Rarity of rocket-driven Penrose extraction in Kerr spacetime

An T. Le[*]

College of Engineering and Computer Sciences, VinUniversity, Hanoi, Vietnam;
Center for Environmental Intelligence, VinUniversity, Hanoi, Vietnam;
and Intelligent Autonomous Systems, TU Darmstadt, Darmstadt, Germany

(Received 28 January 2026; accepted 8 May 2026; published 11 June 2026)

We study rocket-driven Penrose extraction in the test-particle limit on a fixed Kerr background for equatorial prograde flybys under explicit steering prescriptions. A spacecraft ejects exhaust inside the ergosphere; when the exhaust attains negative Killing energy, the remaining spacecraft gains energy by 4-momentum conservation. Across 320,000 simulated trajectories spanning black hole spin, exhaust velocity, and orbital parameters, extraction with escape is rare in broad parameter scans (at most $\sim$1%) and requires high spin ($a/M \gtrsim 0.89$), highly relativistic exhaust ($v_e \gtrsim 0.91c$), and finely tuned initial conditions. Under optimal tuning the success rate reaches $\sim$70% at $a/M = 0.95$. For representative escape trajectories, a single periapsis impulse is more propellant-efficient than the continuous-thrust controllers studied here. All quoted thresholds are empirical and specific to the orbit family, prior, and steering protocol studied.

DOI: 10.1103/f3jn-7wcp

I. INTRODUCTION

The extraction of energy from rotating black holes stands as one of the most remarkable predictions of general relativity. In 1971, Penrose and Floyd described a mechanism whereby particles entering the ergosphere of a Kerr black hole could split, with one fragment falling into the horizon carrying negative Killing energy while the other escapes with energy exceeding that of the original particle [1]. The rotational energy of a Kerr black hole above its irreducible mass amounts to up to 29% of the total mass [2,3], in principle accessible through repeated Penrose events; the maximum single-decay efficiency is $\sim$20.7% for extremal Kerr [4].

The theoretical elegance of the Penrose process has inspired decades of research into its variants and astrophysical manifestations. Piran *et al.* [5] showed that particle collisions in the ergosphere could achieve substantially higher center-of-mass energies. This collisional Penrose process was later shown by Bañados *et al.* [6] to yield arbitrarily high center-of-mass energies near extremal horizons, though practical efficiencies are limited by gravitational redshift of escaping debris [7–9]. Leiderschneider and Piran [10] systematically characterized the maximal efficiency of the collisional process.

In astrophysical contexts, electromagnetic variants of energy extraction have proven more observationally relevant. The Blandford-Znajek mechanism [11] extracts rotational energy through horizon-threading magnetic fields supported by surrounding accretion flows, powering the relativistic jets observed in active galactic nuclei and quasars [12]. More recently, Comisso and Asenjo [13] demonstrated that magnetic reconnection in the ergosphere provides an efficient Penrose-like mechanism, potentially explaining the rapid variability of high-energy emission from supermassive black holes. Koide and Arai [14] had earlier explored similar magnetic reconnection scenarios numerically.

Despite this extensive body of work, a fundamental question remains largely unexplored: How fine-tuned must conditions be for successful Penrose extraction with escape to infinity? The classical Penrose process requires precisely tuned initial conditions—the decaying particle must enter the ergosphere on a trajectory that (i) penetrates deep enough for negative-energy states to be accessible (the captured fragment must carry negative energy where frame dragging is strong), (ii) does not plunge into the horizon, and (iii) allows the positive-energy fragment to escape with sufficient angular momentum. This narrow viable region in parameter space has been noted qualitatively [15,16], but to our knowledge no systematic statistical characterization exists to quantify how rare success is or what parameter thresholds govern extraction.

This paper addresses this gap by studying the Penrose process via rocket propulsion, in which a spacecraft ejects mass (exhaust) inside the ergosphere, analogous to the

[*]Contact author: an.lt@vinuni.edu.vn, an@robot-learning.de; https://anindex.github.io/

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

classical particle decay but with thrust direction and magnitude allowed to vary under explicit steering prescriptions. This is a kinematic test-particle study: the spacecraft mass satisfies $m_0 \ll M$, so the Kerr background is fixed and self-force effects are negligible; inferred changes to the hole's rotational energy are *post hoc* estimates, not dynamical backreaction. We restrict attention to initially unbound equatorial prograde flybys, the natural baseline family for the mission objective of energy gain with escape to infinity: these trajectories enter the ergosphere, interact with the black hole's rotation, and, if conditions are favorable, escape to infinity with a net energy increase. [Later diagnostic sweeps in Sec. VIII C also explore a broader (E_0, L_z) domain that intentionally includes bound and retrograde states to map failure boundaries; those extended-domain results characterize the parameter landscape, not the targeted mission family.] We provide a "computational feasibility map" for rocket-like Penrose extraction, establishing quantitative thresholds, efficiencies, and success probabilities across large parameter sweeps. Our formulation offers three advantages. First, it connects the abstract Penrose process to a concrete engineering scenario, clarifying which physical parameters govern extraction. Second, it enables Monte Carlo exploration of the initial condition space, quantifying the statistical rarity of successful extraction with escape. Third, it permits comparison of thrust strategies (impulsive vs continuous), revealing path-averaging effects that reduce efficiency in extended maneuvers. The novelty of this work lies not in the existence of the Penrose process, but in providing quantified success rates and parameter thresholds under a rocket-propulsion control model, together with a thrust-profile comparison and a constrained optimality argument. Our primary contribution is the impulsive periapsis-ensemble rarity and threshold map; the continuous-thrust comparison is a secondary study of representative matched trajectories under specific implemented controllers.

Our central finding is a set of constraint results: within our equatorial prograde-flyby family and steering prescriptions (hereafter our "baseline family and protocol"), successful Penrose extraction with spacecraft escape is observed only for (i) high spin ($a/M \gtrsim 0.89$), (ii) highly relativistic exhaust ($v_e \gtrsim 0.91c$), and (iii) carefully tuned initial conditions within a narrow "sweet spot." Under maximum tuning the success rate can reach ~70% at $a/M = 0.95$, but this represents the fine-tuned limit, not typical behavior.[1] Single periapsis-impulse ejection is more propellant efficient than continuous thrust for the representative matched escape trajectories studied, consistent with the path-averaging penalty of the fixed-worldline optimality argument and, additionally, with trajectory-deformation effects during extended ergosphere transit. Quantitative details, confidence intervals, and efficiency metrics are given in the Results sections.

Throughout this paper we distinguish three logically separate layers: (i) "local kinematics," whether the exhaust attains negative Killing energy ($E_{\rm ex} < 0$); (ii) "mission outcome," whether the spacecraft both gains energy and escapes to infinity [$\mathcal{S} = (\Delta E > 0) \wedge$ escape, Eq. (18)]; and (iii) "performance metrics," cumulative and traditional efficiencies ($\eta_{\rm cum}$, $\eta_{\rm trad}$). These are defined in Secs. IV D and IV E and applied throughout the results.

The paper is organized as follows. Section II reviews Kerr spacetime geometry. Section III summarizes the classical Penrose process, and Sec. IV introduces the rocket-driven variant. Sections V and VI present the mathematical formulation and thrust model. Section VII describes numerical methods. Section VIII presents numerical results, including a theoretical optimality argument for single-impulse extraction (Sec. VIII D) followed by an illustrative controller comparison. The statistical core of this study is the periapsis-impulse ensemble (Tables I–V; exploratory $a/M = 0.99$ results appear in Appendix B); continuous thrust is examined as a controller comparison on representative matched trajectories (Sec. VIII E), not as a coequal experimental pillar. Section IX summarizes implications.

II. PROBLEM SETTING: KERR SPACETIME

To make the control and energetics statements precise, we begin by fixing notation for Kerr geometry, the ergosphere, and the conserved quantities used throughout.

We consider a rotating (Kerr) black hole of mass M and angular momentum $J = aM$, where $0 \le a \le M$ is the spin parameter. In Boyer-Lindquist coordinates (t, r, θ, ϕ), the Kerr metric takes the form

$$ds^2 = -\left(1 - \frac{2Mr}{\Sigma}\right)dt^2 - \frac{4M\,ar\sin^2\theta}{\Sigma}\,dt\,d\phi + \frac{\Sigma}{\Delta}dr^2 + \Sigma d\theta^2 + \frac{A\sin^2\theta}{\Sigma}d\phi^2, \tag{1}$$

where

$$\Sigma = r^2 + a^2\cos^2\theta, \tag{2}$$

$$\Delta = r^2 - 2Mr + a^2, \tag{3}$$

$$A = (r^2 + a^2)^2 - a^2\Delta\sin^2\theta. \tag{4}$$

The event horizon is located at

$$r_+ = M + \sqrt{M^2 - a^2}, \tag{5}$$

which exists for $a \le M$. The extremal limit $a = M$ corresponds to a maximally rotating black hole with $r_+ = M$.

[1] An exploratory scan at $a/M = 0.99$ yields ~65%–70% under the same prior (independent-seed verification; see Appendix B).

A. Constants of motion

The Kerr spacetime admits two Killing vectors: $\xi^\mu = (\partial/\partial t)^\mu$ (stationarity) and $\psi^\mu = (\partial/\partial\phi)^\mu$ (axisymmetry). For geodesic motion, these yield conserved quantities,

$$E = -p_t = -g_{t\mu}p^\mu, \tag{6}$$

$$L_z = p_\phi = g_{\phi\mu}p^\mu, \tag{7}$$

where E is the Killing energy (as measured by observers at rest at infinity) and L_z is the azimuthal angular momentum. We normalize the spacecraft's initial rest mass to $m_0 = 1$ throughout (see Sec. VII); all sweep parameters $E_0 \equiv E/m_0 = -u_t$ and $L_{z,0} \equiv L_z/m_0 = u_\phi$ denote specific (per-unit-mass) quantities. In particular, an unbound orbit has specific energy $E_0 > 1$ (not merely $E_0 > 0$), and escape to infinity requires $E/m > 1$ at large radius. Throughout this paper, "retrograde" refers to negative L_z (opposing the black hole rotation), not to $u^\phi < 0$ in Boyer-Lindquist coordinates; inside the ergosphere, frame dragging forces all future-directed worldlines to have $u^\phi > 0$.

B. The ergosphere

The stationary limit surface (outer boundary of the ergosphere) is defined by $g_{tt} = 0$, yielding

$$r_{\rm erg}(\theta) = M + \sqrt{M^2 - a^2\cos^2\theta}. \tag{8}$$

At the equatorial plane ($\theta = \pi/2$), we have $r_{\rm erg} = 2M$ for all values of spin. The ergosphere is the region $r_+ < r < r_{\rm erg}(\theta)$ where ξ^μ becomes spacelike, meaning no observer can remain stationary. Inside the ergosphere, timelike worldlines can have $E < 0$ (negative Killing energy); outside, the timelike character of ξ^μ forces $E \geq 0$ for all physical worldlines, so negative-energy states are confined to the ergoregion. Detailed properties of negative-energy geodesics in the Kerr ergosphere are discussed in Ref. [17].

The angular velocity of frame dragging is

$$\omega = -\frac{g_{t\phi}}{g_{\phi\phi}} = \frac{2Mar}{A}. \tag{9}$$

This frame dragging is central to energy extraction via the Penrose mechanism.

III. CLASSICAL PENROSE PROCESS

The original Penrose process [1] provides a mechanism to extract rotational energy from a Kerr black hole using particle decay.

A. Physical mechanism

Consider an incident particle with 4-momentum p_0^μ and energy $E_0 > 0$ entering the ergosphere. Inside the ergosphere, the particle undergoes a decay,

$$\text{particle}_0 \to \text{particle}_1 + \text{particle}_2. \tag{10}$$

Conservation of 4-momentum requires

$$p_0^\mu = p_1^\mu + p_2^\mu. \tag{11}$$

Contracting with the Killing vector $\xi_\mu = (\partial/\partial t)_\mu$,

$$E_0 = E_1 + E_2. \tag{12}$$

Inside the ergosphere, timelike worldlines can have $E < 0$ (negative Killing energy, as measured by observers at rest at infinity) while remaining physical. Specifically, if particle 2 is ejected with large retrograde angular momentum (opposing the black hole rotation), it can achieve $E_2 < 0$. Particle 1 then escapes with

$$E_1 = E_0 - E_2 > E_0, \tag{13}$$

extracting energy from the black hole.

B. Conditions for energy extraction

For successful energy extraction, the escaping particle must satisfy:

(1) The decay occurs inside the ergosphere: $r_+ < r < r_{\rm erg}$.
(2) The captured particle has negative Killing energy: $E_2 < 0$.
(3) The escaping fragment follows a geodesic that reaches spatial infinity; for a massive particle, this requires specific energy $E_1/m_1 > 1$ (unbound) together with an allowed radial turning-point structure. Positive total energy $E_1 > 0$ alone is insufficient, since positive-energy bound orbits exist in Kerr.

A necessary condition for $E_2 < 0$ is that the fragment is created inside the ergosphere and carries sufficiently negative (retrograde) angular momentum. Writing $p^\mu = mu^\mu$ and $u^\phi = \Omega u^t$ with $\Omega \coloneqq d\phi/dt$, the condition $E < 0$ requires

$$g_{tt} + g_{t\phi}\Omega > 0 \quad \Leftrightarrow \quad \Omega < -\frac{g_{tt}}{g_{t\phi}} (g_{t\phi} < 0). \tag{14}$$

Even inside the ergosphere, frame dragging enforces $\Omega \in (\Omega_-, \Omega_+)$ with $\Omega_\pm > 0$, where

$$\Omega_\pm = \frac{-g_{t\phi} \pm \sqrt{g_{t\phi}^2 - g_{tt}g_{\phi\phi}}}{g_{\phi\phi}}. \tag{15}$$

C. Efficiency limits

Two distinct efficiency bounds govern energy extraction. The total extractable rotational energy of a Kerr black hole is [2]

$$\eta_{\mathrm{rot}} = 1 - \sqrt{\frac{1+\sqrt{1-(a/M)^2}}{2}}, \tag{16}$$

which gives $\eta_{\mathrm{rot}} \approx 29.3\%$ for an extremal black hole ($a = M$). This represents the thermodynamic limit on cumulative extraction via repeated processes.

