alpha) = \sec(-\alpha) \geq 0$, meaning that the uncertainty is always positive. This is a rare scenario in which the uncertainty is asymmetric. Recall from Section 2.3, that we can only use the SIQ (Summation in Quadrature) method for symmetric uncertainties. If we were to account for this uncertainty, we recommend adding to the total uncertainty directly, which would result in an asymmetric uncertainty (much like that depicted in Figure 5).

$$F = \hat{F} \pm U_{tot,F} + U_{F,\alpha}$$

2.7.4. Other sources

Note that there is no need to include the systematic uncertainty of the displacement sensor in the uncertainty analysis. If the sensor were biased, so long as the bias is present during both thruster ON and thruster OFF conditions, then this systematic error will cancel when computing the difference between \bar{F}_{ON} and \bar{F}_{OFF} —this is another advantage of taking the difference between the displacement sensor signals as the thrust stand response.

2.8. Total uncertainty

Finally, once the total uncertainties in \bar{F}_{ON} and \bar{F}_{OFF} have been computed, we can compute the total uncertainty in our thrust measurement. Recall from Equation 46 that the thrust measurement is the difference between \bar{F}_{ON} and \bar{F}_{OFF} : $F_T = \bar{F}_{ON} - \bar{F}_{OFF}$. Here we have a function F_T of two variables. From Figure 10, we know that the total uncertainty in F_T is computed by combining the partial uncertainties in F_T due to \bar{F}_{ON} and \bar{F}_{OFF} via the SIQ method. The complete strategy for conducting a thrust stand uncertainty analysis is shown below.

$$U_{tot,F_T} = \sqrt{(U_{F_T,\bar{F}_{ON}})^2 + (U_{F_T,\bar{F}_{OFF}})^2}$$

with

$$U_{F_T,\bar{F}_{ON}} = \frac{\partial F_T}{\partial \bar{F}_{ON}} U_{tot,\bar{F}_{ON}} = U_{tot,\bar{F}_{ON}}$$

$$U_{F_T,\bar{F}_{OFF}} = \frac{\partial F_T}{\partial \bar{F}_{OFF}} U_{tot,\bar{F}_{OFF}} = U_{tot,\bar{F}_{OFF}}$$

with

$$U_{tot,\bar{F}_{ON}} = \sqrt{(U_{sys,\bar{F}_{ON}})^2 + (U_{rand,\bar{F}_{ON}})^2}$$

$$U_{tot,\bar{F}_{OFF}} = \sqrt{(U_{sys,\bar{F}_{OFF}})^2 + (U_{rand,\bar{F}_{OFF}})^2}$$

with

$$U_{sys, \bar{F}_{ON}} = \text{engineering judgment}$$

$$U_{sys, \bar{F}_{OFF}} = \text{engineering judgment}$$

and

$$U_{rand, \bar{F}_{ON}} = t_{v, \alpha} SE(\bar{x}_{ON} \rightarrow \bar{F}_{ON})$$

$$U_{rand, \bar{F}_{OFF}} = t_{v, \alpha} SE(\bar{x}_{OFF} \rightarrow \bar{F}_{OFF})$$

where

$$SE(\bar{x}_{ON} \rightarrow \bar{F}_{ON}) = \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_{ON} - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right]$$

$$SE(\bar{x}_{OFF} \rightarrow \bar{F}_{OFF}) = \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_{OFF} - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right]$$

where

$n = \text{No. of calibration points}$

$N = \text{No. of measurements used to compute } \bar{x}_{ON} \text{ or } \bar{x}_{OFF}$

3. Conceptual Design

In this section, we will examine the different configurations of pendulum thrust stands (inverted, hanging, and torsional) as well as two common modes of operation (null-mode, and displacement-mode). Then, upon selection of the thrust stand configuration and mode of operation, we will provide a series of design recommendations.

3.1. Brief configuration trade study

Recall from Section 1.2 that the three forms of pendulum thrust stands are all governed by the same rotational spring-mass-damper equation of motion with the difference being the effect that gravity has on the effective angular spring constant, as shown below.

$$I\ddot{\theta} + c\dot{\theta} + k\theta = FL$$

with

$$k = \begin{cases} k_{\theta} & \text{for torsional pendulums} \\ k_{\theta} + mgL & \text{for hanging pendulums} \\ k_{\theta} - mgL & \text{for inverted pendulums} \end{cases}$$

Recall further that we derived an expression for the linear sensitivity of the thrust stand as follows.

$$S_x \equiv \frac{x_{ss}}{F} = \frac{L^2}{k}$$

Substituting the expressions for the effective angular spring constant in the linear sensitivity equation, we can express the linear sensitivity of the various pendulum mechanisms as follows.

$$S_x = \begin{cases} \frac{L^2}{k_\theta} & \text{for torsional pendulums} \\ \frac{L^2}{k_\theta + mgL} & \text{for hanging pendulums} \\ \frac{L^2}{k_\theta - mgL} & \text{for inverted pendulums} \end{cases} \quad \text{Equation 55}$$

Equation 55 is an important expression for the preliminary “sizing” of EP thrust stands. The true angular spring constant, k_θ , can be a challenge to ascertain analytically due to complex interfaces between the stationary and moving components of the thrust stand, such as the propellant lines, cabling, etc.; but these parasitic spring forces are usually much smaller than the angular spring stiffness caused by the primary springs. Thus, for preliminary sizing, we can assume that the parasitic spring forces are negligible and model k_θ analytically.

