

Risk-Optimal Pulsed Control of a Variable-Separation Two-Rail Ball Game

A compliant hybrid model of the Space Force / Shoot-the-Moon mechanism

Jude Gomila

Independent Researcher

August 2026

Abstract

A ball supported by two player-controlled rods forms a compact nonlinear mechanical system. The rods rise along the game board, while increasing their separation lowers the ball relative to the rods; this geometric competition permits apparent uphill motion. Prior work derived rigid-body models and feedback controllers for the Shoot-the-Moon apparatus. The present paper develops a complementary model for the high-score playing problem, incorporating rolling resistance, no-slip feasibility, small rail bending, contact loss, pulsed actuation, target-specific recoverability, and risk-sensitive control.

The cross-sectional geometry is represented by the normalized separation $\eta = d/[2(R + r)]$ and support margin $\delta = \sqrt{1 - \eta^2}$. These variables give an exact symmetric rolling radius $R_{\text{eff}} = R\delta$, an effective mass $M = m(1 + \kappa/\delta^2)$, and a reduced Lagrangian equation separating rail inclination from geometry-induced drive. A critical opening condition is derived. Although the geometric driving force grows as δ^{-1} , rotational inertia makes acceleration from rest vanish as $\delta \downarrow 0$, revealing a previously hidden terminal regularization. A static beam-compliance closure has an explicit fold at $\delta_* = \Lambda^{1/3}$, predicting premature loss of the support branch. The terminal drop is formulated as a hybrid transition, and recoverable states are described by a target-specific viability kernel. Under Gaussian release-timing uncertainty, hit probability is proved to decrease strictly with speed. The observed maximum-score strategy is therefore a boundary-riding risk problem. Pontryagin's maximum principle explains why high-authority pulses are natural, while the claim that exactly one intermediate pulse is globally optimal is stated as a falsifiable conjecture.

Keywords. nonholonomic rolling; variable geometry; compliant rails; hybrid mechanics; viability kernel; risk-sensitive optimal control; bang-bang control; dexterity game.

MSC 2020. 70E18, 70F25, 74H15, 49J15, 93C10.

1 Introduction

The apparatus consists of a rigid ball supported by two slender cylindrical rods. The rods are close together near a hinge and can be opened or closed by the player at the opposite end. Their mean direction is inclined upward, yet opening the rods lowers the ball center enough to produce motion toward the high end. The player attempts to release the ball through one of several scoring apertures, with the terminal target demanding both sufficient progress and precise release timing. A schematic is shown in [Figure 1](#).

The central playing observation is a risk paradox: a run aimed at the maximum score appears to benefit from a strong pulse in an intermediate region and a terminal passage at nearly the greatest controllable speed. This simultaneously narrows the timing window and leaves little recovery authority. The best possible score therefore seems to demand the greatest tolerated risk.




Figure 1: Top-view idealization. The included rod angle is 2α , and the center-to-center separation $d(x, t)$ increases with position. A real apparatus can replace the straight-rail law by a measured geometry $d = d(x, u)$.

The basic rigid dynamics are established. Xu et al. (2010, 2012) derived nonlinear rigid-body equations for the Shoot-the-Moon game, identified its underactuated and nonholonomic structure, and constructed model-based controllers; the associated thesis gives an extended treatment (Xu, 2011). Rolling contact in V-grooves and related systems has also been studied in tribology and contact mechanics (De Moerloose et al., 2011; Johnson, 1985; Kalker, 1990). Accordingly, this paper does not claim the first model of the game.

The contribution is a mathematical extension directed specifically at score-optimal human play. The model adds:

- (i) an exact local rolling radius and geometry-dependent effective inertia;
- (ii) rolling resistance, traction feasibility, and normal-load guards;
- (iii) slight rail compliance and an explicit compliance-induced fold;
- (iv) hybrid transitions among rolling, slip, one-contact motion, and free flight;
- (v) target-specific viability and uncertainty-sensitive score optimization;
- (vi) a precise, testable single-pulse conjecture.

2 Geometry of two-rail support

2.1 Coordinates and assumptions

Let $x \in [0, L]$ be distance along the mean rail direction, positive uphill. The rail-center plane rises at angle $\beta \in (0, \pi/2)$, so

$$z_r(x) = z_0 + x \sin \beta. \quad (2.1)$$

The ball has mass m , radius R , and central moment of inertia

$$I = \kappa m R^2, \quad \kappa = \frac{2}{5} \quad \text{for a homogeneous solid sphere.} \quad (2.2)$$

Each rod has radius r , and

$$a := R + r. \quad (2.3)$$

For symmetric straight rods with half-opening angle $\alpha(t)$,

$$d(x, t) = d_0 + 2x \tan \alpha(t). \quad (2.4)$$

The actual apparatus may instead use a calibrated map $d = d(x, u)$, where u is handle displacement.




