

Crumple

An Empirical Crush and Force Envelope for the Modern Australian Vehicle Fleet

Daniel Vaughan (ORCID: [0009-0000-4783-0763](https://orcid.org/0009-0000-4783-0763))

Abstract

Designers of vehicle barriers face a gap between two regimes. At the low end, AS/NZS 1170.1 prescribes fixed car-park barrier forces whose assumed crumple zones (0.10–0.15 m) have no stated empirical derivation and do not scale with vehicle mass or speed. At the high end, hostile-vehicle-mitigation standards certify specific barrier products at 48–80 km/h but publish no calculation method. In between, for impacts at 30–60 km/h or by vehicles heavier than a passenger car, the designer has no empirically grounded estimate of vehicle crush or average impact force.

This paper offers that basis: an empirical research envelope. From 10,876 crash-test records of the United States National Highway Traffic Safety Administration (NHTSA), a documented quality-control chain yields a modern cohort, model year 2004 on, age-matched to the Australian passenger fleet. The 614 tests with directly derived dynamic crush supply the primary fit, a non-crossing quantile envelope in speed and mass, with average force following from the energy identity $F_{avg} = mv^2/(2d)$, checked independently against barrier load cells. The result is a rigid-barrier crush and force envelope supported over 1,007–3,448 kg and 33–57 km/h, read in both directions: the stiff-vehicle 5th percentile bounds force demand, the soft-vehicle 95th percentile bounds penetration. The load cells locate the height at which the force acts. Beyond the tested speeds, external evidence bounds a crush ceiling the data alone cannot locate, keeping the high-speed force treatment conservative. It is offered as a candidate basis for future calibration, not as characteristic design actions.

1. Problem

AS/NZS 1170.1:2002 Clause 3.8 covers vehicle impact on car-park barriers with three prescriptive loads:

Scenario	Prescribed force	Application
Light traffic barrier (Type F)	30 kN	1.5 m length, 0.5 m above floor
Medium traffic barrier (Type G)	40 kN	1.5 m length, 1.0 m above floor
End of long downward ramp (>20 m)	240 kN	—

The Commentary (AS/NZS 1170.1 Supp 1:2002, Clause C3.8) does state the basis: each prescribed force is the output of the energy-balance formula $F = mV^2/(2\Delta)$ with a stated vehicle mass, speed, and assumed crumple zone (Δ):

Scenario	Mass	Speed	Assumed crumple (Δ)	F (computed)	F (prescribed)
(a) Light traffic (Type F)	1,500 kg	2 m/s (7.2 km/h)	0.10 m	30 kN	30 kN

Scenario	Mass	Speed	Assumed crumple (Δ)	F (computed)	F (prescribed)
(b) End of long ramp	2,000 kg	6 m/s (21.6 km/h)	0.15 m	240 kN	240 kN
(c) Medium traffic (Type G)	2,000 kg	2 m/s (7.2 km/h)	0.10 m	40 kN	40 kN

The prescribed forces are thus exact outputs of the energy-balance formula (the same formula this paper uses). The gap is not that the standard is silent on its basis; it is that the Commentary's tabulated crumple zones have no stated empirical derivation and no way to scale with vehicle mass or speed. A designer asked to rate a barrier, bollard line or building edge for a 2,000 kg vehicle at 50 km/h, a realistic errant-vehicle or low-end hostile-vehicle case, finds nothing between Clause 3.8's three fixed points and the hostile-vehicle-mitigation (HVM) certification standards (ISO 22343-1:2023 and ASTM F2656/F2656M-23). Those standards certify specific barrier products against specific test vehicles, publishing no calculation basis and no way to interpolate between rated scenarios.

What is missing is an empirically grounded relationship between vehicle mass, impact speed, vehicle crush, and force on a rigid barrier, along with a realistic picture of how much that relationship varies across the fleet. The literature's standard tool, the Campbell crush-energy model (S1) and its CRASH3 stiffness-category descendants (S2), is anchored to 1970s–1990s vehicle structures. On the 198 modern vehicles (model year, MY, 2003–2020) the citing study itself sampled, a range bracketing this paper's MY $\geq$ 2004 cohort, the Campbell form accounts for only 15.2 % of the speed–crush relationship (S3). Reconstruction studies reach the same conclusion from the other direction: vehicles built since the mid-1990s deform appreciably less at the same speed, so the older coefficients under-predict impact speed for the modern fleet, by around 20 % in one multi-model study (S28) and consistently across a 1,400-test review that names impacts into structures as a direct application of better crash-test data (S29). A review of anti-ram bollard design likewise states that “further investigation of available crash test data would be beneficial” for its velocity-to-crush ratios (S4). That is the gap this work addresses: a modern refit, with its variability quantified.

2. Data

Dataset lineage: from raw records to fitted model



Figure 1: Dataset lineage from raw records to fitted model

Figure 1. Dataset lineage. The main pipeline (dark) runs from raw NHTSA records through QC and cohort selection to the fitted set. Side branches (orange) show the derived dynamic-crush pool, the load-cell validation set, and the corrected-static sensitivity subset. Each box gives the count and a one-line reason for the step.

2.1 Corpus and quality-control chain

The source corpus is the NHTSA vehicle crash test database (live API, verified 2026-07-07). Figure 1 summarises the pull and QC pipeline: raw records are reduced to clean full-width rigid-barrier (FWRB) head-on tests, every rejected row retained with a per-row reject reason, so nothing is dropped silently. The dominant rejection is scope (not a head-on FWRB test); the remainder fail plausibility rules (missing essentials, non-physical mass/velocity/crush, crush exceeding vehicle length).

Barrier rigidity is confirmed from NHTSA's own barrier records rather than inferred from the test configuration. Querying every previously unverified row against the barrier-information endpoint confirmed 2,741 rows as rigid and excluded 142 that had passed the configuration filter but were in fact deformable-barrier tests; 59 rows with no API barrier record remain flagged unverified.

2.2 Dynamic crush: measured where possible, corrected elsewhere

Provenance first, because it recurs throughout the paper. Every crush value used here is one of two kinds. A measured value comes straight from the test's own accelerometer trace. A corrected value starts from the static (residual) crush that NHTSA records and scales it up to an estimated dynamic value. The measured route is the one we trust, and the primary model is fitted on it; the corrected route only fills gaps. A fit that pools both routes is called the hybrid, and it appears only as a sensitivity check bracketing the measured-only result.

NHTSA tabulates only static (residual) crush, the deformation left after rebound (field CRHDST, per the database coding manual, S5), not the peak dynamic crush a barrier designer needs for deflection and stand-off. Two routes supply it:

- **Direct derivation (preferred):** double-integration of the test's X-axis accelerometer time-series to the velocity zero-crossing, yielding peak dynamic displacement. Run over the corpus and QC-filtered, the derived pool spans 7.7–77.3 km/h (Figure 1).
- **Velocity-dependent correction (fallback):** for cohort rows without a usable accelerometer channel, dynamic crush = static crush $\times$ $ratio(v)$, with $ratio(v) = 1.310 + 1.201/v$ (v in m/s; about 1.4 at the test speeds) fitted from QC-filtered tests with both quantities. The Vangi form (additive, velocity-dependent gap between dynamic and residual crush; S6) motivates the $1/v$ term, and an independent set of eighteen frontal rigid-barrier tests on three MY-2001 cars, reporting dynamic-to-residual ratios of about 1.0 to 1.5 (S28), places the correction in the right range, though measured against plastic rather than static deformation.




Figure 2: How one test yields a measured and a corrected crush value

Figure 2. The two crush provenances, shown on one real test (NHTSA test 9808, 56 km/h full-width rigid barrier). Top row (teal, the measured route): the test's own accelerometer record is double-integrated to a displacement curve whose peak is the measured dynamic crush, 615 mm here. Bottom panel (orange, the corrected route): the tabulated static residual crush for the same test (468 mm) is scaled by the velocity-dependent ratio (1.39 at 56 km/h) to an estimated dynamic value of 649 mm. The chips at the foot show which rows each fit uses: measured-only (primary, $n = 614$), corrected-only (sensitivity, $n = 458$), hybrid (both, $n = 1,072$).

The corrected-static path is kept only as a labelled sensitivity, for three reasons tested directly on the cohort rows that carry both values: the correction runs a little high on average, scatters widely from test to test, and implies a different velocity trend that a sound correction would have reconciled (§3.1, Annex A.10); no cohort imbalance in speed, mass, year or size explains the demotion. Figure 3 distinguishes the two provenances by marker.

2.3 The central cohort and its structure

The design cohort is $MY \geq 2004$ (Figure 1), a cut-off chosen so the cohort's mean age approximates the Australian passenger-fleet mean age: per the Bureau of Infrastructure and Transport Research Economics (BITRE) *Road Vehicles Australia*

series (S7), Australian passenger cars average roughly 11 years old, and the cohort's mean age sits within about a year of that. The discontinued Australian Bureau of Statistics (ABS) Motor Vehicle Census (S8) is retained only for its age-band breakdown, which supplies the fleet-weighting sensitivity weights (§3.4).

Because today's fleet is slightly newer than the cohort, and the cohort's own mass-controlled trend implies newer designs crush slightly less (are stiffer), the age gap biases the envelope mildly toward understating force demand and overstating penetration allowance: a small, one-directional effect no refit corrects here.

Two structural facts belong up front rather than in a limitations annex:

- **The velocity coverage is three regulatory clusters, not a continuum.** Tests concentrate at ~40 km/h, ~47–48 km/h, and ~56 km/h (the regulatory test speeds of the US Federal Motor Vehicle Safety Standards, FMVSS, and the New Car Assessment Program, NCAP), with near-empty space between. Within-band interpolation is therefore model-driven; the velocity exponent is identified chiefly by the 40 ↔ 56 km/h contrast.
- **The fleet is US-market.** NHTSA tests US-specification vehicles; the Australian fleet differs in model mix (higher share of utes and light commercial vehicles, LCVs, and a different small-car mix). ABS weighting matches fleet age, not model mix: the §3.4 sensitivity bounds part of this effect, and the model-mix component is unquantified.

The measured-only primary model's data envelope is 9.1–15.8 m/s (33–57 km/h) and 1,007–3,448 kg. The lower edge of the 30–60 km/h band of practical interest (30–33 km/h) sits just below the data floor. That small sliver is extrapolated, and flagged as such rather than rounded into the supported claim.

3. Method: crush and force

3.1 Quantile envelope

Dynamic crush is fitted as a power law in velocity and mass:

$$d = a \cdot v^\beta \cdot m^\gamma$$

where d is crush in mm, v impact velocity in m/s, and m vehicle mass in kg. On the measured-only cohort the velocity exponent $\beta \approx 1.27$ and the mass exponent $\gamma \approx 0.18$. In plain terms: crush grows somewhat faster than proportionally with speed, and a heavier vehicle crushes only slightly more (doubling the mass adds roughly 13 % crush, not 100 %). The full parameter tables and exponent interpretation are in Annex A.2–A.4.

The fit is quantile regression at the 5th, 25th, 50th, 75th and 95th percentiles on log-transformed data, with the velocity and mass slopes fitted once and shared across all five curves, so the curves differ only by a multiplicative constant and cannot cross: a structural requirement for any chart read at extrapolated speeds (Annex A.1).

