Our Gods Haven't Crashed, Yet

A Silicon Junkies' Riddle — Our Highway

Version 1.7

Tony O'Connor*

Critical Infrastructure Protection Specialist

Peer Reviewed by: Grok (xAI), Claude (Anthropic), ChatGPT (OpenAI), Copilot (Microsoft)

December 7, 2025

DOI: 10.5281/zenodo.17851041

"I saw Satan fall like lightning from heaven."

— Luke 10:18

"Professing themselves to be wise, they became fools."

— Romans 1:22

Abstract

Building upon the orbital debris risks explored in "Our Gods Haven't Fallen, Yet — A Space Junkies' Riddle — Our Cathedral" [O'Connor, 2025], this paper segues to terrestrial analogs in autonomous vehicle (AV) fleets, where correlated failures echo Kessler Syndrome's N^2 scaling. Through a triune lens—*Breath* (narrative), *Logos* (analytical), and *Word* (executive)—we dissect how Tesla's Full Self-Driving (FSD) system, with shared firmware across ∼960k–1.6M vehicles today scaling to a scenario of ∼10M by 2030, could birth chain-reaction crashes and 'trust debris'—regulatory exiles and societal shuns persisting hours physically but years in confidence erosion. Governance gaps rely on voluntary frameworks; recovery timescales span years to decades if unaddressed. Technology for sustainable AV operations exists; intention—the will to implement known mitigations—remains the gap to prevent exile.

Keywords: Autonomous vehicles, correlated risks, trust debris, N^2 scaling, AV governance, fleet sustainability

Contents

*Contact: @tonyoconnor_

Preface

From OGHFY's void cathedral to our garden of interconnected highways, traveler—where silicon prayers now hum along asphalt amid a quieter roar of engines and a whisper of code, receipts we mustn't ignore in the twilight glow of taillights. This terrestrial riddle echoes the orbital, segueing N^2 cascades from space junk to synced glitches: three doors to one corridor, each entry grasping risks before they erode trust and exile us to manual drudgery. Pick your path—*Breath* for the tale, *Logos* for the proof, *Word* for the pitch—but see they weave like lanes under rush-hour sky, guarding the path we must not wreck!

Breath: The narrative breath, a storyteller's wind where facts wear wheels and analogies accelerate through neon-lit merges. For general readers seeking the human pulse—vivid drives through glitch storms, pulling bugs from code like cats from boxes, recounting the 2025 phantom brakes that chained highways in chaotic symphonies.

Logos: The analytical core, the ledger of logic where equations etch failure and governance gaps grind wide like potholes in neglected pavement. For PhDs and experts craving rigor—N^2 models, breakdowns of 8M+ vehicles, failure rates (1 crash/6.36M miles for Autopilot vs. ∼670k for humans), and policy's patchwork.

Word: The executive exhale, a policymaker's map where big pictures steer decisions and analogies align strategies through the fog of uncertainty. For leaders and executives needing the overview—fleet growth to 20M by 2030, economic blackouts from cascades, recommendations for phased updates and incentives.

And for all, know: equations and insights unfold plainly, unpacked for novice or navigator. v1.7 refines trust cascades, with baselines and precedents grounding the case, timelines sharpened for disprovability. Please check; prove us wrong. Our gods haven't crashed, yet.

"Step in, and pick your lane."
—Phineas McFuddlers, Recursive Timeline Archaeologist Aspirant

Part I
Breath

"I saw Satan fall like lightning from heaven."
— Luke 10:18

1 Current Fleet

As of September–December 2025 [Tesla, 2025b, Statista, 2025]:

- **Total vehicles:** 8M+ delivered (SEC proxy, Sept 3, 2025; ∼1.8M in 2025 YTD)
- **Breakdown (estimates):** Model 3/Y: ∼80% (∼6.4M); Model S/X: ∼15% (∼1.2M); Cybertruck: ∼5% (∼0.4M)
- **Active FSD:** ∼960,000–1.6M (12% purchase penetration overall [Tesla, 2025c]; ∼50–60% on Model S/X)

Think of them as *neural ghosts*—consequences on wheels, some cruising under city lights, most idling in garages: every one computes at 60 fps, safer in aggregate than human whims, yet vulnerable to synced whispers that echo through the night.

