

LANDING A THRUST VECTOR CONTROLLED ROCKET AFTER MAIN CHUTE DEPLOYMENT

A novel parachute-to-propulsive recovery architecture for propellant-efficient sounding rocket reusability.
Research Project AENG30017

Ethan Sheehan
2400720, fc23871

Supervisor: Prof. Tom Richardson

Department of Aerospace Engineering, University of Bristol,
Queen’s Building, University Walk, Bristol, BS8 1TR, UK

Generative AI Declaration Statement:

During the preparation of this coursework the author used Google DeepMind’s Antigravity IDE for in-line code completions and commenting, Google Gemini 3.1 Pro for MATLAB error explanations and aid with text condensation, and Anthropic’s Claude Opus 4.7 for LaTeX formatting, and text editing assistance. After using these tools/services, the author reviewed and edited the content as needed and takes full responsibility for the content submitted for assessment.

Intellectual Property and Data Management Statement:

Flight IMU data used for plant validation was provided by BristolSEDS under internal team agreement for academic use, and the Waxwing reference airframe parameters were shared by HyPower Bristol under the same basis. Third-party datasets (NASA Langley LaSRS++ turbulence acceptance criteria, MIL-F-8785C specifications) are openly available and cited accordingly.

Report word count: 7138

ARTICLE INFO

Keywords:

- Thrust Vector Control
- Hybrid Parachute-Propulsive Recovery
- Linear Quadratic Regulator
- Apparent Mass Dynamics
- Dryden Turbulence Model
- Coupled System Dynamics
- Monte Carlo Analysis
- Sounding Rocket Reusability
- Powered Descent Guidance
- Parachute Recovery System
- Gain Scheduling
- Optimal Control

ABSTRACT

This research investigates the feasibility of a novel hybrid parachute-to-propulsive recovery architecture in which an active Thrust Vector Control (TVC) system arrests wind-driven pendulum oscillations and delivers a soft vertical touchdown after main chute deployment. A 3-Degree-of-Freedom simulation of a 78kg high-fineness-ratio sounding rocket was developed in MATLAB and Simulink, incorporating coupled mass-inertia dynamics, apparent-mass and Munk-moment effects from potential-flow parachute theory, time-varying propellant depletion, and a NASA-verified Dryden turbulence model with MIL-F-8785C discrete gust and logarithmic shear components. The plant was validated against flight data from a comparably dimensioned airframe, reproducing the characteristic pendulum energy band with statistically equivalent stochastic structure. A phased Linear Quadratic Regulator with dynamic lead compensation, gain-scheduled damping, stabilization, and flare modes, and a hard-coded thrust override for off-nominal recovery, was armed at 200 m altitude. Across 300 Monte Carlo runs under realistic turbulence, the controller achieved a $70 \pm 5.2\%$ success rate against strict terminal constraints ($|v_z| < 3$ m/s, $|\theta| < 5^\circ$). Mean propellant consumption of 12.6 kg represents approximately one-third of the mass required for pure powered-descent guidance in equivalent conditions. The results demonstrate that integrating optimal control with high-fidelity coupled-body and atmospheric modelling provides a propellant-efficient, experimentally accessible recovery architecture for sounding rockets, and establishes a framework in which controller failure reverts to a survivable parachute landing rather than vehicle loss.

1. Introduction

The global transition toward rapid rocket reusability has necessitated a fundamental pivot from passive recovery to precise, autonomous vertical landing architectures, with launcher reusability now widely regarded as the most effective mechanism for reducing the cost of access to space [1][2]. Parachute-based systems remain the historical standard due to their mechanical simplicity and mass efficiency, and continue to see active development for small launchers such as PLD Space’s MIURA 1 and Rocket Lab’s Electron first stage [3][4]. However, they are increasingly viewed as a bottleneck for high-cadence operations [5]. Although aerodynamic drag reliably delivers a survivable terminal velocity, a canopy-borne vehicle lacks the control authority required to regulate descent attitude or navigate to a precise geospatial target [6]. The operational consequence is well documented: SpaceX’s abandonment

of parachute recovery drove refurbishment cycles down from months to days [7], and even PLD Space’s MIURA 1 designates its pyrotechnic initiators, reefing cutters, and riser extensions as single-use despite a three-flight reuse target for the airframe [3]. Compounding this, a slender body suspended under a canopy behaves as a forced pendulum subject to stochastic atmospheric perturbations, producing coupled horizontal and rotational velocities at touchdown that make stable vertical ground contact physically unachievable [8][9]. Passive recovery therefore represents a technical compromise that prioritises hardware survival over the operational precision demanded by true rapid reusability.

To address these limitations, the leading launch providers have pivoted to Powered Descent Guidance (PDG), executed as highly dynamic “suicide burn” or “hooverslam” manoeuvres that exploit Thrust Vector Con-

trol (TVC) for high-authority regulation of both attitude and trajectory [2][10]. While PDG delivers pinpoint recovery, it operates within narrow propellant margins and imposes a substantial payload penalty: reserved landing propellant and landing-hardware structural mass directly erode payload fraction relative to expendable configurations of equivalent gross lift-off weight [2][10][11][12][13]. This trade-off becomes particularly acute at the small-vehicle end of the market, where absolute propellant budgets are tight and every kilogram reserved for landing is a kilogram unavailable for payload which is precisely the regime in which high-power sounding rockets and emerging microlaunchers operate.

This research proposes a hybrid architecture in which the parachute absorbs the energy-intensive bulk deceleration phase, preserving the propellant mass fraction, while the TVC system engages only in the terminal phase to neutralise wind-induced pendulum oscillations and deliver a soft vertical touchdown. Beyond propellant efficiency, the architecture functions as an inherently low-risk testbed for experimental flight software: a TVC malfunction in conventional PDG typically terminates in catastrophic impact, whereas, under the proposed scheme, the canopy guarantees a survivable terminal velocity even in the event of complete propulsive failure. The existing literature has largely developed these two recovery modes in isolation. PDG studies (Botelho et al. [2], Kamath [10]) assume quasi-rigid descent and neglects forced canopy oscillations. Parachute-dynamics work (Dobrokhodov et al. [14], Guglieri [15]) rigorously characterises pendulum and apparent-mass behaviour but treats the payload as unpowered. Small-scale TVC work (dos Santos and Oliveira [16], Linsen et al. [17]) targets ascent-phase stabilisation and reverts to passive recovery for descent. Parachute-to-propulsive recovery has operational heritage in the capsule domain with the Soyuz retrorockets since 1967 [18] and Viking-class throttleable engines since 1976 [19] and was proposed for pendulum-damping the Orion capsule [20] and SpaceX DragonFly capsule [21]. In every case the airframe is aerodynamically blunt, the parachute is jettisoned or uncoupled before propulsive firing, or no closed-loop control of canopy-pendulum dynamics is attempted. No existing framework unifies canopy-suspended slender-body dynamics with a closed-loop thrust-vectorised landing burn. This motivates the central research question of this work: can a phased optimal controller, driving an active TVC system on a slender high-fineness-ratio sounding rocket, autonomously arrest wind-driven pendulum oscillations after main parachute deployment and deliver a survivable vertical touchdown, while consuming a small fraction of the propellant required for equivalent powered-descent guidance?

Accurately characterising this transition requires the integration of stochastic atmospheric modelling with high-fidelity coupled dynamics. A rocket suspended under a canopy is governed by apparent-mass effects, where the inertia of the air entrained by the vehicle, and the Munk moment, a destabilising torque that acts to orient slender bodies perpendicular to the flow [14][22]. This work unifies parachute aerodynamics [15], propulsive-landing guidance [23], and optimal control [24][25] inside a single 3DOF MATLAB/Simulink environment, that resolves propellant depletion, time-varying CoG and inertia, and

apparent-mass coupling, excited by NASA-verified Dryden turbulence [26][27], and validated against 6DOF IMU telemetry from a comparable sounding rocket [28]. Within this framework, a phased Linear Quadratic Regulator (LQR) optimises the recovery manoeuvre, balancing vertical alignment against the saturation limits of the TVC actuators. The control architecture is designed to satisfy the terminal constraints required for hardware survival: a vertical touchdown velocity below 3 m/s and a pitch attitude within $\pm 5^\circ$ at ground contact [2]. The Waxwing airframe [29] serves as the reference case because its 20:1 fineness ratio, 3.49m toroidal canopy, and 65–78kg mass class place it squarely within the regime in which PDG payload penalties are prohibitive but parachute-only touchdown dynamics exceed survivable limits for liquid-propulsion hardware.

1.1. Aims and Objectives

This research investigates the feasibility of a hybrid canopy-to-propulsive recovery architecture for high-fineness-ratio sounding rockets, and aims to establish, through a high-fidelity, flight-data-validated 3DOF simulation, that an active TVC system can autonomously arrest a wind-driven pendulum descent and deliver a soft vertical landing under stochastic atmospheric loading, thereby offering a new perspective on the reusability–payload trade-off for small-scale vehicles.

To achieve this, the project satisfies the following objectives, utilizing a combination of flight-test data, industry-standard atmospheric models, and control theory:

- Correlate the simulated coupled-dynamics plant, including apparent-mass and Munk-moment effects, against 6DOF IMU flight data from a comparably dimensioned airframe, targeting dominant-peak agreement within the 0.1–1Hz pendulum band and matched PSD roll-off.
- Implement a Dryden turbulence model with MIL-F-8785C discrete-gust and logarithmic-shear components, and verify against NASA Langley LaSRS++ acceptance criteria (zero-mean outputs within ± 0.003 , near-unity variance ratios, PSD reconstruction of the ideal spectra).
- Design and tune a phased LQR that neutralizes wind-induced oscillations, and quantify its robustness through a Monte Carlo campaign ($N \geq 300$) against strict terminal constraints: $|v_z| < 3$ m/s, $|\theta| < 5^\circ$ in $\geq 70\%$ of runs.

2. Methods

The terminal recovery manoeuvre is modelled as a coupled planar dynamics problem in with a high-fineness-ratio airframe, suspended from an inflated main canopy, must be actively stabilised by a gimbaled liquid engine and brought to a soft, near-vertical touchdown under stochastic atmospheric disturbances. The simulation is built in MATLAB and Simulink as a 3-Degree-of-Freedom (3DOF) environment covering pitch-plane motion, structured into five loosely-coupled subsystems interacting through a shared state vector at each 1kHz integration step (fig. 1). All variable names and definitions can be found in Appendix A; full derivations and verifications are found in Appendix B.

2.1. Reference Vehicle and Mission Scenario

The vehicle investigated in this study is based on the Waxwing high-power sounding rocket [29] developed by HyPower Bristol for the 2025 European Rocketry Chal-

lenge (fig. 2), augmented with a liquid-propulsion system capable of deep throttling and active Thrust Vector Control. Key parameters are summarised in Table 1.



Figure 2: Waxwing reference airframe. The 3.97m airframe retains the 20:1 fineness ratio and slender-body aerodynamic sensitivities characteristic of high-power sounding rockets. The base configuration was augmented with a liquid-propulsion module providing deep throttling authority and an active TVC mount, enabling the autonomous vertical landing manoeuvre investigated in this study.

Table 1: Reference Vehicle Parameters

Parameter	Symbol	Value	Unit
Total Length	L	3.97	m
Airframe Diameter	d	0.2	m
Maximum Thrust	T_{max}	5000	N
Specific Impulse	I_{sp}	180	s
Max Gimbal Angle	δ_{max}	10	deg
Chute Diameter	d_{chute}	3.49	m

The mission profile (fig. 3) begins at $h = 400\text{m}$ under the fully-inflated toroidal canopy with an initial pendulum displacement of $\theta_0 = 10^\circ$ representing observed post-deployment transients [6]. The LQR arms at 200 m and transitions to terminal flare at 25m; successful touchdown requires $v_z < 3$ and $|\theta| < 5^\circ$ $h = 0$ [30][2].




Figure 3: Mission concept of operations. Phase 1 — Passive canopy-borne descent from the 400 m simulation start at $v_z \approx 8 \text{ m/s}$ terminal velocity. The 400 m headroom allows atmospheric disturbance filters (§2.4) to reach steady state before control activates. Phase 2 — LQR activation at 200 m, chosen to provide a ≈ 25 second time-to-impact budget, sufficient to arrest inherited angular momentum within the linear-authority envelope (§3.1.4.4). Phase 3 — Terminal flare at 25 m with target $v_z = -1.8 \text{ m/s}$ to bleed off residual vertical velocity. Phase 4 — Touchdown with $v_z < 3 \text{ m/s}$, $|\theta| < 5^\circ$, thresholds set by landing-leg survivability and tip-over prevention.




Figure 1: Closed-loop simulation architecture. The Plant (§2.3) propagates the 3DOF state under external forces and moments, with translational and rotational dynamics coupled through the apparent mass of the entrained air. The Wind Model (§2.4) injects a three-component stochastic disturbance validated against NASA acceptance criteria. The Engine & TVC Model (§2.5) reproduces first-order lags, rate limits, and ignition transients of the Waxwing’s liquid engine and TVC mount. The Sensor Model (§2.6) corrupts the true state with noise, bias, and latency matching the on-board IMU before a complementary filter blends the result with GPS corrections. The Controller (§2.7) is a phased LQR with lead-compensated error, gain-scheduled weighting, and two safety overrides above the optimal loop. Colours denote functional role; dashed outlines and § tags mark the report subsection describing each subsystem.

2.2. Reference Frames and State Vector

Three right-handed frames resolve the planar dynamics (fig. 4): a North-East-Down inertial frame $\mathcal{I}$, a body-fixed frame $\mathcal{B}$ at the instantaneous CoG, and a propulsion frame $\mathcal{P}$ rotated from $\mathcal{B}$ by the commanded gimbal angle δ . Earth-rotation effects are neglected given the short landing-burn duration [31][23]. The planar rotation matrices R^{BI} , R^{IB} , and R^{PB} are given in Appendix B.1.




Figure 4: Reference frames and geometry. The inertial frame is fixed at the deployment point with z^I directed downward along local gravity; altitude and touchdown constraints reference this frame. The body-fixed frame has origin at the instantaneous CoG with z^B aligned with the longitudinal axis toward the nose, used to resolve aerodynamic forces, moments, and the coupled-mass matrix (§2.3). The propulsion frame attaches rigidly to the gimbal assembly; the thrust vector is defined along $+z^P$, and transformation back to $\mathcal{B}$ through R^{PB} yields the lateral thrust component driving controllable TVC torque. Pitch angle θ and gimbal angle δ are shown.

The complete dynamics of the system are resolved by the interaction between the state vector $\mathbf{X}$ and the control input vector $\mathbf{u}$. The 3DOF state vector is defined as:

$$\mathbf{X}_{3DOF} = [x^I, z^I, v_x^I, v_z^I, \theta^B, \omega^B]^T \quad (1)$$

This vector is mapped against the control inputs provided by the Thrust Vector Control and throttling systems:

$$\mathbf{u}_{3DOF} = [T, \delta]^T \quad (2)$$

Where T represents the engine thrust magnitude and δ denotes the gimbal deflection angle.(B.2)

2.3. Plant Equations of Motion

The dynamic behavior of the Waxwing vehicle, and the internals of the Plant block in Figure 1, are modelled as a coupled hydrodynamic body rather than an isolated rigid mass. Because the airframe descends immersed in a moving fluid while tethered to a large canopy, its translational and rotational responses are linked through the inertia of the air it entrains, known as the apparent mass effect [14]. Neglecting this coupling, as standard point-mass parachute simulations do, systematically underestimates

the low-frequency pendulum energy and the destabilising Munk moment that the controller must reject.

2.3.1 External Forces

Translational motion is driven by the sum of gimballed thrust, parachute drag, and gravity derived in Appendix B.5. Resolving in the inertial frame:

$$F_x = T \sin(\theta - \delta) - D_x, F_z = -(T \cos(\theta - \delta) + D_z) \quad (3)$$

with drag magnitude $D_{total} = \frac{1}{2} \rho v_{rel}^2 C_{d_{chute}} A_{chute}$ [22] resolved along the instantaneous relative-wind direction. The relative velocity $v_{rel} = \mathbf{v} - \mathbf{v}_{wind}$ combines the three-component wind field supplied by the Wind Model (§2.4).

2.3.2 External Moments

The rotational dynamics of the Waxwing vehicle are governed by the sum of moments acting about the Center of Gravity. To ensure a high-fidelity representation of the landing burn, the model incorporates four physically distinct pitching moments. Full derivations can be found in Appendix B.6.

The total pitching moment ΣM_y is the algebraic sum of the components derived below. This assembly represents the competition between the stabilizing parachute and control torques against the destabilizing Munk effect and atmospheric resistance:

$$\Sigma M_y = \underbrace{\tau_{chute}}_{\text{Stabilizing}} + \underbrace{\tau_T}_{\text{Control}} + \underbrace{\tau_{damping}}_{\text{Resistance}} + \underbrace{\tau_{munk}}_{\text{Destabilizing}} \quad (4)$$

Parachute restoring torque — the drag force, applied at the pivot L_p ahead of the CoG, produces a pendulum restoring moment that acts to return the airframe toward the relative-wind direction:

$$\tau_{chute} = -D_x \cdot L_p \cdot \cos(\theta) - D_z \cdot L_p \cdot \sin(\theta) \quad (5)$$

This is the primary passive stabilising moment in the system [14].

TVC control torque — nozzle deflection generates a lateral thrust component at moment arm L_T :

$$\tau_T = T \cdot \sin(\delta) \cdot L_T \quad (6)$$

This is the only actively controllable moment and what output of the Controller in Figure 1 influences.

Munk moment — the parachute-rocket system entrains different volumes of air depending on its direction of motion relative to the canopy axis. Axial translation (along the canopy's symmetry axis) accelerates a large air mass characterized by the apparent-mass coefficient a_{33} , while transverse translation displaces a smaller volume characterized by a_{11} . When the system has non-zero velocity components in both directions, these two apparent masses experience asymmetric momentum exchange, producing a pitching torque about the centre of gravity:

$$\tau_{munk} = -(a_{33} - a_{11}) v_x^B v_z^B \quad (7)$$

For a canopy-dominated system $a_{33} > a_{11}$ [14], which gives the torque a destabilising sign: any transverse velocity component generates a moment that rotates the airframe further into the relative flow rather than away from it. Although this torque shares its mathematical form with the

classical slender-body Munk moment developed for airships and missiles [32][33], the physical origin here is the canopy's apparent-mass asymmetry rather than the airframe's slender-body aerodynamics. This is the primary aerodynamic instability the controller must actively reject when disturbed.

Body-relative damping — modelling the airframe as a cylinder rotating through still air, the local velocity at station r from the CoG is $v_{local}(r) = v_x^B + \omega r$, and the damping torque follows by integrating elemental drag forces along the body:

$$Q_d = - \int_{-L_T}^{L_p} \frac{1}{2} \rho \cdot d \cdot C_{d_{body}} \cdot (v_x^B + \omega r) |v_x^B + \omega r| \cdot r \, dr \quad (8)$$

The integral is evaluated numerically using a 50-point trapezoidal rule [34]. This term provides amplitude-dependent rotational damping that enters the passive stability analysis of Section 3.1.1.

2.3.3 Apparent Mass Coefficients

The apparent-mass terms are derived from potential-flow results for the parachute-payload system following Dobrokhodov et al. [14], referenced to the canopy's hemispheroidal air slug as characteristic volume. The axial term a_{33} represents the air mass accelerated longitudinally with the canopy; a_{11} represents the air mass displaced in sideways translation; a_{44} is the rotational apparent inertia about the body pitch axis; and the coupling coefficient k_{15} captures the kinematic link between transverse translation and pitch rotation imposed by the canopy tether. Numerical values for k_{11} , k_{33} , k_{44} , and k_{15} follow the Doherr classical set scaled for the Waxwing canopy geometry and are listed in Appendix B.3.

