

# AN EFFICIENT NON-GAUSSIAN CHANCE CONSTRAINT METHOD FOR STOCHASTIC NONLINEAR PROBLEMS IN SPACEFLIGHT

Ethan R. Burnett\*, Spencer Boone<sup>†</sup>, and Niccolò Michelotti<sup>‡</sup>

Standard chance-constrained spacecraft guidance typically relies on the assumption that uncertainties in vehicle states obey Gaussian statistics. In frontier applications such as the cislunar environment or deep space flybys, the dynamics can be particularly nonlinear, and time between measurements or maneuvers can be long, leading to the need to make decisions whose outcomes produce non-Gaussian distributions. This paper demonstrates a non-Gaussian confidence boundary technique for stochastic guidance in such applications. Our approach is to consider the true confidence contour as a perturbation of the one predicted from covariance, then to derive perturbed boundary geometry from computed higher-order statistical moments. Applying this technique to so-called “banana-shaped distributions”, found in orbital mechanics problems, enables a simple parameterization of the confidence contour using the skew and kurtosis tensors. This parameterization is then applied to a stochastic and nonlinear impulsive spacecraft maneuver targeting problem, with special treatment of a relevant non-convex constraint.

## INTRODUCTION

State-of-the-art techniques of guidance and control are often stochastic in nature, whereby control of a nominal trajectory and expected statistical dispersions is jointly enforced. In contrast with many terrestrial robotics applications, spaceflight is often plagued by comparatively large state uncertainties and, sometimes, non-Gaussian statistics. These arise due to operation in dynamic regimes that are nonlinear and chaotic, and long periods without measurements or corrective control maneuvers.

The assumption of Gaussian statistics can be reasonable for some stochastic control problems in spaceflight. Spacecraft rendezvous problems can exploit the nearly linear dynamics of close-proximity relative motion, for which statistical distributions remain very nearly Gaussian. Ref. 1 leverages this property for passively safe spacecraft rendezvous with a linear covariance (“LinCov”) treatment of uncertainty evolution. This enables a “chance-constrained” approach, whereby satisfaction of the path constraint can be certified to a certain probability level. Chance-constrained approaches have become popular for safety-critical problems in spaceflight with sufficiently small and Gaussian dispersions.<sup>2</sup> The prospect of increasingly autonomous operation in space foresees more frequent measurements and more active maneuvering, thus such approaches may often be satisfactory. However, in long time-horizon maneuver planning, or in cases where accurate measurements are not available for long periods, the specter of non-Gaussian statistics cannot be ignored.

In spaceflight, a commonly observed manifestation of non-Gaussian statistics is with the so-called “banana-shaped” distributions. These emerge because the nonlinear dynamics of orbital mechanics

\*Assistant Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin, 2617 Wichita St, Austin TX, 78712, USA.

<sup>†</sup>Unaffiliated, Toulouse, FR.

<sup>‡</sup>PhD Student, Dept. of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milano IT.

tend to first stretch and then bend the ellipsoidal confidence region corresponding to an initially compact (and perhaps also Gaussian) distribution. For Keplerian problems, this problem can be partially avoided via use of “less nonlinear” coordinates such as polar coordinates or orbit elements (see e.g. Ref. 3). However, practical path constraints may not be convenient to express in such coordinates, and also non-Keplerian contexts such as cislunar astrodynamics challenge this solution strategy, as superior native coordinates are typically not available nor easy to identify.

For stochastic guidance and control with non-Gaussian statistics in astrodynamics, Monte Carlo methods are often the default strategy, but they are very slow. Overall the field is still in pursuit of accurate and efficient methods for such problems. One strategy is to probe the statistical moments beyond mean and covariance for geometric insights. Methods of estimating these moments include polynomial chaos expansion (PCE),<sup>4</sup> or the conjugate unscented transform (CUT),<sup>5</sup> which uses comparatively far fewer points and has the benefit of being deterministic. Recently, Ref. 6 explores nonlinear steering of non-Gaussian distribution back to an approximately Gaussian nature by leveraging feedback on the sigma points for the conjugate unscented transform. This extends the more familiar “covariance steering” (e.g., Ref. 7) to a regulation of higher statistical moments.

Active statistical steering methods introduce an additional continuous control feedback term which increases fuel use and might not always be operationally feasible. Instead, long-time horizon maneuvers can be planned, accepting the resulting non-Gaussian statistical distribution, and simply enforcing chance constraints on the non-Gaussian distribution itself. Along this line of thought, Ref. 8 applies a chance constraint approach with Gaussian Mixture Models to render the non-Gaussian distribution, however the approach is still rather numerically intensive. Ref. 9 introduces a promising new technique for analytically approximating confidence boundaries for the non-Gaussian “banana distributions” encountered in astrodynamics. Namely, select components of the higher-order moments of skew and kurtosis are used to directly inform geometric corrections to the ellipsoidal boundary predicted from covariance in the Gaussian limit. The methodology works well for high-fidelity cases we’ve tested,<sup>10</sup> and when paired with an efficient methodology for estimating skew and kurtosis, such as CUT4,<sup>5</sup> is quite numerically efficient. In this paper, which is devoted to application of the method, we focus on planar (two-state) parameterizations, but note that full 6-state distributions can be addressed by linear composition of each pair.<sup>10</sup>

In this paper, we demonstrate how possession of an accurate analytic approximation for a non-Gaussian confidence boundary allows for explicit consideration of the resulting non-convex constraint. We design a stochastic guidance approach that is fast and stable for admissible non-Gaussian cases with banana-shaped distributions, without excessive conservatism. After introducing the stochastic guidance approach, this paper revisits the numerical example of Ref. 9, then furnishes all developments in a timely and relevant end-to-end example: long-horizon re-entry corridor targeting of an Artemis II-like lunar free-return trajectory in the multibody Earth-Moon system.

## DERIVATION OF NON-GAUSSIAN CONFIDENCE CONTOUR

We derive corrections of the confidence bounds in the case of weakly non-Gaussian statistics (i.e. statistics for which the distribution is “banana-shaped” – see e.g. Ref. 3), leveraging these quantities. The approach we will expand on was first outlined in Ref. 9.

### Confidence Bounds: The Gaussian Case, and a Case for Small Corrections

The covariance ellipse can be used to bound Gaussian distributions with a certain  $k\sigma$  confidence level with covariance  $P \in \mathbb{R}^{6 \times 6}$ . A  $2 \times 2$  subset covariance is extracted as

$$\Sigma = \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy} & \Sigma_{yy} \end{bmatrix} = S P S^\top, \quad (1)$$

where  $S$  is a selection matrix composed of ones and zeros. Let  $\boldsymbol{\mu} = (\bar{x}, \bar{y})$  denote the mean position. The confidence ellipse is the set of points  $\mathbf{r} = (x, y)$  satisfying

$$(\mathbf{r} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{r} - \boldsymbol{\mu}) = k^2. \quad (2)$$

For a “3-sigma” scaled ellipse,  $k = 3$ . For a 2D confidence level  $p$ , one may use  $k^2 = \chi_{2,p}^2$ . The covariance is diagonalized as below:

$$\Sigma = R \Lambda R^\top, \quad (3a)$$

$$\Lambda = \text{diag}(\lambda_1, \lambda_2), \quad \lambda_1 \geq \lambda_2 > 0, \quad (3b)$$

$$R = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (3c)$$

The rotation angle  $\theta$  may be written as  $\theta = \frac{1}{2} \text{atan2}(2\Sigma_{xy}, \Sigma_{xx} - \Sigma_{yy})$ . The boundary of a certain statistical confidence level is parameterized by an ellipse with the following principal semi-axes:

$$\bar{a} = k\sqrt{\lambda_1}, \quad \bar{b} = k\sqrt{\lambda_2}. \quad (4)$$

We define local principal coordinates  $\mathbf{z} = (u, v)^\top$ , which point along the long and short axes:

$$u = \bar{a} \cos t, \quad v = \bar{b} \sin t, \quad t \in [0, 2\pi). \quad (5)$$

Thus a point on the elliptical boundary is parameterized as below:

$$\mathbf{r}(t) = \boldsymbol{\mu} + R \begin{pmatrix} u(t) \\ v(t) \end{pmatrix} \quad (6)$$

We additionally define the normalized coordinates  $\hat{u} = u/\sqrt{\lambda_1}$  and  $\hat{v} = v/\sqrt{\lambda_2}$ , which in the case of the parameterization of a Gaussian boundary satisfy  $\hat{u} = k \cos t$ ,  $\hat{v} = k \sin t$ .

We seek methods suitable for when a distribution is “weakly” non-Gaussian, i.e. confidence contours close to the mean still bear strong similarity to those expected from a Gaussian distribution, but further out they lose their accuracy. We seek suitable and tractable first-order corrections (i.e. linear in a parameter). The dominant effect commonly observed in astrodynamics problems is the stretching of the distribution along the so-called “maximum stretching direction”, which corresponds with the maximum eigendirection of the Cauchy-Green tensor (e.g., Ref. 11). Beyond linear effects, the first and most noteworthy key feature is the bend of the distribution: the decoupling of  $\hat{u}$ ,  $\hat{v}$  from the Gaussian case is lost. In particular, sufficiently large departures in  $\hat{u}$  from the mean will carry a bend in  $\hat{v}$ . As a first-order correction of the elliptical assumption, we thus apply the following *ansatz* which breaks the independence of the principal coordinates for large deviations:

$$\hat{v}_q \approx \beta_q + \alpha_q \hat{u}_q^2 \quad (7)$$

with index  $q$  for each possible two-state slice, and parameters  $(\alpha_q, \beta_q)$  are to be determined. The motivation for the quadratic *ansatz* is as follows: Viewing the action of the flow of the nonlinear dynamics on the statistical distribution (and hence its confidence boundary) as a weakly nonlinear map, the first expected contribution after the linear effect should be quadratic.

