Tail-Risk-Aware Information Guidance for Angles-Only Spacecraft Rendezvous Under Markovian Measurement Outages

Archit Das [*]

September 17, 2026

Abstract

Angles-only spacecraft rendezvous couples guidance and navigation because trajectory geometry changes both observability and measurement availability. This paper develops tail-risk-aware guidance for geometry-dependent, burst-correlated optical outages. A planar benchmark uses explicit Markov measurement-arrival scenarios and conditional value at risk to shape terminal navigation uncertainty, followed by independent 10,000-run validation. Relative to expected-availability guidance, the selected trajectory reduces 95% conditional value at risk by 39.2% and 99% conditional value at risk by 15.7%, at 35.8% higher maneuver expenditure. A separate 200-run nonlinear three-dimensional holdout uses two-body-plus-oblateness truth dynamics, azimuth/elevation sensing, maneuver errors, unscented Kalman filtering, and covariance-weighted terminal feedback. The risk-aware trajectory increases realized measurement availability from 20.7% to 56.1%, reduces mean longest outage from 14.39 to 6.80 epochs, and reduces median and 95th-percentile terminal estimation error from 49.64 m to 0.89 m and 223.24 m to 21.93 m, respectively, while mission success remains statistically comparable. The demonstrated operational payoff is navigation robustness, not superior terminal accuracy. Because approach-corridor constraints are monitored but not enforced, the study does not establish safety-compliant proximity operations.

[*]Independent Researcher, Avionics & GNC Engineering.

Nomenclature

$\boldsymbol{x}_p = [x, y, \dot{x}, \dot{y}]^T$	planar relative state in the target LVLH frame
$\boldsymbol{x}_r = [x_R, y_T, z_N, v_R, v_T, v_N]^T$	three-dimensional RTN relative state
$\Delta\boldsymbol{v}_k$	impulsive relative-velocity command at maneuver k
F	discrete Hill–Clohessy–Wiltshire state-transition matrix
P_k	navigation error covariance after the measurement update
H_k	bearing or angle-measurement Jacobian
γ_k	binary measurement-arrival variable
$p_{\text{geo}}(\boldsymbol{x}_k)$	geometry-dependent optical measurement probability
$s_k \in \{G, B\}$	hidden Markov visibility state (good/bad)
T	Markov visibility transition matrix
β_G, β_B	visibility-state multipliers on measurement probability
L_j	terminal 3σ position-major-axis loss in scenario j
VaR_α	value at risk at confidence level α
CVaR_α	conditional value at risk at confidence level α
α	CVaR confidence level (0.95 in the primary study)
$\boldsymbol{r}_c, \boldsymbol{r}_d$	chief and deputy inertial position vectors
$\boldsymbol{q}$	deputy-to-chief optical line-of-sight vector
P_r	RTN position-covariance submatrix
ϵ_{NEES}	normalized estimation error squared
ϵ_{NIS}	normalized innovation squared

I Introduction

Autonomous rendezvous and proximity operations require guidance, navigation, and control to be designed as an integrated system rather than as independent subsystems. The classical Hill–Clohessy–Wiltshire model remains a useful first-order description of relative motion near a circular reference orbit [1], while decades of rendezvous research have established the importance of covariance propagation, passive safety, and autonomous mission planning [2–4]. In an angles-only architecture, however, the guidance problem is more tightly coupled to navigation than in a range-rich sensor suite because the relative trajectory directly determines observability of range and range rate.

This coupling has motivated a substantial literature on observability-aware rendezvous. Woffinden and Geller derived optimal rendezvous maneuvers for angles-only navigation [5], while Grzymisch and Fichter developed analytic observability maneuvers and subsequently embedded enhanced bearings-only observability into optimal rendezvous guidance [6, 7]. Chu et al. optimized bearing-only rendezvous trajectories by combining fuel consumption with covariance- and Fisher-information-based observability indices [8], and Mok et al. developed a one-step guidance law that adds a Fisher-information observability term to closed-loop rendezvous guidance [9]. More recently, Fujiwara and Funase incorporated Fisher information into differential dynamic programming for impulsive observability-aware maneuver design [10]. These contributions establish that guidance should deliberately alter geometry when doing so improves state estimation.

A separate line of work has increased the realism of rendezvous uncertainty models. Linear covariance analysis has long been used to connect navigation errors, trajectory dispersions, and maneuver design [2]. Kalman filtering with intermittent observations established that random measurement loss can fundamentally alter covariance behavior [11]; Markovian loss models further show that burst correlation changes covariance stability and peak-error behavior [26, 27]. Stochastic rendezvous trajectory optimization now includes nonlinear chance constraints on collision avoidance, sensor field of view, docking corridors, and control magnitude [12]. Ra and Oguri recently combined chance-constrained spacecraft autonomy with sensing-optimal angles-only path planning in CWH and cislunar dynamics [13]. Chance-constrained visual-servoing MPC has also been developed for non-cooperative rendezvous under model, disturbance, and sensor uncertainty [14]. Meanwhile, delayed and asynchronous relative-navigation measurements have been addressed directly in filtering architectures and hardware-in-the-loop experiments [15].

These two lines of work leave a narrower intermediate problem. Optical measurements are not only noisy; they may be unavailable. Camera field of view, target occlusion, illumination, bright-body exclusion, target feature quality, and onboard processing can make availability strongly geometry dependent. This issue is not absent from prior rendezvous work: You et al. included a scheduled 30% eclipse data gap in covariance-triggered angles-only rendezvous planning [16], and Lee et al. quantified relative-navigation degradation under intermittent laser measurements [17]. Image-based rendezvous studies have enforced deterministic visibility and occlusion constraints [18, 19], while

Kruger and D'Amico incorporated partial measurement-availability constraints into angles-only observability analysis and distributed-system design [20]. Deole and Mesbahi have further developed estimation-aware trajectory optimization with state-dependent, set-valued measurement uncertainty, including an uncooperative rendezvous example [21]. Very recent work has embedded covariance dynamics directly in MPC for satellite inspection [22], and robust beacon-placement work has optimized P99 relative-navigation performance under line-of-sight and field-of-view measurement loss [23]. Outside spacecraft guidance, active-sensing methods have explicitly studied trajectory optimization under intermittent measurements [24]. Optical-navigation literature also notes that bright-body exclusion angles are an operational constraint, with Sun-exclusion angles on the order of tens of degrees being common for optical sensors [25].

The research question of this paper is therefore narrower than generic observability-aware guidance:

> *How should an angles-only rendezvous trajectory change when measurement availability is both geometry dependent and burst correlated, and when guidance is required to reduce the severity of rare navigation failures rather than only expected covariance?*

The study proceeds in two stages. First, an expected-Markov-availability baseline propagates a differentiable fractional-information surrogate. Second, the proposed planner samples full Markov outage sequences and optimizes the terminal uncertainty distribution using a CVaR-based scenario objective. CVaR is used because it quantifies loss beyond a percentile threshold and admits direct scenario-based optimization [28].

The contributions are: (1) a rendezvous sensing model in which a smooth state-dependent optical availability probability is modulated by a burst-correlated two-state Markov process; (2) a scenario-based risk-sensitive guidance objective that penalizes mean terminal uncertainty, the 95th percentile, and 95%-CVaR under explicit binary measurement arrivals; (3) a terminal-null-space multistart optimization architecture that separates design-scenario randomness from an independent fixed validation set; (4) a 10,000-realization planar comparison that quantifies the maneuver-effort versus navigation-tail-risk trade; and (5) an independent nonlinear three-dimensional closed-loop holdout validation using a UKF, maneuver execution errors, covariance-gated terminal feedback, safety diagnostics, and common random numbers across all guidance methods. The numerical study deliberately distinguishes navigation-tail reduction from closed-loop terminal-state accuracy rather than treating expected information as a safety guarantee.

