

Cost–Information Tradeoffs in Active Diagnosis of Rocket Airbrake Effectiveness Under Atmospheric Uncertainty

Nitish Pendyala

South Brunswick High School
Monmouth Junction, New Jersey, USA
nitishpen007@gmail.com

Braiden Martin

South Brunswick High School
Monmouth Junction, New Jersey, USA
braidenmartin.r@gmail.com

Francesco Venza

Northeastern University
Boston, MA, USA
venza.f@northeastern.edu

Abstract—Rocket airbrakes are commonly used to control apogee during coast, but a weaker-than-expected response does not always mean the airbrake has failed. Changes in drag, atmosphere, wind, sensing, timing, and actuator behavior can produce similar effects. In this study, we use an OpenRocket-based one-dimensional coast model to determine when reduced airbrake effectiveness can be separated from these uncertainties. A predefined set of diagnostic pulses is tested using measurements available during flight, and each pulse is evaluated by the information it provides and the modeled apogee it costs. The method was evaluated on a frozen 2025 ERA5 holdout and a separately defined 2026 ERA5 dataset. In 2026, the screened short pulse resolved 6/84 trials at a mean apogee cost of 0.684 m, while a longer fixed pulse resolved 12/84 at 2.047 m. An offline search of every feasible pulse resolved 15/84 trials. Of the 78 cases unresolved by the online method, 69 also remained unresolved offline, showing that the main limitation was the amount of diagnostic information available under the tested conditions.

Index Terms—active fault diagnosis, rocket airbrake, atmospheric uncertainty, ERA5, IGRA

I. INTRODUCTION

Deployable airbrakes are commonly studied as a way to control rocket apogee by changing drag during coast. Previous work has examined uncertainty-aware control, drag modulation, controller design, and adaptive apogee prediction [1]–[4]. Most of these studies assume the airbrake effect is known well enough for control. In this study, we focus on what happens when a known deployment produces less deceleration than expected. The cause may be reduced airbrake effectiveness, but changes in body drag, density, vertical wind, sensor bias, timing, or actuator response can produce a similar result.

The coast phase gives only a limited amount of time and altitude for diagnosis. With the brakes closed, airbrake effectiveness leaves no direct drag signature, so the measurements may reveal little about the airbrake itself. A short deployment creates a response that can be measured, but it also lowers apogee. Active fault diagnosis (AFD) describes this trade between diagnostic information and control cost [5], [6]. Safety filters and reachability methods provide related ways to reject actions that violate operating constraints [7], [8].

Our goal is to study diagnosability rather than introduce a new controller. Each trial can end as healthy, faulty, or

unresolved. The method screens a bounded set of pulses using only information available at decision time and compares the screened short pulse with a longer fixed pulse. Afterward, an offline replay tests every feasible action in the same library to determine whether a different pulse could have resolved the case. This lets us examine when reduced effectiveness can be separated from the modeled uncertainty and how that changes as more apogee is spent on diagnosis. Any reported minimum cost applies only to the predefined pulse library.

II. REFERENCE SYSTEM, ATMOSPHERE, AND MODEL

A. OpenRocket reference vehicle

URTC-AB-127 is a 0.1270-m-diameter, 2.70-m reference rocket with a 0.55-m von Kármán nose, four trapezoidal fins, and an AeroTech K700W-P. Table I lists the main vehicle and uncertainty parameters, and Fig. 1(a) shows the OpenRocket model. The reference uses lumped dry mass instead of detailed internal geometry because the study is focused on coast-phase airbrake behavior. OpenRocket 24.12 predicts a 1001.410 m AGL apogee at 15.043 s. The diagnosis simulation begins at the verified burnout state: 3.593 s after launch, 319.116 m AGL, and 126.404 m/s upward velocity.