The next feature is less obvious. The probability density towards  $\hat{u} = +g$  and  $\hat{u} = -g$  is no longer even for some  $g \gg 0$ . In other words, the symmetry of the confidence bounds is broken. This effect is more noticeable along the long axis of the distribution, so we seek a correction to  $\hat{u}$  but not  $\hat{v}$ . This is similar to the logic by which we prioritized bending of the form  $\hat{v} \propto \hat{u}^2$  but ignored

the (assumed sub-dominant) analogous bending term  $\hat{u} \propto \hat{v}^2$ . Focusing just on corrections along  $\hat{u}$ , the Cornish-Fisher expansion provides an asymptotic approximation of the quantiles of a univariate non-Gaussian distribution based on its cumulants.<sup>12</sup> Below we provide the first-order term:

$$x(p) \approx \mu + \sigma \left( z + \frac{\gamma_1}{6}(z^2 - 1) \right) \quad (8)$$

where  $x$  follows a slightly non-Gaussian univariate distribution, with mean  $\mu$  and standard deviation  $\sigma$ . Furthermore  $z = \Phi^{-1}(p)$  where  $\Phi$  is the cumulative distribution function of the standard normal distribution, i.e.  $\Phi(3) \approx 0.9987$ . Lastly,  $\gamma_1$  is expressed in terms of skew and standard deviation:

$$\gamma_1 = \frac{\mu_3}{\sigma^3}, \quad \mu_n = \mathbb{E}[(x - \mu)^n] \quad (9)$$

We seek an additive  $t$ -periodic correction to the (whitened) Gaussian parameterization  $\hat{u}(t) = k \cos t$  which obeys the following properties for  $c(k) = \frac{\gamma_1}{6}(k^2 - 1)$ :

$$\delta \hat{u}(0) = \delta \hat{u}(\pi) = c(k) \quad (10a)$$

$$\delta \hat{u}\left(\frac{\pi}{2}\right) = \delta \hat{u}\left(\frac{3\pi}{2}\right) = 0 \quad (10b)$$

$$\delta \hat{u}'(0) = \delta \hat{u}'(\pi) = 0 \quad (10c)$$

$$\delta \hat{u}'\left(\frac{\pi}{2}\right) = \delta \hat{u}'\left(\frac{3\pi}{2}\right) = 0 \quad (10d)$$

The first two conditions state that along the direction  $\hat{u}$ , we expect to recover the univariate correction: The  $k\sigma$  boundaries shift to  $-k + c(k)$  for  $\hat{u} < 0$  and  $k + c(k)$  for  $\hat{u} > 0$ . The second two enforce no net change when  $\hat{u} = 0$ . The last four conditions are for regularity – to avoid unphysical directional biases. Noting that the univariate Cornish-Fisher expansion tells us nothing about how the confidence contour changes except purely along the long axis of the distribution, we furthermore require evenness of  $\delta \hat{u}(t)$  about  $t = 0$  and  $t = \pi$ , as any other choice induces off-axis asymmetries that cannot be justified. The class of all functions satisfying these conditions can be expressed as:

$$\delta \hat{u}_q(t) = \eta_{0,q} + \sum_{n \geq 1} \eta_{n,q} \cos(nt) \quad (11)$$

We will start with the derivation of the proper coefficients for Eq. (7), then revisit this.

### Moment Identities

Let  $\mathbf{X} = (X_1, \dots, X_d)^\top$  be a random vector with finite central moments through order four, mean  $\boldsymbol{\mu}$ , covariance  $\Sigma$ , third-order moment tensor  $M^{(3)}$ , and fourth-order moment tensor  $M^{(4)}$ . The tensors  $\Sigma_{ij}$ ,  $M_{ijk}^{(3)}$ , and  $M_{ijkl}^{(4)}$  are invariant under permutation of their indices. For a multivariate Gaussian distribution,

$$M_{ijk}^{(3)} = 0, \quad (12)$$

and the fourth central moment is generated by the covariance as

$$M_{ijkl}^{(4)} = \Sigma_{ij}\Sigma_{kl} + \Sigma_{ik}\Sigma_{jl} + \Sigma_{il}\Sigma_{jk}. \quad (13)$$

Equivalently, the fourth-order cumulant tensor vanishes for Gaussian statistics:

$$\kappa_{ijkl} \equiv M_{ijkl}^{(4)} - (\Sigma_{ij}\Sigma_{kl} + \Sigma_{ik}\Sigma_{jl} + \Sigma_{il}\Sigma_{jk}) \quad (14)$$

$$\kappa_{ijkl} = 0 \quad \forall i, j, k, l. \quad (15)$$

Thus, deviations from Eqs. (12) and (15) provide third- and fourth-order indicators of non-Gaussianity.

### First Correction: The bend of the banana

Referring the reader again to Eq. (7), for a given slice  $q$ , we solve for parameters  $(\alpha_q, \beta_q)$  minimizing the expected square of the fit error below:

$$(\alpha, \beta) = \arg \min_{\alpha, \beta} \mathbb{E} \left[ (\hat{v} - \beta - \alpha \hat{u}^2)^2 \right]. \quad (16)$$

Applying the first-order conditions of optimality w.r.t.  $\alpha$  and  $\beta$ , we obtain:

$$\mathbb{E}[\hat{v}] = \beta + \alpha \mathbb{E}[\hat{u}^2], \quad (17a)$$

$$\mathbb{E}[\hat{v} \hat{u}^2] = \beta \mathbb{E}[\hat{u}^2] + \alpha \mathbb{E}[\hat{u}^4]. \quad (17b)$$

These are linear equations in  $\alpha, \beta$ . Some further manipulations isolate the corrective coefficients:

$$\alpha = \frac{\mathbb{E}[\hat{v} \hat{u}^2] - \mathbb{E}[\hat{u}^2] \mathbb{E}[\hat{v}]}{\mathbb{E}[\hat{u}^4] - \mathbb{E}[\hat{u}^2]^2}, \quad (18a)$$

$$\beta = \frac{\mathbb{E}[\hat{u}^4] \mathbb{E}[\hat{v}] - \mathbb{E}[\hat{u}^2] \mathbb{E}[\hat{v} \hat{u}^2]}{\mathbb{E}[\hat{u}^4] - \mathbb{E}[\hat{u}^2]^2}. \quad (18b)$$

For a 2-state sample  $\mathbf{r}$  (not necessarily from a Gaussian distribution) and a given slice covariance  $\Sigma_q$  of the form of Eq. (1), we define the following convenient whitening transformation:

$$W_q = \Lambda_q^{-1/2} R_q^\top, \quad (19)$$

$$\begin{pmatrix} \hat{u} \\ \hat{v} \end{pmatrix} = W_q (\mathbf{r} - \boldsymbol{\mu}_q). \quad (20)$$

We seek to apply the transformation of Eq. (20) to the random state from our non-Gaussian distribution, substitute the resulting identities into Eqs. (18a)-(18b), and solve for the coefficients  $\alpha, \beta$ . Further simplification is possible. First it is easy to show that  $\mathbb{E}[\hat{u}] = \mathbb{E}[\hat{v}] = 0$ , leveraging some useful quantities:

$$\mathbf{a} = W_q^\top \begin{pmatrix} 1 \\ 0 \end{pmatrix}; \mathbf{b} = W_q^\top \begin{pmatrix} 0 \\ 1 \end{pmatrix}; \boldsymbol{\delta} = \mathbf{r} - \boldsymbol{\mu}_q. \quad (21)$$

For nonlinear terms we switch to an index notation for convenience:

$$\hat{u}^2 = (\mathbf{a}^\top \boldsymbol{\delta})^2 = \left( \sum_j a_j \delta_j \right) \left( \sum_k a_k \delta_k \right) = \sum_j \sum_k a_j a_k \delta_j \delta_k \quad (22a)$$

$$\hat{v} \hat{u}^2 = \sum_i b_i \delta_i \cdot \hat{u}^2 = \sum_i \sum_j \sum_k b_i a_j a_k \delta_i \delta_j \delta_k \quad (22b)$$

From such expressions it is simple to compute the necessary expected values, e.g.:

$$\mathbb{E}[\hat{v} \hat{u}^2] = \sum_i \sum_j \sum_k b_i a_j a_k \mathbb{E}[\delta_i \delta_j \delta_k] = \sum_{i,j,k} b_i a_j a_k M_{q,ijk}^{(3)} \quad (23)$$

Lastly, we establish one more identity by noting  $\Sigma_q = R_q \Lambda_q R_q^\top$ :

$$\mathbb{E}[\hat{u}^2] = \mathbf{a}^\top \Sigma_q \mathbf{a} = \begin{bmatrix} 1 & 0 \end{bmatrix} \Lambda_q^{-1/2} R_q^\top \Sigma_q R_q \begin{pmatrix} \Lambda^{-1/2} \end{pmatrix}^\top \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 1 \quad (24)$$

The final result for Eqs. (18a)-(18b) is thus obtained, solely as a function of skew and kurtosis:

$$\alpha = \frac{\sum_{i,j,k} b_i a_j a_k M_{q,ijk}^{(3)}}{\sum_{i,j,k,l} a_i a_j a_k a_l M_{q,ijkl}^{(4)} - 1}, \quad (25a)$$

$$\beta = -\alpha. \quad (25b)$$

These satisfy the expected property  $\alpha = \beta = 0$  for a purely Gaussian distribution. Finally, we can establish the first corrected parameterization of the confidence contour. Revisiting Eqs. (5) and (7), we leverage the zero-order solution  $\hat{u} = k \cos t$ , and expect that  $\hat{v}$  should be unmodified from the Gaussian case in the event that  $\alpha = 0$  (in other words, we retain an unbiased reference value). Thus:

$$v = \bar{b} \sin t + \alpha \sqrt{\lambda_2} (k^2 \cos^2 t - 1), \quad t \in [0, 2\pi) \quad (26)$$

where  $k$  is the sigma confidence level and the definitions  $\bar{a} = k\sqrt{\lambda_1}$  and  $\bar{b} = k\sqrt{\lambda_2}$  are reused from earlier and these are unrelated to the components of  $\mathbf{a}$ ,  $\mathbf{b}$ . This result notably recovers the expected “bending” form, via the cosine-squared term in the  $v$  coordinate of the contour.

### Second Correction: Long axis asymmetry

As the first four moments of a non-Gaussian distribution are not enough to uniquely determine the isoprobability contour, we seek the simplest justifiable correction based on known information and constraints. Returning to Eq. (11), we establish the simplest non-trivial satisfier of Eq. (10):

$$\delta \hat{u} = \eta_0 + \eta_2 \cos 2t \quad (27)$$

Applying the constraints, we obtain:

$$\eta_0 = \eta_2 \quad (28a)$$

$$2\eta_0 = c(k) \quad (28b)$$

$$c(k) = \frac{1}{6} \mathbb{E}[\hat{u}^3] (k^2 - 1) \quad (28c)$$

Or, rewritten below in terms of the skew directly:

$$c(k) = \frac{k^2 - 1}{6} \sum_{i,j,k} a_i a_j a_k M_{q,ijk}^{(3)} \quad (29)$$

Thus, applying the double-angle identity and rescaling as  $u(t) = \sqrt{\lambda_1} \hat{u}(t)$ , the final corrected  $u(t)$  is:

$$u(t) = \bar{a} \cos t + c(k) \sqrt{\lambda_1} \cos^2 t \quad (30)$$

As before, this result recovers the Gaussian case if  $c(k) = 0$  (no skew). The final corrected contour is obtained by combining Eqs. (26) and (30) and mapping back to  $\mathbf{X}$  coordinates via Eq. (6).

