



Review Article

From Expendable to Reusable: An Engineering Analysis of Propulsive Landing in SpaceX's Falcon 9

Solmin Kim*

Independent Researcher, Fort Lee High School High school in New Jersey, United States

Publication Information: Accepted: July 29, 2026 **Published:** August 25, 2026

Correspondence to: Solmin Kim, Independent Researcher, Fort Lee High School High school in New Jersey, United States

Citation: Kim S. From Expendable to Reusable: An Engineering Analysis of Propulsive Landing in SpaceX's Falcon 9. American Research Journal of Electronics and Communication Engineering, Vol 3, no. 1, 2026, pp. 1-10.

Copyright © 2026 Kim S. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

The transition from expendable to reusable launch vehicles changed the way engineers think about launch cost and rocket design. This paper examines the propulsive landing technology at the core of SpaceX's first-stage recovery system, analyzing the engineering principles, governing equations, and trade-offs that make booster reusability both technically feasible and commercially viable.

Beginning with the Tsiolkovsky rocket equation and expendable vehicle economics, the paper traces the systematic engineering challenges SpaceX solved to enable vertical propulsive landing. Key technical domains examined—each with their governing mathematical framework—include retropropulsive deceleration, aerodynamic drag and thermal heating during reentry, Guidance Navigation and Control (GNC) through state-space formulation and Kalman filtering, and the hoverslam terminal descent constraint. The paper further develops a quantitative payload-penalty model and a reuse cost model to evaluate the economic impact of propellant reservation and multi-flight amortization.

Quantitative data on launch cost reduction, booster reuse cadence, and payload penalty are presented alongside the theoretical models to evaluate real-world outcomes. The paper concludes that while propulsive landing introduces measurable engineering penalties, the cumulative cost savings over repeated flights—predictable from the derived cost model—substantially change the economics of access to space.

Keywords: Reusable Launch Vehicle; Propulsive Landing; Falcon 9; Tsiolkovsky Equation; Retropropulsion; Drag; Stagnation Heating; GNC; Kalman Filter; RTLS; ASDS; Payload Penalty.

INTRODUCTION

On December 21, 2015, a 14-story rocket booster descended through the Florida night sky, decelerated using its own engines, extended four carbon-fiber landing legs, and settled gently onto a concrete pad at Cape Canaveral. It was the first time an orbital-class rocket booster had ever been successfully recovered in this manner. The feat marked the beginning of a new era in launch vehicle engineering—one in which the most expensive component of a rocket could be flown, recovered, and flown again.

The economic logic is straightforward. A Falcon 9 commercial launch has been publicly priced at roughly

\$67 million, although SpaceX's internal marginal cost is not publicly known; the first-stage booster is estimated to account for roughly 60–70% of vehicle hardware cost and is jettisoned and discarded after a single use in an expendable configuration. The analogy often cited is that of a commercial airline discarding the aircraft after each flight. Reusability offers the prospect of amortizing booster manufacturing costs across dozens of missions, compressing the cost-per-launch to a fraction of its expendable value.

Yet propulsive landing is not simply an economic decision—it is a difficult engineering problem governed by well-defined physical laws. The booster must survive

atmospheric reentry, performing precisely timed engine burns in a degraded aerodynamic environment, navigate to a landing target that may be hundreds of kilometers downrange, and execute a final braking maneuver to touch down at near-zero velocity. This demands mastery across propulsion, aerodynamics, thermodynamics, structural mechanics, and real-time guidance and control—each domain with its own mathematical framework [1].

This paper examines the engineering foundations of Falcon 9's propulsive landing system with attention to the governing equations in each domain. Section II develops the Tsiolkovsky rocket equation and its implications for expendable vehicles. Section III reviews early reusability programs. Sections IV and V detail the Falcon 9 architecture and the step-by-step mechanics of booster recovery, including thrust-to-weight analysis, aerodynamic drag, stagnation heating, the hoverslam constraint, and GNC state-space formulation. Section VI addresses thermal and structural challenges. Section VII develops quantitative payload-penalty and reuse cost models. Section VIII evaluates economic outcomes and future implications.

The thesis of this paper is that Falcon 9's propulsive landing system succeeds not because it solves a single difficult problem, but because it integrates physically rigorous solutions across propulsion, aerodynamics, and control in a manner that makes the aggregate engineering trade-offs economically favorable on scale.

