Trading Robustness Requirements
in Mars Entry Trajectory Design

Jarret M. Lafleur[*]
Georgia Institute of Technology, Atlanta, Georgia 30332

One of the most important metrics characterizing an atmospheric entry trajectory in preliminary design is the size of its predicted landing ellipse. Often, requirements for this ellipse are set early in design and significantly influence both the expected scientific return from a particular mission and the cost of development. Requirements typically specify a certain probability level (σ-level) for the prescribed ellipse, and frequently this latter requirement is taken at 3σ. However, searches for the justification of 3σ as a robustness requirement suggest it is an empirical rule of thumb borrowed from non-aerospace fields. This paper presents an investigation into the sensitivity of trajectory performance to varying robustness (σ-level) requirements. The treatment of robustness as a distinct objective is discussed, and an analysis framework is presented involving the manipulation of design variables to effect trades between performance and robustness objectives. The scenario for which this method is illustrated is the ballistic entry of an MSL-class Mars entry vehicle. Here, the design variable is entry flight path angle, and objectives are parachute deploy altitude performance and error ellipse robustness. Resulting plots show the sensitivities between these objectives and trends in the entry flight path angles required to design to these objectives. Relevance to the trajectory designer is discussed, as are potential steps for further development and use of this type of analysis.

Nomenclature

C_A	=	vehicle axial force coefficient	X	=	given random variable
CDF	=	Cumulative Distribution Function	y_i	=	sampled value of a parameter
k	=	standard deviation multiplier	y_{target}	=	target value of a parameter
MSD	=	Mean Squared Deviation	ΔDOR	=	Delta-Differential One-way Ranging
MSL	=	Mars Science Laboratory	γ_{entry}	=	entry flight path angle
n	=	number of samples	μ	=	distribution mean
TCM	=	Trajectory Correction Maneuver	σ	=	distribution standard deviation

I. Introduction

ARGUABLY, two of the most important metrics in the preliminary design of atmospheric entry trajectories are the lengths of the major and minor axes of the predicted landing ellipse. For robotic Mars missions, during early trajectory design these dimensions serve as surrogate measures of the probability of reaching a scientifically interesting – or at least a low-hazard – landing target. Additionally, these dimensions continue to be important throughout the entire design process since final landing site selection may not occur until just a few months prior to landing. Typically, requirements for this ellipse are set very early in design. For example, the requirement for precision landing on the 2011 Mars Science Laboratory (MSL) mission was introduced a decade before its currently planned launch date. As a result, the stringency of these requirements can have a large influence on both the scientific return from a particular mission (particularly if the landed vehicle is not mobile) and the cost of development (particularly if precision guidance and navigation techniques must be implemented).

In order to completely specify a landing ellipse requirement, not only must major and minor axis lengths be specified, but the probability level (i.e., σ-level) for the ellipse must also be specified. Frequently this latter requirement is taken at 3σ. For example, the MSL mission has a precision landing requirement to land within 10 km

---

[*]Ph.D. Candidate, Daniel Guggenheim School of Aerospace Engineering, Student Member AIAA.

of its target with a 99.87% (3σ) probability.[†,1] While a physical reason may exist for the choice of range or axis length (e.g., limits on the rover's lifetime and ground range capability or, in future applications, the size of particular surface features or distances to pre-emplaced assets), it is unclear whether studies have been performed to assess the sensitivity of mission performance metrics to more or less stringent probability levels (i.e., lower or higher levels of risk in reaching the desired target). That is, the question exists, "What performance benefits could be realized by designing the error ellipse at, for example, 2.9σ or 3.1σ instead of 3σ?"

The objectives of this paper are to introduce this notion of robustness as simply another variable in the design of entry trajectories and to provide a practical example illustrating how this objective could be traded against other performance metrics.

A. Why 3σ?

The popularity of 3σ as a trajectory robustness goal is difficult to trace, and the simplicity of this number suggests that it has developed as a rule of thumb rather than through rigorous analysis. Evidence in the literature supports the hypothesis that the 3σ goal is largely an empirical one.

One potential source for the popularity of the 3σ goal is the well-known "3σ Rule" in statistics. This rule, which was initially empirical, states that Eq. (1) holds "for the overwhelming majority of commonly encountered random variables X".[2] The Vysočanskiĭ-Petunin inequality[2] generalizes this for any unimodal distribution via Eq. (2) (see also Fig. 1).[‡] If $k = 3$ is inserted into Eq. (2), it can be seen that drawing 3σ limits about any unimodal distribution will at most exclude 4/81 (4.94%) of the potential values of x. The significance of 5% is attributed to the fact that this percentile is widely used in applied science fields such as biology and medicine.[2]
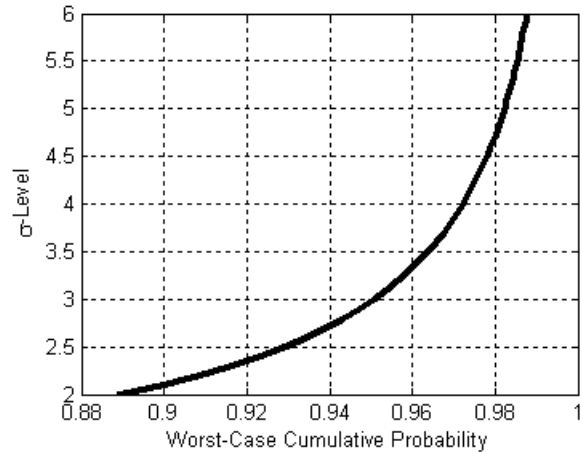

Figure 1. Worst-case cumulative probabilities (see Eq. (2)) for unimodal distributions.

$$P\big(|X - \mu| \geq 3\sigma\big) \leq 0.05 \tag{1}$$

$$P\big(|X - \mu| \geq k\sigma\big) \leq \frac{4}{9}\frac{1}{k^2} \tag{2}$$

Another potential – and perhaps more likely – origin of the 3σ goal lies in the field of quality engineering and the seminal work of Drs. Walter Shewhart and W. Edwards Deming. In 1931, Shewhart suggested the use of 3σ control limits as a way of economically selecting failed instances of a given process for examination.[3] For example, control charts are graphical methods of representing statistical data on the performance of a given process.[4,5,6] Users of control charts typically set upper and lower control limits that are three standard deviations from the mean of the data and use these limits to identify cases which may have resulted from special causes (as opposed to common causes) of variation. Identifying these special causes allows a decision-maker to take action to remove such outliers from future operations of the process. In this context, control limits that are too stringent (i.e. at low σ-levels) simultaneously allow large numbers of special causes to be identified and incur a high cost in terms of the resources required to investigate each outlier. Shewhart noted that, from experience, choosing 3σ control limit provides an economic balance between the value gained from investigating outliers and the cost of conducting the investigation. It is difficult to justify, however, that this balance would not be scenario-dependent (e.g. in scenarios where

---

[†] Because the required MSL landing ellipse is actually circular, the 3σ requirement is implemented in terms of distance from the target site, resulting in a 99.87% definition of 3σ (under the assumption that the data is normally distributed) instead of the 98.89% definition associated with 3σ for a bivariate normal distribution.