**Key Result 1:** In our non-Gaussian setting, the ellipse is replaced by the banana contour

$$u(t) = \bar{a} \cos t + c(k) \sqrt{\lambda_1} \cos^2 t, \quad (31a)$$

$$v(t) = \bar{b} \sin t + \alpha \sqrt{\lambda_2} (k^2 \cos^2 t - 1), \quad (31b)$$

with mapping

$$\mathbf{r}(t) = \boldsymbol{\mu} + R \begin{pmatrix} u(t) \\ v(t) \end{pmatrix}. \quad (32)$$

Thus the confidence contour of the banana distribution is treated as a first-order deformation of the Gaussian linear covariance ellipse. Despite its simplicity, numerical experiments have shown it works well for even very large perturbed distributions in astrodynamics problems. As a first-order correction, it can be applied simultaneously to multiple distinct coordinate pairs (e.g.  $u$ - $v$  and  $u$ - $w$  for principal directions  $u, v, w$ ). This is discussed in Ref. 10, along with high-fidelity test cases.

## GEOMETRIC CONSTRAINTS

Now we introduce the local half-plane constraint

$$\mathbf{n}^\top \mathbf{r} - b_0 \leq 0, \quad (33)$$

where  $\mathbf{n} \in \mathbb{R}^2$  is the outward normal of the local boundary approximation and  $b_0 \in \mathbb{R}$  is the associated offset. Define the contour residual

$$\psi(\boldsymbol{\chi}, t) \equiv \mathbf{n}^\top \mathbf{r}(\boldsymbol{\chi}, t) - b_0, \quad (34)$$

where  $\boldsymbol{\chi}$  denotes the free optimization variables (e.g. initial state/parameter, delta-V, or other control variables). The requirement that the confidence contour lie on the safe side of the half-plane is

$$\max_{t \in [0, 2\pi)} \psi(\boldsymbol{\chi}, t) \leq 0. \quad (35)$$

The basic problem geometry is depicted in Fig. 1(a).

### Gaussian case

We first recover the familiar Gaussian result. Let  $\mathbf{m}(\boldsymbol{\chi}) \equiv R(\boldsymbol{\chi})^\top \mathbf{n} = \begin{pmatrix} m_1(\boldsymbol{\chi}) \\ m_2(\boldsymbol{\chi}) \end{pmatrix}$ . Then, for the Gaussian contour,

$$\psi_G(\boldsymbol{\chi}, t) = \mathbf{n}^\top \boldsymbol{\mu}(\boldsymbol{\chi}) - b_0 + \mathbf{n}^\top R \begin{pmatrix} u_G(t) \\ v_G(t) \end{pmatrix}. \quad (36)$$

Applying the definitions of  $\mathbf{m}(\boldsymbol{\chi})$ , and of  $\bar{a}$  and  $\bar{b}$  from Eq. (4),

$$\psi_G(\boldsymbol{\chi}, t) = \mathbf{n}^\top \boldsymbol{\mu}(\boldsymbol{\chi}) - b_0 + k m_1(\boldsymbol{\chi}) \sqrt{\lambda_1(\boldsymbol{\chi})} \cos t + k m_2(\boldsymbol{\chi}) \sqrt{\lambda_2(\boldsymbol{\chi})} \sin t. \quad (37)$$

For any scalars  $p, q$ ,  $\max_t (p \cos t + q \sin t) = \sqrt{p^2 + q^2}$ , so

$$\max_t \psi_G(\boldsymbol{\chi}, t) = \mathbf{n}^\top \boldsymbol{\mu}(\boldsymbol{\chi}) - b_0 + k \sqrt{m_1(\boldsymbol{\chi})^2 \lambda_1(\boldsymbol{\chi}) + m_2(\boldsymbol{\chi})^2 \lambda_2(\boldsymbol{\chi})}. \quad (38)$$

Furthermore, because  $\mathbf{m} = R^\top \mathbf{n}$  and  $\Sigma = R \Lambda R^\top$ , it is easy to show  $m_1^2 \lambda_1 + m_2^2 \lambda_2 = \mathbf{n}^\top \Sigma \mathbf{n}$ . Therefore, the Gaussian half-plane chance surrogate is recovered as:

$$\max_t \psi_G(\boldsymbol{\chi}, t) = \mathbf{n}^\top \boldsymbol{\mu}(\boldsymbol{\chi}) - b_0 + k \sqrt{\mathbf{n}^\top \Sigma(\boldsymbol{\chi}) \mathbf{n}} \leq 0. \quad (39)$$

This convenient geometric constraint facilitates covariance-based strategies for stochastic control.<sup>8,13</sup>

### Non-Gaussian case

Substituting the banana parameterization (31) into (34) gives

$$\begin{aligned}\psi(\boldsymbol{\chi}, t) &= \mathbf{n}^\top \mu(\boldsymbol{\chi}) - b_0 + m_1(\boldsymbol{\chi}) u(t) + m_2(\boldsymbol{\chi}) v(t) \\ &= \mathbf{n}^\top \mu(\boldsymbol{\chi}) - b_0 + m_1 \left( a \cos t + c(k, \boldsymbol{\chi}) \sqrt{\lambda_1} \cos^2 t \right)\end{aligned}\quad (40)$$

$$+ m_2 \left( b \sin t + \alpha(\boldsymbol{\chi}) \sqrt{\lambda_2} (k^2 \cos^2 t - 1) \right). \quad (41)$$

Grouping terms by trigonometric dependence yields

$$\psi(\boldsymbol{\chi}, t) = A(\boldsymbol{\chi}) + B(\boldsymbol{\chi}) \cos t + C(\boldsymbol{\chi}) \sin t + D(\boldsymbol{\chi}) \cos^2 t, \quad (42)$$

with coefficients defined as below:

$$A(\boldsymbol{\chi}) = \mathbf{n}^\top \mu(\boldsymbol{\chi}) - b_0 - m_2(\boldsymbol{\chi}) \alpha(\boldsymbol{\chi}) \sqrt{\lambda_2(\boldsymbol{\chi})}, \quad (43a)$$

$$B(\boldsymbol{\chi}) = k m_1(\boldsymbol{\chi}) \sqrt{\lambda_1(\boldsymbol{\chi})}, \quad (43b)$$

$$C(\boldsymbol{\chi}) = k m_2(\boldsymbol{\chi}) \sqrt{\lambda_2(\boldsymbol{\chi})}, \quad (43c)$$

$$D(\boldsymbol{\chi}) = m_1(\boldsymbol{\chi}) c(k, \boldsymbol{\chi}) \sqrt{\lambda_1(\boldsymbol{\chi})} + m_2(\boldsymbol{\chi}) \alpha(\boldsymbol{\chi}) k^2 \sqrt{\lambda_2(\boldsymbol{\chi})}. \quad (43d)$$

Note that the banana correction contributes only through the constant and the  $\cos^2 t$  terms. With the addition of this term, we find no way to recover an analog for Eq. (38), blocking recovery of a non-Gaussian analog for Eq. (39). This is due to branching behavior in the angle parameter  $t$  of the contour yielding the worst-case value of the half plane constraint.

Two major points are worth emphasizing before continuing. First, simple *conservative* surrogates are still possible. For example, for  $D \geq 0$ , the following holds:

$$\psi(\boldsymbol{\chi}, t) = A + B \cos t + C \sin t + D \cos^2 t \leq A + \sqrt{B^2 + C^2} + D. \quad (44)$$

Second, and critically, we emphasize that a pragmatic implementation of the non-Gaussian confidence boundary is to simply discretely sample along the boundary of the banana contour and jointly enforce inequality constraints on each point. The downsides are the potential for the sampling to be too sparse, and in the wasted computational effort in evaluating inequalities on many points, when a worst-violating or closest-violating point always exists. However, this approach works for preliminary implementations.

Solving for the worst-violating angle is numerically inexpensive compared to evaluating the constraint function or its gradients. We thus aim to first solve for this worst-case angle  $t^*$  maximizing Eq. (42). A change of variables is convenient for this:

$$\zeta \equiv \cos t, \quad \zeta \in [-1, 1]. \quad (45)$$

$$\sin t = \pm \sqrt{1 - \zeta^2}, \quad (46)$$

$$\psi_{\pm}(\boldsymbol{\chi}, \zeta) = A(\boldsymbol{\chi}) + B(\boldsymbol{\chi}) \zeta \pm C(\boldsymbol{\chi}) \sqrt{1 - \zeta^2} + D(\boldsymbol{\chi}) \zeta^2. \quad (47)$$

The branch ambiguity is in the sign of the square-root term, and the maximizing sign is the one which matches the sign of  $C(\boldsymbol{\chi})$ . Hence,

$$\max_t \psi(\boldsymbol{\chi}, t) = \max_{\zeta \in [-1, 1]} \left( A(\boldsymbol{\chi}) + B(\boldsymbol{\chi}) \zeta + |C(\boldsymbol{\chi})| \sqrt{1 - \zeta^2} + D(\boldsymbol{\chi}) \zeta^2 \right). \quad (48)$$

Thus the continuous half-plane support computation reduces to a one-dimensional maximization over  $\zeta \in [-1, 1]$  for some maximizing  $\zeta_*$ , which is trivially numerically solvable.

By definition, the following hold

$$\cos t^* = \zeta_* \quad (49a)$$

$$\sin t^* = \text{sgn}(C(\chi))\sqrt{1 - \zeta_*^2}, \quad C(\chi) \neq 0. \quad (49b)$$

Therefore

$$t^* = \text{atan2}\left(\text{sgn}(C(\chi))\sqrt{1 - \zeta_*^2}, \zeta_*\right), \quad (50)$$

which is unique whenever  $C(\chi) \neq 0$ . When  $C(\chi) = 0$ , the branching of Eq. (47) disappears and the solution for  $t^*$  is instead to maximize a quadratic in  $\zeta$  then solve for  $t^*$ . Once  $t^*$  is found, no other points on the contour are relevant to enforcing the constraint.

**Solution by Sequential Corrections:** We seek values of the general problem free variables  $\chi$  minimizing some  $J(\chi)$  while also satisfying the inequality constraint  $g(\chi) \leq 0$ , where  $g(\chi) = \max_t \psi(\chi, t) = \psi(\chi, t^*)$ . As this is a nonlinear problem, these are to be obtained by a local sequential correction procedure. When the constraint is violated, once the worst-violating contour point  $t^*$  is identified, a local correction can be defined. Consider for the sake of argument a simple first-order (gradient-based) local correction:

$$\Delta\chi_i = -K_1 \frac{g}{\|\nabla_{\chi} g\|^2} \nabla_{\chi} g - K_2 \left( I - \frac{\nabla_{\chi} g \nabla_{\chi} g^{\top}}{\|\nabla_{\chi} g\|^2} \right) \nabla_{\chi} J, \quad (51)$$

for some  $K_1, K_2 > 0$  and iterate  $i > 1$ , where  $g, J, \nabla g$ , and  $\nabla J$  are evaluated on the prior iterate  $\chi_{i-1}$ . The first term in Eq. (51) reduces the constraint violation, and the second term reduces the cost along an orthogonal direction. The associated gradient is  $\nabla_{\chi} g = \nabla_{\chi} \psi(\chi, t^*)$ :

$$\nabla_{\chi} \psi(\chi, t) = \nabla A(\chi) + \cos t \nabla B(\chi) + \sin t \nabla C(\chi) + \cos^2 t \nabla D(\chi). \quad (52)$$

All gradients are defined in Appendix B. This is the gradient of the active half-plane residual evaluated at the worst contour point. As the worst violator  $t^*$  moves, holding  $\chi$  constant, the term  $\nabla g$  updates as in Eq. (52), which affects corrective schemes such as Eq. (51). Thus, we seek to understand how the nature of changes in  $t^*$  affects the stability of a sequential correction procedure.