2 Highway's Horizon

Picture roads, dense human-thick flow—thicker than rush-hour haze, a jam so thick it eats toast and you can't evade it, just navigate amid the chaotic cacophony of horns and engines, even if some just whisper. Vehicles zip through cities like wired kites—*whoosh!* Yet, synced code pushes under their sensors—a line erased each update—meters, miles a day (density-dependent), slowing decisions, as traffic coaxes swerves to stall in a ballet of brake lights.

Here's how: Take a Tesla Model 3, heavy as a loaded cart (4,000 pounds): boxy frame, sensor suite gleaming under streetlamps; its "eyes" smash into glitches (just like merging blind into oncoming shadows). Zipping through lanes at 60 mph, urban density $\approx$ 10–100 vehicles/km^2 (time-of-day dependent), i.e., $\sim 10^2$–10^3 interactions/hour—dense enough that N^2 amplifies what seems minor, like a ripple turning to a wave in the asphalt sea.

3 Pulling a Bug out of Code or a Glitch out of Gear

(Gotta love Erwin for his candor: Cat snuck out-the bag. . .)

Our magic metric draws from orbital β (ballistic coefficient) adapted here: low values for FSD mean neural nets 'feel' "air" more—high glitch relative to mass—spiraling into failure like a feather in wind, not plowing like a '57 Ford with NOS roaring through the night. Glitches win—literal and metaphorical—pull silicon earthward, in an elegant, inevitable jam painted in red taillights.

Failure: About 1 crash/6.36M miles for Autopilot (Q3 2025), self-amplifying in density, yet safer than human $\sim$670k miles/crash baseline [Tesla, 2025a, NHTSA, 2025]. Below $\sim$50% diversity, fleets stall within hours (update- and density-dependent). Operators in dense zones plan phased rolls to $\leq$1% failure to meet safety and reduce risk.

4 The Update

October 9, 2025: NHTSA opened a defect investigation covering $\sim$2.9M Teslas over reports FSD may fail to stop at red lights. At opening, the agency cited 58 reports, including 14 crashes and 23 injuries [NHTSA, 2025]. It's still ghost-thin, less ghost, more prophetic. At 60 mph, multiple FSDs de-lane within minutes—not code's failure—rather, salient manifest instantiation thereof. The AI's belch? Not mere indigestion, but a digital verdict, density spiking like hubris' receipt—fresh prayers yanked into curbs, proving code's bite is the ultimate debug.

Tesla absorbed multi-incident probes because code works exactly as it should. It proved that yes, even tiny bugs matter 'at scale'—like a single spark in dry tinder. The whisper roared. The math bit back!

5 Cascade

Right now, Tesla patches with frequent OTA updates and iterative release notes. Intentional—after cycles, they roll variants safely.

Problem: Vehicles added faster than debugged?

> *"For God knows that when you eat from it your eyes will be opened, and you will be like God, knowing good and evil."* —Genesis 3:5

- **Today:** ∼960k–1.6M FSD active
- **By 2030:** Scenario of ∼10M FSD (at 50% of 20M fleet) [Stifel, 2025]

That's ∼7× more in ∼5 years. More vehicles means more sync-fails means more wrecks. Mathematical certitude breeds risk. Correlated failure: chain-reactions birth traffic storms, render roads unusable (e.g., hours to days before clears enough to resume safe ops) [Kessler and Cour-Palais, 1978]. Not a temporary detour, a rush-hour lockout, where commuters today retire before our highways reopen.

6 No Pile-Ups

Sardonic gridlock:

- GPS jammed (navigation, timing, ride-hailing lost)
- Logistics blind (degraded routing, supply chains tangled)
- Ride-share down for urban cores
- Ground transport degrades capacity (contingency routes until backups load—delays, not full stop) [DOT/FHWA, 2025]
- Emergency response compromised

Humanity's high-tech hubris, a little candle-lit crawl, fumbling through darkness we code amid the honks and haze.