[‡] The Vysočanskiĭ-Petunin inequality, which is true only for a unimodal distribution (i.e., a distribution with a single mode, or peak), is a more restrictive version of the better-known Chebyshev inequality.

investigations are inexpensive, lower σ-levels might be easily justified). Even as Shewhart wrote, "Obviously, the basis for such limits must be, in the last analysis, empirical."[3]

In summary, the author has not been successful at finding any rigorous justification for using 3σ performance as a strict statistical indicator. The popular statistical "3σ Rule" only guarantees that the conversion from 3σ to percentage probability will not result in probabilities less than 95.06% for any unimodal distribution. It should be emphasized that this conversion from σ-level to a worst-case probability can be performed for any given σ-level (not just 3σ, e.g., see Fig. 1).[§] The field of quality engineering uses 3σ as an empirical economic guideline by which to decide whether to investigate outliers in an existing process, which has little relationship to the use of 3σ as a trajectory design requirement.

B. Robustness as an Objective

The basic notion that this paper highlights is the treatment of robustness as a distinct objective rather than a constraint or a component of a weighted "macro" objective. An example of the former would be a rigid requirement for a trajectory's 3σ error ellipse major axis to be smaller than a given length. The disadvantage to this approach is that it permits no trades to be made on whether 3σ is a reasonable robustness goal for the given ellipse length. An example of the latter is Taguchi's mean squared deviation (MSD).[7] As shown in Eq. (3), MSD combines the variance (i.e. robustness) of a distribution with its mean (i.e. average performance) into a single metric to be minimized for a given problem. Thus, the two objectives of robustness and performance are implicitly weighted within the MSD metric.

$$ MSD \equiv \frac{1}{n}\sum_{i}^{n}\left(y_i - y_{\text{target}}\right)^2 = \sigma^2 + \left(\bar{y} - y_{\text{target}}\right)^2 \tag{3} $$

A simple example of performance vs. robustness trades can be easily visualized in the cumulative distribution function (CDF) for a given Mars entry trajectory Monte Carlo simulation. As shown in Fig. 2, parachute deploy altitudes in this simulation range from 5.3 km to 6.8 km. Often, and particularly in scenarios where high-altitude landing sites or high-ballistic-coefficient, low-lift-to-drag vehicles are involved, it is desirable to deploy the vehicle's supersonic parachute as high as possible. Thus, in Fig. 2, the left tail of the distribution is much more preferable than the right tail, and for performance purposes when interacting with other disciplines, a designer would prefer to quote altitudes as high as possible without making large compromises in risk (i.e., the chance that his quoted altitude will be too high). Figure 2 can be helpful because the presence of steep slopes indicates that a small sacrifice in risk can allow one to quote much higher deploy altitudes. Shallow slopes indicate that a large gain can be made in certainty with only a small sacrifice in quoted altitude.


Figure 2. CDF for a sample Mars trajectory Parachute Deploy Altitude.

In fact, as can be more clearly seen in Fig. 3, steep slopes exist as the quoted percentile approaches 100%. Parachute deploy altitude at the 2σ level (97.72%) is 5.50 km, while it is 5.40 at the 3σ level (99.87%). Thus, in this region, every 0.1σ sacrificed results in approximately 10 m of altitude gain. Knowledge such as this can allow the trajectory designer to make more informed decisions on the level of risk and performance to which a quoted trajectory should be designed.

It should be emphasized that this example is very simple (executing it is trivial if Monte Carlo data is available) and is meant as an introduction to the concept of robustness as a distinct objective. The body of this paper focuses on a more complex multiobjective problem where parachute deploy altitude (performance) is traded against the deploy location error ellipse (robustness). Moreover, decisions on these robustness and performance trades govern the values that design variables will take (such as entry flight path angles and bank angle profiles).

[§] For example, if a guaranteed 97% probability is sought, Eq. (2) indicates that 3.85σ is an appropriate sigma-level
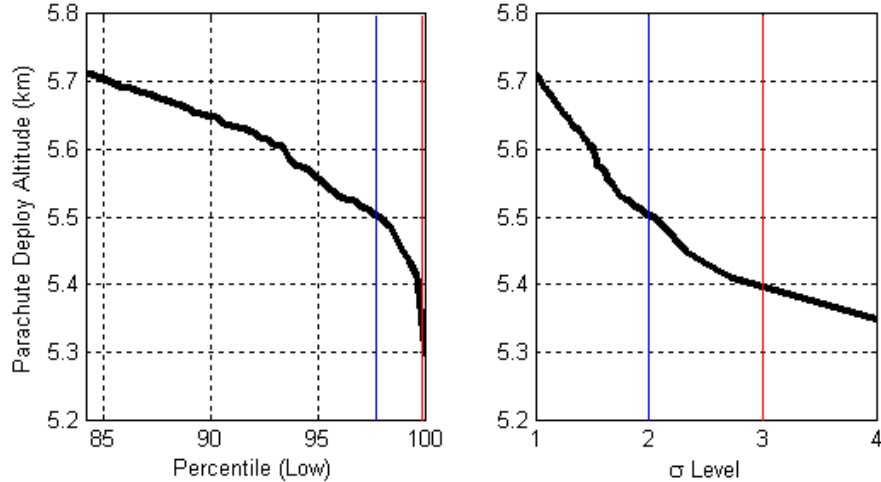
Figure 3. Magnified versions of Figure 1 in terms of Percentile (left) and σ Level (right).
Note that the blue line represents the 2σ level in both plots, and the red line represents 3σ.