**Confidence contour non-convexity and constraint effects:** For a coplanar half-plane constraint, the only possible non-smooth change in  $t^*$  in the present geometry is the competition between the two “tips” of the banana contour. This occurs when the two ends of the confidence boundary are simultaneously equidistant or near equidistant to the half-plane constraint surface – see Fig. 1(b). This is a concern because the corrective equation Eq. (51) is informed directly by the gradient of the constraint violation, Eq. (52), which itself is a trigonometric function of the maximum violating angle, holding  $\chi$  fixed. A non-smooth change in  $t^*$  from one iteration to the next thus invites the possibility of undesirable behaviors such as chatter or optimizer stall. It can be shown e.g. by inspection of Eqs. (31) that when the  $D(\chi)$  term dominates, this occurs near the following angles:

$$t_{\text{crit}} \approx \{0, \pi\}, \quad (53)$$

Let us elaborate further on this problematic orientation. The following holds:

$$B(\chi) = k m_1(\chi) \sqrt{\lambda_1(\chi)}, \quad m_1(\chi) = e_1(\chi)^{\top} \mathbf{n}. \quad (54)$$

A “near-tie” in constraint violation between the two banana tips occurs only when  $B(\chi) \approx 0$  or  $m_1(\chi) = \mathbf{e}_1(\chi)^\top \mathbf{n} \approx 0$ , where  $\mathbf{e}_1$  is the first basis vector of  $R$ . This is the degenerate alignment in which the half-plane normal is nearly orthogonal to the long principal axis. Only in the narrow regime of small  $|B(\chi)|$  can the active parameter jump non-smoothly between  $t^* \approx 0$  and  $t^* \approx \pi$ .



**Figure 1. Non-Gaussian Chance Constraint Geometry**

### Overcoming the Non-Convex Constraint

**I. Direct Gradient Smoothing:** We must first identify exactly the  $t_{\text{crit}}$  angles, and then propose a stable alternative to the discrete jump in gradient between these values. Starting from Eq. (42), we first impose  $B(\chi) = 0$ . Then

$$\psi_0(\chi, t) = A(\chi) + C(\chi) \sin t + D(\chi) \cos^2 t. \quad (55)$$

Differentiating with respect to  $t$ , the critical points can be identified as Eq. (57):

$$\psi'_0(\chi, t) = C(\chi) \cos t - D(\chi) \sin 2t = \cos t \left( C(\chi) - 2D(\chi) \sin t \right). \quad (56)$$

$$\cos t = 0 \quad \text{or} \quad \sin t = \frac{C(\chi)}{2D(\chi)}. \quad (57)$$

The condition  $C(\chi)/(2D(\chi)) < 1$  enables the second set. The two critical angles of interest, namely the ones lying on the two banana-tip branches near 0 and  $\pi$ , are therefore

$$t_{0,\pi}(\chi) = \text{atan2} \left( \frac{C(\chi)}{2D(\chi)}, \pm \sqrt{1 - \left( \frac{C(\chi)}{2D(\chi)} \right)^2} \right). \quad (58)$$

Define the following useful identities:

$$s_*(\chi) := \frac{C(\chi)}{2D(\chi)}, \quad c_*(\chi) := \sqrt{1 - s_*(\chi)^2}. \quad (59)$$

$$\sin t_0 = \sin t_\pi = s_*, \quad \cos t_0 = +c_*, \quad \cos t_\pi = -c_*, \quad \cos^2 t_0 = \cos^2 t_\pi = c_*^2. \quad (60)$$

Evaluating Eq. (42) at these two branch representatives and applying some trigonometry,

$$\psi(\chi, t_0) = \Lambda(\chi) + B(\chi)c_*(\chi), \quad \psi(\chi, t_\pi) = \Lambda(\chi) - B(\chi)c_*(\chi), \quad (61)$$

where

$$\Lambda(\chi) := A(\chi) + D(\chi) + \frac{C(\chi)^2}{4D(\chi)}. \quad (62)$$

The “log-sum-exp” function (see e.g. Ref. 14) gives a surrogate

$$g_{\text{loc},\tau}(\chi) := \tau \log \left( e^{\psi(\chi, t_0)/\tau} + e^{\psi(\chi, t_\pi)/\tau} \right) \quad (63)$$

which smoothly recovers the true maximum of  $\psi(\chi, t)$  at  $t_0$  and  $t_\pi$  as  $\tau \rightarrow 0^+$ , and thus acts as a local proxy for  $\psi$  in the problematic orientation. Substituting Eq. (61) yields

$$g_{\text{loc},\tau}(\chi) = \Lambda(\chi) + \tau \log \left( 2 \cosh \frac{B(\chi)c_*(\chi)}{\tau} \right). \quad (64)$$

**Key Result 2:** For a single half-plane chance constraint  $g$ , in lieu of a discontinuous change in  $\nabla_\chi g$  between  $t_0$  and  $t_\pi$ , we seek a stable alternative. Differentiating Eq. (64) gives a smooth transition of the gradient in the vicinity of the problem orientation near  $|B(\chi)| = 0$ .

$$\nabla_\chi g_{\text{loc},\tau}(\chi) = \nabla_\chi \Lambda(\chi) + \tanh \left( \frac{B(\chi)c_*(\chi)}{\tau} \right) \nabla_\chi (B(\chi)c_*(\chi)). \quad (65)$$

This proxy can be implemented whenever  $|B(\chi)|$  is small, with judicious choice of smoothing parameter  $\tau$ . Away from the problematic orientation, the normal gradient can be used. All necessary gradients defined in Appendix B.

For small  $|C(\chi)/(2D(\chi))|$ , the two critical angles remain close to 0 and  $\pi$ , with

$$t_0(\chi) \approx \frac{C(\chi)}{2D(\chi)}, \quad t_\pi(\chi) \approx \pi - \frac{C(\chi)}{2D(\chi)}. \quad (66)$$

**II. Conservative Surrogate:** Previously we discussed solving for  $t^*$  maximizing  $\psi(\chi, t)$ , and enforcing  $\psi(\chi, t^*) \leq 0$  to satisfy the chance constraint, noting that this process does not produce smooth corrective updates  $\Delta\chi$  on problem free variables. We can obviate the need for computing  $t^*$  entirely via a more elegant but more complex approach, by defining a tunable smooth surrogate for the non-convex constraint. For a fixed value of  $\chi$ , define the maximum angular value of  $\psi$ :

$$M := \max_{t \in [0, 2\pi]} \psi(t). \quad (67)$$

The “log-integral-exp” surrogate, inspired by log-sum-exp, approaches  $M$  as  $\tau \rightarrow 0^+$ :

$$g_{\text{b},\tau}[\psi] := \tau \log \left( \frac{1}{2\pi} \int_0^{2\pi} e^{\psi(t)/\tau} dt \right), \quad \tau > 0, \quad (68)$$

where the “b” subscript denotes “below”. One can show that Eq. (68) always satisfies  $g_{\text{b},\tau}[\psi] \leq M$ . Thus,  $g_{\text{b},\tau}[\psi]$  alone is not a conservative replacement for Eq. (67).

We need a conservative smooth *upper bounding* function. Define the maximum perimetric derivative  $L$ , noting  $\psi$  is smooth and analytic, and  $(\cdot)' = \frac{d}{dt}(\cdot)$ :

$$L := \max_t |\psi'(t)|. \quad (69)$$

As shown in Appendix C, the following conservative upper bound holds:

$$\max_t \psi(t) \leq g_{b,\tau}[\psi] + \mathcal{C}(\tau, L), \quad (70)$$

where


$$\mathcal{C}(\tau, L) := -\tau \log\left(\frac{\tau}{2\pi L} \left(1 - e^{-2\pi L/\tau}\right)\right). \quad (71)$$

This term satisfies  $\mathcal{C}(\tau, L) \geq 0$  and furthermore,  $\mathcal{C}(\tau, L) \rightarrow 0$  as  $\tau \rightarrow 0^+$ .

The original constraint  $\psi(t^*) \leq 0$  is conservatively enforced by the following condition, which side-steps the need to solve for the worst violating angle  $t^*$ :

$$g_{b,\tau}[\psi] + \mathcal{C}(\tau, L) \leq 0. \quad (72)$$

Writing the conservative upper bound as  $g_{a,\tau}[\psi] := g_{b,\tau}[\psi] + \mathcal{C}(\tau, L)$ , where subscript “a” denotes “above”, we observe that enforcing  $g_{a,\tau}[\psi] \leq 0$  conservatively satisfies Eq. (35). Explicitly, Eq. (35) is replaced by a new function of the decision variables  $\chi$ . This is just a new  $g$  function in Eq. (51), for which we need to compute the necessary  $\nabla_{\chi} g$  terms. An example plot of  $g_{a,\tau}[\psi]$  is given in Fig. 2, illustrating the effect of tuning parameter  $\tau$ .



**Figure 2. Example Demonstration of Conservative Surrogate for Various Values of  $\tau$**

**Key Result 3:** The aforementioned tunable conservative chance constraint surrogate  $g_a$  is smooth in decision variables  $\chi$  for all practical purposes, as is its gradient:

$$\nabla_{\chi} g_{a,\tau}[\psi] \approx \nabla_{\chi} g_{b,\tau}[\psi], \quad \tau \rightarrow 0^+. \quad (73)$$

As discussed in Appendix C, the parameter  $L$  may be safely treated as a *fixed* conservative parameter, or updated occasionally. All necessary gradients are defined in Appendices B and C. When properly tuned, this approach offers strong guarantees of performance, but the equations for  $g_{a,\tau}$  and its gradient contain integrals which may become numerically burdensome to compute as  $\tau \rightarrow 0^+$ .

