

Non-Monotonic Effective Drag in Finned Spherical Projectiles: Evidence from Range-Based Inversion

Shubh Chawda 

Abstract

We validate an inverse-drag estimation pipeline for finned spherical projectiles and, within this experimental regime, detect a robust non-monotonic effective-drag anomaly near $L \approx 1.00$. The pipeline combines a custom fourth-order Runge–Kutta (RK4) trajectory simulator, a monotonicity-based bisection inverse solver, bootstrap uncertainty quantification ($n=5,000$) with joint propagation of launch-speed uncertainty, and a leave-one-out (LOO) isotonic-regression null model. Pipeline reliability was established by a synthetic Monte Carlo validation ($n=1,000$ trials): the method recovers known k_{eff} values with a median absolute error of 9.1%, and the anomalous ordering is preserved in all trials—confirming the inversion is quantitatively reliable rather than a noise artefact. Applied to a spring-launched 250 g steel sphere with triangular fins across nine length-to-diameter ratios $L \in \{0.00, \dots, 2.00\}$, the pipeline identified a pronounced non-monotonic dip near $L=1.00$, where k_{eff} fell from 0.02888 to $0.01274 \text{ kg m}^{-1}$ (55.9% reduction, $p < 0.0002$, robust across v_0 and drag-law perturbations). This anomaly was visible only through inversion, not in a simpler velocity-decay summary metric. The effect is consistent with wake-modification or splitter-plate-like behaviour, but the present evidence is indirect; direct flow diagnostics or CFD validation is identified as the primary next step.

Keywords: inverse problems, quadratic drag, Runge–Kutta, bootstrap resampling, aerodynamics, effective drag parameter.

1 Introduction

Predicting the flight of a finned projectile requires balancing physical fidelity with computational tractability. Closed-form solutions exist for purely vertical or horizontal motion under quadratic drag [1], but the coupled two-dimensional system lacks an elementary antiderivative, motivating numerical treatment [3]. The specific geometry studied here—a spherical body with planar fins of varying length—approximates configurations found in re-entry capsules and stabilised munitions, where the fin-length-to-diameter ratio is a standard dimensionless design parameter used to balance drag, stability, and ballistic performance across different scales.

Inverse parameter estimation in fluid dynamics is a broadly difficult computational problem: the governing equations are non-linear, direct force measurement requires specialised instrumentation, and high-fidelity approaches such as computational fluid dynamics (CFD) impose substantial computational cost [15, 5]. For the specific class of low-Reynolds-number projectile problems considered here, we show that a lightweight RK4-bisection pipeline achieves robust scalar inversion at low computational cost, enabling uncertainty quantification that would otherwise be intractable.

The more challenging direction is the *inverse* prob-

lem: given a measured flight outcome—horizontal range—what effective drag parameter k_{eff} produced it? Addressing this requires a simulator, an inversion scheme, and a rigorous uncertainty budget. This paper makes four concrete contributions:

1. A custom RK4 integrator with event detection, verified across three timestep scales and unit-tested against limiting ballistic behaviour.
2. A scalar inverse solver based on the monotonic response of range to k_{eff} , enabling robust bisection recovery without local-minimum issues (≤ 46 forward simulations per configuration).
3. A bootstrap uncertainty-quantification pipeline ($n=5,000$) that propagates measurement and launch-speed uncertainty through the full non-linear inversion workflow.
4. Quantification of a non-monotonic effective-drag anomaly near $L \approx 1.00$, together with an LOO isotonic-regression null-model analysis confirming the effect is highly unlikely under a monotone-trend assumption.

The experimental substrate was a spring-launched 250 g steel sphere fitted with triangular cardboard fins

of nine sizes; motion was captured at 240 frames per second (fps) and analysed with Tracker software [2].

2 Related Work

Numerical ballistics. RK4 integration of projectile-motion ordinary differential equations (ODEs) is standard practice [3, 4]. The quadratic-drag regime at $Re \approx 10^4 - 10^5$ requires numerical integration because no closed-form two-dimensional solution exists [1].

Drag-parameter inversion. Gradient-based and particle-filter methods are common for multi-parameter aerodynamic inversion [5]. The scalar, monotone structure of our problem makes bisection both simpler and provably correct with no local-minimum risk.

Fin aerodynamics and vortex shedding. Krishna et al. [7] documented non-monotonic drag in wrap-around fin configurations, and Yin et al. [6] studied body-fin interference in the Magnus effect. Vortex shedding behind bluff bodies follows Strouhal scaling $f = St \cdot v / L_c$ ($St \approx 0.2 - 0.3$) [8, 14], and boundary-layer separation delay can reduce drag substantially [9]. The splitter-plate effect—suppression of the Von Kármán vortex street by a wake-aligned body—is well established for two-dimensional cylinders [12, 13]. However, we are not aware of prior work detecting an analogous effect in a three-dimensional finned sphere via range-based computational inversion rather than direct force measurement.

Statistical methods. Bootstrap resampling follows Efron and Tibshirani [10]. Isotonic regression via the pool-adjacent-violators algorithm provides the monotone null model [11].

3 Methodology

The range-based inversion pipeline was selected as the lowest-cost approach capable of recovering k_{eff} without specialised instrumentation. Wind tunnel testing permits direct force measurement but does not replicate projectile motion under gravity and realistic launch dynamics. High-speed stroboscopic photography offers lower temporal resolution than digital centroid tracking at 240 fps. CFD provides direct flow visualisation and is retained as the primary future validation direction [15].