Scope and operational interpretation. The primary claim of this paper is deliberately limited to *navigation resilience under intermittent sensing*. In an angles-only system, preventing long measurement outages and suppressing heavy estimation-error tails is operationally valuable because it reduces the probability that downstream guidance, abort logic, or safety monitors must act on a poorly localized vehicle. The present experiments do not establish that the proposed trajectory is terminal-error optimal, nor do they establish safety-compliant proximity operations. In particular, the nonlinear holdout treats the keep-out zone as part of the mission-success criterion but monitors the approach corridor only diagnostically. The latter limitation is retained explicitly so that improved navigation robustness is not conflated with a complete safety guarantee.

II Literature Review and Research Gap

Table 1 summarizes the literature most directly related to the proposed formulation. The distinguishing feature of the present paper is not the use of an observability metric, stochastic uncertainty, or visibility constraints individually. Rather, it is the use of a *state-dependent probability of measurement availability inside the guidance objective*, followed by evaluation under random dropout sequences.

Table 1: Representative literature and the remaining gap addressed here.

Study	Main contribution	Relation to the present problem
Woffinden and Geller [5]; Grzymisch and Fichter [6, 7]	Foundational maneuver design and optimal guidance for enhanced angles-/bearings-only observability	Establish that rendezvous maneuvers should shape sensing geometry; do not model stochastic state-dependent measurement arrival.
You et al. [16]; Lee et al. [17]	Rendezvous/navigation under eclipse data gaps or intermittent ranging/angle measurements	Demonstrate that outages matter, but availability is imposed as an external schedule/outage rather than optimized as a smooth function of rendezvous geometry.
Chu et al. [8]; Mok et al. [9]; Fujiwara and Funase [10]	Covariance/Fisher-information trajectory optimization and observability-aware guidance	Optimize information geometry, but stochastic measurement arrival is not the central decision-dependent quantity.
Zhao and Zhang [18]; Lin et al. [19]	Image-based planning/control with field-of-view and occlusion constraints	Model sensor visibility primarily through deterministic feasibility constraints rather than an arrival probability in the information objective.
Kruger and D'Amico [20]	Angles-only observability analysis/design including measurement-availability constraints	Includes partial visibility, but optimizes distributed-system configuration rather than a chaser rendezvous trajectory under stochastic arrivals.
Deole and Mesbahi [21]; Clees et al. [22]	State-dependent estimation-aware trajectory optimization and covariance-aware MPC	Strongly adjacent estimation-control coupling; neither directly studies stochastic geometry-dependent measurement arrival in angles-only rendezvous.
Zhang et al. [12]; Zhao et al. [14]	Stochastic/chance-constrained rendezvous under state, disturbance, and sensor uncertainty	Treat uncertainty and safety rigorously, but do not optimize the random availability of the relative-navigation measurement itself.
Ra and Oguri [13]	Chance-constrained sensing-optimal angles-only path planning in CWH and cislunar dynamics	Closely couples probabilistic safety and active sensing; the present study isolates stochastic, geometry-dependent measurement arrival as the guidance-dependent quantity.
Bommena and Woollands [23]	Robust beacon-layout optimization under LOS/FOV loss using a P99 navigation metric	Closest recent tail-risk geometry co-design, but the design variable is beacon placement over prescribed trajectories rather than rendezvous guidance.
Huang and Dey [26]; Mo and Sinopoli [27]	Kalman filtering under Markov-correlated measurement losses and analysis of covariance tails	Motivate burst-correlated outage modeling and explicit tail metrics; do not optimize spacecraft guidance geometry.
Rockafellar and Uryasev [28]	Scenario-based conditional value-at-risk for general loss distributions	Provides the tail-risk measure used here to penalize severe terminal navigation outcomes.

This review suggests a defensible but specific gap rather than a claim that intermittent sensing or estimation-aware planning is new. A rendezvous planner can account for scheduled outages, covariance, or hard visibility constraints and still be over-optimistic if the probability of obtaining a useful bearing varies continuously with the trajectory itself. The present model therefore treats measurement arrival as a smooth stochastic function of optical geometry and optimizes the chaser trajectory against that probability. The resulting claim is intentionally narrower than the adjacent 2026 estimation-aware and robust-navigation co-design literature.

III Problem Formulation

A Relative Dynamics

The numerical study uses planar Hill–Clohessy–Wiltshire dynamics about a circular chief orbit of radius r_c. With $n = \sqrt{\mu/r_c^3}$, the continuous relative state satisfies

$$\dot{\boldsymbol{x}} = A\boldsymbol{x}, \tag{1}$$

with

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 3n^2 & 0 & 0 & 2n \\ 0 & 0 & -2n & 0 \end{bmatrix}. \tag{2}$$

Impulsive maneuvers act on the velocity states. For a sample period Δt,

$$\boldsymbol{x}_{k+1} = F\boldsymbol{x}_k + B_d \Delta \boldsymbol{v}_k, \qquad F = e^{A\Delta t}, \tag{3}$$

where $B_d = F[\mathbf{0}_{2\times 2}\ I_2]^T$ for an impulse applied at the beginning of the interval.

The benchmark begins at

$$\boldsymbol{x}_0 = [0, -1000, 0, 0]^T \tag{4}$$

and terminates at the 50 m V-bar hold point

$$\boldsymbol{x}_f = [0, -50, 0, 0]^T. \tag{5}$$

Twenty maneuver opportunities are separated by 60 s, giving a 20 min transfer.

B Bearing-Only Measurement Model

The relative-navigation sensor returns a planar line-of-sight angle

$$z_k = h(\boldsymbol{x}_k) + v_k, \qquad h(\boldsymbol{x}_k) = \operatorname{atan2}(y_k, x_k), \tag{6}$$

where $v_k \sim \mathcal{N}(0, R)$. The linearized measurement matrix is

$$H_k = \begin{bmatrix} -y_k/\rho_k^2 & x_k/\rho_k^2 & 0 & 0 \end{bmatrix}, \qquad \rho_k^2 = x_k^2 + y_k^2. \tag{7}$$

The bearing noise standard deviation in the benchmark is 0.05 deg. The initial covariance is

$$P_0 = \operatorname{diag}(150^2, 150^2, 0.2^2, 0.2^2), \tag{8}$$

with position in meters and velocity in meters per second. White acceleration uncertainty is represented by $\sigma_a = 2 \times 10^{-5}$ m/s^2 over each interval.

C Geometry-Dependent Intermittent Sensing

Measurement reception is modeled by

$$\gamma_k \sim \text{Bernoulli}(p_k), \tag{9}$$

where $\gamma_k = 1$ indicates that a usable bearing is obtained. The probability p_k is a smooth function of the angular separation α_k between the target line of sight and a fixed bright-body direction. A logistic transition is used:

$$p_k = p_{\min} + (p_{\max} - p_{\min}) \left[1 + \exp\left(-\frac{\alpha_k - \alpha_{\text{ex}}}{s_\alpha} \right) \right]^{-1}. \tag{10}$$

This provides a differentiable approximation to a probabilistic transition near an exclusion cone, where image usability degrades because of stray light, target contrast, or feature-processing failures. The nominal study uses $p_{\min} = 0.02$, $p_{\max} = 0.98$, $\alpha_{\text{ex}} = 15$ deg, transition width $s_\alpha = 1.5$ deg, and a Sun azimuth of 72 deg in the LVLH plane.