## NUMERICAL RESULTS

### Remark on Higher-Order Moment Computations

The non-Gaussian confidence contour parameterization given by Eq. (31) requires that we estimate the skew and kurtosis of the distribution. This can be accomplished in a number of ways of varying efficiency. Importantly, because the resulting confidence contour is used in an optimizer, we need a deterministic method for estimating the statistical moments. We use the Conjugate Unscented Transform (CUT),<sup>5</sup> which generally has significant speed advantages over random sampling-based

methods. The equations for implementation are provided in the Appendix for completeness. See Refs. 10, 15 for discussions of different techniques for estimating higher-order statistical moments. In this paper, all numerical results are obtained from Python scripts using open-source optimization tools like SLSQP from `scipy.optimize`, on a 2024 MacBook Pro with Apple M4 Max chip.

## I. Numerical Comparison of Different Methodologies: An Asteroid Orbiter Example

As a first test of the stochastic control techniques discussed previously, we repeat the example of Ref. 8 which was originally solved with a Gaussian mixture model. The spacecraft is initialized about an asteroid with gravitational parameter  $\mu = 5.2 \text{ m}^3/\text{s}^2$  with the state

$$\mathbf{x}_0 = [-1000 \ 0 \ 0 \ 0 \ 0 \ -7.211 \times 10^{-2}]^T$$

with position and velocity components in m, m/s respectively. At the initial time  $t_0$ , an impulsive control  $\mathbf{u}_0$  is selected to transfer the spacecraft toward a reconnaissance trajectory that descends closer to the asteroid surface, as illustrated in Fig. 3. The maneuver is designed subject to a terminal safety requirement at time  $t_1$ : despite uncertainty in the initial state, the propagated spacecraft distribution should not approach the asteroid more closely than the prescribed keep-out limit.

The initial uncertainty is taken to be Gaussian with diagonal covariance  $P_0 = \text{diag}(\sigma_0^2)$ , where

$$\sigma_0 = [1.0 \ 1.0 \ 1.0 \ \epsilon \ \epsilon \ \epsilon]^T \quad (\text{m, m/s}), \quad (74)$$

and  $\epsilon \ll 1$ . This corresponds to a position-dominated dispersion, which is appropriate for the present problem because the escape speed is only on the order of 0.1 m/s. The terminal time  $t_1$  is chosen to occur after 1.5 revolutions of the reconnaissance orbit following the maneuver, approximately 23.6 hours later. By this time, the initially Gaussian uncertainty has evolved into a visibly non-Gaussian distribution in Cartesian components. The scenario is depicted in Figure 3.




Figure 3. Asteroid maneuver targeting scenario

The terminal requirement is imposed as a rectangular keep-in box in position space, represented by six scalar chance constraints below. Each face is assigned the risk allocation,  $\Delta_1 = 0.01$ , so each corresponding constraint must be satisfied by at least 99% of the distribution.

$$p[x_1 \geq 495 \text{ m}] \geq 1 - \Delta_1 \quad (75a)$$

$$p[x_1 \leq 505 \text{ m}] \geq 1 - \Delta_1 \quad (75b)$$


$$p[y_1 \geq -80 \text{ m}] \geq 1 - \Delta_1 \quad (75c)$$

$$p[y_1 \leq 80 \text{ m}] \geq 1 - \Delta_1 \quad (75d)$$

$$p[z_1 \geq -25 \text{ m}] \geq 1 - \Delta_1 \quad (75e)$$

$$p[z_1 \leq 25 \text{ m}] \geq 1 - \Delta_1 \quad (75f)$$

Figure 4 shows the result of banana and LinCov policies in a 5000 run Monte Carlo simulation. The LinCov solution satisfies the keep-out constraint in 88.10% of samples, while the banana solutions improve the satisfaction rate to about 98.35%. Note each *individual* constraint is successfully satisfied to 99.0% probability. Both methods require similar control of  $\|u\| \approx 9.3 \times 10^{-2} \text{ m/s}$ .



**Figure 4. Monte Carlo Outcomes of LinCov vs. Banana Stochastic Control**

Table 1 gives the optimizer runtime for the asteroid targeting problem for all major implementation strategies discussed in this paper. Here  $g$  denotes a constraint-function callback and  $Dg$  denotes a constraint-derivative callback. Entries marked FD use finite-difference constraint derivatives, so derivative information is obtained by repeated calls to  $g$ . The analytic variants instead provide separate  $Dg$  callbacks. The timing column  $t_{UQ}$  denotes the portion of the optimizer runtime spent propagating the covariance or higher-order uncertainty description inside the constraint evaluations. The table shows that the different variants of the non-Gaussian stochastic maneuver design approach have similar runtimes despite different callback counts. The sampled contour method requires the most constraint function calls. The active-angle and integral variants reduce the number of constraint function calls by supplying analytic derivatives, but the derivative-enabled uncertainty construction is more expensive per evaluation. These effects largely offset one another. Thus, for this example, the active-angle and integral surrogate formulations improve the mathematical structure of the constraint enforcement, but they do not substantially reduce optimizer runtime relative to a contour sampling approach.

**Table 1. Runtime breakdown for asteroid targeting optimizations**

| Method          | Derivatives | Calls      | $t_{\text{opt}}$ [s] | $t_{\text{UQ}}$ [s] |
|-----------------|-------------|------------|----------------------|---------------------|
| LinCov          | FD          | $23g$      | 0.604                | 0.598               |
| Banana sampled  | FD          | $18g$      | 9.567                | 9.097               |
| Banana active   | analytic    | $6g + 3Dg$ | 9.018                | 8.377               |
| Banana integral | analytic    | $6g + 3Dg$ | 8.950                | 8.335               |

## II. A Full Stochastic Guidance Problem: Lunar Free-Return Midcourse Corrections

We now introduce an applied example motivated by the recent Artemis II mission, which marked humanity's first crewed visit to the Moon in over 50 years. The example is based on a simple planar circular restricted three-body problem (CR3BP, see Ref. 16) analog of the Artemis II free-return trajectory. Our purpose here is not to re-optimize the Artemis II ConOps, but rather to give a simple preliminary demonstration of how the proposed non-Gaussian chance-constraints might be applied in long-horizon spacecraft targeting. For greater realism, Reference 17 provides an actual NASA study of Artemis II stochastic guidance using a LinCov framework.


**Figure 5. Artemis II-like Planar Free Return**

Figure 5 shows the corresponding nominal planar free-return trajectory. Consider hypothetical midcourse correction maneuvers performed on the post-lunar-flyby return leg. The navigation state error and maneuver execution error effects are treated statistically. Let

$$\mathbf{x}_m^- = \hat{\mathbf{x}}_m^- + \delta \mathbf{x}_m^-, \quad \delta \mathbf{x}_m^- \sim \mathcal{N}(\mathbf{0}, P_m^-), \quad (76)$$

where  $\hat{\mathbf{x}}_m^-$  is the post-navigation estimate immediately before the correction maneuver. Let  $\Delta \bar{\mathbf{v}}_m$  denote the commanded correction maneuver, and model the executed maneuver as

$$\Delta \mathbf{v}_m = \Delta \bar{\mathbf{v}}_m + \mathbf{e}_{\Delta v}, \quad \mathbf{e}_{\Delta v} \sim \mathcal{N}(\mathbf{0}, Q_{\Delta v}). \quad (77)$$

For an impulsive maneuver,

$$\mathbf{x}_m^+ = \mathbf{x}_m^- + B_m \Delta \bar{\mathbf{v}}_m + B_m \mathbf{e}_{\Delta v}, \quad (78)$$

where  $B_m = [0_{3 \times 3}, I_{3 \times 3}]^\top$ . Therefore, the post-maneuver nominal state and covariance are

$$\hat{\mathbf{x}}_m^+ = \hat{\mathbf{x}}_m^- + B_m \Delta \bar{\mathbf{v}}_m, \quad (79)$$

$$P_m^+ = P_m^- + B_m Q_{\Delta v} B_m^\top. \quad (80)$$

The classical JPL Gates error model<sup>18</sup> provides a simplified description of maneuver execution error via distinct modalities of shutoff, resolution, pointing, and autopilot errors. The execution covariance is decomposed into components parallel and perpendicular to the commanded maneuver, based on the various error pathways and statistics. Defining  $V = \|\Delta \bar{\mathbf{v}}_m\|$  and  $\hat{\mathbf{e}} = \frac{\Delta \bar{\mathbf{v}}_m}{\|\Delta \bar{\mathbf{v}}_m\|}$ , we write

$$Q_{\Delta v} = (\sigma_r^2 + V^2 \sigma_s^2) \hat{\mathbf{e}} \hat{\mathbf{e}}^\top + (\sigma_a^2 + V^2 \sigma_p^2) (I - \hat{\mathbf{e}} \hat{\mathbf{e}}^\top). \quad (81)$$

Here  $\sigma_r$  is a fixed along-burn resolution error,  $\sigma_s$  is a fractional shutoff or scale error,  $\sigma_p$  is a pointing error in radians, and  $\sigma_a$  is a fixed transverse autopilot error. See Reference 18 for more information. Thus the cleanup command affects the stochastic problem both by shifting the nominal post-maneuver state and by changing the state covariance through the commanded nominal maneuver. The downstream uncertainty propagation is then posed as

$$\mathbf{x}_f = \varphi_{t_f, t_m}(\hat{\mathbf{x}}_m^+ + \delta \mathbf{x}_m^+), \quad \delta \mathbf{x}_m^+ \sim \mathcal{N}(\mathbf{0}, P_m^+), \quad (82)$$

where  $\varphi_{t_f, t_m}$  denotes the flow of the CR3BP dynamics from the cleanup epoch to the atmospheric entry interface. As these dynamics are nonlinear, the final state distribution at the atmospheric entry point will no longer be Gaussian. Critically, we emphasize the following point: For passive safety after the post-lunar-flyby correction maneuver, this maneuver must account for the expected final non-Gaussian uncertainty, even if navigation updates continue until atmospheric entry or until the next burn. Later navigation updates may reduce epistemic uncertainty in the estimated state, but they do not reduce the physical dispersion of outcomes induced by the preceding correction maneuver.

The entry constraints are in local entry coordinates. Let  $\mathbf{r}$  denote the Earth-relative position at entry and let  $\mathbf{v}$  denote the relative inertial velocity. Defining  $\hat{\mathbf{e}}_r = \mathbf{r}/\|\mathbf{r}\|$ , decompose the velocity:

$$v_R = \hat{\mathbf{e}}_r^\top \mathbf{v}, \quad v_T = \hat{\mathbf{e}}_t^\top \mathbf{v}, \quad (83)$$

where  $\hat{\mathbf{e}}_t$  is the local tangential direction. The entry flight-path angle satisfies  $\tan \gamma = v_R/v_T$ , assuming  $v_T > 0$ . Rather than constraining  $\gamma$  directly, the shallow- and steep-side entry constraints can be written as half-plane constraints in the local velocity plane. Table 2 gives a compact form of the entry corridor constraints. The two position constraints enforce an admissible entry-interface radial window, while the two velocity constraints enforce shallow and steep flight-path-angle limits.