The power-law form itself is chosen on three grounds. It is the simplest form through the origin that nests the two limiting cases of a crushing structure, a linear spring (crush proportional to v) and a constant-force structure (crush proportional to v^2), so the fitted exponent reads directly as a position between those limits. Being linear in logarithms, it lets the five quantile curves share slopes and differ only by intercept, making the non-crossing envelope exact. The tested alternatives fail on the evidence: the Campbell linear-with-intercept form needs a modern low-speed anchor (b_0) that no public corpus reviewed here supplies (§10 item 1), and saturating forms are statistically tied with the power law inside the band while their ceilings run away unidentified, a question treated with the extrapolated tail (§5).

The primary model is the measured-only fit (Figure 1); the corrected-only and pooled-hybrid fits are labelled sensitivity bounds that bracket it. The velocity-trend mismatch flagged in §2.2 is the +0.58 gap between the measured and corrected velocity

exponents: the correction does not reconcile the two provenances (Annex A.2). Figure 3 shows the three fits against the underlying data at cohort median mass.




Figure 3: Three provenance fits at cohort median mass

Figure 3. Measured dynamic crush (teal circles) and corrected static crush (orange triangles) with three median fitted curves at cohort median mass. The measured-only fit (solid) is the primary model; the corrected-only (dashed) and hybrid (dotted) are sensitivity bounds.

Read against the two limiting cases above, the measured-only velocity exponent sits near the linear-spring limit, consistent with progressive stiffening of modern front structures. The mass term is substantially a vehicle-size proxy rather than an independent mass effect: its apparent stiffness-with-mass trend largely disappears when footprint is controlled (Annex A.4).

3.2 Two-sided conservatism rule

The envelope is read in two directions, analogous to percentile-based characteristic loads in wind and seismic codes:

- **5th percentile (stiff vehicle, low crush) → governs FORCE.** A stiff vehicle stops in less distance and transmits more force to the barrier and its anchorage.
- **95th percentile (soft vehicle, high crush) → bounds PENETRATION.** On a rigid barrier the penetration is the vehicle crush. For a deflecting barrier, rigid-wall crush plus barrier deflection is an upper bound on total travel only (§7.2), not an exact partition.

A single check therefore uses both tails of the same fitted envelope.




Figure 4: Measured-only crush envelope across the cohort mass range

Figure 4. The measured-only dynamic-crush envelope: the 5th, 25th, 50th, 75th and 95th quantile curves (drawn at the cohort median mass) against the measured cohort data, coloured by vehicle mass. The green band marks the data-supported speed range (9.1 to 15.8 m/s, 33 to 57 km/h); red shading marks extrapolation beyond it, where the quantile curves are drawn dashed. The 5th-percentile (stiff) curve governs force; the 95th-percentile (soft) curve bounds penetration.

Dynamic-crush envelope at representative vehicle masses (measured-only fit; dashed = beyond the supported speed band)




Figure 5: Per-mass crush envelopes at representative vehicle masses

Figure 5. The dynamic-crush envelope evaluated at six representative masses (1,100, 1,500, 2,000, 2,500, 3,000 and 3,400 kg). Each panel shows the five quantile curves against the cohort tests within 150 kg of that mass, drawn dashed beyond the supported speed band. The green band is the data-supported speed range; red is extrapolation.




Figure 6: Crush contours over the supported mass-speed envelope

Figure 6. Crush contours over the supported envelope (9.1 to 15.8 m/s, 1,007 to 3,448 kg). Left: the 5th-percentile (stiff-vehicle) crush, the force-side read-off. Right: the 95th-percentile (soft-vehicle) crush, the penetration-side read-off. Contours are labelled at 50 mm increments; grey dots are the cohort tests, and the dashed lines mark the 1,100 and 3,400 kg reference masses. Crush rises steeply with speed and only weakly with mass.

3.3 Force model and the checks on it

Where the force number comes from

Average impact force follows from energy balance on a rigid barrier:

$$F_{avg} \text{ [kN]} = m \cdot v^2 / (2 \cdot d)$$

This is the same formula the standard's own Commentary uses (AS/NZS 1170.1 Supp 1:2002, eq C3.8: $F = mV^2/(2\Delta)$, where Δ is "the sum of the deflection of the vehicle and barrier"). The model's d is the vehicle-crush component only, derived on a rigid barrier with zero barrier deflection. The paper's contribution is not a new formula but an empirical replacement for the vehicle-crush component of Δ , which the Commentary assumes at 0.10–0.15 m without stated derivation.

This is an exact identity for the loading-phase average interface force on a rigid barrier: the work done over the loading stroke equals $\frac{1}{2}mv^2$ regardless of rebound energy, so restitution does not enter. The model's real breakdown below ~16–30 km/h is a crush-prediction failure, not a force-formula failure (§6), and the identity's own error terms are second-order at design speeds (Annex A.8).

How the force estimates are checked

Three checks apply; the first two test crush prediction, the third is a genuinely independent force check:

- **(i) Out-of-sample crush prediction.** We refit the model five times, each holding out a fifth of the measured-crush tests and predicting the held-out rows: the median model/measured force ratio is 0.997 (Annex A.9(i)). Because both sides of the ratio share the same m and v , this validates crush prediction, not force.
- **(ii) Cross-channel displacement agreement.** For the clean load-cell tests (§2.2), crush derived from the barrier load cells (double-integrating $F = ma$) is compared against crush from the onboard accelerometers. Median ratio 0.981 (interquartile range, IQR, 0.931–1.05): two independent instrument chains agreeing to ~2 % in the median. This is a displacement check, not a force validation, and the load-cell set spans 32.7–57.0 km/h, so it bears on the high-speed end only.

- **(iii) Independent force validation against barrier load cells.** The model's predicted average force is compared, at each test's own mass and speed, against the force the barrier load cells actually measured: a separate instrument from the accelerometers the model is built on (§2.2). The comparison is made at the model median (q50) deliberately: picking whichever curve happened to track the load cells most closely would turn an independent check into a self-calibration. At the median the model sits about 8 to 17 % below the modern load-cell averages. That is a central-tendency check, not the design reading: force demand is read off the stiff-vehicle q05 curve, a fixed 27 % higher, which more than offsets the central bias. The margin arithmetic and the single assumption behind it are set out in §8.5; the validation itself is Annex A.9(iii) (Figure A3).

One further force caveat matters. F_{avg} is an average over the crush distance; peak forces on load-cell barriers are a median 1.9× higher (IQR 1.8–2.2; robust to the CFC-60, Channel Frequency Class 60, filter correction), and anchorage design should account for this.




Figure 7: Peak versus average interface force

Figure 7. Where the peak-to-average multiplier comes from: a representative modern load-cell force trace at 56 km/h. The peak force is a brief transient (here about 1.9 times the average); the average is the sustained value the energy-balance identity $F = mv^2/(2d)$ returns. Anchorage and local checks must account for the peak (§3.3, Annex A.11).

3.4 Uncertainty characterisation

Three sources of uncertainty bound the envelope:

- **Bootstrap confidence intervals.** 500 resamples give 90 % confidence bands on each quantile curve: at a representative point (10 m/s, 1,500 kg), ± 3 % at the median and ± 3 –6 % at the tails, properly ordered at every velocity (Annex A.5).
- **Fleet-weighting sensitivity.** Weighting the cohort to the Australian fleet age distribution (ABS 2021 age bands, S8) shifts all quantiles downward by 4–7 %, small but real, not within resampling noise: the unweighted cohort is mildly biased toward softer (higher-crush) vehicles (Annex A.6).
- **Onset and initial-condition sensitivity.** Each measured crush depends on where integration starts and on the pinned initial speed. A perturbation study

over the measured fit-ready tests found a per-test worst-case shift of median ~5 % at highway speeds, rising steeply below 30 km/h, dominated by the detection threshold. Measured crush is therefore a band, not a point, and the paper quotes it as such (Annex A.7).

These internal uncertainties are smaller than the documented errors of the damage-based reconstruction methods this refit is meant to supplement. On a modern-fleet sample (93 frontal NCAP tests, MY 2005–2007), the standard reconstruction chain underpredicts frontal change-in-velocity by about 25 % on average, traced largely to crush-measurement error propagating into the stiffness coefficients (S30); a single-crash-point estimate can be off by as much as 50 % (S29). Those figures size the problem the envelope addresses, not the envelope itself.

4. Where the force acts: impact height

4.1 Why height matters

The crush and force envelopes say how much a vehicle deforms and how hard it pushes. A barrier designer also needs to know where that force acts, because the bending moment on a post or anchorage is the force times its height above the base. AS/NZS 1170.1 Clause 3.8 prescribes application heights of 0.5 m (Type F) and 1.0 m (Type G) over a 1.5 m length, but the Commentary says nothing about why those heights were chosen: the same gap as the crumple zones, no stated empirical derivation. This section supplies the empirical distribution; §7.3 reads it into a bending moment.

4.2 What the load cells measure

Each cell in a load-cell wall (§2.2) sits at a known height, so the force-weighted average height of the impact resultant is recoverable test by test. The quantity is a variant of the Average Height of Force (AHOF) used in NHTSA vehicle-compatibility research (S9, S10, S11): where that literature's AHOF400 is force-weighted over the first 400 mm of crush, we compute the impulse-weighted centroid of the per-row forces over the full loading stroke, from the per-cell records of the 386 clean load-cell tests, and call it AHOF without implying an exact match.

Of the 386 tests, 377 yield a computed AHOF. Six hi-res tests whose centroids fall implausibly in the upper half of the wall (658 to 832 mm) are excluded; including them would raise the hi-res 95th percentile from 521 to 550 mm, so the exclusion lowers the reported upper tail rather than trimming it toward a more comfortable answer. Of the rest, 258 tests on walls with a citable cell pitch give ground-referenced heights (154 hi-res at 125 mm pitch, 104 on the older 4×9 walls at 246 mm); the other 113, on a 2×3 layout with no citable pitch, are reported as normalised centroids only.

The two wall generations label their rows in opposite directions: the hi-res walls letter from the top down, whereas the published description of the 4×9 walls places row A at the bottom (S9). Getting either order wrong would shift every height on that wall by hundreds of millimetres, so the assignment was verified three independent ways, by cross-wall model agreement, by rebinning, and by the location of the most-loaded row (Annex A.12).

4.3 What the data show

AHOF depends on the wall that measures it (S12), so the two ground-referenced wall classes are reported separately:

Wall class	n	p50 (mm)	p75 (mm)	p90 (mm)	p95 (mm)	% of tests in US bumper zone
Hi-res 125 mm	154	419	447	478	521	49
4×9 246 mm	104	406	467	499	517	42

The two strata agree closely, and the resolution effect runs downward: rebinning the hi-res tests onto the 246 mm grid lowers the per-test reading by a median 26 mm, so coarseness understates rather than inflates the height. The hi-res subset is the more reliable, and every design read-off is taken from it. One documented 4×9 rig places its wall base 66.6 mm above ground (S9); applied across the stratum as a labelled sensitivity it would lift the 4×9 median from 406 to 473 mm, but it is one-rig evidence, so base-at-ground stays as the headline.

The US bumper zone, 406 to 508 mm above ground, is fixed by the federal bumper regulation (49 CFR Part 581; cited in S10) and anchors these heights physically: the hi-res median (419 mm) sits inside it, and every ground-referenced test lands in the plausible bumper and lower-rail band of 300 to 650 mm; a median far outside that band would have flagged an assignment error, and none did.

Phase 1a — impact-force resultant height from barrier load cells




Figure 8: Impact-force resultant height from barrier load cells

Figure 8. Average Height of Force (AHOF) from the clean load-cell tests, by barrier wall class. Left: ground-referenced AHOF, with the 49 CFR Part 581 bumper zone (406 to 508 mm) shaded. Right: normalised force-height centroid (0 = wall base, 1 = wall top).

