(Preprint) AAS XX-XXX

# LOW-THRUST MINIMUM-FUEL OPTIMIZATION WITH CHANCE CONSTRAINT AND CVAR PENALTIES: AN INDIRECT APPROACH

Carmine Giordano* and Francesco Topputo†

This paper addresses low-thrust minimum-fuel trajectory optimization under uncertainty. The spacecraft dynamics are modeled as a random ordinary differential equation, where thrust magnitude and pointing errors are represented through Gauss–Markov processes. A risk-aware stochastic optimal control problem is formulated by combining fuel minimization with a CVaR penalty on the terminal-position error. The problem is then treated through Pontryagin's Maximum Principle, deriving the Hamiltonian structure, adjoint equations, and optimality conditions. The method is finally applied to interplanetary transfers, showing its suitability for robust low-thrust trajectory design for limited-capability deep-space spacecraft.

## INTRODUCTION

Spacecraft trajectory design has historically been driven primarily by mission objectives and engineering constraints, under the assumption that chemical propulsion provides significant control authority. Because conventional propulsion systems can generate relatively large thrust levels, they have long allowed spacecraft to compensate for off-nominal behavior through short corrective maneuvers. As a result, nominal trajectories are typically optimized first, while uncertainties are treated only afterward. In this traditional workflow, sources of uncertainty, including imperfect dynamical models, such as gravitational mismodeling or poorly predicted solar radiation pressure, navigation errors, and command execution inaccuracies,[1] are not embedded directly in the design phase. Instead, once a nominal solution has been obtained, a subsequent analysis is carried out to evaluate its robustness, estimate the size and frequency of required correction maneuvers, and assess the attainable accuracy of the spacecraft state. For large spacecraft, this sequential strategy is generally acceptable, since high-thrust propulsion and relatively generous propellant margins can absorb the resulting inefficiencies.

In recent years, however, the growing interest in miniaturized spacecraft, such as SmallSats and CubeSats, has challenged this paradigm. These platforms are attractive because they enable relatively inexpensive missions while still offering valuable scientific and technological return.[2,3] Their limited size, though, imposes severe restrictions on trajectory control. Propulsion systems are often characterized by very low thrust levels, and available propellant is typically scarce. Consequently, these spacecraft cannot tolerate the same level of inefficiency as larger missions. In addition, they frequently operate in conditions of greater uncertainty, due to less accurate navigation capabilities and the performance limits of miniaturized subsystems, which tend to increase both state-estimation

*Assistant Professor, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, 20156, Milan (Italy).
†Full Professor, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, 20156, Milan (Italy).

and control-execution errors.

Under these circumstances, the classical separation between nominal trajectory optimization and robustness assessment becomes inadequate. For small spacecraft, correction maneuvers are no longer marginal refinements, but mission-critical actions whose cost may substantially affect feasibility. Trajectory design must therefore account for uncertainty from the beginning, so that the resulting solution is not only optimal in a nominal sense, but also practically achievable and sufficiently robust in the presence of disturbances and imperfect knowledge.

Motivated by these challenges, a number of stochastic optimal control and robust trajectory design methods have been proposed over the last decade. An early contribution is the statistical targeting algorithm introduced in Park et al.,[4] which incorporates probabilistic information directly into the targeting process. In contrast with classical targeting methods, which solve a deterministic boundary-value problem for a nominal trajectory, this formulation seeks a trajectory that reaches the desired target in a statistical sense. Its applicability, however, is limited when the dispersion of the stochastic trajectories cannot be represented adequately through an approximately Gaussian description.

The uncertain Lambert problem has also received considerable attention. Differential algebra has been employed to address this problem,[5–7] while Schumacher et al.[8] proposed a formulation based on first-order variational equations to characterize the effect of uncertainty. This line of work was later extended by explicitly deriving the partial derivatives of the transfer velocities,[9] and subsequently by introducing a derivative-free numerical approach that leverages recent advances in uncertainty quantification.[10] Differential algebra was also used in the context of gravity-assist trajectory design, specifically within the space-pruning algorithm.[11]

Related developments can be found in rendezvous trajectory design. In Li et al.,[12] a multi-objective optimization framework was introduced for linear rendezvous problems, with a robustness metric defined in terms of the final uncertainty and accounting for both navigation and control errors. For short-duration missions, this study also highlighted the relationship among rendezvous time, propellant expenditure, and robustness. Later work extended these ideas to nonlinear rendezvous dynamics and longer mission phases.[13] Stochastic asteroid rendezvous problems have also been investigated by considering uncertainty in both the spacecraft and target states, together with the optimization of correction maneuvers under Lambert-type transfer assumptions.[14]

More recently, broad methodologies for trajectory optimization under uncertainty have been developed. In Xiong et al.,[15] stochastic trajectory optimization was reformulated as a deterministic problem by combining Polynomial Chaos Expansion with an adaptive pseudospectral collocation scheme. A different direction was taken in Greco et al.,[16] where a Belief Markov Decision Process framework was proposed and then applied to the robust design of a Europa Clipper flyby trajectory subject to execution, observation, and knowledge uncertainties. Stochastic Differential Dynamic Programming has also been explored for robust Earth–Mars transfers, with the unscented transform adopted for uncertainty propagation;[17] this methodology was later improved through a hybrid multiple-shooting strategy to mitigate some limitations of Differential Dynamic Programming.[18] Convex optimization combined with covariance steering has further enabled the treatment of several classes of robust continuous-control problems in astrodynamics, including interplanetary transfers[19] and atmospheric reentry.[20]

A more comprehensive framework for trajectory design under uncertainty was recently introduced in,[21] with the goal of systematically incorporating uncertainties in dynamics, navigation, and control. Despite the broad literature on robust and stochastic trajectory design, a methodological gap remains between uncertainty-aware trajectory formulations and indirect optimal-control theory. Most available approaches rely on direct optimization, uncertainty propagation, or covariance-based approximations, whereas the use of Pontryagin's Maximum Principle (PMP) to analyze and solve stochastic low-thrust trajectory optimization problems remains comparatively unexplored. This is particularly relevant when robustness is expressed through risk-sensitive metrics, such as quantile- or CVaR-based criteria, which are well suited to missions with limited control authority and tight propellant margins.

This work addresses this gap by formulating a stochastic low-thrust minimum-fuel problem in which thrust-magnitude errors, pointing errors, and initial-state dispersion are embedded directly in the equations of motion. The thrust execution errors are modeled as exogenous Gauss–Markov processes, so that each uncertainty realization defines a random ordinary differential equation. The terminal requirement is expressed in probabilistic form by requiring the spacecraft to reach prescribed position and velocity balls at the final time with assigned probability levels. This chance-constrained problem is then approximated by a finite-sample, CVaR-penalized formulation.

Within this framework, two indirect optimal-control problems are introduced. The first is a pathwise oracle problem, in which each uncertainty realization is allowed to use its own control history. This problem is not directly implementable, because it assumes perfect knowledge of each realization before the control is selected, but it provides a lower bound on the achievable risk-aware cost. The second is a common-control open-loop problem, in which all realizations share the same commanded throttle and thrust-direction histories. This second problem represents the open-loop control law that would be executed without in-flight corrections. The difference between the two optimal values quantifies the cost of replacing realization-dependent compensation with a single common command history.

The necessary conditions for both problems are derived using PMP. The resulting Hamiltonian structure, adjoint equations, optimal thrust directions, and switching functions are obtained explicitly. The numerical solution is performed through single-shooting schemes combined with an energy-to-fuel homotopy for the throttle law, smoothing and annealing of the CVaR terms, and continuation on the terminal-risk penalty weights. The proposed methodology is assessed on a representative interplanetary transfer inspired by the CubeSat M-ARGO, conceived as a standalone ESA deep-space CubeSat designed to rendezvous with a near-Earth asteroid.

## PROBLEM FORMULATION

Consider a spacecraft equipped with a low-thrust propulsion system and moving in the heliocentric two-body gravitational field. The spacecraft state is

$$\mathbf{x} = \begin{bmatrix} \mathbf{r} \\ \mathbf{v} \\ m \end{bmatrix}, \tag{1}$$

where $\mathbf{r} \in \mathbb{R}^3$ and $\mathbf{v} \in \mathbb{R}^3$ are the position and velocity vectors, respectively, and $m \in \mathbb{R}$ is the spacecraft mass. The commanded controls are the throttle factor $u(t) \in [0, 1]$ and the thrust direction $\boldsymbol{\alpha}(t) \in \mathbb{S}^2$, with $\|\boldsymbol{\alpha}(t)\| = 1$.

The low-thrust acceleration is affected by execution errors in both magnitude and direction. The thrust magnitude is modeled as

$$T^\omega(t) = T\left(1 + z(t, \omega)\right), \tag{2}$$

where $T$ is the nominal maximum thrust and $z(t, \omega)$ is the relative thrust-magnitude deviation. The process $z$ is modeled as an Ornstein–Uhlenbeck process[22]

$$\mathrm{d}z = -\beta_z z \,\mathrm{d}t + \sigma_z \,\mathrm{d}W_z, \tag{3}$$

where $\beta_z$ is the inverse correlation time and $\sigma_z$ is the diffusion intensity. Since the Ornstein–Uhlenbeck process is Gaussian and therefore unbounded, the formulation is restricted to the small-noise regime in which

$$1 + z(t, \omega) > 0, \qquad t \in [t_i, t_f], \tag{4}$$

for the realizations under consideration.

Pointing uncertainty is represented through a random rotation matrix

$$\mathbf{E}(t, \omega) \in SO(3), \tag{5}$$

which maps the commanded thrust direction into the actual thrust direction,

$$\tilde{\boldsymbol{\alpha}}(t, \omega) = \mathbf{E}(t, \omega)\boldsymbol{\alpha}(t). \tag{6}$$

In the numerical implementation, $\mathbf{E}(t, \omega)$ is generated from two independent Gauss–Markov angular-error processes, which define a small random attitude error. This formulation, differently from cone-about-command pointing model[23] and Gates' model,[24] avoids introducing an arbitrary transverse basis and provides a gauge-free description of the achieved thrust direction.