II.   Data Analysis Method

A. Analysis Process

Figure 4 shows the distinctions between input and output variables in a generic Monte Carlo simulation. Typically, inputs include design variables that the user can choose in advance[**] as well as dispersion variables that account for various uncertainties (often these uncertainties are in external factors, such as a vehicle's operating environment). Outputs of interest are the distributions on variables that define the fitness of a given design, or how well the design fulfills its objectives.

If it is assumed that the input dispersions are properties of the operating environment which are not controllable by the designer, then the question for the designer that remains is of how to choose the design variables such that an "optimal" set of output distributions results. Even if only one objective output exists, this is a difficult problem to solve since optimality involves trading the objectives of the mean or average of the distribution (i.e., its representative performance) and its variance (i.e., a measure of its robustness).[††] This problem becomes even more difficult when multiple objective outputs exist.

Figure 5 shows the generic process in Fig. 4 applied to the example ballistic Mars entry problem used for the rest of this paper. In this case, selection of entry flight path angle has implications on the distributions for both parachute deployment location and parachute deployment altitude (both with respect to the Martian surface). From the trajectory designer's perspective, these are perhaps the two most important metrics characterizing the overall performance of a Mars entry trajectory.[‡‡] Since landing site selection typically does not occur until very late in design (in fact, often after launch), the mean of the parachute deployment location distribution (which is a bivariate error ellipse) is not as important as the variance. As illustrated in Fig. 6, this leaves three metrics to trade: parachute deployment location robustness, parachute deployment altitude performance, and parachute deployment altitude robustness. An important altitude-related parameter is altitude performance since change in the entry flight path angle design variable produces competing effects. That is, a steeper entry flight path angle for a ballistic entry lowers the Mach-triggered parachute deployment altitude but reduces landing ellipse size. Furthermore, it is desirable to have a high parachute deployment altitude in order to maximize the amount of accessible surface area

---

[**] In an ideal scenario, the user would be able to set these parameters in advance with zero uncertainty. However, it may be that only a central tendency (e.g. the median) of the design variable may be set in advance. If this is the case, the design variable would refer to the setting of the central tendency, over which the user has direct control.

[††] When more than one objective is involved, optimality has little meaning unless weights on the objectives are defined. In general, only the recognition of Pareto-optimal (non-dominated) designs can be accomplished.

[‡‡] Note, however, that these are not the only performance metrics of interest. For example, maximum heat rate and heat load are also important and in some cases can be trajectory design drivers.

and to provide timeline margin. Thus, although other trades can be made, consideration is currently only given in this paper to the key trade between parachute deployment location robustness and altitude performance.

Figure 7 shows the basic process to properly evaluate the desired location robustness vs. altitude performance trade. The process begins with previously-executed sets of Monte Carlo runs at varying entry flight path angle (γ_{entry}).[8] Each Monte Carlo set contains 1001 runs, and sets exist for 11 different entry flight path angles (resulting in just over 11,000 trajectory runs for this analysis). Recall that the flight path angle distinction is important because this is the design variable for this problem, as indicated in Fig. 5. The output of the process is the best altitude performance attainable as a function of design σ-levels at a given (e.g., 30 km) major axis requirement. The term "best" is used to describe the altitude performance because entry flight path angle is left as a free parameter via which altitude may be maximized (as long as the major axis requirement is not violated). This process is divided into three basic steps as indicated on Fig. 7:

1. **Raw Data Generation.** In the first step, the raw trajectory data for each run is processed such that the vehicle state (in particular, latitude, longitude, and altitude) at the parachute deployment condition is extracted.
2. **Data Extraction as a Function of Design Variable(s).** In the second step, robustness and performance data are computed as a function of entry flight path angle (the design variable for this problem). For a 1001-run data set at a given entry flight path angle, the error ellipse major axis length is first computed over a range of σ-levels. Next, the altitude performance is computed for that data set. Since an altitude distribution exists, altitude performance can in principle be quoted at any desired percentile level. Because here the central tendency of the altitude distribution is desired, the 50th percentile (median) altitude is quoted, although plots will be shown later in this paper of the 0.13th percentile low altitude as well.
3. **Data Extraction in the Objective Space.** In the third and final step, the data from the second step are used as a basis for interpolation for the final plots in the performance vs. robustness space. For each σ-level from Step 2 (i.e., each row of the table in Fig. 7), a value of the γ_{entry} design variable is selected to provide the best possible quoted altitude (i.e., the last row in Fig. 7) subject to an ellipse length constraint (e.g., 30 km). Thus, the last two columns of Fig. 7 are generated. The left column shows the best quoted altitude that can be achieved for a given σ-level, and the right column shows the value of the γ_{entry} design variable that the designer must select to allow achieve this. These data are the subject of the remainder of this paper.
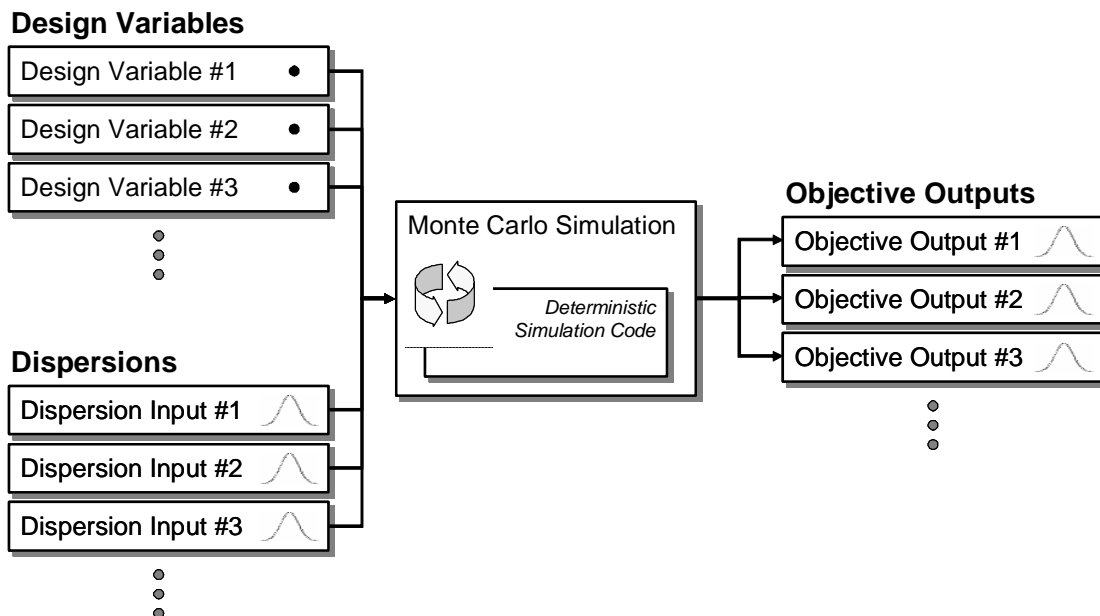