Figure 2: The same support margin δ controls ball height, contact orientation, rolling kinematics, and normal load. The plotted height curves are schematic, not fitted data.

2.2 Cross-sectional support

In transverse cross-section, the ball center is a distance a from either rod center. If h is its height above the rod-center plane,

$$\left(\frac{d}{2}\right)^2 + h^2 = a^2, \quad h(d) = \sqrt{a^2 - \frac{d^2}{4}}. \quad (2.5)$$

Introduce the dimensionless separation and support margin

$$\eta := \frac{d}{2a}, \quad \delta := \sqrt{1 - \eta^2}, \quad h = a\delta. \quad (2.6)$$

The rigid geometric support limit is $\eta \uparrow 1$, equivalently $\delta \downarrow 0$.

Proposition 2.1 (Monotone, concave support height). *For $0 < d < 2a$,*

$$h_d = -\frac{d}{4h} < 0, \quad h_{dd} = -\frac{1}{4h} - \frac{d^2}{16h^3} < 0. \quad (2.7)$$

Thus opening the rods lowers the ball center, and the magnitude of geometric sensitivity increases without bound as $\delta \downarrow 0$.

Proof. Differentiate Equation (2.5) once and twice with respect to d . □

The ball-center height and potential energy are

$$z_B(x, \alpha) = z_0 + x \sin \beta + a\delta(x, \alpha), \quad V(x, \alpha) = mgz_B(x, \alpha). \quad (2.8)$$

3 Exact symmetric rolling reduction

3.1 Contact kinematics

Let $\mathbf{e}_x$ point along the mean rail direction, $\mathbf{e}_y$ transversely, and $\mathbf{e}_z$ normally away from the rail-center plane. The unit normal forces on the ball are directed along

$$\mathbf{n}_{\pm} = \mp \eta \mathbf{e}_y + \delta \mathbf{e}_z, \quad (3.1)$$

where $+$ and $-$ label the right and left rods. The radius vectors from the ball center to the contact points are

$$\boldsymbol{\rho}_{\pm} = -R\mathbf{n}_{\pm}. \quad (3.2)$$

For symmetric longitudinal rolling, take $\boldsymbol{\omega} = \omega \mathbf{e}_y$. On stationary rods the longitudinal no-slip condition is

$$0 = (v\mathbf{e}_x + \boldsymbol{\omega} \times \boldsymbol{\rho}_{\pm}) \cdot \mathbf{e}_x = v - R\delta\omega. \quad (3.3)$$

Therefore

$$v = R_{\text{eff}}\omega, \quad R_{\text{eff}} = R\delta. \quad (3.4)$$

Proposition 3.1 (Geometry-dependent effective mass). *Under symmetric no-slip rolling,*

$$T = \frac{1}{2}M(x, \alpha)v^2, \quad M(x, \alpha) = m \left(1 + \frac{\kappa}{\delta(x, \alpha)^2} \right). \quad (3.5)$$

Proof. Substitute $\omega = v/(R\delta)$ from Equation (3.4) into $T = \frac{1}{2}mv^2 + \frac{1}{2}I\omega^2$ and use Equation (2.2). $\square$

3.2 Reduced Lagrange equation

For prescribed $\alpha(t)$, define

$$\mathcal{L}(x, v; \alpha) = \frac{1}{2}M(x, \alpha)v^2 - V(x, \alpha). \quad (3.6)$$

Let $F_d(v, x, \alpha)$ denote the longitudinal dissipative force magnitude opposing motion. The Euler–Lagrange equation is

$$M\dot{v} + \frac{1}{2}M_x v^2 + M_\alpha \dot{\alpha} v + V_x = -F_d. \quad (3.7)$$

From Equations (2.4) and (2.6),

$$\eta_x = \frac{\tan \alpha}{a}, \quad \eta_\alpha = \frac{x \sec^2 \alpha}{a}, \quad (3.8)$$

$$\delta_x = -\frac{\eta \tan \alpha}{a\delta}, \quad \delta_\alpha = -\frac{\eta x \sec^2 \alpha}{a\delta}. \quad (3.9)$$

Consequently,

$$-V_x = mg \left(\frac{\eta \tan \alpha}{\delta} - \sin \beta \right), \quad (3.10)$$

$$M_x = \frac{2m\kappa\eta \tan \alpha}{a\delta^4}, \quad M_\alpha = \frac{2m\kappa\eta x \sec^2 \alpha}{a\delta^4}. \quad (3.11)$$

The first term in Equation (3.10) is geometry-induced uphill drive, while the second is ordinary downhill gravity.