We consider a hypothetical sequence of two correction maneuvers during mission operations to re-target the entry corridor, computing a policy with high statistical confidence of success. The maneuver error parameters are given in Table 3, with a fairly pessimistic proportional error of 3%. In this example, the nominal re-entry condition has entry flight path angle (EFPA) of  $\gamma = -6^\circ$  and altitude of  $h = 120$  km, but state errors after the lunar flyby are sufficient that trajectory correction is necessary. The first scheduled maneuver occurs shortly after lunar flyby, anticipating significant error and dispersion at the scheduled re-entry time due to navigation and maneuver errors. These error

**Table 2. Lunar Return Earth Arrival/Re-entry Constraints**

| Constraint type                    | Generic form                                | Local EDL corridor version                              |
|------------------------------------|---------------------------------------------|---------------------------------------------------------|
| Position half-plane ( $\times 2$ ) | $\mathbf{n}_r^\top \mathbf{r} - b_0 \leq 0$ | $\hat{\mathbf{e}}_r^\top \mathbf{r} - r_{\max} \leq 0$  |
|                                    |                                             | $-\hat{\mathbf{e}}_r^\top \mathbf{r} + r_{\min} \leq 0$ |
| Velocity half-plane ( $\times 2$ ) | $\mathbf{n}_v^\top \mathbf{v} - b_0 \leq 0$ | $v_R - \tan(\gamma_{\text{shallow}})v_T \leq 0$         |
|                                    |                                             | $-v_R + \tan(\gamma_{\text{steep}})v_T \leq 0$          |

parameters are given in Table 4. Out-of-plane errors are assumed subdominant for this planar example. In this context, due to significant expected dispersions, the final distribution is non-Gaussian, and exactly re-targeting the nominal re-entry condition is not *passively safe*: if a follow-up maneuver cannot be completed in time, the spacecraft is at risk of burning up in the Earth’s atmosphere from a too-steep entry condition. We seek instead a passively safe policy for the first maneuver, which corrects the expected dispersion to be closer to the nominal entry condition, while enforcing four three-sigma chance constraints

$$p[h > 150 \text{ km}] \geq \Phi(3) = 0.99865, \quad (84a)$$

$$p[h < 400 \text{ km}] \geq \Phi(3) = 0.99865, \quad (84b)$$

$$p[\gamma > -6^\circ] \geq \Phi(3) = 0.99865, \quad (84c)$$

$$p[\gamma < 8^\circ] \geq \Phi(3) = 0.99865. \quad (84d)$$

The lower constraints on  $h$  and  $\gamma$  are operationally critical. The upper constraints are enforced simply to prevent solutions that are “safe” but arbitrarily far from the nominal entry condition.

**Table 3. Maneuver Execution Error Parameters**


| Parameter  | Interpretation                                 | Value                    |
|------------|------------------------------------------------|--------------------------|
| $\sigma_s$ | Proportional magnitude/shutoff error           | $3.0 \times 10^{-2}$     |
| $\sigma_r$ | Fixed error along the commanded burn direction | 10.0 mm/s                |
| $\sigma_p$ | Pointing error standard deviation              | $3.0 \times 10^{-4}$ rad |
| $\sigma_a$ | Fixed transverse execution error               | 0.9 mm/s                 |

**Table 4. Artemis Lunar Return Error: Correction Burn 1**


| Quantity                  | Key                                                                  | Value                     |
|---------------------------|----------------------------------------------------------------------|---------------------------|
| Midcourse maneuver 1 time | $t_{m_1}$                                                            | 0.25 days post-flyby      |
| Time to entry             | $t_f - t_{m_1}$                                                      | 3.6163 days               |
| Mean cleanup state error  | $\delta \mathbf{x}_c = (\delta R, \delta T, \delta V_R, \delta V_T)$ | $(-0.5, 50.0, -0.1, 5.0)$ |
|                           | Units                                                                | [km, km, m/s, m/s]        |
| Nav. position uncertainty | $(\sigma_R, \sigma_T, \sigma_N)$                                     | (500.0, 5000.0, 5.0) m    |
| Nav. velocity uncertainty | $(\sigma_{V_R}, \sigma_{V_T}, \sigma_{V_N})$                         | (50.0, 500.0, 0.5) mm/s   |

A naive maneuver, directly targeting a desired  $\gamma = -6^\circ$  and  $h = 120$  km, requires 6.485 m/s and places the nominal state directly on the lower-altitude/steep-entry boundary, which is not passively safe. In a 5000 sample Monte Carlo check, this violates the constraints in 62.86% of samples,

dominated by unacceptable low altitude and steep flight-path angles at the desired arrival time. The LinCov policy requires 3.352 m/s and lowers the Monte Carlo violation percentage to 0.56%, but residual violations remain because the true final dispersion is not well-accounted for by the LinCov ellipse. See Figs. 7 and 8 for both LinCov and banana policy outcomes. The dispersion plots for the latter also show naive linear and nonlinear covariance ellipse predictions along with the banana contour for the same computed statistics. The banana policy, which is warm-started with the LinCov solution, reduces the maneuver from the LinCov solution to 2.867 m/s while lowering the violation percentage to 0.2%. The resulting trajectory is given in Fig. 6 with the final  $3\sigma$  position bounds at the final time. Relative to LinCov, the banana policy saves 0.484 m/s, or 14.4%, and reduces the observed violation fraction by approximately 64%. We note that the banana-shaped confidence boundary fits the final Monte Carlo samples extremely well, even though it only makes use of LinCov-computed covariance and skew/kurtosis components from CUT4. The method does not use Monte Carlo sampling at all. The solve times in a Python script are 0.2s for naive retargeting, 0.25s for LinCov, and 4.29s using the banana solver. While more numerically expensive than LinCov, the approach has  $\sim 84\%$  lower runtime than the Monte Carlo study, as 5000 samples took 27.6 seconds in this example.




**Figure 6. Post-flyby stochastic maneuver planning with banana policy: Correction Burn 1**



**Figure 7. Arrival position constraint satisfaction with Monte Carlo**

We now move on to the second corrective maneuver in our two-burn study. Key parameters given




**Figure 8. Arrival velocity constraint satisfaction with Monte Carlo**


in Table 5, with the mean state error reported with respect to the nominal outcome of maneuver 1. Out-of-plane errors are kept small for this planar example. The proportional maneuver execution error  $\sigma_s$  is relaxed from its prior pessimistic stress-test value to  $5 \times 10^{-3}$ . All other parameters in Table 3 are unchanged. The navigation estimate at maneuver 2 time is chosen as a representative outcome from maneuver 1, close to the boundary of  $1\sigma$  dispersed outcomes from maneuver 1.

**Table 5. Artemis Lunar Return Error: Correction Burn 2**

| Quantity                  | Key                                                                  | Value                             |
|---------------------------|----------------------------------------------------------------------|-----------------------------------|
| Midcourse maneuver 2 time | $t_{m_2}$                                                            | 5 hr before Earth entry interface |
| Time to entry             | $t_f - t_{m_2}$                                                      | 0.2083 days                       |
| Mean cleanup state error  | $\delta \mathbf{x}_c = (\delta R, \delta T, \delta V_R, \delta V_T)$ | $(-3.50, -40.78, -0.0094, 0.500)$ |
|                           | Units                                                                | [km, km, m/s, m/s]                |
| Nav. position uncertainty | $(\sigma_R, \sigma_T, \sigma_N)$                                     | (333.3, 1333.3, 33.3) m           |
| Nav. velocity uncertainty | $(\sigma_{V_R}, \sigma_{V_T}, \sigma_{V_N})$                         | (16.7, 83.3, 1.67) mm/s           |



**Figure 9. Pre-entry stochastic maneuver planning with banana policy: Correction Burn 2**



**Figure 10. Arrival constraint satisfaction with Monte Carlo (banana policy)**

For the second correction maneuver, the final entry corridor is enforced, again in one-sided inequalities at  $3\sigma$  confidence, as the union of altitude range  $118 < h < 132$  km and EFPA range  $-6.2^\circ < \gamma < -5.9^\circ$ . Entry longitude is allowed to vary in this example. The resulting trajectory is given in Fig. 9. As this maneuver executes with significantly lower navigation error and less propagation time, the final dispersion is much closer to Gaussian, so the LinCov and banana policies are essentially the same, with delta-V of 61.37 m/s, and Monte Carlo violation percentage of 0.14%. The final position and velocity confidence bounds and dispersions obtained with the banana policy are provided in Figure 10. The bounds are accurate and respect the desired corridor with some margin. This second maneuver illustrates that the banana policy converges to a LinCov policy for Gaussian dispersions.

## CONCLUSIONS

This paper derives a non-Gaussian chance-constraint approach for stochastic nonlinear spacecraft targeting problems. The method employs a moment-informed “banana” contour, using covariance for scale, and skew and kurtosis to capture the dominant bending and asymmetry of propagated uncertainty distributions. This enables direct geometric enforcement of half-plane constraints, including active-angle and smooth surrogate treatments of the resulting non-convex support problem. The ideas were demonstrated in two different numerical examples. The first was a simple asteroid orbiter control scenario, reproduced from prior literature. The second, more complex example revisited Artemis II stochastic guidance under the influence of navigation and maneuver execution errors and circular restricted three-body dynamics. In both tests, the banana method substantially improved Monte Carlo constraint satisfaction over a linear covariance (LinCov) approach, while remaining computationally tractable at roughly an order-of-magnitude longer runtime than LinCov. This is a useful middle ground between LinCov and Monte Carlo methods in stochastic maneuver design, particularly for runtime-constrained applications, such as rapid trade studies or onboard guidance. This should be used when non-Gaussian distributions challenge the validity of a classical LinCov approach, but the degree of their non-Gaussianity is not extreme (i.e. the distribution is neither multi-modal nor exhibiting complex unmodeled geometry).