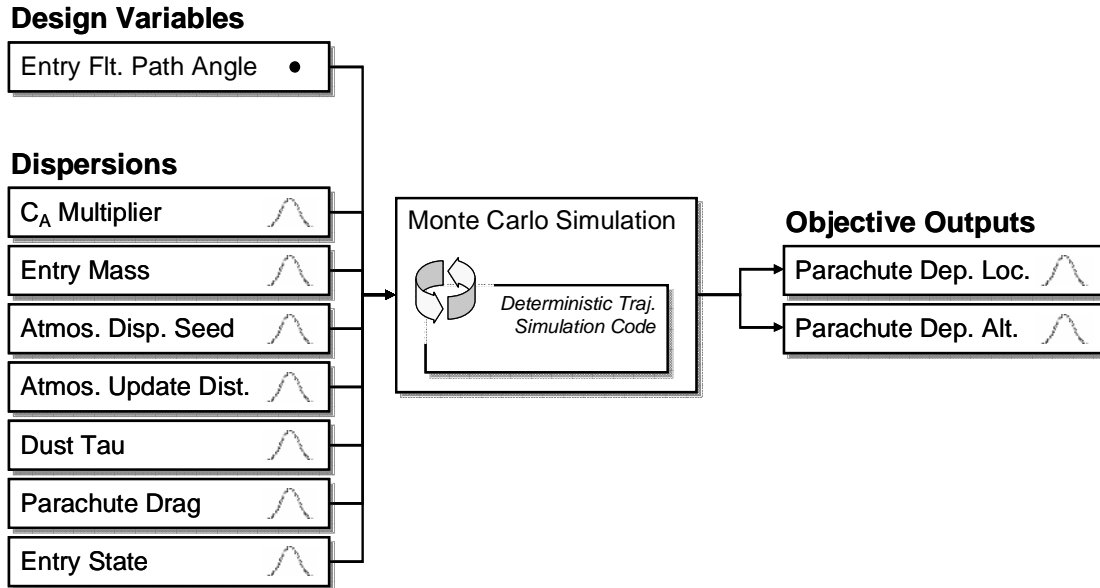**Figure 4. Definition of Monte Carlo input and output terms.**

Figure 5. Definition of Monte Carlo input and output terms in the context of this study.


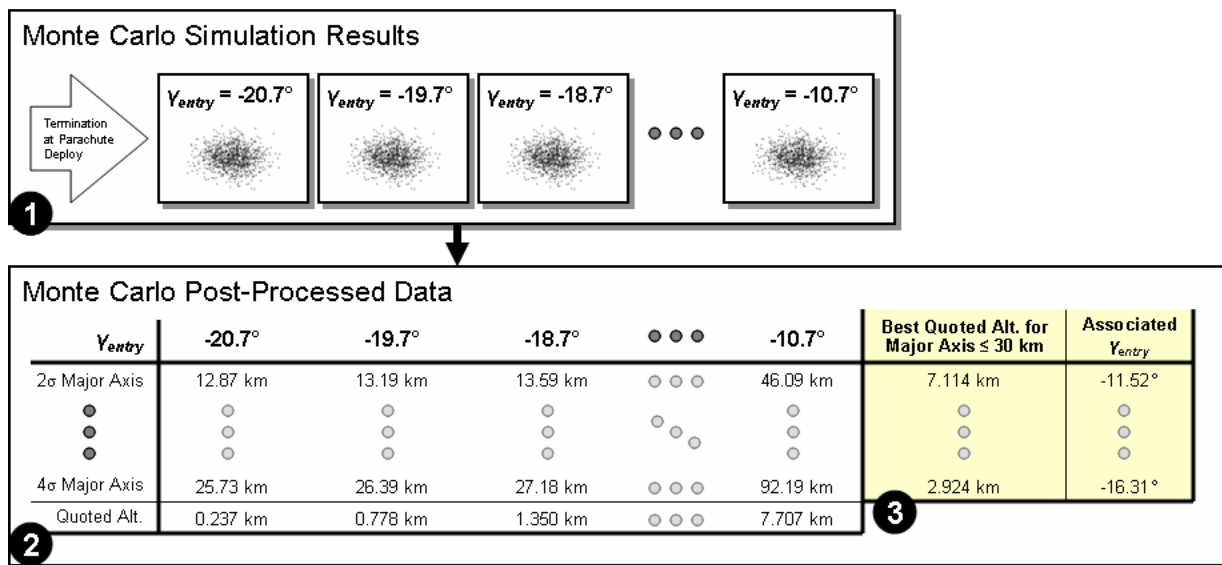**Figure 6. Possible combinations of two-objective trades in the context of this study's setup.**

Figure 7. Data Processing and Analysis Procedure for this study.

B. Scenario Definition

For convenience in conducting this study, raw Monte Carlo data is taken from a previous study of the ballistic entry, descent, and landing of an MSL-class vehicle.[8] The nominal entry mass is 2196 kg, and the hypersonic ballistic coefficient is 84 kg/m². Entry is assumed to occur on July 26, 2010 at -40.6° latitude and -62.9° longitude. The nominal initial state at entry is at an altitude of 120 km with a relative velocity of 5.39 km/s and relative azimuth of 73.5°. Dispersed inputs are indicated by Fig. 5, and details on dispersion ranges can be found in Ref. 8. For example, the atmospheric dust tau dispersion is modeled as uniform between 0.1 and 0.9. Also, the entry state dispersion is modeled from NASA Jet Propulsion Laboratory entry state covariance assuming that state-of-the-art delta differential one-way ranging (ΔDOR) is used and that the fifth trajectory correction maneuver (TCM-5) has been performed. As a result, the 3σ entry flight path angle uncertainty is 0.1°. This uncertainty is taken to be the same regardless of the nominal entry flight path angle (the design variable in the present study). Note also that all flight path angles referenced in this paper refer to relative (not inertial) flight path angles. Additional details on simulation setup parameters can be found in Ref. 8.