2.3.4 Coupled Mass Matrix

To find the instantaneous accelerations, the system of equations is solved by multiplying the inverse of the Coupled Mass Matrix by the vector of summed external forces and moments:

$$\begin{bmatrix} (m + a_{11}) & 0 & K_{total} \\ 0 & (m + a_{33}) & 0 \\ K_{total} & 0 & (I_s + a_{44}) \end{bmatrix} \begin{bmatrix} \ddot{v}_x^B \\ \ddot{v}_z^B \\ \ddot{\omega}^B \end{bmatrix} = \begin{bmatrix} \Sigma F_x^B \\ \Sigma F_z^B \\ \Sigma M_y \end{bmatrix} \quad (9)$$

with off-diagonal coupling $K_{total} = -k_{15} \cdot m_a \cdot L_p$ [14]. The axial row remains decoupled because longitudinal descent does not itself induce rotation, but the (v_x, ω) block is fully coupled: a gust-induced sideways velocity immediately generates a pitching acceleration, and vice versa. The system is solved each time step by direct inversion of M_{mat} ; body-frame accelerations are then transformed back into $\mathcal{I}$ via R^{IB} for numerical integration using a fixed-step RK4 scheme at 1 kHz. See Appendix B.7 for full mass matrix and state derivative assembly.

2.4. Wind Model

The Wind Model block in Figure 1 supplies a disturbance velocity field to the Plant, represented as a superposition of three physically distinct classes added in the inertial frame and ramped from zero over the first 2 seconds of simulation to prevent spurious initial transients:

$$v_{wind} = v_{mean} + v_{gust} + v_{turb} \quad (10)$$

The mean horizontal wind follows a logarithmic boundary-layer profile [35]), ensuring that surface-relative

wind decays realistically near touchdown and avoids the “lateral slide” artefact in which a low-altitude vehicle is carried sideways by a wind that does not physically exist near the ground [36]. Discrete gusts use the MIL-F-8785C 1-cosine spatial profile [37] rather than step inputs, producing a smooth velocity pulse as the vehicle traverses a gust pocket of length L_m . Continuous turbulence uses the Dryden spectral model with altitude-dependent scale lengths from MIL-F-8785B [26][27]: a single first-order shaping filter for the longitudinal component and a cascaded two-stage filter for the vertical. Following the NASA Langley LaSRS++ methodology, a zero-order hold is used on the first filter stage and a first-order hold on subsequent stages; using FOH throughout causes 7–19% variance loss and falsifies the turbulence energy the controller is tested against. Filter forms, coefficients, and full validation against NASA Langley acceptance criteria are given in Appendix C.1.




Figure 5: Dryden turbulence verification against NASA Langley LaSRS++ acceptance criteria. Linear-scale (left) and log-scale (right) PSDs for longitudinal and vertical gust components. Statistical validation across 6×10^7 samples: longitudinal mean 0.0026, variance ratio 1.0010; vertical mean 0.0017, variance ratio 0.9943, both inside the NASA-specified ± 0.003 mean envelope and unity-variance target. Near-perfect overlap of reconstructed and ideal spectra across both scales confirms the ZOH-first implementation preserves low-frequency energy where pendulum excitation is concentrated.

2.5. Engine and TVC Model

The commanded thrust and gimbal angle (T_{cmd}, δ_{cmd}) emitted by the controller are not realised instantaneously at the nozzle. The Engine & TVC Model block in Figure 1 captures the dominant non-idealities of both subsystems. The engine model incorporates an ignition delay t_{delay} , representing valve actuation and manifold priming and, once combustion is established, an ignition pressure spike is modelled as an exponentially-decaying overshoot. Subsequent throttle changes follow a first-order lag of time constant τ_T , with the commanded rate of change subject to a hard throttle rate limit $\dot{T}_{max}$. Similarly, the gimbal is subject to a first-order servo-actuator lag with time constant τ_δ , a slew rate limit $\dot{\delta}_{max} = 45^\circ/s$ reflecting actuator torque limits, and a hard angular saturation at $\delta_{max} = 10^\circ$ representing the structural envelope of the gimbal housing. All parameters fitted directly to BristolSEDS Experi-

mental Propulsion’s Engine-1 and the integrated Waxwing feed system during the 2025 full-stack flight-qualification hotfire campaign [38]. Full equation derivation and constants justification can be found in Appendix C.2.

2.6. Sensor Model

The state feedback reaching the Controller is not the true plant state $\mathbf{X}$ but an estimate $\hat{\mathbf{X}}$ produced by the Sensor Model block in Figure 1, which reproduces the dominant non-idealities of the on-board Waxwing avionics across three subsystems. The IMU model, based on the Murata SCH16T-K01 [39], applies a 250 ms boot-up delay, datasheet-specified Gaussian noise and constant bias on each channel, and hard clipping at $\pm 300^\circ/\text{s}$ and $\pm 80\text{m}/\text{s}^2$; the gravity vector is added in the body frame so that the accelerometer reports specific force rather than kinematic acceleration. The GPS model, based on the u-blox MAX-M10S [40], runs at 10Hz with a 1.0s hot-start time-to-first-fix and position and velocity noise at the datasheet CEP and velocity-accuracy specifications, sample-and-held between update ticks. The navigation computer implements a GPS-aided complementary filter that treats the IMU as a high-frequency short-term source of truth and GPS as a low-frequency correction, with an accelerometer tilt-correction term gated to activate only when the measured specific-force magnitude lies within $0.5\text{m}/\text{s}^2$ of g to reject the term during high-thrust manoeuvres. Full parameter values, filter gains, and numerical integration detail can be found in Appendix C.3.

2.7. Control Law

The Controller block in Figure 1 is a phased Linear Quadratic Regulator (LQR) with gain-scheduled weighting matrices and a lead-compensated error signal. The LQR was selected over PID for its principled handling of multi-input/multi-output coupled dynamics [41] since the Waxwing plant has two inputs (T, δ) and six states with non-trivial cross-coupling through gravity and through the coupled mass matrix and chosen over nonlinear methods such as MPC for computational efficiency on a resource-constrained flight computer [11][12].

The controller architecture presented here, a four-gate redundantly-activated phased LQR with lead-compensated error and hammer override, is, to the author’s knowledge, a novel combination for hybrid parachute-TVC recovery. Prior work on PDG uses single-phase LQR or MPC without parachute coupling and prior work on parachute-coupled vehicles does not include active TVC (§1). The phased architecture is necessitated by the transition between regimes. The main controller (named “L”) executes a deterministic sequence of operations at each time step, shown in Figure 6. Each invocation reads the current estimated state, determines whether the vehicle is in a safe operating envelope, decides whether to activate or remain passive, selects the active control mode and reference, computes the optimal feedback command, applies physical saturation and emergency overrides, and finally emits a control input to the actuator subsystem. The remainder of this section walks through that sequence in execution order. Detailed explanation of alternative, less aggressive controller “R” layout seen in Appendix D.1.

2.7.1 Lead Compensation

At each invocation, the controller first extracts the subset of the estimated state it operates on. To compensate for

the combined actuator and sensing lag, the attitude error is then computed against a predictive pitch estimate rather than the instantaneous value [42]. The lead time is not static, it scales with the angular rate to provide more aggressive anticipation during high-frequency oscillations:

$$\tau_{lead} = 0.12 + \min(0.10, |\omega| \cdot 0.02) \quad (11)$$

$$\theta_{pred} = (\theta + \omega \cdot \tau_{lead}) \quad (12)$$

This ensures the LQR feedback commands are ahead of the physical rotation, effectively neutralising the momentum of the airframe before it reaches peak deflection. The bounds of 0.12–0.22s were selected to match the combined sensor-plus-actuator lag measured on the physical avionics stack.

2.7.2 Predictive Recovery

Before any control computation, the controller checks whether the vehicle is in an operable attitude envelope. If $|\theta| > 50^\circ$ or $|\omega| > 160^\circ/\text{s}$, the vehicle is considered uncontrollable by the LQR and the controller enters a recovery state in which thrust is forced to zero ($u = [0, 0]$). In this state, the passive parachute restoring torque (§2.3.2) is allowed to do the work of bringing the airframe back toward vertical without the destabilizing influence of misaligned thrust [43].

The recovery state exits only when one of two conditions is met. If the attitude and rate return to within the controllable envelope, or if a time-to-impact calculation indicates that a ground collision is imminent within 2.5s, control is relinquished back to the LQR, even if the vehicle has not fully stabilised, since at that point the cost of losing altitude without any active control exceeds the cost of attempting a non-ideal LQR solution. Otherwise the controller holds recovery mode and re-evaluates on the next tick.

2.7.3 Activation Check

If the vehicle is inside the operable envelope, the controller then decides whether to arm the LQR at all. The landing burn can be triggered by any of four independent logical gates, preventing a single-sensor failure from aborting the burn:

1. Flare Proximity: Immediate activation if altitude (h) falls below the critical flare threshold.
2. Existing Combustion: If the engine has already achieved minimum thrust (T_{min}) due to a manual or timed override.
3. Apex Detection: Activation at the “top” of a pendulum swing ($\omega < 3^\circ/\text{s}$) within a secondary altitude window, allowing the controller to begin damping at a point of minimum kinetic energy.
4. Hard Floor: A safety-critical altitude (h_{act_min}) that forces activation regardless of attitude or velocity.

If none of the four gates is triggered [44], the controller emits $u = 0$ and the vehicle remains in passive descent under the parachute alone. If any gate is triggered, execution proceeds to target selection.




Figure 6: Controller architecture. Colours: orange = predictive recovery, violet = activation, blue = mode selection, green = LQR + safety.

2.7.4 Mode Selection

Once armed, the controller switches between three distinct operational modes based on the magnitude of the pitch rate and the current altitude [45]. A hysteresis band around the rate threshold prevents “chatter” (rapid switching) near the transition: damping mode engages when $|\omega| > 10^\circ/\text{s}$ and releases only when $|\omega|$ falls below $3^\circ/\text{s}$.

- Flare mode: engaged at low altitudes ($h \leq h_{\text{flare}}$).

The weighting matrices are set to K_{flare} , which shifts the LQR’s priority onto vertical velocity, targeting a soft-touchdown reference $v_h^{\text{ref}} = -1.8$.

- Damping mode: engaged when the pitch rate is above the hysteresis threshold. The weighting matrices K_{damp} prioritise the neutralisation of angular momentum. Under extreme rotation ($|\omega| > 40^\circ/\text{s}$) the controller temporarily shifts its reference pitch to the instantaneous current angle, $\theta^{\text{ref}} = \theta$, to prevent the LQR from “fighting” a large momentum build-up and focus exclusively on rate reduction.
- Stabilisation mode: the default branch when neither flare nor damping conditions are met. The weighting matrices K_{act} hold the vehicle at $\theta^{\text{ref}} = 0$.

2.7.5 Feedback Law

Once the phase-specific weights and reference are fixed, the controller forms the error vector:

$$X_{\text{err}} = X_{\text{control}} - X_{\text{target}} \quad (13)$$

and applies the state-feedback law

$$\beta_u = -K X_{\text{err}} \quad (14)$$

The gain matrix K is obtained by minimising the infinite-horizon quadratic cost

$$J = \int_0^\infty (x^T Q x + u^T R u) dt \quad (15)$$

which balances state tracking error (via Q) against actuator effort (via R) [25][41]. The solution follows from the Algebraic Riccati Equation $A^T P + P A - P B R^{-1} B^T P + Q = 0$; the linearised A and B matrices about hover capture the gravity–pitch coupling and the TVC moment-arm effectiveness, and are given in Appendix D.2 and D.3 [25].

Before the command reaches the actuator, the gimbal magnitude is clipped to the physical envelope $|\delta_{\text{cmd}}| \leq 10^\circ$ (§2.5) to prevent the LQR from requesting commands the hardware cannot execute.

2.7.6 Hammer Override

The “Hammer” mode is a hard-coded safety override designed to prevent ground impact during failed damping maneuvers. The controller ignores the LQR’s optimized thrust command and forces the engine to Maximum Thrust (T_{max}) under two conditions:

1. Low-Altitude Instability: If the angular rate is high ($> 30^\circ/\text{s}$) while the vehicle is below 50 m.
2. Excessive Sink Rate: If the vertical descent velocity (v_h) exceeds -12m/s, indicating the parachute drag is insufficient to meet touchdown constraints. This hierarchy ensures that while the LQR handles the “elegant” transition to landing, the system possesses the raw authority to prioritize vehicle survival during edge-case scenarios.

3. Results and Discussion

Results are organised as a four-stage cross-validation chain (§3.1.1 analytical, §3.1.2 flight-data, §3.1.3 nominal trajectory, §3.1.4 Monte Carlo), each stage tested against an independent reference rather than internal consistency. Inferential statistics anchor the principal claims: bootstrap 95% CIs on success rate ($N=1000$), χ^2 independence test on L-vs-R outcomes, Welch’s t-test with Cohen’s d on

propellant consumption, and logistic regression on trigger-state failure predictors. The discussion (§3.2) develops the implications across eight subsections, working from physical origin of plant behaviour (§3.2.1) through atmospheric model scope (§3.2.2), controller performance against published benchmarks (§3.2.3), the L-R trade-off (§3.2.4), propellant economics (§3.2.5), failure-mechanism diagnosis (§3.2.6), testbed accessibility (§3.2.7), and limitations (§3.2.8). All significance claims use $\alpha = 0.05$.

3.1. Results

3.1.1 Passive Plant Verification

The uncontrolled, undisturbed plant was characterised first to establish that the coupled mass matrix, apparent-mass terms, and body-relative damping produce physically consistent dynamics before any disturbance or control input is introduced.




Figure 7: Passive aerodynamic damping response. Peak amplitude decays from an initial 19° overshoot to below 2° within 45s. Modal parameters extracted by peak detection and logarithmic decrement. The envelope decay rate (and the derived time-to-half-amplitude 11.08s) carry 95% bootstrap confidence [46] intervals from $N=1000$ resamples of the peak sequence, while the natural frequency $\omega_n = 4.264\text{rad/s}$ carries a parametric 95% CI on the mean peak-to-peak period (Appendix E.1)

Released from a 10° pitch displacement, the plant exhibits the under-damped ring-out of Figure 7, with the phase portrait of Figure 8 confirming monotonic convergence to the origin. The amplitude-banded damping ratio ($\zeta = 0.0134$ at high amplitude, rising to $\zeta = 0.0145$ at low amplitude) is consistent with the quadratic velocity dependence of profile drag in the body-relative damping integral (Eq. 8), and matches the qualitative pattern Aubuchon and Owens [47] report for Orion drogue-phase oscillations. The implied simple-pendulum length corresponding to the measured 0.679Hz natural frequency is 0.54m , much shorter than the geometric 2.19m pivot-to-CoG length, confirming that apparent-mass and aerodynamic-stiffness contributions dominate over gravitational restoring torque. The smooth phase-space contraction and constant-frequency funnel together rule out altitude-coupled forcing artefacts and confirm numerical consistency of the state-derivative assembly.

3.1.2 Flight Data Validation

With the passive plant verified, the full multi-layered atmospheric model (shear, discrete gusts, and Dryden turbulence) was activated and the simulated response compared against 6DOF IMU tilt data from a geometrically-




Figure 8: Phase portrait of angular velocity ω against pitch angle θ across the 47s simulated descent, colour-mapped by altitude. The trajectory traces a monotonically contracting spiral from the 10° release to the origin with no limit cycles or discontinuities, confirming numerical stability of the coupled-mass matrix solution and the absence of spurious energy injection from the frame-transformation or state-derivative assembly.

comparable sounding rocket flight [28]. Because the reference flight vehicle is not identical to the simulated Waxwing configuration, and because the 6DOF flight tilt aggregates pitch, roll, and yaw contributions that the planar simulation cannot represent in isolation, the comparison is framed around qualitative and statistical equivalence rather than deterministic agreement.




Figure 9: FFT comparison of simulated planar pitch against 6DOF flight tilt magnitude across the 47 s validation window. Both signals concentrate energy below 1 Hz, within the pendulum band expected for slender-body parachute recovery. Dominant peaks: simulation 0.610Hz , flight 0.366Hz ; offset consistent with the geometric differences between the two platforms. RMS amplitudes agree to within 39% (simulation 15.99° , flight 11.49°) with a variance ratio of 1.94, a conservative asymmetry for controller validation since a controller tuned against a harsher disturbance environment carries margin when deployed against a milder real-vehicle response.

Both signals exhibit the characteristic low-frequency energy concentration of canopy-borne descent, with pendulum energy dominating below 1 Hz in both cases (fig. 9). The simulated pitch peaks at 0.610Hz and the flight tilt at 0.366Hz , both within the $0.1\text{--}1\text{Hz}$ band reported for

slender-body parachute recovery systems [6][15]; the offset between them is consistent in magnitude with the geometric and configurational differences between the two platforms. In the time domain (fig. 10), the simulation reproduces the matched envelope range and the stochastic character of the flight signal, including the intermittent high-amplitude excursions characteristic of discrete gust encounters superimposed on continuous turbulence-driven oscillation. The simulation is driven by independent noise realizations and does not reproduce specific historical events, but reproduces the statistical properties that matter for controller validation: amplitude distribution, oscillation envelope, and gust-spike prevalence. Full quantitative validation metrics are summarised in Appendix E.2.




Figure 10: Time-domain overlay of simulation ($|\theta|$) against flight (2 Hz low-pass absolute tilt) across the 47s window. Matched envelope range (peak amplitudes approximately 70° flight, 50° simulation during the most energetic phase) and similar stochastic character: both exhibit intermittent high-amplitude gust excursions superimposed on continuous turbulence-driven oscillation. The simulation’s 0.31s autocorrelation time is approximately one-third of the flight’s 0.98s, consistent with the 3D-to-1D aggregation effect discussed in §3.2.2

3.1.3 Nominal Closed-Loop Performance

A single best-performing successful run (Run #59 of the 300-run Monte Carlo batch) is presented in Figure 11 to expose the phased controller’s behaviour in detail before its statistical performance is evaluated. This run was selected as the closest individual realisation to the zero-error touchdown target, with terminal state $\theta_f = -0.63^\circ$ and $v_{z,f} = 0.96\text{m/s}$; later discussion will contextualize it against the full outcome distribution.

The descent decomposes cleanly into two dominant regimes demarcated by the phase shading. At activation ($t \approx 20\text{ s}$, $h = 200\text{ m}$), the vehicle inherits a parachute-borne pendulum in mid-swing: pitch crosses $+15^\circ$ and angular rate peaks above $+55^\circ/\text{s}$ within the first second of control. The gimbal command saturates near both $\pm 10^\circ$ limits almost continuously through the first ten seconds of control, reflecting the controller’s aggressive damping-mode gains fighting the momentum inherited from the preceding 20 seconds of uncontrolled pendulum oscillation. Actual gimbal deflection visibly lags the commanded signal, confirming the first-order actuator model of Section 2.5 is exercising realistic authority limits rather than providing idealised instantaneous response.

Flare transition occurs at $t \approx 68\text{ s}$ when the vehicle crosses the 10 m altitude threshold. The controller’s phase switch is behaviourally distinct: a sharp thrust spike to 1 kN (visible as the ignition-spike transient in panel 6) accompanies a step change in targeted descent velocity from -8 m/s to -1.8 m/s . Descent velocity halves within two seconds of flare onset and continues to bleed off smoothly toward the soft-touchdown target. Gimbal activity increases noticeably as the controller makes fine corrective pulses to maintain near-vertical attitude at reduced thrust margin. Terminal state is $\theta = -0.63^\circ$, $v_z = 0.96\text{ m/s}$ at $h = 0$, well inside the $\pm 5^\circ$ pitch and 3 m/s velocity envelopes.