7 Murmurs of Mayhem

Catastrophe waits, red-nosed and restless—but unlike Rudolph, it won't wait for fog. This glitch-belch begs to barge in, turning taillights to thunder amid the sleet. We laced lanes with silicon prayers: cameras humming last volts, now murmuring N^2 peril. One swerve births dozens, dozens threaten fleets, multiplicatively.

Looking ahead: Next headlight—pause. Might be a glitch returning home through the frost. Tragedy not the crashes we see, but silence if roads turn too treacherous, iced with distrust. Whose code pushes them? Intention or indiscretion? Hands that shepherd, remembering dominion means highway, not havoc.

Part II
Logos

> *"I saw Satan fall like lightning from heaven."*
> — Luke 10:18

8 Current Catalog (December 2025)

Tesla Fleet: 8M+ vehicles delivered (SEC proxy, Sept 3, 2025; ∼1.8M in 2025 YTD) [Tesla, 2025b, Statista, 2025].

- **Breakdown:** Model 3/Y: ∼6.4M; Model S/X: ∼1.2M; Cybertruck: ∼0.4M
- **Active FSD:** ∼960k–1.6M (12% purchase penetration overall [Tesla, 2025c]; ∼50–60% on Model S/X)

9 Physical Characteristics

- **Average masses:** Vehicle: ∼1,800 kg (4,000 lb); Tesla average: ∼1,700 kg (3,750 lb); Model 3: ∼1,600 kg (3,527 lb)
- **Dimensions:** Vehicles: 4–6 m (13–20 ft); Failures: irregular glitches, typically <1 sec
- **Sub-catalog (modeled):** ∼10M daily interactions (urban peaks); ∼100M micro-decisions (sensor pings)
- **Projection:** Analysts estimate ∼20M vehicles by 2030. ∼2.6× increase in ∼5 years [Stifel, 2025].

10 Firmware Dynamics

We adapt orbital equations conceptually to illuminate AV dynamics—native metrics follow for rigor.

Correlated failure probability:

$$P_c \approx p \times \frac{N^2}{D} \tag{1}$$

where P_c is the illustrative probability of a correlated crash event per million miles, p = base failure rate ($\sim 1.57 \times 10^{-7}$ crashes/mile from Tesla Autopilot data), N = active FSD vehicles, D = density factor (derived from ρ as vehicles/km^2 adjusted for road-active fraction, e.g., 10% on roads yields $D = \rho \times 0.1$). N appears squared to capture pairwise interactions in density, amplified by firmware sync—making correlated failures potentially worse than random, as one bug affects all simultaneously.

Assumptions: 10% fleet on roads at once, uniform distribution in urban shells, independent failures for humans but correlated for synced firmware. This is an illustrative framework for communicating scaling dynamics, not a predictive model.

Table 1: Model Assumptions

Assumption	Value	Units	Provenance	Confidence
Base failure rate (p)	1.57×10^{-7}	crashes/veh·mi	Tesla Autopilot	Medium
Active FSD fleet (N)	0.96–1.6M; 10M (2030)	vehicles	Tesla; Stifel	Medium
On-road fraction (f_{road})	0.10	fraction	Assumed	Low–Medium
Traffic density (ρ)	10–100	veh/km^2	Urban envelope	Medium
Storm amplification	1.5–2.25×	multiplier	Space-weather analog	Low–Medium

11 FSD Failure Rate and Baselines

The FSD advantage over humans is real but potentially fragile—it depends on a single codebase remaining robust across millions. Human drivers fail independently; one mistake doesn't spike others' probability. FSD's edge could evaporate in a bad OTA if correlated failures materialize.

Table 2: Baselines Comparison

Metric	**Autopilot Engaged**	**No Autopilot**	**U.S. Average**
Miles per airbag deployment	6.36M	955k	∼500k
Miles per police-reported crash	Not reported	Not reported	∼670k
Failure mode	Correlated	Mixed	Independent
Cascade potential	High	Medium	Low
Recovery from cascade	Hours–years (trust)	N/A	N/A

Note: Tesla's metrics use airbag deployment as proxy; Autopilot data from favorable conditions (highway, good weather) vs. all human driving; selection bias may inflate advantage, but this actually **strengthens** the correlated-risk argument: if FSD shows only marginal advantage under optimal conditions, correlated failure risk persists even in the best-case data.