4 Uphill drive and terminal asymptotics

Define the dimensionless local drive margin

$$\Gamma(x, \alpha) := \frac{\eta \tan \alpha}{\delta} - \sin \beta. \quad (4.1)$$

Proposition 4.1 (Local opening threshold). *At a fixed local separation (η, δ) , with $v = 0$ and negligible dissipation, the ball accelerates uphill if and only if $\Gamma > 0$. Equivalently,*

$$\tan \alpha > \frac{\delta \sin \beta}{\eta} = \frac{2h \sin \beta}{d}. \quad (4.2)$$

When d itself depends on α , Equation (4.2) is an implicit local threshold equation.




Figure 3: The ideal geometric force grows near the support limit, but rotational inertia grows faster. Curves show asymptotic forms with $\kappa = 2/5$.

Proof. At rest the velocity-dependent terms in Equation (3.7) vanish. With $F_d = 0$, one has $M\dot{v} = -V_x = mg\Gamma$. Since $M > 0$, the sign of $\dot{v}$ equals the sign of Γ . $\square$

A naive inspection of Equation (3.10) suggests an unbounded uphill force as $\delta \downarrow 0$. The rolling inertia in Equation (3.5) changes that conclusion.

Proposition 4.2 (Rotational regularization at zero speed). *For fixed $\alpha > 0$, negligible dissipation, and $v = 0$,*

$$\dot{v} = g \frac{\eta \tan \alpha / \delta - \sin \beta}{1 + \kappa / \delta^2}. \quad (4.3)$$

As $\delta \downarrow 0$,

$$\dot{v} \sim \frac{g \tan \alpha}{\kappa} \delta \longrightarrow 0. \quad (4.4)$$

Thus the force singularity is overpowered by the δ^{-2} growth of rotational effective mass.

Proof. Set $v = 0$, $\dot{\alpha} = 0$, and $F_d = 0$ in Equation (3.7), then use Equations (3.5) and (3.10). Since $\eta \rightarrow 1$ as $\delta \rightarrow 0$, the leading numerator is $\tan \alpha / \delta$ and the leading denominator is κ / δ^2 . $\square$

Remark 4.3 (Why the terminal regime remains dangerous). The vanishing of Equation (4.4) does not make the terminal regime benign. At nonzero speed, $\omega = v/(R\delta)$ diverges, while $M_x v^2$ and $M_\alpha \dot{\alpha} v$ scale as δ^{-4} . Normal load and traction demand also grow. The ideal no-slip reduction therefore ceases to be physically credible before the rigid limit is reached.

5 Dissipation and contact feasibility

5.1 Normal force

Under quasi-static balance normal to the rail-center plane, equal rail forces satisfy

$$2N\delta \approx mg \cos \beta, \quad N \approx \frac{mg \cos \beta}{2\delta}. \quad (5.1)$$

Dynamic normal acceleration adds correction terms, but Equation (5.1) gives the principal geometric scaling.

A reduced rolling-resistance model is

$$F_d = \mu_r mg \frac{\cos \beta}{\delta} \operatorname{sgn}(v) + c_1 v + c_2 |v|v, \quad (5.2)$$

where the first term is the sum of the resistance at the two rods. A regularized odd function should replace $\operatorname{sgn}(v)$ in numerical work near $v = 0$.

5.2 No-slip feasibility

Let $\mathbf{F}_{t,i}$ be the tangential contact force at rail i . Sticking requires

$$\|\mathbf{F}_{t,i}\| \leq \mu_s N_i, \quad (5.3)$$

with transition to microslip or sliding when the inequality is saturated. Monitoring normal force and static-friction demand is essential in rolling-ball systems (Putkaradze and Rogers, 2018). Because both the required spin and normal load become extreme near small δ , terminal high-speed play is expected to encounter slip or material compliance before the rigid geometric limit.

It is useful to define admissibility guards

$$G_{\text{geom}} := 1 - \eta, \quad (5.4)$$

$$G_N := N_{\max} - N, \quad (5.5)$$

$$G_\mu := \mu_s N - \|\mathbf{F}_t\|, \quad (5.6)$$

$$G_\omega := \omega_{\max} - \frac{|v|}{R\delta}. \quad (5.7)$$

The intended two-contact rolling mode requires each guard to remain positive.

6 Time-dependent opening and pulse work

The handle motion changes $\alpha(t)$ and may also perturb β or the rod-center position. Even when only α is retained, the constraint manifold is time dependent and the player can exchange energy with the ball.