The clean-body $C_D(M)$ curve comes from the corrected OpenRocket component-analysis export [12]. A separate Python coast model started from the same burnout state predicts 1000.911 m apogee at 15.063 s. Compared with OpenRocket, apogee differs by -0.499 m and time to apogee by $+0.020$ s, with altitude and velocity RMSE of 0.230 m and 0.051 m/s through apogee and no empirical correction factor fitted.

B. Three-panel airbrake assembly

The airbrake concept has three synchronized panels driven by a central linkage. Fig. 1(b)–(c) shows the retracted and deployed CAD states. The CAD is used to define the mechanism layout and deployment convention. It does not provide the aerodynamic coefficients, hinge moment, actuator torque, or verified kinematics used in the simulation.




Fig. 1. Reference system used in the study. (a) OpenRocket URTC-AB-127 side view in the K700W-P configuration with the modeled airbrake location and CG/CP markers. (b)–(c) CAD assembly in the retracted and deployed states. The CAD shows the three-panel mechanism, while aerodynamic and actuator values are defined separately in the simulation.

TABLE I
REFERENCE VEHICLE AND AIRBRAKE MODEL PARAMETERS

Parameter	Value
Vehicle diameter / length	0.1270 m / 2.70 m
Loaded / burnout mass	14.035 / 12.743 kg
Motor	AeroTech K700W-P
OpenRocket apogee / time	1001.410 m / 15.043 s
Burnout state (t, h, v)	3.593 s, 319.116 m, 126.404 m/s
Effectiveness η	1.00, 0.75, 0.50
Body-drag / density scale	0.90–1.10 / 0.98–1.02
Projected-area scale	0.85–1.15
Vertical wind	–3 to +3 m/s

The simulation distinguishes commanded deployment u , actual normalized deployment x_a , and sensed deployment. Projected area is

$$A_{\text{proj}} = 3A_p \sin(x_a \delta_{\text{max}}) \gamma_A, \quad (1)$$

where A_{proj} is total projected area, $A_p = 0.003600 \text{ m}^2$ is one panel's planform area, $x_a \in [0, 1]$ is actual deployment, $\delta_{\text{max}} = 60^\circ$ is maximum panel angle, and $\gamma_A \in [0.85, 1.15]$ is projected-area uncertainty. The panel drag coefficient $C_{D,p} = 0.95$ and hinge/actuator limits are likewise pre-specified model assumptions.

C. Atmospheric data and coast dynamics

ERA5 reanalysis near 40.379° N , 74.5465° W produced 8,768 processed profiles for 2020–2025 [9], [10]. Timestamp-only sampling gave 96 development profiles (2020–2023), 24 validation profiles (2024), and 24 frozen holdout profiles

(2025). NOAA IGRA 2.2 station USM00072501 provided 730 processed 2025 radiosonde profiles [11]. Five selected IGRA profiles lacked density or temperature records, so the historical IGRA result is exploratory.

The prospective 2026 test used 28 ERA5 profiles from January–July (days 1 and 15, 00/12 UTC), giving 84 trials per policy. The timestamp rule, policies, analysis, and 2.5-km vertical-support requirement were fixed before scoring. The inherited pressure-level request did not reach the required altitude, so 700 hPa was added for the same timestamps before outcomes were examined. All corrected profiles passed the finite and monotonic coverage checks.

The post-burnout vehicle follows

$$\dot{h} = v, \quad \dot{v} = -g - D/m, \quad (2)$$

where h is AGL altitude, v vertical velocity, g gravitational acceleration, m vehicle mass, and D signed aerodynamic drag. Relative vertical airspeed is $V_r = v - w$, with vertical wind w . Drag is

$$D = \frac{1}{2} \rho V_r |V_r| [\beta_b C_{D,b}(M) A_{\text{ref}} + \eta C_{D,p} A_{\text{proj}}], \quad (3)$$

where $\rho = \beta_\rho \rho_{\text{atm}}$ and $\beta_\rho \in [0.98, 1.02]$ scales the profile density ρ_{atm} . Clean-body drag uses $\beta_b \in [0.90, 1.10]$, the Mach-dependent coefficient $C_{D,b}(M)$, and reference area A_{ref} . Airbrake effectiveness is $\eta \in \{1.00, 0.75, 0.50\}$, with $C_{D,p}$ and A_{proj} defined above. Weather variation comes directly from the selected profile.