The stochastic equations of motion are therefore

$$\dot{\mathbf{x}}(t, \omega) = \mathbf{f}(\mathbf{x}, u, \boldsymbol{\alpha}; \omega, t) \Rightarrow \begin{bmatrix} \dot{\mathbf{r}}(t, \omega) \\ \dot{\mathbf{v}}(t, \omega) \\ \dot{m}(t, \omega) \end{bmatrix} = \begin{bmatrix} \mathbf{v}(t, \omega) \\ -\mu \dfrac{\mathbf{r}(t, \omega)}{\|\mathbf{r}(t, \omega)\|^3} + \dfrac{T\,(1 + z(t, \omega))}{m(t, \omega)}\, u(t)\, \mathbf{E}(t, \omega)\boldsymbol{\alpha}(t) \\ -\dfrac{T\,(1 + z(t, \omega))}{c}\, u(t) \end{bmatrix}, \tag{7}$$

where $\mu$ is the gravitational parameter of the Sun and $c$ is the effective exhaust velocity. The stochastic processes are assumed to be exogenous, namely independent of the spacecraft state and of the control. Therefore, once a realization $\omega$ is fixed, Eq. (7) is an ordinary differential equation with prescribed time-dependent coefficients. In this sense, the controlled low-thrust dynamics are modeled as a random ordinary differential equation (RODE).

The initial state is uncertain and is modeled as

$$\mathbf{x}(t_i, \omega) = \mathbf{x}_i(\omega), \qquad \mathbf{x}_i \sim \mathcal{N}(\bar{\mathbf{x}}_i, \Sigma_i), \tag{8}$$

where $\bar{\mathbf{x}}_i$ is the mean initial state and $\Sigma_i$ is the initial covariance matrix. The final time $t_f$ is fixed. At $t_f$, the spacecraft is required to reach prescribed terminal position and velocity balls with assigned probability levels,

$$\mathbb{P}\left(\|\mathbf{r}(t_f, \omega) - \mathbf{r}_f\| \le R\right) \ge k_r, \tag{9}$$

and

$$\mathbb{P}\left(\|\mathbf{v}(t_f, \omega) - \mathbf{v}_f\| \le V\right) \ge k_v, \tag{10}$$

where $\mathbf{r}_f$ and $\mathbf{v}_f$ are the target position and velocity, $R$ and $V$ are the admissible tolerances, and $k_r, k_v \in (0, 1)$ are the required probability levels.

To express the propellant consumption in Mayer form, an auxiliary state $y \in \mathbb{R}$ is introduced,

$$\dot{y}(t, \omega) = \frac{T\,(1 + z(t, \omega))}{c}\, u(t), \qquad y(t_i, \omega) = 0. \tag{11}$$

Thus,

$$y(t_f, \omega) = \frac{T}{c} \int_{t_i}^{t_f} (1 + z(t, \omega))\, u(t)\, \mathrm{d}t \tag{12}$$

is the propellant mass consumed along the realization $\omega$.

The stochastic low-thrust optimal-control problem considered in this work can then be stated as follows:

Find an admissible thrust magnitude and direction history that minimize the expected propellant consumption,

$$\min_{u(\cdot),\, \boldsymbol{\alpha}(\cdot)} \mathbb{E}\left[y(t_f, \omega)\right], \tag{13}$$

subject to the stochastic dynamics in Eq. (7), the fuel-state dynamics in Eq. (11), the initial distribution in Eq. (8), the control constraints $u(t) \in [0, 1]$ and $\boldsymbol{\alpha}(t) \in \mathbb{S}^2$, and the terminal chance constraints in Eqs. (9)–(10).

## METHODOLOGY

### Sample-average approximation

The chance-constrained stochastic optimal-control problem in Eqs. (7)–(13) is not solved directly. Instead, a finite set of uncertainty realizations $\{\omega_s\}_{s=1}^{N_s}$ is introduced, with associated non-negative weights $w_s$ satisfying

$$\sum_{s=1}^{N_s} w_s = 1. \tag{14}$$

The weights may represent either Monte Carlo samples, in which case $w_s = 1/N_s$, or more general quadrature weights.

For each realization $\omega_s$, the sampled state is denoted by

$$\mathbf{x}^{(s)} = \begin{bmatrix} \mathbf{r}^{(s)} \\ \mathbf{v}^{(s)} \\ m^{(s)} \end{bmatrix}, \qquad y^{(s)} = y(t, \omega_s), \tag{15}$$

with initial condition

$$\mathbf{x}^{(s)}(t_i) = \mathbf{x}_i^{(s)}. \tag{16}$$

The sampled uncertainty histories are denoted by $z^{(s)}(t) = z(t, \omega_s)$ and $\mathbf{E}^{(s)}(t) = \mathbf{E}(t, \omega_s)$. The corresponding sampled dynamics are

$$\dot{\mathbf{x}}^{(s)} = \mathbf{f}^{(s)}\left(\mathbf{x}^{(s)}, u^{(s)}, \boldsymbol{\alpha}^{(s)}, t\right) \Rightarrow \begin{bmatrix} \dot{\mathbf{r}}^{(s)} \\ \dot{\mathbf{v}}^{(s)} \\ \dot{m}^{(s)} \end{bmatrix} = \begin{bmatrix} \mathbf{v}^{(s)} \\ -\mu \dfrac{\mathbf{r}^{(s)}}{\left\|\mathbf{r}^{(s)}\right\|^3} + \dfrac{T\left(1 + z^{(s)}\right)}{m^{(s)}} u^{(s)}\, \mathbf{E}^{(s)} \boldsymbol{\alpha}^{(s)} \\ -\dfrac{T\left(1 + z^{(s)}\right)}{c} u^{(s)} \end{bmatrix}. \tag{17}$$

The fuel-state dynamics become

$$\dot{y}^{(s)} = \frac{T\left(1 + z^{(s)}\right)}{c} u^{(s)}, \qquad y^{(s)}(t_i) = 0. \tag{18}$$

The dependence of $z^{(s)}$, $\mathbf{E}^{(s)}$, $u^{(s)}$, and $\boldsymbol{\alpha}^{(s)}$ on time has been omitted for compactness.
The notation in Eq. (17) corresponds to the pathwise oracle problem, in which each realization is allowed to use its own control history $u^{(s)}(\cdot)$ and $\boldsymbol{\alpha}^{(s)}(\cdot)$. In the common-control open-loop problem introduced below, the same dynamics are used, but the admissible controls are restricted to a single history shared by all realizations, namely

$$u^{(1)}(\cdot) = \cdots = u^{(N_s)}(\cdot) = u(\cdot), \qquad \boldsymbol{\alpha}^{(1)}(\cdot) = \cdots = \boldsymbol{\alpha}^{(N_s)}(\cdot) = \boldsymbol{\alpha}(\cdot). \tag{19}$$

**Risk-aware ensemble objective**

The hard chance constraints in Eqs. (9)–(10) are treated through soft Conditional Value at Risk (CVaR) penalties. This produces a differentiable, sample-based risk-aware surrogate of the original stochastic problem, while preserving the expected-fuel objective.
For each realization, the terminal position and velocity violations are defined as

$$d_r^{(s)} = \left\| \mathbf{r}^{(s)}(t_f) - \mathbf{r}_f \right\|^2 - R^2 \tag{20}$$

and

$$d_v^{(s)} = \left\| \mathbf{v}^{(s)}(t_f) - \mathbf{v}_f \right\|^2 - V^2. \tag{21}$$

Thus, $d_r^{(s)} \le 0$ and $d_v^{(s)} \le 0$ correspond to satisfaction of the terminal position and velocity balls for the $s$-th realization.
The expected-fuel term is approximated as

$$J_f = \sum_{s=1}^{N_s} w_s \, y^{(s)}(t_f). \tag{22}$$

To penalize the upper tail of the terminal violations, Conditional Value at Risk (CVaR) penalties are adopted.[25] The sample-average approximation of the CVaR of the position violation is

$$\mathrm{CVaR}_{k_r}(d_r) = \eta_r + \frac{1}{1-k_r} \sum_{s=1}^{N_s} w_s \left( d_r^{(s)} - \eta_r \right)^+, \tag{23}$$

where $(a)^+ = \max(a, 0)$ and $\eta_r$ is the corresponding Value-at-Risk auxiliary variable. Similarly, for the terminal velocity violation,

$$\mathrm{CVaR}_{k_v}(d_v) = \eta_v + \frac{1}{1-k_v} \sum_{s=1}^{N_s} w_s \left( d_v^{(s)} - \eta_v \right)^+, \tag{24}$$

where $\eta_v$ is the velocity-side Value-at-Risk auxiliary variable.
The sample-based risk-aware objective is then

$$J = \sum_{s=1}^{N_s} w_s \, y^{(s)}(t_f) + \gamma_r \left[ \eta_r + \frac{1}{1-k_r} \sum_{s=1}^{N_s} w_s \left( d_r^{(s)} - \eta_r \right)^+ \right] + \gamma_v \left[ \eta_v + \frac{1}{1-k_v} \sum_{s=1}^{N_s} w_s \left( d_v^{(s)} - \eta_v \right)^+ \right], \tag{25}$$

where $\gamma_r \ge 0$ and $\gamma_v \ge 0$ are penalty parameters. The first term approximates the expected propellant consumption, whereas the second and third terms penalize the upper tails of the terminal position and velocity violations. Since the terminal requirements are imposed through penalties, the formulation does not enforce the original chance constraints as hard constraints. Instead, $\gamma_r$ and $\gamma_v$ regulate the trade-off between expected fuel consumption and terminal-risk reduction.