For example, the Owens inverted pendulum thrust stand currently uses a linear wave (or M-shaped) spring for the primary spring that pushes back with a spring force F_s against the deflection of the thrust stand θ . Under such angular deflection, the linear deflection is $x = L\sin(\theta) \cong L\theta$, in which case the torque from the linear spring is as follows.

$$\begin{aligned} \tau_s &= (\text{spring force})(\text{moment arm}) \\ &= F_s L \\ &= k_s x L \\ &= k_s L^2 \theta \end{aligned}$$

Therefore, the angular spring constant can be computed from the linear spring constant as follows.

$$k_\theta = k_s L^2 \quad \text{Equation 56}$$

Note that if multiple, say N_s , springs were connected “in parallel,” their spring constants would combine via a direct sum, and the total angular spring constant would be as follows.

$$k_\theta = N_s k_s L^2 \quad \text{Equation 57}$$

More modern thrust stands employ *flexure pivots*, which are purely rotational springs and offer several advantages over planer leaf springs (durability, higher buckling load, etc.), and since they operate “in parallel,” their spring constant also add. The total angular spring constant for N_f flexure pivots would be as follows.

$$k_{\theta} = N_f k_f \quad \text{Equation 58}$$

Note that one could combine linear springs and flexure pivots, in which case, the angular spring constant would simply be the sum of $N_s k_s L^2 + N_f k_f$.

With these expressions for the angular spring constant, we can perform preliminary sizing of the new thrust stand to determine which configuration would best suit our needs. The scope of this trade study will be limited to the three types of pendulum mechanisms and two modes of operation shown below.

		Mechanism		
		Inverted Pendulum	Hanging Pendulum	Torsional Pendulum
Mode of Operation	Displacement Mode	DM-IP	DM-HP	DM-TP
	Null Mode	NM-IP	NM-HP	NM-TP

Figure 18. Thrust stand preliminary design configuration trade space.

3.1.1. Preliminary design constraints

The Owens Chamber will be the workhorse for the next generation of high-power Hall thrusters, with thruster masses approaching 100 kg. Let us take the mass of the thrust stand (hardware and thruster) to be 100 kg ($m = 100 \text{ kg}$). Furthermore, let us strive for a linear sensitivity of one micron of deflection for every millinewton of thrust ($S_x = 1 \mu\text{m}/\text{mN} = 1 \times 10^{-3} \text{ m/N}$). Finally, let us strive for a relatively compact size, i.e., pendulum length less than half a meter ($L < 0.5 \text{ m}$) to fit comfortably inside the existing vacuum chamber.

3.1.2. Torsional pendulum

The torsional pendulum (TP) is appealing because the angular stiffness and dynamics of the thrust stand are independent of gravity, which makes for a more straightforward design process. TPs can thus be designed for very high sensitivities, which, as we have discussed, is advantageous for multiple reasons. TPs are sometimes employed as swinging gates or counterbalanced mechanisms. With the swinging gate mechanism, the 100 kg load would put an uncomfortably large moment on the hinges of the swinging gate. We could combat this large moment via a counterbalance such that the net moment is zero; however, in doing so, we are increasing the standing weight of the thrust stand, which is already considerable given the

100 kg class thruster. Theoretically, the standing weight could be reduced with longer moment arms, but this would increase the pendulum length constraint for our new design. For these reasons, torsional pendulums are generally not suitable for very large thrusters, such as the ones considered for the Owens Chamber.