III. CAUSAL ACTIVE-DIAGNOSIS METHOD

A. Sensors, actuation, and passive baseline

The simulated sensor stream runs at 50 Hz and includes altitude, vertical velocity, vertical acceleration, and normalized deployment. The Gaussian standard deviations are 1.20 m, 0.60 m/s, 0.35 m/s², and 0.015 deployment. Altitude and acceleration bias are bounded by ± 1.0 m and ± 0.20 m/s². Each packet has a measurement time and an availability time, with 0.040-s latency and ± 0.005 -s timestamp jitter, which prevents future data from entering the decision. The actuator model uses a 0.030-s delay, a 0.060-s first-order lag, a 0.667 s⁻¹ rate limit, and a 2.667 s⁻² acceleration limit.

For the passive arm, the airbrake is never deployed. When $x_a = 0$, the airbrake regressor is also zero, so the closed-brake window has no direct excitation of the effectiveness term. We use this arm as a structural identifiability check rather than as a competitive passive estimator. A full-coast no-deployment estimator was not tested, so the 0-resolution result does not show that every passive method would fail.

B. Online action, fixed comparator, and estimate

The online decision occurs 4.50 s after launch. The library includes no probe, 5, 10, 15, and 20% single pulses lasting 0.15 or 0.30 s, and one two-pulse 10% action separated by 0.15 s. This is a bounded engineering test set rather than an optimized or exhaustive library, and no continuous pulse optimization or library ablation was performed. Candidates must pass timing, altitude, positive-velocity, 18-kPa dynamic-pressure, 4-Nm hinge-moment, and 3-m paired-cost limits. They must also reach squared Mahalanobis separation 2.0, which measures predicted hypothesis separation after accounting for modeled measurement covariance. The value 2.0 came from the frozen 2024 information screen rather than a universal significance or safety threshold, and the selector then takes the lowest-cost candidate that passes all screens.

The written freeze also listed a 1.0-s remaining-coast gate that the historical selectors did not implement. A later audit of 3,321 candidates found a 9.923-s finite minimum, so adding the gate retrospectively changed no action or result.

The fixed comparator was chosen using 2024 data only by first requiring zero wrong-confident decisions and then favoring more correct resolutions and lower mean cost. This procedure selected a 10%/0.30-s pulse, which keeps the physical, timing, and cost screens but intentionally omits the online separation screen because the purpose is to test a predetermined longer excitation. The resulting 2025 comparison is retrospective, whereas the 2026 comparison is prospective.

After a pulse begins, the estimator uses only causal samples from 0.40 s before the action through 0.80 s after it ends. Those samples are fit to

$$y_i = \beta_b X_{b,i} + \eta X_{ab,i} + b_a + \epsilon_i, \quad (4)$$

where i indexes a causal sample, a_i is measured vertical acceleration, and $y_i = -a_i - g$ is drag-related deceleration. With $q_i = \frac{1}{2}\rho_i V_{r,i}^2$, $X_{b,i} = q_i C_{D,b}(M_i) A_{\text{ref}}/m$ is the clean-body term per unit β_b , and $X_{ab,i} = q_i C_{D,p} A_{\text{proj},i}/m$ is

the airbrake term per unit η . The offset b_a absorbs constant acceleration bias and ϵ_i collects residual error.