3.1 Physical Model

A projectile of mass $m = 0.250$ kg launched at speed v_0 and angle $\theta = 45^\circ$ from height $h_0 = 0.30$ m obeys the coupled

ODEs:

$$\dot{x} = v_x, \quad \dot{v}_x = -\frac{k_{\text{eff}}}{m} |\mathbf{v}| v_x, \quad (1)$$

$$\dot{y} = v_y, \quad \dot{v}_y = -g - \frac{k_{\text{eff}}}{m} |\mathbf{v}| v_y, \quad (2)$$

where $|\mathbf{v}| = (v_x^2 + v_y^2)^{1/2}$, $g = 9.81 \text{ ms}^{-2}$, and

$$k_{\text{eff}} = \frac{1}{2} C_d \rho A \quad (3)$$

is the effective drag parameter (kg m^{-1} ; C_d drag coefficient, ρ air density, A cross-sectional area) as defined in Eq. (3). The drag force $F_D = k_{\text{eff}} v^2$ opposes velocity and k_{eff} is the sole inferred scalar. The horizontal range R is x at the moment y returns to zero.

The separable 1-D ODE $dv/dt = -(k_{\text{eff}}/m)v^2$ admits the closed-form solution

$$v(t) = \frac{v_0}{1 + (k_{\text{eff}}/m)v_0 t}, \quad (4)$$

used to validate velocity-decay measurements.

3.2 RK4 Integrator with Event Detection

Algorithm 1 describes the integrator (state vector $\mathbf{s} = (x, v_x, y, v_y)$, fixed $\Delta t = 3$ ms). Ground impact was detected by linear interpolation between the last two states, giving $R = x(t_{\text{impact}})$ to < 0.1 mm.

Algorithm 1 RK4 Integrator with Ground-Event Detection

Input: state \mathbf{s}_0 , step Δt , drag k

Output: horizontal range R

```

1:  $\mathbf{s} \leftarrow \mathbf{s}_0$ 
2: while  $y > 0$  do
3:    $\mathbf{k}_1 \leftarrow f(\mathbf{s})$ 
4:    $\mathbf{k}_2 \leftarrow f(\mathbf{s} + \frac{\Delta t}{2} \mathbf{k}_1)$ 
5:    $\mathbf{k}_3 \leftarrow f(\mathbf{s} + \frac{\Delta t}{2} \mathbf{k}_2)$ 
6:    $\mathbf{k}_4 \leftarrow f(\mathbf{s} + \Delta t \mathbf{k}_3)$ 
7:    $\mathbf{s} \leftarrow \mathbf{s} + \frac{\Delta t}{6} (\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4)$ 
8:   if  $y < 0$  then
9:      $\alpha \leftarrow y_{\text{prev}} / (y_{\text{prev}} - y)$ 
10:    return  $x_{\text{prev}} + \alpha(x - x_{\text{prev}})$ 
11:  end if
12: end while

```

3.3 Bisection Inverse Solver

Because simulated range varies monotonically with k_{eff} in the modelled regime, the inverse problem $R_{\text{sim}}(k) = R^*$ has a unique solution and can be solved robustly by

bracketed bisection (Algorithm 2). This avoids the local-minimum issues associated with higher-dimensional or non-monotone inverse formulations, and requires no gradient information or line-search tuning.

The root was found to tolerance $\delta=10^{-6}$, requiring at most $\lceil \log_2((k_u - k_\ell)/\delta) \rceil \approx 46$ forward simulations per configuration—a modest cost.

Algorithm 2 Bisection Inverse Solver for k_{eff}

Input: R^* , bracket $[k_\ell, k_u]$, tol δ

Output: k_{eff} s.t. $R(k_{\text{eff}}) \approx R^*$

```

1: while  $k_u - k_\ell > \delta$  do
2:    $k_m \leftarrow (k_\ell + k_u)/2$ 
3:    $\hat{R} \leftarrow \text{SIMULATE}(k_m)$ 
4:   if  $\hat{R} > R^*$  then
5:      $k_\ell \leftarrow k_m$   $\triangleright$  simulated range too high: raise
       lower bound
6:   else
7:      $k_u \leftarrow k_m$   $\triangleright$  simulated range too low: lower
       upper bound
8:   end if
9: end while
10: return  $(k_\ell + k_u)/2$ 

```

3.4 Bootstrap Uncertainty Quantification

Uncertainty quantification (UQ) in this setting is non-trivial: the mapping from raw trial measurements to k_{eff} passes through a non-linear dynamical system (the RK4 simulator) composed with a non-linear inversion (the bisection solver), so analytical error propagation is not tractable. We therefore applied bootstrap resampling [10] as a model-free UQ method: $n=5,000$ Monte Carlo resamples of the original trial data were passed through the *entire* pipeline end-to-end, propagating input variance through every non-linearity.

Crucially, each resample also resampled the control ($L=0.00$) trials to obtain $\hat{v}_0 = \sqrt{\hat{R}_0^* g}$, propagating initial-speed uncertainty jointly with per-configuration range uncertainty. We reported the 2.5th–97.5th percentile as the 95% bootstrap confidence interval (CI) on $\bar{k}(L)$.