3.1.3. Hanging pendulum

The hanging pendulum (HP) is appealing because the system is inherently stable. HPs are usually either single or parallel linkage systems. Parallel linkage systems are advantageous as they maintain a horizontal thruster vector (perpendicular to the force of gravity). With a single linkage, the thrust vector would tilt with the deflection of the thrust stand. Recall from Equation 55 that the sensitivity of HPs is inversely proportional to the sum of the angular spring constant and gravitational torque constant, $S_x = L^2 / (k_\theta + mgL)$. If we consider the best-case scenario that $k_\theta = 0$ (free-swinging HP), and we want a thrust stand of mass $m = 100 \text{ kg}$ to have a linear sensitivity of $1 \mu\text{m}/\text{mN}$, then the length of the pendulum is roughly estimated as follows.

$$S_x = \frac{L^2}{mgL} = \frac{L}{mg}$$

$$\Rightarrow$$

$$L = mgS_x = (100 \text{ kg}) \left(10 \frac{\text{m}}{\text{s}^2}\right) \left(1 \times 10^{-3} \frac{\text{m}}{\text{N}}\right) \cong 1 \text{ m}$$

This length is larger than our initial design constraint of half a meter, and unfortunately, this length only increases as we add additional springs (linear or rotational). In the absence of springs, this free-swinging hanging pendulum would also suffer from relatively large settling times, which is not desirable for most electric propulsion testing. In short, although the HP is inherently stable, a reasonably sized pendulum would suffer from lower sensitivities, which is not suitable for our purposes.

3.1.4. Inverted pendulum

In terms of sensitivity, inverted pendulums (IP) are the most attractive option. Recall from Equation 55 that the linear sensitivity of IPs is inversely proportional to the difference $k_\theta - mgL$. Notice that as $mgL \rightarrow k_\theta$, $S_x \rightarrow \infty$, which means that we have several design variables available to tune the sensitivity of the thrust stand. For example, suppose we use a combination of linear springs and flexure pivots such that $k_\theta = N_s k_s L^2 + N_f k_f$, then the linear sensitivity would be as follows.

$$S_x = \frac{L^2}{N_s k_s L^2 + N_f k_f - mgL}$$

If S_x , m , and L were fixed, then we could still carefully select the number and stiffnesses of the linear springs as well as the number and stiffnesses of the flexure pivots to maintain our desired linear sensitivity. An inverted pendulum allows for the maximum control of the thrust stand sensitivity, as there are several design variables that we can control to tune the sensitivity of the thrust stand. Thus, the inverted pendulum is the best solution for our purposes as it can be designed both compact and with high sensitivity.

Note that no one pendulum design is inherently more or less accurate than the other. The proverbial devil is in the details. Recall that the total random uncertainty in a thrust measurement is governed by the standard error of the predicted response given by Equation 45 and repeated below for convenience.

$$SE^2(\hat{F}_{\bar{x}_0}) = \frac{\sigma_{RSS}^2}{\hat{\beta}_1^2} \left[\frac{1}{N} + \frac{1}{n} + \frac{(\bar{x}_0 - \bar{x})^2}{\hat{\beta}_1^2 \sum_{i=1}^n (F_i - \bar{F})^2} \right]$$

Further recall that the estimated slope coefficient (for the recommended *x on F* calibration) is functionally the linear sensitivity ($\hat{\beta}_1 = S_x$) and that the total random uncertainty of a thrust measurement is inversely proportional to the sensitivity of the thrust stand. Provided that all thrust stand mechanisms can achieve the same sensitivity, there is no reason to expect that any one pendulum type would exhibit a greater random uncertainty than any other, all else equal. Systematic uncertainties on the other hand, which, if present, typically manifest in the form of thermal effects, have a more pronounced effect on thrust stand accuracy the greater the characteristic length of the pendulum. Long, cantilevered structures, for example, are more susceptible to non-negligible thermal expansion and contraction with temperature variations. While the thrust stand may be thermally controlled, temperature variation can still exist, which may be different between thruster calibration and operation. Desirable TPs and HPs would require longer pendulum mechanisms, given the thruster mass under consideration. For these reasons, the inverted pendulum is our best option.

3.1.5. Displacement-mode vs null-mode

Finally, we must decide whether to implement a displacement mode or a null mode thrust stand. As a reminder, displacement mode thrust stands infer the thrust from the displacement of the thrust stand under some applied force. The calibration process, *x on F*, is the statistical learning of this relationship via a best-fit line. However, recall from Section 1.2 (see Equation 13) that the linear relationship between *x* and *F* was derived under the assumption of small angles. State-of-the-art inverted pendulum thrust stands maintain very small deflections (typically less than 1 degree), so we can assume small angles without significant error, although some systematic error is incurred under this assumption ($\leq 0.13\%$ for $\theta \leq 5^\circ$).

The small angle approximation can be eliminated via implementing a null mode thrust stand. As a reminder, null mode pendulums use a *restoring force*, or *null force*, to counter the deflection of the thrust stand, increasing the null force until the deflection is nulled, i.e., equivalent to the zero-position value when the thruster was OFF. Once the thrust stand is back to the zero-position (potential hysteresis notwithstanding), static equilibrium dictates that the null force is equivalent to the applied force. An additional advantage of null mode thrust stands is that (at steady state) the thruster remains at the same position regardless of thrust, which could be useful for near-field plasma diagnostics. For these reasons, a null-mode inverted pendulum is the recommended configuration.