THE ERA OF EXPENDABLE ROCKETS

For most of the first six decades of the Space Age, orbital launch systems were either fully expendable or only partially reusable at very high refurbishment cost. This paradigm was inherited from ballistic missile development, in which single-use hardware was the only viable design approach given to the technology of the era.

The Tsiolkovsky Rocket Equation

The fundamental governing equation of rocket propulsion is the Tsiolkovsky rocket equation [2], derived from conservation of momentum for a variable-mass system. It relates the velocity increment a rocket can achieve to its propellant mass fraction and engine efficiency:

$$\Delta v = I_{sp} \cdot g_0 \cdot \ln\left(\frac{m_0}{m_f}\right)$$

where Δv = velocity change (m/s), I_{sp} = specific impulse (s), $g_0 = 9.81 \text{ m/s}^2$, m_0/m_f = mass ratio

Eq. 1 Tsiolkovsky Rocket Equation — the foundational equation of rocket propulsion

The specific impulse I_{sp} is a measure of engine efficiency,

defined as thrust produced per unit weight flow of propellant. Measured in seconds, it is engine- and propellant-combination-dependent: the Merlin 1D engine burns RP-1 (kerosene) and liquid oxygen, achieving $I_{sp} \approx 282 \text{ s}$ at sea level and $\approx 311 \text{ s}$ in vacuum [3, 4]. The natural logarithm of the mass ratio (m_0/m_f) means that to achieve the Δv required for LEO orbit (approximately 9,400 m/s including gravity and drag losses), a rocket must expend an enormous fraction of its initial mass as propellant.

Applying the Tsiolkovsky equation to the Falcon 9 Block 5 first stage with published parameters:

$$\Delta v_{stage1} = 311 \text{ s} \times 9.81 \text{ m/s}^2 \times \ln\left(\frac{549,054 \text{ kg}}{25,600 \text{ kg}}\right) \approx 9,360 \text{ m/s}$$

$$\text{Falcon 9 Block 5: } I_{sp}(\text{vac}) \approx 311 \text{ s}, m_0 = 549,054 \text{ kg (full)}, m_f = 25,600 \text{ kg (dry first stage)}$$

Eq. 2 Tsiolkovsky equation applied to Falcon 9 Block 5 first stage — numerical example

This idealized calculation shows the theoretical velocity capability of the stage under simplified assumptions. It should not be interpreted as the actual first-stage contribution to orbital velocity, because real launch performance is divided between stages and reduced by gravity losses, drag losses, staging conditions, throttle profiles, and recovery propellant reserves.

Mass Ratio and the Expendable Cost Penalty


Rearranging the rocket equation, the structural mass fraction $\epsilon = m_f/m_0$ of a typical first stage is approximately 0.047 (the stage is $\sim 95.3\%$ propellant by mass at liftoff). The payload ratio $\lambda = m_L/m_0$, where m_L is the useful payload mass, is even smaller—typically 2–4% of gross liftoff weight for an expendable configuration. This extreme mass sensitivity means that every kilogram of hardware recovered for reuse, rather than discarded, directly reduces the manufacturing cost embedded in each launch.

For an expendable vehicle flying $n = 1$ missions per booster, the full manufacturing cost C_{mfg} is charged to that single mission. The cost-per-flight is simply C_{mfg} . For a reusable vehicle flying n missions, as developed in Section VII, the manufacturing cost is amortized across all flights. This is the fundamental economic leverage of reusability—and it is the reason the rocket equation makes reusability so attractive despite the engineering complexity it introduces.

EARLY ATTEMPTS AT REUSABILITY: DC-X AND THE SPACE SHUTTLE

The Space Shuttle: Partial Reuse at High Cost

The Space Shuttle (1981–2011) was the first operational



system designed with reusability as a central goal. Its Orbiter, main engines, and solid rocket boosters were intended to be refurbished and re-flown, reducing permission costs. In practice, the refurbishment process proved enormously complex. The Orbiter’s 24,000 thermal protection tiles required individual inspection and frequent replacement. The Space Shuttle Main Engines required complete disassembly and recertification after each flight. The net result was an average cost per mission of approximately \$1.5 billion—roughly \$54,500 per kilogram to LEO—equivalent in cost to an expendable system despite the added complexity of recovery [5].

The fundamental flaw was not the concept of reuse, but the failure to control refurbishment labor cost at the design stage. The Shuttle demonstrated the principle of reuse but not its economics.