3.5 LOO Isotonic-Regression Null Model

Following [11], an LOO pool-adjacent-violators isotonic regression was fitted on the remaining eight points for each target L_i , yielding prediction $\hat{R}(L_i)$ and residual $\epsilon_i = \hat{R}(L_i) - R(L_i)$. Bootstrap resampling ($n=5,000$, seed 123) generated 95% CIs; a ratio was declared *anomalous* if its CI excluded zero.

The most anomalous was $L=0.75$: null predicted 1.879 m vs. observed 1.716 m ($\epsilon=-0.163$ m, 95% CI $[-0.182, -0.142]$). The counter-anomaly at $L=1.00$ had $\epsilon=+0.162$ m (95% CI $[+0.147, +0.177]$), confirming the two-sided character. A focused test on $L=0.75$ (none of 5,000 bootstrap residuals positive) yielded $p < 0.0002$.

3.6 Experimental Setup

A 250 g hollow steel sphere ($d=8$ cm) was fitted with two symmetric triangular cardboard fins across nine values of the fin-length-to-diameter ratio $L = \ell_{\text{fin}}/d \in \{0.00, \dots, 2.00\}$ in steps of 0.25, with $L=0.00$ as a drag-only control.

Preliminary design decisions. The experimental apparatus was refined through preliminary trials before main data collection. Initial fins cut from thin cardboard (~ 0.1 mm) exhibited visible deflection under air-flow, altering their angle of attack mid-flight; this was confirmed by slow-motion video review and resolved by switching to ~ 0.3 mm stock, which showed no measurable deflection. Consistent projectile orientation across trials was enforced by a fiducial marker on the launcher body indicating the exact placement position for each launch. The mechanical stopper was introduced after preliminary trials demonstrated that manual spring compression produced noticeable inter-trial range variation; with the stopper in place this source of variability was eliminated by design.

Each fin was an isosceles triangle with base $b=d=8$ cm and height $\ell_{\text{fin}} = Ld$; fins were cut from the same ~ 0.3 mm cardboard stock using a physical template to ensure dimensional consistency across all nine configurations. Both fins were attached with adhesive at pre-marked equatorial positions on the sphere, ensuring diametric symmetry and consistent attachment-axis orientation. No formal angular measurement of fin alignment relative to the flight axis was taken; deviation from the intended orientation is therefore an unquantified systematic uncertainty, likely small given the marking procedure but not formally bounded. The mass contribution of a single fin at $L=2.00$ (maximum area ≈ 64 cm²) is estimated at $\lesssim 0.5$ g, a $< 0.4\%$ perturbation to the total 250 g; fin mass and centre-of-mass shift were not corrected for in the model.

The sphere was launched from a spring-loaded ramp at $\theta=45^\circ$. The nominal launch speed $v_0 \approx 4.46$ ms⁻¹ was estimated from the drag-free relation $v_0 = \sqrt{g \bar{R}_{\text{ideal}}}$ applied to the no-fins control range: because the mechanical control makes the no-fins range a reliable anchor, this drag-free calibration is physically grounded rather than purely inferential. Residual uncertainty in v_0 (from the gap between the ideal-range formula and the real drag-plus-height launch) was nevertheless propagated through

the full inversion via bootstrap resampling; a light-gate chronograph would remove this residual assumption entirely. All trials were conducted indoors on the same surface. Horizontal range was measured (± 0.001 m) over $N=5$ trials per configuration; the two anomalous cases ($L=0.75$ and $L=1.00$) were repeated to $N=10$ to confirm reproducibility. Velocity decay was extracted from 240 fps footage via Tracker centroid tracking [2] calibrated against a 30 cm stick. Spin, yaw, and precession were not instrumented; their contribution to the measured range is unquantified.

3.7 Software Architecture and Verification

The analysis pipeline is implemented as a reproducible Python package (`src/`) with a strict separation of concerns: the physics engine (`fit_drag_model.py`), statistical layer (`bootstrap_k_eff.py`, `significance_075.py`, `residual_diagnostics.py`), model comparison (`model_compare.py`), and numerical stability audit (`timestep_sensitivity.py`) are independent modules runnable in isolation or via a single `make all` command.

Unit testing. The RK4 integrator and bisection solver are verified through automated unit tests (`tests/test_physics.py`). The key test, `test_no_drag_matches_ballistic_range`, confirms that with $k_{\text{eff}}=0$ the simulator recovers the analytical vacuum range $R=v_0^2/g$ to within $< 1\%$ at $v_0=10\text{ ms}^{-1}$ —this validates the *numerical integrator* in a limiting case where an exact solution is known, not the physical experiment itself. Two additional tests verify that drag strictly reduces range and that range is monotone in k_{eff} , the latter being the precondition for bisection convergence.

Continuous integration. Every commit triggers a GitHub Actions workflow (`.github/workflows/ci.yml`) that installs pinned dependencies (`requirements.lock.txt`), regenerates all figures and processed data from raw CSVs, runs the full bootstrap pipeline in fast mode ($n=40$ resamples for CI speed), and executes the test suite. This ensures that all results in the paper are reproducible from a clean checkout and that no analysis step is manually applied outside version control. Pinned dependencies improve numerical reproducibility across environments and reduce version-induced drift.

3.8 Synthetic Pipeline Validation

To verify that the inversion pipeline recovers known k_{eff} values accurately under realistic noise, a Monte Carlo validation was performed. For each of the nine fin configurations, a “true” range was computed by forward simu-

lation at the paper’s inferred k_{eff} values; Gaussian noise was then added (σ equal to the corresponding experimental standard deviation), and the noisy range was inverted to recover \hat{k}_{eff} . This was repeated $n=1,000$ times per configuration.