Define the reduced mechanical energy

$$E(x, v, \alpha) = \frac{1}{2} M(x, \alpha) v^2 + V(x, \alpha). \quad (6.1)$$

Proposition 6.1 (Reduced energy identity). *Along solutions of Equation (3.7),*

$$\dot{E} = \left(V_\alpha - \frac{1}{2} M_\alpha v^2 \right) \dot{\alpha} - F_d v. \quad (6.2)$$

Proof. Differentiate Equation (6.1), substitute Equation (3.7), and cancel the $M_x v^3/2$ terms. Equivalently, use the standard identity $\dot{E} = -\partial \mathcal{L}/\partial t + Qv$ with $Q = -F_d$. $\square$

For an open-close pulse over $[t_0, t_1]$,

$$\Delta E = \int_{t_0}^{t_1} \left(V_\alpha - \frac{1}{2} M_\alpha v^2 \right) \dot{\alpha} \, dt - \int_{t_0}^{t_1} F_d v \, dt. \quad (6.3)$$

The first integral can be nonzero even if $\alpha(t_1) = \alpha(t_0)$ because x and v change during the cycle. A finite pulse may be parameterized as

$$\alpha(t) = \alpha_0 + A f\left(\frac{t - t_p}{\tau_p}\right), \quad (6.4)$$

where A , τ_p , and t_p are amplitude, duration, and center time.

In the complete contact model, the instantaneous actuator power is more directly

$$P_{\text{contact}} = \sum_i \mathbf{F}_i \cdot \mathbf{v}_{r,i}, \quad (6.5)$$

where $\mathbf{v}_{r,i}$ is the local rod material velocity. Equation (6.2) is a symmetric reduction of this moving-contact work.

If a pulse supplies net useful energy ΔE_{use} while M varies little over the pulse, then

$$v_+^2 \approx v_-^2 + \frac{2\Delta E_{\text{use}}}{M}. \quad (6.6)$$

This equation formalizes the player’s sensation of “pulsing the ball.”

7 Small rail compliance

7.1 Beam equation

Let $w_i(s, t)$ be the outward transverse deflection of rod i , measured along rod coordinate $s \in [0, L]$. Euler–Bernoulli theory gives

$$\rho_b A_b w_{i,tt} + c_b w_{i,t} + EI_b w_{i,ssss} = q_i(s, t), \quad (7.1)$$

where EI_b is flexural rigidity and c_b is distributed damping (Meirovitch, 2001). The lateral component of a localized contact load is approximately

$$q_i(s, t) = N_i(t)\eta(t)\delta_D(s - x(t)) + q_{h,i}(s, t). \quad (7.2)$$

For a simply supported rod, the static compliance at the load point is

$$C_b(x) = \frac{x^2(L - x)^2}{3EI_b L}, \quad (7.3)$$

although the actual boundary conditions should be measured.

7.2 Static compliance closure

Let $\eta_c = d_c/(2a)$ be the commanded rigid separation. Each rail deflects outward by approximately $w = C_b N \eta$, so the total separation change is $2w$. Using Equation (5.1),

$$\eta = \eta_c + \Lambda \frac{\eta}{\delta}, \quad \Lambda := \frac{C_b m g \cos \beta}{2a}. \quad (7.4)$$

Equivalently,

$$\eta_c(\eta) = \eta \left(1 - \frac{\Lambda}{\delta}\right). \quad (7.5)$$

Proposition 7.1 (Compliance-induced fold). *Assume $0 < \Lambda < 1$. The quasi-static branch Equation (7.5) has a fold at*

$$\delta_* = \Lambda^{1/3}, \quad \eta_* = \sqrt{1 - \Lambda^{2/3}}, \quad \eta_{c,*} = \left(1 - \Lambda^{2/3}\right)^{3/2}. \quad (7.6)$$

Thus finite compliance terminates the monotone support branch before the rigid command value $\eta_c = 1$.




Figure 4: Compliance is not merely a small correction. The reduced static closure develops a fold, providing a mechanism for premature release or snap-through.

Proof. Since $d(\eta/\delta)/d\eta = \delta^{-3}$,

$$\frac{d\eta_c}{d\eta} = 1 - \frac{\Lambda}{\delta^3}. \quad (7.7)$$

The fold occurs at $\delta^3 = \Lambda$. Substitution into $\eta^2 + \delta^2 = 1$ and Equation (7.5) gives the remaining formulas. $\square$

The first-order height perturbation caused by a small separation error Δd is

$$\Delta h \approx -\frac{d}{4h} \Delta d. \quad (7.8)$$

Because $N = O(\delta^{-1})$ and $|h_d| = O(\delta^{-1})$, the formal composition of load-induced deflection and geometric height sensitivity scales as $O(\delta^{-2})$. This explains why visually slight bending can dominate the terminal region.

8 Hybrid contact and terminal release

The game does not remain in one smooth dynamical mode. Let $g_i(q)$ be the normal gap at contact i and N_i the corresponding compressive normal force. Unilateral contact is represented by

$$g_i \geq 0, \quad N_i \geq 0, \quad g_i N_i = 0. \quad (8.1)$$

Together with Equation (5.3), these relations define a nonsmooth mechanical system (Brogliato, 2016). Relevant modes are:

- M1: symmetric two-contact rolling;
- M2: two-contact rolling with microslip;
- M3: asymmetric or one-contact motion;
- M4: free flight after release;
- M5: impact with the target board or collection tray.