Figure 11: Best-performing closed-loop landing (Monte Carlo Run 59).

3.1.4 Monte Carlo Performance

The stochastic nature of the Dryden model makes deterministic performance claims meaningless; robustness must be evaluated in distribution. Three hundred independent runs were executed for the primary (“L”) controller, with 100-run campaigns for the conservative (“R”) variant. Convergence of the success-rate estimator with run count is documented in Appendix E.4: the 95% confidence interval on the L-controller success rate is $70 \pm 5.2\%$ [48].

3.1.4.1 Ensemble Behavior

The trajectory visualization in Figure 12 yields 209 successes (70%), 22 out-of-spec (OOS) landings (7%), and 69 spin/tumble events (23%) from the 300-run L-controller batch. The success envelope widens to a $\pm 1\sigma$ half-width peak of approximately 180 m near 150 m altitude, then narrows sharply as the controller converges toward touchdown with a 55% reduction in distribution spread between 150 m and 50 m. This convergence-in-distribution is a

stronger robustness guarantee than any single-trajectory demonstration could provide, confirming that the phased LQR architecture generalises beyond the nominal realisation of §3.1.3. Mean horizontal drift during the LQR phase is 42m, with the overall success population reaching a median touchdown displacement of approximately 450 m from release, a drift attributable to steady-state wind shear integrated over descent duration rather than controller deficiency. Failure modes are spatially distinct. OOS traces (orange) exit the success corridor predominantly below 100m altitude, representing runs where late-arriving turbulent gusts or residual pitch oscillations prevent full convergence before flare initiation; these vehicles reach touchdown intact and survivable, merely exceeding the strict $\pm 5^\circ/3$ m/s terminal envelope. Spin traces (red) originate predominantly above the 200m activation threshold, confirming that catastrophic rotations accumulate during the uncontrolled parachute-borne phase rather than under active LQR authority, a distinction developed quantitatively in Section 3.1.4.4.




Figure 12: Trajectory spaghetti for the 300-run L-controller Monte Carlo batch, classified as success (green, $n=209$), out-of-spec (orange, $n=22$), and spin/tumble (red, $n=69$). Success median and $\pm 1\sigma$ envelope overlaid.

3.1.4.2 Actuator Effort and Control Margin

Every successful run reached the $\pm 10^\circ$ gimbal limit at some point during the active phase, and every run's peak gimbal rate reached the $45^\circ/\text{s}$ actuator slew limit. This initially appears to suggest chronic saturation, but the interpretation is subtler. The gimbal limit is reached momentarily at peak disturbance response, the commanded gimbal routinely exceeds the physical envelope, with the actuator saturating at $\pm 10^\circ$ and slew rate limiting at $45^\circ/\text{s}$ during these transient overshoots, but the time-averaged deflection remains well inside the linear envelope. This pattern, universal saturation at peaks combined with moderate average effort, is the signature of a controller operating at the correct authority level: the LQR is using the full actuator envelope when required to reject large disturbances, but not operating chronically against the limit

in steady state. An undersized actuator would manifest as sustained rather than transient saturation, and failure modes would cluster at runs with high peak effort; the failure analysis in §3.1.4.4 shows that spin events instead cluster by trigger-state angular rate rather than actuator limit proximity, confirming that when the controller fails, it fails because the disturbance exceeds the physical time-to-arrest budget, not because the gimbal hardware is inadequate.

3.1.4.3 Propellant Economics and Flare Altitude

Mean propellant consumed across the 300-run L-controller batch (25 m flare) is 12.63kg ($\sigma = 4.35\text{kg}$), corresponding to 16.% of the 78.27kg wet mass. Reducing the flare altitude to 10 m (100-run batch) decreases this to 11.24kg ($\sigma = 3.90\text{kg}$; 14.4% mass fraction), an 11.0% mean reduction of 1.39kg per mission. Welch's t-test confirms statistical significance ($t = 2.99$, $p = 0.003$), with a modest effect size (Cohen's $d = 0.34$) and 95% CI on the difference of $[0.47, 2.30]\text{kg}$. Per-run distribution spread ($\sigma \approx 4\text{kg}$) is substantial relative to the mean shift, reflecting atmosphere-driven variability rather than controller inconsistency. The reduction is achieved without measurable change in classification outcomes (Appendix E.3).




Figure 13: Propellant consumption comparison between 25m flare ($N=300$) and 10m flare ($N=100$) Monte Carlo batches. Error bars show 95% confidence interval on the mean.

3.1.4.4 Predictor Analysis of Failure Outcomes

To identify the physical drivers of mission failure, a logistic regression of failure probability on two candidate predictors was fitted to the L-controller Monte Carlo batch: the vehicle's attitude ($|\theta_{\text{trigger}}|$) and angular rate ($|\omega_{\text{trigger}}|$) at the moment of LQR activation [49]. Trigger-state variables were selected in preference to global peak values to isolate the controller's inherited initial condition rather than post-activation dynamics. The resulting failure probability surface is shown in Figure 14.

The regression identifies angular rate as the dominant predictor and attitude as statistically insignificant once angular rate is controlled. The standardised coefficients show angular rate dominates attitude by a factor of 2.7 in predictive contribution. The model's modest AUC of 0.67

indicates trigger-state angular rate captures the inherited-momentum component of failure, while a substantial residual fraction of failure variance arises from post-activation atmospheric disturbance not anticipated by the activation snapshot. The $P = 0.5$ authority boundary falls at $|\omega_{\text{trigger}}| \approx 119^\circ/\text{s}$ at $|\theta_{\text{trigger}}| = 0^\circ$, sloping gently downward to approximately $68^\circ/\text{s}$ at the 30° attitude limit. Full breakdown in Appendix E.5.




Figure 14: Logistic regression of failure probability on trigger-state attitude $|\theta_{\text{trigger}}|$ and angular rate $|\omega_{\text{trigger}}|$ for the 300-run L-controller batch. White contour lines mark $P(\text{fail}) = 0.25$ and $P(\text{fail}) = 0.5$.

3.1.4.5 Controller Architecture Comparison

To assess the sensitivity of landing outcomes to controller design, the nominal L-controller batch ($N = 300$) was compared against an alternative R-controller batch ($N = 100$) executed under identical atmospheric, vehicle, and mission parameters. Both controllers were classified using a consistent rule: spin ($|\theta|_{\text{peak}} > 90^\circ$ with non-compliant terminal state), success ($|\theta_{\text{term}}| \leq 5^\circ$ and $|v_{z,\text{term}}| \leq 3 \text{ m/s}$), or out-of-spec (OOS) otherwise. The comparison is shown in Figure 15. The L controller produced 70% successes, 7% OOS failures, and 23% spin failures. The R controller produced 48% successes and 52% OOS failures with no spin events. A chi-squared test of independence confirms the outcome distributions differ between controllers ($\chi^2 = 109.2$, $\text{df} = 2$, $p < 0.001$), reflecting a redistribution of failure modes rather than a uniform change in overall reliability. The terminal-state scatters visualise this redistribution directly: L-controller OOS failures cluster tightly around the success envelope at $|v_z| \approx 2 \text{ m/s}$ with $|\theta|$ spread up to 25° , while R-controller OOS failures occupy a distinct band at $|v_z| \approx 7\text{--}9 \text{ m/s}$, indicating insufficient terminal deceleration rather than attitude drift.

3.2. Discussion

3.2.1 Plant Dynamics

The amplitude-banded damping ratio ($\zeta = 0.0134$ at high amplitude, rising to 0.0145 at low amplitude) lines up with the quadratic velocity component of the profile drag in the body-relative damping integral (Eq. 8). At larger swing amplitudes, integrated local velocities along the airframe are higher, and quadratic drag produces proportionally greater resistive torque that normalises to a slightly lower damping ratio. This amplitude ordering is weak (a 0.0011 difference between bands) but consistent across repeated simulation runs with different initial conditions (Appendix E.1), ruling out extraction noise as the sole explanation. Aubuchon and Owens [47] report similar amplitude characteristics in Orion drogue-phase oscillations, and Guglieri [15] identifies it as a characteristic signature of quadratic aerodynamic damping in parachute-payload trajectory simulations.

The phase portrait of Figure 8 verifies that these modal parameters are properties of the plant itself rather than artefacts of the release transient. Internal verification cannot, however, validate that the numerical parameter values correspond to physical reality, that requires comparison against an independent reference.

3.2.2 Atmospheric Model Validity

The flight-data comparison establishes that the simulated plant produces dynamics of the correct type, stochastic low-frequency pendulum oscillation with gust-driven high-amplitude excursions, but with quantitative offsets that merit careful interpretation before the controller performance results of Section 3.1.4 are themselves interpreted.

Two distinct mechanisms contribute to the observed simulation-flight divergence, and their separation is important to the validation argument. The first is platform difference. The flight-data vehicle differs from the simulated Waxwing configuration in its physical parameters and recovery hardware, and each platform's canopy-pendulum dynamics reflect its own geometry. The observed frequency offset is consistent in magnitude with the differences between the two vehicles: a planar simulation of the reference vehicle at its own parameters would be expected to oscillate closer to the flight-observed frequency rather than Waxwing's. The frequency offset is therefore not a modelling defect but a consequence of running the simulation on a different vehicle than the reference flight.

The second mechanism is dimensional aggregation. The flight tilt data is a 6DOF magnitude of attitude error that combines pitch, roll, and yaw departures from vertical into a single scalar quantity. The 3DOF simulation produces only a planar pitch angle. In 6DOF canopy-borne descent, roll-yaw coupling and coning motions contribute meaningfully to the measured tilt magnitude, typically smoothing and slowing the aggregate signal compared to a pure pitch-plane oscillation. The flight data's longer autocorrelation time (0.98 s versus the simulation's 0.31 s) and slower apparent envelope dynamics are consistent with this 3D-to-1D aggregation effect meaning the flight signal reflects an average over three rotational degrees of freedom while the simulation reflects only one.

Critically, neither mechanism undermines the validity of the controller. The atmospheric model is demonstrably producing stochastic disturbances of the correct statisti-




Figure 15: L-controller (N=300) and R-controller (N=100) outcome comparison. Left and centre: terminal-state scatter plots with success envelope overlaid; L-controller inset shows peak pitch distribution for spin events. Right: outcome-category distribution.

cal character (RMS amplitude within 39%, variance ratio within a factor of two), at the correct frequency band for canopy-pendulum dynamics, with the correct stochastic structure (intermittent gust excursions on continuous background). These are the properties the LQR controller must handle. The class of disturbance is created, not any specific historical flight-path. A controller validated against this simulation has been tested against realistic canopy-pendulum dynamics driven by NASA-compliant turbulence, which is the relevant fidelity claim in the next sections of this discussion.

The simulation's more energetic oscillation represents a conservative design environment (fig. 10): a controller tuned against a harsher disturbance carries margin when deployed against a milder real response, so the Monte Carlo success rates of Section 3.1.4 are, if anything, underestimates of likely performance on the physical Waxwing rocket.

3.2.3 Controller Benchmarking

The 70% success rate achieved by the L configuration sits between the two established benchmarks in the literature. SpaceX's operational Falcon 9 booster recovery campaigns report terminal success rates exceeding 98% under PDG, while drop-test campaigns of novel sub-orbital reusable vehicles typically report first-iteration success rates below 60% [9]. The result here, obtained on a high-fineness-ratio sounding rocket an order of magnitude smaller than either reference case, is therefore within the expected band for small-vehicle hybrid architectures. Importantly, the failures are not distributed uniformly: the ensemble identifies angular rate at trigger as the dominant predictor (§3.1.4.4), with attitude and turbulence secondary. This indicates that the remaining 30% failure margin is not a stochastic-disturbance-limited ceiling but a control-authority-bounded envelope or the LQR's time-to-arrest-momentum budget from the 200m trigger altitude. Decreasing the inherited angular rate at the hand-off from the parachute phase, rather than improving atmospheric robustness, is the optimization predicted to yield the most significant gains. This distinguishes the bottleneck of hybrid recovery from that of pure-PDG where terminal-phase gust rejection dominates [2][50].

3.2.4 L–R Trade-off

The statistically significant redistribution of failure modes between the L and R controllers (§3.1.4.5) is structurally a Pareto trade: the R controller eliminates the catastrophic spin failures by sacrificing 22 percentage points of overall success. In pure-PDG architectures, this trade is one-sided, an uncontrollable spin produces a high-energy impact and an OOS landing would impact at the high terminal velocity of the parachute-free vehicle and would be unrecoverable. The hybrid architecture modifies this trade significantly. Because the parachute ensures a survivable terminal velocity even in the absence of controller authority, an R-controller OOS corresponds not to a destroyed vehicle but to a hard-but-survivable parachute landing equivalent to the unmodified Waxwing baseline. The L–R decision therefore becomes not a reliability trade but a cost trade: the L path retains a small catastrophic-loss probability; the R path converts all failures into routine parachute recoveries with increased refurbishment. Which controller is preferable depends on the ratio of vehicle replacement cost to refurbishment cost, a parameter that varies by orders of magnitude between development vehicles and operational fleets.

3.2.5 Propellant Efficiency

The 11.0% propellant reduction achieved by lowering the flare altitude from 25m to 10m (§3.1.4.3) quantifies the hybrid architecture's core mass-fraction advantage over pure-PDG and shows further fuel-efficiency improvements are possible. Applying the Tsiolkovsky equation to the same vehicle executing a hypothetical PDG deceleration from the bare-airframe terminal velocity (200m/s at sea level) to a soft touchdown gives a required propellant mass of approximately 40kg (full derivation in Appendix G.1), 3.2 times the 12.6kg consumed by the L controller. This factor of three is the engineering cost of the hybrid architecture's parachute mass penalty versus the propellant-mass penalty it avoids. For a high-fineness-ratio sounding rocket with tight mass margins, the parachute represents a fixed 3.5kg penalty while PDG would require dedicating approximately 40kg of propellant, a mass fraction cost of 60% of the dry vehicle. The hybrid architecture therefore inverts the Dresia [13] framing of the PDG trade: for

small-scale vehicles, propellant is the binding constraint, not precision, and parachute-assisted controlled descent dominates this frontier wherever mass fraction is tight.

3.2.6 Linearization Validity

The failure-predictor analysis (§3.1.4.4) identifies an internally consistent boundary: the LQR loses authority approximately where its linearisation assumptions themselves break down. The A-matrix small-angle approximation $\sin \theta \approx \theta$ retains $< 5\%$ error at $|\theta| = 15^\circ$ but exceeds 20% by 45° . Similarly, rigid-body angular momentum $J\omega$ becomes a poor approximation of the coupled-mass moment at high ω due to apparent-mass nonlinearities (§2.3.3). The logistic regression $P = 0.5$ authority boundary falls at $|\omega_{\text{trigger}}| \approx 119^\circ/\text{s}$ with shallow θ -dependence, matching the regime where both approximations degrade. The controller is thus operating at a boundary set by the fidelity of its own internal model rather than by actuator limits, a self-consistent failure mechanism rather than a hardware one. This suggests a Model-Predictive Control or nonlinear architecture which would retain accuracy beyond the small-angle regime, would be the highest-leverage controller upgrade with an expected benefit directly in the regime where the current LQR fails.

The same self-consistency argument applies to the navigation stack: filtered state-estimate error from the Section 2.6 sensor model ($\approx \pm 0.4^\circ$ attitude, $\pm 2^\circ/\text{s}$ rate at GPS update) sits roughly two orders of magnitude below the trigger-state authority boundary of Section 3.1.4.4, ruling out sensor noise as a contributor to the observed failure distribution (Appendix C.3).

3.2.7 Testbed Accessibility

A central motivation for this research was lowering the barrier to entry for precision-recovery experimentation. The results substantiate this claim quantitatively. In pure-PDG, a control malfunction produces hardware loss so the only controlled variable is the probability of failure. In the demonstrated hybrid architecture, a control malfunction produces an unplanned parachute landing at the baseline Waxwing terminal velocity which is a survivable outcome already validated on every recovery of this vehicle class. The maximum cost of an experimental control law is therefore a refurbishment cycle, not a vehicle replacement, irrespective of controller performance. This asymmetry is the substantive accessibility gain: small research groups, where vehicle capital is scarce and development iteration rate is the binding constraint, can now trial aggressive propulsive control strategies without the catastrophic-failure risk profile that has historically restricted such work. The approach does not replace PDG for operational systems where payload mass fraction and landing-pad accuracy dominates, but it does open the design space of experimental precision-recovery control to vehicles and institutions that were previously excluded from the problem.

3.2.8 Limitations, and Scope

The principal strength of this study is its methodological rigor: the plant dynamics were validated against independent 6DOF flight IMU data rather than self-referentially (§3.2.1), the atmospheric model reproduces NASA-verified spectral statistics to within a 2% variance ratio (§2.4), and the failure classification treats the ensemble as a statistical sample rather than a worst-case envelope, enabling the

chi-squared and logistic inferences that anchor the central claims. The failure-mode analysis is mechanistic rather than descriptive, identifying both the dominant predictor (angular rate at trigger, §3.1.4.4) and its underlying physical origin (linearisation breakdown, §3.2.6), which together yield a concrete design path forward.

Three limitations bound the inference. First, the 3DOF planar formulation cannot capture 6DOF coupled modes such as roll-yaw coupling, coning oscillations, and asymmetric canopy instabilities so conclusions therefore apply rigorously only to pitch-plane dynamics under symmetric single-chute descent (§3.2.2). Second, the static-gravity, flat-Earth assumptions preclude direct extrapolation to orbital-scale recovery where coordinate rotation and altitude-variable gravity become first-order effects. Third, the controller was evaluated against a harsher simulated disturbance environment than the reference flight exhibited (§3.2.2), which provides conservative margin but does not certify performance against extreme events outside the Dryden model's design envelope, such as wind-shear events at convective cell boundaries. These boundaries define the regime within which the present results constitute a valid proof-of-concept and, beyond them, the framework is a starting point for higher-fidelity study rather than a definitive answer.

4. Conclusion

This study set out to test whether an active Thrust Vector Control system could autonomously arrest a wind-driven pendulum descent and deliver a soft vertical touchdown for a slender-body sounding rocket, and the evidence assembled here supports that it can, within a bounded envelope. The phased LQR achieved a $70 \pm 5.2\%$ success rate across 300 Monte Carlo runs under NASA-verified Dryden turbulence, with terminal states consistently satisfying the $|v_z| < 3 \text{ m/s}$ and $|\theta| < 5^\circ$ survivability constraints in successful realisations. Set against this, the propellant-mass comparison quantifies the architectural claim that motivated the project: 12.6kg consumed by the hybrid landing burn against an estimated 40kg required for pure powered-descent from the same activation state, a factor-of-three reduction that inverts the Dresia mass-fraction trade for the small-vehicle regime. The flight-data-validated plant, NASA-compliant atmospheric model, and the L-R controller comparison together substantiate the core thesis that parachute-to-propulsive recovery is a viable, propellant-efficient, and experimentally accessible alternative to pure PDG for high-fineness-ratio sounding rockets.