The key difference between the Shuttle and Falcon 9 was not simply that one reused hardware and the other did not. The Shuttle reused the most complex part of the system—the crewed orbiter—and required large teams to inspect

and refurbish it. Falcon 9 reused the first stage, which was expensive but comparatively simpler to inspect, especially after the Block 5 design changes. This made the economics of reuse more practical.

DC-X: Proving Propulsive Vertical Landing

A more direct technological precursor to Falcon 9’s landing system is the McDonnell Douglas DC-X (Delta Clipper Experimental), developed from 1991 to 1996. The DC-X was a subscale prototype—approximately 12 meters tall, designed to demonstrate vertical takeoff and vertical landing (VTVL). Between 1993 and 1996, it conducted twelve test flights, demonstrating throttleable engines, in-flight attitude control, and vertical landings from altitudes up to 2,500 meters [6,7].

The key technical lessons from DC-X that informed later development included: (1) the viability of deep-throttle engine operation for terminal descent; (2) the importance of redundant landing leg deployment mechanisms; and (3) the need for robust attitude control authority at low

Table 1: Comparative overview of launch vehicles and reusability programs.

Program	Period	Recovery Method	Outcome	Cost/kg LEO
Saturn V	1967–73	Expendable	Successful (lunar)	~\$54,500
Space Shuttle	1981–2011	Partial reuse	Expensive refurb	~\$54,500
DC-X	1993–1996	VTVL subscale	Proof of concept	N/A
Falcon 9 v1.0	2010–2013	Expendable	Operational	~\$6,700
Falcon 9 Block 5	2018–pres.	Propulsive VTVL	Operationally mature	~\$1,500

dynamic pressure. The DC-X proved VTVL was physically achievable; SpaceX’s engineering task was to scale and integrate this capability into an orbital-class vehicle [8,9].

FALCON 9 ARCHITECTURE OVERVIEW

The Falcon 9 is a two-stage, liquid-fueled orbital launch vehicle standing 70 m tall with a 3.7 m diameter. The Block 5 configuration can deliver approximately 22.8 metric tons to Low Earth Orbit in expendable configuration [10].

First Stage

The first stage is powered by nine Merlin 1D engines in an octa web arrangement, producing approximately 7,607 kN of thrust at sea level. The engines burn RP-1 and liquid oxygen (LOX), with a vacuum Isp of approximately 311 s [10]. The first stage carries approximately 394 metric tons of propellant at liftoff (RP-1: ~147 t, LOX: ~247 t) against a dry mass of approximately 25,600 kg—a propellant mass fraction of approximately 93.9%, consistent with the rocket equation analysis of Section II.

A propellant reserve of approximately 10–15% of initial load is retained at stage separation specifically to enable the recovery maneuvers. This reserve is the source of the payload penalty analyzed in Section VII.

What Is Recovered and Why

SpaceX recovers only the first stage, not the second stage or fairings (except for separate fairing recovery operations). The first stage, comprising nine Merlin engines and the main propellant tanks, accounts for approximately 60–

70% of total vehicle hardware cost. The second stage must complete orbital insertion and is left in a low-perigee orbit to naturally deorbit, making its recovery geometrically and thermally prohibitive relative to its fractional cost.

Block 5 Reusability Improvements

The Block 5 variant (2018) incorporated targeted engineering changes for maximum reusability: titanium grid fins replacing earlier aluminum versions for better thermal resistance; Merlin engine turbopumps rated for ten or more reflights without overhaul; a PICA-X thermal coating on the interstage; and rapid-connect avionics interfaces for faster turnaround. Turnaround time between flights was reduced from approximately 3–4 months (early flights) to as low as 21 days for mature boosters.

ENGINEERING PRINCIPLES OF PROPULSIVE LANDING

The propulsive landing sequence is a multi-phase, multi-discipline process. Each phase is governed by distinct physical equations, and success depends on integrating these across propulsion, aerodynamics, and control. The full sequence—from stage separation to touchdown—takes approximately eight to ten minutes.