A fit needs at least eight informative deployed samples and full rank. The effectiveness estimate $\hat{\eta}$ uses $I_\eta = [(\hat{\eta} - zSE)/1.15, (\hat{\eta} + zSE)/0.85]$, where SE is its least-squares standard error and $z = 1.95996$. The area bounds widen the interval within the model. A trial is healthy when the lower bound exceeds 0.875, faulty when the upper bound is below 0.875, and unresolved when the interval crosses it. The threshold 0.875 is the midpoint between modeled healthy ($\eta = 1.00$) and moderate fault ($\eta = 0.75$). Truth is used only after the run for scoring.

Diagnostic cost is paired within each physical trial,

$$J_{\text{apo}} = h_{\text{apo}}^{\text{no probe}} - h_{\text{apo}}^{\text{diagnostic}}, \quad (5)$$

where the two h_{apo} terms are matched no-probe and diagnostic apogees. A positive J_{apo} therefore represents altitude sacrificed for diagnosis, and each paired trajectory shares the same atmosphere, fault, nuisance realization, and seed.

C. Experiment and statistical analysis

The original frozen study contains 504 physical trials: 288 for development and 72 each for 2024 validation, the 2025 ERA5 holdout, and historical IGRA. Trial seeds are generated deterministically from master seed 12720260807, the partition, profile ID, and effectiveness state. To check that the frozen setup did not drift, 13 hashes were stored for the code, configurations, atmospheric processing, data splits, drag data, and reference trajectory. All 13 matched after the final runs, and the final regression and consistency suite passed 114 tests.

Trial-level binomial rates are reported with 95% Wilson intervals. For policy comparisons, we use 50,000 fixed-seed profile-cluster bootstrap replicates and keep the three effectiveness states from each atmospheric profile together. The active-versus-fixed McNemar tests for 2024–2026 are treated as one family and Holm-corrected, while the sensitivity runs are descriptive. After each simulation, the offline replay uses truth to test every feasible action under the same nuisance realization. It therefore shows the best result available from this finite library rather than the performance of an online policy. Representative cases are also rerun at 0.020, 0.010, and 0.005 s to check numerical convergence.

IV. RESULTS

Fig. 2(a)–(b) shows that the Python coast model closely follows OpenRocket, with a 0.499 m apogee difference, 0.230 m altitude RMSE, and 0.051 m/s velocity RMSE. The active method resolved 7/72 trials in 2024 validation and 2/72 in the untouched 2025 ERA5 holdout, while the closed-brake arm resolved 0/72 in both. The fault level also affected the result. Healthy, moderate, and severe counts were 0/24, 2/24, and 5/24 in 2024 and 0/24, 1/24, and 1/24 in 2025. This shows some severity dependence even though most severe faults were still unresolved. Historical IGRA recorded 9/72 active resolutions, but 13/72 trials had non-finite apogees, so that result remains exploratory.




Fig. 2. Model and atmospheric inputs. (a) OpenRocket and Python nominal coast altitude. (b) Python-minus-OpenRocket altitude and velocity residuals. (c)–(e) Example timestamp-selected ERA5 and IGRA density, temperature, and horizontal-wind profiles, with panels (c)–(e) sharing the 0–2.5 km AGL altitude scale.

TABLE II
PROSPECTIVE 2026 DIAGNOSTIC OUTCOME COMPOSITION

Policy	Correct	Wrong	Unresolved	Mean cost (m)
Passive	0	0	84	0
Screened low-cost	6	0	78	0.684
Fixed 10%/0.30 s	12	0	72	2.047
Offline oracle	15	–	69	–

Wrong-confident outcomes were 0 for passive, screened, and fixed. The oracle is truth-scored and evaluated offline.

For the prospective 2026 set, the screened pulse resolved 6/84 trials, the longer fixed pulse resolved 12/84, and the offline replay resolved 15/84. Their profile-cluster 95% intervals were 2.38–11.90%, 7.14–21.43%, and 10.71–26.19%. Active-minus-fixed resolution was -7.14 percentage points $[-14.29, 0]$, while the cost difference was -1.362 m $[-1.389, -1.337]$. The resolution-difference interval included zero in every partition. Raw exact McNemar p values of 0.146, 0.125, and 0.109 each became $p_{\text{adj}} = 0.327$ after Holm correction. These pooled results should therefore not be treated as severity-independent performance.