Given another two to three months, the highest-leverage extension would be replacing the linearised LQR with a Model Predictive Controller operating on the full nonlinear coupled-mass dynamics. The failure analysis (§3.2.6) identifies linearization breakdown, the LQR losing authority precisely where its small-angle assumption degrades and not where actuators saturate or disturbances peak, as the binding constraint on success rate. Model Predictive Control, by re-linearizing around the current operating point at each horizon step rather than relying on a single hover-point linearization, addresses this failure mechanism directly rather than trimming around its edges. The expected benefit therefore lies in the regime where the present LQR fails, not where it already succeeds. In parallel, extending the simulation from 3DOF

to 6DOF would capture the roll-yaw coupling and coning modes that the planar formulation cannot represent and that partly account for the dimensional-aggregation offset observed against the 6DOF flight reference. If starting the project again, the principal change would be sequencing: the plant, atmospheric model, and controller were developed largely in series, whereas building a low-fidelity end-to-end closed loop in the first fortnight and then progressively substituting higher-fidelity subsystems would have surfaced the angular-rate-at-trigger bottleneck earlier and allowed controller architecture decisions to be informed by it, rather than discovered after tuning was complete.

Acknowledgments

I would like to thank my supervisor, Prof. Tom Richardson, for his guidance throughout this project and for supporting the idea from the outset. I am also grateful to Jack Cunningham for his expert advice on control systems and thrust vectoring, which shaped several of the key design decisions in this work.

References

- [1] Antoine Patureau de Mirand, Jean-Marc Bahu, and Eric Louaas. "Ariane Next: A Vision for a Reusable Cost Efficient European Rocket". In: (2019), 9 pages.
- [2] Afonso Botelho et al. "Design of the Landing Guidance for the Retro-Propulsive Vertical Landing of a Reusable Rocket Stage". In: *CEAS Space Journal* 14 (Feb. 2022).
- [3] Mariagrazia Bibiano, Francisco Garcia, and Raul Torres. "MIURA1: The Reusable Sounding Rocket. Recovery and Reusability Strategies". In: (2019), 14 pages.
- [4] Shane Glynn. "Reusable Electron: Analysis of Progress toward the World's First Reusable Commercial Small Rocket". In: *Small Satellite Conference*. Lakewood, CA: Rocket Lab, 2022, p. 7.
- [5] Shraddha C., Pankaj Priyadarshi, and Devendra Prakash Ghatge. "A Survey of Launch Vehicle Recovery Techniques". In: *Progress in Aerospace Sciences* 155 (May 2025), p. 101092.
- [6] Ricardo Machin and Eric Ray. "Pendulum Motion in Main Parachute Clusters". In: *23rd AIAA Aerodynamic Decelerator Systems Technology Conference*. Daytona Beach, FL: American Institute of Aeronautics and Astronautics, Apr. 2015.
- [7] Harry W. Jones. "The Recent Large Reduction in Space Launch Cost". In: *48th International Conference on Environmental Systems*. Albuquerque, New Mexico: NASA Ames Research Center, July 2018, p. 10.
- [8] Christine E. Stewart. "Spacecraft Parachute Recovery System Testing from a Failure Rate Perspective". In: *AIAA Aerodynamic Decelerator Systems (ADS) Conference*. Daytona Beach, Florida: American Institute of Aeronautics and Astronautics, Mar. 2013.
- [9] Geun-Jeong Lee, Min-Seon Jo, and Jeong-Yeol Choi. "Comparison of Second-Stage Recovery Methods for Reusable Launch Vehicles Across Vehicle Scales". In: *Aerospace* 13.1 (Jan. 2026).
- [10] Abhinav Kamath. "Robust Thrust Vector Control for Precision Rocket-Landing". PhD thesis. UC Davis, 2021.
- [11] Zhenbo Wang. "A Survey on Convex Optimization for Guidance and Control of Vehicular Systems". In: *Annual Reviews in Control* 57 (Jan. 2024), p. 100957.
- [12] Carlo Alberto Pascucci, Samir Bennani, and Alberto Bemporad. "Model Predictive Control for Powered Descent Guidance and Control". In: *2015 European Control Conference (ECC)*. Linz, Austria: IEEE, July 2015, pp. 1388–1393.
- [13] Kai Dresia et al. "Multidisciplinary Design Optimization of Reusable Launch Vehicles for Different Propellants and Objectives". In: *Journal of Spacecraft and Rockets* 58 (Feb. 2021), pp. 1–13.
- [14] Vladimir Dobrokhodov, Oleg Yakimenko, and Christopher Junge. "Six-Degree-of-Freedom Model of a Controlled Circular Parachute". In: *Journal of Aircraft* 40 (May 2003), pp. 482–493.
- [15] Giorgio Guglieri. "Parachute-Payload System Flight Dynamics and Trajectory Simulation". In: *International Journal of Aerospace Engineering* 2012 (May 2012).
- [16] Pedro dos Santos and Paulo Oliveira. "Thrust Vector Control and State Estimation Architecture for Low-Cost Small-Scale Launchers". In: *Acta Astronautica* 222 (Sept. 2024), pp. 52–68.
- [17] Raphael Linsen et al. "Optimal Thrust Vector Control of an Electric Small-Scale Rocket Prototype". In: *2022 International Conference on Robotics and Automation (ICRA)*. Philadelphia, PA: IEEE, May 2022, pp. 1996–2002.
- [18] Steven J. Isakowitz, Joseph P. Hopkins, and Joshua B. Hopkins. *International Reference Guide to Space Launch Systems*. Reston, VA: American Institute of Aeronautics and Astronautics, 2004.
- [19] N. Holmberg, R. P. Faust, and H. M. Holt. "Viking '75 Spacecraft Design and Test Summary. Volume 1: Lander Design". In: Nov. 1980.
- [20] Mark A. Kane and Roger Wacker. "Orion GN&C Detection and Mitigation of Parachute Pendulosity". In: *American Astronautical Society*. Springfield, VA, Feb. 2016.
- [21] George C. Nield. *Final Environmental Assessment for Issuing an Experimental Permit to SpaceX for Operation of the DragonFly Vehicle at the McGregor Test Site*. Tech. rep. Federal Aviation Administration, Aug. 2014.
- [22] David J. Cockrell. *The Aerodynamics of Parachutes*. AGARDograph 295. Neuilly sur Seine: North Atlantic Treaty Organization, Advisory Group for Aerospace Research & Development, 1987.
- [23] Jacopo Carradori. "6-DOF Atmospheric Rocket Landing Guidance Using Meta-Reinforcement Learning". In: (Sept. 2024), p. 172.
- [24] Gene Franklin and J. D. Powell. "Digital Control of Dynamic Systems". In: (Jan. 1980).
- [25] Mingzhen Zhang et al. "Computing the Weight Matrix R of Linear Quadratic Control Based on Riccati Equation". In: *IFAC-PapersOnLine* 59.20 (Jan. 2025), pp. 1890–1894.
- [26] Hugh L. Dryden. "A Review of the Statistical Theory of Turbulence". In: *Quarterly of Applied Mathematics* 1.1 (Apr. 1943), pp. 7–42.
- [27] Michael M. Madden. "Verifying Implementation of the Dryden Turbulence Model and MIL-F-8785 Gust Gradient". In: *2018 Modeling and Simulation Technologies Conference*. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 2018.
- [28] BristolSEDS. *Mach24 Blackbird External Design Review*. Tech. rep. Apr. 2024.
- [29] HyPower Bristol. *Team 16 Technical Report for the 2025 EuRoC*. Tech. rep. BristolSEDS, Oct. 2025.
- [30] Alicia Dwyer Cianciolo et al. "Defining Navigation Requirements for Future Precision Lander Missions". In: ().
- [31] Duoqi Gao et al. "Adaptive Pseudospectral Successive Convex Optimization for Six-Degree-of-Freedom Powered Descent Guidance". In: *Aerospace Science and Technology* 155 (Dec. 2024), p. 109544.
- [32] Max M. Munk. "The Aerodynamic Forces on Airship Hulls". In: ().
- [33] Jack N. Nielsen. *Missile Aerodynamics*. Reprinted on the occasion of the 70th birthday of Jack N. Nielsen. Mountain View, California: Nielsen Engineering & Research, Inc., 1988.
- [34] Guangzhong Gao et al. "Modelling Nonlinear Aerodynamic Damping during Transverse Aerodynamic Instabilities for Slender Rectangular Prisms with Typical Side Ratios". In: *Journal of Wind Engineering and Industrial Aerodynamics* 197 (Feb. 2020), p. 104064.
- [35] Emil Simiu and DongHun Yeo. *Wind Effects on Structures: Modern Structural Design for Wind*. 4th. Hoboken, NJ: John Wiley & Sons, 2019.
- [36] *An Introduction to Boundary Layer Meteorology*.
- [37] David J. Moorhouse and Robert J. Woodcock. *MIL-F-8785C*. Tech. rep. AD A119421. Air Force Wright Aeronautical Laboratory, July 1982, p. 255.
- [38] BristolSEDS Experimental Propulsion. *Warwing Integrated Hot-fire Campaign*. Tech. rep. Astron Systems, Aug. 2025.
- [39] muRata. *SCH16T-K01*. Datasheet.
- [40] u-blox. *MAX-M10S*. Datasheet. Jan. 2026.
- [41] Richard C. Dorf and Robert H. Bishop. *Modern Control Systems*. 14th, Global Edition. Harlow: Pearson, 2022.
- [42] Huanshui Zhang et al. "Linear Quadratic Regulation and Stabilization of Discrete-Time Systems With Delay and Multiplicative Noise". In: *IEEE Transactions on Automatic Control* 60.10 (Oct. 2015), pp. 2599–2613.
- [43] Karl J. Astrom. *Computer Controlled Systems*. 3rd. Prentice Hall, 1997.
- [44] Daniel Liberzon. "Switched Systems: Stability Analysis and Control Synthesis". In: ().
- [45] Sara McNamara et al. "Gain Scheduling for the Orion Launch Abort Vehicle Controller". In: *AIAA Guidance, Navigation, and Control Conference*. Portland, Oregon: American Institute of Aeronautics and Astronautics, Aug. 2011.
- [46] B. Efron and R. Tibshirani. "Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy". In: *Statistical Science* 1.1 (Feb. 1986), pp. 54–75.
- [47] Vanessa V. Aubuchon and Bruce Owens. "Damping Effects of Drogue Parachutes on Orion Crew Module Dynamics". In: *AIAA SPACE 2016*. Long Beach, California: American Institute of Aeronautics and Astronautics, Sept. 2016.
- [48] J. M. Hanson. *Applying Monte Carlo Simulation to Launch Vehicle Design and Requirements Analysis*. Tech. rep. NASA/TP—2010-216447. Marshall Space Flight Center, Alabama: NASA, Sept. 2010, p. 134.
- [49] Andrew Cucchiaro. "Applied Logistic Regression". In: *Technometrics* 34 (Mar. 2012), pp. 358–359.
- [50] Behcet Acikmese, John M. Carson, and Lars Blackmore. "Lossless Convexification of Nonconvex Control Bound and Pointing Constraints of the Soft Landing Optimal Control Problem". In: *IEEE Transactions on Control Systems Technology* 21.6 (Nov. 2013), pp. 2104–2113.
- [51] Nihad Daidzic. "Global Optimized Isothermal and Nonlinear Models of Earth's Standard Atmosphere". In: *International Journal of Aviation, Aeronautics, and Aerospace* (2014).
- [52] Marcel J. Sidi. *Spacecraft Dynamics and Control: A Practical Engineering Approach*. Cambridge; New York: Cambridge University Press, 1997.

5. Personal Reflection

5.1. Summary of Activities

The aspect of this project I am most proud of is the phased controller architecture I eventually arrived at. I tried a considerable number of approaches before this one, a single LQR with static gains, a gain-scheduled variant without mode switching, and a few others but none of them coped with the tight $\pm 10^\circ$ gimbal envelope and $45^\circ/\text{s}$ slew rate once realistic disturbances were introduced. Finding a four-gate activation layer feeding a three-mode phased LQR, with the predictive-recovery and hammer overrides sitting above the optimal loop as independent safety layers, was the point at which the Monte Carlo success rate moved out of the 20–40% regime and into the 70% band reported in Section 3.1.4. The satisfaction was less in the final structure itself than in the fact that every element of it earned its place against a physical failure mode I had observed in an earlier version. The controller is, to my knowledge, a novel combination for hybrid parachute-TVC recovery, and the fact that it works within such restrictive actuator limits is the result I am most proud of.

5.2. Key Learnings

The project developed a set of technical skills I did not have at the start. I learnt how to derive state-space equations of motion for a complex coupled system from first principles, rather than adapting a textbook rigid-body template, and to assemble the coupled mass matrix with apparent-mass and Munk-moment terms in a form that a numerical integrator could actually use. I learnt how to synthesise results from very different research domains and combine their governing equations inside a single self-consistent simulation rather than treating them as isolated literatures which I had never done previously to this extent. I built a MATLAB and Simulink model structured around modular, iterable subsystems with a shared state vector, which made the difference between a simulation I could develop over months and one that would have collapsed under its own complexity. I also learnt how to conduct Monte Carlo experimental design and the statistical inference techniques that turn an ensemble of runs into a defensible claim: sample-size justification, confidence intervals on success rate, Welch's t-tests for propellant comparisons, logistic regression for failure-mode attribution, and chi-squared tests for controller architecture comparison.

5.3. Challenges and Solutions

The dominant non-technical challenge was compressing everything into the 15-page report limit. I had enough material for at least twice that length, and deciding what was load-bearing enough for the main body and what belonged in appendices forced a much sharper view of what the central argument actually was. Roughly half of the methods chapter now lives in appendices, which initially felt like I was hiding work, but ultimately made the main body readable and the argument visible. The dominant technical challenge was quite different: establishing which modelling assumptions would actually hold under the conditions I cared about, and letting that judgement shape the direction of the work rather than fixing a scope upfront and discovering the holes afterwards. Questions like whether the small-angle linearisation would survive the activation transient, whether the potential-flow

apparent-mass coefficients remained valid once the canopy was oscillating, and whether a 3DOF planar projection could be reconciled with 6DOF flight data, each had to be resolved before I could trust any downstream result. Working through these, rather than assuming them away, is what pushed the linearisation-breakdown discussion in Section 3.2.6 and the 3D-to-1D aggregation argument in Section 3.2.2 into the report as substantive findings rather than caveats.

5.4. Feedback and Reflection

The project evolved substantially from my initial framing. I wanted to go all-out from the start with a 6DOF simulation and an MPC controller, on the assumption that higher fidelity everywhere would produce a more defensible result. Early feedback pushed me back to a 3DOF plant with a simple LQR controller, on the grounds that a simpler architecture would be far easier to tune, benchmark, and diagnose when something went wrong, and that the novel contribution of the project lay in the hybrid architecture rather than in controller complexity. That turned out to be exactly right: the tuning work in Section 2.7 and the failure-mode attribution in Section 3.1.4.4 would have been practically impossible on an MPC-6DOF stack within the project timeline. The second piece of feedback that shaped the final form of the report was the push to articulate the “so what”: why a hybrid architecture is actually useful rather than merely novel. Taking that seriously is what produced the quantified propellant-efficiency argument in Section 3.2.5, the testbed-accessibility argument in Section 3.2.7, and the L–R trade-off discussion in Section 3.2.4, and it is those three sections together that carry the thesis of the discussion chapter. Without that reframing, I would have written a report that demonstrated a controller worked and stopped there.

5.5. Plan for Next Week

If I were starting the project again, the first change I would make is modelling the wind deterministically before making it stochastic. The Dryden implementation is the element that took longest to trust, partly because a stochastic disturbance makes every other subsystem harder to debug since you cannot tell whether a strange plant response is a modelling error or a gust realisation until you have the plant independently verified against something repeatable. Building a deterministic wind field first, validating the plant and controller against it, and only then layering in Dryden turbulence would have given me a far more consistent foundation to work from. The second change would be moving to an inner-outer loop controller architecture: an inner pitch-stabilisation loop feeding from an outer loop that commands pitch targets to generate horizontal velocity toward a specified landing zone. The current controller does not attempt lateral guidance, and an inner-outer structure would have opened that capability without changing the optimal core. The single piece of advice I would give myself at the start of week one is to narrow the scope. A deeper, more detailed treatment of a smaller problem would have produced a stronger report than the broader treatment I initially set out to write, and almost every good decision I made over the course of the project was, in effect, a narrowing of scope that I should have made earlier.

A. Appendices

A.1. Nomenclature (Acronyms)

Table 2: List of Acronyms and Meanings

Acronym	Meaning
AGL	Above Ground Level
ARE	Algebraic Riccati Equation
AUC	Area Under the (ROC) Curve
CEP	Circular Error Probable
CI	Confidence Interval
CoG	Centre of Gravity
CoM	Centre of Mass
DOF	Degree of Freedom
FFT	Fast Fourier Transform
FOH	First-Order Hold
GPS	Global Positioning System
IMU	Inertial Measurement Unit
LQR	Linear Quadratic Regulator
MC	Monte Carlo
MCI	Mass, Centre of Gravity and Inertia
MIL	U.S. Military Specification standard
MIMO	Multi-Input, Multi-Output
MPC	Model Predictive Control
NASA	National Aeronautics and Space Administration
NED	North-East-Down inertial frame
OOS	Out-Of-Specification landing outcome
PBL	Planetary Boundary Layer
PDF	Probability Density Function
PDG	Powered Descent Guidance
PID	Proportional-Integral-Derivative controller
PSD	Power Spectral Density
RK4	4th-order Runge-Kutta numerical integration
RMS	Root Mean Square
SEDS	Students for the Exploration and Development of Space
TFF	(GPS) Time To First Fix
TVC	Thrust Vector Control
VKC	von Kármán Constant
ZOH	Zero-Order Hold

Table 4: Mass and Geometric Properties

Symbol	Description	Value	Unit
m_{wet}	Wet mass (propellant-loaded)	78.286	kg
m_{dry}	Dry mass (propellant-depleted)	65.417	kg
L	Total airframe length	3.97	m
d	Airframe diameter	0.20	m
$z_{CoG,wet}$	Wet-mass CoG station (from aft)	2.31	m
$z_{CoG,dry}$	Dry-mass CoG station	2.18	m
z_p	Parachute attachment station	0.116	m
z_T	Engine gimbal station	3.59	m
L_p	Distance from CoG to chute pivot	2.19	m
L_T	Distance from CoG to gimbal	1.28	m
I_s	Transverse (pitch/yaw) MOI	103.0	kg·m ²
I_{zz}	Longitudinal (roll) MOI	0.392	kg·m ²
I_p	Parallel-axis pivot inertia	479.9	kg·m ²

Table 5: Propulsion and Actuator Properties

Symbol	Description	Value	Unit
T_{max}	Maximum thrust	5000	N
T_{min}	Min. controllable thrust	100	N
T_{hov}	Nominal hover thrust	$m \cdot g$	N
I_{sp}	Specific impulse	180	s
τ_T	Throttle time constant	0.05	s
T_{max}	Throttle rate limit	5000	N/s
t_{delay}	Ignition delay	0.02	s
A_{spike}	Ignition spike magnitude	1.1	–
τ_{spike}	Ignition spike decay time	0.05	s
δ_{max}	Max. gimbal deflection	10	deg
$\dot{\delta}_{max}$	Gimbal slew rate limit	45	deg/s
τ_{δ}	Gimbal servo time constant	0.02	s
T_s	Control loop period (50 Hz)	0.02	s

Table 6: Aerodynamic and Parachute Properties

Symbol	Description	Value	Unit
d_{chute}	Parachute diameter	3.49	m
A_{chute}	Parachute reference area	9.57	m ²
$C_{d,chute}$	Parachute drag coefficient	2.0	–
$C_{d,body}$	Body drag coefficient	1.2	–
ε	Canopy shape ratio	0.5	–
R_p	Inflated canopy radius	1.745	m
η_V	Effective canopy volume factor	0.65	–
k_{11}	Transverse apparent mass coeff.	0.30	–
k_{33}	Axial apparent mass coeff.	0.65	–
k_{44}	Apparent MOI coefficient	0.18	–
k_{15}	Trans.-rot. coupling coeff.	0.40	–
m_a	Apparent (entrained) air mass	derived	kg
I_{pp}	Apparent rotational inertia	derived	kg·m ²

Table 3: State and Control Variables

Symbol	Description	Unit
x^I, z^I	Inertial downrange / depth	m
h	Altitude AGL (nose-bottom)	m
v_x^I, v_z^I	Inertial translational velocities	m/s
v_h	Altitude rate ($-v_z^I$, positive upward)	m/s
θ	Pitch angle (body w.r.t. inertial)	rad
ω	Pitch rate	rad/s
T	Actual engine thrust	N
δ	Actual gimbal deflection	rad
T_{cmd}, δ_{cmd}	Commanded thrust and gimbal	N, rad
$v_{turb,x}, v_{turb,z}$	Dryden turbulence states	m/s

Table 7: Environmental Parameters

Symbol	Description	Value	Unit
g	Gravitational acceleration	9.81	m/s^2
ρ	Sea-level air density	1.225	kg/m^3
v_{ref}	Reference wind at 10 m	5	m/s
VKC	von Kármán constant (terrain)	0.143	–
$v_{gust,p}$	Peak discrete gust velocity	5	m/s
L_m	Gust pocket spatial thickness	30	m
t_{gust}	Gust onset time	5	s
σ_x	Longitudinal turbulence intensity	5	m/s
σ_w	Vertical turbulence intensity	0.5	m/s
L_{FA}	Free-atmosphere turbulence scale	533.4	m

B. Appendix

B.1. Reference Frames

The global environment is characterized by a North–East–Down (NED) inertial frame, denoted $\mathcal{I}$, fixed at the initial point of the vehicle’s descent in the air. In this convention, the z^I axis is directed downward, so that an increase in z^I corresponds to a decrease in altitude [31]. While this frame is not strictly inertial due to the rotation of the Earth, the short duration of the landing burn (approximately 65s of simulated descent from 400m) allows Earth rotation to be neglected without compromising simulation fidelity [23]. Within this 3DOF model, the state variables $[x, z]$ track the North displacement and depth respectively, with $h = z_0 - z^I - (L - z_{CoG})$ representing the positive vertical altitude of the rocket base above the ground.