For barrier design the upper percentiles are the relevant read-off. The AS/NZS 1170.1 Type F application height of 500 mm sits near the 92nd percentile of the hi-res marginal distribution. The Type G height of 1,000 mm is well above any observed AHOF, consistent with its intent for a different engagement scenario.

The distribution is stable across model-year cuts: restricting to MY 2013 or later (n = 120) moves the hi-res median from 419 to 426 mm and leaves the 95th percentile unchanged, so the full modern dataset gives the best estimate; the tests' model-year range (2005 to 2016, median 2012) aligns with the BITRE fleet-age anchor of \$2.3.

4.4 Height conditioned on mass

These marginal percentiles pool all vehicle masses, the wrong summary for design: an engineer sizing against a critical accidental impact or a deliberate hostile one designs to a vehicle selected by the threat or consequence scenario, so AHOF is reported conditioned on mass.

Mass also has the strongest claim among readily available keys: on the hi-res subset the Spearman correlation between mass and AHOF is 0.47 ($p < 0.001$), while speed

shows no detectable relationship across the narrow 40 to 57 km/h span ($p = 0.06$). The table below gives a quantile regression of AHOF on test mass, fitted over the hi-res subset, evaluated at six masses; the edge-of-support rows rest on two and three tests within 150 kg.

Test mass (kg)	q50 (mm)	q75 (mm)	q90 (mm)	q95 (mm)
1,100 (edge of support)	370	371	382	394
1,500	396	407	423	435
2,000	427	453	474	487
2,500	459	498	525	539
3,000	490	544	576	591
3,400 (edge of support)	515	580	617	632

Read this way, the Type F height of 500 mm covers the conditional 95th percentile only up to about 2,100 kg; above that the table exceeds it. Heavier vehicles apply force both harder and higher, so moment demand grows faster than mass alone would suggest; a heavy design vehicle should be read from its own mass row, not the fleet-wide median.

The fit is valid only over the hi-res cohort's mass range, roughly 1,100 to 3,400 kg, the top populated by the full-size pickups of §4.5. Hostile-vehicle-mitigation standards specify threat vehicles up to about 6,800 kg (ASTM F2656); for a heavier design vehicle an engineer should read the highest tabulated row and carry judgement rather than extrapolate: the mass-to-height relationship is not established for classes the cohort does not contain.




Figure 9: Impact-force height versus vehicle mass

Figure 9. Average height of force against vehicle test mass for the hi-res load-cell tests ($n = 154$), with the fitted 50th/75th/90th/95th quantile lines, the 49 CFR Part 581 bumper zone (406 to 508 mm), and the AS/NZS 1170.1 Type F application height (500 mm). The on-figure formulas give the straight-line read-offs. The lines are supported over the hi-res cohort's mass range only.

Does a vehicle dimension explain AHOF better than mass? Overall vehicle height does (NHTSA Canadian Vehicle Specifications database, 137 of the 154 hi-res tests): its correlation with AHOF is 0.70 against 0.54 for mass, and in a bivariate model mass drops out ($p = 0.58$), height being the proximate variable and mass a correlated proxy ($r = 0.75$). Mass stays as the design key anyway. The engineer already selects for mass because it enters the force calculation, and even with height half the variance is unexplained, governed by load-path architecture (engine block, frame rails, battery tray) that no external dimension captures. The height relationship also breaks down exactly where it matters: the Tesla Model S, no taller than a Honda Accord, records AHOF 326 mm, 80 mm below its height prediction, because its stiffest structure is a floor-mounted battery tray. A covariate wrong by 80 mm for the fastest-growing powertrain class is not a reliable design input; mass is coarser but more robust.

4.5 The Australian fleet reading

Cohort composition governs how these numbers transfer to the Australian fleet. Classified by body type, the 154 hi-res tests are 54% passenger cars (median test mass 1,743 kg), 31% SUVs and crossovers (1,944 kg), 10% pickups (2,743 kg), with the remainder vans. AHOF rises with body class, from a median of 391 mm for cars to 440 mm for SUVs and 510 mm for pickups. Excluding the ten full-size United States pickups (the Ford F-250, RAM 2500 and similar, all above 2,500 kg) lowers the marginal 95th percentile from 521 to 509 mm; excluding every vehicle at or above 2,300 kg lowers it to 471 mm.

This is where the cohort and the Australian fleet diverge. BITRE's analysis of the 2018 ABS Motor Vehicle Census (S13) puts average passenger-vehicle tare near 1,450 kg and light-commercial tare near 1,800 kg; test mass runs roughly 150 to 250 kg above tare (dummies and instrumentation), placing the cohort car median (1,743 kg) on the Australian passenger band. The cohort SUV share (31%) sits between the 2018 Australian fleet share (25%) and the new-vehicle sales share (45%), a mildly forward-looking match. The mismatch is confined to the top of the range: the Australian heavy segment is the 2,200 to 2,500 kg mid-size ute (Toyota HiLux, Ford Ranger), and Australia carries no equivalent volume of the 2,700 to 3,400 kg US full-size pickups that populate the cohort upper tail. For an accidental-impact scenario the Australian-relevant range therefore tops out near 2,500 kg, and the table is read at or below that mass; for a hostile-impact scenario the heavy tail remains relevant, the design vehicle being chosen by threat rather than fleet frequency (§4.4).

The fleet is also shifting under these numbers, in two directions. The Australian fleet has been moving toward taller SUVs and light trucks and vans (LTVs), which apply force higher on the barrier, so the upper-percentile AHOF values may become more representative over time, not less. The shift is modest, though, because AHOF is impulse-weighted, not geometric: an SUV's tall upper body is comparatively soft structure, while the bumper beam, frame rails, engine block and subframe that carry the bulk of the loading impulse sit at roughly the same height across vehicle types, constrained by ground clearance, suspension geometry and the Part 581 zone. A 3,300 kg SUV might read 50 to 80 mm higher than a 1,200 kg hatch, not 300 mm, consistent with the vehicle-compatibility literature (S9, S16): the stiff structures have not moved up as far as the roofline has.

The second shift runs the other way. Battery-electric vehicles carry their mass in a floor-mounted pack, so their stiffest structure sits lower than a combustion vehicle's of the same roofline (S14, S15). The cohort can count them but not yet characterise them: four sit in the hi-res height set, all at or below their mass-predicted median, and six in the measured-crush fit, straddling the crush median. Both counts are too small to carry a fleet-level claim, so the reading stays directional, and for the mass key it errs

safe, because mass over-predicts the AHOF of a floor-pack vehicle. Battery-electric vehicles remain near 1.6% of registered Australian passenger vehicles in early 2025, though growing by more than half year on year (S7): a forward-looking caveat rather than a present-day correction.

4.6 What this changes for a barrier

The structural implications go beyond the bending moment. Base shear is unchanged by height: the horizontal force is the same whether it acts at 0.5 m or 1.0 m. But a barrier designed to AS/NZS 1170.1 typically has a horizontal rail or strong beam at the prescribed height, sized accordingly. Force arriving lower may bypass that rail and engage a post or wall section not designed for direct loading, losing the rail's force-sharing across posts. A lower height of force reduces the moment, which is favourable, but it changes the load path in ways the fixed-height assumption does not contemplate: capacity should be checked at the actual engagement height, not just the prescribed one.

The low-speed car-park case needs the same caution in the other direction. The load-cell tests span 40 to 57 km/h; nothing in them licenses extrapolation to the low-speed impacts the Type F case describes, where a vehicle may engage only its bumper crossbeam. Stated without a number: a bumper-only engagement concentrates load toward the lower part of the Part 581 zone rather than the higher stiff load paths that raise AHOF at test speed, placing the effective height at or below the bottom of that zone, below the 500 mm Type F prescription. The dataset does not sample this regime; this is a directional expectation, not a tested value.

4.7 Caveats and bias direction

Three caveats apply to all of these numbers, and the first two share a direction. The load cells in essentially all 386 tests are single-axis, measuring X-force only (moment-measuring cells exist for six tests), and NHTSA's own finding is that single-axis cells underestimate the height of force (S10, S17). The base-elevation assumption, that the wall base sits at ground level, is corroborated by the bumper-zone check in §4.3 but not confirmed from test records (the one documented exception, the 66.6 mm 4×9 rig clearance, is carried there as a sensitivity). Both effects push the true AHOF upward, which is not conservative here: the reported percentiles understate the moment demand and should be treated as a floor. No figure is attached to the bias: the single-axis studies establish its direction, not a magnitude transferable to these vehicles and wall geometries, so the margin above the floor is an engineering-judgement decision.

The third caveat concerns combination: pairing an upper-percentile force with an upper-percentile height multiplies two marginal percentiles from different distributions, since the same vehicle need not sit at both tails at once; the product is a deliberately conservative bound, not a joint percentile (§7.3 carries it into the worked moment). Finally, these AHOF values describe force against a flat rigid wall at ground datum: a kerb or wheel-stop raises the effective engagement height, and a rail at a fixed height changes the engagement geometry, so the datum should be checked against the actual installation.

5. The crush ceiling: a bounded envelope beyond the tested range

The validity envelope (§6) flags everything above about 57 km/h as extrapolation. For most quantities that flag is the end of the story: the data do not reach there, so no claim is made. The high-speed crush tail is the one exception where the shape of the response can still be bounded, and bounding it changes how conservative the force treatment is at the top end. This section sets out that bound and the evidence for it; the

power law of §3 remains the in-band relationship, and nothing here is read into a design calculation inside the supported band.

5.1 Why the high-speed tail matters

The deployed crush model is a power law in speed, and a power law has no ceiling: it lets crush grow without limit as speed rises. That is a fair description inside the tested band, where crush stays well below the size of the vehicle. Extended far enough, though, it predicts crush larger than the car itself. A vehicle cannot crush more than its own length, and long before that limit the front structure is exhausted and the occupant compartment begins to collapse. This two-stage collapse is not only a physical expectation. A review of more than 1,400 frontal rigid-barrier tests reports that the force-crush response is linear only to about 56 km/h and then turns bi-linear between roughly 56 and 80 km/h, as the front structure saturates and the stiffer occupant compartment begins to engage, the mechanism first set out in the foundational crush-energy literature (S29, tracing to S31). That literature observes the transition only to about 80 km/h, far below the speeds treated here, so it supports the *existence* of a stiffening transition, not its *location*. The true crush-versus-speed curve must therefore bend away from the power law somewhere above the tested range. The engineering question is not whether a ceiling exists but where it sits and how sharply the curve bends toward it.

This matters for design because force and crush are reciprocal on a rigid barrier. For a given impact energy, less crush means more average force, from the same identity $F = mv^2/(2d)$ used throughout this paper. A model that lets crush rise unchecked under-predicts force at high speed, which is non-conservative; one that caps crush too low over-predicts force, conservative but wasteful. Bracketing the true crush ceiling therefore brackets the high-speed force.

5.2 What the crash-test data can and cannot fix

We tested whether the cohort itself can locate a ceiling by fitting two saturating curves, a logistic and a Gompertz, alongside the power law on the same measured data; three findings result.

First, inside the tested band the three forms are statistically indistinguishable. On the usual model-selection criterion the gap between them is below the threshold at which the data can prefer one over another. This is expected: the tested speeds sit on the near-straight middle of any S-curve, before it bends, so the data cannot see a ceiling from there.