The single-decay efficiency was bounded by Wald [4]. For an incident particle from rest at infinity in extremal Kerr ($a = M$), the maximum fractional energy gain of the escaping fragment is

$$\eta_{\mathrm{Wald}} = \frac{E_1 - E_0}{E_0} \approx 20.7\%. \tag{17}$$

Our single-impulse results should be compared to this classical benchmark, while cumulative processes may approach the rotational bound through multiple extractions.

IV. ROCKET-DRIVEN PENROSE VARIANT

We formulate a general rocket-driven Penrose model with two thrust modes: single periapsis impulses and continuous thrust. Most ensemble statistics reported below use the impulsive mode; continuous thrust is used mainly for matched-trajectory comparisons (Sec. VIII E). The continuous formulation reveals quantitative path-averaging effects absent from single-decay analyses.

A. Physical concept

A spacecraft with initial rest mass m_0 and specific energy $E_0 = -u_t > 1$ (unbound) enters the ergosphere on a prograde-flyby orbit (angular momentum $L_z > 0$, corotating with the black hole). The orbit is chosen such that its periapsis lies inside the ergosphere. While passing through the ergosphere near periapsis, the spacecraft activates a rocket engine that:

(1) ejects mass at exhaust velocity v_e (in the spacecraft's rest frame);
(2) chooses a thrust orientation such that the exhaust carries retrograde angular momentum ($L_{z,\mathrm{ex}} < 0$); and
(3) inside the ergosphere, the exhaust can then have $E_{\mathrm{ex}} < 0$, transferring energy to the remaining spacecraft.

Conceptually, each infinitesimal exhaust element plays the role of the captured negative-energy fragment in the classical single decay, while the remaining spacecraft plays the role of the escaping fragment; the rocket engine provides control over the split direction and timing.

In successful extraction-with-escape cases, the ejection of negative-energy exhaust results in the spacecraft gaining energy at infinity, emerging from the flyby with $E_f > E_0$. This is the rocket-driven Penrose process: many infinitesimal mass ejections, each of which can carry away negative Killing energy when the exhaust direction is suitably chosen.

Critical insight. In our parameter sweeps, successful extraction with escape arises from "prograde, initially unbound" flybys ($L_z > 0$, $E_0 > 1$). Retrograde ($L_z < 0$) or initially bound ($E_0 < 1$) initial conditions did not yield successful extraction with escape in our sampled domain, which may reflect their more restrictive turning-point structure for deep-ergosphere transits.

B. Advantages over classical process

Our continuous formulation offers several modeling advantages: rocket propulsion is a well-understood technology analogy; thrust magnitude and direction can be modulated; and active propulsion can be used to attempt to prolong ergosphere dwell time and reduce capture risk under explicit steering prescriptions.

C. Definition of extraction-with-escape success

We define a "successful extraction-with-escape event" (hereafter "extraction-with-escape success," denoting the mission-level criterion, not local Penrose extraction alone) as the conjunction of two conditions:

$$\mathcal{S} \coloneqq (\Delta E > 0) \wedge (\text{spacecraft escapes to infinity}). \tag{18}$$

Operationally, as a finite-radius escape proxy, we diagnose "escape" when the spacecraft reaches $r > 50M$ with outward radial velocity $dr/d\tau > 0$ and specific energy $E/m > 1$. We diagnose "capture" when $r < r_+ + 0.02M$ (horizon approach) or when integration fails due to stiffness near the horizon. We conservatively count integration failures (11.2% of samples) as captures.

Under our throttle and sign-selection policy, all single-impulse runs satisfying $\Delta E > 0$ with escape also exhibit $E_{\mathrm{ex}} < 0$ during the burn: within the sampled single-impulse ensemble, $E_{\mathrm{ex}} < 0$ is a necessary condition for extraction-with-escape success. However, $E_{\mathrm{ex}} < 0$ is not sufficient for mission success. The continuous-thrust capture example in Table IV achieves $E_{\mathrm{ex}} < 0$ for nearly all thrust events and $\Delta E > 0$, yet the spacecraft does not escape. More generally, negative-energy exhaust is confined to the ergoregion and is captured (Sec. II). The spacecraft must also navigate out of the ergosphere without plunging. We therefore use $\Delta E > 0 \wedge$ escape as the primary mission-level criterion, with $E_{\mathrm{ex}} < 0$ serving as the local Penrose signature.

D. Efficiency metrics

We define the "instantaneous efficiency" $\eta_{\rm inst} = dE/d(-m)$, where $\eta_{\rm inst} > 1$ indicates propellant rest-mass breakeven is exceeded (the Killing-energy gain per unit rocket mass loss exceeds unity). However, this alone does not prove Penrose extraction; a direct signature is negative exhaust Killing energy $E_{\rm ex} < 0$ inside the ergosphere.

The "cumulative efficiency" over the maneuver is

$$\eta_{\rm cum} = \frac{E_f - E_0}{m_0 - m_f} = \frac{\Delta E}{\Delta m}. \tag{19}$$

For later parameter sweeps we report propellant expenditure via the ejected rocket mass fraction

$$\delta m := \frac{m_0 - m_f}{m_0}, \tag{20}$$

so that $\Delta m = m_0 - m_f = \delta m m_0$.

For relativistic exhaust velocity v_e with Lorentz factor $\gamma_e = 1/\sqrt{1 - v_e^2}$, the differential relation $-dm = \gamma_e d\mu$ between rocket mass loss and exhaust rest mass is exact in the instantaneous rest frame and applies exactly to continuous-thrust simulations. For a finite (impulsive) burn, however, the exact mass-shell norm of $p'_\mu = p_\mu - \delta\mu\, u_{{\rm ex},\mu}$ gives

$$m_f^2 = m_0^2 - 2\gamma_e m_0 \delta\mu + \delta\mu^2, \tag{21}$$

so the exact exhaust rest mass for a desired mass fraction $\delta m = (m_0 - m_f)/m_0$ is

$$\delta\mu = m_0\left[\gamma_e - \sqrt{\gamma_e^2 - 2\delta m + \delta m^2}\right]. \tag{22}$$

Our single-impulse simulations parametrize by a nominal burn parameter $\delta m_{\rm nom}$ (not a physical mass fraction; see below) and compute $\delta\mu = \delta m_{\rm nom} \cdot m_0/\gamma_e$ (the differential-limit approximation), then apply exact 4-momentum conservation [Eq. (53)]. Because this approximate $\delta\mu$ exceeds the exact value from Eq. (22), the actual rocket mass fraction consumed $\delta m_{\rm actual} = (m_0 - m_f)/m_0$ exceeds $\delta m_{\rm nom}$: at $v_e = 0.95c$ and $\delta m_{\rm nom} = 0.4$, the discrepancy reaches ~34% (see Table IV caption for actual consumed fractions). The cumulative efficiency $\eta_{\rm cum} = \Delta E/(m_0 - m_f)$ is computed from the exact postimpulse mass and is unaffected by this labeling convention. For the per-exhaust-rest-mass ratio $\Delta E/\delta\mu$ reported in Table IV, we use the exact $\delta\mu$ from Eq. (22) applied to the actual consumed fraction $\delta m_{\rm actual}$; this is the physically meaningful normalization. We also define the "traditional Penrose efficiency," analogous to the single-decay efficiency used for the classical process,

$$\eta_{\rm trad} = \frac{E_f - E_0}{E_0}, \tag{23}$$

which measures the fractional energy gain relative to the initial energy. In Table IV we report $\eta_{\rm cum}$ and $\Delta E/\delta\mu$ for the representative trajectories; $\eta_{\rm trad}$ is retained as a complementary energy-normalized diagnostic and as the quantity directly comparable to Wald's classic single-decay benchmark. The cumulative efficiency $\eta_{\rm cum}$ [Eq. (19)] is an engineering propellant-efficiency metric: it quantifies how much Killing energy is gained per unit of rocket mass expended. It is the natural figure of merit for the rocket-propulsion context but mixes Penrose extraction physics with propellant expenditure. For cleaner Penrose-physics diagnostics, $E_{\rm ex} < 0$ (the local kinematic signature of negative-energy exhaust) and $\eta_{\rm trad}$ (the fractional energy gain) are more appropriate; in particular, $\eta_{\rm trad}$ enables direct comparison with Wald's classic single-decay benchmark $\eta_{\rm Wald} \approx 20.7\%$ for extremal Kerr [4].

E. Local kinematic threshold estimate

A key prediction of our Monte Carlo study is a sharp velocity threshold near $v_e \approx 0.91$–$0.92c$, below which extraction with escape largely fails (local $E_{\rm ex} < 0$ can still occur without the spacecraft escaping). Before formalizing the thrust model (Secs. V and VI), we preview the physical insight behind this threshold by connecting it to the local requirement that the exhaust (not the spacecraft) attain negative Killing energy. The exact exhaust-energy expression derived in Sec. VI D is

$$E_{\rm ex} = -\gamma_e(u_t - v_e s_t), \tag{24}$$

where $u_t = g_{t\mu}u^\mu$ is the covariant time component of the spacecraft 4-velocity and $s_t = g_{t\mu}s^\mu$ is the covariant time component of the unit spacelike thrust direction (the exhaust is emitted opposite to s^μ).

Negative-energy exhaust requires

$$E_{\rm ex} < 0 \quad \Leftrightarrow \quad u_t - v_e s_t > 0. \tag{25}$$

For the unbound, escaping trajectories of interest we have $E/m = -u_t > 1$, hence $u_t < 0$. Equation (25) then implies a simple but important sign constraint: the thrust direction must satisfy $s_t < 0$ so that the $-v_e s_t$ term can overcome the negative u_t. In our setup, this is achieved by selecting thrust orientations that eject exhaust with sufficiently retrograde angular momentum (cf. the convention defined in Sec. II).

When $u_t < 0$ and $s_t < 0$, Eq. (25) becomes a lower bound on exhaust speed,

$$v_e > v_e^{\rm crit} := \frac{|u_t|}{|s_t|}, \tag{26}$$

with u_t and s_t evaluated at the burn event. We stress that $v_e > v_e^{\mathrm{crit}}$ is a necessary local condition for $E_{\mathrm{ex}} < 0$ at a given burn event, not a sufficient condition for mission-level extraction-with-escape success [$\mathcal{S}$, Eq. (18)]: the spacecraft must still escape to infinity with $\Delta E > 0$. This makes the origin of the sharp onset transparent: as v_e increases, the set of spacetime points and thrust angles satisfying $v_e > v_e^{\mathrm{crit}}$ (hence $E_{\mathrm{ex}} < 0$) expands rapidly.

Evaluating Eq. (26) at periapsis for representative sweet-spot states ($a/M = 0.95$, $r_p \approx 1.5M$, $E_0 \approx 1.2$, $L_z \approx 3.0$) yields $v_e^{\mathrm{crit}} \approx 0.85$–$0.92c$, depending on the orbital phase and steering angle. This estimate is consistent with the Monte Carlo observation (Fig. 5) that success rates transition sharply around $v_e \approx 0.91$–$0.92c$: below this value, only rare geometries achieve $E_{\mathrm{ex}} < 0$ while retaining escape; above it, negative-energy exhaust becomes robust over a wide range of periapsis configurations.

Important caveat. The threshold $v_e \approx 0.91$–$0.92c$ is *not* a universal Kerr-invariant constant. It is an empirical threshold specific to our baseline Gaussian sweet-spot prior (Gaussian sampling around $E_0 = 1.22$, $L_z = 3.05$), our greedy thrust policy (maximize instantaneous $-dp_t/d\tau$), and the finite angular discretization of the thrust direction search (74 candidate directions; Sec. VI G). Different orbit families, steering laws, spin parameters, or finer direction scans will yield different critical velocities. The threshold is thus a geometric constraint arising from the Killing-energy decomposition $v_e > |u_t|/|s_t|$ [Eq. (26)], evaluated for a particular class of trajectories and controller.

We now formalize this controlled-ejection Penrose variant as a variable-mass, forced trajectory problem in Kerr spacetime.

V. MATHEMATICAL FORMULATION

We now present the mathematical formulation. Geodesic (thrust-off) phases follow Hamilton's equations in Kerr spacetime; thrust-on phases are governed by a forced variable-mass equation of motion (Sec. VI).

A. Equatorial restriction

For simplicity, we restrict to equatorial motion ($\theta = \pi/2$, $p_\theta = 0$). The dynamics reduce from 8 to 6 dimensions. At the equator,

$$\Sigma = r^2, \tag{27}$$

$$\Delta = r^2 - 2Mr + a^2, \tag{28}$$

$$g_{tt} = -\left(1 - \frac{2M}{r}\right), \tag{29}$$

$$g_{t\phi} = -\frac{2Ma}{r}, \tag{30}$$

$$g_{rr} = \frac{r^2}{\Delta}, \tag{31}$$

$$g_{\phi\phi} = r^2 + a^2 + \frac{2Ma^2}{r}. \tag{32}$$

B. State variables

The state vector for the spacecraft is

$$\boldsymbol{y} = (r, \phi, p_t, p_r, p_\phi, m), \tag{33}$$

where

(i) r is the Boyer-Lindquist radial coordinate;
(ii) ϕ is the azimuthal angle;
(iii) p_t is the covariant time momentum (energy, $E = -p_t$);
(iv) p_r is the covariant radial momentum;
(v) p_ϕ is the covariant azimuthal momentum (angular momentum, $L_z = p_\phi$); and
(vi) m is the rest mass (evolving due to propellant consumption).

C. Mass-shell constraint

The spacecraft 4-momentum must satisfy the mass-shell condition

$$g^{\mu\nu} p_\mu p_\nu = -m^2. \tag{34}$$

Expanding at the equator,

$$g^{tt} p_t^2 + 2g^{t\phi} p_t p_\phi + g^{rr} p_r^2 + g^{\phi\phi} p_\phi^2 = -m^2. \tag{35}$$

This constrains the relationship between p_t, p_r, and p_ϕ for given m.