III.   Results and Discussion

A. Primary (50th Percentile) Altitude Plots

The result of plotting the data from Step 3 in the data processing procedure described earlier (and in Fig. 7) is shown in the six plots in Fig. 8 below. The first row of plots in Fig. 8 show the median altitude as a function of the robustness required of the parachute deploy ellipse major axis. The second row of Fig. 8 shows the derivative of the first row of curves, and the third row shows the entry flight path angle selections that allowed the altitudes in the first row to be achieved. The left plot in each row shows this median altitude as a function of σ-level, while the right plot shows this median altitude as a function of percentile.[§§] In all plots, the 3σ robustness level is denoted by a vertical red line. Note that the 30 km major axis requirement suggested by Figure 7 is represented by the cyan curve; other notional requirements (which might be considered as allowable depending on, for example, rover ground range abilities) are shown by additional curves. A discussion of these plots follows.

1. Altitude Performance Plots

The first row of plots is perhaps the most important because it shows the fundamental trade that exists between designing the trajectory for error ellipse robustness and designing it for altitude performance. For a given ellipse length, median altitude performance decreases with increasing robustness requirements as expected. For a 30 km ellipse length requirement (e.g., if a rover were designed to have a maximum ground range of 15 km), the difference between designing for 2.5σ (95.61%) and 3.0σ (98.89%) is a decrease in median parachute deploy altitude by 930 m. Thus, in this example, a 3% increase in the risk of falling outside the 30 km ellipse results in a 16% gain in parachute deploy altitude. This illustrates how plots such as in Fig. 8 may be used to trade altitude performance with error ellipse robustness.

Also regarding the first row of Fig. 8, an interesting note to make is that, because of the nonlinear relationship between σ-level and percentile, curves take a very different shape when plotted against σ-level compared to when plotted against percentile (despite the fact that the same data is plotted). In particular, the near-vertical slopes near the 100% robustness requirement show that large penalties in altitude occur as the parachute deploy location success requirement approaches 100%. Conversely, the shallower slopes at lower percentiles indicate that large increases in robustness are possible with little penalty in parachute deploy altitude.

2. Altitude Performance Derivative Plots

The second row of plots shows the derivatives of the first row. These plots are of interest because they directly show the sensitivities of altitude to small changes in error ellipse robustness requirements. For example, it can be seen in the left plot of the second row that for a 30 km ellipse length requirement, small deviations about a 3σ requirement result in 2.15 km of altitude for every 1σ of relaxation in error ellipse robustness (or, more conceptually accurate, 215 m of altitude for every 0.1σ of relaxation in error ellipse robustness). The right plot of the second row shows that for a 30 km ellipse length requirement, small deviations about a 3σ (98.89%) requirement result in 640 m of altitude for every 1% of relaxation in error ellipse robustness.

Of particular interest in these plots is the fact that these derivatives become larger in magnitude as the parachute deploy location success requirement approaches 100%. In fact, the derivatives with respect to percentile appear to

---

[§§] Note that the percentile shown is converted from the σ-level assuming bivariate normally-distributed data.
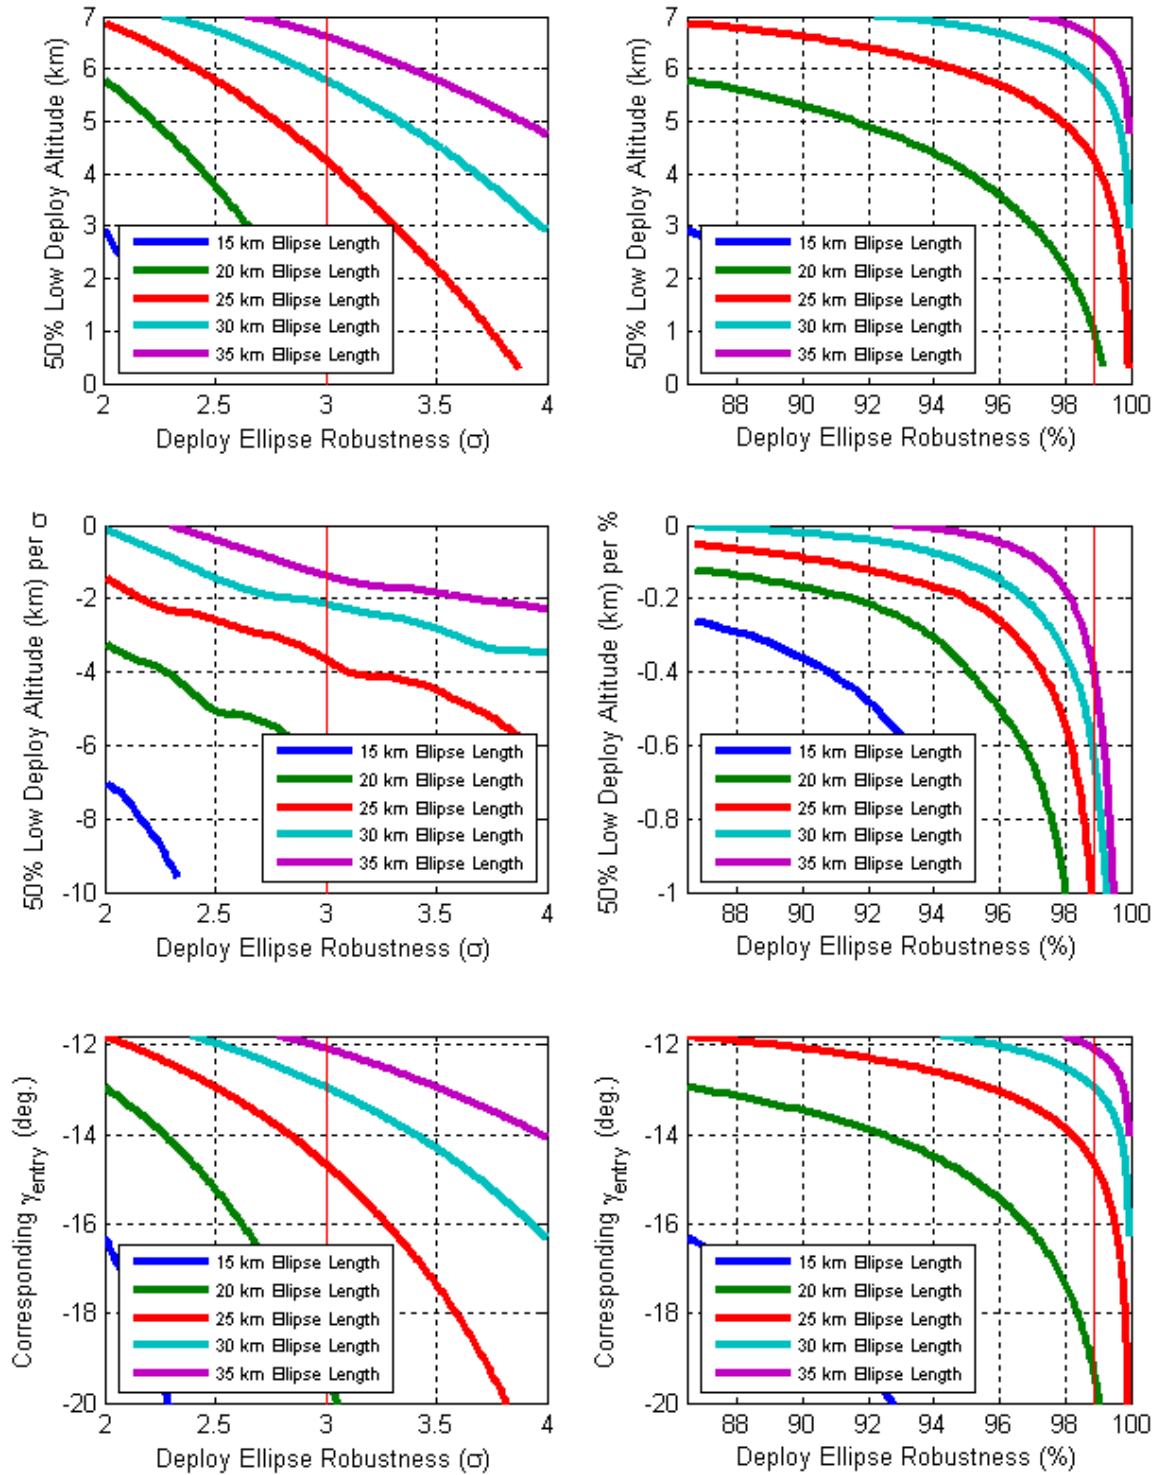
Figure 8. 50th Percentile Altitude Plots.