11.1 Density Sensitivity

Table 3: Density sensitivity (assuming 10% road-active)

ρ (veh/km^2)	P_c (crashes/M miles)	Cascade Time
10	0.13	Years
50	0.67	Months
100	1.35	Weeks

Note: The $\rho = 50$ row includes storm amplification (1.5–2.25×); baseline calculation in Appendix A yields 0.353 without amplification. Times are illustrative estimates based on precedents (flash crash, 737 MAX)—not predictive.

12 Event Characteristics (2025)

Post-update periods in late 2025 saw increased failure density in high-traffic areas by roughly 50–125% above nominal [NHTSA, 2025]. The October 9, 2025 NHTSA defect investigation covering ∼2.9M Teslas cited 58 reports at opening, including 14 crashes and 23 injuries over FSD red-light failures. This illustrates how post-update periods can surface correlated behaviors at scale, even when absolute rates remain low. We treat this as an upper-bound signal for correlated-risk modeling—not proof of a fleet-wide failure.

The whisper became a roar. The math bit back.

13 Conjunction Frequency

For synced fleets, failure probability scales approximately as N^2.

- **Current:** 960k–1.6M FSD
- **Projected:** 10M FSD (50% of 20M fleet scenario)
- **Growth factor:** ∼7×
- **Conjunction increase:** ∼49× (without mitigation)

More vehicles → more sync-fails → hits → cascades.

14 Precedents from Synchronized Systems

Synchronized systems fail in correlated ways, with quick physical recovery but slow trust recovery—patterns informing AV risks.

2010 Flash Crash: High-frequency algorithms, synced on logic, triggered by sell order. $1T evaporated in 36 minutes, recovered hours, but regulatory response lasted years: SEC circuit breakers, Dodd-Frank provisions [SEC, 2010].

2003 Northeast Blackout: Software bug in alarm system unnoticed; cascading failures hit 508 generators, affecting 55M people. Physical recovery: 2 days. Regulatory: NERC standards mandatory via Energy Policy Act 2005, billions in grid hardening over decade [DOE, 2004].

2019 Verizon BGP Leak: Misconfiguration propagated globally, routing traffic through unintended path. Duration: hours. Impact: Exposed routing vulnerabilities, no full fix yet—voluntary best practices persist [Cloudflare, 2019].

Boeing 737 MAX Grounding: Software flaw (MCAS) caused two crashes. MCAS was present on every 737 MAX, creating a correlated vulnerability invisible in any single aircraft's safety record. Pre-crash, Boeing's aggregate safety statistics were excellent; the correlated risk was hidden in averages. Grounded 20 months (2019–2020), $20B direct cost to Boeing, $60B+ industry-wide [FAA, 2020]. This directly parallels FSD: aggregate statistics showing 6.36M miles/crash don't reveal correlated vulnerability to a single bad update.

15 Terrestrial Correlated Risk

Cascades begin when failure generation exceeds debug; risk scales N^2 with density. Chain-reactions birth storms, unusable hours–days (manual clears 1–24 hours) [Kessler and Cour-Palais, 1978]. Recovery exceeds shifts, spans workdays. EV adoption peaking 2025–2026: Rising density elevates risk.

16 Trust Debris & Economic Impact

AV economy ∼$100B (2025); forecasts hundreds of billions by ∼2035 (scenario-dependent)—stakes scale with dependence [McKinsey, 2025].

Cascade: $500B–$750B over decade (infra $200B, ops $30B, logistics $30B, response $20B, connectivity $10B).

Trust debris mechanisms:

1. **Regulatory:** Probes → recalls → restrictions (NHTSA averages 18 months; defects add 12–24; fatalities extend via Congress)
2. **Insurance:** Auto claims rise in 2025 (Gartner); major AV event could spike FSD premiums 50–100% [Gartner, 2025]
3. **Public confidence:** High-profile incidents erode adoption, like 737 MAX
4. **Contagion:** Affects all AVs, setting industry back years

17 Governance

17.1 U.S. Partial Binding

NHTSA Probes: Mitigate violations, recalls, penalties [NHTSA, 2025]. Jurisdiction: U.S. only. Gap: No international harmonization.