IV Risk-Aware Guidance Formulation

A Expected Markov Availability Baseline

The smooth geometry-only probability from Eq. (10) is denoted $p_{\mathrm{geo}}(\boldsymbol{x}_k)$. To represent burst-correlated visibility degradation, a hidden state $s_k \in \{G, B\}$ evolves according to

$$T = \begin{bmatrix} 0.97 & 0.03 \\ 0.20 & 0.80 \end{bmatrix}, \tag{11}$$

where G and B denote good and bad visibility. Conditional on state s_k, the binary optical measurement arrives with probability

$$q_k = p_{\mathrm{geo}}(\boldsymbol{x}_k)\,\beta_{s_k}, \qquad \beta_G = 1, \quad \beta_B = 0.05, \tag{12}$$

and $\gamma_k \sim \mathrm{Bernoulli}(q_k)$. The bad-state persistence $T_{BB} = 0.80$ creates sustained outages rather than independent missed detections, consistent with the qualitative behavior of Markov-correlated intermittent observations [26, 27].

For a differentiable baseline planner, the Markov-state occupancy distribution $\boldsymbol{\pi}_k = [\Pr(G), \Pr(B)]$ is propagated as $\boldsymbol{\pi}_{k+1} = \boldsymbol{\pi}_k T$, initialized in the good state. The expected availability is

$$\bar{q}_k = p_{\mathrm{geo}}(\boldsymbol{x}_k)\left(\pi_{G,k} + 0.05\pi_{B,k}\right). \tag{13}$$

After covariance prediction $P_k^- = FP_{k-1}F^T + Q$, the fractional-information update is

$$P_k = P_k^- - \frac{P_k^- H_k^T H_k P_k^-}{H_k P_k^- H_k^T + R/\bar{q}_k}. \tag{14}$$

This expected-Markov update is a planning surrogate, not $E[P_k]$. The continuous-sensing information-aware baseline uses the same recursion with $\bar{q}_k = 1$.

The first three planners minimize a smoothed total maneuver cost plus a terminal log-determinant information term,

$$J_{\mathrm{info}} = \sum_{k=0}^{N-1} \sqrt{\Delta\boldsymbol{v}_k^T \Delta\boldsymbol{v}_k + \epsilon} + \lambda \log\det(P_{r,N}), \tag{15}$$

with $\lambda = 0.8$, terminal equality $\boldsymbol{x}_N = \boldsymbol{x}_f$, and componentwise impulse bounds of $\pm 0.5\,\mathrm{m/s}$.

B Scenario-Based CVaR Guidance

The proposed risk-sensitive planner uses explicit binary measurement realizations during trajectory design. For design scenario j, a Markov state sequence $s_{j,1:N}$ and independent uniform variates are generated once and held fixed during optimization. The arrival sequence follows Eq. (12). For each arrival realization, covariance is propagated through the intermittent Riccati recursion

$$P_{j,k} = P_{j,k}^{-} - \gamma_{j,k} P_{j,k}^{-} H_k^T \left(H_k P_{j,k}^{-} H_k^T + R \right)^{-1} H_k P_{j,k}^{-}. \tag{16}$$

The scenario loss is the terminal 3σ major axis of position uncertainty,

$$L_j = 3\sqrt{\lambda_{\max}\left(P_{r,N}^{(j)}\right)}. \tag{17}$$

With M fixed design scenarios and $\alpha = 0.95$, the empirical tail measure is

$$\mathrm{CVaR}_\alpha(L) = \frac{1}{|\mathcal{T}_\alpha|} \sum_{j \in \mathcal{T}_\alpha} L_j, \quad \mathcal{T}_\alpha = \{ j : L_j \geq \mathrm{VaR}_\alpha(L) \}. \tag{18}$$

The dimensionless numerical objective is

$$\bar{J} = \frac{\Delta v}{1\,\mathrm{m/s}} + 0.01\,\overline{\left(\frac{L}{1\,\mathrm{m}}\right)} + 0.03\ \mathrm{VaR}_{0.95}\left(\frac{L}{1\,\mathrm{m}}\right) + 0.04\ \mathrm{CVaR}_{0.95}\left(\frac{L}{1\,\mathrm{m}}\right). \tag{19}$$

The primary design additionally imposes $\Delta v \leq 3.25\,\mathrm{m/s}$.

C Terminal-Feasible Multistart Optimization

Let $\boldsymbol{u}_E$ denote the expected-Markov trajectory and let M_E be the linear endpoint map from the stacked impulsive control sequence to terminal state. Candidate CVaR trajectories are parameterized as

$$\boldsymbol{u} = \boldsymbol{u}_E + B\boldsymbol{c}, \qquad M_E B = 0, \tag{20}$$

so every candidate preserves the terminal equality to numerical precision. Six low-frequency sinusoidal impulse modes are projected into the endpoint null space and scaled to form the columns of B. The coefficient bounds are $c_i \in [-0.08, 0.08]$.

Because explicit binary dropout makes Eq. (19) nonsmooth and nonconvex, differential evolution is used with population size six and a fixed budget of 30 generations. Five independent optimizer restarts use the same 500 design outage scenarios but different population seeds. Among feasible restarts, the minimum design-objective trajectory is selected *before* independent validation. This separation prevents selection on test-set performance.

V Numerical Experiment

A Benchmark and Controllers

The benchmark is a planar HCW rendezvous about a 500 km circular orbit. The chaser begins at $(x, y) = (0, -1000)$ m with zero relative velocity and reaches a hold point at $(0, -50)$ m with zero relative velocity after 20 one-minute maneuver/measurement intervals. Four guidance laws are compared:

1) **Fuel-only**: minimum total impulsive Δv.

2) **Continuous-sensing information-aware**: information-aware guidance with every future bearing assumed available.

3) **Expected-Markov-availability-aware**: differentiable fractional-information planning using the expected Markov visibility probability.

4) **CVaR-Markov-availability-aware**: the proposed scenario-risk formulation in Eq. (19).

Table 2: Primary numerical-study parameters.

Parameter	Value
Chief altitude	500 km
Transfer duration	20 min
Maneuver/measurement interval	60 s
Initial / terminal relative position	$(0, -1000)$ m / $(0, -50)$ m
Bearing noise, σ_θ	0.05 deg
Acceleration uncertainty, σ_a	2×10^{-5} m/s^2
Initial position standard deviation	150 m per axis
Sun azimuth / exclusion half-angle	72 deg / 15 deg
Geometry probability range	$[0.02, 0.98]$
Markov transition matrix	$[[0.97, 0.03], [0.20, 0.80]]$
Bad-state probability multiplier	0.05
CVaR confidence level	0.95
CVaR design scenarios	500
Optimizer restarts	5
Independent validation realizations	10,000
Primary CVaR Δv cap	3.25 m/s

B Independent Validation and Metrics

The four trajectories are evaluated on the same fixed set of 10,000 Markov-state paths and Bernoulli-arrival random variates, generated from a validation seed distinct from the design seed. Validation is performed only after the multistart selection rule has fixed the proposed trajectory. The reported metrics are median, P95 and P99 terminal L; $\mathrm{CVaR}_{0.95}(L)$ and $\mathrm{CVaR}_{0.99}(L)$; $\Pr(L < 2\,\mathrm{m})$; mean measurement availability; mean number of received measurements; and the mean and P95 duration of the longest contiguous dropout burst.

A secondary maneuver-cap sensitivity study uses caps of 2.50, 2.75, 3.00, 3.25 and 3.50 m/s. Lower-cap points use three multistart runs and the two highest-cap points use five. Because the optimization is nonconvex and generation limited, these points are interpreted as a sensitivity study rather than a certified Pareto frontier.

