CanLaunch
Canadian Space Launch · Compliance & Risk

CANLAUNCH TECHNICAL SERIES · VOLUME 1

Quantitative Risk Assessment for Canadian Space Launches

A textbook treatment with a complete worked example: *Cyclone-4M lifting off from Spaceport Nova Scotia*.

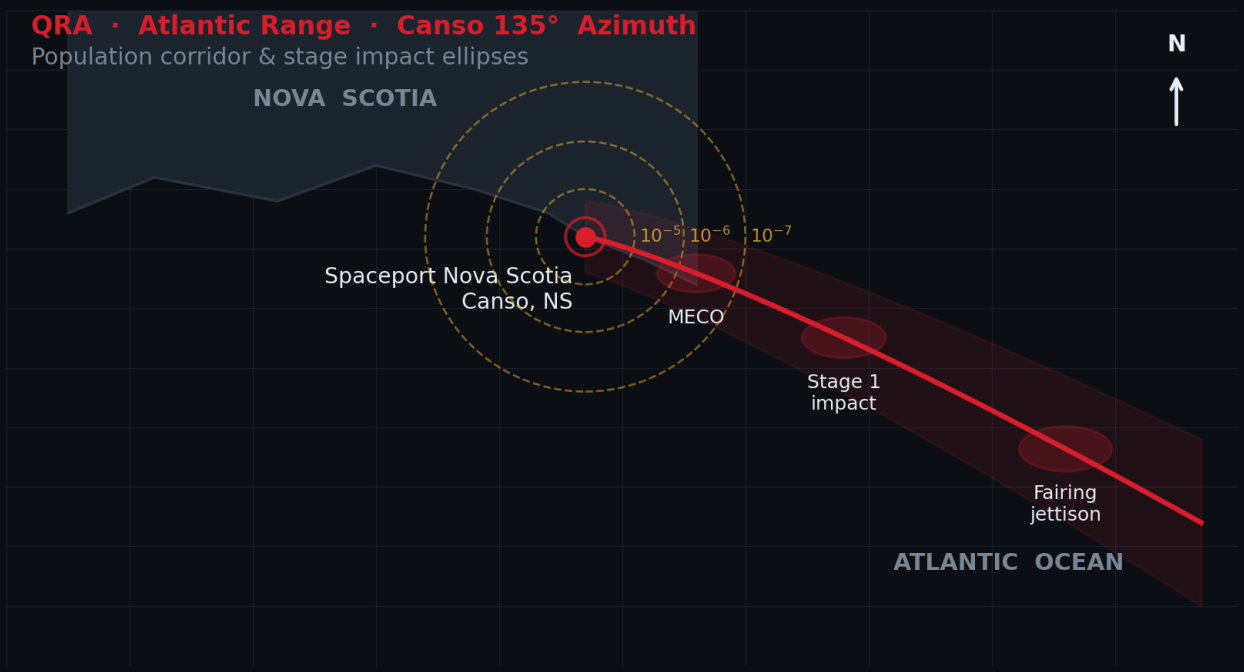

Aligned to

Transport Canada LSO Framework 2024

Canadian Space Agency Act (S.C. 2023, c. 9)

CARs Part VI · Space Activities

Reference frameworks

FAA 14 CFR Part 450 (comparison)

RCC Standard 321-07

DNV-RP-G101 · IEC 61508 (concepts)

For

Range safety officers

Launch operators & integrators

Regulators & assessors

Contents

Foreword

Canada is, for the first time, building a sovereign orbital launch capability. Spaceport Nova Scotia near Canso is under construction; the Atlantic Spaceport Complex in Newfoundland is flying suborbital test articles; and Spaceport Québec is moving through proposal stages. Until 2023, Canadian operators had no domestic regulatory framework distinct from foreign jurisdictions; the *Canadian Space Agency Act* (S.C. 2023, c. 9) and the Transport Canada Launch Safety Office framework changed that.

This volume teaches Quantitative Risk Assessment — the formal, numerical discipline that converts "is this launch safe enough?" into specific numbers you can defend before a regulator. It is intended as a textbook: each chapter introduces theory, then applies it to a single worked example — a Cyclone-4M lifting off from Canso on a sun-synchronous trajectory — that we carry from the first hazard list all the way through to a final, acceptable-or-not verdict.