The auxiliary variables $\eta_r$ and $\eta_v$ are optimized together with the control histories. For fixed terminal violations, their optimality conditions are given, in the subgradient sense, by

$$1 - \frac{1}{1-k_r}\sum_{s=1}^{N_s} w_s \, \mathbf{1}_{\{d_r^{(s)} > \eta_r\}} = 0 \tag{26}$$

and

$$1 - \frac{1}{1-k_v}\sum_{s=1}^{N_s} w_s \, \mathbf{1}_{\{d_v^{(s)} > \eta_v\}} = 0. \tag{27}$$

Thus, $\eta_r$ and $\eta_v$ correspond to the sample-based Value at Risk of the terminal position and velocity violations, respectively, up to the usual non-uniqueness associated with discrete weighted samples.

**Pathwise oracle bound**

The first optimization problem considered in this work is the pathwise oracle problem. It is obtained by minimizing the sample-based objective in Eq. (25) while allowing each uncertainty realization to use its own control history. Therefore, the admissible controls are

$$u^{(s)}(\cdot) \in [0,1], \qquad \boldsymbol{\alpha}^{(s)}(\cdot) \in \mathbb{S}^2, \qquad s = 1,\ldots,N_s. \tag{28}$$

The corresponding optimal value is denoted by $J_1$. Since the complete histories $z^{(s)}(t)$ and $\mathbf{E}^{(s)}(t)$ are assumed to be known before the optimization is performed, the resulting control is anticipative. Thus, $J_1$ is not the performance of an implementable guidance law, but a perfect-information lower bound associated with the sampled uncertainty ensemble.

The pathwise oracle problem can be written as

$$J_1 = \min_{\substack{u^{(s)}(\cdot),\, \boldsymbol{\alpha}^{(s)}(\cdot) \\ s=1,\ldots,N_s \\ \eta_r,\, \eta_v}} J. \tag{29}$$

For fixed values of $\eta_r$ and $\eta_v$, the sample trajectories are decoupled. Therefore, the necessary conditions can be derived independently for each realization.

For the $s$-th realization, the augmented state is

$$\mathbf{x}_a^{(s)} = \begin{bmatrix} \mathbf{r}^{(s)} \\ \mathbf{v}^{(s)} \\ m^{(s)} \\ y^{(s)} \end{bmatrix}. \tag{30}$$

For brevity's sake, the superscript $(s)$ is omitted in the derivation when no ambiguity arises. The Hamiltonian associated with the $s$-th pathwise problem is

$$\mathcal{H}_1^{(s)} = \boldsymbol{\lambda}_r \cdot \mathbf{v} + \boldsymbol{\lambda}_v \cdot \left( -\mu \frac{\mathbf{r}}{\|\mathbf{r}\|^3} + \frac{T(1+z)}{m}\, u\, \mathbf{E}\boldsymbol{\alpha} \right) - \lambda_m \frac{T(1+z)}{c}\, u + \lambda_y \frac{T(1+z)}{c}\, u, \tag{31}$$

where $\boldsymbol{\lambda} = \left[\boldsymbol{\lambda}_r, \boldsymbol{\lambda}_v, \lambda_m, \lambda_y\right]^T$ is the vector of costates. The corresponding adjoint dynamics are

$$\dot{\boldsymbol{\lambda}} = -\frac{\partial \mathcal{H}_1^{(s)}}{\partial \mathbf{x}_a} \Rightarrow \begin{bmatrix} \dot{\boldsymbol{\lambda}}_r \\ \dot{\boldsymbol{\lambda}}_v \\ \dot{\lambda}_m \\ \dot{\lambda}_y \end{bmatrix} = \begin{bmatrix} \mu\left(\frac{1}{\|\mathbf{r}\|^3}\mathbf{I} - 3\frac{\mathbf{r}\mathbf{r}^\top}{\|\mathbf{r}\|^5}\right)\cdot\boldsymbol{\lambda}_v \\ -\boldsymbol{\lambda}_r \\ \frac{T(1+z)}{m^2}\, u\,\boldsymbol{\lambda}_v\cdot\mathbf{E}\boldsymbol{\alpha} \\ 0 \end{bmatrix}. \tag{32}$$

The explicit dependence of $z$, $\mathbf{E}$, $u$, and $\boldsymbol{\alpha}$ on time and on the realization index has been omitted for compactness.

The terminal adjoint conditions follow from the terminal part of the sample objective. Defining the active-set indicators

$$\chi_r^{(s)} = \mathbf{1}_{\{d_r^{(s)} > \eta_r\}}, \qquad \chi_v^{(s)} = \mathbf{1}_{\{d_v^{(s)} > \eta_v\}}, \tag{33}$$

one obtains

$$\boldsymbol{\lambda}_r^{(s)}(t_f) = \frac{\gamma_r}{1-k_r}\,\chi_r^{(s)}\,2\left(\mathbf{r}^{(s)}(t_f) - \mathbf{r}_f\right), \tag{34}$$

$$\boldsymbol{\lambda}_v^{(s)}(t_f) = \frac{\gamma_v}{1-k_v}\,\chi_v^{(s)}\,2\left(\mathbf{v}^{(s)}(t_f) - \mathbf{v}_f\right), \tag{35}$$

and

$$\lambda_m^{(s)}(t_f) = 0, \qquad \lambda_y^{(s)}(t_f) = 1. \tag{36}$$

The final mass is free, whereas the terminal value of $y^{(s)}$ enters the expected-fuel objective. At the nonsmooth points $d_r^{(s)} = \eta_r$ or $d_v^{(s)} = \eta_v$, Eqs. (34)–(35) are interpreted in the subgradient sense.

The optimal commanded thrust direction is obtained by minimizing the Hamiltonian with respect to $\boldsymbol{\alpha}$. Since $\mathbf{E}^{(s)} \in SO(3)$, the pointing-error map preserves unit norms. Therefore, for the oracle problem, the known pointing error can be exactly pre-compensated. The optimal direction is

$$\boldsymbol{\alpha}^{\star,(s)} = -\left(\mathbf{E}^{(s)}\right)^T \frac{\boldsymbol{\lambda}_v^{(s)}}{\lambda_v^{(s)}}, \tag{37}$$

where $\lambda_v^{(s)}$ is the norm of $\boldsymbol{\lambda}_v^{(s)}$, so that the achieved thrust direction satisfies

$$\mathbf{E}^{(s)}\boldsymbol{\alpha}^{\star,(s)} = -\frac{\boldsymbol{\lambda}_v^{(s)}}{\lambda_v^{(s)}}. \tag{38}$$

Consequently,

$$\boldsymbol{\lambda}_v^{(s)}\cdot\mathbf{E}^{(s)}\boldsymbol{\alpha}^{\star,(s)} = -\lambda_v^{(s)} \tag{39}$$

Substituting Eq. (39) into Eq. (31), the minimized Hamiltonian can be written as

$$\mathcal{H}_1^{(s)} = \boldsymbol{\lambda}_r^{(s)}\cdot\mathbf{v}^{(s)} - \boldsymbol{\lambda}_v^{(s)}\cdot\mu\,\frac{\mathbf{r}^{(s)}}{\left\|\mathbf{r}^{(s)}\right\|^3} + \frac{T\left(1+z^{(s)}\right)}{c}\,u^{(s)}\,S_1^{(s)}, \tag{40}$$

where the oracle switching function is

$$S_1^{(s)} = -\frac{c}{m^{(s)}}\lambda_v^{(s)} - \lambda_m^{(s)} + \lambda_y^{(s)}. \tag{41}$$

The optimal throttle is therefore

$$u^{\star,(s)} = \begin{cases} 0, & S_1^{(s)} > 0, \\ 1, & S_1^{(s)} < 0, \\ \text{undetermined}, & S_1^{(s)} = 0. \end{cases} \tag{42}$$

Outside the singular case $S_1^{(s)} = 0$, the optimal throttle has bang–bang structure.
Once $\boldsymbol{\alpha}^{\star,(s)}$ and $u^{\star,(s)}$ are determined as functions of the states, costates, and prescribed uncertainty histories, the pathwise equations can be integrated implicitly. Thus, for fixed $\eta_r$ and $\eta_v$, each realization defines a two-point boundary value problem. The unknown initial costates $\boldsymbol{\lambda}_i^{(s)}$ are determined by enforcing the terminal adjoint conditions in Eqs. (34)–(36). The outer variables $\eta_r$ and $\eta_v$ are then updated so that the optimality conditions in Eqs. (26)–(27) are satisfied.
The value $J_1$ obtained in this way is an oracle bound. Since each realization uses its own control history and can pre-compensate the known thrust-magnitude and pointing-error histories, the admissible control set of the oracle problem contains the admissible control set of the common-control problem introduced next.