VI Results

A Primary Fixed-Validation Comparison

Table 3 reports the independent 10,000-realization comparison. The fuel-only trajectory has the lowest maneuver cost but produces severe navigation tails under burst-correlated measurement loss. The two information-aware baselines reduce nominal uncertainty substantially. The expected-Markov baseline has a median terminal 3σ major axis of 0.760 m and satisfies the 2 m threshold in 83.46% of realizations, but its extreme tail remains large: CVaR95 is 638.0 m and CVaR99 is 1263.7 m.

The selected CVaR trajectory uses 3.245 m/s, 35.8% more Δv than the expected-Markov baseline. In exchange, P95 decreases by 27.8%, from 51.79 m to 37.41 m; P99 decreases by 4.3%; CVaR95 decreases by 39.2%, from 638.04 m to 387.77 m; and CVaR99 decreases by 15.7%, from 1263.74 m to 1064.88 m. The median changes only from 0.760 m to 0.747 m, and $\Pr(L < 2\,\mathrm{m})$ increases by 1.77 percentage points. Thus the primary benefit is a reduction in the severity of rare failures, not a large change in nominal accuracy.

Table 3: Independent 10,000-realization terminal navigation results.

Method	Δv [m/s]	Median [m]	P95 [m]	P99 [m]	CVaR95 [m]	CVaR99 [m]
Fuel-only	1.879	3.979	769.32	1057.62	1072.99	1458.75
Continuous-sensing information-aware	2.233	1.335	50.26	831.84	682.60	1279.28
Expected-Markov-availability-aware	2.389	0.760	51.79	833.10	638.04	1263.74
CVaR-Markov-availability-aware	3.245	0.747	37.41	796.87	387.77	1064.88

Figure 1 normalizes the principal tail metrics by the expected-Markov baseline. The largest improvement occurs in CVaR95, the quantity most directly targeted by the design objective; the simultaneous reduction in CVaR99 is an out-of-objective benefit.

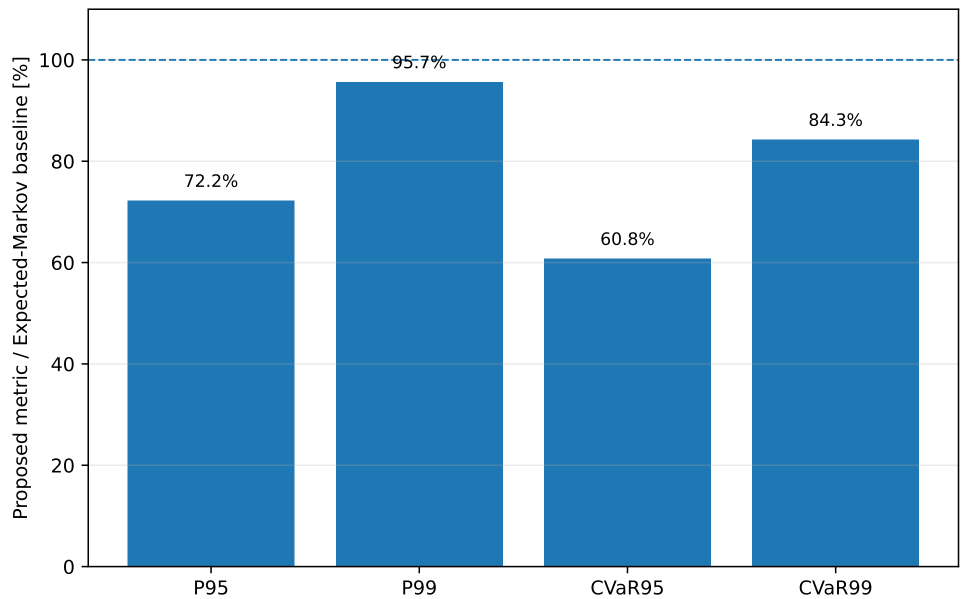

Figure 1: Primary CVaR trajectory tail metrics normalized by the expected-Markov baseline. Values below 100% indicate improvement.

B Measurement Availability and Outage Duration

The risk reduction is accompanied by a substantial change in sensing statistics. Table 4 shows that the proposed trajectory raises mean measurement availability from 0.403 to 0.526, a 30.6% relative increase, and raises the mean received-measurement count from 8.06 to 10.52 of 20 opportunities. The mean longest contiguous dropout burst falls from 10.37 to 7.87 epochs, a 24.1% reduction, while the P95 longest burst falls from 14 to 11 epochs. These changes support a physical interpretation in which the optimizer reshapes the trajectory toward optical geometries that reduce sustained periods of poor measurement availability.

Table 4: Independent validation of sensing availability and outage duration.

Method	Mean availability	Mean measurements	Mean longest outage	P95 longest outage
Fuel-only	0.228	4.56	12.23	18
Continuous-sensing information-aware	0.376	7.53	10.66	14
Expected-Markov-availability-aware	0.403	8.06	10.37	14
CVaR-Markov-availability-aware	0.526	10.52	7.87	11

C Multistart Robustness

All five CVaR optimization restarts at the 3.25 m/s cap produced terminal-feasible trajectories. Their design objectives ranged from 26.056 to 27.081; the three best restarts were within 0.47% of one another. The selected restart had $\bar{J} = 26.056$, $\Delta v = 3.245$ m/s, design-scenario P95 of 76.15 m, and design-scenario CVaR95 of 506.19 m. The solver status for each restart reported that the pre-specified generation limit was reached, so global or local convergence is not claimed. Nevertheless, all candidate terminal residuals were below 2×10^{-11} and the repeated low-objective solutions indicate that the reported trajectory is not a single-population anomaly.

Table 5: Five-restart diagnostics for the selected 3.25 m/s-cap CVaR design.

Restart	Design objective	Δv [m/s]	Design P95 [m]	Design CVaR95 [m]
1	26.531	3.241	84.96	511.45
2	26.175	3.249	78.06	507.59
3	27.081	3.235	79.02	529.60
4	26.177	3.244	78.31	507.61
5 (selected)	26.056	3.245	76.15	506.19

D Maneuver-Cap Sensitivity

Figure 2 summarizes the cap sensitivity. The trend is not monotonic because the scenario objective is nonsmooth, the candidate parameterization is low dimensional, and the derivative-free search is generation limited; consequently the points are not presented as a certified Pareto frontier. Nevertheless, increasing maneuver authority generally reduces tail severity. The 3.50 m/s-cap study selected a 3.437 m/s trajectory with CVaR95 of 373.55 m, only 3.7% below the 3.25 m/s design, while CVaR99 was essentially unchanged (1065.63 m versus 1064.88 m). It also produced a worse median and a lower probability of satisfying the 2 m threshold. The 3.25 m/s cap is therefore retained as the balanced primary design point rather than selecting the lowest training objective or the largest maneuver budget.