Second, left free to find their own ceiling, the saturating curves do not find one. The fitted ceiling runs away to physically meaningless values of tens of metres, and the curves simply track the rising power law instead of bending over. Saturation has to be imposed as a physical constraint; the cohort cannot discover it. The external evidence for a bi-linear stiffening transition (S29, S31; §5.1) confirms such a constraint is physically real, but it fixes the mechanism, not the knee speed, so it does not lift this limitation.

Third, when a ceiling is imposed, the saturating curve bends the high-speed crush below the power law at a small in-band cost. Higher imposed ceilings cost progressively less in fit quality, and every imposed-ceiling variant stays statistically tied with the power law where the data live.

The consequence, which governs all the wording that follows, is that a bounded curve is data-compatible, not data-preferred. The crash data neither favour a saturating curve over the power law nor tell us where the knee is. What follows is a physically-motivated bound, tagged as such, not a claim that the data show crush follows an S-curve.

5.3 An external anchor for the far tail

One observation from outside the cohort constrains the far tail. A televised crash test drove a small hatchback (a Ford Focus, about 4.18 m long and roughly 1,150 kg as tested) head-on into a solid concrete block at a nominal 120 mph, about 193 km/h. Three independent estimates of the crush converge on about 2.0 m, with an uncertainty near ± 0.25 m after a photogrammetric scale check against the known vehicle length in the footage: kinematic (the reported 68 ms stopping time), geometric (crush reaching back to the B-pillar), and energy (impact energy over a plausible mean force). The test mass is inferred rather than measured, and the speed is a nominal public figure rather than an instrumented reading.

A single, roughly-read data point would be worthless for fitting a curve; it is valuable here because at this speed the candidate forms disagree by metres. The raw power law predicts about 4.7 m of crush, which exceeds the length of the car and is impossible independent of the video. Imposed ceilings below 1 m predict well under the observed value. Only ceilings near 2 m and above reproduce it. An observation with ± 0.25 m uncertainty easily separates hypotheses that differ by metres: the point serves as a discriminator, not a fitted point. Appending it to the fit, even weighted as though it were fifty tests, moves the in-band curve (at or below 100 km/h) by less than the fit's own scatter, so it constrains the tail and not the design range.

The configuration is what makes the observation usable: a head-on, full-width impact against an effectively rigid block, exactly the configuration the cohort lacks above about 65 km/h, where every high-speed frontal test in the corpus is against a deformable barrier whose energy absorption confounds the crush reading. The vehicle predates the modern cohort, and it is a single uninstrumented test, so every use of the anchor carries those caveats.

5.4 The bounded envelope

Rather than draw one saturating curve, which would imply a precision the data do not support, Figure 10 draws the envelope of every bounded curve that is at once tied with the power law inside the tested band and passes through the extreme-speed observation. Both saturating families are included so the result reads as boundedness rather than an artefact of one functional form.

The envelope has three readable features. Where the data live, the bounded curves and the power law coincide, so the deployed model is unchanged in-band. In the far tail, near the extreme-speed observation, the bounded curves converge again onto the observation, roughly 1.9 to 2.6 m of crush at the 1,500 kg reference mass. Between the two, across the roughly 60 to 190 km/h range where there is no data, the envelope is wide, spreading to about 350 to 480 mm by 144 to 180 km/h. That width is the honest measure of what is not known: the ceiling magnitude is bounded at both ends, but the knee speed is under-determined. The transition is a band, not a line.

The ceiling scales only weakly with vehicle mass. Re-evaluated across weight bands from 1,100 to 3,400 kg (Figure 11), the bounded ceiling rises from roughly 1.7 to 2.3 m at the light end to roughly 2.4 to 3.3 m at the heavy end, while the power-law tail grows implausibly large at every mass. The weak mass dependence is consistent with the in-band finding that mass acts largely as a vehicle-size proxy (§3.1, Annex A.4).

Bounded crush envelope: two saturating fits threaded through the same constraints (reference mass 1500 kg)
 Shaded = every bounded curve that is statistically tied to the power law where data exist ($\Delta AIC \leq 2$) and also passes through the ~ 2.0 m crush measured in a 120 mph rigid-block test.
 The power law coincides with them in-band and grows past the vehicle length in the tail. The band width above ~ 100 km/h (no data) is the under-determined transition: the ceiling magnitude is evidenced, its onset speed is not.




Figure 10: Bounded crush envelope at reference mass

Figure 10. The bounded crush envelope at 1,500 kg reference mass. Each shaded band spans every saturating curve, across two functional families, that is statistically tied with the power law where the cohort data exist and also passes through the extreme-speed observation at about 193 km/h. Left: the full range out to that observation. Right: the in-band zoom, where the bounded curves and the power law coincide. The power law (solid) matches the data in-band but grows past the vehicle length in the tail. The wide band between the data ceiling and the observation is the under-determined transition: the ceiling magnitude is bounded, its onset speed is not.

Bounded crush envelope across vehicle weight bands
The bounded envelope re-evaluated at each mass; the ceiling scales only weakly with mass. The power law stays the in-band model and grows implausibly large in the tail at every weight band. Star = the ~2.0 m extreme-speed observation scaled to each mass.




Figure 11: Bounded crush envelope across weight bands

Figure 11. The bounded crush envelope re-evaluated across vehicle weight bands from 1,100 to 3,400 kg, using the fitted mass exponent with no refit. The bounded ceiling scales only weakly with mass, from roughly 1.7 to 2.3 m at 1,100 kg to roughly 2.4 to 3.3 m at 3,400 kg. The power law (solid) grows implausibly large in the tail at every mass. The star marks the extreme-speed observation scaled to each mass. In-band data are sparse at the lightest and heaviest bands, so those panels are the most extrapolated.

5.5 What this changes, and what it does not

Crush does not grow without limit; it knees toward a ceiling bounded to roughly 0.9 to 1.7 m on the corpus-derived dynamic statistics and anchored near 2.0 m at extreme speed by the external observation. Because capping crush raises the implied force, treating the high-speed tail as bounded makes the force treatment conservative at the top end, which is the direction a barrier designer wants.

Three limits gate every use of this result. As a plateau the bounded form is valid only to about 80 to 100 km/h; beyond the knee the true curve resumes rising as the occupant compartment collapses, and the external observation itself shows about 2 m of crush at 193 km/h, far above any plateau, so the flat part must not be read at

extreme speed. The ceiling onset speed is not established from data and remains a physical argument only. And the whole construction is data-compatible, not data-preferred: it is the shape that expresses a physical bound, not a claim that the crash data select a saturating curve over the power law. The deployed in-band model, the design charts, and the worked examples are unchanged by this section.

6. Validity envelope

The model is offered for research use and interpretation only inside the following envelope, and not as a design basis even inside it. Everything outside is flagged as outside the tested range; the one exception, bounded rather than merely flagged, is the high-speed crush ceiling of §5.

Dimension	Supported	Basis
Velocity	33–57 km/h (9.1–15.8 m/s)	Measured-only fit data band; clusters at 40/48/~56 km/h
Mass	1,007–3,448 kg	Measured-only fit mass range
Impact	Head-on, full-width rigid barrier	Scope filter + API-confirmed rigidity
Vehicle era	MY ≥ 2004	Cohort cutoff

Explicit exclusions:

- **Below ~30 km/h.** No confirmed-rigid MY ≥ 2004 test exists below 30 km/h, and no defensible modern low-speed anchor exists (§10 item 1). The power-law form goes through the origin, physically correct for dynamic crush but unconstrained by data below the floor; restitution dominates below ~16 km/h (S18). Very-low-speed force values should be treated as artefacts of the form, not predictions.
- **Narrow impactors (poles, bollards).** The envelope is for full-width engagement. A pole engages the structure locally and crushes it more deeply than a full-width wall at the same speed, so published velocity-to-crush ratios for rigid poles are materially lower (S19, S4). Applying a full-width crush value to a narrow impactor is therefore conservative for force but non-conservative for penetration depth. The direction is noted only to bound the error; narrow-impactor crush is a different regime.
- **Heavy vehicles (> 3,448 kg).** Rows for HVM threat vehicles (6,800–30,000 kg) extrapolate mass 2–8× beyond the fitted fleet, using an exponent estimated on passenger cars and LTVs; heavy trucks are ladder-frame, cab-over or bonneted structures that do not resemble that fleet. Such rows are illustrative of the method, not design values, until heavy-vehicle FWRB or HVM test data anchors them.
- **Oblique, offset, and deformable-barrier impacts.** Out of scope by construction.

7. Using the envelope

Everything in this section is a research reading of the fitted envelope: something an engineer can compute and weigh, not a characteristic design action (§8.4, §10 item 10). The readings stay inside the validity envelope of §6 except where a flag says otherwise.

7.1 Read-off formulas

For hand calculation the fitted envelope reduces to two read-off formulas, each valid only inside the supported window. Crush is

$$d \text{ (mm)} = a \cdot v^{1.27} \cdot m^{0.18} \text{ (v in m/s, m in kg)}$$

with $a = 4.09$ at the 5th percentile, 5.20 at the 50th, and 7.19 at the 95th. Rounding the exponents to two places holds the crush estimate within about 2 % of the full fit;

carrying the fitted slopes (1.272 and 0.182) removes even that. Impact height (§4.4) reads off a line in mass:

$$\text{AHOF } q_{95} \text{ (mm)} \approx 280 + 104 \cdot m_t \text{ (} m_t = \text{vehicle mass in tonnes)}$$

which sits within about 0.3 % of the quantile fit. The crush formula applies over $v = 9.1\text{--}15.8$ m/s (33–57 km/h) and $m = 1,007\text{--}3,448$ kg; the AHOF line applies over the hi-res cohort's roughly 1,100–3,400 kg range. Average force then follows from the energy identity $F_{\text{avg}} = mv^2/(2d)$ (§3.3). These formulas are the paper's primary read-off artefact, and the design charts (§3.2, Figures 4 to 6) give the same numbers graphically.

7.2 Worked example: 2,000 kg at 50 km/h

A worked illustration in the style of an AS 1170.1 annex, on a rigid barrier: $m = 2,000$ kg, $v = 50$ km/h = 13.89 m/s, both inside the supported envelope. Crush uses the measured-only primary fit (§3.1); the pooled-hybrid sensitivity fit is more conservative (higher force) here and brackets the answer on the safe side (Annex A.2, A.10).

Step 1: Crush bounds (the §7.1 read-off formula evaluated at $a = 4.09/5.20/7.19$):

Percentile	Crush d
5th (stiff, force governs)	463 mm
50th (median)	589 mm
95th (soft, penetration governs)	814 mm

Step 2: Kinetic energy.

$$\text{KE} = \frac{1}{2} \cdot 2,000 \cdot 13.89^2 = \mathbf{192.9 \text{ kJ}}$$

Step 3: Average force bounds ($F = \text{KE} / d$):

Percentile	Crush	F_avg
5th	0.463 m	416 kN
50th	0.589 m	327 kN
95th	0.814 m	237 kN

Step 4: Readings.