Every ERA5 trial left the same 20%/0.15-s pulse as the only online-admissible action, so the minimum-cost rule never chose between simultaneous candidates. The fixed 10%/0.30-s pulse had predicted separation below 2.0 in every ERA5 trial (0.630–0.907), but it is still a useful comparator because the fixed policy intentionally omits that screen. With actuator limits included, the short pulse reached 6.0% maximum deployment versus 12.91% for the fixed pulse, which explains much of the apogee-cost gap.

The offline replay gives a useful check on the failures. Of 65, 70, and 78 online failures in 2024–2026, 53 (81.5%), 61

(87.1%), and 69 (88.5%) stayed unresolved after every feasible library action was tried. Only 12, 9, and 9, respectively, could have been rescued by another action from the same library.

The estimator usually had enough data, since all active fits were full rank with about 26 deployed samples and median condition number near 37. Even so, median interval widths were 1.57–1.60 and 90.3–97.2% crossed the 0.875 boundary. Lag-1 residual correlation stayed near zero, Durbin–Watson remained near 2, and within-simulation interval coverage was 96.4–98.6%. These checks support the simulator analysis but do not establish physical calibration.

The sensitivity runs in Fig. 3(c) moved the result in different directions. Reducing panel drag by 15% gave 11/72 correct resolutions with two healthy false alarms, while increasing it by 15% gave 5/72 with none. Wider area uncertainty gave 9/72, ± 5 m/s wind gave 6/72, and $1.5\times$ acceleration noise gave 5/72 with one false alarm. The matched half-noise case fell to 0/72, so these runs are best read as descriptive sensitivity checks rather than robustness certification.

A post-hoc finite-record screen retained 23/24 IGRA profiles and gave 10/69 active resolutions, no wrong decisions, and a 0.6825 m mean cost without replacing the frozen 9/72 result. Numerical convergence was also stable, since changing the integration step from 0.010 to 0.005 s changed apogee by at most 0.085 m, cost by 0.0023 m, and $\hat{\eta}$ by 0.00794 without changing any classification.

V. DISCUSSION

The main result is the tradeoff between altitude spent on diagnosis and the information gained from it. The longer fixed pulse cost about three times as much modeled apogee and resolved more observed cases, but the resolution difference




Fig. 3. Strengthened results. (a) Correct-resolution rates with profile-cluster 95% intervals. (b) Mean paired apogee cost. (c) Descriptive ERA5-2024 sensitivity, where orange labels show wrong-confident counts. (d) Fraction of active intervals crossing 0.875, with labels giving median width.

was not clear after profile clustering and Holm correction. Every ERA5 trial also had only one admissible online action, so adaptive choice among several pulses was not tested. The closed-brake arm is an identifiability check, and the pulse library was not optimized or ablated.

The offline replay helps explain why most cases remained unresolved. In 2026, 69/78 online failures were still unresolved after every feasible pulse in the library was tried, so a different pulse from the same set would not have fixed most failures. Under the tested sensor model, uncertainty ranges, and observation window, the measurements often did not separate the competing explanations well enough. Since the replay uses truth after the run and only the tested actions, it represents the best result from this library rather than a physical limit.

For a rocketry team, the method is not yet a standalone flight fault classifier. The longer fixed pulse spent about 2.05 m of modeled apogee and resolved only 12/84 prospective cases. Future work should test a competitive passive estimator, different pulse libraries, richer sensing or excitation, and CFD, bench, wind-tunnel, or flight data. The current study remains simulation-specific because the airbrake aerodynamics and actuator behavior are provisional.