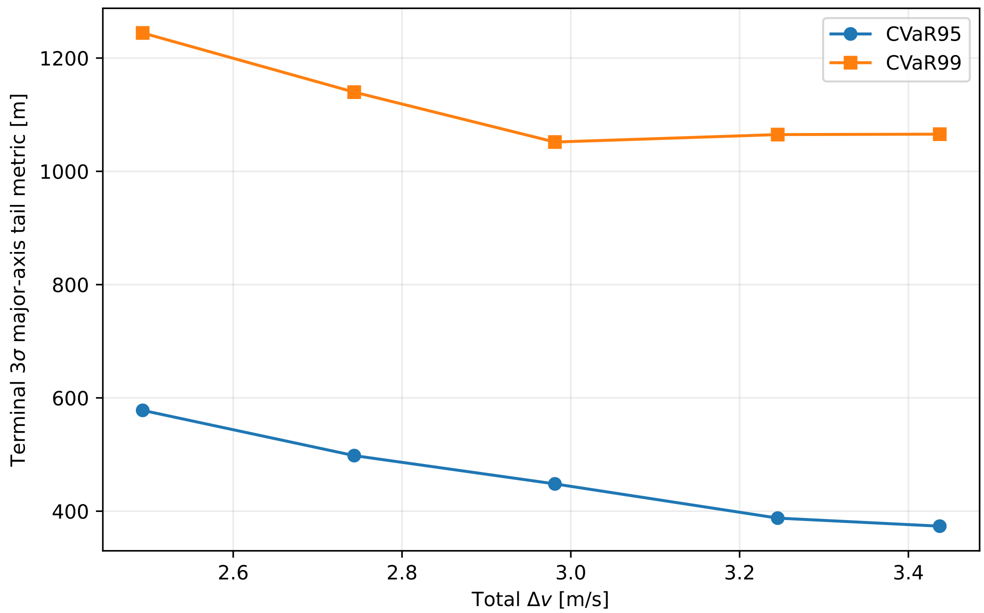

Figure 2: Sensitivity of CVaR95 and CVaR99 to the achieved maneuver expenditure under the tested total-Δv caps. Because restart counts and local solutions differ across cap values, the curve is interpreted as a sensitivity study rather than a formal Pareto frontier.

VII Nonlinear Three-Dimensional Closed-Loop Holdout Validation

A Truth Dynamics and Frozen Validation Protocol

The planar study isolates the stochastic sensing mechanism and supports large-sample tail statistics, but it does not test whether the navigation benefit survives nonlinear dynamics, estimator nonlinearities, maneuver execution error, and state-feedback coupling. A separate three-dimensional validation was therefore frozen after development and evaluated on a new 200-run holdout set. No feedback, estimator, guidance-weight, or visibility parameter was retuned after examining the holdout outcomes.

The chief is initialized on a 500 km, 51.6-deg circular orbit. Chief and deputy truth states are propagated independently in ECI coordinates with fourth-order Runge–Kutta integration and

$$\ddot{\boldsymbol{r}} = -\mu \frac{\boldsymbol{r}}{\|\boldsymbol{r}\|^3} + \boldsymbol{a}_{J_2}(\boldsymbol{r}), \tag{21}$$

where $\boldsymbol{a}_{J_2}$ is the standard first-order oblateness acceleration. The initial and terminal RTN relative states are

$$\boldsymbol{x}_{r,0} = [0, -1000, 75, 0, 0, 0]^T, \tag{22}$$

$$\boldsymbol{x}_{r,f} = [0, -50, 0, 0, 0, 0]^T, \tag{23}$$

with distance in meters and velocity in meters per second. The transfer retains twenty 60 s maneuver/measurement epochs.

Nominal guidance trajectories remain three-dimensional HCW designs so that the nonlinear experiment tests plant/model mismatch rather than replacing the planner with the truth model. The nonlinear CVaR nominal design uses 100 fixed outage scenarios, two differential-evolution restarts, ten search generations, the same $\alpha = 0.95$ risk metric, and a 3.25 m/s total-Δv cap. Both restarts were terminal-feasible; the selected restart had design objective 25.655 and nominal $\Delta v = 2.824$ m/s. The nonlinear holdout uses a previously unseen validation seed (314159265) to generate 200 common scenario realizations that are shared by all four guidance methods.

B Optical Measurement, UKF, and Maneuver Error

The camera supplies noisy azimuth and elevation of the deputy-to-chief line of sight $\boldsymbol{q} = [q_x, q_y, q_z]^T$,

$$\boldsymbol{z}_k = \begin{bmatrix} \operatorname{atan2}(q_y, q_x) \\ \operatorname{atan2}\Big(q_z, \sqrt{q_x^2 + q_y^2}\Big) \end{bmatrix} + \boldsymbol{v}_k, \qquad \boldsymbol{v}_k \sim \mathcal{N}(\boldsymbol{0}, \sigma_\theta^2 I_2), \tag{24}$$

with $\sigma_\theta = 0.05$ deg. The same Sun-exclusion probability and two-state Markov visibility process used in the planning model determine whether each measurement is actually returned.

A six-state ECI UKF is used for deputy navigation, with the chief treated as known. The

UKF propagates its sigma points through the same two-body+J_2 mean dynamics used by the estimator and applies circular statistics to the two angle measurements [29]. The initial 1σ navigation uncertainty is 50 m per position axis and 0.05 m/s per velocity axis. White acceleration process uncertainty is 10^{-4} m/s^2, with 1.01 covariance inflation per prediction interval. Commanded maneuver uncertainty is injected into the velocity covariance. Truth maneuvers include independent 1% magnitude error and 0.20 deg pointing error.

Estimator consistency is summarized using

$$\epsilon_{\text{NEES}} = \boldsymbol{e}^T P^{-1} \boldsymbol{e} \tag{25}$$

for the six-state error and

$$\epsilon_{\text{NIS}} = \boldsymbol{\nu}^T S^{-1} \boldsymbol{\nu} \tag{26}$$

for the two-dimensional innovation. For a perfectly calibrated filter their expectations are 6 and 2, respectively; the reported values below are empirical averages over the holdout campaign and should be interpreted as diagnostics rather than formal consistency proofs.

C Covariance-Weighted Terminal Feedback and Safety Metrics

At each guidance epoch, the frozen nominal maneuver is augmented by the first control from a regularized finite-horizon terminal correction. Let $\boldsymbol{e}$ be the current RTN estimate error relative to the nominal state and let M_k map the remaining impulsive corrections to the terminal state. The correction solves

$$\min_{\Delta U} \left(F^{N-k}\boldsymbol{e} + M_k \Delta U\right)^T Q_f \left(F^{N-k}\boldsymbol{e} + M_k \Delta U\right) + \rho \|\Delta U\|^2, \tag{27}$$

where position errors are normalized by 20 m, velocity errors by 0.10 m/s, velocity receives an additional factor of four, and $\rho = 0.10$. Only the first impulse is applied and the problem is resolved at the next epoch. Each feedback component is limited to 0.04 m/s.

Feedback authority is also reduced when the UKF position covariance is poor. With

$$\sigma_{3r} = 3\sqrt{\lambda_{\max}(P_r)}, \tag{28}$$

the covariance authority is one for $\sigma_{3r} \leq 30\,\mathrm{m}$, zero for $\sigma_{3r} \geq 100\,\mathrm{m}$, and linearly interpolated between these limits. This is multiplied by a time-varying factor with a 0.35 late-stage floor and an overall feedback gain of 0.35.

Mission success is defined by terminal position error no larger than 20 m, terminal velocity error no larger than 0.10 m/s, and no penetration of a 30 m keep-out zone. A 30-deg approach-corridor diagnostic is monitored inside 300 m, but is not imposed as a guidance constraint and is not part of the success definition.

D Independent 200-Run Holdout Results

Table 6 reports the true closed-loop terminal performance. All methods retain high mission-success probability: 98.0% for fuel-only, 99.0% for continuous-information, 99.0% for expected-Markov, and 98.5% for CVaR-Markov. The corresponding Wilson 95% confidence intervals overlap, so the 200-run experiment does not support a claim that CVaR improves mission-success probability. The proposed controller also does not minimize true terminal position error: its P95 position error is 15.50 m, compared with 15.21 m for expected-Markov and 12.97 m for continuous-information. No P99 or CVaR99 statistic is reported from this 200-run holdout; those extreme-tail claims are reserved for the independent 10,000-realization planar experiment. Instead, the main nonlinear benefit appears in sensing and navigation robustness.