D. Hamiltonian

The Hamiltonian for a free particle is

$$\mathcal{H} = \frac{1}{2} g^{\mu\nu} p_\mu p_\nu = \frac{1}{2}(g^{tt} p_t^2 + 2g^{t\phi} p_t p_\phi + g^{rr} p_r^2 + g^{\phi\phi} p_\phi^2). \tag{36}$$

For massive particles with $p_\mu = m u_\mu$, the mass-shell constraint gives $\mathcal{H} = -m^2/2$.

For evolution in proper time τ (rather than affine parameter), we define

$$\mathcal{H}_\tau = \frac{1}{2m}(g^{\mu\nu} p_\mu p_\nu + m^2), \tag{37}$$

which satisfies $\mathcal{H}_\tau = 0$ on shell. Hamilton's equations in proper time are

$$\frac{dr}{d\tau} = \frac{1}{m}\frac{\partial \mathcal{H}}{\partial p_r} = \frac{g^{rr}p_r}{m}, \tag{38}$$

$$\frac{d\phi}{d\tau} = \frac{1}{m}\frac{\partial \mathcal{H}}{\partial p_\phi} = \frac{g^{\phi\phi}p_\phi + g^{t\phi}p_t}{m}, \tag{39}$$

$$\frac{dp_t}{d\tau} = -\frac{1}{m}\frac{\partial \mathcal{H}}{\partial t} = 0, \tag{40}$$

$$\frac{dp_r}{d\tau} = -\frac{1}{m}\frac{\partial \mathcal{H}}{\partial r}, \tag{41}$$

$$\frac{dp_\phi}{d\tau} = -\frac{1}{m}\frac{\partial \mathcal{H}}{\partial \phi} = 0. \tag{42}$$

The $1/m$ factor converts covariant momentum to four-velocity in proper-time evolution; geodesic motion remains mass independent, consistent with the equivalence principle. Note that $p_t = -E$ and $p_\phi = L_z$ are conserved for geodesic motion.

VI. THRUST MODEL AND EQUATIONS OF MOTION

Notation reminder. Throughout this section: m denotes the rocket's (decreasing) rest mass; μ is the cumulative exhaust rest mass ejected; $\delta m_{\rm nom}$ is the nominal burn-size parameter supplied to each run, while $\delta\mu$ and $\delta m_{\rm actual}$ are the exact exhaust rest mass and realized rocket mass loss, respectively (Sec. IV). The exhaust speed in the rocket frame is v_e with Lorentz factor $\gamma_e = (1 - v_e^2)^{-1/2}$. Killing energy is $E = -p_t$; specific energy is E/m; exhaust Killing energy is $E_{\rm ex} = -u_{{\rm ex},t}$ [Eq. (55)].

A. Variable-mass rocket dynamics

For a rocket with proper acceleration $a_{\rm prop}$ and exhaust velocity v_e, the mass evolution follows the rest-frame mass-flow relation

$$\frac{dm}{d\tau} = -\frac{T}{v_e}, \tag{43}$$

where $T = m \cdot a_{\rm prop}$ is the thrust magnitude.

With v_e denoting the exhaust speed in the rocket rest frame and Lorentz factor $\gamma_e = 1/\sqrt{1 - v_e^2}$, the rocket mass loss per unit exhaust rest mass is $-dm = \gamma_e d\mu$ (accounting for both rest mass and kinetic energy in the rocket frame), i.e., $dm/d\tau = -\gamma_e d\mu/d\tau$. In the rocket rest frame, the exhaust carries spatial momentum $\gamma_e v_e d\mu$ per ejection event, so the thrust magnitude is $T = \gamma_e v_e d\mu/d\tau$. Substituting $d\mu/d\tau = (-dm/d\tau)/\gamma_e$ from the mass relation, the γ_e factors cancel,

$$T = \gamma_e v_e \frac{d\mu}{d\tau} = \gamma_e v_e \cdot \frac{-dm/d\tau}{\gamma_e} = v_e\left(-\frac{dm}{d\tau}\right), \tag{44}$$

consistent with Eq. (43).

The nongravitational 4-force on the spacecraft is defined by

$$f^\mu := \frac{Dp^\mu}{D\tau} - u^\mu \frac{dm}{d\tau} = ma^\mu, \tag{45}$$

where $D/D\tau$ denotes the covariant derivative along the worldline and a^μ is the 4-acceleration. The force satisfies the orthogonality condition

$$u_\mu f^\mu = 0, \tag{46}$$

which preserves the 4-velocity normalization $u_\mu u^\mu = -1$; the evolving mass shell $p_\mu p^\mu = -m^2(\tau)$ changes consistently through $dm/d\tau$.

B. Thrust direction

To maintain physically meaningful thrust magnitude, we construct an orthonormal spatial basis in the rocket's rest frame via Gram-Schmidt orthogonalization in the (t, r, ϕ) subspace. Starting from the rocket's 4-velocity u^μ and coordinate basis vectors $\hat{r}^\mu = (0, 1, 0)$ and $\hat{\phi}^\mu = (0, 0, 1)$, we project each onto the subspace orthogonal to u^μ using $P^\mu{}_\nu = \delta^\mu_\nu + u^\mu u_\nu$ (where $u_\mu u^\mu = -1$). The radial leg is

$$\tilde{e}^\mu_{(r)} = P^\mu{}_\nu \hat{r}^\nu = \hat{r}^\mu + (u_\alpha \hat{r}^\alpha)u^\mu, \tag{47}$$

$$e^\mu_{(r)} = \frac{\tilde{e}^\mu_{(r)}}{\sqrt{g_{\alpha\beta}\tilde{e}^\alpha_{(r)}\tilde{e}^\beta_{(r)}}}. \tag{48}$$

The azimuthal leg additionally removes the component along the unnormalized radial leg,

$$\tilde{e}^\mu_{(\phi)} = P^\mu{}_\nu \hat{\phi}^\nu - \frac{g_{\alpha\beta}\tilde{e}^\alpha_{(r)}(P^\beta{}_\nu \hat{\phi}^\nu)}{g_{\alpha\beta}\tilde{e}^\alpha_{(r)}\tilde{e}^\beta_{(r)}}\tilde{e}^\mu_{(r)}, \tag{49}$$

$$e^\mu_{(\phi)} = \frac{\tilde{e}^\mu_{(\phi)}}{\sqrt{g_{\alpha\beta}\tilde{e}^\alpha_{(\phi)}\tilde{e}^\beta_{(\phi)}}}, \tag{50}$$

satisfying $u_\mu e^\mu_{(i)} = 0$ and $e_{(i)\mu}e^\mu_{(j)} = \delta_{ij}$. The unit thrust direction is

$$s^\mu = \sin\alpha e^\mu_{(r)} + \sigma\cos\alpha e^\mu_{(\phi)}, \qquad \sigma \in \{+1, -1\}, \tag{51}$$

and $f^\mu = Ts^\mu$, enforcing both $u_\mu f^\mu = 0$ and $f_\mu f^\mu = +T^2$.

C. Optimal thrust selection

The azimuthal thrust sign ($\pm$) is chosen to maximize instantaneous energy gain,

$$\text{Choose sign}(\pm) \quad \text{to maximize} -\frac{dp_t}{d\tau}. \tag{52}$$

Since $E = -p_t$, this maximizes the rate of energy extraction.

D. Exact 4-momentum conservation

For exact energy accounting, we use 4-momentum conservation rather than approximate thrust models. When the rocket ejects exhaust of rest mass $\delta\mu$ with 4-velocity $u^\mu_{\rm ex}$, conservation requires

$$p'_\mu = p_\mu - \delta\mu \cdot u_{{\rm ex},\mu}, \tag{53}$$

where p'_μ is the new rocket 4-momentum. The exhaust 4-velocity is

$$u^\mu_{\rm ex} = \gamma_e(u^\mu - v_e s^\mu), \tag{54}$$

where s^μ is the unit spacelike thrust direction (the direction of the reaction force on the rocket) and $\gamma_e = 1/\sqrt{1-v_e^2}$. Since the exhaust is ejected opposite to the thrust, the minus sign ensures correct physics. The covariant exhaust 4-velocity is $u_{{\rm ex},\mu} = g_{\mu\nu}u^\nu_{\rm ex}$.

The exhaust Killing energy is

$$E_{\rm ex} = -u_{{\rm ex},t} = -\gamma_e(u_t - v_e s_t). \tag{55}$$

Here $E_{\rm ex}$ is the exhaust Killing energy per unit exhaust rest mass, since $u^\mu_{\rm ex}$ is a 4-velocity (unit timelike vector). Negative $E_{\rm ex}$ is the local kinematic signature that the exhaust occupies a negative-Killing-energy state; such exhaust is confined to the ergoregion and is captured (cf. Sec. II and Ref. [17]). Absorption of this negative Killing energy would, if backreaction were included, decrease the hole's mass parameter and rotational energy.

For continuous thrust, the rates are

$$\frac{dp_\mu}{d\tau} = -\frac{d\mu}{d\tau}u_{{\rm ex},\mu}, \tag{56}$$

$$\frac{dm}{d\tau} = -\gamma_e\frac{d\mu}{d\tau}, \qquad -dm = \gamma_e d\mu, \tag{57}$$

where $d\mu/d\tau$ is the exhaust rest-mass ejection rate. The relation between rocket mass loss and thrust magnitude is

$$T = \gamma_e v_e\frac{d\mu}{d\tau} = v_e\left(-\frac{dm}{d\tau}\right). \tag{58}$$

The last equality in Eq. (58) follows directly by substituting Eq. (57) into $T = \gamma_e v_e d\mu/d\tau$. Similarly, substituting Eq. (54) into Eq. (56) gives

$$\frac{dp_\mu}{d\tau} = -\gamma_e\frac{d\mu}{d\tau}(u_\mu - v_e s_\mu).$$

Taking the t component and using $E = -p_t$ and Eq. (55), one obtains $dE/d\tau = -(d\mu/d\tau)E_{\rm ex}$. Integrating over the burn gives the continuous form of energy conservation, which in the discrete implementation is verified by

$$\Delta E + \sum_i E_{{\rm ex},i}\cdot\delta\mu_i = 0, \tag{59}$$

where the sum is over all thrust events. This relation is satisfied to residuals $< 10^{-9}$ in our simulations, as it follows from the exact 4-momentum conservation imposed at each thrust event.

E. Full equations of motion

Between thrust events, the spacecraft follows geodesic Hamiltonian evolution [Eqs. (38)–(42)]. In continuous-thrust mode, we add the covariant momentum source implied by exact 4-momentum conservation with the exhaust,

$$\left(\frac{dp_\mu}{d\tau}\right)_{\rm thrust} = -\frac{d\mu}{d\tau}u_{{\rm ex},\mu}, \qquad \frac{dm}{d\tau} = -\gamma_e\frac{d\mu}{d\tau}, \tag{60}$$

with $u_{{\rm ex},\mu} = \gamma_e(u_\mu - v_e s_\mu)$ from Eq. (54). The full forced system can be written compactly as

$$\frac{dp_\mu}{d\tau} = \left(\frac{dp_\mu}{d\tau}\right)_{\rm geo} + \left(\frac{dp_\mu}{d\tau}\right)_{\rm thrust}, \tag{61}$$

where $(dp_t/d\tau)_{\rm geo} = (dp_\phi/d\tau)_{\rm geo} = 0$ and $(dp_r/d\tau)_{\rm geo} = -(1/m)\partial\mathcal{H}/\partial r$ [Eq. (41)]. Position updates use $p^\mu = g^{\mu\nu}p_\nu$ and $dx^\mu/d\tau = p^\mu/m$. In impulsive-thrust mode, we apply the discrete version of Eq. (53) at the burn event and then resume geodesic evolution.

F. Throttle function

We activate thrust only inside the extraction zone where Penrose extraction ($E_{\rm ex} < 0$) is achievable. We compute the extraction limit radius $R_{\rm extraction}$ dynamically as the maximum radius where $E_{\rm ex} < 0$ is geometrically possible given the current orbit parameters. We implement this using a smooth throttle function $\chi(r)$,

$$\chi(r) = \begin{cases} 0, & r \le r_{\rm in} \text{ or } r \ge r_{\rm out} \\ S\left(\frac{r-r_{\rm in}}{w}\right)\left[1 - S\left(\frac{r-r_{\rm out}}{w}\right)\right], & \text{otherwise} \end{cases}, \tag{62}$$

where $r_{\rm in} \coloneqq r_+ + \epsilon_{\rm in}$ and $r_{\rm out} \coloneqq R_{\rm extraction} - \epsilon_{\rm out}$ enforce safety margins near the horizon and the extraction boundary; $S(x)$ is a quintic smoothstep providing C^2-continuous ramp transitions, and $w = 0.15(r_{\rm out} - r_{\rm in})$ sets the ramp width proportional to the extraction zone thickness. We use $\epsilon_{\rm in} = 5 \times 10^{-3} M$ and $\epsilon_{\rm out} = 5 \times 10^{-3} M$.

We determine $R_{\rm extraction}$ by scanning radii outward from r_+ and, at each radius, scanning over all thrust directions (α, σ) to find the minimum achievable $E_{\rm ex}$. The largest radius at which $\min_{\alpha,\sigma} E_{\rm ex} < 0$ defines $R_{\rm extraction}$. This depends on the current orbit parameters (E, L_z), the current mass m, exhaust velocity v_e, and spin a; for single-impulse burns (applied at periapsis during geodesic infall), these coincide with the initial values $(E_0, L_{z,0})$, while for continuous thrust $R_{\rm extraction}$ is recomputed from the instantaneous state. Note that $R_{\rm extraction}$ represents the region where negative-energy exhaust is kinematically accessible for at least one thrust direction (the one minimizing $E_{\rm ex}$); the actual controller [proportional–integral–derivative (PID) for α, greedy sign choice for σ] does not guarantee $E_{\rm ex} < 0$ at every active step within this region.