Figure 5: Hybrid mode graph. The scoring release is a controlled guard crossing into free flight; an early or asymmetric crossing is a failure.

During free flight,

$$m\ddot{\mathbf{r}}_C = m\mathbf{g} + \mathbf{F}_{\text{air}}, \quad (8.2)$$

followed by an impact map or compliant impact law. The full terminal event is therefore a transition from constrained rolling to ballistic motion, not merely a continuation of Equation (3.7).

The compliance fold adds another guard,

$$G_{\text{fold}} := \delta^3 - \Lambda. \quad (8.3)$$

A conservative safe set can require

$$G_{\text{geom}} > 0, \quad G_{\text{fold}} > 0, \quad G_N > 0, \quad G_\mu > 0, \quad G_\omega > 0. \quad (8.4)$$

9 Target-specific viability and speed risk

9.1 Recoverable states

Let the state include rigid, actuator, and retained beam variables,

$$z = (x, v, \alpha, \dot{\alpha}, q_b, \dot{q}_b, \dots). \quad (9.1)$$

Let $\mathcal{S}$ be the state set in which the desired contact mode and all actuator/contact guards are valid, and let $\mathcal{T}_k$ be the terminal set corresponding to scoring target k .

Definition 9.1 (Target-specific viability kernel). For horizon T , define

$$\mathcal{K}_k(T) := \{z_0 \in \mathcal{S} : \exists u(\cdot) \in \mathcal{U} \text{ such that } z(t) \in \mathcal{S} \text{ until } \mathcal{T}_k \text{ is reached}\}. \quad (9.2)$$

This is a controlled viability/capture-basin problem in the sense of Aubin (1991); Saint-Pierre (1994). In a projection onto (x, v) , part of the boundary may be summarized by a target-specific critical curve $v_{\text{crit},k}(x)$. A high-performance path uses a small reserve

$$\varepsilon(x) = v_{\text{crit},k}(x) - v(x) > 0. \quad (9.3)$$

The curve is global: it represents remaining distance, available control, friction, compliance, and terminal timing, not merely a local detachment speed.




Figure 6: Schematic target-specific phase portrait. Record-seeking play approaches the viability boundary from below, whereas reliability-oriented play retains a larger reserve.

9.2 Release timing uncertainty

Suppose the desired release position is centered in a longitudinal interval of width Δx . Let residual position error and motor timing error be independent Gaussians,

$$e_x \sim \mathcal{N}(0, \sigma_x^2), \quad e_t \sim \mathcal{N}(0, \sigma_t^2). \quad (9.4)$$

The linearized release-position error is $e = e_x + v e_t$, hence

$$\sigma_{\text{eff}}^2(v) = \sigma_x^2 + v^2 \sigma_t^2. \quad (9.5)$$

The probability of releasing inside the target window is

$$p_{\text{hit}}(v) = \mathbb{P}\{|e| \leq \Delta x/2\} = \text{erf}\left(\frac{\Delta x}{2\sqrt{2}\sqrt{\sigma_x^2 + v^2 \sigma_t^2}}\right). \quad (9.6)$$

Proposition 9.2 (Timing risk increases strictly with speed). *If $\sigma_t > 0$ and $\Delta x > 0$, then $p'_{\text{hit}}(v) < 0$ for every $v > 0$.*

Proof. The denominator in the argument of erf is strictly increasing for $v > 0$, and erf is strictly increasing. Their composition is therefore strictly decreasing. $\square$

The elementary release-time window obeys

$$\Delta t_{\text{release}} \approx \frac{\Delta x}{v}. \quad (9.7)$$

Thus speed can improve reach while unavoidably reducing release robustness.




Figure 7: The same action that increases reach can move the trajectory toward the viability boundary and reduce terminal hit probability. Curves are illustrative, not fitted to a particular machine.

9.3 Score objectives

Let $S(v)$ denote score conditional on an aggressiveness class and $P_s(v)$ its success probability. Expected-score play maximizes

$$J(v) = P_s(v)S(v). \quad (9.8)$$

At an interior optimum,

$$\frac{S'(v)}{S(v)} = -\frac{P'_s(v)}{P_s(v)}. \quad (9.9)$$

The fractional gain in score is balanced by the fractional loss of reliability.

A record-seeking objective is different. For target k with per-attempt success probability p_k , the probability of at least one success in n approximately independent attempts is

$$P_{k,n} = 1 - (1 - p_k)^n. \quad (9.10)$$

A player seeking the maximum target at least once can rationally accept lower per-attempt reliability.