Table 6: Nonlinear 3-D closed-loop holdout performance over 200 common scenarios.

Method	Mean cmd. Δv [m/s]	Median pos. [m]	P95 pos. [m]	P95 vel. [m/s]	Success [%]	KOZ viol. [%]
Fuel-only	1.892	5.83	14.72	0.065	98.0	1.0
Continuous-information	2.287	4.44	12.97	0.061	99.0	0.0
Expected-Markov	2.491	6.03	15.21	0.038	99.0	0.5
CVaR-Markov	2.829	5.18	15.50	0.049	98.5	0.0

Table 7 shows a markedly different ordering for navigation. Relative to expected-Markov guidance, CVaR-Markov increases the realized measurement fraction from 20.7% to 56.1%, reduces the mean longest outage from 14.39 to 6.80 epochs, reduces median terminal position-estimation error from 49.64 m to 0.89 m, and reduces P95 estimation error from 223.24 m to 21.93 m. These correspond to a 171.5% relative increase in realized sensing rate, a 52.7% reduction in mean longest outage, a 98.2% reduction in median estimation error, and a 90.2% reduction in P95 estimation error, for a 13.5% increase in mean commanded maneuver expenditure. Relative to continuous-information guidance, the CVaR trajectory still provides 36.3% higher mean measurement availability and 50.4% lower P95 terminal estimation error.

Table 7: Nonlinear holdout navigation and sensing performance.

Method	Median est. [m]	P95 est. [m]	Meas. frac. [-]	Mean outage [epochs]	NEES [-]	NIS [-]
Fuel-only	7.36	65.35	0.344	11.19	34.42	2.59
Continuous-information	1.77	44.23	0.404	9.81	82.32	2.57
Expected-Markov	49.64	223.24	0.207	14.39	743.89	40.14
CVaR-Markov	0.89	21.93	0.561	6.80	12.72	1.62

Because the methods use common random numbers, paired comparisons can be made scenario by scenario. CVaR-Markov has lower terminal estimation error than expected-Markov in 97.0% of matched holdout cases, a shorter longest outage in 95.5%, and a higher average arrival probability in all 200 cases. Against continuous-information guidance, the corresponding fractions are 72.0%, 77.5%, and 100%. By contrast, CVaR has lower true terminal position error than continuous-information guidance in only 39.0% of matched cases. This separation is important: the nonlinear holdout strongly supports the sensing/navigation mechanism, but does not show that the current feedback architecture converts the full navigation advantage into a lower terminal-position-error distribution.

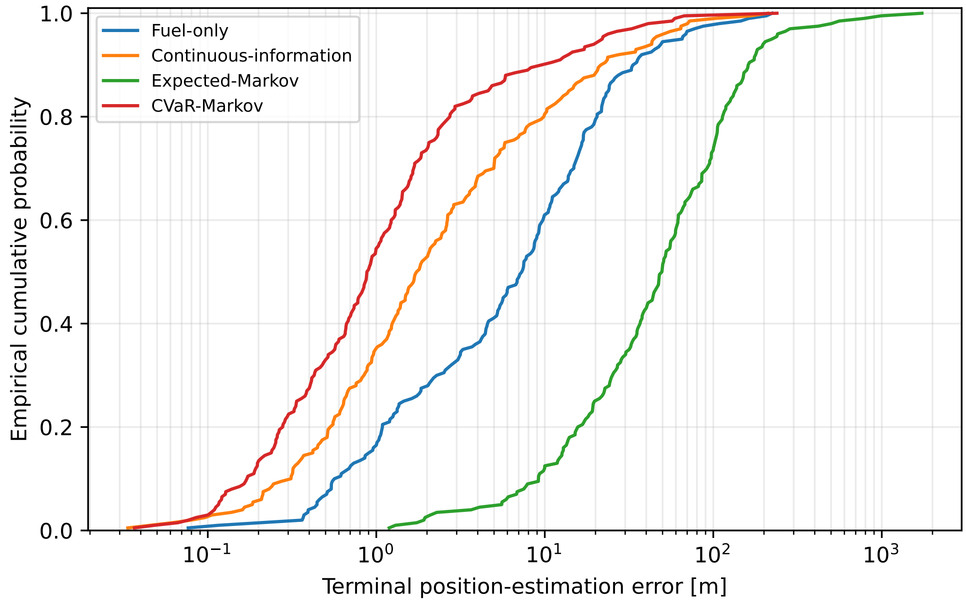

Figure 3: Empirical CDF of terminal position-estimation error in the independent 200-run nonlinear holdout. The horizontal axis is logarithmic to expose the heavy tails associated with poor observability.

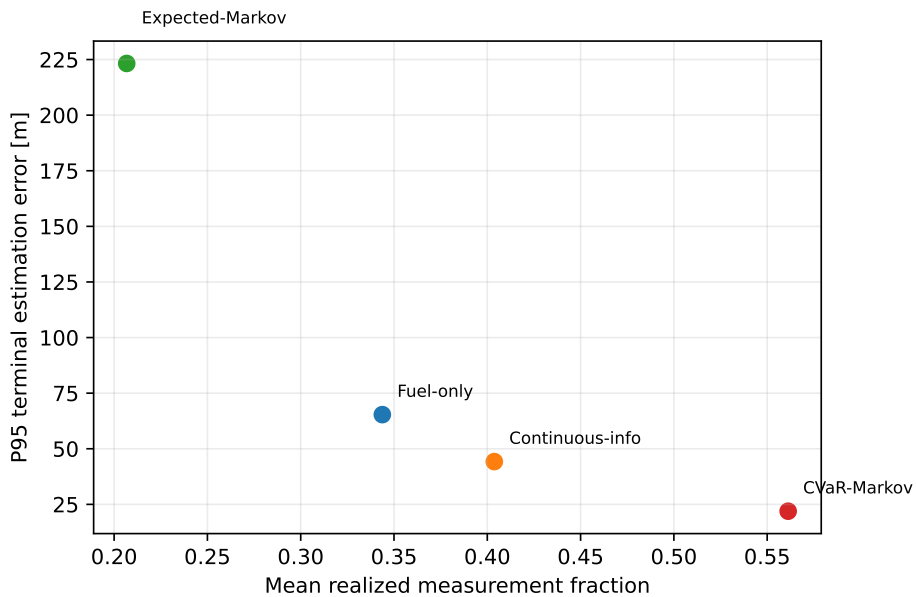

Figure 4: Nonlinear holdout sensing/navigation trade: mean realized measurement fraction versus P95 terminal position-estimation error. The CVaR trajectory occupies the favorable high-availability, low-tail-error region.

The CVaR trajectory has no keep-out-zone violations in the 200-run sample. The monitored approach corridor is not enforced by any of the planners: violation rates are 100% for fuel-only, 88.0% for continuous-information, 100% for expected-Markov, and 40.5% for CVaR-Markov. The lower CVaR violation rate shows that the risk-aware trajectory happens to align more often with the diagnostic corridor, but 40.5% remains far from a safety-compliant result. Because the corridor was not included as a design constraint, these rates should not be interpreted as a fair comparison of corridor-constrained guidance performance. They instead identify a missing requirement: corridor constraints must be embedded directly in the guidance problem before the method can support a claim of safety-compliant proximity operations.

VIII Discussion

The combined experiments support four conclusions.