VI. CONCLUSION

In this paper, an OpenRocket-based coast model was used to study whether a short airbrake deployment can help separate reduced effectiveness from other sources of uncertainty. The frozen method resolved 2/72 trials in the untouched 2025 ERA5 holdout at 0.679 m mean modeled apogee cost. In prospective 2026, the screened pulse resolved 6/84 trials at 0.684 m, the longer fixed pulse resolved 12/84 at 2.047 m, and the offline replay found 15/84 resolvable cases. Only one online action was admissible, and the Holm-corrected active-versus-fixed comparisons did not show a clear resolution advantage.

Of 78 prospective cases left unresolved online, 69 also remained unresolved in the offline replay. The estimator had enough samples and full-rank fits, but its effectiveness intervals were usually too wide for a confident decision. The results therefore show a modeled cost–information tradeoff under the tested conditions rather than flight readiness or hardware validation.

REFERENCES

- [1] T. Lew, F. Lyck, and G. Müller, “Chance-Constrained Optimal Altitude Control of a Rocket,” in *Proc. 8th European Conf. Aeronautics and Space Sciences (EUCASS)*, 2019, doi: 10.13009/EUCASS2019-388.
- [2] E. Bregin and M. K. M. Jaffar, “The Development of a Drag-Modulating Closed-loop Feedback System to Control a Rocket’s Ascent,” in *AIAA SciTech 2024 Forum*, AIAA 2024-2388, 2024, doi: 10.2514/6.2024-2388.
- [3] S. Jack, “The Evaluation of Various Controller Architectures for an Air Brake on a High-Powered Model Rocket,” undergraduate honors thesis, Univ. Maryland, College Park, MD, USA, 2026, doi: 10.13016/bepq-so9q.
- [4] J. Johnston, J. Gibbs, C. Fraser, and K. Worrall, “Real-Time Adaptive Apogee Prediction for a Suborbital Student Rocket,” in *75th Int. Astronautical Congress*, IAC-24-D2-IPB-30, paper 84285, 2024.
- [5] T. A. N. Heirung and A. Mesbah, “Input design for active fault diagnosis,” *Annu. Rev. Control*, vol. 47, pp. 35–50, 2019, doi: 10.1016/j.arcontrol.2019.03.002.
- [6] I. Punčochář and J. Škach, “A Survey of Active Fault Diagnosis Methods,” *IFAC-PapersOnLine*, vol. 51, no. 24, pp. 1091–1098, 2018, doi: 10.1016/j.ifacol.2018.09.726.
- [7] K. P. Wabersich and M. N. Zeilinger, “A predictive safety filter for learning-based control of constrained nonlinear dynamical systems,” *Automatica*, vol. 129, Art. no. 109597, 2021, doi: 10.1016/j.automatica.2021.109597.
- [8] J. Borquez, K. Chakraborty, H. Wang, and S. Bansal, “On Safety and Liveness Filtering Using Hamilton–Jacobi Reachability Analysis,” *IEEE Trans. Robot.*, vol. 40, pp. 4235–4251, 2024, doi: 10.1109/TRO.2024.3454470.
- [9] H. Hersbach *et al.*, “The ERA5 global reanalysis,” *Q. J. R. Meteorol. Soc.*, vol. 146, pp. 1999–2049, 2020, doi: 10.1002/qj.3803.
- [10] Copernicus Climate Change Service, “ERA5 hourly data on pressure levels from 1940 to present,” Climate Data Store, doi: 10.24381/cds.bd0915c6, accessed Aug. 7, 2026.
- [11] I. Durre, X. Yin, R. S. Vose, S. Applequist, S. Arnfield, B. Korzeniewski, and B. Hundermark, “Integrated Global Radiosonde Archive (IGRA), Version 2,” NOAA NCEI, 2016, doi: 10.7289/V5X63K0Q, accessed Aug. 7, 2026.
- [12] OpenRocket Project, “OpenRocket Documentation,” accessed Aug. 7, 2026. [Online]. Available: <https://openrocket.readthedocs.io/>