**Common-control open-loop bound**

The second optimization problem considered in this work is the common-control open-loop problem. It is obtained by minimizing the same sample-based objective in Eq. (25), but with a single control history shared by all uncertainty realizations. Therefore, the admissible controls satisfy

$$u^{(1)}(\cdot) = \cdots = u^{(N_s)}(\cdot) = u(\cdot), \qquad \boldsymbol{\alpha}^{(1)}(\cdot) = \cdots = \boldsymbol{\alpha}^{(N_s)}(\cdot) = \boldsymbol{\alpha}(\cdot), \tag{43}$$

with

$$u(t) \in [0,1], \qquad \boldsymbol{\alpha}(t) \in \mathbb{S}^2. \tag{44}$$

The corresponding optimal value is denoted by $J_2$. Differently from the oracle problem, the control cannot be tailored to each realization. The same commanded thrust direction and the same throttle profile must be used for the whole ensemble.
The common-control problem can be written as

$$J_2 = \min_{\substack{u(\cdot),\,\boldsymbol{\alpha}(\cdot) \\ \eta_r,\,\eta_v}} J. \tag{45}$$

For fixed values of $\eta_r$ and $\eta_v$, the state and adjoint equations remain decoupled across the samples. However, the Hamiltonian minimization is now coupled, because the same control variables $u$ and $\boldsymbol{\alpha}$ appear in all sample Hamiltonians.
For the $s$-th realization, the Hamiltonian has the same structure as in Eq. (31), with the realization-dependent controls replaced by the common controls,

$$\mathcal{H}_2^{(s)} = \boldsymbol{\lambda}_r^{(s)} \cdot \mathbf{v}^{(s)} + \boldsymbol{\lambda}_v^{(s)} \cdot \left( -\mu \frac{\mathbf{r}^{(s)}}{\left\|\mathbf{r}^{(s)}\right\|^3} + \frac{T\left(1+z^{(s)}\right)}{m^{(s)}} \, u \, \mathbf{E}^{(s)} \boldsymbol{\alpha} \right) - \lambda_m^{(s)} \frac{T\left(1+z^{(s)}\right)}{c} u + \lambda_y^{(s)} \frac{T\left(1+z^{(s)}\right)}{c} u. \tag{46}$$

The ensemble Hamiltonian to be minimized with respect to the common controls is then

$$\mathcal{H}_2 = \sum_{s=1}^{N_s} w_s \, \mathcal{H}_2^{(s)}. \tag{47}$$

The adjoint equations associated with each realization are

$$\begin{bmatrix} \dot{\boldsymbol{\lambda}}_r^{(s)} \\ \dot{\boldsymbol{\lambda}}_v^{(s)} \\ \dot{\lambda}_m^{(s)} \\ \dot{\lambda}_y^{(s)} \end{bmatrix} = \begin{bmatrix} \mu \left( \frac{1}{\left\|\mathbf{r}^{(s)}\right\|^3}\mathbf{I} - 3\frac{\mathbf{r}^{(s)}\left(\mathbf{r}^{(s)}\right)^\top}{\left\|\mathbf{r}^{(s)}\right\|^5} \right) \cdot \boldsymbol{\lambda}_v^{(s)} \\ -\boldsymbol{\lambda}_r^{(s)} \\ \frac{T\left(1+z^{(s)}\right)}{\left(m^{(s)}\right)^2} u\, \boldsymbol{\lambda}_v^{(s)} \cdot \mathbf{E}^{(s)}\boldsymbol{\alpha} \\ 0 \end{bmatrix}, \qquad s = 1, \ldots, N_s. \tag{48}$$

The terminal adjoint conditions are the same as in the oracle problem, namely

$$\boldsymbol{\lambda}_r^{(s)}(t_f) = \frac{\gamma_r}{1-k_r}\, \chi_r^{(s)}\, 2\left(\mathbf{r}^{(s)}(t_f) - \mathbf{r}_f\right), \tag{49}$$

$$\boldsymbol{\lambda}_v^{(s)}(t_f) = \frac{\gamma_v}{1-k_v}\, \chi_v^{(s)}\, 2\left(\mathbf{v}^{(s)}(t_f) - \mathbf{v}_f\right), \tag{50}$$

and

$$\lambda_m^{(s)}(t_f) = 0, \qquad \lambda_y^{(s)}(t_f) = 1. \tag{51}$$

where $\chi_r^{(s)}$ and $\chi_v^{(s)}$ are defined in Eq. (33). As before, these conditions are interpreted in the subgradient sense at the nonsmooth points $d_r^{(s)} = \eta_r$ and $d_v^{(s)} = \eta_v$.
The optimal common thrust direction is obtained by minimizing the ensemble Hamiltonian in Eq. (47) with respect to $\boldsymbol{\alpha}$. The terms depending on $\boldsymbol{\alpha}$ can be written as

$$u\,T\,\boldsymbol{\alpha} \cdot \sum_{s=1}^{N_s} w_s \frac{1+z^{(s)}}{m^{(s)}} \left(\mathbf{E}^{(s)}\right)^T \boldsymbol{\lambda}_v^{(s)}. \tag{52}$$

Therefore, defining

$$\mathbf{p}_2 = \sum_{s=1}^{N_s} w_s \frac{1+z^{(s)}}{m^{(s)}} \left(\mathbf{E}^{(s)}\right)^T \boldsymbol{\lambda}_v^{(s)}, \tag{53}$$

the optimal common commanded direction is

$$\boldsymbol{\alpha}^\star = -\frac{\mathbf{p}_2}{\|\mathbf{p}_2\|}, \qquad \mathbf{p}_2 \neq \mathbf{0}. \tag{54}$$

If $\mathbf{p}_2 = \mathbf{0}$, the Hamiltonian is locally independent of $\boldsymbol{\alpha}$ and the common direction is not determined by the first-order condition. This degenerate case is non-generic and physically irrelevant, and is therefore disregarded.
Substituting Eq. (54) into Eq. (47), the minimized ensemble Hamiltonian can be written as

$$\mathcal{H}_2 = \sum_{s=1}^{N_s} w_s \left[ \boldsymbol{\lambda}_r^{(s)} \cdot \mathbf{v}^{(s)} - \boldsymbol{\lambda}_v^{(s)} \cdot \mu \frac{\mathbf{r}^{(s)}}{\left\|\mathbf{r}^{(s)}\right\|^3} \right] + \frac{T}{c}\, u\, S_2, \tag{55}$$

where the common-control switching function is

$$S_2 = -c\,\|\mathbf{p}_2\| + \sum_{s=1}^{N_s} w_s \left(1 + z^{(s)}\right)\left(-\lambda_m^{(s)} + \lambda_y^{(s)}\right). \tag{56}$$

The optimal common throttle is then

$$u^\star = \begin{cases} 0, & S_2 > 0, \\ 1, & S_2 < 0, \\ \text{undetermined}, & S_2 = 0. \end{cases} \tag{57}$$

Thus, also in the common-control problem, the optimal throttle has bang–bang structure outside the singular case $S_2 = 0$. However, the switching function is now an ensemble quantity. It depends on all sample states and costates through the weighted vector $\mathbf{p}_2$ and through the weighted contribution of the mass and fuel costates.

For fixed $\eta_r$ and $\eta_v$, the common-control open-loop solution is obtained by solving the coupled two-point boundary value problem formed by the $N_s$ state and adjoint systems, the shared control laws in Eqs. (54) and (57), and the terminal conditions in Eqs. (49)–(51). The outer variables $\eta_r$ and $\eta_v$ are updated so that Eqs. (26)–(27) are satisfied.

## NUMERICAL SOLUTION STRATEGY

### Pathwise oracle problem

The pathwise oracle problem is solved by exploiting the decoupled structure of the necessary conditions derived in previous Section. Once the uncertainty realizations are fixed, each sample defines a deterministic two-point boundary-value problem in the corresponding state and costate variables. The only coupling among samples is introduced by the CVaR terms, through the auxiliary variables $\eta_r$ and $\eta_v$ and through the terminal tail weights. Therefore, the oracle solution is obtained through an outer ensemble-level risk iteration and a set of independent single-sample shooting problems.

Each uncertainty realization is generated before the optimization is performed. The sampled initial condition $\mathbf{x}_i^{(s)}$, the thrust-magnitude history $z^{(s)}(t)$, and the pointing-error history $\mathbf{E}^{(s)}(t)$ are then kept fixed throughout the solve. In the pathwise oracle problem, the pointing error can be pre-compensated by the realization-dependent commanded direction. Consequently, the achieved thrust direction is aligned with the negative velocity costate, and the pointing-error rotation does not enter the propagated oracle dynamics. The thrust-magnitude history $z^{(s)}(t)$, however, remains in the dynamics and in the fuel consumption.

The fuel-optimal throttle law is discontinuous in the limit problem. To improve the robustness of the shooting procedure, the bang–bang problem is approached through an energy-to-fuel homotopy. The auxiliary fuel dynamics are replaced, during continuation, by the regularized accumulator

$$\dot{y}_\epsilon^{(s)} = \frac{T\left(1 + z^{(s)}\right)}{c}\left[u^{(s)} - \epsilon u^{(s)}\left(1 - u^{(s)}\right)\right], \qquad \epsilon \in [0, 1] \tag{58}$$

The corresponding regularized throttle is

$$u_\epsilon^{\star,(s)} = \Pi_{[0,1]}\left(\frac{\epsilon - S_1^{(s)}}{2\epsilon}\right), \tag{59}$$

where $\Pi_{[0,1]}$ denotes projection onto the interval $[0, 1]$. As $\epsilon$ decreases, Eq. (59) approaches the bang–bang law in Eq. (42). The continuation starts from a comparatively smooth problem and progressively decreases $\epsilon$ to a small final value $\epsilon_f$.

For a fixed realization, fixed values of $\eta_r$ and $\eta_v$, and fixed terminal tail weights, the unknowns of the shooting problem are the seven initial costates

$$\mathbf{p}_0^{(s)} = \begin{bmatrix} \boldsymbol{\lambda}_r^{(s)}(t_i) \\ \boldsymbol{\lambda}_v^{(s)}(t_i) \\ \lambda_m^{(s)}(t_i) \end{bmatrix}. \tag{60}$$

The state-costate system is propagated forward from $t_i$ to $t_f$, and the shooting residual enforces the terminal transversality conditions. During the $\epsilon$-continuation, each continuation step is warm-started from the previous solution. A tangent predictor is used to estimate the change in $\mathbf{p}_0^{(s)}$ induced by the change in $\epsilon$, and the resulting prediction is corrected by solving the nonlinear shooting residual. If a continuation step fails, the step in $\epsilon$ is bisected until either convergence is recovered or a minimum continuation step is reached.