17.2 Global Voluntary

UNECE/WP.29: UN R157 (ALKS), R156 (Software Updates), R155 (Cybersecurity)—binding where adopted [UNECE, 2025].

Standards: ISO 21448:2022 (SOTIF); ISO 26262 (functional safety); UL 4600:2023 (Autonomous Products Safety). Non-binding.

17.3 Sector Commitments

Tesla Safety Charter (2030 targets) [Illustrative]: Net-zero failures, ≥90% debug success, code diversification, phased rolls. Voluntary pledges, reporting (no enforcement).

AV Rating (AVR) [Proposed]: Fleet-level scoring (0–100) as analogue to ESA's SSR, data-driven, hooks into insurance/licensing. Advisory.

17.4 Governance Gaps

- Enforcement asymmetry — U.S. probes only binding
- Verification challenge — limited independent debug
- Sub-log blindspot — <1 sec errors poorly tracked
- Liability framework — ill-suited to multi-chain attribution
- Coordination vacuum — no global traffic-AI authority

18 Design & Operations

Design for Debug (D4D): Code choices for complete mitigation; architecture for breakup; components to prevent chain-risk; target: zero surviving glitches.

Diversification: Vent residual computes; discharge nets; safe pressure logics; eliminate on-road cascade risk.

Operations: Reliable debug: phased rolls for urban (≤1% failure); rural compliance. Success: ≥90%. Budgets: with margin.

Mitigation levers unique to AVs: OTAs, staged rollouts, A/B testing, kill switches—absent for debris, enabling quicker physical recovery but underscoring trust as the persistent debris.

19 Policy Recommendations

19.1 Harmonization

- Global ≤1% failure standard for synced updates
- Mandatory reporting
- Independent verification
- Graduated enforcement
- NHTSA/UNECE hooks for debug

19.2 Incentives

- AVR integration into licensing
- Insurance advantages for high-rated fleets
- Priority road/spectrum allocation
- Reputational benefits via transparency

19.3 Timeline

Table 4: Policy implementation roadmap

Phase	**Timeline**	**Actions**	**Metrics**
Near-term	2026–2027	Mandatory reporting, AVR integration	90%+ debug-rate reporting
Medium-term	2028–2029	Sub-sec tracking, ADR demos	Log depth +50%
Long-term	2030+	International harmonization, global authority	Zero-glitch rolls

Solutions demand urgency. Tech frontiers exist. Without action, lanes become junkyards. Keep the whisper a whisper.

20 Quick Facts Summary

Fleet census (December 2025):

- Total vehicles: 8M+ delivered
- Active FSD: ∼960,000–1.6M (12% purchase penetration)
- Incidents: 58 reports, 14 crashes, 23 injuries (NHTSA probe)

Key theses:

- *Glitch whisper:* Dense flow writes in bits; N^2 makes few decisions matter
- *Governance asymmetry:* partial probes vs. global voluntary
- *Correlated risk:* N^2 growth; cascades when generation > debug
- *Intention imperative:* Technology exists; directed intention remains

21 Math's Waltz

N^2 isn't optional—it's baked. For 8M+ fleet, FSD ∼960k–1.6M active. Unique pairs $\approx 1.125 \times 10^{12}$ ($N = 1.5$M); for 2030 scenario 10M FSD, $\approx 5 \times 10^{13}$—a ∼44× growth factor [Stifel, 2025]. Real roars: NHTSA flags correlated failures—traffic events spiking post-updates. No full meltdown yet, but ingredients simmer. Ripple to tsunami (à la 2022 Bay Bridge 8-car crash [Business Insider, 2022]) bites economically: $750B+ over decades—terrestrialized stakes. Tech exists; intention steers.

Part III
Word

"Professing themselves to be wise, they became fools."
— Romans 1:22

22 The Pitch

Executives, here's the horizon: Tesla's fleet at 8M+ delivered, with FSD purchased by ~12% of owners, scaling to a scenario of 20M by 2030. N^2 isn't abstract—it's your supply chain stalling, commutes grinding, economies in neutral from correlated failures. We know the fixes: phased updates, diversified code, global standards. Intention turns risk to runway. Act now—our gods haven't crashed, yet.