First, explicit tail-risk optimization matters when sensing losses are burst correlated. In the 10,000-realization planar study, the expected-Markov trajectory achieves excellent median covariance but retains severe rare-event tails. The CVaR planner spends additional maneuver effort to reshape the uncertainty distribution, reducing CVaR95 by 39.2% and CVaR99 by 15.7%. This behavior is consistent with intermittent-observation theory, in which temporal loss correlation materially affects covariance peaks and tail behavior [26, 27].

Second, the physical mechanism survives a substantial increase in model fidelity. In the nonlinear 3-D holdout, CVaR guidance more than doubles the realized sensing rate relative to expected-Markov guidance and roughly halves the longest outage duration. The resulting UKF estimate is dramatically more robust: median and P95 terminal estimation errors decrease by 98.2% and 90.2%, respectively. The paired common-random-number results further show that this is not driven by a small number of favorable trials; CVaR produces lower terminal estimation error than expected-Markov in 97% of the same scenarios. This is the principal operational payoff established by the study: the vehicle spends substantially less of the rendezvous arc in navigation-poor geometry, reducing the incidence and severity of poorly localized states available to downstream guidance and safety logic.

Third, improved navigation and improved terminal guidance performance are not equivalent. Despite its much stronger sensing and estimation performance, CVaR-Markov does not beat continuous-information or expected-Markov guidance on P95 true terminal position error. The velocity-aware covariance-gated feedback law preserves high mission success, but the residual terminal-error distribution is still shaped by execution errors, model mismatch, and how estimated-state corrections are mapped into late rendezvous burns. The nonlinear result therefore supports the sensing/navigation claim more strongly than a closed-loop accuracy-dominance claim.

Fourth, the maneuver premium is mission dependent. In the planar large-sample benchmark the balanced CVaR solution costs 35.8% more Δv than expected-Markov guidance. In the nonlinear closed-loop holdout the mean commanded premium is 13.5%, partly because the three-dimensional nominal designs and feedback corrections differ from the planar benchmark. In both stages the proper interpretation is a trade between maneuver expenditure and resilience to poor sensing geometry, not a uniformly dominating controller.

The estimator diagnostics also merit caution. The CVaR holdout has the smallest mean NEES of the four nonlinear cases, but its value of 12.7 still exceeds the nominal six-state expectation of 6, while its mean NIS of 1.62 is below the two-dimensional expectation of 2. The filter is therefore substantially better behaved than the poorly observed expected-Markov case but is not perfectly calibrated. These statistics reinforce the need to report true estimation error together with covariance-based metrics.

IX Limitations and Remaining Work

The nonlinear holdout materially strengthens the study but does not make the present formulation flight ready.

First, the nominal guidance optimizer still uses three-dimensional HCW dynamics while truth propagation uses nonlinear two-body+J_2 dynamics. This mismatch is intentional for validation, but future work should solve the risk-aware trajectory problem directly in nonlinear relative dynamics. Finite-duration burns, attitude dynamics, plume constraints, drag, solar-radiation pressure, target ephemeris uncertainty, and chief-state uncertainty are also omitted.

Second, the optical-availability model remains stylized. The holdout uses azimuth/elevation measurements and an explicit Sun-exclusion probability, but does not model camera field of view, target attitude and feature visibility, Earth-limb exclusion, exposure/saturation, image-processing confidence, or a calibrated detection-probability curve for a particular sensor. Those effects should replace the generic logistic availability law before mission-specific interpretation.

Third, the 200-run nonlinear holdout is adequate for median and P95-level comparisons but is too small for stable P99 or CVaR99 claims. Extreme-tail claims therefore remain anchored in the independent 10,000-realization planar experiment. A future nonlinear campaign should use substantially more holdout cases or variance-reduction/importance-sampling methods if 99th-percentile closed-loop risk is a primary requirement.

Fourth, safety is monitored more completely than it is enforced. The keep-out-zone metric is included in the success definition and CVaR has no KOZ violations in the holdout, but the approach corridor is only diagnostic and remains violated in 40.5% of CVaR runs. The next guidance formulation should impose corridor, field-of-view, passive-safety, and collision-probability constraints directly inside the optimization [4, 12].

Fifth, the UKF is not perfectly calibrated across all geometries. CVaR-Markov produces much better NEES/NIS behavior than the poorly observed expected-Markov trajectory, yet its mean NEES remains above the six-state expectation. The current process-noise and covariance-inflation choices were frozen before holdout validation and should not be retrospectively tuned on these results. A future study should perform separate calibration data generation or consistency-constrained filter design.

Finally, the optimization is still generation limited and not globally certified. The planar study uses five multistart runs over 500 design scenarios, while the nonlinear nominal CVaR design uses two restarts over 100 design scenarios. Sequential convexification, stochastic direct transcription, sample-average approximation with convergence diagnostics, or policy-based receding-horizon risk optimization could provide stronger numerical guarantees.

X Conclusion

This paper formulated angles-only rendezvous guidance under geometry-dependent, burst-correlated optical-navigation outages as a tail-risk optimization problem and evaluated the concept at two levels of fidelity. A two-state Markov visibility process was combined with a smooth geometry-dependent detection probability, and explicit binary-arrival covariance scenarios were embedded in a CVaR-based trajectory objective. A terminal-null-space parameterization and fixed-scenario multistart optimization separated trajectory selection from independent validation.

In the 10,000-realization planar benchmark, the selected 3.25 m/s-cap CVaR trajectory required 3.245 m/s, 35.8% more maneuver expenditure than expected-Markov guidance, while reducing terminal position-uncertainty CVaR95 by 39.2% and CVaR99 by 15.7%. It also increased mean measurement availability by 30.6% and reduced the mean longest outage by 24.1%. These results establish the large-sample tail-risk mechanism.

A separate 200-run nonlinear 3-D holdout then tested the frozen architecture with two-body+J_2 truth, azimuth/elevation sensing, UKF navigation, Markov losses, maneuver execution errors, and covariance-weighted velocity-aware feedback. Relative to expected-Markov guidance, CVaR increased the realized measurement fraction from 20.7% to 56.1%, reduced mean longest outage duration from 14.39 to 6.80 epochs, reduced median terminal estimation error from 49.64 m to 0.89 m, and reduced P95 terminal estimation error from 223.24 m to 21.93 m, for a 13.5% increase in mean commanded Δv. Mission success remained statistically comparable at 98.5%, with zero observed keep-out-zone violations.

The nonlinear study also establishes an important limitation: the navigation advantage does not automatically produce a lower terminal-position-error distribution. The current CVaR trajectory has P95 terminal position error of 15.50 m, comparable to 15.21 m for expected-Markov guidance and higher than 12.97 m for continuous-information guidance. Moreover, the diagnostic approach corridor is violated in 40.5% of CVaR holdout runs because that corridor is not enforced by the planner. The strongest supported claim is therefore that tail-risk-aware guidance improves sensing geometry and navigation resilience under burst-correlated outages while preserving rendezvous performance; it is not yet a safety-compliant proximity-guidance law. The next step is joint nonlinear risk-aware guidance and explicit safety-constraint design so that the navigation advantage is translated directly into closed-loop terminal and proximity-safety guarantees.

Reproducibility Statement

The planar primary comparison uses 500 fixed CVaR design scenarios, five differential-evolution restarts, design seed 20260913, and a separate 10,000-realization validation set with seed 987654321. Every restart uses the same design scenario set; only the optimizer population seed changes. The minimum feasible design-objective restart is selected before validation.

The nonlinear 3-D holdout uses the frozen UKF/feedback-v2 architecture, 100 fixed CVaR design scenarios, two CVaR restarts, design seed 20260913, and 200 previously unseen validation scenarios generated with seed 314159265. All four guidance methods are evaluated using the same holdout scenario bank (hidden Markov states, arrival uniforms, camera noise, initial-state perturbations, and maneuver-error draws). The accompanying result CSVs, restart diagnostics, and simulation scripts reproduce the numerical values reported in this draft.