The vehicle’s physical state is defined within a Body frame, denoted $\mathcal{B}$, fixed to the rocket with its origin at the instantaneous Centre of Gravity (CoG). The z^B axis is aligned with the longitudinal symmetry of the airframe, pointing toward the nose. Under this orientation the primary thrust vector is directed along the $-z^B$ axis as it exits the engine nozzle. The body frame is the primary reference for calculating the moments of inertia and the aerodynamic forces acting upon the slender body of the rocket [23].

To model the active redirection of thrust, a Propulsion frame $\mathcal{P}$ is attached directly to the gimbal assembly. This frame defines the specific direction of the thrust vector relative to the body frame, accounting for the mechanical deflection of the nozzle. The orientation of the $\mathcal{P}$ frame relative to $\mathcal{B}$ is parametrised solely by the gimbal angle δ .

Transforming vectors between the NED inertial frame and the body-centric perspective is achieved through planar rotation matrices. The matrix R^{BI} transforms a vector from the world frame into the rocket’s perspective using the pitch angle θ :

$$\begin{bmatrix} x^B \\ z^B \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x^I \\ z^I \end{bmatrix} \quad (16)$$

Conversely, R^{IB} projects forces such as the thrust generated in the body frame back into the inertial environment for state integration:

$$\begin{bmatrix} x^I \\ z^I \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x^B \\ z^B \end{bmatrix} \quad (17)$$

The relationship between the intended thrust and its actual gimballed direction is handled by R^{PB} , mapping the thrust from the propulsion frame into the body frame as a function of δ :

$$\begin{bmatrix} x^B \\ z^B \end{bmatrix} = \begin{bmatrix} \cos \delta & -\sin \delta \\ \sin \delta & \cos \delta \end{bmatrix} \begin{bmatrix} x^P \\ z^P \end{bmatrix} \quad (18)$$

The integration of these three matrices ensures that the 3DOF model accurately reflects the horizontal and vertical components of motion in the x – z vertical plane throughout the landing sequence.

B.2. State Vector Definition

The position and velocity of the vehicle are described using a Cartesian representation relative to the inertial reference frame. Because the flight takes place at a low altitude below 2km, a flat-Earth assumption is adopted, making it convenient to describe motion in a surface-relative frame. The Cartesian components express the North displacement and depth with respect to the NED inertial frame established at the deployment point. In this implementation, x^I and z^I track global position, while v_x^I and v_z^I represent the translational velocity components required for state-derivative assembly. Cartesian coordinates provide a more intuitive interpretation of low-altitude flight than spherical coordinates, where small variations in latitude or longitude are numerically difficult to discern [23].

The attitude of the rocket is defined as the orientation of the body-fixed reference frame relative to the NED inertial frame. While higher-fidelity simulations often use quaternions due to their computational efficiency and absence of gimbal-lock singularities, this planar 3DOF model uses the classical pitch Euler angle θ for both simulation and visualization [23]. This choice is preferred for interpreting the “pendulum” behaviour of the vehicle suspended under a parachute, where the pitch angle directly corresponds to the swing magnitude. The angular rate ω^B represents the rotational velocity of the body frame with respect to the inertial frame, expressed in body-frame components. In the Waxwing model this rate is the pitch rate the controller must damp prior to touchdown.

B.3. Recovery Hardware and Apparent-Mass Model

The Waxwing recovery system is a single main parachute deployed from the nose-cone compartment following apogee drogue separation. The main canopy is a 3.49 m toroidal parachute selected for its combination of high drag coefficient ($C_{d,chute} = 2.0$) and controlled inflation transient. The parachute mass breakdown comprises the canopy ($\approx 2.1\text{kg}$), suspension lines and riser ($\approx 0.8\text{kg}$), and pyrotechnics ($\approx 0.6\text{kg}$), for a total recovery-system mass of approximately 3.5 kg.

The aerodynamic coupling between the canopy and the airframe is captured through four apparent-mass coefficients drawn from the classical Doherr set scaled to the canopy geometry following Dobrokhodov et al. [14]. When a bluff body accelerates through a fluid it must accelerate not only itself but also a volume of the surrounding air that is kinematically entrained by its motion. For a canopy of inflated radius R_p and shape ratio ε (minor-to-major axis ratio of the inflated hemi-spheroid), the reference entrained-air mass is

$$m_a = \frac{2}{3} \pi \rho R_p^3 \varepsilon \eta_V \quad (19)$$

where $\eta_V = 0.65$ is an empirical effective-volume factor accounting for the fact that the inflated canopy is not a perfect hemispheroid and that porosity reduces the effectively-coupled air volume. Using $R_p = 1.745$ m, $\varepsilon = 0.5$, $\rho = 1.225$ kg/m³ and $\eta_V = 0.65$ gives $m_a \approx 8.87$ kg, approximately 12% of the vehicle wet mass, confirming that the apparent-mass terms are not negligible perturbations but first-order contributions to the coupled dynamics. The reference apparent rotational inertia about the body pitch axis, treating the same air slug as a spheroid, is

$$I_{pp} = \frac{1}{5} m_a R_p^2 (1 + \varepsilon^2) \quad (20)$$

The four dimensionless coefficients scale this reference mass / inertia into the specific directional and rotational components that populate the coupled mass matrix:

- $k_{11} = 0.30$ — the transverse apparent mass coefficient. It scales m_a into the additional inertial resistance the system experiences when translated sideways (perpendicular to the canopy axis). This is the air that must be shoved out of the way when the parachute-rocket assembly swings laterally under a gust.
- $k_{33} = 0.65$ — the axial apparent mass coefficient. It scales m_a into the additional inertial resistance for motion along the canopy symmetry axis (i.e. during vertical descent acceleration). A toroidal canopy couples to a much larger volume of air axially than transversely, so $k_{33} > k_{11}$ is the physical origin of the Munk moment asymmetry discussed in Section 2.3.2.
- $k_{44} = 0.18$ — the apparent moment-of-inertia coefficient. It scales I_{pp} into the rotational resistance to pitching, representing the additional rotational inertia from air that must be spun up when the vehicle rotates about its CoG.
- $k_{15} = 0.40$ — the translation-rotation coupling coefficient. This is the key term that makes the mass matrix non-diagonal: it captures the fact that sideways translation of the canopy under its tether necessarily induces a rotation of the suspended body (and vice versa). It enters the off-diagonal entries of the coupled mass matrix through $K_{total} = -k_{15} m_a L_p$.

B.4. Gravitational Model

For the purposes of this study a constant gravity model is employed to calculate the weight of the vehicle and its influence on the pendulum dynamics. Gravitational acceleration is assumed uniform with the standard value $g = 9.81$ m/s². This is justified by the mission scenario taking place below 2 km altitude, within which the variation in gravitational acceleration due to the inverse-square law is of order 10^{-3} and has no discernible effect on control authority or pendulum frequency [51].

The gravitational force F_G^I is expressed in the NED inertial frame as a vector acting exclusively along the z^I axis. Because the origin of the NED frame is fixed at the deployment point in the air, the gravitational vector is directed downward, consistent with the convention that increasing z^I corresponds to increasing depth. This force is applied at the CoG in the state-derivative assembly (Appendix B.7) and therefore generates no moment about the CoG in the model; rotational moments arise solely from the four torque sources enumerated in Appendix B.5.

B.5. External Forces

The translational motion of the rocket in the vertical pitch plane is driven by the interaction of gimballed thrust, passive aerodynamic drag on the parachute canopy, and gravity. Resolving the vector sum in the inertial frame gives

$$F_x = T \sin(\theta - \delta) - D_x, \quad F_z = -(T \cos(\theta - \delta) + D_z) \quad (21)$$

in which F_z is directed primarily along the negative depth axis, counteracting the downward gravitational force (applied separately in the acceleration solver) to effect the controlled landing burn.

Aerodynamic drag: The primary source of aerodynamic resistance is the toroidal parachute, which generates a drag force opposing the vehicle's motion relative to the local air mass. The total magnitude, using the instantaneous relative-velocity magnitude,

$$D_{total} = \frac{1}{2} \rho v_{rel, mag}^2 C_{d, chute} A_{chute} \quad (22)$$

is then decomposed into its horizontal and vertical inertial components along the direction of the relative wind [22]:

$$D_x = D_{total} \frac{v_{rel, x}}{v_{rel, mag}}, \quad D_z = D_{total} \frac{v_{rel, z}}{v_{rel, mag}} \quad (23)$$

A small additive body-drag term could be written for the bare airframe but for a $C_{d, body} \cdot A_{body}$ product of order 0.04, m² against the canopy's $C_{d, chute} \cdot A_{chute} \approx 19$, m² the body drag contributes less than 0.25% of the total axial force and is instead accounted for only through its moment contribution in Appendix B.6.

Propulsion: The propulsive force is generated in the propulsion frame $\mathcal{P}$ with the base thrust vector defined along the nozzle axis as $T^P = (0, T)^T$. Composing the three rotation matrices of Appendix B.1 (i.e. $R^{IB} R^{BP}$) yields the inertial components

$$T_x^I = T \sin(\theta - \delta), \quad T_z^I = -T \cos(\theta - \delta) \quad (24)$$

which combine with the drag terms above to give the total translational force inputs driving the coupled-mass solve of Appendix B.7.

B.6. External Moments

The rotational dynamics of the Waxwing are governed by the sum of four physically distinct moments acting about the instantaneous CoG.

Parachute restoring torque: The canopy drag force, applied at the pivot L_p ahead of the CoG, produces a pendulum restoring moment that acts to return the airframe toward the relative-wind direction [14]:

$$\tau_{chute} = -D_x \cdot L_p \cdot \cos(\theta) - D_z \cdot L_p \cdot \sin(\theta) \quad (25)$$

This is the primary passive stabilising moment, and its amplitude-dependent coupling with body-relative damping (below) sets the quasi-linear ring-down envelope observed in Section 2.7.4.

TVC control torque: Nozzle deflection produces a lateral thrust component at moment arm L_T [23]:

$$\tau_T = T \cdot \sin(\delta) \cdot L_T \quad (26)$$

This is the only actively controllable moment in the model.

Munk moment: The Munk moment (sometimes called the “slender-body destabilising moment” or simply the “cross-flow inertia moment” in the aerospace literature [33][32]) arises whenever a submerged body moves through a fluid with asymmetric apparent masses in its principal directions. The effect is most easily visualised for a slender airship: if the body has large a_{33} (hard to accelerate along its axis, because a long column of air must be accelerated with it) and small a_{11} (easier to accelerate sideways), then any velocity component at an angle to the body axis produces unbalanced momentum exchange between the two directions, and the resulting pressure distribution generates a torque that tries to rotate the body perpendicular to the flow. In the present system the physical origin of the effect is not the slender-body asymmetry of the airframe itself, but the canopy’s hemispheroidal apparent-mass asymmetry ($a_{33} \gg a_{11}$ for a bluff drag device). The torque about the CoG retains the classical form

$$\tau_{munk} = -(a_{33} - a_{11})v_x^B v_z^B \quad (27)$$

with $a_{ii} = k_{ii}, m_a$. Because $a_{33} > a_{11}$ in the canopy-dominated regime, the coefficient is negative, giving the torque a destabilising sign: a sideways velocity component couples with the axial descent velocity to produce a moment that rotates the airframe further into the relative flow rather than away from it. This is the primary aerodynamic instability the TVC must reject when disturbed, and it is the physical mechanism that converts a steady side-wind into a divergent pendulum swing if the controller is absent.

Body-relative damping: Modelling the bare airframe as a cylinder rotating through still air (no canopy contribution, which is accounted for separately through τ_{chute}), the local transverse velocity at station r from the CoG is $v_{local}(r) = v_x^B + \omega r$. The drag force on an elemental strip dr opposes the local velocity (enforced by the $|v_x^B + \omega r|$ factor) and produces a moment r, dF about the CoG. Integrating from the gimbal station $-L_T$ to the chute-pivot station L_p [34]

$$Q_d = - \int_{-L_T}^{L_p} \frac{1}{2} \rho \cdot d \cdot C_{d_{body}} \cdot (v_x^B + \omega r) |v_x^B + \omega r| \cdot r \, dr \quad (28)$$

This integral is evaluated each time step using a 50-point trapezoidal quadrature. The resulting torque is intrinsically amplitude-dependent through the quadratic $v, |v|$ kernel, which is the physical origin of the weak amplitude-ordering of the damping ratio reported in Section 3.1.1 and discussed in Section 3.2.1.

B.7. Coupled Mass Matrix and State-Derivative Assembly

The final stage of the equations of motion resolves the translational and rotational accelerations by inverting the system’s inertia matrix augmented with the apparent-mass and coupling terms introduced in Appendix B.3. Because Waxwing descends immersed in a moving fluid while tethered to a large canopy, its translational and rotational responses are physically linked: a sideways velocity of the airframe drags the attached canopy sideways, which in turn swings the vehicle’s line of suspension and imparts a rotational acceleration, and conversely a pitching rotation about the CoG displaces the canopy’s air slug laterally and so imparts a transverse translational acceleration.

The coupling coefficient k_{15} quantifies the strength of this kinematic link [14].

The system’s effective inertia is expressed as the Coupled Mass Matrix

$$M_{mat} = \begin{bmatrix} (m + a_{11}) & 0 & K_{total} \\ 0 & (m + a_{33}) & 0 \\ K_{total} & 0 & (I_s + a_{44}) \end{bmatrix} \quad (29)$$

in which the first row represents transverse translational inertia augmented by $a_{11} = k_{11}, m_a$, the second represents axial translational inertia augmented by $a_{33} = k_{33}, m_a$, and the third represents rotational inertia augmented by $a_{44} = k_{44}, I_{pp}$. The axial row remains decoupled because vertical descent along the canopy axis does not itself induce pitching about the CoG. The coupling coefficient

$$K_{total} = -k_{15} \cdot m_a \cdot L_p \quad (30)$$

derived by Dobrokhodov from potential-flow analysis of the payload-canopy system, directly ties transverse acceleration to pitching acceleration through the off-diagonal entries. A positive transverse acceleration produces a proportional pitching acceleration (and vice versa) with a sign set by the direction of the moment arm L_p .

The accelerations are recovered each time step by inverting the 3×3 system

$$\begin{bmatrix} (m + a_{11}) & 0 & K_{total} \\ 0 & (m + a_{33}) & 0 \\ K_{total} & 0 & (I_s + a_{44}) \end{bmatrix} \begin{bmatrix} \dot{v}_x^B \\ \dot{v}_z^B \\ \dot{\omega}^B \end{bmatrix} = \begin{bmatrix} \Sigma F_x^B \\ \Sigma F_z^B \\ \Sigma M_y \end{bmatrix} \quad (31)$$

with M_{mat} being small, dense, and well-conditioned (condition number $\sim 10^2$ across the flight envelope), so that direct mldivide inversion is numerically stable at every time step and iterative solvers are unnecessary.

The final step is assembling the state derivative vector

$$\dot{\mathbf{X}} = [\dot{x}^I, \dot{z}^I, \dot{v}_x^I, \dot{v}_z^I, \dot{\theta}, \dot{\omega}, \dot{m}]^T \quad (32)$$

in which the body-frame translational accelerations are transformed back into the inertial frame through R^{IB} with gravity added along $+z^I$, the pitch rate is identically ω , and $\dot{m} = -T/(g_0 I_{sp})$ as derived in Appendix B.8. The full vector is then integrated forward using a fixed-step RK4 scheme at $\Delta t = 1$ ms.

B.8. Mass, CoG and Inertia Model

The fidelity of the simulation is enhanced by treating the Mass, CoG and Inertia (MCI) as time-varying rather than fixed at their wet-mass values. During the landing burn the vehicle transitions from a wet mass of 78.286 kg to a minimum achievable dry mass of 65.417kg, a 16.4% change that is small but non-negligible given the sensitivity of the LQR gains to $T_{hov} = mg$. Over the full 5–15s of powered flight in a typical landing, continuous re-evaluation of the mass-dependent terms prevents gradual accumulation of model error that would otherwise appear as a slow thrust trim bias.

The evolution of mass depletion is governed by

$$\dot{m} = -\frac{T}{g_0 I_{sp}} \quad (33)$$

with T the instantaneous thrust and $I_{sp} = 180$ s. The propellant fraction γ and the linearly-interpolated CoG

station are

$$\gamma = \frac{m - m_{dry}}{m_{wet} - m_{dry}} \quad (34)$$

$$z_{CoG} = z_{CoG_{dry}} + \gamma(z_{CoG_{wet}} - z_{CoG_{dry}}) \quad (35)$$

For the purposes of this analysis the CoG is taken as co-incident with the centre of mass ($z_{CoG} = z_{CoM}$) [52], providing a single point through which gravity acts and about which moments are summed. The CoG shift between wet and dry states is approximately 13 cm, which produces a 10% swing in the gimbal moment arm L_T and a comparable correction to the parachute pivot arm L_p ; both are re-computed at every time step.