- **Force-demand reading (barrier strength, anchorage):** the stiff-vehicle 5th-percentile average force is 416 kN; peak force is empirically $\sim 1.9\times$ the average (§3.3).
- **Penetration (rigid barrier):** on a rigid barrier the penetration is the 95th-percentile vehicle crush, 814 mm.
- **Deformable barriers: what this envelope can and cannot say (scope limit).** The paper's scope is rigid barriers. For a barrier that deflects, vehicle and barrier form a coupled system that a single rigid-wall crush quantile cannot solve, and two asymmetric statements are all the data supports. (a) Travel: rigid-wall crush plus barrier deflection is a defensible upper bound on total travel, because a vehicle absorbs less energy against a compliant barrier, so its actual crush is smaller. (b) Force: the interface force against a compliant barrier is governed by the softer element and cannot be recovered from a rigid-wall crush quantile at all. Solving the coupled problem needs validated vehicle force–displacement families: a separate research programme.

For context, 416 kN average force is $\sim 1.7\times$ the AS 1170.1 ramp-end load (240 kN), for a scenario (50 km/h) far beyond anything Clause 3.8 contemplates and well below HVM certification speeds: precisely the gap region the envelope is built for.

7.3 The bending-moment combination

The impact-height results (§4) are meant to be read into a calculation. Select the design vehicle by mass and speed from the threat or consequence scenario; read its average force from the force envelope (§3.3); read the AHOF for the same mass from the §4.4 table, at the percentile the target reliability calls for; form the bending moment as force times height above the barrier base; and apply the §4.7 bias allowance, because the reported heights are a downward-biased floor, not a central estimate.

A worked combination shows the size of the numbers. For a 2,000 kg design vehicle at 50 km/h, the stiff-vehicle 5th-percentile average force from §7.2 is 416 kN and the conditional 95th-percentile AHOF from the §4.4 table is 487 mm, so the bending moment at the base is 416 kN times 0.487 m, about 203 kN·m, before the peak-force multiplier (§3.3, roughly 1.9 times the average) and the height-bias allowance are applied. Pairing the upper-tail force with the upper-tail height is the deliberately conservative bound of §4.7, offered as a demand ceiling for a first sizing pass.

7.4 Checks an engineer can run

The checks below let a structural engineer test the envelope against their own experience or code requirements: “something to check,” not assertions to accept. All values read from the measured-only design charts (q05 = the stiff-vehicle, force-governing curve; peak = $\times 1.9$ the average, §3.3); where a scenario leaves the validity envelope the governing flag is stated.

(a) A car-park barrier above the speed Clause 3.8 tabulates: 2,000 kg at 36 km/h. The standard’s most severe car-park case (the ramp end) is 2,000 kg at 21.6 km/h $\rightarrow$ 240 kN, yet a vehicle running away down a descending ramp can reach appreciably higher speeds. At 36 km/h (inside the supported band):

Reading	Value
Crush (q05, stiff)	304 mm
Average force (q05)	329 kN
Peak force (q05 $\times$ 1.9)	$\approx$ 625 kN

The q05 average force (329 kN) already exceeds the 240 kN prescribed for the much slower ramp-end case, and the peak the anchorage must resist ($\sim$ 625 kN) is $\sim$ 2.6 $\times$ the prescribed value. An engineer can weigh this against the capacity of a barrier and its fixings detailed only to Clause 3.8.

(b) The three scenarios the standard itself tabulates. The Commentary (Supp 1, Clause C3.8) gives the mass, speed, and assumed crumple zone behind each prescribed force, so its three scenarios run through the model with no added assumptions. This is a contextual comparison, not an independent validation: standard and model use the same energy formula, so agreement is arithmetic consistency rather than corroboration. One difference matters throughout: the Commentary’s Δ includes barrier deflection where the model’s crush is vehicle-only (§3.3), so the model returns a higher force wherever the standard credits barrier deflection.

Only one Commentary scenario sits near the model’s supported band, the ramp end (2,000 kg at 21.6 km/h, assumed $\Delta = 0.15$ m, prescribed 240 kN), and it is the one worth reading:

Read-off (2,000 kg, 21.6 km/h)	Crush	F_avg
q05 (stiff vehicle)	159 mm	227 kN

Read-off (2,000 kg, 21.6 km/h)	Crush	F_avg
q50 (median vehicle)	202 mm	178 kN

The stiff-vehicle q05 crush lands close to the Commentary's assumed 0.15 m, and the q05 average force within about 5 % of the prescribed 240 kN; the median vehicle crushes further, which would lower the implied force. The Commentary's two light-traffic scenarios sit at 7.2 km/h, far below the 30 km/h data floor, where the origin-constrained power law drives crush toward zero as speed falls and so overstates force: the model returns roughly twice the prescribed 30 and 40 kN there, an artefact of the fitted form rather than an evidenced prediction, and those readings must not be quoted as findings (§6, §10 item 1).

The value of the comparison is directional: the same formula and empirical crush envelope reproduce the right order of magnitude at the one Commentary scenario near the supported band, then extend, with their uncertainty stated, into the 33–57 km/h band where the standard is silent; that extension is what this paper proposes.

8. Broader applicability: gaps in other barrier design methods

The vehicle-crush gap this paper addresses is not unique to AS/NZS 1170.1. The same energy-balance formula underlies barrier design methods in multiple codes and research traditions, and in all cases the vehicle-crush component (δc) is the weak link. The dataset addresses three gaps identified in the literature:

8.1 The empirical crush update that Iqbal's method has called for since 2008

Iqbal (S20, S21) proposed the same energy-balance formula for barrier design, with vehicle crush approximated as $\delta c = \sqrt[3]{(v/3)}$ feet (v in mph), derived from “a limited survey” of NHTSA tests at 35 mph. Both articles state that “the car crush data need to be updated accordingly” as vehicle designs change; the CMAA summary by Williams (S22) reproduces the formula and repeats the call. BS 6399-1:1996 and EN 1991-1-7:2006 use the same $F = mV^2/(2\Delta)$ structure, with δc fixed at 100 mm unless “better evidence is available.”

The Crumple model is that better evidence. The power-law fit from the modern crash-test cohort replaces the single-variable $\sqrt[3]{(v/3)}$ formula with a mass-dependent, statistically bounded envelope, and the quantile bands address the fleet variability that Iqbal and Williams both flag but neither quantifies.

8.2 Statistical calibration for reduced vehicle models in academic research

Xiang et al. (S23, S24) and He et al. (S25) developed reduced vehicle models (elastoplastic springs representing vehicle stiffness and energy absorption) for analysing parking structure columns under impact, calibrated from as few as three NHTSA tests: a 2001 Toyota Echo, a 2002 Ford Explorer, and a 1992 Chevy C1500 pickup (S23). The authors introduce a “vehicular energy absorption ratio” to improve Eurocode 1's rigid-impact assumption, but the ratio itself is derived from the same three tests.

The Crumple dataset provides modern data for statistical calibration of these spring properties: a fleet-level stiffness distribution with confidence bands rather than hand-picked tests.

8.3 Validation benchmark for finite-element vehicle models

Li et al. (S26) states the vehicle-model gap directly: rigid or spring simplifications “fail[] to accurately replicate vehicle impacts”, so “the use of detailed vehicle models is imperative.” Their LS-DYNA simulations use detailed FE vehicle models, but validation against empirical crush data is not reported.

The Crumple dataset offers an independent validation benchmark: an FE vehicle model that predicts crush within the empirical quantile envelope at the same mass and speed is population-plausible; one that lands outside it is not. This is a population plausibility check, not a pass/fail calibration (calibration requires matching a specific vehicle’s force history, intrusion, and energy partition), and it is a use the dataset can serve without prescribing how the engineer designs the barrier.

8.4 Australian standards and reliability context

Three further documents frame where this envelope sits for an Australian structural audience, and what it is not:

- **AS/NZS 1170.0:2002 (general principles): the reliability framework any “design value” must live in.** A code-ready characteristic action requires a defined target reliability, load-combination role, consequence class, and model-uncertainty allowance. A 5th-percentile crush value is not automatically the correct percentile for a characteristic force, and AS/NZS 1170.0 is the document against which a future calibration of this envelope would have to be argued (§10).
- **EN 1991-1-7:2006 Annex C: the nearest codified analogue.** Annex C gives an equivalent-static vehicle-impact force model with published vehicle stiffness values and the same $F = \frac{1}{2}mv^2/\delta$ structure (100 mm default crush unless better evidence exists); the update §8.1 describes answers its default as well.
- **AS/NZS 3845.1/2, AASHTO MASH, and the ANZCTC guidance: the adjacent barrier worlds, scoped out.** Road-safety-barrier containment and the Australian *Hostile Vehicle Guidelines for Crowded Places* already codify crash-test/containment frameworks. This paper does not enter them: it models rigid-barrier vehicle response, not barrier containment performance; the ANZCTC/HVM certification route (ISO 22343-1:2023) remains the practitioner’s reference for hostile-vehicle work.

8.5 Quantified design margin at q05

How the layers stack

The independent load-cell check (§3.3(iii)) is made at the model median (q50), where the model sits a small distance below modern-fleet force. Force demand, however, is read off the stiff-vehicle q05 curve, so the design-relevant question is whether q05, with the standard’s own factors, brackets the force a real modern vehicle exerts, not whether q50 is biased. The layers stack as follows (median values; force $\propto$ 1/crush throughout):

Layer	Effect on force demand	Basis
Read at q05, not q50	$\times 1.272$ (+27.2 %)	the q05/q50 crush ratio is fixed across every point of the envelope
Model q50 vs modern load-cell force (the bias)	± 0.921 (distance-avg) / ± 0.835 (time-avg)	independent load-cell force validation (§3.3(iii))
Net: model q05 vs modern-fleet average force	$\times 1.17$ (distance) / $\times 1.06$ (time)	product of the two rows above
Peak conversion: model $\times 1.9$ vs reality $\times 1.93$	$\times 0.98$ ($\approx$ neutral; the peak factor is a demand, not a margin)	LC peak/average 1.93
AS/NZS 1170.0 ultimate-limit-state (ULS) action factor	$\times 1.2 - 1.5$	load-combination factor on the action
Capacity-reduction factor (φ)	further margin on the resistance side	AS/NZS 1170.0 / material standard

Layer	Effect on force demand	Basis
Combined (before ϕ)	$\times 1.25 - 1.72$	product of the rows above

Read together, selecting q05 more than offsets the q50 bias, and the ULS action factor supplies the working margin: $\times 1.25$ (worst case: time-average, minimum factor) to $\times 1.72$, with ϕ on top. Whether that is adequate is a judgement against the consequence class and the governing reliability framework (§8.4).

What could break the margin

Two caveats keep this a margin for the engineer to close, not a property the paper certifies. First, it is a margin measured at the middle of the fleet. Against the stiffest vehicles themselves (the 5th percentile of the per-test load-cell ratio), the q05 reading lands only at parity ($\times 0.99$, distance) or about 12 % below ($\times 0.88$, time); for those vehicles it is the ULS load factor, not the choice of percentile, that provides the margin. Second, the whole argument assumes the load-cell bias lowers the stiff q05 curve by the same proportion as the median, which holds only if the modern, and Australian, fleet has the same spread of vehicle stiffness as the cohort. If today's fleet carries proportionally more very stiff vehicles (plausible, given the unquantified US-to-Australian model-mix differences, §10 item 3), q05 would sit lower than assumed and the load factor alone might not close the gap. This is the single assumption the force reading in §3.3(iii) and Annex A.9(iii) point back to; for these reasons the margin is laid out for the engineer to weigh, not certified (§8.4, §10 item 10).