References

[1] W. H. Clohessy and R. S. Wiltshire, "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, Vol. 27, No. 9, 1960, pp. 653–658. doi:10.2514/8.8704.

[2] D. K. Geller, "Linear Covariance Techniques for Orbital Rendezvous Analysis and Autonomous Onboard Mission Planning," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 6, 2006, pp. 1404–1414. doi:10.2514/1.19447.

[3] D. C. Woffinden and D. K. Geller, "Navigating the Road to Autonomous Orbital Rendezvous," *Journal of Spacecraft and Rockets*, Vol. 44, No. 4, 2007, pp. 898–909. doi:10.2514/1.30734.

[4] L. Breger and J. P. How, "Safe Trajectories for Autonomous Rendezvous of Spacecraft," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 5, 2008, pp. 1478–1489. doi:10.2514/1.29590.

[5] D. C. Woffinden and D. K. Geller, "Optimal Orbital Rendezvous Maneuvering for Angles-Only Navigation," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 4, 2009, pp. 1382–1387. doi:10.2514/1.45006.

[6] J. Grzymisch and W. Fichter, "Analytic Optimal Observability Maneuvers for In-Orbit Bearings-Only Rendezvous," *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 5, 2014, pp. 1658–1664. doi:10.2514/1.G000612.

[7] J. Grzymisch and W. Fichter, "Optimal Rendezvous Guidance with Enhanced Bearings-Only Observability," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 6, 2015, pp. 1131–1139. doi:10.2514/1.G000822.

[8] X. Chu, Z. Liang, and Y. Li, "Trajectory Optimization for Rendezvous with Bearing-Only Tracking," *Acta Astronautica*, Vol. 171, 2020, pp. 311–322. doi:10.1016/j.actaastro.2020.03.017.

[9] S.-H. Mok, J. Pi, and H. Bang, "One-Step Rendezvous Guidance for Improving Observability in Bearings-Only Navigation," *Advances in Space Research*, Vol. 66, No. 11, 2020, pp. 2689–2702. doi:10.1016/j.asr.2020.07.035.

[10] M. Fujiwara and R. Funase, "Observability-Aware Differential Dynamic Programming with Impulsive Maneuvers," *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 9, 2024, pp. 1905–1919. doi:10.2514/1.G007798.

[11] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman Filtering with Intermittent Observations," *IEEE Transactions on Automatic Control*, Vol. 49, No. 9, 2004, pp. 1453–1464. doi:10.1109/TAC.2004.834121.

[12] Y. Zhang, M. Cheng, B. Nan, and S. Li, "Stochastic Trajectory Optimization for 6-DOF Spacecraft Autonomous Rendezvous and Docking with Nonlinear Chance Constraints," *Acta Astronautica*, Vol. 208, 2023, pp. 62–73. doi:10.1016/j.actaastro.2023.04.004.

[13] M. A. P. Ra and K. Oguri, "Chance-Constrained Sensing-Optimal Path Planning for Safe Angles-Only Autonomous Navigation," *AAS/AIAA Astrodynamics Specialist Conference*, 2024, pp. 1–20, AAS 24-436.

[14] Z. Zhao, F. Wu, and J. Chen, "Chance-Constrained Visual Servoing MPC for Non-Cooperative Spacecraft Rendezvous," *Aerospace Science and Technology*, Vol. 169, 2026, Art. 111237. doi:10.1016/j.ast.2025.111237.

[15] H. Frei, M. Burri, F. Rems, and E.-A. Risse, "A Robust Navigation Filter Fusing Delayed Measurements from Multiple Sensors and Its Application to Spacecraft Rendezvous," *Advances in Space Research*, Vol. 72, No. 7, 2023, pp. 2874–2900. doi:10.1016/j.asr.2022.10.025.

[16] Y. You, H. Wang, C. Paccolat, V. Gass, J.-P. Thiran, and J. R. Li, "Time and Covariance Threshold Triggered Optimal Uncooperative Rendezvous Using Angles-Only Navigation," *International Journal of Aerospace Engineering*, Vol. 2017, 2017, Art. 5451908. doi:10.1155/2017/5451908.

[17] J. Lee, S. Y. Park, and D. E. Kang, "Relative Navigation with Intermittent Laser-Based Measurement for Spacecraft Formation Flying," *Journal of Astronomy and Space Sciences*, Vol. 35, No. 3, 2018, pp. 163–173. doi:10.5140/JASS.2018.35.3.163.

[18] X. Zhao and S. Zhang, "Image-Feature-Based Integrated Path Planning and Control for Spacecraft Rendezvous Operations," *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 5, 2022, pp. 830–845. doi:10.2514/1.G006325.

[19] Z. Lin, B. Wu, and D. Wang, "Occlusion Avoidance in Image-Based Control of Multiple Spacecraft," *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 6, 2024, pp. 1150–1166. doi:10.2514/1.G007856.

[20] J. Kruger and S. D'Amico, "Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems," *Advances in Space Research*, Vol. 73, No. 11, 2024, pp. 5464–5483. doi:10.1016/j.asr.2023.08.055.

[21] A. Deole and M. Mesbahi, "Estimation-Aware Trajectory Optimization with Set-Valued Measurement Uncertainties," *Journal of Guidance, Control, and Dynamics*, Vol. 49, No. 5, 2026, pp. 1263–1281. doi:10.2514/1.G009299.

[22] S. E. Clees, S. Phillips, and C. Petersen, "Information-Aware Model Predictive Control for Satellite Inspection Missions," arXiv:2608.07765, 2026, preprint AAS 26-682. doi:10.48550/arXiv.2608.07765.

[23] R. Bommena and R. M. Woollands, "Fault-Tolerant Beacon Placement Optimization with Line-of-Sight and Field-of-View Constraints for Robust Relative Navigation," *Acta Astronautica*, Vol. 249, 2026, pp. 28–47. doi:10.1016/j.actaastro.2026.07.004.

[24] O. Napolitano, D. Fontanelli, L. Pallottino, and P. Salaris, "Gramian-Based Optimal Active Sensing Control Under Intermittent Measurements," *Proc. IEEE International Conference on Robotics and Automation*, 2021, pp. 9680–9686. doi:10.1109/ICRA48506.2021.9561008.

[25] J. A. Christian, "StarNAV: Autonomous Optical Navigation of a Spacecraft by the Relativistic Perturbation of Starlight," *Sensors*, Vol. 19, No. 19, 2019, Art. 4064. doi:10.3390/s19194064.

[26] M. Huang and S. Dey, “Stability of Kalman Filtering with Markovian Packet Losses,” *Automatica*, Vol. 43, No. 4, 2007, pp. 598–607. doi:10.1016/j.automatica.2006.10.023.

[27] Y. Mo and B. Sinopoli, “Kalman Filtering with Intermittent Observations: Tail Distribution and Critical Value,” *IEEE Transactions on Automatic Control*, Vol. 57, No. 3, 2012, pp. 677–689. doi:10.1109/TAC.2011.2166309.

[28] R. T. Rockafellar and S. Uryasev, “Conditional Value-at-Risk for General Loss Distributions,” *Journal of Banking & Finance*, Vol. 26, No. 7, 2002, pp. 1443–1471. doi:10.1016/S0378-4266(02)00271-6.

[29] S. J. Julier and J. K. Uhlmann, “Unscented Filtering and Nonlinear Estimation,” *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422. doi:10.1109/JPROC.2003.823141.