The smoothstep function is

$$S(x) = \begin{cases} 0 & x \le 0 \\ 6x^5 - 15x^4 + 10x^3 & 0 < x < 1 \\ 1 & x \ge 1 \end{cases}. \tag{63}$$

1. Continuous-thrust amplitude

The proper acceleration is modulated spatially by the throttle function,

$$a_{\rm prop}(r) = a_{\max} \chi(r), \tag{64}$$

where $a_{\max}$ is the nominal maximum proper acceleration (a simulation parameter; we use $a_{\max} = 0.15$ in escape-optimized runs and $a_{\max} = 0.35$ in capture-mode runs, in units where $M = c = 1$). Combining with the mass-flow relation (Sec. IV), the exhaust rest-mass ejection rate is

$$\frac{d\mu}{d\tau} = \frac{m a_{\max} \chi(r)}{\gamma_e v_e}, \tag{65}$$

and the thrust magnitude reads

$$T(r) = \gamma_e v_e \frac{d\mu}{d\tau} = m a_{\max} \chi(r), \tag{66}$$

which decreases with fuel consumption ($m \to m_f$) and vanishes outside the extraction zone ($\chi = 0$ for $r \notin [r_{\rm in}, r_{\rm out}]$). Burning continues while $\chi(r) > 0$ and $m > m_{\min}$ (fuel budget); we use $m_{\min} = 0.1 m_0$. Equations (64)–(66) close the continuous-thrust ordinary differential equation system.

G. Steering and protocol summary

Our Monte Carlo experiments use explicit steering prescriptions rather than solving a global optimal-control problem. For reproducibility, we summarize the protocol used in both impulse and continuous-thrust simulations:

(1) Thrust modes. (i) Periapsis impulse: a single burn is executed at the first periapsis passage, expending a prescribed nominal burn-size parameter $\delta m_{\rm nom}$ (a labeling parameter; see Sec. IV for the label/actual distinction). (ii) Continuous thrust: thrust is applied continuously while the throttle $\chi(r)$ is nonzero (Sec. VI F) until the fuel budget is expended or a termination event occurs (escape/capture/constraint failure).

(2) Thrust direction selection (impulse mode). For periapsis-impulse burns, the thrust direction (α, σ) is jointly optimized by exhaustive scan over 37 uniformly spaced angles $\alpha \in [-\pi/2, \pi/2]$ with both azimuthal signs $\sigma \in \{+1, -1\}$ (74 candidate directions total), selecting the direction that minimizes the exhaust Killing energy $E_{\rm ex}$ subject to the postimpulse unbound-energy filter $E > m_{\rm after}$ (necessary but not sufficient for escape; actual escape is verified by trajectory integration, Sec. IV). This greedy direction selection is the "minimize $E_{\rm ex}$" prescription referenced throughout.

(3) Thrust direction selection (continuous mode). The azimuthal sign $\sigma \in \{+1, -1\}$ is chosen greedily at each step to maximize instantaneous Killing-energy gain $(-dp_t/d\tau)$. The radial steering angle α is set by a PID controller targeting $r_{\rm set} \sim 1.5M$ to extend ergosphere dwell time while avoiding horizon approach (Sec. VII D).

(4) Exact energy accounting. Burn updates use exact 4-momentum conservation with an explicit exhaust 4-velocity (Sec. VI D). Exhaust with $E_{\rm ex} < 0$ cannot reach spatial infinity, since negative-energy worldlines cannot reach the region where ∂_t is timelike.

Why both strategies? A simple optimal-allocation argument (Sec. VIII D) shows that, for a fixed trajectory and exhaust-direction schedule, concentrating all ejection at the point where $E_{\rm ex}$ is most negative maximizes energy extraction. Single-impulse thrust approximates this optimum. Continuous thrust is studied to (i) provide numerical evidence for the path-averaging penalty predicted by this argument and (ii) connect to realistic rocket operation where truly instantaneous thrust is an idealization.

VII. NUMERICAL METHODS

We work in geometric units $G = c = M = 1$ throughout; exhaust velocities quoted as fractions of c (e.g., $v_e = 0.95c$) are dimensionless with $c = 1$ implicit. We adopt the test-particle limit ($m_0 \ll M$) so that backreaction and self-force are negligible. We normalize the spacecraft's

initial rest mass to $m_0 = 1$; the initial specific energy $E_0 = -p_t^{(0)}/m_0$ and specific angular momentum $L_{z,0} = p_\phi^{(0)}/m_0$ are therefore dimensionless. These are independent normalizations: $G = c = M = 1$ sets the spacetime scale, while $m_0 = 1$ is a bookkeeping convention for the spacecraft sector; the physical test-particle assumption $m_0/M \to 0$ is maintained throughout. As mass evolves, p_t and p_ϕ are the total (not specific) covariant momenta; we track the specific energy $E/m = -u_t$ to diagnose escape ($E/m > 1$ at large radius).

A. Solver selection

For parameter sweeps (single-impulse trajectories), we use SciPy's eighth-order Dormand-Prince (DOP853) integrator (eighth-order Dormand-Prince) with adaptive step control, achieving near conservation of Killing quantities during geodesic phases (mass-shell residual $< 10^{-9}$; energy-balance residual $< 10^{-9}$; see Appendix A). For continuous-thrust simulations, we employ classical fourth-order Runge-Kutta (RK4) with fixed time step $\Delta\tau = 0.005M$ and mass-shell projection after each step (enabling transparent debugging of near-horizon physics and thrust events). The projection modifies p_r to satisfy $g^{\mu\nu}p_\mu p_\nu + m^2 = 0$ while preserving $E = -p_t$.

Integrator validation. We verified that DOP853, RK45, and Radau methods yield identical trajectory outcomes (escape vs capture) for all tested initial conditions. Convergence tests confirm that success rates are stable across time steps $\Delta\tau \in [0.002, 0.02]M$ when projection is applied (see Appendix A).

B. Termination events

We terminate integration upon horizon approach ($r < r_+ + 0.02M$), mass depletion ($m < 0.1$), constraint violation ($|\mathcal{C}| > 10^{-1}$), or escape ($r > 50M$ with $dr/d\tau > 0$ and $E/m > 1$).

C. Initial conditions

Initial conditions must represent a valid geodesic. Given $(r_0, E_0, L_{z,0}, m_0)$ with $r_0 = 15M$ (outside the ergosphere, allowing infall to be geodesic), we set

$$p_t^{(0)} = -E_0, \tag{67}$$

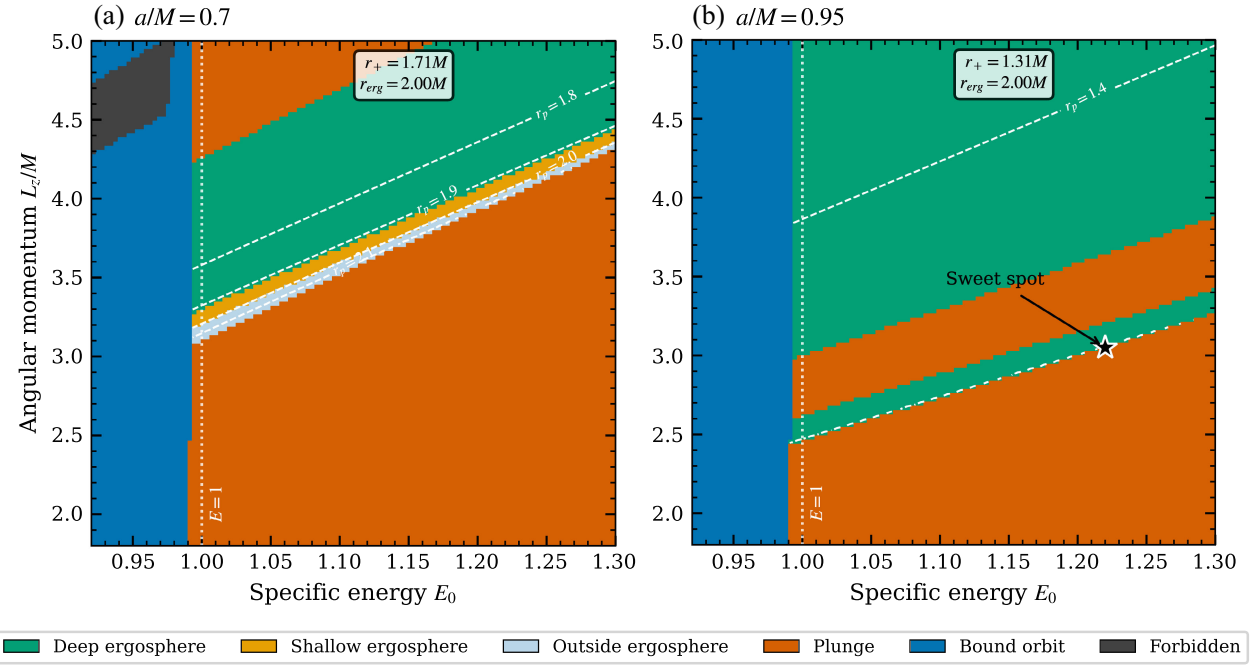


FIG. 1. Orbit classification in (E_0, L_z) parameter space comparing moderate and high spin. (a) $a/M = 0.70$: Six classification regions are present—forbidden (dark gray, inaccessible configurations), bound (blue, $E_0 < 1$), plunge (red, no turning point), outside ergosphere (light blue, $r_{\rm peri} > r_{\rm erg}$), shallow ergosphere (orange, outer 30% of ergosphere), and deep ergosphere (green, deep-ergosphere diagnostic region). The deep-ergosphere region is negligible due to the narrow-ergosphere width ($0.29M$). (b) $a/M = 0.95$: Higher spin dramatically expands the deep-ergosphere region [43.7% of the plotted (E_0, L_z) domain] at the expense of the outside region (now absent) and the shallow-ergosphere band (greatly reduced). The wider ergosphere ($0.69M$) and stronger frame dragging enable more trajectories to reach the deep-ergosphere region. Note: the actual $E_{\rm ex} < 0$ boundary ($R_{\rm extraction}$) is state-dependent and computed dynamically (Sec. VI F); this fixed-depth diagnostic cut provides a geometric proxy, not the physical extraction condition. Dashed white contours show constant periapsis r_p; vertical dotted line marks $E_0 = 1$ (bound/unbound threshold). Star marker in (b) indicates the sweet spot $(E_0, L_z) = (1.22, 3.05)$.

$$p_\phi^{(0)} = L_{z,0}, \tag{68}$$

$$p_r^{(0)} = -\sqrt{\frac{-(g^{tt}p_t^{(0)2} + 2g^{t\phi}p_t^{(0)}p_\phi^{(0)} + g^{\phi\phi}p_\phi^{(0)2} + m_0^2)}{g^{rr}}}, \tag{69}$$

where the negative sign for $p_r^{(0)}$ indicates an initially infalling trajectory (moving toward smaller r).

For the initial state to be physical (i.e., within the allowed region for geodesic motion), we require

$$g^{tt}p_t^2 + 2g^{t\phi}p_t p_\phi + g^{\phi\phi}p_\phi^2 + m^2 < 0, \tag{70}$$

ensuring $p_r^2 > 0$ (real radial momentum).

D. Radial controller

To extend ergosphere duration, a PID controller adjusts the radial steering angle: $\alpha = \text{clamp}[k_p(r_{\text{set}} - r) - k_d\dot{r} + k_i\int(r_{\text{set}} - r)d\tau]$, with gains $k_p = 50$, $k_d = 30$, $k_i = 5$, maximum angle $\alpha_{\max} = 60°$, and target radius $r_{\text{set}} \sim 1.5M$.

VIII. NUMERICAL RESULTS

We present systematic numerical experiments characterizing the Penrose process. The main mission family is initially unbound equatorial prograde flybys; broader scans also include nearby bound and retrograde initial conditions to map failure boundaries. Our central finding is that successful energy extraction with spacecraft escape is statistically rare in broad, untuned scans, occurring in a narrow region of parameter space, but is numerically verified to the stated tolerances when achieved. In the focused Gaussian sweet-spot prior, success rates can reach ∼70% (Sec. VIII G). All simulations use geometric units $G = c = M = 1$.

Scope notes. Most ensemble statistics in this section are for single periapsis impulses; continuous-thrust trajectories are analyzed separately in Sec. VIII E on representative matched trajectories only. All quoted thresholds (spin, velocity onset, success rates) are specific to the equatorial prograde-flyby family, the Gaussian sweet-spot prior, and the steering prescriptions described in Sec. VI G; different orbit families, priors, or steering laws may shift these thresholds.

Organization of results. Our analysis proceeds as follows: Sec. VIII A validates numerical accuracy; Sec. VIII B identifies the extraction window using orbit classification (Fig. 1); Sec. VIII C quantifies success rates via Monte Carlo sampling (Tables I–III and Fig. 2); Sec. VIII E compares thrust strategies (Table IV and Fig. 3); Sec. VIII F examines spin dependence (Fig. 4); Sec. VIII G characterizes the velocity-onset threshold (Tables V–VIII, Figs. 5, and 6); and Sec. VIII H synthesizes these findings. The optimality argument (Sec. VIII D) motivates the thrust strategy comparison by establishing, *a priori*, why single-impulse ejection should outperform continuous thrust.