Note that the calibration process for a null-mode thrust stand regresses the *null force* on the *applied force* (F_N on F), i.e., with the null force F_N on the “y-axis” as the response variable, and the applied force F on the “x-axis” as the covariate. The analytics for conducting random uncertainty analysis of null mode pendulums is exactly the same as for displacement mode presented in Section 2, with *displacement x* replaced by *null force* F_N .

3.2. Preliminary design recommendations

There are two central themes that the new design should focus on: thermal control and inclination control—as these are likely the dominant systematic uncertainties. What follows is a short list of design recommendations, followed by a brief explanation, that may be implemented in the new thrust stand design.

1. A wider base. As mentioned in Section 1.3, the Owens parallel linkage thrust stand is uncomfortably narrow for large thrusters. With the variability of thruster and cabling harness masses, the center of mass of the thrust stand will shift for different thrusters. Were the center of mass to shift outside the parallel linkages, the thrust stand would simply topple over. A wider base would afford lateral freedom in the location of the center of mass.
2. Use of flexure pivots wherever frictionless rotation is desired. As mentioned in Section 1.3, the Owens thrust stand currently employs thin metal strips to limit the motion of the thrust stand to one direction. These thin metal strips are a mechanical weak point for thrust stands as their high aspect ratio leads to a relatively low buckling point. Flexure pivots are a relatively modern technology that improve on the planar flexures by incorporating multiple thin metal strips inside a rotating assembly. Flexure pivots thus provide the same frictionless rotation while supporting large compressive loads and are thus recommended as an alternative to planar metal strips.
3. Multiple linear extension springs with fine-tune adjustability. Recall that the Owens currently uses a linear compressive wave spring as the primary spring. A single compressive spring offers little freedom to fine tune the stiffness of the system for different thruster masses as well as the zero position of the thrust stand. Multiple linear, and opposing, extension springs would afford this adjustability. Furthermore, if the extension springs are made with tension adjustability, the thrust stand could be “zeroed” or vertically balanced more accurately.
4. Optical sensors for displacement sensing. Optical sensors generally have higher accuracies and would thus reduce the number of data points required in the computation of \bar{x}_{ON} and \bar{x}_{OFF} to achieve a sufficiently low random uncertainty.
5. Flexure pivot on inclination pivot. Flexure pivots come in various sizes and the larger ones can withstand the shear force caused by 100-kg thrust stands. A pair of these can serve as a replacement pivoting mechanism to the weakened threaded rod design currently in place.
6. Inclination stepper motor and encoder. The inclination of the thruster can cause a significant systematic uncertainty in the thrust measurement due to the $mg \sin(\theta)$ component of gravity as discussed. With such heavy thrust stands, care should be taken to ensure that the motorized assembly used for inclination control has enough control authority to maintain the upright position without hysteresis. One solution is the use of a lead-screw design driven by a stepper motor and encoder system.
7. Inclination pivot downstream of center of mass and inclination motor upstream (or vice versa). Related to the above, the inclination pivot point and inclination control system should be on either side of the thrust stand center of mass. Currently, the pivot point is more-or-less directly beneath the center of mass, which is inherently unstable, but reduces the load on the inclination control. Shifting the inclination pivot point and inclination control system to either side of the center of mass would demand more control authority from the inclination control motors, but would be stable. Also, if the thrust stand were to tilt, we could design and plan for the direction it would tilt. If the inclination control can be accommodated this extra load, the stable setup is preferred.
8. Displacement lever. A notable idea is to amplify the displacement of the thrust stand using a mechanical lever. By amplifying the change in displacement, the minute displacements caused by thermal variation would be a lesser portion of the overall displacement, thus improving the

resolution. Put another way, it would bring the displacement caused by the thruster “out of the noise” of the displacements caused by thermal variation.

A notional depiction of a modernized thrust stand that incorporates these design recommendations is shown below.




Figure 19. Notional schematic of a modernized thrust stand. The calibration pulley mechanism is not shown to avoid excessive clutter.

4. Conclusions

We have tried to disseminate the many insights on thrust stands for electric propulsion; the most notable are the mechanics for conducting uncertainty analysis of these devices, and selection of thrust stand configuration. In this work, we first provided an overview of thrust stands for electric propulsion testing and presented the governing equations as well as the linearized sizing equations for their conceptual design. Second, we reviewed best practices for conducting uncertainty analysis in general and provided a review of the statistical mechanisms behind them. When calibrating the thrust stand, we recommend using inverse prediction method and train x on F , then to quantify the random uncertainty, we recommend using the (inverted) prediction intervals. Finally, some back-of-the-envelope calculations were shown to determine which configuration of thrust stand would be most suitable for future JPL Hall thruster programs. We recommend a parallel-linkage inverted pendulum, and we made several design recommendations that should, at a minimum, be thought-provoking.