The nonsmooth CVaR positive part is also regularized. For a smoothing temperature $\tau > 0$, define

$$(a)_\tau^+ = \tau \log\left(1 + \exp\left(\frac{a}{\tau}\right)\right), \tag{61}$$

and the associated smooth tail indicator

$$\sigma_\tau(a) = \frac{1}{1 + \exp(-a/\tau)}. \tag{62}$$

The smoothed sample CVaR of a generic terminal violation $d^{(s)}$ is then written as

$$\mathrm{CVaR}_{k,\tau}(d) = \eta + \frac{1}{1-k} \sum_{s=1}^{N_s} w_s \left(d^{(s)} - \eta\right)_\tau^+. \tag{63}$$

For fixed terminal violations, the smoothed VaR variable is computed from

$$\sum_{s=1}^{N_s} w_s \sigma_\tau \left(d^{(s)} - \eta\right) = 1 - k. \tag{64}$$

The smoothed terminal tail weights are evaluated at the terminal state generated by the current shooting iterate,

$$\rho_r^{(s)} = \sigma_{\tau_r}\left(d_r^{(s)} - \eta_r\right), \qquad \rho_v^{(s)} = \sigma_{\tau_v}\left(d_v^{(s)} - \eta_v\right). \tag{65}$$

The corresponding terminal transversality conditions become

$$\boldsymbol{\lambda}_r^{(s)}(t_f) = \frac{\gamma_r}{1-k_r} \rho_r^{(s)}\, 2\left(\mathbf{r}^{(s)}(t_f) - \mathbf{r}_f\right), \tag{66}$$

$$\boldsymbol{\lambda}_v^{(s)}(t_f) = \frac{\gamma_v}{1-k_v} \rho_v^{(s)}\, 2\left(\mathbf{v}^{(s)}(t_f) - \mathbf{v}_f\right), \tag{67}$$

with

$$\lambda_m^{(s)}(t_f) = 0. \tag{68}$$

Evaluating the tail weights inside the shooting residual makes the transversality conditions state-consistent and avoids freezing the active set from a previous outer iteration. The smoothing temperatures $\tau_r$ and $\tau_v$ are annealed progressively. At each annealing level, they are chosen proportional to the current dispersion of the terminal violations, with a prescribed lower bound to avoid degeneracy when the ensemble becomes tightly clustered.

Because the terminal requirements are imposed through soft CVaR penalties, the penalty coefficients $\gamma_r$ and $\gamma_v$ are increased through continuation. At each penalty level, the smoothed oracle problem is solved to convergence. The terminal violations are then evaluated using the nonsmoothed CVaR definitions in Eqs. (23)–(24). If the hard CVaRs are non-positive, up to the prescribed tolerance, the continuation is stopped. Otherwise, both penalty coefficients are multiplied by a constant factor and the process is repeated, using the previous costates and VaR variables as warm starts. The final oracle value $J_1$ is reported using the nonsmoothed CVaR metrics at the final penalty level.

The numerical implementation is reported in Algorithm 1 and in Figure 1.

### Common-control open-loop problem

The common-control open-loop problem is solved using the same regularization and risk-continuation strategy adopted for the pathwise oracle problem. The main difference is that the sample problems are no longer independent. Although each realization has its own state and costate variables, all realizations share the same throttle history and the same commanded thrust direction. Therefore, the Hamiltonian minimization must be performed at the ensemble level, and the resulting two-point boundary-value problem is solved as one coupled shooting problem.

For a fixed ensemble of uncertainty realizations, the open-loop unknown is the stacked vector of initial costates

$$\mathbf{P}_0 = \left[\left(\mathbf{p}_0^{(1)}\right)^T \quad \left(\mathbf{p}_0^{(2)}\right)^T \quad \cdots \quad \left(\mathbf{p}_0^{(N_s)}\right)^T\right]^T \in \mathbb{R}^{7N_s}, \tag{69}$$

where

$$\mathbf{p}_0^{(s)} = \begin{bmatrix} \boldsymbol{\lambda}_r^{(s)}(t_i) \\ \boldsymbol{\lambda}_v^{(s)}(t_i) \\ \lambda_m^{(s)}(t_i) \end{bmatrix}. \tag{70}$$

Given $\mathbf{P}_0$, all sample state-costate systems are propagated forward simultaneously. At every integration time, the common commanded direction and common throttle are reconstructed from the ensemble optimality conditions derived in the previous Section. Thus, the control is not introduced as an additional discretized unknown; it is evaluated pointwise from the current ensemble state and costate values.

The shooting residual is obtained by stacking the terminal residuals of all samples

$$\mathbf{R}_2(\mathbf{P}_0) = \left[\left(\mathbf{R}_2^{(1)}\right)^T \quad \left(\mathbf{R}_2^{(2)}\right)^T \quad \cdots \quad \left(\mathbf{R}_2^{(N_s)}\right)^T\right]^T \in \mathbb{R}^{7N_s}. \tag{71}$$

For each realization, $\mathbf{R}_2^{(s)}$ enforces the smoothed terminal transversality conditions associated with the CVaR-penalized objective, together with the free-final-mass condition $\lambda_m^{(s)}(t_f) = 0$. The terminal weights entering the transversality conditions are evaluated consistently at the terminal state reached by the current shooting iterate. Therefore, as in the oracle case, the active-tail information is not frozen from a previous outer iteration.

The same energy-to-fuel homotopy introduced in Eq. (58) is used for the open-loop problem. For each value of $\epsilon$, the stacked shooting system

$$\mathbf{R}_2\left(\mathbf{P}_0; \epsilon, \eta_r, \eta_v, \tau_r, \tau_v\right) = \mathbf{0} \tag{72}$$

**Algorithm 1:** Numerical solution of the pathwise oracle problem

**Input:** Problem data, sample size $N_s$, $\epsilon$-schedule, $\tau$-annealing schedule, penalty-growth factor, CVaR tolerance

**Output:** Oracle value $J_1$, sample costates $\mathbf{p}_0^{(s)}$, terminal violations, fuel consumption

Generate the uncertainty ensemble $\{\omega_s, w_s\}_{s=1}^{N_s}$

Initialize the costate guesses $\mathbf{p}_0^{(s)}$

**for** $s = 1, \ldots, N_s$ **do**

  Solve a full-pull single-sample problem by $\epsilon$-homotopy

  Store $\mathbf{p}_0^{(s)}$, $d_r^{(s)}$, $d_v^{(s)}$, and $y^{(s)}(t_f)$

**end**

Set the penalty multiplier $m_\gamma = 1$

**for** *each penalty-continuation level* **do**

  Scale $\gamma_r$ and $\gamma_v$ by $m_\gamma$

  **for** *each smoothing level* $\tau$ **do**

    Compute $\eta_r$ and $\eta_v$ from the smoothed VaR conditions

    Initialize $J_1^{\mathrm{old}} = \infty$

    **for** *outer iteration* $= 1, \ldots, N_{\max}$ **do**

      **for** $s = 1, \ldots, N_s$ **do**

        Solve the smoothed single-sample TPBVP at $\epsilon_f$

        **if** *the direct solve fails* **then**

          Recover the sample solution by the full $\epsilon$-homotopy

        **end**

        Update $d_r^{(s)}$, $d_v^{(s)}$, and $y^{(s)}(t_f)$

      **end**

      Evaluate the smoothed CVaRs and the smoothed objective $J_1$

      **if** *the relative change in* $J_1$ *is below tolerance* **then**

        break

      **end**

      Update $\eta_r$ and $\eta_v$ by damped smoothed-VaR iteration

    **end**

  **end**

  Evaluate the hard CVaRs from the nonsmoothed definitions

  **if** $\mathrm{CVaR}_{k_r}(d_r) \le 0$ *and* $\mathrm{CVaR}_{k_v}(d_v) \le 0$ **then**

    break

  **end**

  Increase $m_\gamma$

**end**

Compute the final hard-metric objective $J_1$

**return** $J_1$, $\mathbf{p}_0^{(s)}$*, terminal violations, and fuel values*

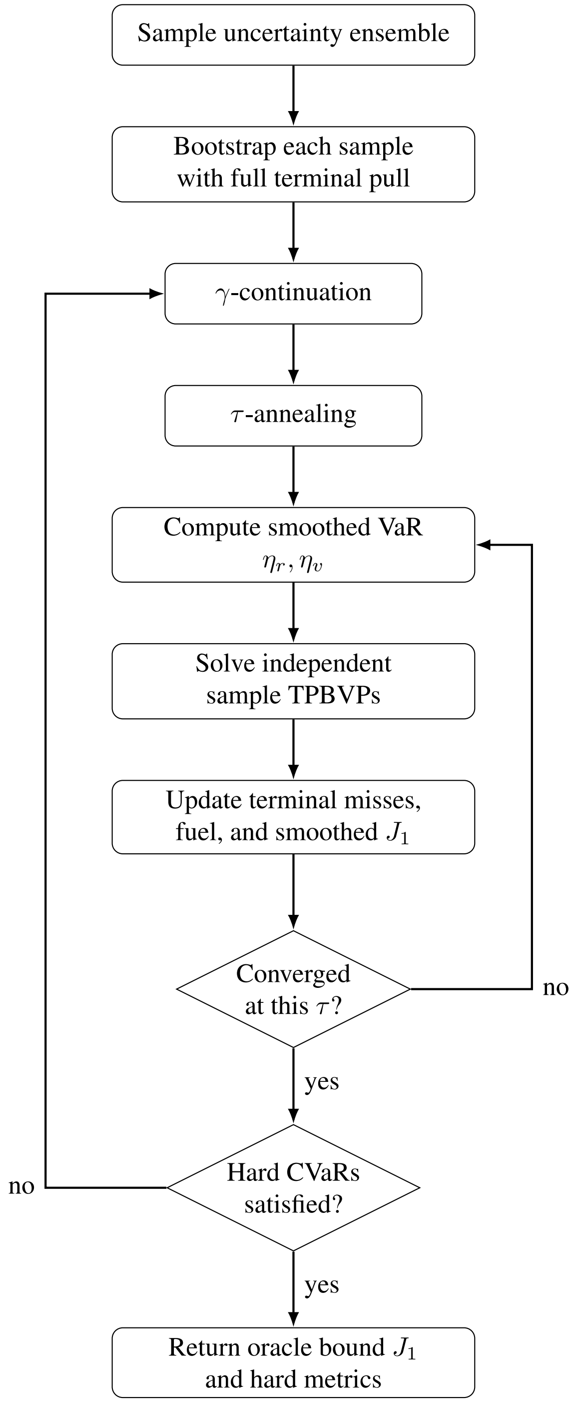


**Figure 1. Numerical workflow used to compute the pathwise oracle bound.**

is solved for the initial costates. The solution at one value of $\epsilon$ is used to initialize the next continuation step. As $\epsilon$ approaches the final value $\epsilon_f$, the regularized open-loop control approaches the bang–bang common-control solution.