Design-margin ladder at q05




Figure 12: Design-margin ladder at q05

Figure 12. The design-margin ladder, reading the layers of the table above as a cumulative force margin relative to the modern-fleet average force. The model median (q50) sits below the fleet average; reading at the stiff q05 curve ($\times 1.272$) lifts it just above; the AS/NZS 1170.0 ULS action factor ($\times 1.2$ to 1.5) supplies the working margin. Each bar spans the time-average load-cell ratio (low edge) to the distance-average (high edge). Combined before the capacity-reduction factor: $\times 1.25$ to 1.72 .

9. Research artefacts

The deliverable is a register of primary artefacts, each issued as a printable sheet in Annex C:

Primary artefact	What it gives	Where	Annex C sheet
Crush read-off formula	crush quantiles in closed form at any in-window mass and speed	§7.1	restated on every sheet
Crush-versus-speed envelope	the five quantile curves against the cohort data	Figure 4	A
Per-mass envelope panels	the same envelope at six design masses	Figure 5	B1–B6
Crush contours	stiff-vehicle (q05, force side) and soft-vehicle (q95, penetration side) read-offs over the supported mass and speed window	Figure 6	C1, C2
Impact-height line (AHOF versus mass)	force application height conditioned on mass, with the q95 line in closed form	§4.4, §7.1, Figure 9	D
Peak-to-average factor	the empirical $\times 1.9$ multiplier from average force to peak	§3.3, Annex A.11, Figure 7	—
Design-margin ladder	how q05 with the standard's own factors sits against modern-fleet force	§8.5, Figure 12	—

The impact-height line is named a primary artefact deliberately: it supplies the height component of the bending-moment reading (§7.3), which neither AS/NZS 1170.1 nor the HVM standards derive empirically. All in-band artefacts are drawn on the measured-only fit; the hybrid fit is a labelled sensitivity set. Supporting figures complete the set: the impact-height distribution by wall class (Figure 8), the bootstrap ribbons (Figure A1), the crush-validation scatter (Figure A2), the load-cell force validation (Figure A3), and the bounded-envelope figures (Figures 10 and 11, drawn on the saturating fits since they describe the extrapolated tail). No energy-partition figure is offered: a single rigid-barrier crush quantile cannot resolve the coupled vehicle and barrier energy split (§7.2).

The validity windows of §7.1 are the paper's supported envelope; outside them a value is an extrapolation of the form, flagged as such throughout.

10. Limitations and data calls

Limitations (stated, quantified where possible):

1. **The low-speed gap is real and unfillable from public US data.** Zero confirmed-rigid MY ≥ 2004 tests below 30 km/h; equivalent-barrier-speed translation from other configurations adds zero modern rows; no defensible modern Campbell- b_0 anchor exists. The reconstruction literature offers a modern low-speed intercept of roughly 6 km/h (S28), but that is a curve-fit

value, not a rigid-barrier measurement. The origin-constrained form overestimates crush at very low speed, where a genuine no-damage threshold exists (traced to the foundational crush-energy work at about 5 mph, S31; FMVSS Part 581 bumpers): conservative for penetration, but it makes the low-speed force values an artefact.

2. **Velocity coverage is clustered** at three regulatory speeds; interpolation between clusters is model-driven (§2.3).
3. **US fleet, not Australian.** Age-matched via ABS weights; model-mix differences unquantified. The weighted-fit sensitivity shifts all quantiles down by 4–7 %, a small but real adjustment, not within noise (§3.4, Annex A.6). Separately, the q05 force reading rests on the fleet-stiffness-spread assumption set out in full at §8.5, which remains unquantified.
4. **Average force, not peak.** Load-cell peak forces are empirically a median 1.9× the energy-balance average (IQR 1.8–2.2, over the clean load-cell set); anchorage design must account for this.
5. **Static-crush data quality in the corrective path.** The NHTSA static-crush field has known quality issues in older tests (the QC flagged 364 physically impossible dynamic<static ratios). The measured-only model bypasses static crush entirely; the corrected-static path, which carries this exposure, enters only as a labelled sensitivity.
6. **Heavy-vehicle rows are illustrative** (§6). The method extends naturally; the data does not yet.
7. **Rigid-barrier scope only.** The coupled deformable-barrier problem is not solved here: the only defensible carry-over to compliant barriers is an upper bound on total travel, and force demand against a compliant barrier is out of scope (§7.2).
8. **Provenance: measured-only is primary.** The design charts and figures are built on the measured-only fit (§3.1). The corrected-static path is a labelled sensitivity, demoted for a small but significant level bias, large per-row scatter, and an unreconciled velocity exponent (§2.2, Annex A.10); the pooled hybrid fit is retained as a sensitivity set and quoted in the §7.2 worked example as a conservative bracket.
9. **Every measured crush is a band, not a point.** The impact-onset detection and the pinned initial speed give a worst-case per-test uncertainty of median ~5 % (95th percentile ~49 %), larger below 30 km/h, with the detection threshold dominant (§3.4).
10. **This is not a design basis.** The output is an empirical research response envelope, not a code-ready characteristic action. An independent force validation now exists (§3.3(iii)) and shows the design-relevant q05 reading brackets the modern fleet with a small margin the standard's load factors are expected to cover; but converting the envelope to a characteristic action still requires the AS/NZS 1170.0 reliability calibration (target reliability, load-combination role, model-uncertainty allowance; §8.4), which this work does not perform.
11. **The high-speed crush ceiling is bounded but not located.** Beyond the tested band crush must bend toward a physical ceiling, bounded to roughly 0.9 to 1.7 m on the corpus-derived statistics and anchored near 2.0 m by a single external observation (§5). The bounded form is data-compatible, not data-preferred; it is valid as a plateau only to about 80 to 100 km/h, and its knee onset speed is not established from data.

Data calls: what would most improve this envelope.

1. **Barrier load-cell force histories: addressed.** The modern NHTSA fleet was processed to the clean load-cell set (32.7–57.0 km/h); the load-cell-vs-accelerometer crush cross-check gives a median ratio 0.981 (§3.3). Sample size is limited mainly by a server-side serving issue (157 of 274 unprocessable tests return byte-identical per-channel files despite distinct request URLs; a further 110 carry no load-cell channels). The loss is largely unrecoverable, not a selection choice.
2. **Bollard/narrow-impactor test data from barrier suppliers:** crush and force from certified ISO 22343-1:2023 (formerly PAS 68 / IWA 14-1) and ASTM F2656/F2656M-23 tests, most of which are commercially held and unpublished.
3. **Euro NCAP / ANCAP FWRB records:** these widen fleet coverage toward the AU-relevant model mix, though their 50–55 km/h speeds do not lower the velocity floor.
4. **Modern low-speed rigid-barrier tests** (e.g. structural tests at 15–30 km/h adjacent to those of the Research Council for Automobile Repairs, RCAR): the only route to an evidenced low-speed anchor.
5. **Heavy-vehicle FWRB or instrumented HVM tests:** to anchor the mass extrapolation.

11. Disclosures

- **Research disclosure.** This is independent, unfunded, single-author research. Generative-AI tools were used, under the author’s direction and review, for literature search and triage, code and data-pipeline development, figure generation, and copy-editing. All analytical decisions, interpretation, and final wording are the author’s own. Every quantitative claim is traceable to the reproducible scripts and source data described under Data availability. The author is solely responsible for the content.
- **NHTSA attribution.** Crash-test data are drawn from the U.S. DOT/NHTSA Vehicle Crash Test Database; NHTSA’s terms of use and disclaimer apply, and NHTSA is not responsible for any analyses or conclusions drawn from the data here.

Annex A: Statistical methods

This annex supplies the full statistical detail behind the methods summarised in §§3–4. It is written for a reader who wants to verify the fitting procedure, the uncertainty quantification, or the force-validation arithmetic.

A.1 Quantile regression and shared-slope design

Dynamic crush is fitted as a power law:

$$d \text{ [mm]} = a \cdot v^{b_v} \cdot m^{b_m} \text{ (} v \text{ in m/s, } m \text{ in kg)}$$

by quantile regression at $q = 0.05, 0.25, 0.50, 0.75, 0.95$ on log-transformed data (statsmodels QuantReg). The velocity and mass slopes (b_v, b_m) are fitted once, at the median ($q = 0.50$), and shared across all five quantile curves. Each quantile then receives its own intercept from the residual distribution. Because the five curves differ only by a multiplicative constant, they cannot cross, which is a structural requirement for any table read at extrapolated speeds.

A.2 Three-provenance fit table

Provenance (median fit)	Role	a	b_v	b_m	R²(log)	n
Measured-only	PRIMARY	5.202	1.272	0.182	0.597	614
Hybrid (measured	sensitivity	1.865	1.477	0.245	0.576	1,072

Provenance (median fit)	Role	a	b_v	b_m	R ² (log)	n
+ corrected)						
Corrected-only	sensitivity	0.136	1.849	0.460	0.565	458

The measured and corrected velocity exponents differ by +0.58 (1.27 vs 1.85). An adequate global static-to-dynamic correction would have reconciled them; it does not. Because the correction ratio(v) = $A + B/v$ is a decreasing function of v , the corrected subset inherits static crush's steeper velocity dependence. The corrected path is therefore a different population, not a substitute for the measured one.

A.3 Per-quantile parameter table (measured-only, shared slopes)

Quantile	a	b_v (shared)	b_m (shared)	pseudo-R ²
5th	4.089	1.272	0.182	0.418
25th	4.723	1.272	0.182	0.459
50th	5.202	1.272	0.182	0.362
75th	6.004	1.272	0.182	0.300
95th	7.185	1.272	0.182	0.213

A.4 Exponent interpretation

Two limiting cases bracket the velocity exponent b_v : a constant-force (rigid-plastic) structure gives $d \propto v^2$ ($b_v = 2$); a linear spring gives $d \propto v$ ($b_v = 1$). All three fitted velocity exponents (measured 1.27, hybrid 1.48, corrected 1.85) sit inside [1, 2], and the measured-only primary fit lies nearest the linear-spring limit, consistent with progressive stiffening of modern front structures (force rises as the structure crushes).

The mass exponent $b_m < 0.5$ (measured 0.18) implies stiffness rises with mass, consistent with the fleet-design literature and with independent findings that mass has limited independent influence on barrier-test crush within the passenger fleet (S27). On the measured cohort, b_m largely collapses when vehicle size (footprint) is controlled (b_m 0.18 $\rightarrow$ -0.04, $b_{\text{footprint}} \approx 0.41$), so the mass term is substantially a vehicle-size proxy rather than an independent mass effect. Structural body class (unibody vs body-on-frame) is unavailable from the NHTSA API; vehicle size is the tested proxy.

A.5 Bootstrap confidence intervals

B = 500 resamples give 90 % confidence bands on each quantile curve (shared-slope fit). At 10 m/s and 1,500 kg:

Quantile	CI (mm)	Half-width
q50	326–344	±3 %
q05	202–228	±6 %
q95	440–471	±3 %

All quantile CIs are properly ordered (q05 < q50 < q95) at every velocity.




Figure A1: Bootstrap confidence ribbons on the quantile curves

Figure A1. Bootstrap confidence ribbons on the quantile curves (shared-slope, measured-only fit). The bands are the 90 % confidence intervals from 500 resamples.

A.6 Fleet-weighting sensitivity

Repeated weighted bootstrap ($B = 200$, ABS age-distribution weights, zero-weight rows excluded) shows a consistent downward shift of 4–7 % across all quantiles:

Quantile	Shift	Resampling SD	Shift / SD
q05	–6.4 %	3.7 %	~1.7
q50	–6.7 %	3.9 %	~1.7
q95	–4.1 %	3.5 %	~1.2

Each shift is 1.2–1.7 $\times$ its own bootstrap resampling standard deviation, so the effect is small but real, not within noise. ABS age-weighting biases the unweighted cohort mildly toward softer (higher-crush) vehicles; the correction is a genuine ~5 % downward adjustment.