TABLE I. Broad parameter sweep results ($N = 6400$ samples per spin, $E_0 \in [0.95, 2.0]$, $L_z \in [-3.0, 6.0]$, $v_e = 0.95c$, $\delta m_{\text{nom}} = 0.20$). All confidence intervals (CI) use the Clopper-Pearson exact method. The success rate quantifies trajectories achieving extraction with escape [Eq. (18)]: $\Delta E > 0$ and escape to infinity. Integration uses DOP853 with relative and absolute tolerances rtol $= 10^{-9}$, atol $= 10^{-11}$.

a/M	N	Escape (%)	Success (%)	95% CI
0.50	6400	8.11	0.00	[0.00%, 0.06%]
0.70	6400	8.34	0.00	[0.00%, 0.06%]
0.90	6400	13.94	0.03	[0.00%, 0.11%]
0.95	6400	19.52	0.95	[0.73%, 1.22%]
0.99	6400	27.22	1.41	[1.13%, 1.73%]

TABLE II. Spin threshold analysis ($N = 10,000$ LHS samples per spin in rectangular focused region: $E_0 \in [1.1, 1.3]$, $L_z \in [2.5, 3.5]$; single periapsis-impulse burns, $v_e = 0.95c$, $\delta m_{\text{nom}} = 0.3$). For compactness, the row $a/M \le 0.88$ aggregates nine spins $a/M = 0.80, 0.81, \ldots, 0.88$ (90,000 samples total). No successful extractions for $a/M \le 0.88$ and nonzero rates at $a/M = 0.89$ constrain the practical threshold to $0.88 < a_{\text{crit}}/M \lesssim 0.89$.

a/M	N	Success rate (%)	95% CI
≤ 0.88	90,000	0.00	[0.00%, 0.004%]
0.89	10,000	0.09	[0.04%, 0.17%]
0.90	10,000	0.30	[0.20%, 0.43%]
0.92	10,000	2.47	[2.18%, 2.79%]
0.95	10,000	8.24	[7.71%, 8.80%]
0.99	10,000	10.92	[10.32%, 11.55%]

TABLE III. Focused rectangular sweet-spot sweep results (uniform grid, $N = 3600$ samples per spin, $E_0 \in [1.1, 1.4]$, $L_z \in [2.5, 3.8]$, $v_e = 0.95c$, $\delta m_{\text{nom}} = 0.20$; single periapsis-impulse burns). The ∼12× increase in extraction-with-escape success rate (0.95% → 11.19% for $a/M = 0.95$) and ∼10× increase (1.41% → 13.69% for $a/M = 0.99$) confirm that successful extraction with escape depends critically on precise tuning of initial conditions. (Note: rates differ from Table II due to different sampling domains and methods. The Gaussian sweet-spot prior used in Sec. VIII G employs a different sampling method.).

a/M	N	Escape (%)	Success (%)	95% CI
0.50	3600	0.00	0.00	[0.00%, 0.10%]
0.70	3600	2.78	0.00	[0.00%, 0.10%]
0.90	3600	18.64	0.58	[0.36%, 0.89%]
0.95	3600	26.33	11.19	[10.18%, 12.27%]
0.99	3600	65.14	13.69	[12.59%, 14.86%]

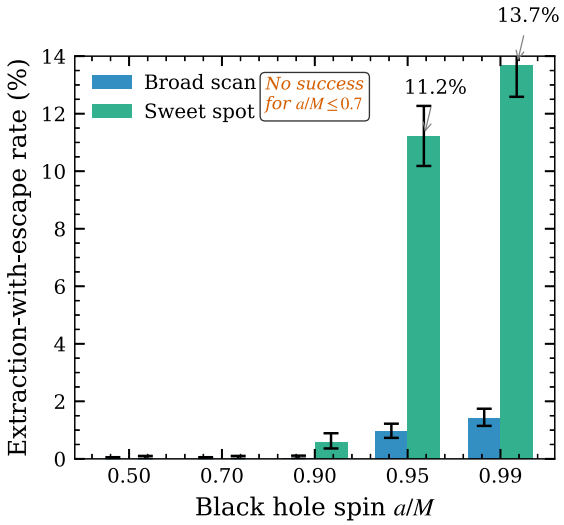


FIG. 2. Extraction-with-escape success rate vs black hole spin from systematic parameter sweep: broad scan ($E_0 \in [0.95, 2.0]$, $L_z \in [-3.0, 6.0]$; blue, $N = 6400$ per spin) and rectangular sweet-spot region (uniform grid, $E_0 \in [1.1, 1.4]$, $L_z \in [2.5, 3.8]$; green, $N = 3600$ per spin). Error bars show 95% Clopper-Pearson exact confidence intervals. No successful extractions were observed for $a/M \leq 0.7$ across 12,800 trajectories. Focusing on the sweet spot yields amplification factors of 19× at $a/M = 0.9$, 12× at $a/M = 0.95$, and 10× at $a/M = 0.99$, demonstrating the essential role of precise initial condition tuning.

TABLE IV. Thrust strategy comparison for $a/M = 0.95$, $v_e = 0.95c$. Each column reports a representative single trajectory at the listed initial conditions. "Cont. (Esc.)" and "Cont. (Capture)" refer to the spacecraft outcome, not the exhaust fate. We report $\eta_{\text{cum}} = \Delta E/\Delta m$ [propellant normalized; Eq. (19)], the traditional Penrose efficiency $\eta_{\text{trad}} = \Delta E/E_0$ [Eq. (23)], and the per-exhaust-rest-mass ratio $\Delta E/\delta\mu$ where $\delta\mu$ is computed via Eq. (22) for the single-impulse case and via the exact differential relation $-dm = \gamma_e d\mu$ for continuous thrust. The single-impulse $\delta m = 0.22$ is the actual consumed mass fraction (see Sec. IV for the finite-impulse discussion). These are illustrative matched trajectories; broader single-impulse ensemble statistics appear in Tables V–VIII. Efficiency metrics (η_{cum}, $\Delta E/\delta\mu$) are computed from full-precision simulation values; displayed ΔE and δm_{actual} are independently rounded.

Metric	Single impulse	Continuous (Escape)	Continuous (Capture)
(E_0, L_z)	$(1.20, +3.0)$	$(1.20, +3.0)$	$(1.20, +2.8)$
δm_{actual}	0.22	0.07	0.06
ΔE	+0.013	+0.002	+0.003
η_{cum}	5.7%	3.7%	5.0%
η_{trad}	1.1%	0.2%	∼0.3%
$\Delta E/\delta\mu$	∼20%	∼12%	∼16%
Fraction of events $E_{\text{ex}} < 0$	100%	100%	∼100%
Outcome	Escape	Escape	Capture

Broad rarity claims derive from the grid/Latin hypercube sampling (LHS) sweeps (Tables I–III), threshold claims from the focused Gaussian scans (Tables V–VIII), and continuous-thrust comparisons from representative matched trajectories (Table IV and Fig. 3).

A. Numerical validation

Our implementation preserves the mass-shell constraint to $|\mathcal{C}|_{\max} < 10^{-9}$ (single-impulse) and $< 10^{-6}$ (continuous thrust with projection). We verify future-directedness ($u^t > 0$) and timelike normalization throughout. Energy conservation via exact 4-momentum conservation yields residuals $< 10^{-9}$. Detailed diagnostics are provided in Appendix A.

B. Extraction window in parameter space

We require prograde-flyby orbits satisfying: (1) $E_0 > 1$ (unbound), (2) periapsis $r_{\text{peri}} < r_{\text{erg}} = 2M$ (ergosphere penetration), and (3) exhaust direction chosen to minimize E_{ex} (in the successful sampled cases, the minimizing direction is locally retrograde and carries sufficiently negative angular momentum $L_{z,\text{ex}} < 0$ to permit $E_{\text{ex}} < 0$).

For diagnostics and visualization (Figs. 1 and 4), we further classify ergosphere-penetrating flybys by periapsis depth. Let $\Delta r_{\text{erg}} \coloneqq r_{\text{erg}} - r_+$ denote the equatorial ergosphere thickness. We call a trajectory "shallow" if its periapsis lies in the outer 30% of the ergosphere, $r_{\text{peri}} > r_{\text{erg}} - 0.3\Delta r_{\text{erg}}$, and "deep" otherwise; deep-ergosphere flybys correspond to the diagnostic region where exhaust with sufficiently negative angular momentum ($L_{z,\text{ex}} < 0$) can most robustly attain $E_{\text{ex}} < 0$, though the actual extraction boundary $R_{\text{extraction}}$ depends on the instantaneous orbital state (Sec. VI F).

For $a/M = 0.95$, we find the viable region to be remarkably narrow: specific energy $E_0 \in [1.15, 1.35]$ and specific angular momentum $L_z \in [2.9, 3.3]$, yielding periapsis at $r_{\text{peri}} \approx 1.5M$. Orbits outside this window fail in different ways: initially bound conditions ($E_0 < 1$) remain bound under our baseline burns; low-L_z cases are captured; high L_z places periapsis outside the ergosphere; and retrograde ($L_z < 0$) cases do not yield extraction with escape in the sampled domain. Figure 1 shows how orbit classification depends critically on spin: at $a/M = 0.70$, the deep-ergosphere region is negligible (confined to a thin band), while at $a/M = 0.95$, it expands to dominate 44% of the plotted (E_0, L_z) diagnostic domain [Fig. 1(b)].

C. Statistical ensemble analysis

To quantify the rarity of successful extraction with escape, we performed systematic grid and LHS studies over the parameter space (E_0, L_z, a).

Unless stated otherwise, the grid/LHS sweeps in this section use a single impulsive burn applied at periapsis. The

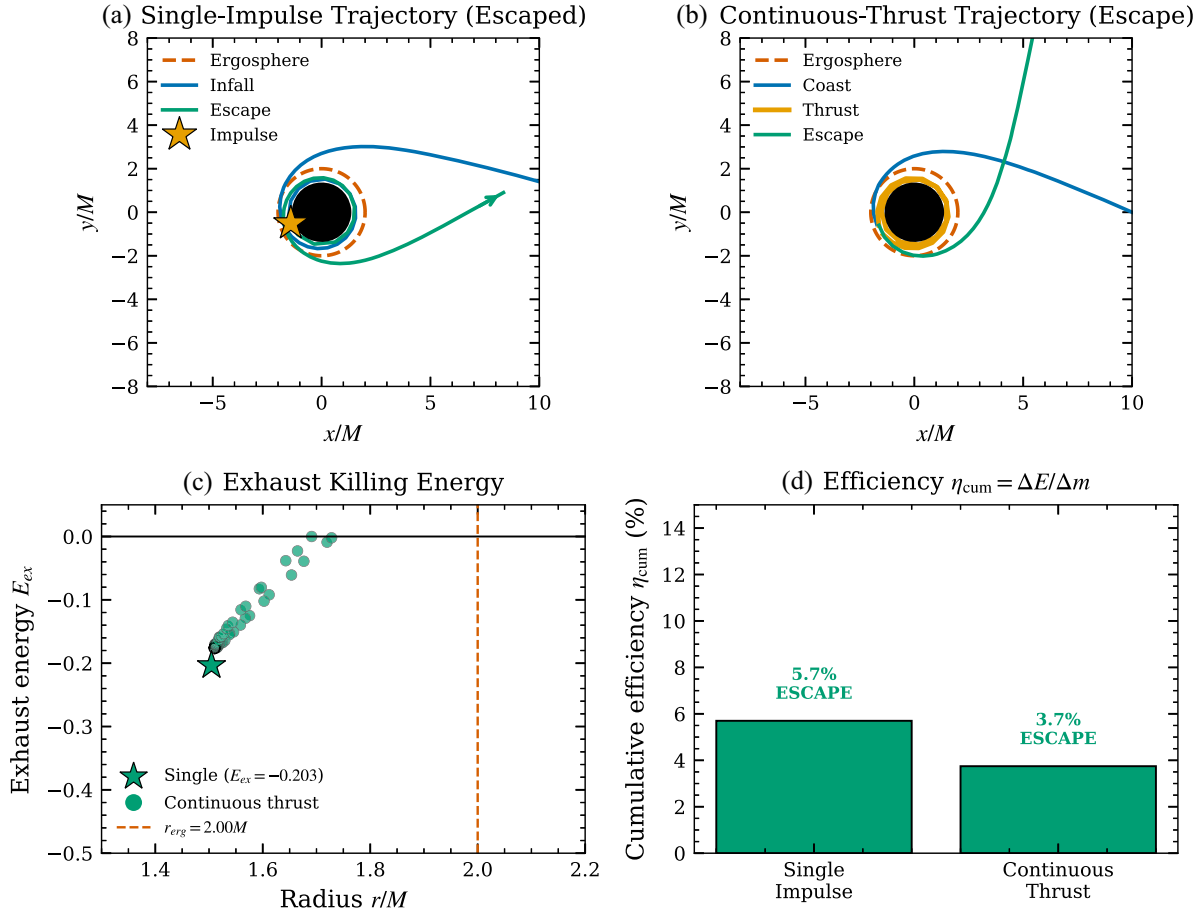


FIG. 3. Thrust strategy comparison at $a/M = 0.95$, $v_e = 0.95c$. (a) Single-impulse trajectory ($E_0 = 1.20$, $L_z = 3.0$): geodesic infall (blue), impulse at periapsis (orange star), geodesic escape (green). (b) Continuous-thrust trajectory ($E_0 = 1.20$, $L_z = 3.0$): coasting infall (blue), sustained thrust with ∼33,000 discrete events inside ergosphere (orange), escape (green). (c) Exhaust Killing energy $E_{\rm ex}$ vs radius: large star marks the representative single-impulse event ($E_{\rm ex} \approx -0.203$); small circles show continuous-thrust events. All events satisfy $E_{\rm ex} < 0$. Dashed line marks ergosphere boundary ($r_{\rm erg} = 2M$). (d) Cumulative efficiency $\eta_{\rm cum} = \Delta E/\Delta m$ [Eq. (19)] for these representative trajectories: single impulse achieves $\eta_{\rm cum} = 5.7\%$, continuous thrust 3.7% (exact values in Table IV). Note that the Wald single-decay benchmark (20.7%) uses energy normalization $\eta_{\rm trad} = \Delta E/E_0$ rather than propellant normalization $\eta_{\rm cum} = \Delta E/\Delta m$; the two metrics are not directly comparable.

broad scan domain ($E_0 \in [0.95, 2.0]$, $L_z \in [-3.0, 6.0]$) intentionally includes initially bound ($E_0 < 1$) and retrograde ($L_z < 0$) conditions to map the failure boundaries around the mission-relevant prograde unbound family, not because those conditions belong to the mission family itself. Continuous-thrust trajectories are analyzed separately in Sec. VIII E.