The CVaR treatment is also identical to the oracle case. The nonsmooth positive part is replaced by the softplus approximation in Eq. (61), and the smoothed tail indicators are computed with Eq. (62). The VaR variables $\eta_r$ and $\eta_v$ are updated from the smoothed VaR conditions, while the smoothing temperatures $\tau_r$ and $\tau_v$ are gradually reduced. At each smoothing level, the coupled open-loop shooting problem is solved until the smoothed objective becomes stationary within the prescribed tolerance.

A penalty continuation is then used to recover satisfaction of the original terminal chance constraints. Starting from an initial value of the penalty multiplier, the open-loop problem is solved with the smoothed CVaR objective. The resulting terminal violations are then evaluated using the nonsmoothed CVaR definitions. If the hard CVaRs are non-positive, up to tolerance, the continuation is stopped. Otherwise, the penalty multiplier is increased and the coupled shooting problem is solved again using the previous stacked costate vector and VaR variables as warm starts. The final value $J_2$ is evaluated using the same hard CVaR metrics used for the oracle bound.

Because the common-control formulation couples all samples through the control law, the open-loop shooting problem is substantially more ill-conditioned than the pathwise oracle problem. In the present implementation, the complete ensemble is solved by monolithic single shooting. This is adequate for the small ensembles considered here and preserves a direct connection with the indirect necessary conditions. For larger ensembles or longer transfers, the same formulation could be improved by replacing the monolithic shooting system with a multiple-shooting scheme, in which intermediate continuity constraints are introduced to improve conditioning.

## RESULTS

The proposed approach is assessed on a representative deep-space low-thrust transfer inspired by the M-ARGO mission concept. The Miniaturised Asteroid Remote Geophysical Observer (M-ARGO) was conceived as a standalone deep-space CubeSat for rendezvous with a near-Earth asteroid.[26] The mission concept considered a small spacecraft equipped with electric propulsion, released as a secondary payload and subsequently operated in an autonomous heliocentric cruise. This scenario is representative of missions for which thrust authority is limited and uncertainty in the executed control can have a non-negligible effect on terminal accuracy. The objective of the numerical test is not to perform a full mission design campaign, but to assess the behavior of the proposed stochastic indirect formulation on a single representative transfer. The same set of uncertainty realizations is used to compute both the pathwise oracle bound $J_1$ and the common-control open-loop bound $J_2$. This allows the difference $J_2 - J_1$ to be interpreted as the cost of enforcing a single open-loop control history across the ensemble.

### Test case

The spacecraft parameters used in the simulation are reported in Table 1.

The baseline trajectory to asteroid 2000 SG344, having the characteristics in Table 3, is used as test case scenario in this work. Keplerian parameters for the target asteroid are given in Table 2. Asteroid 2000 SG344 is chosen as nominal target due to the repeatability of its porkchop pattern and the needed low propellant. States for both $L_2$ and Asteroid are retrieved using SPICE kernels, while target covariance at the final time is retrieved by exploiting a telnet connection with NASA Horizon

**Algorithm 2:** Numerical solution of the common-control open-loop problem

**Input:** Problem data, uncertainty ensemble, $\epsilon$-schedule, $\tau$-annealing schedule, penalty-growth factor, CVaR tolerance
**Output:** Open-loop value $J_2$, stacked initial costates $\mathbf{P}_0$, terminal violations, fuel consumption

```
Generate or reuse the uncertainty ensemble {ω_s, w_s}_{s=1}^{N_s}
Initialize the stacked costate vector P_0
Set the penalty multiplier m_γ = 1
for each penalty-continuation level do
    Scale γ_r and γ_v by m_γ
    for each smoothing level τ do
        Compute η_r and η_v from the smoothed VaR conditions
        Initialize J_2^old = ∞
        for outer iteration = 1, ..., N_max do
            for each ϵ in the continuation schedule do
                Propagate the coupled ensemble state-costate system
                At each time, reconstruct the common u* and α* from the ensemble
                  Hamiltonian
                Solve the stacked shooting system for P_0
                Warm-start the next ϵ level from the current solution
            end
            Update d_r^(s), d_v^(s), and y^(s)(t_f) for all samples
            Evaluate the smoothed CVaRs and the smoothed objective J_2
            if the relative change in J_2 is below tolerance then
                break
            end
            Update η_r and η_v by damped smoothed-VaR iteration
        end
    end
    Evaluate the hard CVaRs from the nonsmoothed definitions
    if CVaR_{k_r}(d_r) ≤ 0 and CVaR_{k_v}(d_v) ≤ 0 then
        break
    end
    Increase m_γ
end
Compute the final hard-metric objective J_2
return J_2, P_0, terminal violations, and fuel values
```

**Table 1. Spacecraft data used in the numerical test.**

| Quantity | Value |
|---|---|
| Initial mass, $m_0$ | 22.6 kg |
| Maximum thrust, $T$ | 2.0 mN |
| Specific impulse, $I_{sp}$ | 3000 s |

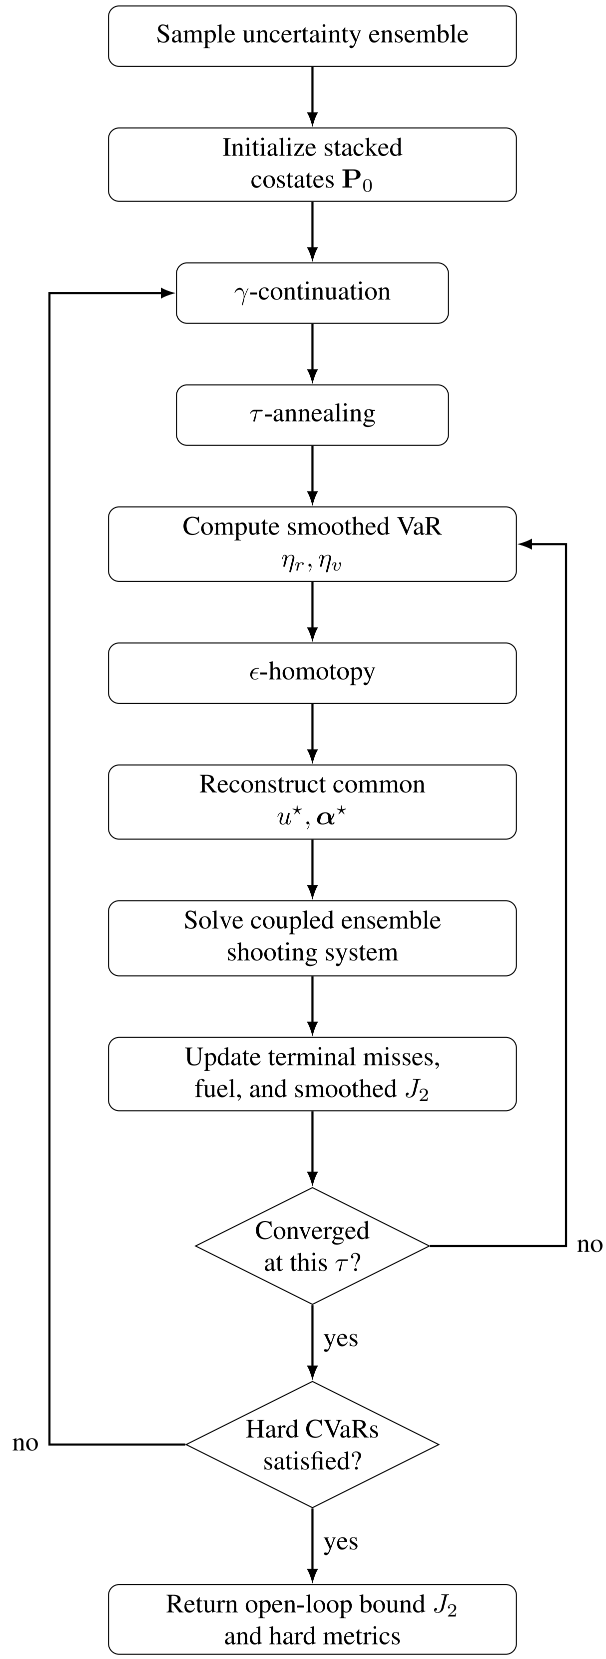


**Figure 2. Numerical workflow used to compute the common-control open-loop bound.**

system.[27]

**Table 2. Orbital elements for the asteroid 2000 SG344.**

| Name | a [AU] | e [-] | i [deg] | $\omega$ [deg] | $\Omega$ [deg] |
|---|---|---|---|---|---|
| 2000 SG344 | 0.9775 | 0.0669 | 0.1121 | 275.3026 | 191.9599 |

**Table 3. Characteristics of baseline solution for M-ARGO.**

| **Name** | **Departure Date** | **TOF [d]** | $m_p$ [kg] |
|---|---|---|---|
| 2000 SG344 | 09 NOV 2023 | 900 | 0.857 |

The uncertainty and terminal-risk parameters are summarized in Table 4. The initial position and velocity errors are modeled as independent zero-mean Gaussian variables in Cartesian components. The spacecraft mass is assumed deterministic at departure. The thrust-magnitude error and the two pointing-error components are modeled as independent Gauss–Markov processes. The two pointing components define the sample-dependent attitude-error rotation used in the stochastic thrust-direction model.