A.7 Onset and initial-condition sensitivity

Each measured dynamic crush depends on where integration starts (the detected impact onset) and on the pinned initial speed. A perturbation study over the 1,412 measured fit-ready tests, driving the same production integrator, applied onset ± 2 ms, detection-threshold $\times 0.5/\times 2$, and initial-speed ± 2 %.

Metric	Value
Per-test worst-case shift, median	5.1 %
Per-test worst-case shift, 95th percentile	49 %
Per-test worst-case shift, maximum	140 %
Dominant axis (mean shift)	Detection threshold: 7.4 %
Onset window (mean shift)	3.9 %
Initial speed (mean shift)	3.3 %
Worst-case median, < 30 km/h	8.9 %
Worst-case median, ≥ 30 km/h	5.1 %

The effect is speed-dependent: a fixed onset shift admits or omits a roughly fixed pre-onset travel ($\approx v_0 \Delta t$) that is a larger fraction of a small low-speed crush. For the operative highway-speed regime the bound is modest ($\sim 5\%$); at low speed it is not. Sampling uniformity holds (0 of 1,412 traces exceed a 1 % dt-nonuniformity tolerance), so the fixed-dt integration assumption is safe.

A.8 Force identity and error terms

Maximum dynamic crush is defined as the displacement at the velocity zero-crossing, exactly the integration endpoint used to derive it. By the work-energy theorem, the work done on the vehicle over the loading stroke equals $\frac{1}{2}mv^2$ regardless of rebound energy, so $F_{avg} = mv^2/(2 \cdot d_{dynamic})$ is an exact identity for the loading-phase average interface force on a rigid barrier. Restitution does not enter it. This is a definitional property of the loading-stroke work, not an approximation, but it is worth noting that the rebound energy is small in any case: an independent set of frontal rigid-barrier tests reports coefficients of restitution between about 0.05 and 0.22 (S28), so even the total rebound energy, which the loading-phase identity already excludes, is a minor fraction of the impact energy.

Restitution genuinely matters elsewhere: for the static-to-dynamic crush relationship, for Δv and occupant severity, and for whether a low-speed impact leaves measurable residual crush. The model's real breakdown below $\sim 16\text{--}30$ km/h is a crush-prediction failure (origin-constrained form, no low-speed data, a genuine no-damage threshold), not a force-formula failure (§6).

The error terms the identity does carry are:

- **Vehicle pitch/rotation:** the accelerometer measures one point, not the barrier-normal travel of the centre of gravity (CG).
- **Accelerometer location and mounting compliance.**
- **Barrier compliance:** near-zero for a FWRB (≈ 0.1 mm at peak force, estimated in this work from the reinforced-concrete barrier's stiffness).
- **Lateral/vertical energy paths** not captured by the longitudinal channel.

On a rigid barrier these are second-order at design speeds.

A.9 Force-validation detail

(i) Out-of-sample crush prediction (expressed in force units). A 5-fold hold-out cross-validation: the 614 measured-crush tests were split into folds, the median (measured-only) fit retrained on each complement, and force predicted on unseen rows. Result: median model/measured force ratio 0.997 (essentially unbiased; mean 1.034), 5th–95th percentile spread 0.78–1.37. Because both “model force” and “measured force” are $mv^2/(2d)$ with the same m and v , the ratio is exactly the inverse crush-prediction ratio. This validates crush prediction, not force; no independent force channel enters. It is out-of-sample within the accelerometer measurement channel.




Figure A2: Out-of-sample crush-validation scatter in force units

Figure A2. Out-of-sample crush-validation scatter expressed in force units (§A.9(i)): a crush-prediction check under 5-fold hold-out, not an independent force validation.

(ii) Cross-channel displacement agreement. The barrier load-cell processing batch produced 386 clean modern-fleet tests. For each, dynamic crush derived from the barrier load cells (double-integrating $F = ma$) is compared against crush derived independently from the onboard accelerometers. Median load-cell/accelerometer crush ratio 0.981 (IQR 0.931–1.05, 5th–95th 0.877–1.52). Two physically independent instrument chains agree in the median to ~2 % (71 % within ± 10 %, 87 % within ± 20 %). The load-cell energy-balance ratio and the model-vs-LC-average-force self-consistency ratio remain self-consistency checks, not validation (both ≈ 1 by construction because displacement is integrated from the same force signal). The load-cell set spans 32.7–57.0 km/h with no tests below 30 km/h.

(iii) Independent force validation against barrier load cells. The model’s predicted average force (measured-only q50 fit) is compared, at each load-cell test’s own mass and speed, against force measured by the barrier load cells, a channel independent of the accelerometer crush the model is built on (386 clean tests; 344 share a test number with the fit cohort, which does not compromise force independence because the force channel is distinct from the crush channel).

Two load-cell averages are reported (neither a sole primary):

Average	Definition	Median model/LC ratio	5th–95th
Distance-average	Energy / LC crush	0.921	0.78–1.40
Time-average	Impulse / duration	0.835	0.69–2.03

The distance-average is matched to $F = mv^2/2d$ but partly crush-dependent; the time-average is crush-free and genuinely independent, but is a time- rather than distance-average.

This comparison is made at the model q50 (central-tendency validation, not a design reading). The model median sits ~8–17 % below the modern load-cell averages (the ~17 % time-average figure is inflated by the pulse-shape offset; the definition-matched distance-average ~8 % is the better estimate of the force bias).

The design reading is taken off the stiff-vehicle q05 curve. Force demand $\propto 1/\text{crush}$ and the q05/q50 crush ratio is fixed, so q05 carries a flat +27.2 % over q50. Model q05 sits +6 to +17 % above the modern-fleet average force (1.272×0.835 – 0.921). Measured against the stiff tail itself (5th percentile of the per-test ratio: 0.78 distance / 0.69 time), it lands at parity to ~12 % below the stiffest modern vehicles ($1.272 \times 0.78 = 0.99$; $\times 0.69 = 0.88$). Erring to q05 closes almost the entire central bias on its own, with the AS/NZS

1170.0 ULS action factor and capacity-reduction factor supplying the residual tail margin (§8.4). This rests on the fleet-stiffness-spread assumption stated canonically at §8.5.




Figure A3: Independent force validation against barrier load cells

Figure A3. Independent force validation against barrier load cells (measured-only fit). The model median (q_{50}) is compared with the load-cell distance-average and time-average force at each test's own mass and speed, a force channel independent of the accelerometer crush the model is built on. The bottom-right panel shows the load-cell peak-to-average ratio (median 1.93).

A.10 Correction-residual evidence

The corrected-static crush path is demoted for three evidenced reasons, established by a correction-residual test on the 614 dual-value cohort rows:

- **(a) Level bias.** The correction is a least-squares (mean) fit to a right-skewed dynamic/static ratio, so on the operative cohort it over-predicts the median by 3.1 % (median residual -0.043 ratio units, 90 % bootstrap CI $[-0.063, -0.004]$, excludes zero).
- **(b) Large per-row scatter.** Cohort residual 5th–95th percentile $[-0.328, +0.780]$, approximately 1.1 ratio units.
- **(c) Unreconciled exponent divergence.** The measured and corrected velocity exponents (1.27 vs 1.85) that an adequate global correction would have reconciled but does not.

The corrected path is not demoted for cohort covariate non-homogeneity. On the operative cohort no single covariate (speed, mass, year, size) is materially structured (mass $p = 0.15$, footprint $p = 0.11$).

A.11 Peak-to-average force ratio

Peak forces on load-cell barriers are a median $1.93\times$ the distance-average force (IQR 1.77–2.16, 5th–95th 1.5–2.5, over the 386 clean load-cell tests; robust to the CFC-60 filter correction). This independently reproduces the $\sim 1.9\times$ headline from a channel not used to derive the model.

A.12 AHOF wall-row order verification

The hi-res walls letter their rows from the top down; the published description of the 4×9 walls places row A at the bottom (S9). The row-order assignment behind §4.2 was checked three ways. (i) Cross-wall agreement: same-model vehicles tested on both wall types agree under the documented orders (mean difference 10 mm across 39 models) and disagree by an average 136 mm under the reverse. (ii) Rebinning: collapsing the hi-res force records onto the coarse 246 mm 4×9 grid reproduces that stratum's median to within about 20 mm. (iii) Load-path location: under the documented order the most heavily loaded 4×9 row sits at 369 mm, at bumper and rail height, not at 615 mm, where no primary load path runs.

Annex B: References

- S1 Campbell, K.L. (1974), "Energy Basis for Collision Severity", SAE 740565.
- S2 Siddall, D.E. & Day, T.D. (1996), "Updating the Vehicle Class Categories", SAE 960897.
- S3 Ait Syad et al. (2021), "Computational Method of the Stiffness Coefficients A and B...", Int. J. ITS Research 19:587–593.
- S4 Hu, Li, Chen, Shi (2011), "State-of-the-Art Review on Anti-Ram Bollards", Applied Mechanics and Materials 90–93:3206–3213.
- AS/NZS 1170.1:2002, *Structural design actions — Part 1*, Clause 3.8.
- AS/NZS 1170.1 Supp 1:2002, *Structural design actions — Commentary* (Supplement to AS/NZS 1170.1:2002), Clause C3.8: energy-balance formula (eq C3.8) and scenario basis (mass, speed, assumed crumple zone) for the prescribed barrier forces.
- Wen, Y.K. & Yeo, G.L. (2001), "Design Live Loads for Passenger Cars Parking Garages", *J. Structural Engineering*, ASCE, 127(3):280–289; referenced by the Commentary C3.8; covers floor live loads (not barrier forces); vehicle weight survey (mean 1,551 kg, lognormal).
- ISO 22343-1:2023, *Vehicle security barriers — Part 1: Performance requirement, vehicle impact test method and performance rating* (current HVM certification standard; adopted in Australia as AS ISO 22343.1:2024). Supersedes the now-superseded PAS 68:2013 (whose historical ratings the paper's scenario labels reuse) and IWA 14-1:2013.
- ASTM F2656/F2656M-23, *Standard Test Method for Crash Testing of Vehicle Security Barriers*.
- ANZCTC, *Hostile Vehicle Guidelines for Crowded Places* (Australian HVM practice guidance).
- S5 NHTSA (2014), "Test Reference Guide Version 5, Volume I: Vehicle Tests".
- S6 Vangi, Cialdai, Gulino (2019), "Vehicle stiffness assessment for energy loss evaluation in vehicle impacts", *Forensic Science International* 300:136–144.
- S7 BITRE (2022/2023/2025), *Road Vehicles, Australia* (statistical report series); successor fleet-age anchor to S8 (§2.3).
- S8 Australian Bureau of Statistics (2021), "Motor Vehicle Census, Australia", cat. 9309.0. (*Discontinued after 2021; succeeded by S7.*)
- S9 Digges, K. & Eigen, A. (2000), "Analysis of Load Cell Barrier Data to Assess Vehicle Compatibility", SAE 2000-01-0051, SAE Transactions Vol. 109 Sec. 6.
- S10 Brewer, J., Patel, S., Summers, S., Prasad, A., Mohan, P. (2011), "Accuracy of AHOF400 with a Moment-Measuring Load Cell Barrier", 22nd ESV Paper 11-0335.
- S11 Jerinsky, M. & Hollowell, W.T. (2003), "NHTSA's Review of a Vehicle Compatibility Performance Metric Through Computer Simulation", ASME IMECE2003-44045.
- S12 Jerinsky, M.B. & Hollowell, W.T. (2003), "NHTSA's Review of High-Resolution Load Cell Walls' Role in Designing for Compatibility", 18th ESV Paper 393.
- S13 BITRE (2021), "Australia's light vehicle fleet — some insights", Information Sheet 108 (source data: ABS Motor Vehicle Census 2018).