Statistical methodology. We use deterministic grids for systematic exploration and LHS for Monte Carlo validation. Confidence intervals for proportions use the Clopper-Pearson exact method [18]. For deterministic grid sweeps, these intervals quantify uncertainty with respect to the implied uniform measure over the chosen finite domain, not ordinary random-sampling uncertainty. Efficiency uncertainties were cross-checked via bias-corrected and accelerated bootstrap with 10^4 resamples [19]; reported efficiency values are sample means over successful escapes. Integration failures (11.2%) are conservatively counted as captures (Table VII).

1. Broad parameter study

We performed systematic parameter sweeps across a broad (E_0, L_z) domain for five spin values: $E_0 \in [0.95, 2.0]$, $L_z \in [-3.0, 6.0]$, using an 80×80 grid (6400 samples per spin). We choose this domain to encompass all orbits from marginally bound ($E_0 \approx 1$) to moderately relativistic ($E_0 = 2$), and from strongly retrograde to strongly prograde angular momentum, thereby encompassing the regime where ergosphere-penetrating flybys are expected for the spins in our range. We emphasize that success rates are conditional on this finite sampling domain; since E_0 and L_z range over noncompact intervals, the absolute probability depends on the chosen domain. Our purpose is to locate the viable extraction region and quantify its relative extent, not

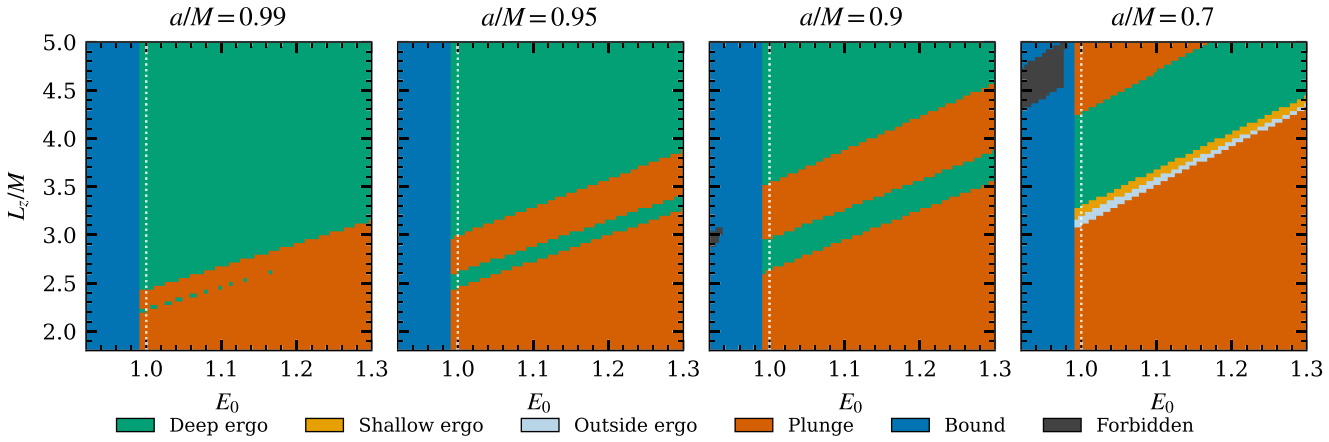


FIG. 4. Spin dependence of the extraction window. Panels show orbit classification in (E_0, L_z) space for $a/M = 0.99$, 0.95, 0.9, and 0.7 (left to right). Color coding follows Fig. 1: green indicates deep-ergosphere flyby, orange shallow ergosphere, light blue outside ergosphere, red plunge, blue bound ($E_0 < 1$), and gray forbidden. The deep-ergosphere region (deep-ergosphere diagnostic region; see Fig. 1 caption for the proxy/actual distinction) expands dramatically with increasing spin: from negligible at $a/M = 0.7$ to dominating the plotted (E_0, L_z) domain at $a/M = 0.99$. The bound region (left of vertical dotted line at $E_0 = 1$) is present at all spins but not relevant for extraction in our baseline impulsive scans, since these initial conditions are bound ($E_0 < 1$) and do not satisfy the escape criterion. Within this prograde-flyby family, high spin ($a/M \gtrsim 0.89$) is observed to be necessary for practical extraction with escape. At high spin (e.g., $a/M = 0.99$), the deep-ergosphere region can appear as two disjoint zones: one at moderate L_z where periapsis reaches moderate ergosphere depth, and another at lower L_z where orbits penetrate very close to the horizon. A thin plunge band separates them; some isolated points in the high-spin panel reflect finite grid resolution rather than truly disconnected regions. The $a/M = 0.95$ and 0.7 panels reproduce Fig. 1 at reduced scale for side-by-side spin comparison; see that figure for detailed zone annotations.

to assign an unconditional probability. Table I summarizes the results.

Because integration failures (11.2% of samples) are conservatively counted as captures (Appendix A), the reported success rates throughout this section are conservative lower estimates; exact percentages near marginal thresholds should be interpreted accordingly.

TABLE V. Thrust parameter sensitivity for $a/M = 0.95$ in the sweet-spot region, with $N = 500$ initial conditions per configuration sampled from a Gaussian distribution centered at $(E_0, L_z) = (1.22, 3.05)$ with $\sigma_E = 0.03$, $\sigma_{L_z} = 0.08$. Extraction-with-escape success rate as a function of exhaust velocity v_e and nominal burn-size parameter δm_{nom} (labeling parameter only; the actual consumed fraction $\delta m_{\mathrm{actual}}$ exceeds this value due to the finite-impulse correction discussed in Sec. IV; at $v_e = 0.95c$ the discrepancy ranges from ∼5% at $\delta m_{\mathrm{nom}} = 0.1$ to ∼34% at $\delta m_{\mathrm{nom}} = 0.4$). Representative actual consumed fractions: at $v_e = 0.95c$, $\delta m_{\mathrm{actual}} \approx 0.105$, 0.223, 0.361, 0.536 for $\delta m_{\mathrm{nom}} = 0.1$, 0.2, 0.3, 0.4, respectively; at $v_e = 0.98c$, $\delta m_{\mathrm{actual}} \approx 0.105$, 0.224, 0.365, 0.546. A sharp onset occurs around $v_e \approx 0.91$–$0.92c$; below this threshold, success is of order 1% or below. Peak success of 69.6% achieved at $v_e = 0.98c$, $\delta m_{\mathrm{nom}} = 0.4$. Corresponding cumulative efficiency trends are summarized in Fig. 5(b).

v_e/c	$\delta m_{\mathrm{nom}} = 0.1$ (%)	$\delta m_{\mathrm{nom}} = 0.2$ (%)	$\delta m_{\mathrm{nom}} = 0.3$ (%)	$\delta m_{\mathrm{nom}} = 0.4$ (%)
0.80	0.0	0.0	0.0	0.0
0.90	0.4	1.0	0.4	0.4
0.95	39.4	47.8	49.0	63.6
0.98	48.0	56.4	63.2	**69.6**

2. Spin threshold analysis

To precisely characterize the spin dependence, we performed a sweep over 14 spin values: $a/M = 0.80, 0.81, \ldots, 0.88$ plus $a/M \in \{0.89, 0.90, 0.92, 0.95, 0.99\}$, with 10,000 LHS samples per spin in the focused region. Table II shows the results.

Across this spin threshold sweep (140,000 additional LHS trajectories, separate from the main 104,000 in

TABLE VI. Simulation accounting breakdown. Phases 1–4 total 104,000 trajectories; phase 5 adds 140,000 LHS trajectories for spin threshold characterization; phases 6 and 7 add 76,000 additional sweep trajectories. Grand total: 320,000. Parameter sweeps use DOP853 (rtol $= 10^{-9}$, atol $= 10^{-11}$); continuous thrust uses classical RK4 with 4-momentum projection.

Phase	Description	Trajectories
1	Broad grid (80×80, 5 spins)	32,000
2	Focused grid (60×60, 5 spins)	18,000
3	Thrust sensitivity (4×4 grid)	24,000
4	Monte Carlo validation ($3 \times 2 \times 5000$)	30,000
5	Spin threshold ($14 \times 10,000$ LHS)	140,000
6	Velocity sweep, Fig. 5	40,000
7	Ultrarelativistic sweep, Fig. 6	36,000
	Total	320,000

Table VI), no successful Penrose extractions with escape were observed for $a/M \leq 0.88$. Combined with the onset of successful extraction at $a/M = 0.89$, this constrains the practical spin threshold for rocket-propelled Penrose extraction.

3. Focused sweet-spot study

We define the "sweet spot" as the region of (E_0, L_z) space corresponding to deep-ergosphere flybys with escape-capable angular momentum, identified from the orbit classification in Fig. 1. Physically, this is the intersection of two constraints: (i) periapsis deep enough inside the ergosphere for $E_{\rm ex} < 0$ to be robustly achievable and (ii) sufficient angular momentum for the postburn orbit to remain unbound. For $a/M = 0.95$, this region is approximately $E_0 \in [1.1, 1.4]$, $L_z \in [2.5, 3.8]$, centered near $(E_0, L_z) \approx (1.22, 3.05)$. We apply this fixed benchmark box (derived at $a/M = 0.95$) at all spins in Table III for comparability, not because it is optimal for each spin; a spin-optimized box would likely shift the sweet-spot center and could alter cross-spin success-rate rankings. Using a 60×60 grid (3600 samples per spin) restricted to this domain yields substantially higher success rates, as expected:

The $\sim 12\times$ amplification in success rate (see Fig. 2) confirms that precise initial condition tuning is essential.

D. Optimality of single-impulse extraction

Before comparing thrust strategies numerically, we establish a theoretical prediction: for a fixed realized trajectory and exhaust-direction schedule, hence a fixed profile $E_{\rm ex}(\tau)$ along the worldline, the total energy change is

$$\Delta E = -\int_0^{\tau_f} E_{\rm ex}(\tau) \cdot \frac{d\mu}{d\tau} d\tau, \tag{71}$$

where $E_{\rm ex}(\tau)$ is the exhaust Killing energy and $d\mu/d\tau$ is the exhaust rest-mass ejection rate (here μ denotes expelled rest mass, not a spacetime index). For fixed total exhaust mass and fixed $E_{\rm ex}(\tau)$ profile, maximizing ΔE requires concentrating ejection where $E_{\rm ex}$ is most negative.

Define normalized mass flow $\rho(\tau) = (1/\delta\mu_{\rm total})(d\mu/d\tau)$ with $\int \rho d\tau = 1$. Then $\Delta E = -\delta\mu_{\rm total}\langle E_{\rm ex}\rangle_\rho$, where the path average satisfies $\langle E_{\rm ex}\rangle_\rho \geq \min_\tau E_{\rm ex}(\tau)$. Maximum extraction occurs when $\rho(\tau) = \delta(\tau - \tau^*)$, concentrated at the minimum, precisely a single impulse at optimal location.

Caveats. (i) The δ-function optimality assumes no upper bound on instantaneous mass flow; if a constraint $(d\mu/d\tau) \leq \dot{\mu}_{\rm max}$ is imposed, the optimizer becomes "burn at maximum rate in an interval around τ^*" rather than a true impulse. (ii) This argument proves optimal allocation of a fixed fuel budget along a given worldline; global optimality over all steering laws is an optimal-control problem that we do not address here.

The exhaust Killing energy varies with radius: $E_{\rm ex} \sim 0$ at the ergosphere boundary, ~ -0.1 to -0.2 at $r \sim 1.5M$ (most favorable for escape in the sampled set), and more negative at smaller radii (but with increasing capture risk). Continuous thrust averages over this profile, suffering a path-averaging penalty. Additionally, continuous thrust modifies (E, L_z) and the periapsis geometry throughout the ergosphere transit; extended burning can move the spacecraft out of the narrow escape corridor before the maneuver completes, increasing capture probability. We now illustrate this comparison with representative matched-initial-state trajectories (noting that actual fuel consumption differs between strategies; see Table IV). Because the impulse and continuous controllers follow different realized trajectories and consume different actual mass fractions, the numerical comparison below illustrates the path-averaging penalty rather than constituting a controlled test of the fixed-worldline optimality result.

E. Thrust strategy comparison

We compare single-impulse and continuous-thrust strategies to test whether concentrating ejection near periapsis—where $|E_{\rm ex}|$ is largest—outperforms distributed thrust, as predicted by the optimality argument above (Sec. VIII D). The representative escape trajectories use $(E_0, L_z) = (1.20, 3.0)$ for both the single-impulse and continuous-thrust cases shown in Fig. 3; Table IV additionally includes a continuous-capture example at (1.20, 2.8) to illustrate the efficiency-capture trade-off. The actual mass fraction consumed $\delta m := (m_0 - m_f)/m_0$ differs between strategies because continuous thrust distributes ejection over the ergosphere transit while the impulse concentrates it at periapsis. Table IV summarizes the results.