## APPENDIX A: CONJUGATE UNSCENTED TRANSFORM

The equations below summarize CUT4,<sup>5</sup> which estimates up to 4<sup>th</sup>-order moments given an  $N$ -dimensional state with mean  $\boldsymbol{\mu}$  and covariance  $P$ , subject to some (assumed nonlinear) process  $\mathbf{g}$ :

$$P = SS^\top, \quad \mathcal{X}_i = \boldsymbol{\mu} + S\mathbf{z}_i, \quad i = 0, \dots, K-1, \quad K = 1 + 2N + 2^N. \quad (\text{A.1a})$$

$$\mathbf{z}_0 = \mathbf{0}, \quad w_0 = 1 - 2Nw_1 - 2^N w_2 \quad (\text{can freely choose } w_0 = 0 \text{ if } N > 2) \quad (\text{A.1b})$$

$$\mathbf{z}_k^\pm = \pm r_1 \mathbf{e}_k, \quad k = 1, \dots, N, \quad (\text{uses weight } w_1) \quad (\text{A.1c})$$

$$\tilde{\mathbf{z}}_k^\pm = \pm r_2 \boldsymbol{\zeta}_{N,k}, \quad \|\boldsymbol{\zeta}_{N,k}\| = \sqrt{N}, \quad k = 1, \dots, 2^N, \quad (\text{uses weight } w_2) \quad (\text{A.1d})$$

$$r_1 = \sqrt{\frac{N+2}{2}}, \quad r_2 = \sqrt{\frac{N+2}{N-2}}, \quad w_1 = \frac{4}{(N+2)^2}, \quad w_2 = \frac{(N-2)^2}{2^N (N+2)^2} \quad (\text{A.1e})$$

$$w_i = \begin{cases} w_1, & 1 \leq i \leq 2N \\ w_2, & 2N+1 \leq i \leq K \end{cases} \quad (\text{A.1f})$$

$$\mathcal{Y}_i = \mathbf{g}(\mathcal{X}_i), \quad \boldsymbol{\mu}_y = \sum_{i=1}^K w_i \mathcal{Y}_i \quad (\text{A.1g})$$

$$P_y = \sum_{i=1}^K w_i (\mathcal{Y}_i - \boldsymbol{\mu}_y)(\mathcal{Y}_i - \boldsymbol{\mu}_y)^\top \quad (\text{A.1h})$$

$$M_{abc}^{(3)} = \sum_i w_i (\mathcal{Y}_{i,a} - \mu_a)(\mathcal{Y}_{i,b} - \mu_b)(\mathcal{Y}_{i,c} - \mu_c) \quad (\text{A.1i})$$

$$M_{abcd}^{(4)} = \sum_i w_i (\mathcal{Y}_{i,a} - \mu_a)(\mathcal{Y}_{i,b} - \mu_b)(\mathcal{Y}_{i,c} - \mu_c)(\mathcal{Y}_{i,d} - \mu_d) \quad (\text{A.1j})$$

Unlike an unscented transform (“UT”), which samples twice along each coordinate direction as  $\pm \mathbf{e}_k$ , CUT4 samples additionally along “conjugate” axes  $\boldsymbol{\zeta}_{N,k}$  which are weighted composite directions of the original basis. In general, for an  $N$ -dimensional state, this includes the directions formed by assigning each coordinate either “+” or “−”, resulting in  $2^N$  extra off-axis samples in addition to the usual  $2N$  axis-aligned ones. For example, in three dimensions these additional directions correspond to the eight diagonals of a cube whose faces are normal to the three basis vectors and the six principal directions along these vectors. For  $N = 6$ , to recover the first four statistical moments, CUT4 requires  $2N + 2^N + 1 = 77$  sigma points (or, by the free choice  $w_0 = 0$ , 76 points). CUT4 improves on the error properties of UT: For a weakly nonlinear transformation  $\mathbf{g}(\mathbf{X})$ , the expected covariance error is  $\mathcal{O}(\|P\|^{5/2})$ . For full details on the moment propagation schemes, please see References 5, 19.

## APPENDIX B: PARTIAL DERIVATIVES FOR CHANCE CONSTRAINT CORRECTIONS

For each control component, define

$$\partial_j \equiv \frac{\partial}{\partial \chi_j}, \quad j = 1, 2, 3. \quad (\text{A.2})$$

The active support-gradient is

$$\nabla_{\mathbf{X}} \psi(\mathbf{X}, t^*) = \nabla_{\mathbf{X}} A + \cos t^* \nabla_{\mathbf{X}} B + \sin t^* \nabla_{\mathbf{X}} C + \cos^2 t^* \nabla_{\mathbf{X}} D. \quad (\text{A.3})$$

The required Jacobian row is obtained from the componentwise partials of  $A, B, C, D$ :

$$\partial_j A = \mathbf{n}^\top \partial_j \boldsymbol{\mu} - \left[ (\partial_j m_2) \alpha \sqrt{\lambda_2} + m_2 (\partial_j \alpha) \sqrt{\lambda_2} + m_2 \alpha \partial_j \sqrt{\lambda_2} \right], \quad (\text{A.4})$$

$$\partial_j B = k \left[ (\partial_j m_1) \sqrt{\lambda_1} + m_1 \partial_j \sqrt{\lambda_1} \right], \quad (\text{A.5})$$

$$\partial_j C = k \left[ (\partial_j m_2) \sqrt{\lambda_2} + m_2 \partial_j \sqrt{\lambda_2} \right], \quad (\text{A.6})$$

and

$$\begin{aligned} \partial_j D = & (\partial_j m_1) c \sqrt{\lambda_1} + m_1 (\partial_j c) \sqrt{\lambda_1} + m_1 c \partial_j \sqrt{\lambda_1} \\ & + (\partial_j m_2) \alpha k^2 \sqrt{\lambda_2} + m_2 (\partial_j \alpha) k^2 \sqrt{\lambda_2} + m_2 \alpha k^2 \partial_j \sqrt{\lambda_2}. \end{aligned} \quad (\text{A.7})$$

Equations (A.4)–(A.7) are the components of  $\nabla_{\mathbf{x}} A$ ,  $\nabla_{\mathbf{x}} B$ ,  $\nabla_{\mathbf{x}} C$ , and  $\nabla_{\mathbf{x}} D$ .

The constituent partials are as follows:

$$\partial_j \mathbf{m} = (\partial_j R)^\top \mathbf{n}. \quad (\text{A.8})$$

$$\partial_j \boldsymbol{\mu} = S \boldsymbol{\Phi}_0 B_{\mathbf{x},j} + S \partial_j \boldsymbol{\mu}_Y, \quad (\text{A.9})$$

where  $B_{\mathbf{x},j}$  is column  $j$  of

$$B_{\mathbf{x}} = \frac{\partial \bar{\mathbf{X}}_0}{\partial \boldsymbol{\chi}} = \frac{1}{V} \begin{bmatrix} \mathbf{0}_{3 \times 3} \\ I_3 \end{bmatrix}, \quad V = \sqrt{\frac{\mu}{L}}. \quad (\text{A.10})$$

The eigenvalue partials associated with  $\Sigma = R \Lambda R^\top$  are

$$\partial_j \lambda_i = \mathbf{e}_i^\top (\partial_j \Sigma) \mathbf{e}_i, \quad \partial_j \sqrt{\lambda_i} = \frac{\partial_j \lambda_i}{2 \sqrt{\lambda_i}}. \quad (\text{A.11})$$

For distinct eigenvalues,

$$\partial_j \mathbf{e}_i = \sum_{\ell \neq i} \mathbf{e}_\ell \frac{\mathbf{e}_\ell^\top (\partial_j \Sigma) \mathbf{e}_i}{\lambda_i - \lambda_\ell}, \quad \partial_j R = [\partial_j \mathbf{e}_1 \quad \partial_j \mathbf{e}_2]. \quad (\text{A.12})$$

In a code implementation one can, if desired, use a tuning parameter  $\rho$  to interpolate between the theoretical LinCov-predicted covariance as the baseline, or the CUT-predicted. Using the latter produces a double-counting phenomenon which adds a natural “buffer” to the banana prediction. While not completely rigorous, it could be a useful and natural tuning parameter in some circumstances. The covariance partial in this case is thus

$$\partial_j \Sigma = (1 - \rho) \partial_j \Sigma_{\text{lin}} + \rho \partial_j \Sigma_{\text{CUT}}, \quad (\text{A.13})$$

with

$$\partial_j \Sigma_{\text{lin}} = S (\partial_j P_{\text{lin}}) S^\top, \quad \partial_j \Sigma_{\text{CUT}} = S (\partial_j P_Y) S^\top. \quad (\text{A.14})$$

The LinCov covariance partial is

$$\partial_j P_{\text{lin}} = \boldsymbol{\Psi}_j P_0 \boldsymbol{\Phi}_0^\top + \boldsymbol{\Phi}_0 P_0 \boldsymbol{\Psi}_j^\top, \quad \boldsymbol{\Psi}_j = \partial_j \boldsymbol{\Phi}_0. \quad (\text{A.15})$$

In the implementation used here,  $\rho = 0$ , so the covariance scaling of the banana is LinCov-based, while the skewness and kurtosis corrections are CUT-based.

For CUT point  $i$ ,

$$\partial_j \mathbf{Y}_i = (\boldsymbol{\Phi}_i - \boldsymbol{\Phi}_0) B_{\chi,j}. \quad (\text{A.16})$$

Thus, with CUT weights  $w_i$ ,

$$\partial_j \boldsymbol{\mu}_Y = \sum_i w_i \partial_j \mathbf{Y}_i, \quad \partial_j \mathbf{Z}_i = \partial_j \mathbf{Y}_i - \partial_j \boldsymbol{\mu}_Y. \quad (\text{A.17})$$

The CUT covariance and moment partials are

$$\partial_j P_Y = \sum_i w_i \left[ (\partial_j \mathbf{Z}_i) \mathbf{Z}_i^\top + \mathbf{Z}_i (\partial_j \mathbf{Z}_i)^\top \right], \quad (\text{A.18})$$

$$\partial_j M^{(3)} = \sum_i w_i [\partial_j \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \mathbf{Z}_i + \mathbf{Z}_i \otimes \partial_j \mathbf{Z}_i \otimes \mathbf{Z}_i + \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \partial_j \mathbf{Z}_i], \quad (\text{A.19})$$

and

$$\begin{aligned} \partial_j M^{(4)} = \sum_i w_i & [\partial_j \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \mathbf{Z}_i + \mathbf{Z}_i \otimes \partial_j \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \mathbf{Z}_i \\ & + \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \partial_j \mathbf{Z}_i \otimes \mathbf{Z}_i + \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \mathbf{Z}_i \otimes \partial_j \mathbf{Z}_i]. \end{aligned} \quad (\text{A.20})$$

Let  $W = \Lambda^{-1/2} R^\top$ , with first and second rows  $\mathbf{a}^\top$  and  $\mathbf{b}^\top$ . Then

$$\partial_j W = \partial_j (\Lambda^{-1/2}) R^\top + \Lambda^{-1/2} (\partial_j R)^\top, \quad (\text{A.21})$$

which gives  $\partial_j \mathbf{a}$  and  $\partial_j \mathbf{b}$ .