The numbers in the worked example are illustrative, chosen to be plausible and to teach the method. They are not the real numbers behind any actual Maritime Launch Services or Reaction Dynamics filing. A genuine licence submission requires manufacturer-supplied reliability data, the actual fragment catalogue, validated wind soundings, and population datasets that we cannot reproduce here without proprietary access. What this book does give you is the framework — every equation, every input class, every decision point — so that when you have those proprietary inputs, you know exactly what to do with them.

How to read this volume

Chapters 1–2 are conceptual scaffolding. Chapters 3–7 are the technical core: each one covers a single building block of QRA in depth, then bolts it onto the running example. Chapter 8 is the assembly — every piece comes together into a single end-to-end calculation. Chapters 9–10 cover what to do with the result: how to mitigate when risk is too high, and how to package the calculation for Transport Canada review.

Two visual conventions appear throughout. **Red-bordered boxes** hold definitions, equations, and concept summaries. **Black boxes** are the running worked example — every black box continues the Cyclone-4M calculation from where the last one left off. Read those in order to follow the calculation as a single thread.

A note on units

QRA literature mixes SI and US Customary units freely. FAA Part 450 uses feet and pounds; ICAO and Canadian sources use metres and kilograms. We use SI throughout, with US Customary equivalents in parentheses where the source material requires it. Probabilities are dimensionless; expected casualties are dimensionless (count of people); ballistic coefficients are kg/m^2.

Chapter 1

Why QRA, and why Canada needs its own

Every launch — every one — is a controlled detonation strapped to a guidance system. A Cyclone-4M carries roughly 290 tonnes of liquid oxygen and refined kerosene at lift-off; a Falcon 9, more than 500 tonnes; even a small suborbital sounding rocket carries enough propellant to flatten a neighbourhood. The question "is this launch safe enough?" is not rhetorical, and it is not answerable by inspection. The risk is small but not zero, and decisions about it are taken under uncertainty.

Quantitative Risk Assessment is the formal discipline of putting a number on that small-but-not-zero risk, comparing it to a published threshold, and documenting the comparison in a way a regulator can audit. Done well, it lets a launch operator say to Transport Canada: "the chance that this launch will hurt any specific member of the public is at most one in five million, and the expected number of casualties across all the public is at most three in a hundred thousand — here is exactly how we got those numbers."

1.1 The Canadian regulatory stack

Until very recently, Canadian space activity sat in a curious legal vacuum. Canada is a party to the Outer Space Treaty (1967) and the Liability Convention (1972), but there was no domestic statute that authorised a private actor to launch an orbital vehicle from Canadian soil and specified what they had to do to obtain that authorisation. The 2023 *Canadian Space Agency Act* closed that gap. It empowered Transport Canada to issue Launch Safety Office (LSO) licences and tied them to amendments in the *Canadian Aviation Regulations* (CARs Part VI, subparts addressing space activities) and to the *Aeronautics Act*.

Layer	Instrument	What it governs
Treaty	Outer Space Treaty (1967), Liability Convention (1972)	International obligations of Canada as a launching state
Statute	Canadian Space Agency Act (S.C. 2023, c. 9); Aeronautics Act	Authority of Transport Canada to license launches
Regulation	CARs Part VI — Space Activities (as amended)	Pre-flight requirements, range safety, environmental
Framework	TC LSO Framework 2024	Specific licence conditions, risk thresholds, evidence
Standard	Nav Canada NOTAM, ICAO Annex 2, RCC 321-07 (referenced)	Airspace coordination, range safety practice
Operator	.canlaunch.json record · Flight Safety Analysis report	The evidence package the operator submits

Table 1.1 — The Canadian launch regulatory stack. QRA results sit in the operator layer at the bottom, but every layer above constrains them.

Why does this matter for QRA? Because the acceptance thresholds — the specific numerical values your QRA result is compared against — live in the framework layer, not the statute. They can be revised without Parliament having to act, and they are *not* identical to the FAA's. A Canadian operator who relies on FAA-aligned compliance software is, at best, approximating the actual rule.