Key findings. (1) Single impulse is most efficient among the strategies studied, achieving $\eta_{\rm cum} \approx 5.7\%$ (here $\delta m_{\rm actual} = 0.22$ for $\delta m_{\rm nom} = 0.20$; see Sec. IV for the label/actual distinction; $\Delta E/\delta\mu \approx 20\%$ per exhaust rest mass)[2] by concentrating fuel at periapsis, consistent with the optimal-allocation argument of Sec. VIII D. The traditional Penrose efficiency is $\eta_{\rm trad} \approx 1.1\%$ for the single impulse (0.2% for continuous escape), well below Wald's single-decay benchmark $\eta_{\rm Wald} \approx 20.7\%$, consistent with subextremal spin $(a/M = 0.95)$ and a nonoptimal orbit. (2) Continuous thrust enables escape with reduced efficiency ($\sim 60\%$ of single-impulse $\eta_{\rm cum}$), consistent with the path-averaging penalty. (3) Deeper periapsis increases $|E_{\rm ex}|$ but also captures probability; the continuous-capture trajectory achieves higher local precapture per-mass efficiency than continuous

[2]The Wald single-decay bound $\eta_{\rm Wald} \approx 20.7\%$ normalizes by incident particle energy $(\Delta E/E_0)$, not exhaust rest mass, so the two figures are not directly comparable despite their numerical proximity. We cite $\eta_{\rm Wald}$ as a benchmark for single-decay scenarios; it is not a strict upper bound on the controlled rocket setup studied here.

escape because all burning occurs deep in the ergosphere where $|E_{\rm ex}|$ is largest; however, this is not an extraction-with-escape success under our mission criterion [Eq. (18)], since the spacecraft does not escape to infinity (cf. the discussion of $\mathcal{S}$ in Sec. IV). (4) All illustrated escape trajectories achieve $E_{\rm ex} < 0$ for all thrust events inside the ergosphere; the continuous-capture case achieves $E_{\rm ex} < 0$ for nearly all thrust events [$P(E_{\rm ex} < 0) \sim 100\%$]. Figure 3 shows the trajectories reported in Table IV.

F. Spin dependence

This section complements the spin threshold analysis of Sec. VIII C by examining the geometric origin of the spin dependence through orbit classification maps, rather than success-rate statistics.

The extraction window depends sensitively on black hole spin. In the broad parameter study (Table I), no successful extractions were observed for $a/M \leq 0.7$ across 12,800 trajectories, while in the dedicated spin threshold analysis (Table II), the finer-grained result is $a_{\rm crit}/M \in (0.88, 0.89]$ using the narrower rectangular sweet-spot domain. These two findings are consistent: the broad study's coarser domain and lower spin values produce zero successes below $a/M = 0.7$, while the focused study, sampling more densely within the sweet spot, resolves the transition at higher spin. We stress that negative-energy timelike states inside the ergosphere exist for any nonzero Kerr spin; what vanishes at low spin in our analysis is the overlap between those states and the escape-capable prograde flybys in our sampled domain, compounded by finite Monte Carlo sampling resolution. Figure 4 illustrates the contraction of the extraction window with decreasing spin.

G. Thrust parameter sensitivity

Having established global conditional rarity and the sweet-spot amplification effect, we now turn to local sensitivity around the benchmark injection point. A critical finding from our systematic sweep is the sharp dependence on exhaust velocity v_e. To characterize this, we perform a local sensitivity study around a fixed benchmark prior at the $a/M = 0.95$ sweet spot, not a generic survey of all orbit families or steering laws: initial conditions are sampled from a Gaussian distribution centered on $(E_0, L_z) = (1.22, 3.05)$ with standard deviations $\sigma_E = 0.03$ and $\sigma_{L_z} = 0.08$. This Gaussian represents finite-precision orbit injection: a spacecraft targeting the optimal orbit will have small uncertainties in its orbital elements, and the Gaussian widths probe sensitivity to these uncertainties. We emphasize that this distribution is anchored to the $a/M = 0.95$ sweet spot and is not a spin-optimized prior; different spins may have different optimal centers. Our velocity analysis reveals a sharp onset around $v_e \approx 0.91$–$0.92c$: below this range, success is of order 1% or below in this Gaussian scan (Table V shows 0.4%–1% at $v_e = 0.90c$), while above it

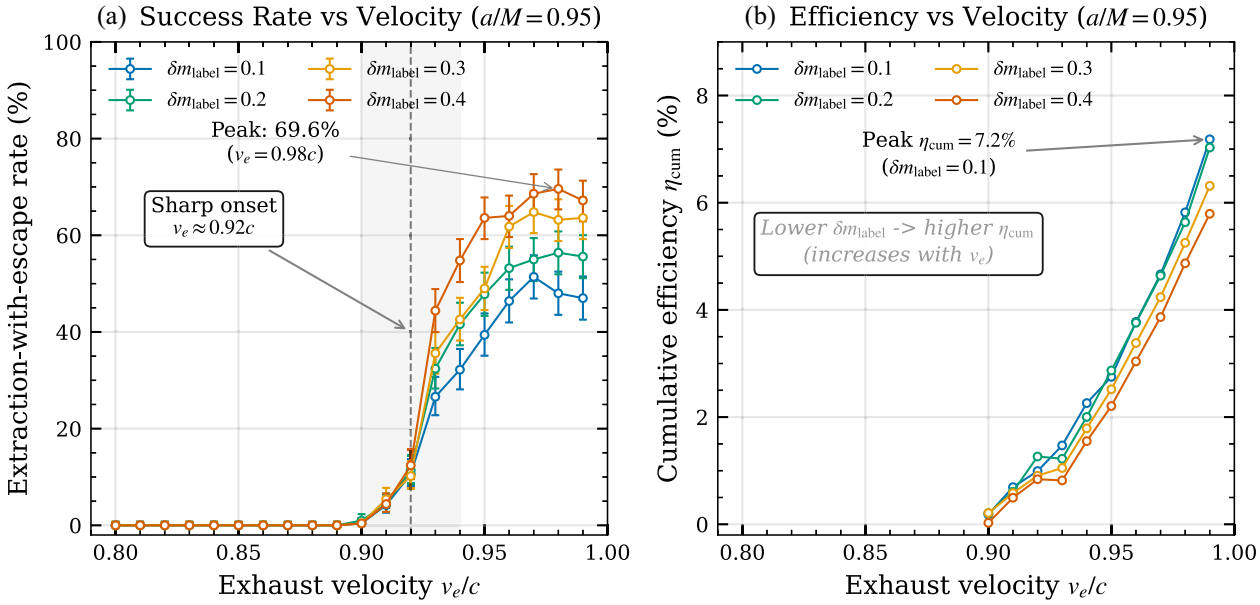


FIG. 5. Thrust parameter sensitivity for $a/M = 0.95$ at the sweet spot ($E_0 = 1.22$, $L_z = 3.05$), sampled at $0.01c$ velocity increments [$N = 500$ initial conditions per configuration from a Gaussian at (1.22,3.05) with $\sigma_E = 0.03$, $\sigma_L = 0.08$]. (a) Extraction-with-escape success rate vs v_e for four $\delta m_{\rm nom} \in \{0.1, 0.2, 0.3, 0.4\}$ (labeling parameters only; $\delta m_{\rm actual}$ exceeds these values—see Table V caption for actual consumed fractions) with 95% Clopper-Pearson confidence intervals. Sharp onset near $v_e \approx 0.91$–$0.92c$; peak success 69.6% at $v_e = 0.98c$, $\delta m_{\rm nom} = 0.4$. (b) Cumulative efficiency $\eta_{\rm cum}$ [Eq. (19)], sample mean over successful escapes only (bins with zero successes are omitted, as the conditional mean is undefined there). Efficiency is inversely related to $\delta m_{\rm nom}$: smaller nominal burn sizes yield higher per-mass energy gain, reaching $\eta_{\rm cum} \approx 7.2\%$ at $v_e = 0.99c$, $\delta m_{\rm nom} = 0.1$. Success rate and efficiency trends oppose each other, presenting a design trade-off. Figure 6 extends (b) to the ultrarelativistic limit.

success increases rapidly with both v_e and $\delta m_{\rm nom}$. Table V shows representative results for $a/M = 0.95$.

Exploratory results at $a/M = 0.99$, now verified with independent random seeds, appear in Appendix B; they are not part of our headline findings.

The cumulative efficiency $\eta_{\rm cum}$ [Eq. (19)], computed as the sample mean over successful escapes (conditioning on $\mathcal{S}$), increases with exhaust velocity: for $a/M = 0.95$ it rises from $\eta_{\rm cum} \approx 0.1\%$–$0.2\%$ just above threshold ($v_e \approx 0.91$–$0.92c$) to $\eta_{\rm cum} \approx 5\%$–$7\%$ at $v_e = 0.98c$ [Fig. 5(b)]. At this spin, higher $\delta m_{\rm nom}$ values increase success rates but decrease per-mass efficiency, presenting a design trade-off.

Within this Gaussian sweet-spot prior and our baseline family and protocol, the observed velocity threshold acts as a practical constraint on material Penrose propulsion in our simulations. Combining higher exhaust velocity ($v_e = 0.98c$) with larger nominal burn-size parameter ($\delta m_{\rm nom} = 0.4$) yields success rates of 69.6% for $a/M = 0.95$, compared to ∼39% at $v_e = 0.95c$, $\delta m_{\rm nom} = 0.1$ in the same domain. Figure 5 visualizes this sharp onset and the associated trade-off between success probability and per-mass efficiency.

To probe the asymptotic limit $v_e \to c$—and test whether $\eta_{\rm cum}$ continues to grow once the sharp onset is passed—we extend the same single-impulse sweet-spot study by pushing the exhaust velocity into the ultrarelativistic regime, up to $v_e = 0.99999c$ ($\gamma \approx 224$). Figure 6 continues Fig. 5(b) and shows that the efficiency gain beyond $v_e \sim 0.99c$ is modest, approaching a plateau for this spin parameter.

H. Discussion

Figures 5 and 6 show two stages: a sharp onset of viable extraction near $v_e \approx 0.91$–$0.92c$, then diminishing efficiency gains as $v_e \to c$. This is consistent with the local kinematic estimate of Sec. IV E: the Killing-energy decomposition $E_{\rm ex} = -\gamma_e(u_t - v_e s_t)$ requires $v_e > v_e^{\rm crit} := |u_t|/|s_t| \approx 0.85$–$0.92c$ (depending on periapsis depth) for negative-energy exhaust. The Monte Carlo results agree: success rates remain below ∼1% for $v_e \leq 0.90c$ but rise steeply for $v_e \geq 0.93c$. Beyond this onset, $\eta_{\rm cum}$ approaches a plateau for $a/M = 0.95$ (Fig. 6). The behavior at $a/M = 0.99$ is discussed in Appendix B.

In our sampled prograde-flyby family, the observed spin threshold is consistent with ergosphere geometry: at $a/M = 0.7$, the ergosphere is shallow and the parameter space allowing deep penetration with escape is negligible in our sampled domain. The $10\times$ – $19\times$ amplification from sweet-spot focusing confirms that initial condition precision dominates over other factors. We note that the velocity-onset location ($v_e \approx 0.91$–$0.92c$) has been characterized only for a single Gaussian prior width and center; varying these or using a different prior family may shift the threshold. A systematic prior-sensitivity study is left for future work.

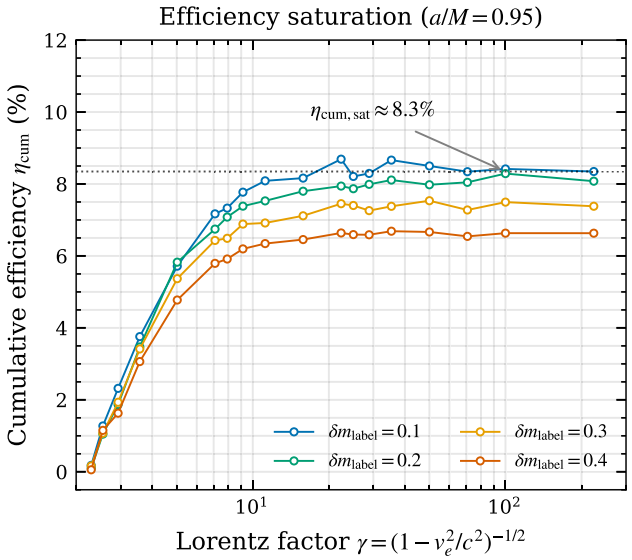


FIG. 6. Continuation of Fig. 5(b) into the extreme-γ regime for $a/M = 0.95$ in the same sweet-spot region [Eq. (19)]. All $\eta_{\rm cum}$ values are sample means over successful escapes (conditioning on $\mathcal{S}$). Exhaust velocities span $v_e = 0.90c$ ($\gamma = 2.3$) to $v_e = 0.99999c$ ($\gamma = 224$). Below $\gamma \approx 10$, $\eta_{\rm cum}$ grows approximately linearly with γ; above $\gamma \approx 20$–30, it saturates at $\eta_{\rm sat} \approx 8.3\%$, indicating diminishing returns at this spin. Lower $\delta m_{\rm nom}$ consistently yields higher $\eta_{\rm cum}$ ($\delta m_{\rm nom}$ is a nominal labeling parameter; see Table V caption for actual consumed fractions). This saturation plateau is an empirical ceiling for our Gaussian sweet-spot prior and thrust policy. We plot against Lorentz factor γ rather than v_e to resolve the saturation regime, where fractional changes in v_e become negligible (e.g., $v_e = 0.99c$ and $v_e = 0.999c$ differ by only 0.009).

Comparing our results with electromagnetic extraction mechanisms is instructive. The Blandford-Znajek process [11] achieves sustained extraction through magnetic field lines without requiring ultrarelativistic material exhaust, helping contextualize its astrophysical prevalence. Magnetic reconnection in the ergosphere [13] similarly bypasses a material exhaust-velocity threshold by accelerating charged particles electromagnetically. Our results are consistent with the view that material-based Penrose extraction, while theoretically elegant, faces more stringent parameter requirements than electromagnetic channels in realistic astrophysical settings.

Related relativistic rocket work has examined the Oberth effect and continuous fuel ejection in Schwarzschild spacetime [20]. That problem shares the variable-mass rocket accounting used here, but it is not a direct analog of the present Kerr calculation: Schwarzschild spacetime has no ergosphere and no exterior negative-Killing-energy states, so its energy gain is an Oberth/redshift effect rather than extraction of rotational energy. The comparison is nevertheless useful because it separates generic relativistic rocket bookkeeping from the specific Kerr ingredient in

this paper, namely controlled exhaust entering negative-energy states in the ergoregion.

Summary of experimental evidence. The principal quantitative results, synthesized in Sec. IX, derive from the following analyses: parameter space structure (Fig. 1); baseline success rates (Table I); spin threshold (Table II); sweet-spot amplification (Table III and Fig. 2); thrust strategy efficiency (Table IV and Fig. 3); spin dependence (Fig. 4); velocity onset (Table V and Fig. 5); ultrarelativistic efficiency saturation (Fig. 6); and exploratory $a/M = 0.99$ results (Appendix B).