Using the contraction notation e.g.

$$M^{(3)}[\mathbf{p}, \mathbf{q}, \mathbf{r}] = \sum_{a,b,c} p_a q_b r_c M_{abc}^{(3)}, \quad (\text{A.22})$$

the whitened-moment partials are

$$\partial_j E_{uuu} = 3M^{(3)}[\partial_j \mathbf{a}, \mathbf{a}, \mathbf{a}] + (\partial_j M^{(3)})[\mathbf{a}, \mathbf{a}, \mathbf{a}], \quad (\text{A.23})$$

$$\partial_j E_{uuv} = M^{(3)}[\partial_j \mathbf{b}, \mathbf{a}, \mathbf{a}] + 2M^{(3)}[\mathbf{b}, \partial_j \mathbf{a}, \mathbf{a}] + (\partial_j M^{(3)})[\mathbf{b}, \mathbf{a}, \mathbf{a}], \quad (\text{A.24})$$

and

$$\partial_j E_{uuuu} = 4M^{(4)}[\partial_j \mathbf{a}, \mathbf{a}, \mathbf{a}, \mathbf{a}] + (\partial_j M^{(4)})[\mathbf{a}, \mathbf{a}, \mathbf{a}, \mathbf{a}]. \quad (\text{A.25})$$

Finally, the two banana-correction partials appearing in Eqs. (A.4) and (A.7) are

$$\partial_j c = \frac{k^2 - 1}{6} \partial_j E_{uuu}, \quad (\text{A.26})$$

and

$$\partial_j \alpha = \frac{(\partial_j E_{uuv})(E_{uuuu} - 1) - E_{uuv}(\partial_j E_{uuuu})}{(E_{uuuu} - 1)^2}. \quad (\text{A.27})$$

Substituting Eqs. (A.4)–(A.7) into Eq. (A.3) gives the Jacobian row associated with one half-plane support constraint.

## APPENDIX C: NOTES ON CONSERVATIVE SURROGATE

For a maximizer  $t^* \in [0, 2\pi]$ ,  $\psi(t^*) = M$ . We make no requirement that there is a single unique maximizing angle  $t^*$ , only that a maximum value  $M$  exists. By the derivative bound (69), the following holds due to a Lipschitz condition:

$$\psi(t) \geq M - L|t - t^*|, \quad t \in [0, 2\pi]. \quad (\text{A.28})$$

Exponentiation and integration of (A.28) gives

$$\frac{1}{2\pi} \int_0^{2\pi} e^{\psi(t)/\tau} dt \geq \frac{e^{M/\tau}}{2\pi} \int_0^{2\pi} e^{-L|t-t^*|/\tau} dt. \quad (\text{A.29})$$

Define the useful quantity  $I(t^*) := \int_0^{2\pi} e^{-L|t-t^*|/\tau} dt$  and integrate:

$$I(t^*) = \frac{\tau}{L} \left( 2 - e^{-Lt^*/\tau} - e^{-L(2\pi-t^*)/\tau} \right). \quad (\text{A.30})$$

Eq. (A.30) is minimized when  $t^*$  is at an endpoint, namely  $t^* = 0$  or  $t^* = 2\pi$ . Consequently,

$$I(t^*) \geq \frac{\tau}{L} \left( 1 - e^{-2\pi L/\tau} \right). \quad (\text{A.31})$$

Substituting (A.31) into (A.29) gives

$$\frac{1}{2\pi} \int_0^{2\pi} e^{\psi(t)/\tau} dt \geq e^{M/\tau} \frac{\tau}{2\pi L} \left( 1 - e^{-2\pi L/\tau} \right). \quad (\text{A.32})$$

Taking  $\tau \log(\cdot)$  of both sides of (A.32) and rearranging gives the conservative upper bound

$$M \leq g_{b,\tau}[\psi] - \tau \log \left( \frac{\tau}{2\pi L} \left( 1 - e^{-2\pi L/\tau} \right) \right). \quad (\text{A.33})$$

Using the definition of  $\mathcal{C}(\tau, L)$  in Eq. (71), Eq. (A.33) becomes

$$\max_t \psi(t) \leq g_{b,\tau}[\psi] + \mathcal{C}(\tau, L). \quad (\text{A.34})$$

We must first note that the lack of smooth differentiability of Eq. (35) is traded for that of Eq. (69). However, this new discontinuity is much less severe, which we now show. Noting that in practice a very small  $\tau$  should be chosen to minimize conservatism, we examine the behavior of Eq. (71) as  $\tau \rightarrow 0^+$ :

$$\mathcal{C}(\tau, L) \sim \tau \log \left( \frac{2\pi L}{\tau} \right), \quad \tau \rightarrow 0^+. \quad (\text{A.35})$$

This observation is significant for two reasons. First, there is extremely weak dependence on  $L$ , so it can be over-estimated conservatively without ill effect. To see this, let  $L_2 = \rho L_1$  for some factor  $\rho \gg 1$ . Then Eq. (A.35) gives

$$\mathcal{C}(\tau, L_2) - \mathcal{C}(\tau, L_1) \sim \tau \log \left( \frac{2\pi \rho L_1}{\tau} \right) - \tau \log \left( \frac{2\pi L_1}{\tau} \right) = \tau \log \rho. \quad (\text{A.36})$$

Hence even large changes in  $L$  modify  $\mathcal{C}$  only by an amount proportional to  $\tau$ , where  $\tau \ll 1$ . Numerical experiments show that even multiple order-of-magnitude overestimates of  $L$  are inconsequential to the value of  $\mathcal{C}$  when  $\tau$  is chosen sufficiently small.

Second, differentiating Eq. (A.35) with respect to  $L$  gives

$$\frac{\partial \mathcal{C}}{\partial L} \sim \frac{\tau}{L}, \quad \tau \rightarrow 0^+. \quad (\text{A.37})$$

We remind the reader that  $\tau$  is chosen as small as possible, and  $L$  may be freely replaced with a highly conservative overestimate. Given the relative unimportance of  $L$  for small  $\tau$ , we adopt a practice whereby  $L$  may be treated as a *fixed* or sequentially updated conservative bound, in which case the resulting surrogate is smooth in  $\chi$  for the purposes of the corrective step. In other words:

$$\nabla_{\chi} g_{a,\tau}[\psi] \approx \nabla_{\chi} g_{b,\tau}[\psi], \quad \tau \rightarrow 0^+. \quad (\text{A.38})$$

To compute the gradients, the partial derivative of the log-integral term follows as:

$$\partial_j g_{b,\tau}[\psi] = \frac{\int_0^{2\pi} e^{\psi(t)/\tau} \partial_j \psi(t) dt}{\int_0^{2\pi} e^{\psi(t)/\tau} dt}. \quad (\text{A.39})$$

Here

$$\partial_j \psi(\chi, t) = \partial_j A + \cos t \partial_j B + \sin t \partial_j C + \cos^2 t \partial_j D, \quad (\text{A.40})$$

using the coefficient partials collected in Appendix B. In our implementation,  $L$  is treated as a fixed or sequentially updated conservative bound, so  $\partial_j \mathcal{C}(\tau, L)$  is neglected,  $\partial_j g_{a,\tau} \approx \partial_j g_{b,\tau}$ .

## REFERENCES

- [1] A. W. Berning Jr., E. R. Burnett, and S. Bieniawski, “Chance-Constrained, Drift-Safe Guidance for Spacecraft Rendezvous,” *AAS Rocky Mountain Guidance, Navigation, and Control Conference*, American Astronautical Society, 2023, 10.48550/arXiv.2401.11077.
- [2] K. Oguri, “Chance-Constrained Control for Safe Spacecraft Autonomy: Convex Programming Approach,” *2024 American Control Conference (ACC)*, 2024, pp. 2318–2324, 10.48550/arXiv.2403.04062.
- [3] J. L. Junkins, “Adventures on the Interface of Dynamics and Control,” *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 20, Nov.–Dec. 1997, pp. 1058–1071, 10.2514/2.4176.
- [4] B. A. Jones, A. Doostan, and G. H. Born, “Nonlinear Propagation of Orbit Uncertainty Using Non-Intrusive Polynomial Chaos,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 2, 2013, pp. 430–444, 10.2514/1.57599.
- [5] N. Adurthi, P. Singla, and T. Singh, “The Conjugate Unscented Transform - An approach to evaluate multi-dimensional expectation integrals,” *2012 American Control Conference (ACC)*, 2012, pp. 5556–5561, 10.1109/ACC.2012.6314970.
- [6] D. C. Qi, K. Oguri, P. Singla, and M. R. Akella, “Non-Gaussian Distribution Steering in Nonlinear Dynamics with Conjugate Unscented Transformation,” 2025, 10.48550/arXiv.2510.12946.
- [7] T. Sial and A. Halder, “Fixed Horizon Linear Quadratic Covariance Steering in Continuous Time with Hilbert-Schmidt Terminal Cost,” 2025, 10.48550/arXiv.2510.21944.
- [8] S. Boone and J. W. McMahon, “Non-Gaussian Chance-Constrained Trajectory Control Using Gaussian Mixtures and Risk Allocation,” *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, 10.1109/CDC51059.2022.9993274.
- [9] E. Burnett and S. Boone, “Analytic Non-Gaussian Confidence Boundary Method for Chance-Constrained Trajectory Control,” arXiv, 2026, 10.48550/arXiv.2604.04304.
- [10] N. Michelotti, E. R. Burnett, and F. Topputo, “Analytical Confidence Boundaries for Non-Gaussian Uncertainty in Perturbed Spacecraft Dynamics,” *AAS/AIAA Astrodynamics Specialist Conference*, No. AAS 26-909, American Astronautical Society, July 2026.
- [11] S. Boone and J. McMahon, “Directional State Transition Tensors for Capturing Dominant Nonlinear Effects in Orbital Dynamics,” *Journal of Guidance, Control, and Dynamics*, Vol. 46, March 2023, pp. 431–442, 10.2514/1.G006910.
- [12] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover Publications, Inc., 1965.

- [13] K. Oguri, M. Ono, and J. W. McMahon, “Convex Optimization over Sequential Linear Feedback Policies with Continuous-time Chance Constraints,” *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 6325–6331, 10.1109/CDC40024.2019.9029604.
- [14] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York: Cambridge University Press, 2004.
- [15] E. R. Burnett and S. Boone, “Efficient Nonlinear Uncertainty Quantification for Spaceflight Leveraging Nonlinear Expansions,” *International Symposium on Space Flight Dynamics*, 2026, 10.48550/arXiv.2605.24147.
- [16] W. S. Koon, M. W. Lo, J. E. Marsden, and S. D. Ross, *Dynamical Systems, The Three-Body Problem, and Space Mission Design*. New York: Springer, 2017.
- [17] D. Woffinden, R. Eckman, and S. Robinson, “Optimized Trajectory Correction Burn Placement for the NASA Artemis II Mission,” *AAS/AIAA Spaceflight Mechanics Meeting*, No. AAS 23-062, Austin, TX, January 2023.
- [18] C. R. Gates, “A Simplified Model of Midcourse Maneuver Execution Errors,” No. JPL-TR-32-504, NASA-CR-53032, 1963.
- [19] E. Wan and R. Van Der Merwe, “The unscented Kalman filter for nonlinear estimation,” *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, 2000, pp. 153–158, 10.1109/ASSPCC.2000.882463.