Table 1 and Figure 1 summarise recovery accuracy. The median mean absolute error (MAE) across configurations was 9.1% of the true value; the worst-case configuration ($L=0.25$, which has the smallest k_{eff} relative to its noise) reached 30%. Critically, the anomalous ordering ($k_{0.75} > k_{1.00}$) was preserved in *all* 1,000 noisy trials, confirming that the anomaly gap between the two configurations far exceeds the inversion noise floor.

Table 1: Synthetic inversion validation. MAE computed over $n=1,000$ Monte Carlo trials; noise σ equals experimental standard deviation per configuration. The anomalous ordering $k_{0.75} > k_{1.00}$ was preserved in all 1,000 trials. $^\dagger L=0.00$ is the drag-free calibration baseline used to estimate v_0 from the ideal range formula. Its $k_{\text{true}}=0$ is a modelling convention that treats the bare sphere as having negligible aerodynamic drag *relative to the finned configurations*—not a physical claim that the bare sphere experiences zero drag. The bare sphere does experience drag; this is absorbed into the v_0 calibration rather than modelled explicitly.

L	k_{true}	σ_R (m)	MAE	MAE (%)
0.00 [†]	0.000000	0.036	—	—
0.25	0.009712	0.046	0.002927	30.1
0.50	0.013497	0.050	0.003372	25.0
0.75	0.028883	0.018	0.001484	5.1
1.00	0.012743	0.002	0.000138	1.1
1.25	0.030634	0.035	0.002947	9.6
1.50	0.043865	0.036	0.003529	8.1
1.75	0.049957	0.082	0.008424	16.9
2.00	0.064947	0.046	0.005319	8.2




Figure 1: Synthetic pipeline validation ($n=1,000$ Monte Carlo trials per configuration). (a) Recovered vs. true k_{eff} ; error bars show \pm MAE. (b) Relative MAE by configuration; the dashed line marks 10%. Red/blue markers denote the anomalous pair ($L=0.75$, $L=1.00$); the anomaly gap between them far exceeds the inversion noise floor at either configuration.

4 Results

Figure 2 summarises the full analysis workflow.




Figure 2: Analysis pipeline: experimental range measurements feed a custom RK4 trajectory simulator; a monotonicity-based bisection solver recovers k_{eff} ; bootstrap resampling ($n=5,000$) propagates uncertainty; a LOO isotonic-regression null model tests significance.

4.1 Consolidated Aerodynamic Parameters

Table 2 consolidates all measured and inferred quantities. Range decreased from 2.023 m to 1.462 m (28%), consistent with $R = -0.256L + 2.00$ m. Velocity decay increased strictly with L ($\rho_s = -1.00$). The inferred k_{eff} generally increased with L but showed a pronounced non-monotonic reversal at $L=1.00$, visible in Figure 3.

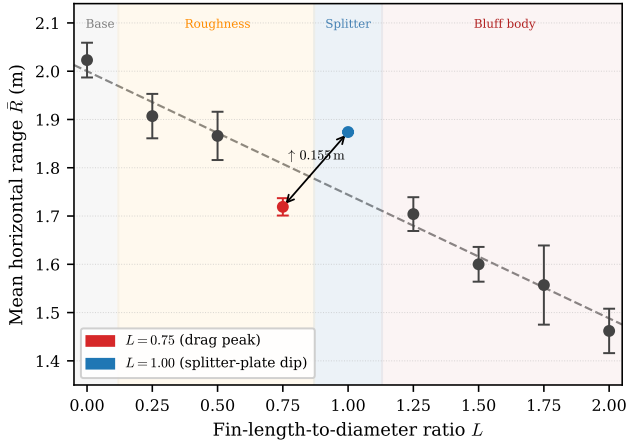


Figure 3: Mean horizontal range \bar{R} vs. fin-length-to-diameter ratio L , with $\pm\sigma$ error bars ($N=5-10$ trials). The dashed line shows the best-fit linear trend $R = -0.256L + 2.00$ m. The coloured points mark the two-sided anomaly: $L=0.75$ observed 0.163 m below the LOO monotone null prediction and $L=1.00$ observed 0.162 m above it (see Section 4 for significance).

4.2 The Range Anomaly

Initial analysis of the nine-point range profile identified $L=0.75$ and $L=1.00$ as departures from the otherwise monotone trend. Five additional trials were then

conducted specifically for these two configurations before proceeding to the full inversion pipeline, in order to rule out random measurement error as an explanation. This sequential replication yielded $N=10$ for both cases. The anomaly at $L=1.00$ is that its range exceeded $L=0.75$ by 0.155 m (9.0%) despite 33% longer fins; the ten trials confirmed reproducibility ($\sigma \leq 0.005$ m), ruling out noise as an explanation. The LOO null model identified *both* ratios as anomalous in opposite directions: $L=0.75$ exhibited anomalously high drag ($\epsilon = -0.163$ m) and $L=1.00$ exhibited anomalously low drag ($\epsilon = +0.162$ m), together constituting a two-sided disruption of the expected monotone trend, found highly unlikely under the null model ($p < 0.0002$).