IX. CONCLUSIONS

We have presented a Monte Carlo characterization of the Penrose process using rocket propulsion in Kerr spacetime, framing controlled exhaust ejection as a controlled-ejection variant of the classical Penrose decay. The statistical core of this study is the periapsis-impulse ensemble (Tables I–V); the continuous-thrust comparison is a secondary study of representative matched trajectories under specific implemented controllers (Table IV and Fig. 3). All constraints below apply to our equatorial prograde-flyby family under the steering prescriptions of Sec. VI G; different orbit families or steering laws may shift thresholds.

(1) Extraction requires fine-tuning. At most 1.4% of trajectories in a deliberately broad diagnostic domain ($E_0 \in [0.95, 2.0]$, $L_z \in [-3, 6]$, which includes bound and retrograde states outside the mission family; Table I) achieve extraction-with-escape success [$\Delta E > 0$ with escape to infinity, Eq. (18)], with success of order 1% only at the highest sampled spins. Under maximal tuning ($v_e = 0.98c$, $\delta m_{\rm nom} = 0.4$, corresponding to $\delta m_{\rm actual} \approx 0.55$), success rates reach 70% for $a/M = 0.95$ (Gaussian sweet-spot prior). Exploratory results at $a/M = 0.99$ yield comparable rates (∼65%–70%) under independent-seed verification (Appendix B). These represent the fine-tuned limit, not generic behavior. All reported success rates are conservative lower bounds, as integration failures (11.2% of samples) are counted as captures; because these failures cluster near the horizon, this conservative counting may also slightly bias estimated threshold locations (Appendix A).

(2) High spin is observed to be necessary within the explored family. Within our baseline family and protocol and sampled (E_0, L_z) domain, successful extraction with escape is observed only at high spin: no successes occur for $a/M \leq 0.88$ (0/90,000 trajectories at $v_e = 0.95c$, $\delta m_{\rm nom} = 0.3$), constraining the practical threshold to $0.88 < a_{\rm crit}/M \lesssim 0.89$ within the present Boyer-Lindquist numerics. This is an empirical boundary of the sampled region, not a coordinate-independent or universal lower bound on Penrose extraction.

(3) Highly relativistic exhaust is empirically required within the Gaussian sweet-spot prior. Under the Gaussian sweet-spot prior (centered at $E_0 = 1.22$, $L_z = 3.05$) and implemented steering prescriptions, we observe a sharp empirical onset near $v_e \approx 0.91$–$0.92c$ within the present Boyer-Lindquist numerics, below which success is of order 1% or below in the sampled prior and above which it increases rapidly. Negative exhaust Killing energy requires $v_e > v_e^{\rm crit} = |u_t|/|s_t|$ [Eq. (26)], a condition that becomes satisfiable over a wide range of periapsis configurations only at sufficiently relativistic speeds. Different orbit families or steering laws may yield different critical velocities.

(4) Single impulse at periapsis is most efficient among the strategies studied. Concentrating fuel at the point where $E_{\rm ex}$ is most negative maximizes extraction (an optimal-allocation result for a fixed worldline, exhaust-direction schedule, and fuel budget). For the specific implemented controllers and representative escape trajectories started from the same initial condition (Table IV), impulsive thrust achieves $\eta_{\rm cum} \approx 5.7\%$ (∼20% per exhaust rest mass; $\eta_{\rm trad} \approx 1.1\%$). Continuous thrust yields $\eta_{\rm cum} \approx 3.7\%$ for the representative matched escape trajectory and implemented controllers (∼65% of the impulse value), consistent with path averaging and trajectory deformation under extended burning; this comparison is between specific controllers, not a general theorem about continuous thrust. Global optimality over all possible steering laws remains an open optimal-control problem (Sec. VIII D).

(5) Efficiency saturates empirically. Over the explored range up to $v_e = 0.99999c$ ($\gamma \approx 224$), $\eta_{\rm cum}$ levels off near 8.3% for $a/M = 0.95$ (Fig. 6), with diminishing returns above $v_e \approx 0.999c$. This is an empirical ceiling for the Gaussian sweet-spot prior and implemented steering prescriptions, not a theoretical upper bound on Penrose extraction efficiency.

Our results characterize the controlled-exhaust regime of the Penrose process, a controlled-ejection variant distinct from the collisional mechanisms of Piran *et al.* [5] and Leiderschneider and Piran [10]. The stringent requirements on spin, velocity, and orbital geometry are consistent with the comparative difficulty of material-exhaust Penrose channels relative to electromagnetic ones (Blandford-Znajek, magnetic reconnection), which are generally more astrophysically relevant for energy extraction from rotating black holes.

Engineering summary. Within our equatorial prograde-flyby family, Gaussian sweet-spot prior, and steering prescriptions, non-negligible success is observed only under: (i) $a/M \gtrsim 0.89$, (ii) $v_e \gtrsim 0.91c$, and (iii) initial conditions tuned to $(E_0, L_z) \approx (1.2, 3.0)$. Single-impulse efficiency is $\eta_{\rm cum} \approx 5.7\%$ (∼20% per exhaust rest mass;

TABLE VII. Conservative extraction-with-escape success rates illustrating the impact of integration failures for representative high-spin cases in a focused sampling domain ($E_0 \in [1.1, 1.3]$, $L_z \in [2.5, 3.5]$) under the periapsis-impulse protocol with $v_e = 0.95c$ and $\delta m_{\text{nom}} = 0.20$. P1: "success (excl. int. failures)" excludes integration failures from the denominator; P2: "success (failures as captures)" counts all failures as nonsuccesses. The differences correspond to the observed ∼11.2% failure fraction and support conservatively classifying integration failures as captures near the horizon.

a/M	P1 (%)	P2 (%)	Difference (%)
0.90	0.38	0.33	−0.05
0.95	8.34	7.41	−0.93
0.99	10.84	9.63	−1.21

$\eta_{\text{trad}} \approx 1.1\%$); continuous thrust achieves $\eta_{\text{cum}} \approx 3.7\%$ for the representative matched escape trajectory (∼65% of the impulse value). All efficiencies are conditional on extraction-with-escape events.

ACKNOWLEDGMENTS

This work is financially supported by VinUniversity under the Environmental Intelligence (CEI) Grant (No. VUNI.CEI.FS_0009).

DATA AVAILABILITY

The data that support the findings of this article are openly available [21].

APPENDIX A: NUMERICAL VALIDATION

We verified numerical accuracy through multiple diagnostics. We maintain the mass-shell constraint $\mathcal{C} = g^{\mu\nu} p_\mu p_\nu + m^2$ to $|\mathcal{C}| < 10^{-9}$ for single-impulse trajectories and $< 10^{-6}$ for continuous thrust (with projection restoring $< 10^{-14}$). We verify future-directedness ($u^t > 0$) and timelike normalization throughout all trajectories. Energy conservation via 4-momentum balance yields residuals $< 10^{-9}$.

We cross-validated a representative subset of continuous-thrust trajectories with RK45, finding consistent escape/capture outcomes and ΔE within statistical uncertainty. Boyer-Lindquist coordinates become numerically stiff near r_+; integration failures (11.2% of samples) occur exclusively for trajectories approaching the horizon ($r < 1.1 r_+$), and we conservatively treat such failures as captures. We find that treating failures as captures yields consistent escape probabilities (Table VII). Because these failures cluster near the horizon rather than distributing uniformly, the conservative counting may slightly bias estimated threshold locations (e.g., the velocity-onset value), not only lower overall success rates. Future work could validate near-horizon outcomes in a horizon-regular coordinate system (e.g., Kerr-Schild) to further distinguish coordinate-induced integration failures from physical captures.

TABLE VIII. Exploratory: Thrust parameter sensitivity for $a/M = 0.99$ in the sweet-spot region centered at ($E_0 = 1.22$, $L_z = 3.05$), with $N = 1000$ initial conditions per configuration and distinct random seeds per $(v_e, \delta m_{\text{nom}})$ pair (same Gaussian sampling protocol, δm_{nom} convention, and δm_{actual} discrepancies as Table V). The original reused-ensemble scan yielded uniformly 86.3% at $v_e = 0.98c$ across all δm_{nom}; independent-seed verification reveals ∼65%–70% with mild δm_{nom} dependence (4–6 percentage-point spread), confirming that the reused-ensemble rates were upward biased by a particularly favorable seed. Peak success of 70.0% at $v_e = 0.98c$, $\delta m_{\text{nom}} = 0.4$.

v_e/c	$\delta m_{\text{nom}} = 0.1$ (%)	$\delta m_{\text{nom}} = 0.2$ (%)	$\delta m_{\text{nom}} = 0.3$ (%)	$\delta m_{\text{nom}} = 0.4$ (%)
0.80	0.0	0.0	0.0	0.0
0.90	0.0	0.1	0.1	0.0
0.95	33.3	32.7	31.7	36.2
0.98	68.3	63.7	65.1	**70.0**

APPENDIX B: EXPLORATORY $a/M = 0.99$ RESULTS

Table VIII shows thrust parameter sensitivity results for $a/M = 0.99$, now updated with independent-seed verification [$N = 1000$ per configuration, distinct random seed per $(v_e, \delta m_{\text{nom}})$ pair]. These results are exploratory and not part of the main evidentiary spine.

Under the fixed $a/M = 0.95$-centered baseline prior, the relative ranking of $a/M = 0.99$ and $a/M = 0.95$ is not monotone across the parameter space. At $v_e = 0.90c$, $a/M = 0.99$ achieves ∼0% success compared to ∼1% for $a/M = 0.95$; at $v_e = 0.95c$, the ordering depends on δm_{nom} (Tables V–VIII). At $v_e = 0.98c$, $a/M = 0.99$ achieves ∼65%–70% compared to 69.6% for $a/M = 0.95$, indicating comparable performance under the $a/M = 0.95$-centered prior. This nonmonotone behavior is a fixed-prior artifact: the sweet-spot center $(E_0, L_z) = (1.22, 3.05)$ was derived at $a/M = 0.95$ and need not be optimal for $a/M = 0.99$, whose escape corridors occupy a different region of phase space. A spin-optimized prior centered on the $a/M = 0.99$ sweet spot would likely shift the balance.

Independent-seed verification confirms mild δm_{nom} dependence (4–6 percentage-point spread at $v_e = 0.95c$ and $v_e = 0.98c$), in contrast to the reused-ensemble artifact of identical rates across all δm_{nom} columns.

[1] R. Penrose and R. M. Floyd, Extraction of rotational energy from a black hole, Nat. Phys. Sci. **229**, 177 (1971).
[2] D. Christodoulou, Reversible and irreversible transformations in black-hole physics, Phys. Rev. Lett. **25**, 1596 (1970).
[3] J. M. Bardeen, W. H. Press, and S. A. Teukolsky, Rotating black holes: Locally nonrotating frames, energy extraction, and scalar synchrotron radiation, Astrophys. J. **178**, 347 (1972).
[4] R. M. Wald, Energy limits on the Penrose process, Astrophys. J. **191**, 231 (1974).
[5] T. Piran, J. Shaham, and J. Katz, High efficiency of the Penrose mechanism for particle collisions, Astrophys. J. **196**, L107 (1975).
[6] M. Bañados, J. Silk, and S. M. West, Kerr black holes as particle accelerators to arbitrarily high energy, Phys. Rev. Lett. **103**, 111102 (2009).
[7] M. Bejger, T. Piran, M. Abramowicz, and F. Håkanson, Collisional Penrose process near the horizon of extreme Kerr black holes, Phys. Rev. Lett. **109**, 121101 (2012).
[8] E. Berti, R. Brito, and V. Cardoso, Ultrahigh-energy debris from the collisional Penrose process, Phys. Rev. Lett. **114**, 251103 (2015).
[9] J. D. Schnittman, Revised upper limit to energy extraction from a Kerr black hole, Phys. Rev. Lett. **113**, 261102 (2014).
[10] E. Leiderschneider and T. Piran, Maximal efficiency of the collisional Penrose process, Phys. Rev. D **93**, 043015 (2016).
[11] R. D. Blandford and R. L. Znajek, Electromagnetic extraction of energy from Kerr black holes, Mon. Not. R. Astron. Soc. **179**, 433 (1977).
[12] A. Tchekhovskoy, R. Narayan, and J. C. McKinney, Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole, Mon. Not. R. Astron. Soc. **418**, L79 (2011).
[13] L. Comisso and F. A. Asenjo, Magnetic reconnection as a mechanism for energy extraction from rotating black holes, Phys. Rev. D **103**, 023014 (2021).
[14] S. Koide and K. Arai, Energy extraction from a rotating black hole by magnetic reconnection in the ergosphere, Astrophys. J. **682**, 1124 (2008).
[15] S. Chandrasekhar, *The Mathematical Theory of Black Holes* (Clarendon Press, Oxford, 1983).
[16] M. Bhat, S. Dhurandhar, and N. Dadhich, Energetics of the Kerr-Newman black hole by the Penrose process, J. Astrophys. Astron. **6**, 85 (1985).
[17] A. A. Grib, Y. V. Pavlov, and V. D. Vertogradov, Geodesics with negative energy in the ergosphere of rotating black holes, Mod. Phys. Lett. A **29**, 1450110 (2014).
[18] C. J. Clopper and E. S. Pearson, The use of confidence or fiducial limits illustrated in the case of the binomial, Biometrika **26**, 404 (1934).
[19] B. Efron, Better bootstrap confidence intervals, J. Am. Stat. Assoc. **82**, 171 (1987).
[20] Y. V. Pavlov and O. B. Zaslavskii, Particle decay, Oberth effect and a relativistic rocket in the Schwarzschild background, Int. J. Mod. Phys. D **32**, 2250143 (2023).
[21] A. Le, Simulation code, GitHub, https://github.com/aninde x/penrose_process (2026), Animated trajectory visualizations are provided in the codebase.