10 Risk-sensitive optimal control

A rate-limited reduced state can be written

$$z = (x, v, \alpha, a, q_b, \dot{q}_b), \quad a = \dot{\alpha}, \quad (10.1)$$

with jerk $j = \dot{a}$ as the bounded control. Retaining one symmetric beam mode gives

$$\dot{x} = v, \quad (10.2)$$

$$M\dot{v} = -V_x - \frac{1}{2}M_x v^2 - M_\alpha a v - F_d + Q_b, \quad (10.3)$$

$$\dot{\alpha} = a, \quad (10.4)$$

$$\dot{a} = j, \quad (10.5)$$

$$\ddot{q}_b + 2\zeta_b \omega_b \dot{q}_b + \omega_b^2 q_b = B_b N \eta, \quad (10.6)$$

subject to bounds on α , a , j , normal force, friction demand, and the hybrid guards.

A risk-sensitive objective is

$$\max_{j(\cdot)} \left\{ \mathbb{E}[S(X_{\text{impact}})] - \lambda \mathbb{P}\{\text{failure}\} - \gamma \int_0^T j(t)^2 dt - \xi T \right\}. \quad (10.7)$$

The weight λ selects a point on the speed-risk frontier.

Proposition 10.1 (Bang-bang tendency). *Suppose a reduced model is control-affine,*

$$\dot{z} = f(z) + g(z)u, \quad u \in [u_{\min}, u_{\max}], \quad (10.8)$$

with a minimum-time or terminal-payoff problem whose Pontryagin Hamiltonian is affine in u . Away from singular arcs, an optimal control takes values only at $u_{\min}$ or $u_{\max}$ almost everywhere.

Proof. The Hamiltonian has the form $H = \lambda_z^T f(z) + \lambda_z^T g(z)u + \ell(z)$. For fixed state and costate it is affine in u , so its extremum over a compact interval occurs at an endpoint unless the switching function $\lambda_z^T g(z)$ vanishes on an interval. This is the standard maximum-principle argument (Pontryagin et al., 1962). $\square$

This result explains why a short high-authority action is natural, but not how many switches are optimal. A schematic pulse-location merit function is

$$\mathcal{Q}(x_p) = B_{\text{energy}}(x_p) - L_{\text{remaining}}(x_p) - \lambda R_{\text{terminal}}(x_p). \quad (10.9)$$

Early pulses leave more distance for dissipation and error growth. Late pulses leave little useful travel and perturb the most sensitive region. These competing effects make an interior x_p^* plausible without implying $x_p^* = L/2$.

Conjecture 10.2 (Single-dominant-pulse regime). *There exists a nonempty parameter regime of rail slope, opening authority, rolling loss, compliance, actuator rate limits, and terminal-risk weight in which the maximum-score optimal policy contains one dominant open-close pulse in an intermediate region, followed by a high-speed arc lying just inside the target-specific viability boundary.*

The conjecture is deliberately conditional. Strong dissipation, weak actuation, accurate high-bandwidth feedback, or a sharply varying viability boundary can favor several smaller pulses. The decisive test is to solve the calibrated problem without constraining pulse count and determine whether one dominant pulse emerges.

11 Dimensionless groups and mechanical regimes

Use the rail length L , gravitational speed $\sqrt{gL}$, and time $\sqrt{L/g}$ to define

$$X = \frac{x}{L}, \quad U = \frac{v}{\sqrt{gL}}, \quad \tau = t\sqrt{\frac{g}{L}}. \quad (11.1)$$

Important dimensionless groups are

$$\rho_r = \frac{r}{R}, \quad \Pi_\alpha = \frac{L \tan \alpha}{R + r}, \quad (11.2)$$

$$\Pi_b = \frac{mgL^3}{EI_b R}, \quad \Pi_c = \frac{c_1}{M_0 \sqrt{g/L}}, \quad (11.3)$$

$$\text{St}_p = \tau_p \sqrt{\frac{g}{L}}, \quad \Lambda = \frac{C_b mg \cos \beta}{2(R + r)}. \quad (11.4)$$

Together with μ_r , μ_s , β , and actuator limits, these parameters distinguish several regimes:

- (a) *geometry dominated*: large Π_α , low loss, stiff rods;
- (b) *dissipation dominated*: rolling loss erases early pulse energy;
- (c) *compliance dominated*: Λ is large enough that the fold precedes the desired release;
- (d) *actuator limited*: finite opening rate or jerk prevents impulsive control;
- (e) *risk dominated*: timing variance and narrow terminal geometry determine the optimum.

The conjectured one-pulse strategy is expected only in a subset of this parameter space.

12 System identification and falsification protocol

A credible validation requires measurement of one physical apparatus rather than guessed parameters. The following protocol identifies the hierarchy of models and tests [Proposition 10.2](#).