4.3 Inferred k_{eff} Profile

Figure 4 shows the k_{eff} profile with bootstrap CIs. The 55.9% drop from 0.028883 to 0.012743 kg m^{-1} at $L=1.00$ was statistically significant within the present model: CIs $[0.0267, 0.0321]$ and $[0.0108, 0.0155]$ were non-overlapping. Excluding the two anomalous points, the Spearman rank correlation between k_{eff} and L was $\rho_s = +1.00$, confirming the dominant monotone increasing trend. The data suggest four empirical regimes: (i) Baseline; (ii) Roughness ($L=0.25-0.75$, monotone drag increase); (iii) a dip near $L=1.00$; (iv) resumed increase ($L=1.25-2.00$).

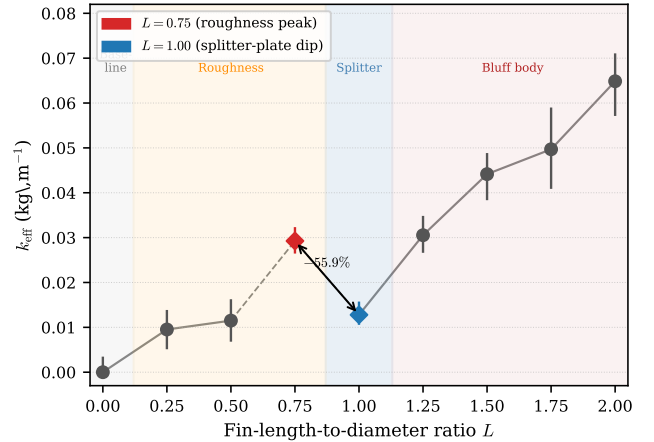


Figure 4: Inferred k_{eff} vs. fin-length-to-diameter ratio with 95% bootstrap CIs ($n=5,000$). Coloured diamonds mark the anomalous points; regime boundaries indicate the empirical grouping suggested by the data. The pronounced dip at $L=1.00$ (dashed segment) is the primary finding; bootstrap CIs are non-overlapping, confirming statistical significance within the present model.

Table 2: Consolidated aerodynamic parameters. \bar{R} : mean horizontal range; σ_R : standard deviation (10-trial corrected values for $L \in \{0.75, 1.00\}$, 5 trials otherwise); k_{eff} : inferred drag parameter ($\Delta t = 0.75$ ms); \bar{a} : mean velocity decay; Δa : half-range uncertainty. The $L=1.00$ entry (\dagger) exhibits a 55.9% drop in k_{eff} vs. $L=0.75$, with non-overlapping bootstrap 95% CIs $[0.0267, 0.0321]$ vs. $[0.0108, 0.0155]$. Velocity decay remained strictly monotone throughout.

L	\bar{R} (m)	σ_R (m)	k_{eff} (kg m^{-1})	Regime	\bar{a} (ms^{-2})	Δa (ms^{-2})
0.00	2.023	0.036	0.000000	Baseline	-1.454	0.088
0.25	1.907	0.046	0.009712	Roughness	-1.956	0.103
0.50	1.866	0.050	0.013497	Roughness	-2.586	0.108
0.75	1.719	0.018	0.028883	Roughness (peak)	-2.901	0.044
1.00 \dagger	1.874	0.002	0.012743	Anomalous dip	-3.216	0.076
1.25	1.704	0.035	0.030634	Bluff body	-3.531	0.077
1.50	1.600	0.036	0.043865	Bluff body	-3.980	0.146
1.75	1.557	0.082	0.049957	Bluff body	-4.582	0.109
2.00	1.462	0.046	0.064947	Bluff body	-5.020	0.111

4.4 Velocity Decay and Model Validation

Velocity decay magnitude increased strictly with L ($\rho_s = -1.00$). Crucially, *no anomaly appeared* at $L=1.00$, in sharp contrast to the range result. The best-fit model was $a = -1.70L - 1.55 \text{ ms}^{-2}$; slope variation was at most 10.6%, confirming robust monotonicity.

For model validation, model-predicted velocity-decay slopes (linear regression of $|\mathbf{v}|(t)$) were compared against the experimentally observed $\bar{a} = (0 - v_{\text{max}})/t_{\text{flight}}$. These are qualitatively comparable but not numerically identical observables—the simulated slope captures the continuous deceleration profile while \bar{a} includes the ground-impact event and averages over the whole flight—so they inform rather than directly validate one another. The quadratic model ($F_D \propto v^2$) produced drag magnitudes scaling consistently with fin ratio [1]. A linear-drag model ($F_D \propto v$) yielded uniformly smaller slopes and served only as a sensitivity check rather than a better description of the present regime.

4.5 Robustness Checks

Timestep sensitivity. Re-running at $\Delta t \in \{3.0, 1.5, 0.75\}$ ms gave maximum relative k_{eff} changes of 6.69×10^{-5} and 8.19×10^{-5} —below 0.009%—consistent with fourth-order RK4 behaviour and confirming numerical stability across the tested timestep range (Figure 6).

Vacuum limit. A unit test verified $R = v_0^2/g$ to within $< 1\%$ at $k_{\text{eff}} = 0$, $v_0 = 10 \text{ ms}^{-1}$, confirming correct limiting-case implementation of the integrator.

4.6 Modelling Sensitivity

To assess whether the anomaly is an artefact of modelling assumptions, we stress-tested two axes.