Boostback Burn and Thrust-to-Weight Ratio

For Return-to-Launch-Site (RTLS) missions, shortly after stage separation the booster performs a “boostback burn” using three engines to reverse its downrange velocity. The key governing parameter for any propulsive maneuver is the Thrust-to-Weight Ratio (TWR) [4]:




Figure 2: Schematic flight profile of Falcon 9 first stage from MECO to touchdown. Phase boundaries correspond to distinct engineering regimes analyzed in Sections 5.1–5.5.

$$TWR = \frac{F_{thrust}}{m \cdot g_0}$$

Merlin 1D (single, min throttle, landing): $F \approx 360 \text{ kN}$, $m_{dry} \approx 25,600 \text{ kg} \rightarrow TWR \approx 1.43$

Eq. 3 Thrust-to-Weight Ratio — governs acceleration capability at each phase of recovery

At liftoff, all nine Merlins produce 7,607 kN against a vehicle weight of approximately 5,385 kN ($549,054 \text{ kg} \times 9.81 \text{ m/s}^2$), giving $TWR \approx 1.41$ —barely sufficient to lift off, which is by design to minimize gravity losses. During the boostback burn with three engines and a significantly lighter (post-MECO) first stage, TWR rises substantially, enabling rapid velocity reversal. The boost-back burn lasts approximately 30 seconds and requires a 180° pitch maneuver executed by cold-gas nitrogen attitude thrusters before engine ignition [11].

Entry Burn, Aerodynamic Drag, and Stagnation Heating

As the first stage reenters the upper atmosphere, two physical phenomena become critical: aerodynamic drag and aerothermal heating. The drag force acting on the descending booster is given by [4]:

$$F_D = \frac{1}{2} \rho v^2 C_D A$$

ρ = air density (kg/m^3), v = velocity (m/s), C_D = drag coefficient, A = reference area (m^2)

Eq. 4 Aerodynamic drag equation — governs deceleration force during atmospheric reentry

As atmospheric density ρ increases with decreasing altitude and velocity v remains high (approximately 1,500–2,000 m/s at entry interface), the drag force increases rapidly. Without active deceleration, peak dynamic pressure ($q = \frac{1}{2}\rho v^2$) would impose structural loads exceeding the booster's design limits. The entry burn—firing three Merlin engines for approximately 20–25 seconds at $\sim 70 \text{ km}$ altitude—reduces velocity before the atmosphere becomes dense enough to cause structural damage, trading propellant expenditure for reduced structural and thermal loads.

The thermal environment during reentry is governed by aerodynamic heating. The stagnation point heat flux—the most severe heating on a blunt body—is approximated by the Fay-Riddell relationship [4]:

$$\dot{q}_{stag} \approx k \cdot \sqrt{\frac{\rho_\infty}{R_n}} \cdot v_\infty^3$$

k = empirical constant, ρ_∞ = freestream density, R_n = nose radius, v_∞ = freestream velocity

Heating scales with v^3 — tripling velocity increases stagnation heat flux 27×

Eq. 5 Fay-Riddell stagnation heating approximation — governs thermal protection system requirements

The critical insight from this equation is the cubic dependence on velocity: tripling the freestream velocity v_∞ increases stagnation heat flux by a factor of 27. This is why the entry burn—which reduces velocity from $\sim 2,000 \text{ m/s}$ to $\sim 700 \text{ m/s}$ before significant atmospheric density—has such a disproportionate effect on thermal environment. A 65% velocity reduction corresponds to approximately a 96% reduction in peak heat flux, dramatically reducing thermal protection requirements. The remaining heating is managed by PICA-X ablative coating on the interstage (Section VI).

Grid Fin Aerodynamic Control

Grid fins are the primary aerodynamic control effector during the high-dynamic-pressure phase of descent, when engine thrust is absent. Mounted at the top of the first stage, the four independently actuated grid fins generate lift and drag forces through deflection, providing pitch, yaw, and roll control authority. Their lattice structure offers comparable control effectiveness to planar fins at a fraction of the structural span, enabling compact stowage during ascent [12].

The aerodynamic control moment generated by each fin is a function of the same drag equation framework (Eq. 4), applied to the fin's reference area and adjusted for deflection angle. The net control moment is the vector sum of contributions from all four fins, commanded by the GNC algorithm described in Section 5.5. Block 5 titanium grid fins withstand reentry heating without refurbishment; earlier aluminum fins experienced significant erosion at the high velocities of initial recovery flights.