Each row of plots shows identical data; the left column shows data as a function of robustness expressed in σ-levels, while the right column shows the same data expressed as a function of robustness in terms of percentile. The first row shows the performance vs. robustness trade, and the second row shows the derivative of this. The third row shows the entry flight path angle (design variable) required to achieve the altitude given in the first row.

American Institute of Aeronautics and Astronautics

asymptote near 100%, indicating that if a robustness requirement is set too high, very large benefits can be obtained in terms of performance if even a small relaxation in robustness is permitted. This highlights the potential value of conducting sensitivity studies with respect to robustness. In the case of a 30 km error ellipse and a 3σ requirement for this scenario, these derivative plots do not show a compelling case for a change from the 3σ requirement, but it is easy to see from Fig. 8 how a larger initial robustness requirement (e.g. 4σ) or smaller required ellipse length (e.g. 20 km) could indeed involve high sensitivities to the robustness requirement.

3. Entry Flight Path Angle Plots

The third row of plots in Fig. 8 shows the entry flight path angles required to achieve the altitudes given in the first row (recall that this is the design variable that could be manipulated to yield the maximum attainable altitude for a given robustness requirement). These plots are important because they indicate to a designer *how* to design a trajectory to achieve the altitudes given in the first row. For example, achieving the maximum possible (5.8 km) median deploy altitude for a 30 km ellipse at the 3σ level requires an entry flight path angle of -12.9°.

Of some interest is the fact that, much like the other sets of plots, the entry flight path angle also has an asymptote near the 100th percentile in deploy ellipse robustness. That is, as expected, the vehicle must enter the Martian atmosphere much steeper as the deploy ellipse robustness requirement becomes more stringent.

B. Secondary (99.87th Percentile) Altitude Plots

It was indicated earlier in this paper that the quoted altitude performance was nominally chosen to be the median (50th percentile) value associated with the Monte Carlo set for a given entry flight path angle. However, depending on the user, it may be more appropriate to quote a more conservative altitude. Although this now begins to touch upon the trade between ellipse robustness and altitude robustness (see Fig. 6), a different quoted altitude can certainly be accommodated within the framework set forth here. For the design of MSL trajectories, the 99.87th percentile (low) parachute deploy altitude is typically quoted (i.e. nearly the lowest altitude detected throughout all Monte Carlo simulations). For demonstration purposes, this 99.87th percentile result is tracked in Fig. 9, which takes the same format as Fig. 8.

1. Altitude Performance Plots

Again, in Fig. 9, the first row of plots shows the fundamental trade between designing the trajectory for error ellipse robustness and designing it for altitude performance. The shapes of the curves in the first row are virtually identical to those in Fig. 8 but are shifted downward in altitude by roughly 800 m. For a 30 km ellipse length requirement, the difference between designing for 2.5σ (95.61%) and 3.0σ (98.89%) is a decrease in 99.87th percentile parachute deploy altitude by 940 m. In this example, a 3% increase in the risk of falling outside the 30 km ellipse results in a 18% gain in parachute deploy altitude. Note that the 940 m altitude gain in this case is nearly identical to the 930 m gain for the median altitude mentioned earlier; however, this results in a slightly larger percentage increase because the 99.87th percentile altitudes are shifted downward by several hundred meters with respect to the 50th percentile altitudes.

2. Altitude Performance Derivative Plots

Interestingly, the derivative curves in the second row of plots in Fig. 9 show little deviation from those of Fig. 8. This is significant because it suggests that the absolute sensitivities to robustness requirements are independent of the quoted altitude percentile. This is reflected in the example above where the difference in altitude performance between 2.5σ and 3.0σ was 940 m for the 99.87th percentile altitude and 930 m for the 50th percentile altitude. In other words, in the example above, this 0.5σ (3%) relaxation produces roughly a 930 m altitude gain regardless of the quoted percentile.