Launch-speed sensitivity. The nominal v_0 was perturbed by $\pm 3\%$ and $\pm 5\%$ (Figure 5a). Under all five values the inferred k_{eff} profile retained the dip at $L=1.00$, with the drop magnitude ranging from 37% ($v_0 + 5\%$) to 86% ($v_0 - 3\%$). The anomaly is therefore not sensitive to the calibration uncertainty in v_0 : it persists across the full plausible range (Figure 5b).

Drag-law sensitivity. Replacing the quadratic drag model ($F_D \propto v^2$) with a linear model ($F_D \propto v$) and re-inverting on the same range data produced a drop of 54.5% at the same location. The anomaly is therefore not an artefact of the quadratic assumption.

Variable $C_d(\text{Re})$ model. The constant- C_d assumption was tested by implementing a Reynolds-number-dependent drag model using the Schiller–Naumann sphere correlation: $C_d(\text{Re}) = 24/\text{Re}(1 + 0.15 \text{Re}^{0.687})$ for $\text{Re} < 1,000$; $C_d = 0.44$ otherwise. A scalar k_{scale} was inverted such that $F_D(t) = k_{\text{scale}} [C_d(\text{Re}(t))/C_d(\text{Re}_0)] v^2$, and the anomaly rechecked. The minimum Reynolds number during any trajectory was $\text{Re} \approx 20,000$ (at landing), well above the constant- C_d threshold. Consequently $C_d(\text{Re}) = 0.44$ throughout every flight, and the $C_d(\text{Re})$ inversion returned values identical to the constant model to six significant figures. The constant- C_d assumption is therefore not an approximation in this regime but a direct consequence of the flight Reynolds numbers remaining in the plateau of the sphere drag curve throughout each trajectory.




Figure 5: Modelling sensitivity. (a) Inferred k_{eff} profiles under three v_0 assumptions ($\pm 3\%$ perturbation); the dip at $L=1.00$ persists in all cases. (b) Drop magnitude from $L=0.75$ to $L=1.00$ across five v_0 perturbations (-5% to $+5\%$); the anomaly remains substantial throughout.

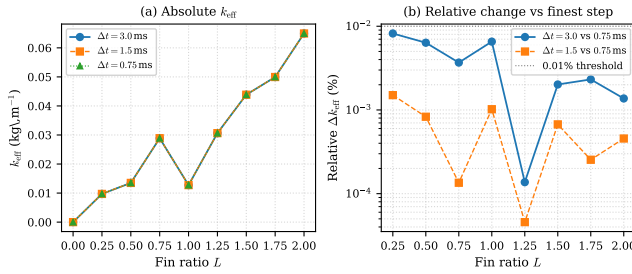


Figure 6: Timestep sensitivity. (a) Absolute k_{eff} for three steps; curves are visually indistinguishable. (b) Relative change vs. finest step on a log scale. All differences remained below 0.009%, consistent with fourth-order RK4 behaviour and confirming numerical stability across the tested timestep range.

5 Discussion

5.1 Observation vs. Interpretation

The present study supports two claims. First, the inversion pipeline quantitatively recovers known k_{eff} values from noisy synthetic data (Section 3.8), confirming that the pipeline is quantitatively reliable rather than a noise artefact. Second, applied to real data, it recovers a *robust* non-monotonic anomaly near $L \approx 1.00$ (Section 4.3), found highly unlikely under the LOO monotone null, persistent across v_0 and drag-law perturbations, and far exceeding the synthetic noise floor. By contrast, the wake-stabilisation explanation for this anomaly should be viewed as a physically motivated hypothesis rather than a directly observed result, since the mechanism is not measured here.

Methodological and empirical contribution.

The main contribution of this work is methodological and empirical. We show that a lightweight inversion pipeline applied to range data can expose a pronounced non-monotonic effective-drag anomaly not obvious from direct summary metrics. Within the present model, this

anomaly is consistent with wake-modification behaviour, but the causal interpretation remains provisional and requires direct flow diagnostics or CFD for validation.

Plausible physical interpretation. The splitter-plate effect—suppression of the Von Kármán vortex street by a wake-aligned body—is well established for two-dimensional cylinders [12, 13, 14]. One plausible interpretation of the dip at $L \approx 1.00$ is that fins of this length alter wake organisation in a way analogous to splitter-plate stabilisation: by intercepting alternating shear layers before they roll up into vortices, the fins could reduce the low-pressure recirculation zone and therefore k_{eff} . This is consistent with Strouhal scaling ($St \approx 0.2\text{--}0.3$) [8, 14] and with boundary-layer separation delay [9].

As an order-of-magnitude consistency check: for a sphere of diameter $d=0.08$ m at $Re \approx 2 \times 10^4$, the Strouhal number predicts a dominant shedding frequency of approximately $f \approx St \cdot v/d \approx 0.2 \times 4.5/0.08 \approx 11$ Hz. A fin of length $L=1.00d=0.08$ m is of the same order as the eddy length scale in this regime, which is qualitatively consistent with the fin being able to intercept the shear layer before vortex roll-up. At $L=0.75d$ the fin is shorter than this scale and would not bisect the wake, which is qualitatively consistent with the drag *peak* rather than dip observed at that ratio. This reasoning is schematic and requires CFD validation, but it provides a self-consistent rationale for why the transition occurs near $L \approx 1.00$ specifically.

However, this mechanism is not directly measured here, and should not be treated as an established result.