Landing Burn: The Hoverslam Constraint

The terminal descent phase—known colloquially as the “hoverslam”—is the most demanding phase of recovery. The Merlin 1D engine at its minimum throttle setting (40% of rated thrust) must decelerate the booster from approximately 250 m/s to near-zero velocity at touchdown. The critical engineering constraint is revealed by comparing minimum engine thrust to vehicle weight at landing:

$$F_{engine,min} = 0.40 \times F_{rated} = 0.40 \times 934 \text{ kN} = 373 \text{ kN}$$

$$W_{booster,dry} = m_{dry} \cdot g_0 = 25,600 \text{ kg} \times 9.81 \text{ m/s}^2 = 251 \text{ kN}$$

Since $F_{engine,min}$ (373 kN)
 $> W_{booster}$ (251 kN), the booster cannot hover
 — it must execute a continuous “hoverslam”

Eq. 6 Hoverslam constraint — minimum Merlin thrust exceeds booster dry weight, precluding hover

Because minimum engine thrust (373 kN) exceeds booster dry weight (≈ 251 kN), the vehicle cannot decelerate to zero and maintain altitude cannot hover. Instead, the GNC system must compute a trajectory that reaches exactly zero velocity at exactly the moment of ground contact. The landing burn is timed to begin at an altitude where the available deceleration profile precisely zeros out velocity at touchdown. Any deviation slightly early start wastes altitude; a late start risks contact before deceleration completes—must be handled in real time by the GNC system.

This “hoverslam” constraint is a direct consequence of the Merlin engine’s thrust-to-weight characteristics and the near-empty booster mass. It places extremely tight requirements on the GNC algorithm, navigation accuracy, and engine throttle response time.

Guidance, Navigation and Control (GNC)

The GNC system integrates navigation sensors, state estimation, and trajectory optimization to guide the booster from stage separation to touchdown. The vehicle’s motion is described by a state-space model:

$\dot{x} = Ax + Bu$

State vector x : [position, velocity, attitude, angular rate]
– Control input u : [thrust magnitude, gimbal angle]

Eq. 7 State-space formulation of Falcon 9 booster dynamics — foundation of the GNC framework

The state vector x contains the booster’s position, velocity, attitude (quaternion), and angular rate—twelve state variables in three-dimensional space. The control input vector u specifies the thrust magnitude and engine gimbal angles (which vector the thrust direction). The matrix A captures the natural dynamics (including gravity and aerodynamic forces as functions of state), and B captures how control inputs translate to state derivatives.

Because the true state cannot be measured directly—GPS provides position, IMUs provide acceleration and angular rate, and radar altimeters provide altitude—a Kalman filter is used to fuse noisy sensor measurements into a

best-estimate of the vehicle state [9,13]:
 $\hat{x}_k = \hat{x}_k^- + K_k(z_k - H\hat{x}_k^-)$
 $K_k = P_k^- H^T (H P_k^- H^T + R)^{-1}$ (Kalman Gain)
 z_k : sensor measurement (GPS + IMU), H : observation matrix, R : measurement noise covariance

Eq. 8 Kalman filter state update — fuses GPS, IMU, and radar altimeter measurements in real time

The Kalman filter operates in two steps: prediction (propagating the state estimate forward using the dynamics model) and update (correcting the prediction using the latest sensor measurement z_k , weighted by the Kalman gain K_k). The gain K_k is computed from the estimated state covariance P_k and the measurement noise covariance R , optimally balancing model prediction against sensor measurement. This produces a minimum-variance state estimate at each timestep, running at several hundred hertz on the vehicle’s flight computers.

Given the estimated state, the trajectory optimization problem—finding the minimum-propellant path to a zero-velocity touchdown at the landing target—is solved using convex optimization (specifically, a technique called Successive Convexification) [8,14]. This allows real-time re-planning of the trajectory in response to wind gusts, sensor errors, or engine performance deviations, generating updated thrust commands at a rate of several hertz throughout the descent.

THERMAL AND STRUCTURAL CHALLENGES
Aerodynamic Heating and Thermal Protection

The stagnation heating analysis of Section 5.2 (Eq. 5) establishes that peak heat flux scales with v^3 . The Falcon 9 interstage and engine section are coated with PICA-X (Phenolic Impregnated Carbon Ablator, SpaceX-developed derivative of NASA’s Apollo-era PICA). Ablative materials manage heating by charring and pyrolyzing on the outer surface, absorbing energy through the ablation process rather than conducting it to the structure. The char layer also has low thermal conductivity, providing additional insulation.

Table 2: Structural and refurbishment comparison between early and Block 5 Falcon 9 boosters.

Parameter	Block 3/4 (2015–17)	Block 5 (2018–present)
Grid fin material	Aluminum	Titanium (solid)
Grid fin refurb between flights	Required	Not required
Min. certified reuse flights	~2–3	10+
Record reuse flights (2026)	~8	35 (as of 2026)
Merlin overhaul interval	Each flight	Every 10+ flights
Min. turnaround time	~3–4 months	21 days

Grid fin thermal management illustrates the engineering iteration in the program. Early aluminum grid fins (melting point 660°C) experienced significant ablation at entry velocities of 1,500+ m/s. Block 5's solid titanium grid fins (melting point 1,668°C, yield strength maintained to ~500°C) survive reentry without refurbishment, a design change driven directly by the empirical evidence from early recovery flights.