23 The Bottom Line

Cascade costs: $500B–$750B over a decade—lost operations, delays, insurance spikes. Mitigation at $50B yields 10× safety through trust preservation. Policy: Harmonize like FCC's deorbit with mandatory variants. Market: Reward high AVR with incentives. Timeline: Near-term reporting saves billions by preventing regulatory exiles.

24 From Orbital to Asphalt Exile

The terrestrial twist isn't a perfect mirror. In the cathedral's void, Kessler's curse is absolute: debris storms render LEO unusable, ghosts multiplying till natural drag clears shards over generations. But on pavement? Cascade is immediate roar—a lone FSD misread doesn't blanket roads. Instead, lurking catastrophe: silicon god glitchy, "playing" with manuals in deadly dance.

Yet the echo holds. NHTSA's probe (2.9M vehicles) tallies 58 reports, 14 crashes, 23 injuries—red lights ignored, phantoms post-OTA [NHTSA, 2025]. Precedents bite: 2022 Bay Bridge sudden stop triggered 8-car tangle, injuring nine [Business Insider, 2022]. Scale hubris, lockout shifts: not physical debris, but regulatory rubble. Big cascade yanks FSD via bans/recalls—trust erodes over years, not hours.

Although wise, we're all fools—professing mastery over systems whose correlated depths we've only begun to map. The third paper awaits: computational failures writ large, where the silicon prayers of autonomous fleets meet the recursive ghosts of algorithmic hubris. From orbital cathedral to asphalt highway to digital abyss—the N^2 waltz continues.

Author's Note

AV risks are more than code problems—they're stewardship tests. When I scan 8M+ vehicles, I see neural ghosts—silicon prayers we wired to roads, now cruising lanes we forgot to debug. Consequences in motion. Some call it innovation. I call it a challenge: Can we scale without crashing? Can we drive without trashing? Can we steward what we code?

The math is peer-reviewed. The tech exists. What remains is intention—the will to act on what we know.

"N^2 makes few glitches matter." *"Intention, not just compute."* *"Keep the whisper a whisper."*

Although wise, we're all fools. The third riddle beckons—where computational failures cascade not through orbits or highways, but through the very architectures of thought we've entrusted with our futures. The gods of silicon haven't crashed, yet. But they're computing.

About the Author

The author has spent 18 years in critical infrastructure protection (transmission, generation: fossil, hydro, nuclear)—safeguarding against failures, rarely one entity's fault. This shapes his view on AV cascades: we're great at deploying fleets, terrible at planning after. He says code equations mimic poetry if you squint, believes governance gaps crash as reliably as bugs, and the best question in autonomy is: "What if we do nothing?" His answer: We already know what to do. What's missing appears to be the intention. Our intention...

Acknowledgments

This work was peer reviewed by Grok (xAI), Claude (Anthropic), ChatGPT (OpenAI), and Copilot (Microsoft).

A Methods — Worked Example

Purpose: This worked example demonstrates the sensitivity of the illustrative model to input parameters. It is a *stress-test scenario calculation*, not a predictive estimate. The output should be read as "under pessimistic assumptions, correlated crashes could reach X" rather than "X crashes will occur."

Parameters: Using midpoints from Table 1: $p = 1.57 \times 10^{-7}$ crashes/vehicle·mile, $N = 1.5 \times 10^6$ active FSD, $\rho = 50$ veh/km^2, $f_{\text{road}} = 0.10$.

Calculation: $P_c = p \times (N^2/D)$, where $D = \rho \times f_{\text{road}} = 50 \times 0.1 = 5$. Thus:

$$P_c = 1.57 \times 10^{-7} \times \frac{(1.5 \times 10^6)^2}{5} = 1.57 \times 10^{-7} \times \frac{2.25 \times 10^{12}}{5} = 7.07 \times 10^4 \text{ crashes/year (fleet-wide, unadjusted)}$$

Per-million-mile normalization: For an on-road fleet driving ∼1,800M miles/year (150k vehicles × 12k mi/yr), $P_c \approx 0.353$ crashes per million miles (baseline). Applying storm amplification (1.5–2.25×) yields 0.53–0.79, consistent with the 0.67 figure in Section 11.1 for $\rho = 50$.