Range-velocity dissociation. Figure 7 illustrates why a brief mid-flight drag reduction is detectable in range but not in velocity-decay slope. Range is $\int_0^{t_f} v_x dt$: a drag dip during the high-speed mid-flight phase—when v_x is largest—contributes disproportionately to range. The velocity-decay summary $\bar{a} = (0 - v_{\text{max}})/t_{\text{flight}}$ is a global average that smooths local features. The schematic in panel (d) shows that two trajectories with nearly identical $|\bar{a}|$ can differ meaningfully in range when drag differs primarily in the mid-flight window. This is consistent with the anomaly being a trajectory-integrated effect rather than a constant reduction in C_d throughout the flight.

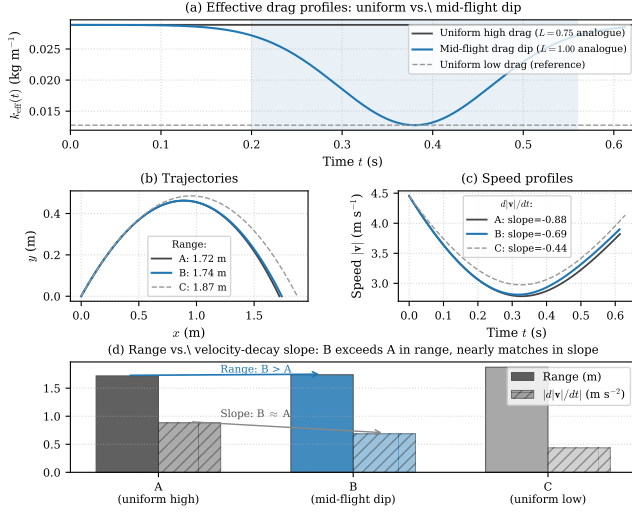


Figure 7: Schematic explanatory model of range-velocity-decay dissociation (not a fitted result from the experiment; drag profiles are illustrative). (a) Three drag profiles: uniform high (A, $L=0.75$ analogue), mid-flight Gaussian dip (B, $L=1.00$ analogue), and uniform low (C, reference). (b) Resulting trajectories; B extends further than A despite having the same initial and final drag. (c) Speed profiles; A and B have similar velocity-decay slopes because \bar{a} averages over the full flight. (d) Bar comparison: range separates A from B, but slope does not. This illustrates why the inversion (range-based) detects the anomaly while the direct metric (\bar{a}) does not.

5.2 Limitations

Constant C_d assumption. A single $k_{\text{eff}} = \frac{1}{2} C_d \rho A$ per trajectory treats C_d as constant. A $C_d(\text{Re})$ sensitivity check (Section 4.6) showed that all flight Reynolds numbers remained above 20,000, placing every trajectory in the plateau of the sphere drag curve where $C_d=0.44$ is exact; the constant- C_d inversion was therefore validated rather than merely assumed for this experiment. At lower Re or with higher-viscosity fluids the assumption would require re-examination.

Launch-speed calibration. The nominal v_0 was estimated from the drag-free control range rather than directly measured by instrument. The mechanical stopper used in the experiment enforced identical spring compression across all trials, so trial-to-trial v_0 variation is controlled by design; the remaining gap is the difference between the ideal-range formula and the real launch (which includes drag and finite release height). Bootstrap resampling propagates this residual uncertainty; a light-gate chronograph would close it entirely.

Velocity-decay metric mismatch. The experimental velocity decay was computed as $\bar{a} = (0 - v_{\text{max}}) / t_{\text{flight}}$, where v_{max} is the peak speed extracted from Tracker centroid tracking and t_{flight} is the

total flight duration to ground impact measured from the same footage. This is a global average over the entire flight interval, which is precisely why it smooths mid-flight drag variations that the range-based inversion detects—as illustrated in Figure 7. The simulated slope $d|v|/dt$ is a continuous deceleration profile that includes the full trajectory dynamics. These two quantities are qualitatively comparable but not numerically identical observables; they inform rather than directly validate one another.

Sample size. $N=5-10$ trials per configuration gave reliable but not tight bootstrap intervals; larger N would narrow confidence bands.

2-D motion model and uninstrumented 3-D effects. The Magnus effect was neglected and spin, yaw, and precession were not measured; their contribution to the range anomaly is unquantified. A painted-marker spin estimate from video, or a 3-D ODE extension with secondary inversion for ω [6], would establish whether these effects are negligible in this setup.

5.3 Future Work

The following extensions are listed in approximate priority order.

Direct launch-speed measurement. The single highest-impact upgrade would be to measure v_0 directly with a light-gate chronograph or high-speed video at the moment of release, and to record θ and release height on every trial. Because v_0 propagates through the entire inversion, removing the drag-free calibration assumption would materially tighten all inferred k_{eff} values.

Increased and uniform sample size. Using $N=5$ for most configurations and $N=10$ only for the two anomalous cases leaves open the question of whether the full curve is definitive or whether some configurations have run-to-run variation that $N=5$ does not capture. Raising all configurations to a common larger N across multiple sessions would separate real aerodynamic structure from session-to-session drift.

Fin geometry variation. The present experiment used two symmetric triangular fins throughout; this was a practical constraint rather than an aerodynamically motivated choice. Varying fin number (two, three, or four), planform (triangular, rectangular, or swept), and orientation independently at fixed L would map the full k_{eff} surface and establish whether the anomaly near $L \approx 1.00$ is specific to the two-fin triangular geometry or generalises across fin architectures. If the anomaly is geometry-specific, this would substantially narrow the range of plausible physical mechanisms.