The inertia tensor is modelled as time-varying through the cylindrical-shell approximations

$$I_s(t) = \frac{1}{12}m(t)L^2 + \frac{1}{4}m(t)(d/2)^2 \quad (36)$$

$$I_{zz}(t) = \frac{1}{2}m(t)(d/2)^2 \quad (37)$$

with $I_p(t) = I_s(t) + m(t)L_p(t)^2$ providing the pendulum-pivot inertia used in the angular-momentum check of Appendix F.1. Although I_s and I_{zz} could be held constant for simplified analysis, the implementation allows these values to update dynamically so that the LQR weight matrices (re-solved off-line, Appendix D.2) remain within 2% of the instantaneous optimal gains across the full mass range.

C.1. Wind Model

To validate the robustness of the TVC system, the simulation incorporates a multi-layered atmospheric environment composed of steady-state wind shear, discrete gusts, and continuous stochastic turbulence. These components are combined to determine the instantaneous relative velocity acting on the vehicle, which directly influences the aerodynamic drag and stability moments.

To prevent unrealistic initial transients at the commencement of the simulation, all wind components are linearly ramped from zero to their full calculated magnitude over the first two seconds of flight. The total wind velocity components along the horizontal and vertical axes are

$$v_{wx} = v_{mean_x} + v_{gust_x} + v_{turb_x} \quad (38)$$

$$v_{wz} = v_{mean_z} + v_{gust_z} + v_{turb_z} \quad (39)$$

and the instantaneous relative velocity follows as

$$v_{rel_x} = v_x - v_{wx}, \quad v_{rel_z} = v_z - v_{wz} \quad (40)$$

$$v_{rel_{mag}} = \sqrt{v_{rel_x}^2 + v_{rel_z}^2} \quad (41)$$

which is the primary input for all subsequent aerodynamic force and moment calculations.

Wind shear: Due to surface friction, wind velocity is non-constant with altitude. The mean horizontal wind is modelled through a logarithmic power law

$$v_{mean_x} = v_{ref} \left(\frac{h}{10} \right)^{VKC} \quad (42)$$

with v_{ref} the 10m reference wind speed and $VKC = 0.143$ the von Kármán constant for open terrain [35]. As $h \rightarrow 0$ the background wind naturally dissipates toward zero, avoiding the “lateral slide” artefact in which a low-altitude

vehicle is carried sideways by a wind that does not physically exist near the ground [36].

Discrete gusts: Transient disturbances are modelled using the MIL-F-8785C 1-cosine spatial profile [37], producing a smooth velocity pulse as the vehicle traverses a gust pocket of length L_m :

$$v_{gust_x}(t) = \frac{v_{gust_p}}{2} \left[1 - \cos \left(\frac{2\pi v_{z_{nom}}(t - t_{gust})}{L_m} \right) \right] \quad (43)$$

with v_{gust_p} the peak magnitude, $v_{z_{nom}}$ the nominal descent speed that sets the traversal rate, and the bump active only inside the time window $t_{gust} < t < t_{gust} + L_m/v_{z_{nom}}$.

Continuous Dryden turbulence: The Dryden turbulence model is implemented following the NASA Langley LaSRS++ methodology [26][27]. The raw white-noise input to each filter is scaled by the simulation time step so that its spectral power is exactly unity before entering the shaping filter:

$$W_x = \eta_x \sqrt{\pi/\Delta t}, \quad W_z = \eta_z \sqrt{\pi/\Delta t} \quad (44)$$

As the vehicle descends, the characteristic turbulence scale lengths shrink according to MIL-F-8785B Planetary Boundary Layer scales [37]:

$$L_{turb,x} = L_{FA}^{2/3} h^{1/3}, \quad L_{turb,z} = h \quad (45)$$

The digitisation strategy is the critical point. NASA research shows that using a First-Order Hold (FOH) on uncorrelated white noise causes variance losses of 7–19%, which would falsify the turbulence energy the controller is tested against. Consequently the simulation uses a Zero-Order Hold (ZOH) for the first filter stage, after which the output is correlated and FOH is used for subsequent stages without incurring variance loss.

The longitudinal gust is a single first-order low-pass

$$G_x(s) = \frac{b_x}{s + d_x} \quad (46)$$

$$b_x = \sigma_x \sqrt{\frac{2, v_{mag}}{\pi L_{turb,x}}}, \quad d_x = \frac{v_{mag}}{L_{turb,x}} \quad (47)$$

The vertical gust is the second-order Dryden spectrum

$$\Phi_w(\Omega) = \sigma_w^2 \frac{L_{turb_z}}{\pi v_{mag}} \frac{1 + 3(L_{turb_z}\Omega/v_{mag})^2}{[1 + (L_{turb_z}\Omega/v_{mag})^2]^2} \quad (48)$$

factored into a ZOH first-order stage driven by white noise in cascade with an FOH first-order stage that introduces the $\sqrt{3}$ numerator term:

$$G_z(s) = \frac{s + \sqrt{3}v_{mag}/L_{turb_z}}{s + v_{mag}/L_{turb_z}} \quad (49)$$

Discretisation uses the NASA discrete algorithms with ZOH constants $C_1 = e^{-d\Delta t}$; $C_3 = (b/d)(1 - C_1)$ and FOH constants $C_1 = e^{-d\Delta t}$; $C_2 = b/d + (1 - C_1)/(d\Delta t) \cdot (a - b/d)$; $C_3 = (b/d)(1 - C_1) - C_2$. By updating b and d using the instantaneous v_{mag} and h at every integration step the simulation preserves turbulence energy across the full flight envelope rather than assuming a single fixed operating point.

The turbulence implementation was validated against the NASA Langley LaSRS++ acceptance criteria [27] across three metrics: zero-mean time-series, unity variance ratio against the specified σ , and reconstruction of the ideal Dryden PSD. Using a dataset of 6.1×10^7 samples sufficient to satisfy 95% confidence intervals on the PSD reconstruction, the longitudinal component achieved a mean of 0.002626 (target 0 ± 0.003) and variance ratio of 1.0010 (target 1.000), and the vertical component a mean of 0.001697 and variance ratio of 0.9943, both inside the NASA-specified bounds. Figure 5 of the main report body overlays the reconstructed and ideal PSDs on both linear and log scales; their near-perfect agreement confirms that the ZOH-first implementation preserves the low-frequency energy where pendulum excitation concentrates.

C.2. Engine and TVC Model

The commanded thrust and gimbal angle (T_{cmd}, δ_{cmd}) emitted by the controller are not realised instantaneously at the nozzle. The Engine & TVC model captures the dominant non-idealities of both subsystems, with parameters fitted directly to BristolSEDS Experimental Propulsion's Engine-1 and the integrated Waxwing feed system during the 2025 full-stack flight-qualification hotfire campaign [38].

Liquid engine dynamics: The engine response is characterised by three mechanisms. On receipt of a non-zero thrust command, the engine produces no thrust for an ignition delay $t_{delay} = 20$ ms, representing valve travel and manifold priming. Once combustion is established, an ignition pressure spike is modelled as an exponentially-decaying overshoot:

$$S(t) = 1 + (A_{spike} - 1)e^{-(t-t_{delay})/\tau_{spike}} \quad (50)$$

with $A_{spike} = 1.1$ the peak magnitude and $\tau_{spike} = 50$ ms the decay time. Subsequent throttle changes follow a first-order lag of time constant $\tau_T = 50$ ms:

$$\dot{T}_{req} = \frac{T_{target} - T_{prev}}{\tau_T} \quad (51)$$

with the commanded rate clipped at the hard slew limit $\dot{T}_{max} = 5000$ N/s prior to integration. The τ_T and $\dot{T}_{max}$ values were fitted to the 4% settling time and peak chamber-pressure rate observed during the Engine-1 tests.

Gimbal actuator dynamics: The TVC assembly redirects thrust through a commanded gimbal angle δ_{cmd} subject to three mechanical constraints applied in cascade: a first-order servo lag of time constant $\tau_\delta = 20$ ms, a slew-rate limit $\dot{\delta}_{max} = 45^\circ/\text{s}$ reflecting actuator torque limits, and a hard angular saturation at $\delta_{max} = 10^\circ$ set by the gimbal housing envelope. The combined update law is

$$\Delta\delta = \text{sat} \left(\frac{\delta_{target} - \delta_{prev}}{\tau_\delta}, \dot{\delta}_{max} \right) \cdot \Delta t \quad (52)$$

with the result subsequently clipped to $\pm\delta_{max}$.

Together the engine and gimbal lags and rate limits define the actuator plant that the controller's lead compensation (§2.7.1) is designed to anticipate. Sustained gimbal saturation during high-rate disturbances is one of the failure modes the hammer-override logic (§2.7.6) exists to handle, and the near-universal transient saturation documented in Section 3.1.4.2 is the signature of the controller correctly sizing its actuator authority for the observed disturbance envelope.

C.3. Sensor Model and Navigation

The state feedback reaching the Controller is not the true plant state $\mathbf{X}$ but an estimate $\hat{\mathbf{X}}$ produced by three cascaded subsystems: an Inertial Measurement Unit (IMU), a Global Positioning System (GPS) receiver, and a complementary filter that fuses them. All three were parameterised to match the avionics stack carried on Waxwing.

The IMU model reproduces the dominant non-idealities of the on-board Murata SCH16T-K01 six-axis sensor [39]. On power-up the sensor applies a 250ms boot delay during which no valid output is emitted, preventing the controller from receiving uninitialised accelerometer readings during the parachute-deployment transient. The accelerometer model takes the true inertial acceleration from the plant, rotates it into the body frame through the instantaneous $R^{BI}(\theta)$, adds the body-frame gravity vector so that it reports specific force rather than kinematic acceleration, and finally applies sensor bias, Gaussian white noise, and hard clipping to match the datasheet. The equivalent chain for the gyroscope is

$$\hat{a}_x = a_x^B - g \sin \theta + b_a + n_a \sqrt{1/\Delta t} \quad (53)$$

$$\hat{a}_z = a_z^B + g \cos \theta + b_a + n_a \sqrt{1/\Delta t} \quad (54)$$

$$\hat{\omega} = \omega + b_g + n_g \sqrt{1/\Delta t} \quad (55)$$

with bias $b_a = 0.01\text{m/s}^2$, $b_g = 0.1^\circ/\text{s}$, noise densities $n_a = 0.0008\text{m/s}^2/\sqrt{\text{Hz}}$, $n_g = 0.0004^\circ/\text{s}/\sqrt{\text{Hz}}$ taken directly from the Murata datasheet, and hard clipping at $\pm 80\text{ m/s}^2$ and $\pm 300^\circ/\text{s}$ applied to both channels before output. Noise variance scales with $\sqrt{1/\Delta t}$ to preserve the continuous-time power spectral density under the 1 ms integration step.

The GPS model reproduces the u-blox MAX-M10S receiver [40] at a 10Hz update rate with a 1.0s hot-start Time-to-First-Fix during which no position or velocity update is available. Between update ticks the most recent measurement is sample-and-held. Position and velocity noise are applied as zero-mean Gaussians with standard deviations at the datasheet Circular Error Probable (1.5m horizontal, 2.0m vertical) and velocity accuracy specification (0.05m/s). No multipath or ionospheric bias is modelled, reflecting the short duration of the landing burn over open terrain.

The navigation computer implements a GPS-aided complementary filter that treats the IMU as a high-frequency short-term source of truth and the GPS as a low-frequency absolute correction. Position and velocity states are propagated by integrating the de-biased IMU accelerations in body axes and transforming back to the inertial frame, then corrected by a weighted sum against the latest GPS fix with weights of $K_{pos} = 0.02$ and $K_{vel} = 0.05$ per 100ms GPS cycle. The attitude state is propagated by integrating the de-biased gyro, then corrected by an accelerometer tilt measurement $\theta_{acc} = \arctan(a_x^{body}/a_z^{body})$ with weight $K_\theta = 0.02$. Critically, the accelerometer correction is gated off whenever the measured specific-force magnitude differs from g by more than 0.5extm/s^2 , which would indicate either the landing burn is active (thrust dominates specific force) or the vehicle is undergoing high-frequency gust buffeting (non-gravitational specific force in the body x axis). This gating ensures the tilt correction only operates during the quasi-static canopy-borne descent where




Figure 16: Navigation filter output (red dashed) overlaid on true plant state (black solid) for a representative L-controller landing run. The complementary filter tracks horizontal and vertical position to within visible plotting tolerance throughout descent, with velocity tracking similarly tight. Pitch and angular rate channels show the noise injection most clearly during the high-rate damping phase ($t \approx 5 - 20$ s, $|\omega| > 100^\circ/\text{s}$), where IMU gyro noise and accelerometer tilt-correction gating produce the visible jitter on the estimated pitch angle. After the controller arrests the inherited pendulum oscillation ($t > 20$ s) the filter converges back to near-truth tracking through the stabilisation and flare phases. The brief divergence near $t \approx 65$ s coincides with the flare ignition transient where the gated tilt correction is suppressed by the high specific-force magnitude during throttle-up.

$\theta_{acc} \approx \theta$, and the filter relies on pure gyro integration through the active-thrust phase.

Together the three subsystems introduce the sensor-induced noise, bias, and latency against which the controller’s lead compensation (§2.7.1) is calibrated. The 0.12s base lead time and the ω -scaled extension to 0.22s were selected to match the combined round-trip latency of IMU acquisition (1 sample at $\Delta t = 1$ ms), complementary-filter propagation (nominally zero but subject to 50Hz control loop quantisation), and gimbal-actuator response ($\tau_\delta = 20$ ms plus up to 100ms of slew-rate-limited rise time at full deflection).

D.1. Controller Architectures

The L controller of Section 2.7 was developed iteratively in dialogue with a more conservative variant, denoted R, which shares the same phased-LQR backbone, identical Q and R weight matrices, identical activation gates, and identical recovery-mode entry conditions, but differs in six specific implementation choices that together convert the controller from an aggressive terminal-recovery scheme into a passive-fallback-biased one. The L–R outcome comparison of Section 3.1.4.5 quantifies the consequences of these differences directly: 70 to 48% success, 23 to 0% spin, 7 to 52% OOS. This appendix enumerates the differences and ties each to its observed effect on outcome distribution. “L” and “R” names are purely how I classified them, “L” standing for Lead-Neutral-Impact-Apex-Dual-Hysteresis-Augmented Trigger and “R” for Reboot-Lead-Neutral-Apex-Dual-Hysteresis-Augmented Trigger which doesn’t mean much besides the author trying to come up with professional-sounding names for controller archetypes.

1. Recovery-mode exit logic: The L controller exits the predictive-recovery state under either of two conditions: (i) attitude and rate return to the control-

lable envelope ($|\theta| < 15^\circ \wedge |\omega| < 10^\circ/\text{s}$), *or* (ii) a time-to-impact estimate $t_{imp} = h / \max(0.1, |v_h|)$ falls below 2.5 s. The R controller removes the second clause entirely: once in recovery, it re-arms only when the parachute has actually restabilised the airframe. The asymmetry is decisive: under L, a vehicle still spinning at $|\omega| > 10^\circ/\text{s}$ at low altitude relinquishes control to the LQR and at least attempts a non-ideal landing burn, whereas under R the same vehicle holds zero thrust all the way to ground impact at the parachute terminal velocity (7–9 m/s). This is the single structural difference that produces the L–R OOS-velocity-band signature in Figure 15: R-controller OOS failures cluster at $|v_z| \approx 7\text{--}9$ m/s because the controller has elected to ride the parachute to ground rather than fire a marginal LQR solution.

2. Hammer override removed. The L controller carries a hard-coded thrust override that forces $T_{cmd} = T_{max}$ whenever ($|\omega| > 30^\circ/\text{s} \wedge h < 50$ m) or $v_h < -12$ m/s, prioritising raw deceleration authority over LQR optimality during edge-case scenarios. The R controller omits this override entirely; thrust is set exclusively by the LQR feedback law throughout the active-control window. The override is the dominant contributor to the L controller’s spin tail: an aggressive thrust spike during a high-rate low-altitude transient can amplify rather than arrest a developing rotation if the gimbal is already saturated. Removing it eliminates spin failures in the R batch (0 of 100), but at the cost of leaving the controller with no recourse against a high-sink-rate fault, which then manifests as additional OOS landings.
3. Damping-mode re-targeting under extreme rotation. When the active mode is damping and $|\omega| > 40^\circ/\text{s}$, the L controller temporarily reassigns its target attitude from $\theta_{ref} = \theta_{pred}$ to $\theta_{ref} = \theta$ (the current angle),




Figure 17: Estimation residuals (filter output minus true state) across all six navigation channels for the run shown in Figure 16. Position residuals remain bounded within ± 2 m throughout descent, well inside the GPS CEP of the u-blox MAX-M10S with stationary noise statistics consistent with the 1.5m horizontal and 2.0m vertical position-noise standard deviations applied at the 10 Hz GPS update rate. Horizontal velocity residuals stay within ± 0.5 m/s after the initial transient, matching the datasheet velocity-accuracy specification. The vertical-velocity residual shows a single-sample initialisation transient at $t \approx 0$ (filter convergence from cold start) before settling to negligible error. Pitch and angular rate residuals are dominated by the $0.1^\circ/\text{s}$ gyro bias and the high-rate IMU noise during the damping phase ($t < 20\text{s}$); after the controller arrests the rotation, both channels return to within $\pm 2^\circ$ pitch and $\pm 1^\circ/\text{s}$ rate, confirming the complementary filter recovers cleanly once the tilt-correction gate re-activates in the lower-acceleration regime.

so the LQR stops trying to steer the airframe back to vertical and focuses purely on rate reduction. The R controller does not perform this reassignment: even at extreme rotation it continues to target the predicted-vertical state, which means the LQR allocates a portion of its actuator authority to attitude correction that is physically futile against the inherited momentum. Without this re-targeting, the R controller is more likely to over-saturate the gimbal during severe rotational transients, but it also makes the controller's behaviour more predictable and avoids the L controller's transient mode-switch which can occasionally synchronise with the gimbal lag and induce a small phase-driven oscillation.

4. Lead compensation: dynamic versus static. The L controller scales its lead time with the instantaneous angular rate, $\tau_{lead} = 0.12 + \min(0.10, |\omega| \cdot 0.02)$ s, ranging from 0.12 s at small rates to 0.22 s at high rates. The R controller uses a fixed $\tau_{lead} = 0.15$ s. The dynamic form gives L more anticipatory authority during high-rate excursions (where the actuator lag of $\tau_T + \tau_\delta + T_s \approx 90$ ms compounds with the descent dynamics), but couples the controller's effective phase margin to a measured state, which is the very state most contaminated by IMU noise during high-rate transients. The static R value is robust against gyro noise amplification at high $|\omega|$ but under-compensates the lag during the $|\omega| > 60^\circ/\text{s}$ regime where it would matter most.
5. Apex-gate strictness. The L controller's apex-detection gate fires whenever $h \leq h_{act,max}$ and $|\omega| \leq 3^\circ/\text{s}$, allowing activation anywhere in the descent below the upper altitude limit. The R controller adds a second altitude constraint requiring $h > h_{act,min}$

(i.e. the gate fires only inside an explicit *window* between the upper and lower altitude bounds) and tightens the rate threshold to $|\omega| \leq 2^\circ/\text{s}$. The R apex gate therefore only activates when the controller is genuinely catching the apex of a clean pendulum swing in the intended altitude band, never by accident at low altitude or during a still-energetic oscillation. This trades a small loss of activation opportunity for a guarantee that activation conditions are physically clean.

6. Engine-on hysteresis margin. The L controller declares the engine "on" whenever $T_{actual} > T_{min}$, with no margin. The R controller adds a 50 N hysteresis band: $T_{actual} > T_{min} + 50$. This prevents engine-state chatter near the minimum-thrust threshold during throttle-down transients, and is the same logic philosophy applied to the damping-mode hysteresis ($10^\circ/\text{s}$ engage, $3^\circ/\text{s}$ release) extended to the engine state.