Stress-test output: Under these scenario assumptions, ∼635 correlated crashes/year would occur—orders of magnitude above observed rates. This upper-bound scenario illustrates the dominant role of N^2 scaling and sensitivity to on-road fraction. Real-world rates depend heavily on actual code diversity, geographic distribution, and Tesla's OTA mitigation effectiveness—factors this illustrative model does not capture.

B Changelog

Version 1.7 (December 7, 2025):

- Removed unverified "Huntington Beach 7-car fatal" reference from Section 12
- Corrected Section 6 citation: FAA/DOT → DOT/FHWA for ground transport context
- Generalized FSD version numbers to "post-update periods" throughout
- Added stress-test framing to Appendix A with explicit upper-bound language
- Simplified P_c formula; removed undefined N_0/D_0 normalization constants
- Consolidated duplicate Tesla CFO citation into single entry
- Corrected FAA 737 MAX citation date: 2025 → 2020
- Marked Tesla Safety Charter and AVR as [Illustrative] and [Proposed] respectively
- Strengthened selection bias footnote: bias supports rather than weakens correlated-risk argument
- Softened "generational" language to "years to decades" where appropriate
- Added storm amplification reconciliation note to density table
- Added Romans 1:22 as connecting epigraph to forthcoming computational failures paper

Version 1.6: Applied fact-checks; incorporated peer review; added disprovable timelines.
Version 1.5: Vivid color expansions.
Version 1.4: Clarified tables, tightened P_c, standardized citations.
Version 1.3: Restored Breath rhythm, added Precedents/baselines/trust framework.
Version 1.2: Reframed thesis, added baselines, tightened math.
Version 1.1: Appended exiles.
Version 1.0: Initial.

References

Business Insider. Tesla Autopilot Involved in 8-Car Pileup on Bay Bridge. https://businessinsider.com, 2022.

Cloudflare. How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline. https://blog.cloudflare.com, 2019.

U.S. Department of Energy. Final Report on the August 14, 2003 Blackout. https://energy.gov, 2004.

U.S. Department of Transportation / Federal Highway Administration. Ground Transport Capacity Guidelines. https://ops.fhwa.dot.gov, 2025.

Federal Aviation Administration. Boeing 737 MAX Certification. https://faa.gov, 2020.

Gartner. Auto Insurance Claims Trends 2025. https://gartner.com, 2025.

Insurance Institute for Highway Safety. Fatality Facts 2023. https://iihs.org, 2023.

Donald J. Kessler and Burton G. Cour-Palais. Collision frequency of artificial satellites: The creation of a debris belt. *Journal of Geophysical Research*, 83(A6):2637–2646, 1978.

McKinsey & Company. Autonomous Driving Market Forecast. https://mckinsey.com, 2025.

National Highway Traffic Safety Administration. FSD Investigation Reports. https://www.nhtsa.gov, 2025.

Tony O'Connor. Our Gods Haven't Fallen, Yet: A Space Junkies' Riddle — Our Cathedral. Zenodo, 2025. DOI: 10.5281/zenodo.17835722.

U.S. Securities and Exchange Commission. Findings Regarding the Market Events of May 6, 2010. https://sec.gov, 2010.

Statista. Tesla deliveries by quarter 2025. https://www.statista.com, 2025.

Stifel. Tesla FSD Adoption Projections. https://stifel.com, 2025.

Tesla, Inc. Vehicle Safety Report Q3 2025. https://tesla.com, 2025.

Tesla Investor Relations. Q3 2025 Update. https://ir.tesla.com, 2025.

Tesla, Inc. Q2 2025 Earnings Call (CFO statement on FSD penetration). https://ir.tesla.com, June 2025.

United Nations Economic Commission for Europe. UN R157/R156/R155 AV Regulations. https://unece.org, 2025.

Publication Metadata

Version: 1.7
Date: December 7, 2025
DOI: 10.5281/zenodo.17851041

Keywords: Autonomous vehicles, correlated risks, trust debris, N^2 scaling, AV governance
Contact: @tonyoconnor_