- **S14** Fuerbeth, U. (2024), "Centre of gravity of motor vehicles", *Forensic Science International* 361:112073.
 - **S15** Baráth, B. & Józsa, D. (2026), "Determination of the Centre of Gravity of Electric Vehicles Using a Static Axle-Load Method", *Future Transportation* 6(1):22.
 - **S16** FIMCAR consortium (2013), "Frontal Impact and Compatibility Assessment Research (FIMCAR) — Final Report Summary", EU FP7 CORDIS Project 234216.
 - **S17** NHTSA (2007), "NHTSA's Recent Vehicle Crash Test Program on Compatibility in Front-to-Front Impacts", 20th ESV Paper 07-0231-O.
 - **S18** Smith, R. (1998), "The formula commonly used to calculate velocity change in vehicle collisions", *Proc IMechE Part D* 212:73–78.
 - **S19** Harrison, T.L. (2004), "Design of Fixed Ram-Resistant Vehicle Barriers for Perimeter Security", ASCE Structures 2004.
 - **S20** Iqbal, M. (2008), "Designing Edge Barriers in Parking Structures", *STRUCTURE magazine*, October 2008, pp. 24–26.
 - **S21** Iqbal, M. (2010), "A Rational Method to Design Vehicular Barriers", *STRUCTURE magazine*, September 2010, pp. 22–24.
 - **S22** Williams, R. (2024), "Barrier Loads for Parking Garages", CMAA Member Communication, New York City Department of Buildings.
 - **S23** Xiang, S., He, Y., Zhou, X. (2022), "Performance assessment of steel parking structure columns subjected to frontal collision based on reduced vehicular models", *Engineering Structures* 265:114517.
 - **S24** Xiang, S., He, Y., Zhou, X. (2022), "Simplified analytical procedure to calculate the impact behaviour of steel parking-structure columns", *Structures* 46:1938–1954.
 - **S25** He, Y., Xiang, S., Zhou, X. (2023), "A reinforcing protection for steel parking-structure column under vehicular impact and design method", *Thin-Walled Structures* 183:110368.
 - **S26** Li, H., Zhou, M., Feng, L., Wen, X., Fan, G. (2025), "Performance evaluation method of precast parking structure vertical members under vehicle impact", *Structural Concrete* 26(4):4030–4050.
 - **S27** Poliak et al. (2026), "Non-Linear Method of Vehicle Velocity Determination...", *Applied Sciences* 16(1):401.
 - **S28** Żuchowski, A. (2015), "The use of energy methods at the calculation of vehicle impact velocity", *The Archives of Automotive Engineering / Archiwum Motoryzacji* 68(2):85–111.
 - **S29** Jiang, T., Grzebieta, R.H., Rechnitzer, G., Richardson, S., Zhao, X.L. (2003), "Review of car frontal stiffness equations for estimating vehicle impact velocities", 18th ESV Paper 439.
 - **S30** Wang, Q. & Gabler, H.C. (2007), "Accuracy of vehicle frontal stiffness estimates for crash reconstruction", 20th ESV Paper 07-0513.
 - **S31** Strother, C.E., Woolley, R.L., James, M.B., Warner, C.Y. (1986), "Crush energy in accident reconstruction", SAE 860371, SAE Transactions Vol. 95:740–756.
 - BS 6399-1:1996, *Loading for buildings — Part 1: Code of practice for dead and imposed loads*, British Standards Institution.
 - EN 1991-1-7:2006, *Eurocode 1: Actions on structures — Part 1-7: General actions — Accidental actions* (incl. Annex C equivalent-static vehicle-impact model), European Committee for Standardization.
 - AS/NZS 1170.0:2002, *Structural design actions — Part 0: General principles* (reliability/characteristic-action framework; §8.4).
 - AS/NZS 3845.1:2015 & 3845.2:2017, *Road safety barrier systems and devices* (adjacent barrier-containment framework; §8.4, scoped out).
 - AASHTO MASH (2016), *Manual for Assessing Safety Hardware* (crash-test/containment levels; §8.4, scoped out).
-

Annex C: Design artefact sheets

The ten sheets below collect the paper's read-off artefacts in printable A4 form, one sheet per page: the all-masses crush envelope (A), the per-mass envelopes at six design masses (B1–B6), the stiff- and soft-vehicle crush contours (C1, C2), and the impact-height line (D). Each sheet restates its own read-off formula, fitted parameters, and validity window, so it stands alone once separated from the paper; the research-artefacts section (§9) describes the set.

Design artefact A - dynamic-crush envelope, all masses

Measured-crush cohort with shared-slope quantile power-law fits; curves at the cohort median mass



READ-OFF FORMULA AND FITTED PARAMETERS (shared slopes, per-quantile level):

crush_mm = a_q x v^1.272 x m^0.182 (v in m/s, m in kg)

quantile	a_q
q05 5th (stiff – conservative for force)	4.0886
q25 25th	4.7226
q50 50th (median)	5.2019
q75 75th	6.0042
q95 95th (soft – conservative for penetration)	7.1853


Read the crush for a design vehicle mass and speed directly from the formula, or from the per-mass sheets (artefacts B1-B6). q05 is the stiff-vehicle (force-governing) read-off; q95 is the soft-vehicle (penetration-bounding) read-off.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004). Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact A: dynamic-crush envelope, all masses

Design artefact B - dynamic crush at 1,100 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 1,100 kg: NOTE: 1100 kg sits near the cohort mass floor (1007 kg).

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	14.6	312	394	415	479
q25	16.9	361	455	479	553
q50	18.6	397	501	528	609
q75	21.4	458	578	609	703
q95	25.6	549	692	729	842

crush_{mm} at 1100 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004). Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B1: crush envelope at 1,100 kg

Design artefact B - dynamic crush at 1,500 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 1,500 kg:

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	15.4	330	416	439	507
q25	17.8	381	481	507	585
q50	19.6	420	530	558	645
q75	22.7	485	612	644	744
q95	27.1	580	732	771	890

crush_{mm} at 1500 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004).
Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B2: crush envelope at 1,500 kg

Design artefact B - dynamic crush at 2,000 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 2,000 kg:

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	16.3	348	439	462	534
q25	18.8	402	507	534	617
q50	20.7	443	558	588	679
q75	23.9	511	644	679	784
q95	28.6	612	771	812	938

crush_{mm} at 2000 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004). Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B3: crush envelope at 2,000 kg

Design artefact B - dynamic crush at 2,500 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 2,500 kg:

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	16.9	362	457	481	556
q25	19.6	419	528	556	642
q50	21.6	461	581	612	707
q75	24.9	532	671	707	816
q95	29.8	637	803	846	977

crush_{mm} at 2500 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004). Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B4: crush envelope at 2,500 kg

Design artefact B - dynamic crush at 3,000 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 3,000 kg:

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	17.5	375	472	498	575
q25	20.2	433	546	575	664
q50	22.3	477	601	633	731
q75	25.7	550	694	731	844
q95	30.8	658	830	874	1010

crush_{mm} at 3000 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004).
Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B5: crush envelope at 3,000 kg

Design artefact B - dynamic crush at 3,400 kg

Quantile crush curves at one design mass; nearby cohort tests overlaid



READ-OFF VALUES AT 3,400 kg: NOTE: 3400 kg sits near the cohort mass ceiling (3448 kg).

quantile	$A_q = a_q \times m^{bm}$	40 km/h	48 km/h	50 km/h	56 km/h
q05	17.9	383	483	509	588
q25	20.7	443	558	588	679
q50	22.8	488	615	648	748
q75	26.3	563	710	747	863
q95	31.5	673	849	894	1033

crush_{mm} at 3400 kg = $A_q \times v^{1.272}$ (v in m/s); table values in mm


q05 = stiff-vehicle crush (use for FORCE via $F = m \times v^2 / (2 \times d)$); q95 = soft-vehicle crush (bounds penetration depth). 40/48/50/56 km/h are common assessment speeds inside the supported band.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year >= 2004).
Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact B6: crush envelope at 3,400 kg

Design artefact C - q05 crush (stiff vehicle) - governs force

Labelled crush contours (50 mm increments) over the supported mass x speed envelope



HOW TO READ:


1. Enter at the design speed (bottom axis) and vehicle mass (left axis).
2. Interpolate the crush d between the labelled 50 mm contours.
3. Average impact force follows from $F = m \times v^2 / (2 \times d)$ - the q05 (small-crush) read-off maximises that force, so it is the conservative choice for barrier strength.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, n=614 tests (model year $\geq$ 2004).
Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact C1: stiff-vehicle (q05) crush contours, force side

Design artefact C - q95 crush (soft vehicle) - bounds penetration

Labelled crush contours (50 mm increments) over the supported mass x speed envelope



HOW TO READ:


1. Enter at the design speed (bottom axis) and vehicle mass (left axis).
2. Interpolate the crush d between the labelled 50 mm contours.
3. The q95 (large-crush) read-off bounds how deep the vehicle front deforms; use it when clearance or standoff distance governs.

VALIDITY WINDOW - measured-crush support: 9.1-15.8 m/s (33-57 km/h), 1007-3448 kg, $n=614$ tests (model year ≥ 2004).
Dashed curves / values beyond this window are extrapolation. Measured-only primary fit on public full-width rigid-barrier crash tests; research envelope, not a certified design basis.

Artefact C2: soft-vehicle (q95) crush contours, penetration side

Design artefact D - height of the resultant force vs vehicle mass

Ground-referenced AHOF from high-resolution load-cell walls, with mass-conditional quantile lines



AHOF READ-OFF TABLE (mm above ground):

mass kg	p50	p75	p90	p95	n +/-150kg	flag
1100	370	371	382	394	2	edge-of-support
1500	396	407	423	435	30	
2000	427	453	474	487	43	
2500	459	498	525	539	9	
3000	490	544	576	591	6	
3400	515	580	617	632	3	edge-of-support

p95 is the design-governing height: a barrier engaging up to this height spans the resultant of 95% of same-mass vehicles. Single-axis load cells bias AHOF low, so these heights are conservative as lower bounds on engagement height.

VALIDITY WINDOW - hi-res ground-referenced AHOF cohort: 1115-3367 kg, n=154 tests, MY 2005-2016. 1,100 and 3,400 kg table rows sit at the edge of support (see n column). Heights assume wall base = ground; single-axis cells read low. Research envelope, not a certified design basis.

Artefact D: impact-force height versus mass

Data availability

The primary data are public records from the U.S. DOT/NHTSA Vehicle Crash Test Database. The derived dataset (cleaned crash-test rows, the dynamic-crush derivations, the cohort selection, and the fitted design tables), together with the quality-control pipeline, the analysis scripts that produce the quantile envelope, design charts, and figures, and the machine-generated QC and validation reports, are deposited at Zenodo: [doi:10.5281/zenodo.21508011](https://doi.org/10.5281/zenodo.21508011). The raw cached NHTSA records are not redistributed; they are re-buildable from the public database with the deposited fetch scripts.