**Table 4. Uncertainty and terminal-risk parameters.**

| **Quantity** | **Value** |
|---|---|
| Initial position $1\sigma$ error, per component | 100 km |
| Initial velocity $1\sigma$ error, per component | 0.1 m/s |
| Thrust-magnitude stationary standard deviation | $1\%$ |
| Pointing-component stationary standard deviation | 1 deg |
| Thrust-magnitude correlation time | 1 d |
| Pointing-error correlation time | 1 d |
| Position tolerance, $R$ | 2500 km |
| Velocity tolerance, $V$ | 5 m/s |
| Position probability level, $k_r$ | 0.9 |
| Velocity probability level, $k_v$ | 0.9 |
| Initial CVaR penalties, $(\gamma_r, \gamma_v)$ | $(50,\ 50)$ |
| Number of samples, $N_s$ | 100 |

### Oracle and open-loop solutions

The two optimization problems are solved on the same sampled ensemble. The pathwise oracle problem allows each realization to use its own control history and therefore provides the lower bound $J_1$. The common-control open-loop problem enforces a unique commanded throttle and thrust-direction history across the ensemble and provides the implementable open-loop bound $J_2$. Since the admissible set of the common-control problem is a subset of the oracle admissible set, the expected ordering is

$$J_1 \leq J_2. \tag{73}$$

The main numerical quantities used to compare the two solutions are summarized in Table 5. The fuel is reported both as a fraction of the initial mass and in kg. The CVaR quantities are evaluated using the nonsmoothed terminal violations, after the smoothing and penalty continuations have converged. A non-positive value of the terminal CVaR indicates satisfaction of the corresponding sample-based chance constraint.

**Table 5. Comparison between the pathwise oracle and the common-control open-loop solution.**

| **Solution** | $J$ | $\bar{m}_p/m_0$ | $\bar{m}_p$ [kg] | $\mathrm{CVaR}_{k_r}(d_r)$ | $\mathrm{CVaR}_{k_v}(d_v)$ | $P_{\mathrm{hit}}$ |
|---|---|---|---|---|---|---|
| Pathwise oracle | 0.04438 | 0.04582 | 1.03558 | $-1.1937 \times 10^{-10}$ | $-2.8047 \times 10^{-8}$ | (1.0, 1.0) |
| Open-loop | 0.05748 | 0.05478 | 1.23802 | $-2.8857 \times 10^{-6}$ | $-1.5749 \times 10^{-10}$ | (0.90, 0.94) |
| Gap, $J_2 - J_1$ | 0.01310 | 0.00896 | 0.20244 | – | – | – |

For clarity, the hit probability in Table 5 can be reported as a pair,

$$P_{\mathrm{hit}} = \left(P\left(\|\mathbf{r}(t_f) - \mathbf{r}_f\| \leq R\right), P\left(\|\mathbf{v}(t_f) - \mathbf{v}_f\| \leq V\right)\right), \tag{74}$$

computed from the same finite ensemble used in the optimization. Since the probability estimate is sample-based, it should be interpreted together with the adopted value of $N_s$.
Figure 3 shows the optimized heliocentric trajectories. In the oracle case, each realization has its own pathwise optimal thrust history. In the open-loop case, all realizations are propagated under the same commanded throttle and thrust direction. Therefore, the terminal dispersion in the open-loop solution directly reflects the effect of uncertainty on a single precomputed control history.

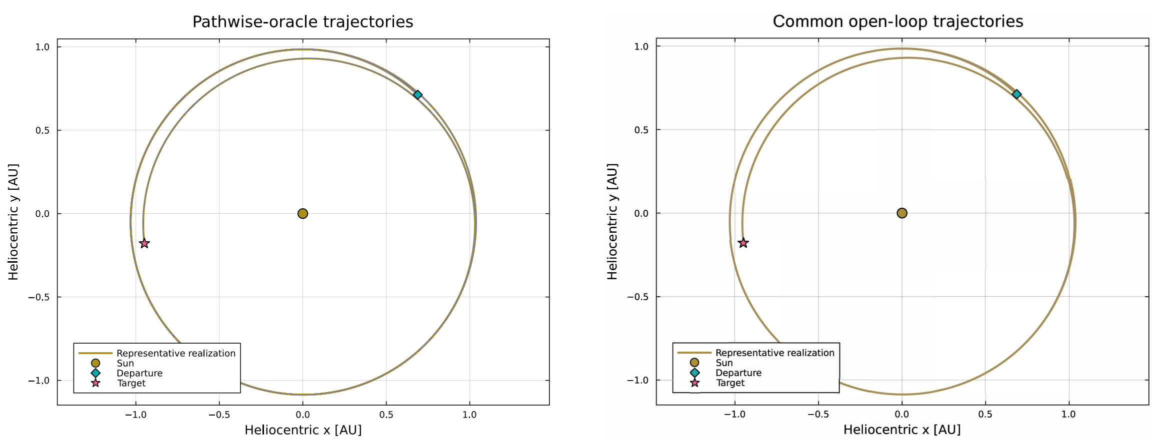


(a) Pathwise oracle (b) Common-control open-loop

**Figure 3. Optimized heliocentric trajectories for the pathwise oracle and common-control open-loop solutions. The markers at the final time indicate the terminal states of the sampled realizations.**

Figure 4 reports the corresponding control histories. For the oracle solution, a representative sample is shown, because each realization has its own throttle and thrust-direction history. For the open-loop solution, the plotted control is the single control history shared by all realizations. The switching functions are also shown to illustrate the bang–bang structure recovered as $\epsilon$ approaches $\epsilon_f$.

The terminal dispersion is shown in Fig. 5. The dashed circles represent the admissible position and velocity balls used in the chance constraints. The comparison between the oracle and open-loop terminal clouds highlights the role of feedback-like pathwise compensation in the oracle solution. Since the oracle can tailor the control to each realization, it is expected to achieve a lower fuel consumption and a tighter terminal distribution. The open-loop solution instead trades additional

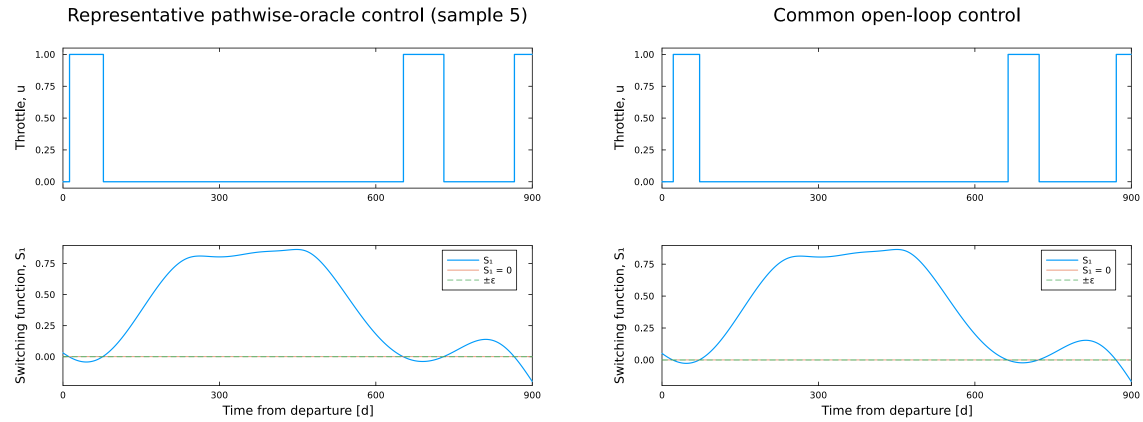


(a) Representative oracle sample

(b) Common open-loop control

**Figure 4. Throttle histories and switching functions for the oracle and open-loop solutions. The oracle plot refers to one representative realization, whereas the open-loop plot reports the common control applied to the full ensemble.**

fuel and terminal dispersion against the requirement that all samples share the same commanded control history.

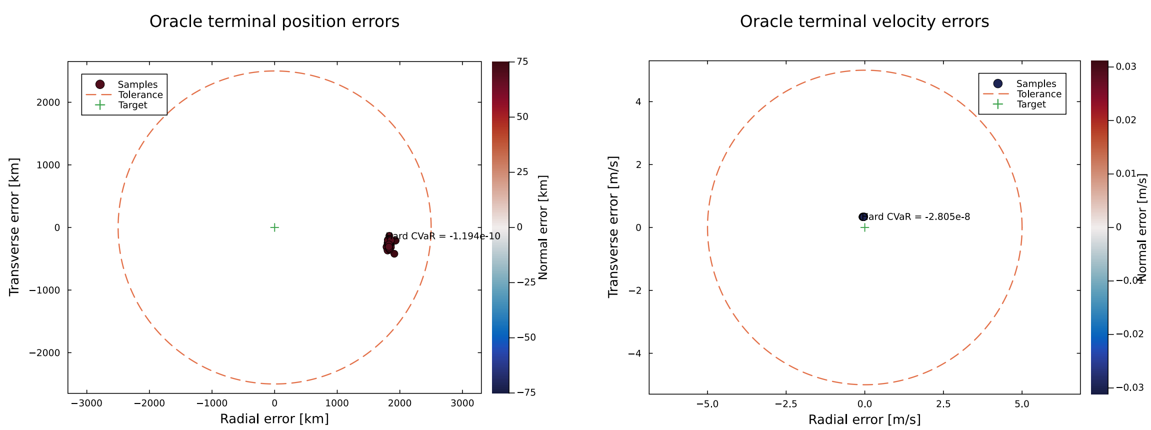


(a) Terminal position error

(b) Terminal velocity error

**Figure 5. Terminal error clouds for the optimized ensemble. The admissible balls correspond to the tolerances $R$ and $V$ used in the chance constraints.**

The results are expected to satisfy the ordering in Eq. (73). The size of the gap $J_2 - J_1$ quantifies the loss associated with replacing pathwise, realization-dependent controls with a single open-loop command history. A small gap indicates that the uncertainty can be handled with limited additional cost by a common control history, whereas a large gap indicates that the oracle exploits sample-dependent compensation in a substantial way. In this sense, the proposed comparison provides not only a risk-aware trajectory design, but also a diagnostic measure of how valuable in-flight correction or feedback would be for the considered transfer.