Structural Fatigue and Multi-Flight Certification

Structural fatigue—the progressive degradation of material strength under cyclic loading—is a principal engineering concern for a reusable booster. Each flight subjects the vehicle to propellant pressurization loads, acoustic and vibration load during ascent, aerodynamic and inertial loads during reentry, and impact loads at touchdown. Over multiple flights, cyclic stresses can initiate and propagate cracks in structural members, particularly at welds and stress concentrations.

SpaceX certifies boosters for repeated flight through a combination of design margin (structural members sized to withstand loads substantially exceeding expected maxima), non-destructive inspection between flights (ultrasonic testing, X-ray weld evaluation, visual inspection of thermal protection), and a conservatively expanding flight-count envelope validated by empirical data. As of mid-2026, the record for most flights by a single booster stands at 35, reported by SpaceX and independently tracked launch manifests; Block 5 was initially certified for a minimum of 10 flights without major refurbishment.

ENGINEERING TRADE-OFFS

Propulsive landing is not a free engineering benefit. The decision to recover the first stage introduces measurable penalties in payload capacity, adds system complexity, and imposes operational constraints. A rigorous engineering analysis must account for these trade-offs quantitatively.

Payload Penalty Model

Falcon 9 recovers the first stage in one of two ways. In a Return-to-Launch-Site (RTLS) recovery, the booster reverses course with a boostback burn and flies back to a landing pad near the launch site, which costs a large amount of propellant. In an Autonomous Spaceport Drone Ship (ASDS) recovery, the booster lands on a drone ship downrange instead, needing little or no boostback burn. Because ASDS uses less propellant for recovery, it leaves more propellant, and therefore more payload capacity, available for the mission.

The payload penalty of reuse arises from two sources: (1) the structural mass of recovery hardware (landing legs, grid fins, actuation systems), and (2) the propellant reserved for recovery burns that is not available for orbital velocity increment. The second source dominates. A simplified model of the payload penalty attributable to propellant reservation is:

$$\Delta m_L \approx m_{\text{reserve}} \cdot \frac{I_{sp} g_0}{\Delta v_{\text{orbit}}}$$

RTLS penalty: $m_{\text{reserve}} \approx 35 \text{ t} \rightarrow \Delta m_L \approx 10 \text{ t}$ ($\approx 44\%$ of 22.8 t expendable capacity)

Larger propellant reserve reduces available Δv for orbit, directly reducing payload mass delivered

Eq. 9 Payload penalty model — relates propellant reserve mass to reduction in orbital payload capacity

The model follows from the Tsiolkovsky equation: each kilogram of propellant diverted to recovery burns reduces the Δv available for orbit, which translates to a reduction in deliverable payload mass. For RTLS recovery, the boostback burn requires approximately 35 metric tons of propellant reserve, producing a payload reduction of approximately 10 metric tons—consistent with the observed RTLS payload capacity of 12.8 t versus the expendable baseline of 22.8 t (a 44% reduction) [10]. ASDS recovery requires only a landing burn reserve of approximately 10–15 metric tons (no boost-back burn), reducing the penalty to approximately 24–32%.

The Failure Record and Engineering Learning

Between January 2015 and January 2016, SpaceX attempted seven booster recovery landings and achieved only two successes. The failure modes were highly informative engineering data:

CRS-5 (Jan. 2015): Insufficient hydraulic fluid for grid fin actuators — excessive lateral velocity at touchdown.

CRS-6 (Apr. 2015): Single sticking throttle valve caused thrust imbalance in terminal descent — booster oscillated and fell.

Jason-3 (Jan. 2016): Landing leg latch failure — booster landed, then collapsed.

Each failure yielded diagnostic data driving specific engineering corrections. The rapid iteration—accepting early losses as test data—is itself an engineering methodology. The failure-to-success trajectory illustrates that propulsive landing is a system integration challenge: any sub-system failure at the margin of performance can cause total loss. As of 2026, the landing success rate on attempted recoveries is reported to be around 97–98%—




Figure 3: Falcon 9 payload capacity to LEO by recovery mode. RTLS imposes the largest penalty consistent with the payload penalty model (Eq. 9). ASDS recovery balances payload performance against recovery feasibility for heavier missions.