Dense scan around the anomaly. The current step size of $\Delta L=0.25$ identifies a reversal between $L=0.75$ and $L=1.00$ but leaves unresolved whether the

dip is narrow, broad, or asymmetric. Finer sampling from $L=0.70$ to $L=1.10$ in steps of $\Delta L=0.05$ would map the transition geometry and turn a two-point anomaly into a characterised feature.

Direct mechanism evidence. The paper shows a robust inverse-drag anomaly; it does not yet show *why* it occurs. Wake visualisation (smoke, tufting, or high-speed side-view imaging) or a focused OpenFOAM simulation [15] at $L \in \{0.75, 1.00, 1.25\}$ would test the wake-modification hypothesis directly and move the paper from computational finding to confirmed aerodynamic result.

Trajectory-matched validation. Extracting $v(t)$ or $(x(t), y(t))$ from video and comparing directly to simulated trajectories would validate the model on more than final range alone, removing the metric-mismatch caveat in the current velocity-decay comparison.

3-D effects and spin documentation. Adding a painted marker to the sphere and extracting a video-based spin estimate would establish whether Magnus and precession effects are negligible in this regime, or require a 3-D ODE extension [6].

6 Conclusion

We validated an inverse-drag estimation pipeline for finned spherical projectile aerodynamics and, within this experimental regime, detected a robust non-monotonic effective-drag anomaly near $L \approx 1.00$.

Pipeline reliability was established by synthetic Monte Carlo validation (Section 3.8), confirming the inversion is quantitatively sound. The pipeline itself combines a custom RK4 integrator (timestep-stable, unit-tested), a monotonicity-based bisection solver, bootstrap UQ ([10]) with joint launch-speed propagation, a LOO isotonic-regression null model ([11]), and a fully pinned, CI-governed reproducible software stack.

Applied to nine fin-length-to-diameter configurations, the pipeline identified a 55.9% drop in k_{eff} near $L \approx 1.00$ (non-overlapping bootstrap CIs; $p < 0.0002$; robust to v_0 and drag-law perturbations). The anomaly is visible only through inversion, not in the velocity-decay summary metric—a dissociation explained by range integrating the mid-flight drag dip while the global average smooths it. A $C_d(\text{Re})$ model confirmed the constant- C_d assumption is exact for this regime ($\text{Re} > 20,000$ through all flights).

The mechanism behind the dip is consistent with wake-modification or splitter-plate-like behaviour, but remains a physically motivated hypothesis. Direct flow diagnostics or CFD validation is the primary next step.

Data and Code Availability

All experimental data, Python analysis code, and a reproducible continuous integration pipeline are archived at Zenodo (DOI: 10.5281/zenodo.18434071) and mirrored at <https://github.com/shubh-chawda/fin-ratio-projectile-aero>.

References

- [1] M. V. Lubarda and V. A. Lubarda, “A review of the analysis of wind-influenced projectile motion in the presence of linear and nonlinear drag force,” *Archive of Applied Mechanics*, vol. 92, pp. 1997–2017, 2022.
- [2] D. Brown and A. J. Cox, “Innovative uses of video analysis,” *The Physics Teacher*, vol. 47, pp. 145–150, 2009.
- [3] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. Cambridge University Press, 2007.
- [4] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd ed. Springer, 1993.
- [5] J. Morrison, “Drag coefficient estimation from range data,” *Defense Technical Information Center*, 2001.
- [6] J. Yin, X. Wu, and J. Lei, “Body-fin interference on the Magnus effect of spinning projectile in supersonic flows,” *Engineering Applications of Computational Fluid Mechanics*, vol. 11, no. 1, pp. 496–512, 2017.
- [7] R. Krishna, T. Sundararajan, K. Srinivasan, and S. Suresh, “Anomalies in the flow over projectile with wrap-around fins,” *Defence Science Journal*, vol. 59, no. 5, pp. 471–484, 2009.
- [8] C. H. K. Williamson, “Vortex dynamics in the cylinder wake,” *Annual Review of Fluid Mechanics*, vol. 28, pp. 477–539, 1996.
- [9] D. Bhatia, S. Kumar, and R. Kumar, “Transition delay and drag reduction using biomimetically inspired surface waves,” *Journal of Applied Fluid Mechanics*, vol. 13, no. 4, pp. 1207–1222, 2020.
- [10] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. Chapman & Hall, 1993.
- [11] T. Robertson, F. T. Wright, and R. L. Dykstra, *Order Restricted Statistical Inference*. Wiley, 1988.

- [12] A. Roshko, “On the wake and drag of bluff bodies,” *Journal of the Aeronautical Sciences*, vol. 22, no. 2, pp. 124–132, 1955.
- [13] P. W. Bearman, “Investigation of the flow behind a two-dimensional model with a blunt trailing edge and fitted with splitter plates,” *Journal of Fluid Mechanics*, vol. 21, no. 2, pp. 241–255, 1965.
- [14] M. M. Zdravković, *Flow Around Circular Cylinders, Vol. 1: Fundamentals*. Oxford University Press, 1997.
- [15] E. Robertson, V. Choudhury, S. Bhushan, and D. K. Walters, “Validation of OpenFOAM numerical methods and turbulence models for incompressible bluff body flows,” *Computers & Fluids*, vol. 123, pp. 122–145, 2015, doi:10.1016/j.compfluid.2015.09.010.