## CONCLUSIONS

This work presented an indirect approach to low-thrust minimum-fuel trajectory optimization under initial-state and control-execution uncertainties. The spacecraft dynamics were modeled as a random ordinary differential equation, with thrust-magnitude and pointing errors represented through Gauss–Markov processes. Terminal position and velocity requirements were expressed as chance constraints and treated through finite-sample CVaR penalties, thereby incorporating terminal risk directly into the trajectory optimization problem.

Two related optimal-control problems were formulated. The pathwise oracle problem allows each uncertainty realization to employ an independent control history and therefore provides a perfect-information lower bound on the achievable risk-aware performance. The common-control open-loop problem instead imposes a single throttle and commanded thrust-direction history on the entire ensemble, providing an implementable open-loop solution. Pontryagin's Maximum Principle was applied to both formulations, yielding the corresponding adjoint equations, optimal thrust directions, and bang–bang switching functions. The resulting boundary-value problems were solved through continuation in the throttle regularization, CVaR smoothing parameters, and terminal-risk penalty coefficients.

The methodology was demonstrated on a representative 900-day M-ARGO-inspired transfer to asteroid 2000 SG344 using an ensemble of 100 uncertainty realizations. Both solutions satisfied the sample-based terminal-risk requirements.

The results confirm the expected ordering $J_1 \leq J_2$ and show that the difference between the oracle and common-control solutions provides a quantitative measure of the value of realization-dependent correction. The proposed framework therefore offers both a risk-aware trajectory-design method and a diagnostic tool for assessing whether a purely open-loop strategy is adequate or whether in-flight guidance corrections are likely to provide a significant benefit.

Future work will focus on extending the formulation to closed-loop and receding-horizon guidance, so that measurements acquired during the transfer can be used to update the commanded control. The numerical treatment can also be improved through multiple shooting to increase robustness for longer transfers and larger uncertainty ensembles. Additional developments will consider richer dynamical models, navigation uncertainties, and alternative risk measures for both terminal accuracy and propellant consumption.


## ACKNOWLEDGMENT

This work was carried out within the Space It Up! project funded by the Italian Space Agency (ASI) and the Italian Ministry of University and Research (MUR) under contract No. 2024-5-E.0, CUP No. I53D24000060005.


## REFERENCES


[1] W. Fehse, *Automated rendezvous and docking of spacecraft*, Vol. 16. Cambridge University Press, 2003.

[2] A. Poghosyan and A. Golkar, "CubeSat evolution: Analyzing CubeSat capabilities for conducting science missions," *Progress in Aerospace Sciences*, Vol. 88, 2017, pp. 59–83, 10.1016/j.paerosci.2016.11.002.

[3] R. Walker, D. Binns, C. Bramanti, M. Casasco, P. Concari, D. Izzo, D. Feili, P. Fernandez, J. G. Fernandez, P. Hager, *et al.*, "Deep-space CubeSats: thinking inside the box," *Astronomy & Geophysics*, Vol. 59, No. 5, 2018, pp. 24–30, 10.1093/astrogeo/aty232.

[4] R. S. Park and D. J. Scheeres, "Nonlinear mapping of Gaussian statistics: Theory and applications to spacecraft trajectory design," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 6, 2006, pp. 1367–1375, 10.2514/1.20177.

[5] P. Di Lizia, R. Armellin, A. Ercoli-Finzi, and M. Berz, "High-order robust guidance of interplanetary trajectories based on differential algebra," *Journal of Aerospace Engineering, Sciences and Applications*, Vol. 1, No. 1, 2008, pp. 43–57.

[6] P. Di Lizia, R. Armellin, F. Bernelli-Zazzera, and M. Berz, "High order optimal control of space trajectories with uncertain boundary conditions," *Acta Astronautica*, Vol. 93, 2014, pp. 217–229, 10.1016/j.actaastro.2013.07.007.

[7] P. Di Lizia, R. Armellin, A. Morselli, and F. Bernelli-Zazzera, "High order optimal feedback control of space trajectories with bounded control," *Acta Astronautica*, Vol. 94, No. 1, 2014, pp. 383–394, 10.1016/j.actaastro.2013.02.011.

[8] P. W. Schumacher Jr, C. Sabol, C. C. Higginson, and K. T. Alfriend, "Uncertain Lambert Problem," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 9, 2015, pp. 1573–1584, 10.2514/1.G001019.

[9] G. Zhang, D. Zhou, D. Mortari, and M. R. Akella, "Covariance analysis of Lambert's problem via Lagrange's transfer-time formulation," *Aerospace Science and Technology*, Vol. 77, 2018, pp. 765–773, 10.1016/j.ast.2018.03.039.

[10] N. Adurthi and M. Majji, "Uncertain Lambert Problem: A Probabilistic Approach," *The Journal of the Astronautical Sciences*, Vol. 67, 2020, pp. 361–386, 10.1007/s40295-019-00205-z.

[11] R. Armellin, P. Di Lizia, F. Topputo, M. Lavagna, F. Bernelli-Zazzera, and M. Berz, "Gravity Assist Space Pruning based on Differential Algebra," *Celestial mechanics and dynamical astronomy*, Vol. 106, No. 1, 2010, pp. 1–24, 10.1007/s10569-009-9235-0.

[12] H.-y. Li, Y.-Z. Luo, G.-J. Tang, *et al.*, "Optimal multi-objective linearized impulsive rendezvous under uncertainty," *Acta Astronautica*, Vol. 66, No. 3-4, 2010, pp. 439–445, 10.1016/j.actaastro.2009.06.019.

[13] Z. Yang, Y.-z. Luo, and J. Zhang, "Robust Planning of Nonlinear Rendezvous with Uncertainty," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 8, 2017, pp. 1954–1967, 10.2514/1.G002319.

[14] M. Balducci and B. A. Jones, "Asteroid Rendezvous Maneuver Design Considering Uncertainty," *Advances in the Astronautical Sciences* (Univelt, ed.), Vol. 168, 2019, pp. 2951–2967.

[15] F. Xiong, Y. Xiong, and B. Xue, "Trajectory Optimization under Uncertainty based on Polynomial Chaos Expansion," *AIAA Guidance, Navigation, and Control Conference*, 2015, p. 1761, 10.2514/6.2015-1761.

[16] C. Greco, S. Campagnola, and M. L. Vasile, "Robust Space Trajectory Design using Belief Stochastic Optimal Control," *AIAA Scitech 2020 Forum*, 2020, p. 1471, 10.2514/6.2020-1471.

[17] N. Ozaki, S. Campagnola, and R. Funase, "Tube stochastic optimal control for nonlinear constrained trajectory optimization problems," *Journal of Guidance, Control, and Dynamics*, 2020, 10.2514/1.G004363.

[18] N. Marmo, A. Zavoli, N. Ozaki, and Y. Kawakatsu, "A hybrid multiple-shooting approach for covariance control of interplanetary missions with navigation errors," *33rd AAS/AIAA Space Flight Mechanics Meeting*, Austin, TX, 2023.

[19] B. Benedikter, A. Zavoli, Z. Wang, S. Pizzurro, and E. Cavallini, "Convex Approach to Covariance Control with Application to Stochastic Low-Thrust Trajectory Optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 11, 2022, pp. 2061–2075, 10.2514/1.G006806.

[20] J. Ridderhof, P. Tsiotras, and B. J. Johnson, "Stochastic entry guidance," *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 2, 2022, pp. 320–334, 10.2514/1.G005964.

[21] C. Giordano and F. Topputo, "Analysis, Design, and Optimization of Robust Trajectories in Cislunar Environment for Limited-Capability Spacecraft," *The Journal of the Astronautical Sciences*, Vol. 70, No. 6, 2023, p. 53.

[22] B. Schutz, B. Tapley, and G. H. Born, *Statistical orbit determination*. Elsevier, 2004.

[23] B. A. Jones, N. Parrish, and A. Doostan, "Postmaneuver collision probability estimation using sparse polynomial chaos expansions," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 8, 2015, pp. 1425–1437.

[24] C. R. Gates, "A simplified model of midcourse maneuver execution errors," tech. rep., 1963.

[25] R. T. Rockafellar, S. Uryasev, *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, Vol. 2, 2000, pp. 21–42.

[26] R. Walker, D. Koschny, C. Bramanti, *et al.*, "Miniaturised Asteroid Remote Geophysical Observer (M-ARGO): A stand-alone deep space CubeSat system for low-cost science and exploration missions," *6th Interplanetary CubeSat Workshop, Cambridge, UK*, 2017, pp. 1–20.

[27] J. Giorgini and D. Yeomans, "On-line system provides accurate ephemeris and related data," tech. rep., NASA, 1999.