Figure 4: Falcon 9 annual launch cadence and booster reuse rate, 2013–2024. Acceleration of reuse post-2018 (Block 5) reflects engineering maturation and the turnaround improvements quantified in Table 2. Compiled from publicly available SpaceX launch records.

a figure reflecting both engineering maturity and the statistical confidence of hundreds of flights.

Reuse Cost Model

The economic viability of reuse depends on whether the amortized manufacturing cost plus per-flight refurbishment cost is less than the cost of a new booster [5]. A simple model for the effective per-flight booster cost as a function of reuse count n is:

$$C_{reuse}(n) = \frac{C_{mfg}}{n} + C_{refurb}$$

$n = 1$ (expendable): $C = \$32M + \$0 = \$32M$ $n = 10$: $C = \$3.2M + \$4M = \$7.2M$

$C_{mfg} \approx \$32M$ (booster manufacturing cost), $C_{refurb} \approx \$3 - 5M$ per flight (Block 5)

Eq. 10 Reuse cost model — per-flight booster cost as a function of reuse count and refurbishment cost

With estimated booster manufacturing cost $C_{mfg} \approx \$32$ million and Block 5 refurbishment cost $C_{refurb} \approx \$3$ – 5 million per flight, the model predicts: at $n = 1$ (expendable), cost = $\$32M$; at $n = 5$, cost $\approx \$11.4M$; at $n = 10$, cost $\approx \$7.2M$; at $n = 20$, cost $\approx \$4.6M$. The diminishing returns are evident—the largest cost reduction occurs in the first few reflights, after which refurbishment cost dominates. This model explains why SpaceX’s commercial strategy focuses on high flight rates per booster and why Block 5 refurbishment cost reduction was a critical engineering objective.

ECONOMIC AND FUTURE IMPLICATIONS

Quantifying the Economic Impact

The reuse cost model (Eq. 10) provides a framework for evaluating the economic impact of the program. Industry analysts triangulating from published pricing, launch




Figure 5: Estimated cumulative cost savings per booster as a function of reuse count, consistent with the cost model (Eq. 10). Savings approach a maximum of nearly 55% at high reuse counts as refurbishment cost becomes the binding constraint.

cadence, and component cost data estimate that a Block 5 booster flying 10 missions generates approximately \$25–30 million in cumulative manufacturing cost savings over its operational life—a significant fraction of the publicly listed \$67 million commercial mission price [5].

The broader market impact has been structural. Prior to Falcon 9, the international launch market was characterized by prices of \$100–200 million per mission for comparable payload classes (Ariane 5, Atlas V, Proton). Falcon 9’s pricing pressure forced a systemic reassessment of competitor economics. Arianespace accelerated development of the Ariane 6; ULA pursued the Vulcan Centaur with engine recovery concepts; national programs in Europe, Japan, and India intensified their own reusability research. The result was not only a better rocket, but also a major change in the launch market.

Implications for the Starship Program

Falcon 9’s propulsive landing system is a direct technological predecessor to the Starship program—SpaceX’s next-generation launch system designed for full reusability of both stages. The Super Heavy booster employs scaled-up Raptor engines (a more advanced full-flow staged combustion cycle burning liquid methane and LOX, $I_{sp} \approx 363$ s in vacuum) and a novel “mechazilla” catch system in which the launch tower’s arms catch the descending booster mid-air, eliminating landing mass entirely.

The engineering lessons from Falcon 9, particularly in GNC algorithm development (Eqs. 7–8), thermal protection adequacy (Eq. 5), and engine deep-throttle capability (Eq. 3/6)—directly informed Starship’s design. The critical

unproven element in Starship is upper-stage reentry: the Ship must survive reentry from full orbital velocities (approximately 7,800 m/s) using a belly-flop orientation and active flap control, then flip to vertical for a landing burn. The stagnation heating at these velocities—scaling with v^3 per Eq. 5—is approximately 60 times greater than Falcon 9 first-stage reentry conditions, representing the primary thermal engineering frontier as of 2025.

Broader Implications for Space Access

The cost reduction enabled by Falcon 9’s reusability has had measurable second-order effects. Lower launch costs made large satellite constellations economically viable; SpaceX’s Starlink constellation, which has grown to over 10,000 satellites as of 2026, would not be economically feasible at pre-reuse launch prices. The proliferation of small satellites has, in turn, enabled a new generation of Earth observation, communications, and scientific missions previously inaccessible to all but the largest government agencies. The impact of propulsive landing has thus gone beyond cost savings alone, changing what is economically achievable in space.