1. **Rigid geometry.** Measure R , r , L , β , and the full handle map $d(x, u)$. Compare side-view center height with [Equations \(2.5\)](#) and [\(2.8\)](#).
2. **Rolling kinematics.** Mark the ball and estimate spin from high-speed video. Test $v = R\delta\omega$ in low-slip trials.
3. **Dissipation.** Run coast-down experiments at several fixed openings to identify μ_r , c_1 , and c_2 .
4. **Compliance.** Apply known transverse loads at several positions to identify $C_b(x)$ and the first beam modes. Test the fold scaling in [Equation \(7.6\)](#) by changing ball mass or rod stiffness.
5. **Pulse work.** Record synchronized handle motion, separation, ball position, speed, spin, and rod vibration during controlled open-close cycles. Compare measured energy change with [Equation \(6.3\)](#).
6. **Hybrid guards.** Identify onset of slip, one-contact motion, and release as functions of $(x, v, \alpha, \dot{\alpha})$.
7. **Human strategies.** Compare smooth conservative, many-small-pulse, one-pulse, and unrestricted expert trials.
8. **Unconstrained optimization.** Solve [Equations \(10.6\)](#) and [\(10.7\)](#) using direct collocation, dynamic programming, or stochastic model-predictive control without prescribing pulse count.

The one-pulse conjecture gains support only if the independently calibrated optimizer repeatedly discovers a dominant intermediate pulse and predicts how its location changes after controlled perturbations to mass, friction, stiffness, slope, or target geometry. If the optimizer instead finds multiple pulses or a smooth singular arc, that outcome refutes the conjecture for the tested regime.

13 Discussion

13.1 What the model establishes

Several conclusions follow directly from the stated assumptions.

First, the geometry-induced drive is exact within the symmetric cross-sectional model: opening the rods lowers the ball center, and uphill acceleration begins when [Equation \(4.2\)](#) is satisfied.

Second, the symmetric no-slip geometry has an exact local rolling radius $R\delta$. This causes effective rotational inertia to grow as δ^{-2} and regularizes the apparent terminal force singularity at zero speed.

Third, small compliance can have a non-small qualitative effect. The static closure in Equation (7.5) develops a calculable fold before the rigid geometric limit. This provides a concrete mechanism for early release, snap-through, and extreme sensitivity to rod stiffness.

Fourth, the speed-risk tradeoff is not merely verbal. Under the release-error model, hit probability is strictly decreasing in speed. High speed becomes rational only because the scoring geometry or reach constraint changes the payoff set.

13.2 What remains conjectural

The model does not prove that a particular commercial machine has a one-pulse optimum. That conclusion depends on measured dissipation, control limits, compliance, terminal geometry, and the player's objective. Nor does the reduced model replace the full three-dimensional contact equations near slip or release. It provides a structured hierarchy:

$$\text{rigid reduced model} \subset \text{compliant modal model} \subset \text{full nonsmooth contact model}. \quad (13.1)$$

The simpler levels are useful only while validated against the next level and against experiment.

13.3 Interpretation of maximum risk

“Maximum risk” should not mean arbitrary recklessness. It means operation near the boundary of the target-specific viability kernel. For a given target and uncertainty model, there is a largest useful aggressiveness at which the fractional score gain still compensates for the fractional loss of success probability. A record-seeking player with many attempts chooses a different boundary reserve from a player maximizing average score per attempt.

14 Conclusion

The Space Force / Shoot-the-Moon mechanism is a compact example of controlled transport on a moving, compliant constraint manifold. A single support variable δ controls center height, contact orientation, rolling radius, normal load, and effective inertia. Rod opening creates geometry-induced uphill drive, while time-dependent opening allows the player to exchange energy with the ball. Rolling resistance and traction constraints limit this mechanism. Slight rod bending is amplified near release and, in the reduced static model, creates an explicit fold before the rigid support limit. Contact loss then turns the problem into a hybrid transition from rolling to free flight.

For scoring, the relevant object is the target-specific viability kernel. Maximum-performance trajectories approach its boundary; timing uncertainty makes terminal hit probability decrease strictly with speed. This formalizes the observed paradox that the route to the maximum score can also be the route with the smallest robustness margin. Bang-bang theory makes a short high-authority pulse mechanically plausible. Whether the optimal strategy contains exactly one intermediate pulse is now a precise, falsifiable question for system identification and unconstrained optimal-control computation.