3. Entry Flight Path Angle Plots

As should be expected, the third row of plots in Fig. 9 is identical to the corresponding row in Fig. 8. As described earlier in this paper, the procedure for selecting the proper entry flight path angle involves selecting the angle which produces the highest parachute deploy altitude while still meeting a given ellipse length requirement. However, for this scenario (and for all practical ballistic entry scenarios), a steepening entry flight path angle simultaneously produces a monotonically decreasing ellipse length and monotonically decreasing deploy altitude. As a result, the entry flight path angle chosen is always the angle that produces exactly the maximum allowable ellipse length (e.g. 30 km). Thus, for this scenario, this selection does not involve consideration of altitude (the altitude is automatically maximized when the maximum allowable ellipse length is met) and, as a result, the entry flight path angle is not dependent on the quoted altitude percentile.
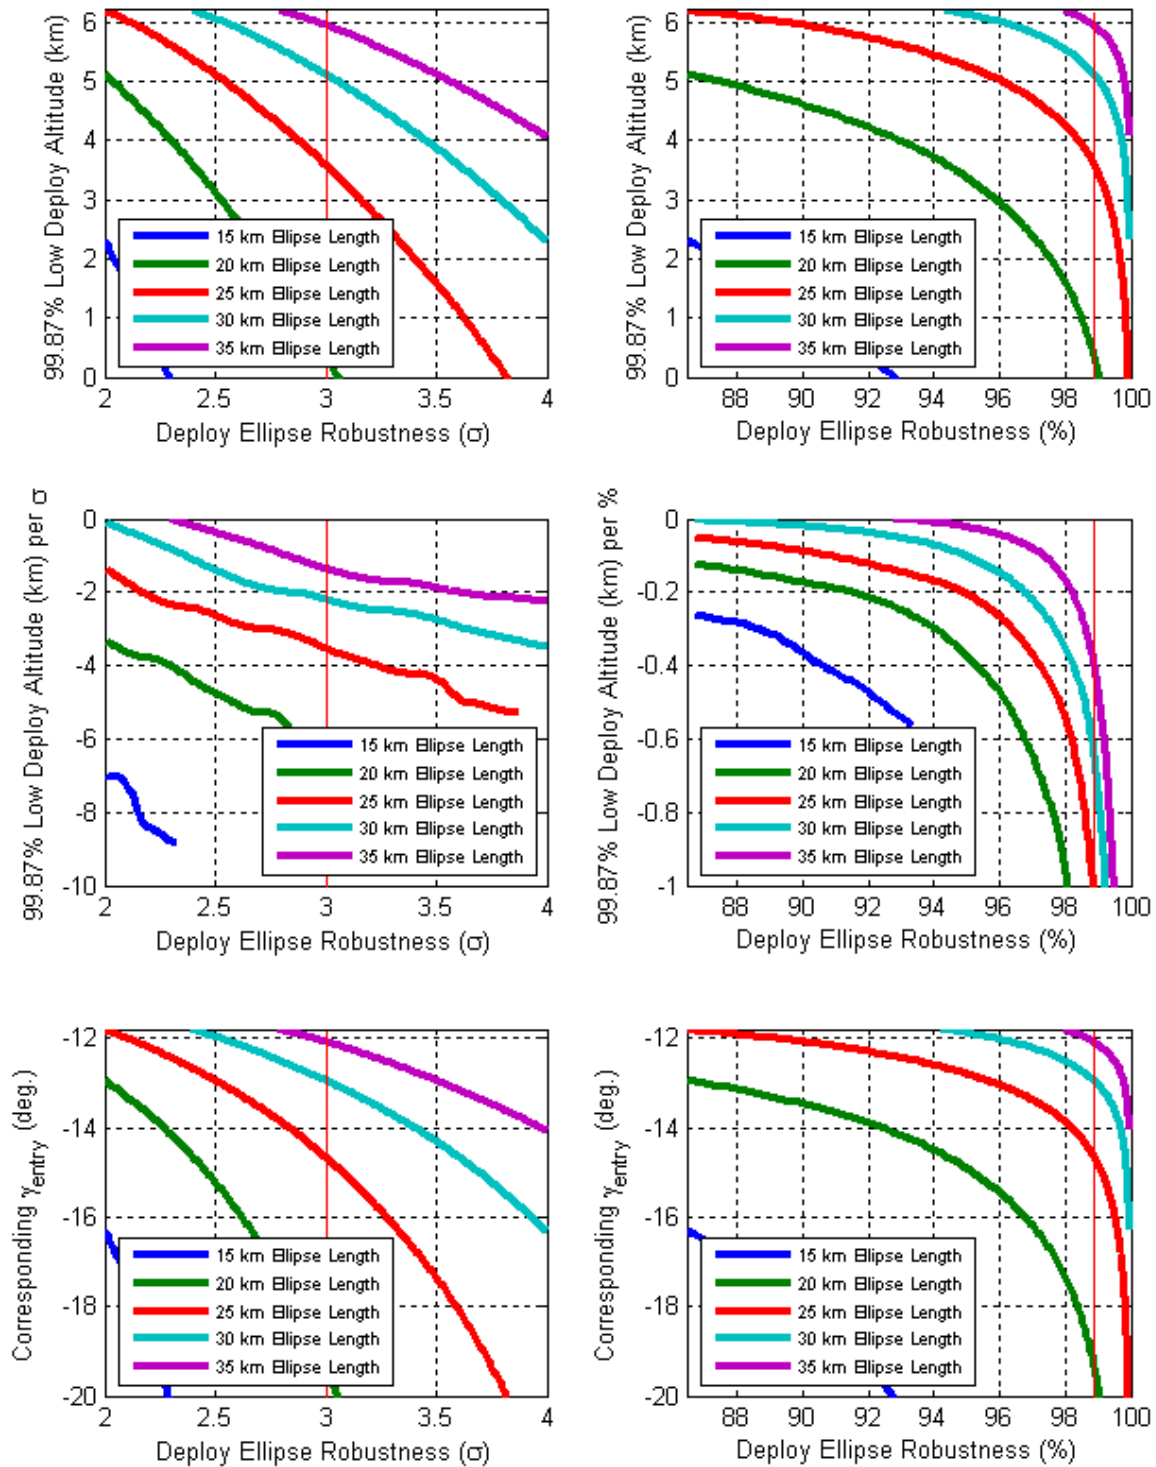
Figure 9. 99.87th Percentile (Low) Altitude Plots.

Each row of plots shows identical data; the left column shows data as a function of robustness expressed in σ-levels, while the right column shows the same data expressed as a function of robustness in terms of percentile. The first row shows the performance vs. robustness trade, and the second row shows the derivative of this. The third row shows the entry flight path angle (design variable) required to achieve the altitude given in the first row.