To summarise, very R modification is in the direction of reduced controller aggressiveness in favour of the parachute-borne fallback. The result is a controller that fails predictably and survivably (OOS at parachute terminal velocity) rather than aggressively and catastrophically (controller-induced spin). In a pure-PDG architecture this would be a strictly inferior design, an OOS landing at 8 m/s under PDG implies hardware loss. Under the hybrid architecture of this work, the parachute guarantees the OOS velocity is survivable, which inverts the trade structure: R becomes the higher-reliability hardware-recovery configuration even though it has the lower headline success rate. The L–R Pareto discussion of Section 3.2.4 in the main text develops the cost-trade implications of this asymmetry.

D.2. System Linearization and LQR Gain Synthesis

The LQR weight matrices are solved against an 8-state linearised plant augmented with the two actuator states from Appendix C.2. Defining the control state

$$\mathbf{x}_c = [N, h, v_x, v_h, \theta, \omega, \Delta T, \delta]^T \quad (56)$$

where $N = x^I - x_0$ is downrange position relative to the release point, h is altitude above ground, $v_h = -v_z^I$ is altitude rate (positive upward), $\Delta T = T - T_{hov}$ is thrust error and δ is actual gimbal deflection, the linearisation about hover ($\theta = 0, \omega = 0, \Delta T = 0, \delta = 0, T = T_{hov} = m, g$) yields the augmented state-space form $\dot{\mathbf{x}}_c = A, \mathbf{x}_c + B, \mathbf{u}_c$ with $\mathbf{u}_c = [\Delta T_{cmd}, \delta_{cmd}]^T$. The non-zero entries of A and B are:

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & g & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1/m & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & M_\theta/I_{tot} & 0 & 0 & T_{hov} L_T/I_{tot} \\ 0 & 0 & 0 & 0 & 0 & 0 & -1/\tau_T & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1/\tau_\delta \end{bmatrix} \quad (57)$$

$$B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1/\tau_T & 0 \\ 0 & 1/\tau_\delta \end{bmatrix} \quad (58)$$

where

$$M_\theta = -m g L_p \\ I_{tot} = I_s + k_{44} I_{pp}$$

The rows have direct physical interpretations. Rows 1–2 are the kinematic chain ($\dot{N} = v_x, \dot{h} = v_h$). Row 3 captures the gravity–pitch coupling ($\dot{v}_x = -g \sin \theta \approx -g\theta$) together with the TVC lateral-force coupling through the gimbal. Row 4 reflects that at hover a small thrust perturbation produces an equal and opposite altitude-rate perturbation ($\dot{v}_h = \Delta T/m$). Row 5 is the integrator relating pitch to pitch rate. Row 6 is the dominant TVC control torque about the pitch axis; the linearised A matrix intentionally omits the small parachute restoring stiffness because that term is already accounted for in the plant and not needed for the controller’s internal model to be stabilising. Rows 7–8 are the first-order actuator dynamics of Appendix C.2.

The LQR gain K is obtained by solving the Algebraic Riccati Equation

$$A^T P + P A - P B R^{-1} B^T P + Q = 0$$

for the positive-definite P and then forming $K = R^{-1} B^T P$. This is performed in MATLAB using ‘lqr(A, B, Q, R)’ which internally calls the Schur-decomposition solver ‘care’, three separate times to produce the stabilisation, damping and flare gain matrices $K_{act}, K_{damp}, K_{flare}$.

D.3. LQR Weight Matrix Selection

The Q and R matrices were chosen using Bryson’s rule, $q_{ii} = 1/(\Delta x_{i,\max})^2$ and $r_{jj} = 1/(\Delta u_{j,\max})^2$, where $\Delta x_{i,\max}$ is the maximum permissible excursion of state i and $\Delta u_{j,\max}$ the maximum desired actuator effort. This produces gains in a natural dimensional ratio before any manual tuning, with per-mode adjustments restricted to the angular states that dominate the operational regime of each phase.

Stabilisation mode (default branch, default target $\theta_{ref} = 0; v_{h,ref} = -8$ m/s):

$$Q_{act} = \text{diag} \begin{pmatrix} 1/100^2 \\ 1/100^2 \\ 1/2^2 \\ 1/1^2 \\ 1/(2^\circ)^2 \\ 1/(0.5^\circ/\text{s})^2 \\ 1/100^2 \\ 1/(1^\circ)^2 \end{pmatrix} \quad (59)$$

$$R_{act} = \text{diag}(10^{-3}, 10) \quad (60)$$

The thrust input weight 10^{-3} reflects that thrust commands are cheap ($\sim$ kN range, permissible), while the gimbal weight 10 penalises aggressive gimbaling to prevent actuator saturation against the $\pm 10^\circ$ envelope.

Flare mode (engaged at $h \leq 25$ m, target $v_{h,ref} = -1.8$ m/s):

$$Q_{flare} = \text{diag} \begin{pmatrix} 1/100^2 \\ 1/100^2 \\ 1/2^2 \\ 1/0.2^2 \\ 1/(0.2^\circ)^2 \\ 1/(0.1^\circ/\text{s})^2 \\ 1/100^2 \\ 1/(1^\circ)^2 \end{pmatrix} \quad (61)$$

$$R_{flare} = R_{act} \quad (62)$$

The key change is a $25\times$ increase in vertical-velocity weight (from ± 1 m/s tolerance to ± 0.2 m/s) and a $100\times$ increase in angular-state weights, reflecting the tight terminal constraints of Section 2.1.

Damping mode (engaged when $|\omega| > 10^\circ/\text{s}$, released at $|\omega| < 3^\circ/\text{s}$):

$$Q_{damp} = \text{diag} \begin{pmatrix} 10^{-3} \\ 10^{-3} \\ 10^{-3} \\ 10^{-3} \\ 10^{-3} \\ 1/(0.1^\circ/\text{s})^2 \\ 1/100^2 \\ 1/(1^\circ)^2 \end{pmatrix} \quad (63)$$

$$R_{damp} = \text{diag}(10^{-3}, 1) \quad (64)$$

All state weights except angular rate are set to negligible values (10^{-3}), and R on the gimbal input is reduced by an order of magnitude. The result is a gain matrix that behaves as a pure rotational shock absorber, using the full actuator envelope to kill angular momentum without worrying about position, velocity or even attitude error.

E.1. Passive-Plant Damping Extraction Methodology

The natural frequency, damping ratio and time-to-half-amplitude reported in Section 3.1.1 were extracted from the simulated passive pitch response (fig. 7) by the following procedure.

1. Release the plant from $\theta_0 = 10^\circ$, $\omega_0 = 0$, no wind, no control, and integrate for 50 s at $\Delta t = 1$ ms. The resulting time series $\theta(t)$ is the raw under-damped ring-down.
2. Detect peaks using MATLAB ‘findpeaks’ with a minimum peak distance of 0.3 s and minimum amplitude of 0.1° to reject noise. This yields a sequence of (t_k, θ_k) local maxima and minima, of which only the positive peaks are retained for the period and log-decrement computations.
3. Fit an exponential envelope $A(t) = A_0 e^{-\zeta\omega_n t}$ to the absolute peak amplitudes by linear regression of $\log|\theta_k|$ against t_k , with the slope giving the decay rate $\zeta\omega_n$ directly.
4. Compute the damped period as the mean peak-to-peak interval $\bar{T} = \text{mean}(\Delta t_k)$, the damped natural frequency as $\omega_d = 2\pi/\bar{T}$, and the undamped frequency as $\omega_n = \omega_d/\sqrt{1-\zeta^2}$.
5. Compute the time-to-half-amplitude from the fitted envelope as $t_{1/2} = \ln 2/(\zeta\omega_n)$, and the 2%-envelope settling time as $t_{\text{settle},2\%} = \ln(50)/(\zeta\omega_n)$.

Two distinct statistical procedures are applied. The envelope decay rate $\zeta\omega_n$ is bootstrapped by resampling the peak sequence (t_k, θ_k) with replacement $N = 1000$ times, repeating the linear-regression envelope fit on each resample, and taking the [2.5%, 97.5%] percentiles of the resulting distribution as the 95% CI; the time-to-half-amplitude CI is derived directly from this via $\tau_{1/2} = \ln 2/(\zeta\omega_n)$. The natural-frequency CI is computed parametrically as $\bar{T} \pm 1.96 \text{SE}(\bar{T})$ from the peak-to-peak period distribution, propagated to ω_n through $\omega_n = 2\pi/\bar{T} \cdot 1/\sqrt{1-\zeta^2}$. The amplitude-banded damping ratios ζ_{high} and ζ_{low} are computed once on the top- and bottom-tercile peak subsequences using the logarithmic-decrement formula $\zeta = \delta/\sqrt{4\pi^2 + \delta^2}$, without resampling; their ~ 0.001 separation is the more meaningful uncertainty signal here than statistical sampling error, since it reflects physical amplitude-dependence rather than peak-detection noise.

To test whether the damping ratio is amplitude-dependent as the quadratic form of Eq. 8 predicts, the positive-peak sequence was partitioned into top-tercile (high-amplitude) and bottom-tercile (low-amplitude) subsequences, and the log-decrement computation was repeated on each band separately.

To verify that the extracted modal parameters are properties of the plant itself rather than artefacts of altitude-dependent forcing, the pitch ring-down was also visualised against altitude (fig. 18). The trajectory traces a monotonically contracting envelope from the 10° release down to small-angle convergence near the ground, with constant peak-to-peak period throughout the descent and no visible coupling between altitude and oscillation frequency. The smooth funnel structure rules out altitude-dependent stiffness or damping artefacts in the coupled-mass solution and confirms that the modal parameters tabulated below are characteristic of the plant rather than of the specific drop trajectory.




Figure 18: Passive pitch ring-down plotted against altitude AGL, colour-mapped by descent time. Released from $\theta_0 = 10^\circ$ at $h = 400$ m, the airframe traces a monotonically contracting envelope from $\pm 19^\circ$ peak (immediately following release transient) down to sub-degree oscillation by touchdown. The constant peak-to-peak spacing across the full altitude range and the smooth envelope contraction confirm the absence of altitude-coupled forcing artefacts, supporting the interpretation of the modal parameters in Table E.1 as intrinsic plant properties.

The amplitude-banded damping ratios $\zeta_{\text{high}} = 0.0134$ and $\zeta_{\text{low}} = 0.0145$ are consistent with the quadratic-velocity dependence of the body-relative damping integral (Eq. 8 of the main text): at larger swing amplitudes, integrated local velocities along the airframe are higher and the quadratic drag produces proportionally greater dimensional resistive torque, which normalises to a slightly lower dimensionless damping ratio. The 0.0011 difference between bands is small but consistent across repeated runs with different initial conditions, ruling out extraction noise as the dominant source. The implied simple-pendulum length of 0.54 m is much shorter than the geometric 2.19 m pivot-to-CoG separation, confirming that the apparent-mass and aerodynamic-stiffness contributions dominate over the gravitational restoring torque, consistent with the canopy-dominated coupled dynamics established in §2.3.

E.2. Flight-Data Validation Metrics

The validation of the plant against 6DOF IMU telemetry from the BristolSEDS Mach24 Blackbird flight [28] uses a 47s window starting 85s after launch, corresponding to the fully-inflated main-chute descent phase. Because the flight vehicle differs from the simulated Waxwing configuration, the 6DOF flight tilt aggregates pitch, roll and yaw, and the two signals are independent stochastic realisations driven by different atmospheric noise sequences, the comparison is framed around statistical and distributional equivalence rather than deterministic agreement.

The flight tilt channel is low-pass filtered at 2Hz to reject IMU noise without attenuating the sub-1 Hz pendulum band of interest. The simulated pitch is taken as $|\theta|$ to match the flight signal’s magnitude form. Both signals are de-trended to zero mean before spectral analysis.

Power spectral densities of both signals were computed using Welch’s method with Hann-windowed segments of

Table 8: Passive modal characterisation. Confidence intervals are reported only where the underlying statistical procedure produced them; entries marked “—” use point estimates from a single fit on the full or banded peak sequence.

Parameter	Value	95% CI	Source
Natural frequency ω_n	4.264 rad/s	[4.245, 4.283]	parametric SE on period
Natural frequency f_n	0.679 Hz	[0.676, 0.682]	parametric SE on period
Damping ratio ζ (high-amp)	0.0134	—	log-decrement, top tercile
Damping ratio ζ (low-amp)	0.0145	—	log-decrement, bottom tercile
Damping ratio ζ (full range)	0.0143	—	log-decrement, all peaks
Logarithmic decrement δ	0.0901	—	full-envelope mean
Time to half-amplitude $\tau_{1/2}$	11.084 s	[10.929, 11.347]	bootstrap on envelope decay rate
2%-envelope settling time t_s	62.559 s	—	$\ln(50)/(\zeta\omega_n)$
Implied simple-pendulum length L	0.54 m	—	g/ω_n^2

length $N = 8 f_s$ samples and 50% overlap. Figure 19 overlays the two spectra on log-log axes spanning the full IMU bandwidth, with the 0.3–2Hz canopy-pendulum band shaded.




Figure 19: Power spectral density comparison of simulated pitch (orange) and flight tilt (blue) across the 0.05–20 Hz band. The shaded region (0.3–2 Hz) marks the canopy-pendulum band of interest. Both signals concentrate energy in this band, the simulation peaking at 0.61 Hz and the flight at 0.37 Hz, with the offset consistent with geometric and configurational differences between the two platforms (§3.2.2). Roll-off rates above 2 Hz are visibly comparable. Below 0.3 Hz both spectra flatten as the detrending procedure removes long-period drift.

Two distinct families of spectral comparison metrics were computed; their interpretation differs and both are reported for completeness.

Distributional metrics: dominant frequency, RMS amplitude, variance ratio, autocorrelation time. These characterise the bulk statistics of each signal independently and compare them. These are the appropriate metrics for validating that two stochastic processes are of the same kind, irrespective of whether their specific time-series realizations are correlated.

Pointwise metrics: magnitude-squared coherence, log-PSD R^2 , Kolmogorov–Smirnov test, peak cross-correlation. These measure deterministic relationships between paired signals. These are the appropriate metrics for validating that one signal predicts another. They are reported here for completeness but are not the primary validation criterion: the simulation and flight signals are independent realisations of similar stochastic processes driven by different atmospheric noise sequences and recorded on

different vehicles, so deterministic agreement is neither expected nor required for the validation argument of Section 3.1.2 to hold. Two independentThe distributional metrics support the Section 3.1.2 validation argument: both signals occupy the same canopy-pendulum frequency band (0.3–1 Hz), with comparable RMS amplitudes, comparable variance, and an autocorrelation-time ratio that matches the 3D-to-1D aggregation effect. The dominant-peak offset, RMS difference, and variance ratio all are more conservative since the simulation represents a harsher disturbance environment than the reference flight, so a controller tuned against it carries margin when deployed against the real-vehicle response.

The distributional metrics support the Section 3.1.2 validation argument: both signals occupy the same canopy-pendulum frequency band (0.3–1 Hz), with comparable RMS amplitudes, comparable variance, and an autocorrelation-time ratio that matches the 3D-to-1D aggregation effect. The dominant-peak offset, RMS difference, and variance ratio all are more conservative since the simulation represents a harsher disturbance environment than the reference flight, so a controller tuned against it carries margin when deployed against the real-vehicle response.

The pointwise metrics confirm what the comparison’s design implies: the two signals are independent stochastic realisations of similar processes, not deterministic predictions of one another. Low coherence and negative log-PSD R^2 indicate that the simulation cannot reproduce the specific flight tilt time series but doing so was never the goal because the simulation is driven by an independent NASA-Langley-validated turbulence noise sequence (§2.4) rather than the actual atmospheric realisation experienced by the reference flight. The K–S test rejects exact distributional equivalence with high confidence, again as expected for two independent stochastic processes from different vehicles. None of these results challenges the validation argument, which rests on the class of dynamics being correctly captured rather than on point-wise prediction.

E.3. Propellant Consumption Statistical Analysis

The statistical comparison of propellant consumption between the 25m and 10m flare altitudes (fig. 13 of the main text) uses two independent Monte Carlo batches: 300 runs of the L controller at a 25m flare and 100 runs of the L controller at a 10 m flare, under identical vehicle, controller and atmospheric parameters with only the flare-trigger altitude changed.

Table 9: Quantitative validation metrics across the 47s comparison window.

Metric	Value	Interpretation
<i>Distributional</i>		
Dominant frequency, simulation	0.610 Hz	Both inside the 0.3–2 Hz canopy-pendulum band; offset attributable to vehicle configuration.
Dominant frequency, flight	0.366 Hz	
Frequency error	66.7%	
RMS amplitude, simulation	15.99°	
RMS amplitude, flight	11.49°	Bounded by geometric difference between platforms. Simulation slightly higher; conservative direction for controller validation.
RMS amplitude error	39.1%	
Variance ratio (sim/flight)	1.94	
Autocorrelation time, simulation	0.31 s	
Autocorrelation time, flight	0.98 s	Within validation band; sim is harsher disturbance. Within factor of 2; conservative direction. Ratio ~3:1 consistent with 3D-to-1D aggregation of three rotational DoF in the flight signal.
<i>Pointwise</i>		
Mean coherence (0.3–2 Hz)	0.103	Low by construction: independent noise realisations.
Log-PSD R^2 (0.1–3 Hz)	−0.65	Spectra diverge in shape; expected for different vehicles.
Peak cross-correlation	0.154	Low; signals are not deterministically related.
Peak cross-correlation lag	9.18 s	Non-physical; confirms independent realisations.
K–S statistic	0.106	Distributional shapes differ;
K–S p -value	$< 10^{-4}$	rejecting deterministic equivalence as expected.

Table 10: Per-batch summary of landing mass and propellant fractions.

Batch	N	Mean [kg]	SD [kg]	Median [kg]	Mass frac.
25 m flare	300	12.63	4.35	12.41	16.1%
10 m flare	100	11.24	3.90	11.02	14.4%
Difference	—	−1.39	—	−1.39	−1.7%

The two-sided Welch’s t -test was selected because the batch sizes and variances differ. The test statistic is

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2/n_1 + s_2^2/n_2}} = 2.99 \quad (65)$$

$$df_{\text{Welch}} \approx 184.4, \quad p = 0.003 \quad (66)$$

confirming statistical significance at the conventional $\alpha = 0.05$ threshold. The 95% confidence interval on the mean difference, computed as $\bar{x}_1 - \bar{x}_2 \pm t_{0.975}, \text{SE}$, is $[0.47, 2.30]$ kg. Cohen’s d effect size: Using the pooled standard deviation

$$d = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{((n_1 - 1)s_1^2 + (n_2 - 1)s_2^2)/(n_1 + n_2 - 2)}} = 0.34 \quad (67)$$

This is a small-to-moderate effect size by Cohen’s conventional thresholds ($d < 0.2$ small, $0.2 \leq d < 0.5$ moderate, $d \geq 0.8$ large). This means the altitude reduction produces a real and statistically significant propellant saving but it is modest in the context of per-run variability ($\sigma \approx 4\text{kg}$) which is dominated by atmospheric disturbance realisation rather than controller choice.

A chi-squared contingency test of the outcome classifications (success / OOS / spin) between the two flare altitudes showed no significant difference ($\chi^2 = 1.84, df = 2, p = 0.40$), confirming that the propellant reduction is achieved without measurable change in success rate.