A Full vector rolling equations

Let $\mathbf{r}_C$ be the ball-center position, $\boldsymbol{\omega}$ its angular velocity, and $\boldsymbol{\rho}_i$ the vector from the center to contact i . The sphere contact-point velocity is

$$\mathbf{v}_{c,i} = \dot{\mathbf{r}}_C + \boldsymbol{\omega} \times \boldsymbol{\rho}_i. \quad (\text{A.1})$$

If the rod material velocity is $\mathbf{v}_{r,i}$, ideal rolling requires

$$\mathbf{P}_{T,i}(\mathbf{v}_{c,i} - \mathbf{v}_{r,i}) = \mathbf{0}, \quad \mathbf{P}_{T,i} = \mathbf{I} - \mathbf{n}_i \mathbf{n}_i^T. \quad (\text{A.2})$$

The Newton–Euler equations are

$$m\ddot{\mathbf{r}}_C = m\mathbf{g} + \sum_i (N_i \mathbf{n}_i + \mathbf{F}_{t,i}) + \mathbf{F}_{\text{air}}, \quad (\text{A.3})$$

$$\mathbf{I}_C \dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (\mathbf{I}_C \boldsymbol{\omega}) = \sum_i \boldsymbol{\rho}_i \times \mathbf{F}_{t,i} + \boldsymbol{\tau}_{\text{ext}}. \quad (\text{A.4})$$

Together with [Equations \(5.3\)](#) and [\(8.1\)](#), these form a differential-algebraic or complementarity system. The scalar equation [Equation \(3.7\)](#) is its symmetric centerline reduction.

B Compact reduced model

A minimal simulation model is

$$d = d_0 + 2x \tan \alpha, \quad \eta = \frac{d}{2(R+r)}, \quad \delta = \sqrt{1 - \eta^2}, \quad (\text{B.1})$$

$$V = mg(x \sin \beta + (R+r)\delta), \quad M = m \left(1 + \frac{\kappa}{\delta^2}\right), \quad (\text{B.2})$$

$$M\dot{v} = -V_x - \frac{1}{2}M_x v^2 - M_\alpha \dot{\alpha} v - F_d, \quad \dot{x} = v, \quad (\text{B.3})$$

$$F_d = \mu_r mg \frac{\cos \beta}{\delta} \text{sgn}(v) + c_1 v + c_2 |v|v. \quad (\text{B.4})$$

Add the beam equation [Equation \(7.1\)](#) or a modal reduction, monitor the guards in [Equation \(8.4\)](#), and switch to [Equation \(8.2\)](#) at release.

C Notation summary

Symbol	Meaning
R, r	ball radius and rod radius
m, I	ball mass and central moment of inertia
x, L	longitudinal position and rail length
α	half of the included rod opening angle
β	mean uphill inclination of the rail-center plane
d	center-to-center rod separation
$a = R + r$	ball-center to rod-center distance at contact
$\eta = d/(2a)$	normalized separation
$\delta = \sqrt{1 - \eta^2}$	dimensionless support margin
$h = a\delta$	ball-center height above the rod-center plane
M	geometry-dependent translational effective mass
$N_i, \mathbf{F}_{t,i}$	normal and tangential contact forces
w_i, q_b	rod deflection and retained beam mode
Λ	dimensionless static compliance parameter
$\mathcal{K}_k(T)$	target- k viability kernel over horizon T
ε	speed reserve below a projected viability boundary

References

- P. Xu, R. E. Groff, and T. C. Burg. The rigid body dynamics of Shoot-the-Moon game and model-based controller design. In *Proceedings of the 2010 American Control Conference*, pages 396–401, 2010. doi:10.1109/ACC.2010.5530797.

- P. Xu, R. E. Groff, and T. C. Burg. Dynamics and control of the Shoot-the-Moon tabletop game. *Journal of Dynamic Systems, Measurement, and Control*, 134(5):051006, 2012. doi:10.1115/1.4006223.
- P. Xu. *Dynamics and Control of the Shoot-the-Moon Tabletop Game*. M.S. thesis, Clemson University, 2011.
- K. De Moerlooze, F. Al-Bender, and H. Van Brussel. Modeling of the dynamic behavior of systems with rolling elements. *International Journal of Non-Linear Mechanics*, 46(1):222–233, 2011. doi:10.1016/j.ijnonlinmec.2010.09.003.
- K. L. Johnson. *Contact Mechanics*. Cambridge University Press, 1985.
- J. J. Kalker. *Three-Dimensional Elastic Bodies in Rolling Contact*. Kluwer Academic Publishers, 1990.
- V. Putkaradze and S. Rogers. On the normal force and static friction acting on a rolling ball actuated by internal point masses. arXiv:1810.04010, 2018.
- B. Brogliato. *Nonsmooth Mechanics: Models, Dynamics and Control*. Springer, third edition, 2016.
- J.-P. Aubin. *Viability Theory*. Birkhäuser, 1991.
- P. Saint-Pierre. Approximation of the viability kernel. *Applied Mathematics and Optimization*, 29:187–209, 1994. doi:10.1007/BF01204182.
- L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. Interscience Publishers, 1962.
- L. Meirovitch. *Fundamentals of Vibrations*. McGraw-Hill, 2001.