American Institute of Aeronautics and Astronautics

IV. Conclusions and Implications

This paper has presented an approach and a realistic example application for evaluating the performance implications of robustness requirements for Mars mission entry trajectory design. The study began with a discussion of the potential origins of 3σ as a robustness requirement and highlighted the empirical "rule of thumb" nature of this requirement from non-aerospace applications. A discussion was also included on the treatment of robustness as a distinct objective. An analysis framework and method was presented which involved the manipulation of design variables to effect trades between performance and robustness objectives. The scenario for which this method was illustrated was the ballistic entry of an MSL-class Mars entry vehicle. In this case, the design variable was entry flight path angle and parachute deploy altitude performance and error ellipse robustness were the objectives of interest. Resulting plots show the sensitivities between these objectives and trends in the entry flight path angles required to design to these objectives.

A. Choosing the "Best" Trajectory

As discussed earlier, trajectory performance should be recognized as distinct from trajectory robustness. The implication of this is that robustness should be treated as a tradable parameter in trajectory design (for example, not stringently set to 3σ without traceable justification). As a result of this tradability, the definition of a "best" trajectory depends on the designer. However, plots such as in Figs. 8 and 9 can facilitate the narrowing of the trade space by allowing the designer to visualize performance vs. robustness trades. Key features of such plots are near-vertical or near-horizontal lines, which indicate regions of the trade space where small sacrifices in one objective can produce large gains in another. In the case of the scenario examined here, the requirement of a 30 km, 3σ parachute deploy error ellipse fell in a region of the trade space without a clear, compelling reason for modification; that is, modest decreases in the robustness stringency produced modest increases in the parachute deploy altitude. Thus, in the example case, one might consider a 3σ requirement reasonable. However, it was clear that some regions of the trade-space exist – particularly at small-ellipse, high-robustness regions – where small relaxations in the robustness requirement produce large gains in parachute deploy altitude. Overall, this suggests that when robustness requirements are too stringent, large gains in performance can be made for small sacrifices in robustness. Plots such as Figs. 8 and 9 allow the designer to identify on a mission-by-mission basis how stringent "too stringent" actually is.

Additionally, a note should be made regarding limits to the relaxation of robustness. In reality, it is highly unlikely that a designer would relax robustness past a certain point even if large performance gains were possible. In part, this is because entry is one of a series of events involved in mission success; even if this total mission chain were as short as five events, an 80% probability of success for each event would result in an overall probability of success of just 33%. No reputable entry trajectory designer would design to the 80[th] percentile (i.e., that one in five cases would fall outside the error ellipse), and it is unlikely that the 90[th] percentile would be acceptable. Of course, this is the nature of objectives – if robustness could be relaxed without bound, it would not be a design objective. A more realistic lower bound might be the 95[th] percentile, but the exact limit is likely a soft one that depends on what performance gains can be obtained, the visualization and analysis of which have been outlined here.

B. Future Considerations

Demonstrated here was a realistic but relatively simple example of trading performance and robustness. This example is considered realistic in the sense that the vast majority of past missions to the surface of Mars have utilized a ballistic entry, and this procedure could be easily re-executed for any ballistic entry. Although conceptually nearly identical, practical complications increase as the number of free design variables increases (e.g., in lifting entry trajectories). For example, the true MSL entry will involve hypersonic guidance using a three-segment bank profile as a reference. Several variables are required to define this bank profile, and these design variables must be included when constructing the framework of Fig. 5. In order to execute the procedure outlined to construct Figs. 8 and 9, Monte Carlo simulations must be executed for many combinations of these design variables (constructed in the form of a grid or, more intelligently, through a design of experiments). The computational power that this requires could place practical limits on the complexity of the problem that is considered. However, a simple next step might be the analysis of a two-design-variable problem. For example, the Viking missions of the 1970s utilized a lifting entry but with a full-lift-up bank profile. In this case, the two design variables would be entry flight path angle (just as in this study) and vehicle lift-to-drag ratio (controlled by the designer through, for example, center-of-gravity placement).

Another avenue for pursuit farther in the future may be the addition of objectives to the trade-space. For example, altitude robustness was not explicitly considered here, and considerations such as heat rate and heat load

objectives were not addressed. Finally, it should be emphasized that the concepts and frameworks presented in this paper are applicable not only to Mars entry but to entry at other planets and moons as well.

Overall, this study has accomplished its original goal of tracing inherent trades between performance and robustness in Mars entry trajectory design. In the process, a general analysis framework was developed and a realistic example case was evaluated. It is hoped that these methods and ideas will find use within the broader trajectory design community.

Acknowledgments

References

[1]Braun, R.D. and Manning, R.M., "Mars Exploration Entry, Descent, and Landing Challenges." *Journal of Spacecraft and Rockets*, Vol. 44, No. 2, March-April 2007, pp. 310-323.

[2]Vysočanskiĭ, D.F. and Petunin, J.I., "Justification of the 3σ Rule for Unimodal Distributions." *Theory of Probability and Mathematical Statistics*, No. 21, 1980, pp. 25-36.

[3]Shewhart, W.A. *Economic Control of Quality of Manufactured Product*. New York: D. Van Nostrand Co., Inc., 1931.

[4]Finison, L.J., Finison, K.S., and Bliersbach, C.M., "The Use of Control Charts to Improve Healthcare Quality." *Journal for Healthcare Quality*, Vol. 15, No. 1, Jan.-Feb. 1993, pp. 9-23.

[5]Finison, L.J. and Finison, K.S., "Applying Control Charts to Quality Improvement." *Journal for Healthcare Quality*, Vol. 18, No. 6, Nov.-Dec. 1996, pp. 32-41.

[6]Wheeler, D.J. and Chambers, D.S., *Understanding Statistical Process Control*. 2nd ed. Knoxville: SPC Press, 1992.

[7]Taguchi, G., Elsayed, E.A., and Hsiang, T.C., *Quality Engineering in Production Systems*. New York: McGraw-Hill, 1989.

[8]Grant, M.J., Steinfeldt, B.A., Braun, R.D., and Barton, G.H. "Smart Divert: A New Entry, Descent, and Landing Architecture." AIAA 2009-0522. 47th AIAA Aerospace Sciences Meeting, Orlando, 5-8 Jan. 2009.