E.4. Monte Carlo Convergence

The sample size for the primary L-controller batch was justified by explicit convergence analysis of the success-rate

estimator. Starting from the full $N = 300$ run pool, the cumulative success rate $\hat{p}(n) = k(n)/n$ was computed for $n = 1, 2, \dots, 300$, where $k(n)$ is the number of successful landings observed in the first n runs.

The 95% Wilson-score CI was computed at each n rather than the simpler normal-approximation CI, because Wilson-score remains accurate at small n and at proportions far from 0.5. The Wilson interval centre and half-width are

$$\tilde{p} = \frac{k + z^2/2}{n + z^2} \quad (68)$$

$$w = \frac{z}{n + z^2} \sqrt{\frac{z^2}{4} + n\hat{p}(1 - \hat{p})} \quad (69)$$

with $z = 1.96$ for 95

The success rate converges to a final value of $\hat{p}_\infty = 69.7$ (209 of 300 successful landings). The cumulative estimate $\hat{p}(n)$ enters and remains within ± 5 percentage points of $\hat{p}_\infty$ by $n = 31$, indicating that the central tendency is well-resolved with relatively few samples. The 95% CI half-width contracts from approximately $\pm 15\%$ at $n = 50$ to $\pm 7.2\%$ at $n = 150$, $\pm 5.7\%$ at $n = 250$, and converges to $\pm 5.2\%$ at $n = 300$. The convergence behaviour is shown in Figure 20. The early stabilisation of $\hat{p}(n)$ at $n = 31$ should not be interpreted as adequate evidence for a 30-run Monte Carlo batch. Point-estimate stability indicates that the central value is unlikely to drift further with additional samples, but the 95% CI half-width at $n = 31$ remains approximately ± 16 percentage points which is far too wide to support fine-grained claims about controller performance. The 300-run batch is justified not by the need to determine where $\hat{p}$ lies (which is clear by $n = 50$) but by the need to bound it tightly enough to support comparisons such as the L–R chi-squared test of Appendix E.6 and the propellant-comparison Welch’s t -test of Appendix E.3.

The 100-run R-controller batch yields a wider $\pm 9.8\%$ half-width on its own success-rate estimate, which remains sufficient to reject the null hypothesis of equal L/R performance at $p < 0.001$ (Appendix E.6) but does not support




Figure 20: Convergence of the L-controller success-rate estimator. The cumulative estimate $\hat{p}(n)$ (orange) is plotted against run count, with the 95% Wilson-score CI (shaded band) tightening as n grows. The dashed horizontal line marks the final estimate $\hat{p}_\infty = 70\%$; the vertical dotted line marks the stabilisation point at $n = 31$, beyond which $\hat{p}(n)$ remains within $\pm 5\%$ of $\hat{p}_\infty$. The 300-run batch size delivers a final 95% CI half-width of $\pm 5.2\%$, sufficient for the precision claims of §3.1.4.

fine-grained claims about the R-controller success rate to better than 10% precision. The smaller batch was deemed acceptable for the comparison purpose because the L–R difference is much larger than the R-controller’s individual sampling uncertainty (22 percentage points versus a 9.8% half-width), and an additional 200 R-controller runs do not change the qualitative conclusion of 3.1.4.5.

E.5. Failure-Mode Logistic Regression

The failure-predictor analysis of Section 3.1.4.4 fits a logistic regression of binary failure (1 = OOS or spin, 0 = success) against two candidate predictors evaluated at the instant of LQR activation: attitude $|\theta_{trigger}|$ (degrees) and angular rate $|\omega_{trigger}|$ (degrees/s). Trigger-state variables are used in preference to peak values throughout descent to isolate the controller’s *inherited initial condition* rather than its post-activation dynamics. The activation instant is detected behaviourally as the first sustained non-zero gimbal command ($\geq 0.5^\circ$ sustained over 100 ms), avoiding dependence on logged controller flags.

The full 300-run L-controller batch contributes 209 successes (69.7%) and 91 failures (30.3%). Outliers above the 98th percentile of $|\theta_{trigger}|$ or $|\omega_{trigger}|$ are excluded from the regression to prevent leverage effects from a small number of extreme trigger states, leaving 289 runs in the fit.

The logistic model is

$$\log \frac{P(\text{fail})}{1 - P(\text{fail})} = \beta_0 + \beta_\theta |\theta_{trigger}| + \beta_\omega |\omega_{trigger}| \quad (70)$$

fitted by maximum likelihood using MATLAB’s ‘fitglm’ with binomial distribution and logit link.

Raw coefficients (per unit of predictor) and standardised coefficients (per one standard deviation of predictor) are reported below. The standardised form makes the relative predictive importance of the two predictors directly comparable on a common scale.

The likelihood-ratio test against the constant-only null model rejects with $\chi^2(2) = 17.7$, $p = 1.4 \times 10^{-4}$, confirm-

ing the joint predictors carry significant information about failure outcome. The angular-rate coefficient is highly significant individually ($p < 10^{-3}$); the attitude coefficient is not ($p = 0.221$). The standardised coefficients ($\beta_\omega^* = 0.46$ vs $\beta_\theta^* = 0.17$) show that angular rate at trigger dominates attitude by a factor of approximately 2.7 in predictive contribution. Controlling for angular rate removes essentially all of the predictive value of attitude, consistent with the near-horizontal $P = 0.5$ boundary visible in the failure-probability surface (fig. 14).

The training-set ROC AUC is 0.668, the McFadden pseudo- R^2 is 0.052, and the Nagelkerke pseudo- R^2 is 0.086. The Hosmer–Lemeshow goodness-of-fit test returns $\chi^2(8) = 5.32$, $p = 0.72$, indicating no significant lack of fit and that the model is well-calibrated to the observed failure frequencies across the predictor range. The modest AUC and low pseudo- R^2 values together indicate that trigger-state angular rate is the single dominant predictor that can be extracted from the activation snapshot alone, but a substantial fraction of failure variance arises from post-activation atmospheric disturbances (gust sequences, late-arriving turbulence) that the trigger state cannot anticipate. This is consistent with the Section 3.1.4.4 framing that the 30% failure margin is divided between inherited momentum (captured here) and downstream stochastic disturbance (not captured), with this regression bounding the inherited-momentum contribution.

The $P = 0.5$ contour follows

$$|\omega_{trigger}|_{P=0.5} = -\frac{\beta_0 + \beta_\theta |\theta_{trigger}|}{\beta_\omega} \quad (71)$$

which evaluates to $118.5^\circ/\text{s}$ at $|\theta_{trigger}| = 0$ and slopes gently downward to $68.3^\circ/\text{s}$ at $|\theta_{trigger}| = 30^\circ$, reproducing the boundary visualised in Figure 14.

E.6. Controller Architecture Chi-Squared Comparison

The L-controller (300 runs, 25m flare) and R-controller (100 runs, same flare) outcome distributions were tested for equivalence using Pearson’s chi-squared test of independence on the 3-category contingency table.

Table 12: Observed contingency table for landing outcomes.

Controller	Success	OOS	Spin	Total
L controller	210	22	68	300
R controller	48	52	0	100
Total	258	74	68	400

Table 13: Expected counts under the independence null hypothesis, computed as $E_{ij} = (\text{row total})_i \times (\text{column total})_j / (\text{grand total})$

Controller	Success	OOS	Spin
L controller	193.5	55.5	51.0
R controller	64.5	18.5	17.0

Test statistic:

$$\chi^2 = \sum_{ij} \frac{(O_{ij} - E_{ij})^2}{E_{ij}} = 109.18 \quad (72)$$

Table 11: Multiple regression results for landing performance coefficients.

Coefficient	Raw estimate	SE	<i>t</i> -statistic	Std. estimate	<i>p</i> -value
β_0 (intercept)	-1.621	0.251	-6.45	—	1.1×10^{-10}
β_θ (per deg)	0.0229	0.0187	1.22	0.169	0.221
β_ω (per deg/s)	0.01368	0.00405	3.37	0.455	7.4×10^{-4}

$$df = (3 - 1)(2 - 1) = 2, \quad p < 10^{-15} \quad (73)$$

The null hypothesis of identical outcome distributions is rejected at any conventional significance level. The R controller is not merely a weaker version of the L controller: its 22-percentage-point reduction in success rate is redistributed into OOS failures and away from spin failures, confirming the architectural interpretation of Section 3.2.4 that the L/R trade is structurally a Pareto shift of failure modes rather than a uniform scaling of reliability.

All expected counts exceed 5 (minimum 17.0 in the R-spin cell), satisfying the standard rule-of-thumb minimum-cell-size requirement for the chi-squared approximation to be valid. No correction (e.g. Yates' continuity correction or Fisher's exact test) is required.

The peak attitude reached by L-controller spin runs has a median of $\sim 2900^\circ$ and a maximum of $\sim 5400^\circ$. These spin events are not borderline tumbles but multi-revolution catastrophic rotations (8–15 full revolutions before hitting the ground). This separation between OOS (near the 5° boundary) and spin (thousands of degrees) confirms that the three-category outcome classification reflects a genuine physical discontinuity in failure modes rather than an arbitrary thresholding of a continuous distribution. The inset histogram in Figure 15 of the main text visualises this distribution on a log scale.

F.1. Linearisation Validity Analysis

The LQR gain calculation (Appendix D.2) assumes the small-angle approximations $\sin \theta \approx \theta$, $\cos \theta \approx 1$, and rigid-body angular momentum $J\omega$ as a proxy for the full coupled-mass momentum. §3.2.6 of the main text argues that the controller's $P = 0.5$ failure boundary coincides with the regime in which these approximations themselves lose accuracy. This appendix quantifies that claim.

Table 14: Relative error of small-angle approximations for $\sin \theta \approx \theta$ and $\cos \theta \approx 1$.

θ [deg]	$\frac{ \sin \theta - \theta }{ \sin \theta }$	$\frac{ \cos \theta - 1 }{ \cos \theta }$
5°	0.13%	0.38%
10°	0.51%	1.54%
15°	1.15%	3.53%
20°	2.06%	6.42%
25°	3.25%	10.4%
30°	4.72%	15.5%
45°	11.1%	41.4%
60°	20.9%	100%

The sin approximation retains sub-5% error up to $|\theta| \approx 30^\circ$ and sub-20% error to $|\theta| \approx 60^\circ$; the cos approximation degrades considerably faster, exceeding 15% by $|\theta| = 30^\circ$ and reaching 100% relative error at $|\theta| = 60^\circ$ (where $\cos \theta = 0.5$ and the unit-approximation has half the true value). The asymmetry matters because the body-to-inertial transformation in Appendix B.1 weights both terms, so the controller's effective view of gravity-pitch

coupling degrades faster than the pitch-state observation alone would suggest.

The rigid-body angular momentum $J\omega$ underestimates the true coupled-mass moment by a factor of $(1 + a_{44}/I_s + a_{11}L_p^2/I_s) \approx 5.7$ in the canopy-dominated regime, independently of angular rate. At small rates this scaling is absorbed into the LQR gain calibration and has no adverse effect since the controller solves the Riccati equation against a plant model whose moment-arm calibration matches the true plant in the linear regime. The nonlinearity enters through two distinct mechanisms not represented in the linearised A-matrix of Appendix D.2: the $v|v|$ dependence of the body-relative damping integral (Eq. 8), which the controller treats as a linear $-c\omega$ term, and the cross-product structure of the Munk moment $\tau_{munk} = -(a_{33} - a_{11})v_x^B v_z^B$, which grows quadratically in body-frame velocity components and is treated as zero by the linearised model. Both terms are negligible in the small-amplitude regime where $v_x^B \approx 0$ and ω is bounded, but become first-order contributions to the moment balance once the airframe is rotating fast enough that ωL_p dominates the body-frame transverse velocity.

The trigonometric and dynamic nonlinearities both become controller-relevant in the same broad regime of pitch angles above $\sim 30^\circ$ and angular rates above ~ 80 – $100^\circ/\text{s}$, in which the LQR is acting on a substantially incorrect internal model of the plant. The logistic-regression failure boundary identified in Appendix E.5 falls at $|\omega_{trigger}| \approx 119^\circ/\text{s}$ with shallow θ -dependence, sitting comfortably inside the regime where both approximation classes have degraded by tens of percent. The coincidence is qualitative rather than precisely numerical, but it supports the Section 3.2.6 interpretation that the controller's failure envelope is set by the fidelity of its own internal model rather than by actuator saturation or disturbance magnitude. The implication for future work is that an MPC or feedback-linearisation architecture retaining accuracy beyond the small-angle regime would directly address the dominant failure mechanism, whereas atmospheric-robustness improvements would not.

G.1. Propellant Requirement for Pure Powered-Descent Recovery

The factor-of-three propellant saving cited in Section 3.2.5 compares the observed hybrid-architecture propellant consumption (mean 12.63 kg, Appendix E.3) against the propellant required for the same vehicle to execute a precision-controlled powered-descent recovery from the bare-airframe terminal velocity, with no parachute deployed at any altitude. This appendix derives the ballpark 40 kg figure quoted in the main text.

The body drag coefficient $C_{d,body} = 1.2$ used in the plant model (Appendix B.6) is the crossflow value, applicable to a cylinder rotating sideways through the air, and is the appropriate choice for the body-relative damping integral. For a hoverslam analysis the relevant value is the *axial*

Table 15: Operational reserves and propellant multipliers for Powered Descent Guidance (PDG).

Reserve	Multiplier	Justification
Divert margin	1.30	Correction for off-nominal initial conditions, wind drift accumulated during the high-velocity descent, and navigation error; standard PDG practice [10].
Terminal soft-flare	1.10	Final 5–10 m flown at low thrust with active throttle modulation; cannot be captured in constant-thrust analysis.
Throttle-down I_{sp} losses	1.05	Empirical loss of $\sim 5\%$ in effective I_{sp} when the engine is throttled below 30% capacity for terminal-phase precision.
Combined factor	1.50	

drag coefficient, which describes the bluff-body resistance of the airframe descending engine-first along its longitudinal axis. NASA Glenn’s reference data place a typical axial C_d for slender model rockets at approximately 0.75 in nose-first flight (from NASA Glenn reference drag data). Tail-first descent — as required for any hoverslam, since the engine must point downward to brake — adds significant base drag from the bluff aft surface (engine bell, fin roots), raising the effective coefficient to approximately $C_{d,axial} \approx 1.0$, which is adopted as the reference value below.

Reference scenario: Without a canopy, the bare airframe descends under gravity and body drag alone. The drag-balance terminal velocity, with $C_{d,axial} = 1.0$ and $A_{body} = \pi d^2/4 = 0.0314\text{m}^2$, is

$$v_{\text{term}} = \sqrt{\frac{2mg}{\rho C_{d,axial} A_{body}}} \approx 200 \text{ m/s} \quad (74)$$

which the vehicle reaches within approximately 25s of free fall, traversing the first 1.5km of descent. With no parachute deployed at any altitude, this is the velocity the vehicle is descending at when the landing burn must arrest its motion. The PDG comparison is therefore a single controlled hoverslam from v_{term} to soft touchdown. A naive minimum-propellant solution is a maximum-thrust suicide burn at $T = T_{\text{max}}$, which arrests $v_{\text{term}} = 200\text{m/s}$ in 3.7s over 370m for a Tsiolkovsky-minimum $m_p \approx 9\text{kg}$. This figure is misleading as a precision-PDG benchmark for three reasons. First, control authority requires thrust headroom above hover for divert manoeuvres; without it, any horizontal velocity component carried into the burn cannot be cancelled. Second, the engine of Appendix C.2 has a 50:1 throttle range ($T_{\text{min}} = 100\text{N}$ to $T_{\text{max}} = 5000\text{N}$) that supports deep throttling for altitude and velocity precision, and a precision lander uses this range rather than a single thrust setting. Third, the slender high-fineness-ratio airframe with a $\pm 10^\circ$ gimbal envelope has limited lateral authority, so the vertical thrust component must be moderated to keep lateral acceleration capability within useful bounds.

The operational benchmark for this regime is the SpaceX Falcon 9 booster landing burn. Its Merlin 1D engine throttles to 40% of maximum thrust, with a single-engine sea-level thrust at minimum throttle of approximately 338 kN against a typical landing mass in the 25–28t range, giving an end-of-burn T/W of 1.2–1.4 (from public Merlin 1D specifications). Falcon 9 thus operates a controlled hoverslam at the lower end of the deep-throttle envelope precisely to retain divert authority and avoid uncontrolled deceleration. A representative $\bar{\alpha} = 1.4$

is adopted here as the equivalent assumption for the Waxwing PDG comparison.

For a constant-thrust burn arresting v_{term} at average T/W ratio $\bar{\alpha}$, the kinematics give

$$a_{\text{net}} = (\bar{\alpha} - 1)g, \quad t_{\text{burn}} = \frac{v_{\text{term}}}{a_{\text{net}}} \quad (75)$$

$$\Delta v_{\text{thrust}} = \bar{\alpha} g t_{\text{burn}} = v_{\text{term}} \frac{\bar{\alpha}}{\bar{\alpha} - 1} \quad (76)$$

Substituting $v_{\text{term}} = 200 \text{ m/s}$ and $\bar{\alpha} = 1.4$:

$$a_{\text{net}} = 3.92 \text{ m/s}^2, \quad t_{\text{burn}} = 51.0 \text{ s}, \quad \Delta v_{\text{thrust}} = 700 \text{ m/s} \quad (77)$$

of which the gravity-loss component is $g \cdot t_{\text{burn}} = 500 \text{ m/s}$, approximately 71% of the total engine Δv . This is the inherent cost of precision PDG: the same low T/W that buys controllability extends burn duration and accumulates gravity losses. A naive suicide burn avoids this penalty by completing the deceleration in 3s rather than 51s, but at the cost of forfeiting precision control authority entirely.

Tsiolkovsky propellant mass: The base propellant required to deliver $\Delta v_{\text{thrust}} = 700 \text{ m/s}$ with exhaust velocity $v_e = g_0 I_{sp} = 1766 \text{ m/s}$ and dry-landed mass $m_f = 65 \text{ kg}$ is

$$m_{p,\text{base}} = m_f \left(e^{\Delta v_{\text{thrust}}/v_e} - 1 \right) \approx 32 \text{ kg} \quad (78)$$

Applying the combined operational reserve factor:

$$m_{p,\text{total}} = 32 \cdot 1.50 \approx 47 \text{ kg} \quad (79)$$

which is rounded to 40kg as referenced in Section 3.2.5.

The hybrid architecture’s parachute performs the equivalent of the entire 51-second high-gravity-loss braking phase passively, for the fixed 3.5kg parachute hardware mass of Appendix B.3, while the LQR consumes propellant only during the final 25m terminal flare. The comparison is therefore

$$\frac{m_{p,\text{PDG}}}{m_{p,\text{hybrid}}} \approx 3.2 \quad (80)$$

The more nuanced framing is that the hybrid converts a 51-second, gravity-loss-dominated, controlled-descent burn into a 5-second, gravity-loss-bounded terminal-flare burn, and pays for this conversion with 3.5kg of parachute hardware that would otherwise have been some tens of kilograms of propellant.

This is an extremely rough estimate built on multiple assumptions such as the axial drag coefficient, the chosen control T/W, the operational reserve factor, and the

simplification of the descent to a constant-thrust hover-slam and should be read as an order-of-magnitude indication rather than a precision figure. However, since the hybrid's 12.6kg consumption is itself not optimised for minimum fuel burn (the L controller prioritises terminal-state precision and includes the hammer override and large flare margins of Section 2.7), the comparison is fair: both sides carry conservative operational realism, and the resulting factor-of-three-to-four saving reflects an architectural rather than tuning-driven advantage.