CONCLUSION

This paper examined the engineering foundations of SpaceX’s Falcon 9 propulsive landing system through the governing physical equations of each technical domain. The analysis spans six distinct mathematical frameworks: the Tsiolkovsky rocket equation establishing the mass-ratio imperatives of reusability; the thrust-to-weight ratio governing maneuver capability; aerodynamic drag equations describing the reentry deceleration

environment; the Fay-Riddell stagnation heating relation quantifying thermal protection requirements; the hoverslam constraint revealing the fundamental tension in terminal descent; and the state-space GNC formulation with Kalman filter fusion enabling real-time trajectory optimization. Two engineering economics models—payload penalty and reuse cost—connect the physical trade-offs to quantitative economic outcomes.

The central trade-off is well-characterized by these equations: propulsive landing imposes a payload penalty of 24–44% depending on recovery mode (Eq. 9), adds complexity through recovery hardware and GNC systems, and requires propellant reserves that subtract from mission performance. Against this, it enables amortization of first-stage manufacturing costs across ten or more flights, reducing effective per-flight booster cost by an estimated 55%+ at high reuse counts (Eq. 10). The engineering feasibility of this trade is not merely theoretical—it has been demonstrated empirically across hundreds of recovery operations since 2016.

Several engineering frontiers remain open. The structural fatigue limits of boosters beyond 20+ flights are not yet fully characterized, and the Starship upper-stage reentry problem—with stagnation heating approximately 60 times more severe than Falcon 9 recovery conditions per the v^3 heating relation—represents the next major thermal engineering challenge. The economics of full reusability at Starship scale, with ambitions for sub-\$10 million per-launch costs, remain to be validated in operation.

For the engineering student, Falcon 9's recovery system is a compelling demonstration that real-world aerospace engineering is not the application of any single equation but the simultaneous satisfaction of multiple governing relationships across fluid mechanics, thermodynamics, structural mechanics, and control theory. The booster that lands itself is the physical expression of this integration—and the cost savings it produces are a direct, quantifiable consequence of the engineering principles examined in this paper.

References

1. Wertz, J. R., Everett, D. F., & Puschell, J. J. (Eds.). (2011). *Space mission engineering: The new SMAD*. Microcosm Press.
2. Tsiolkovsky, K. E. (1903). The exploration of cosmic space by means of reaction devices. *Scientific Review*, 5. (Translated in NASA TT F-237, 1967.)
3. Braeunig, R. A. (2012). *Rocket propulsion basics*. Rocket and Space Technology.
4. Cantwell, B. (2021). *Rocket propulsion* (AA283 course notes). Stanford University, Department of Aeronautics and Astronautics.
5. Fay, J. A., & Riddell, F. R. (1958). Theory of stagnation point heat transfer in dissociated air. *Journal of the Aeronautical Sciences*, 25(2), 73–85.
6. Kang, S.-Y., Jo, M.-S., Choi, J.-Y., & Yang, S. S. (2025). Cost effectiveness of reusable launch vehicles depending on the payload capacity. *Aerospace*, 12(5), 364.
7. Voglewede, S. D. (1994). Single Stage Rocket Technology's real time data system. In NASA Goddard Space Flight Center, Third International Symposium on Space Mission Operations and Ground Data Systems, Part 1.
8. NASA. (2020). *Space Shuttle program cost and schedule*. NASA Office of Inspector General, IG-20-012.
9. Rabinovitch, J., Cotting, C., & D'Souza, S. (2019). Planetary landing guidance using convex optimization. *AIAA SciTech Forum*, AIAA 2019-0667.
10. Simon, D. (2006). *Optimal state estimation: Kalman, H-infinity, and nonlinear approaches*. John Wiley & Sons.
11. SpaceX. (2024). *Falcon 9 launch vehicle payload user's guide* (Rev. D). SpaceX Publications.
12. Urie, D. M., & Johnson, C. R. (1995). *DC-X program summary report*. McDonnell Douglas Aerospace / BMDO.
13. Welch, G., & Bishop, G. (2006). *An introduction to the Kalman filter* (Tech. Rep. TR 95-041). University of North Carolina at Chapel Hill.
14. Acikmese, B., & Ploen, S. R. (2007). Convex programming approach to powered descent guidance for Mars landing. *Journal of Guidance, Control, and Dynamics*, 30(5), 1353–1366.