1.2 How QRA differs from qualitative hazard analysis

Most hazard work in industry is qualitative: a HAZOP study, a risk matrix that puts "likelihood" on one axis and "severity" on the other, and an output of "high / medium / low". That is fine for prioritising engineering effort, but it cannot answer the regulatory question, which is specifically numerical. QRA differs in three concrete ways:

It produces numbers, not labels. The output is an expected casualty count (Ec), an individual probability of casualty (Pc), and a frequency–consequence (F–N) curve — all dimensionless, all comparable to a written threshold.

It propagates uncertainty. Every input — failure rate, fragment mass, wind speed, population density — is a distribution, not a single value. QRA carries those distributions through to the answer, so the final number comes with confidence bounds.

It is auditable. Every input has a source. Every model has a reference. A regulator opening the file two years later can reproduce the calculation, change one assumption, and see how the answer moves.

Definition · QRA

Quantitative Risk Assessment is a formal, systematic process for estimating, in numerical terms, the likelihood and severity of harmful events, propagating uncertainty in inputs through to uncertainty in the output, and comparing the result to a published acceptance criterion.

1.3 Where QRA sits in the launch licensing workflow

A Transport Canada LSO licence application has many parts — environmental review, financial responsibility, organisational fitness, range safety, communications, and a Flight Safety Analysis. The Flight Safety Analysis is where QRA lives. It is the section of the application where the operator demonstrates, in numbers, that the public risk from the proposed operation is below the LSO threshold.

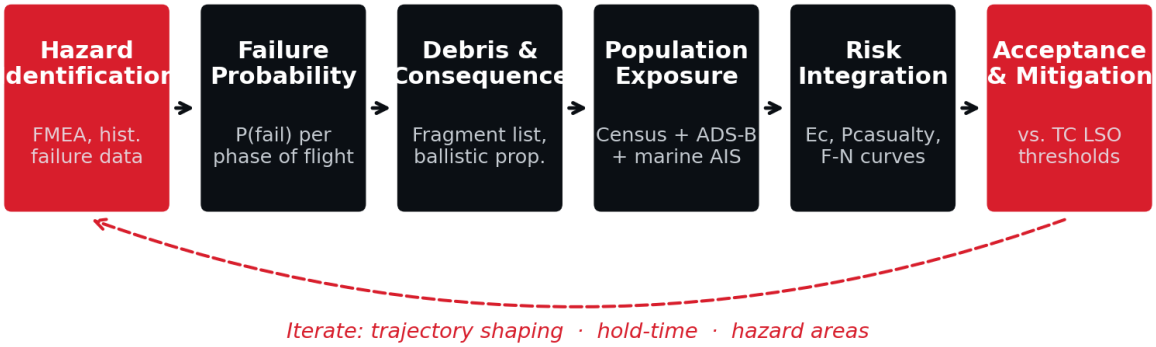

Figure 1.1 — The six-step QRA process. Steps 1–4 build the inputs; step 5 integrates them into risk metrics; step 6 compares against thresholds and triggers iteration if the answer is unacceptable. Chapters 3–8 of this book follow the same six steps.

Three things follow from the diagram. First, QRA is not a one-pass calculation: if the answer comes back over threshold, you go back and change something — the trajectory, the hold-times, the hazard areas — and re-run. Second, the QRA is the *engine* of mitigation decisions, not just an after-the-fact judgement. Third, every step in the loop produces evidence, and that evidence is what TC will eventually audit.

Chapter 2

The QRA framework

Before we touch a single number, we need shared vocabulary. The literature uses three risk metrics — collective Ec, individual Pc, and societal F–N — and three acceptance criteria for each. They are easy to confuse and they answer different questions. This chapter pins them down.

2.1 Risk metrics: Ec, Pc, F–N curves, individual vs societal

Collective risk: Expected Casualties (Ec)

Ec is the expected number of casualties — across the entire public, across the entire mission — from a single launch attempt. It is the weighted sum of "how often" times "how bad" over every failure mode and every population element. "Casualty" in this context means any injury severe enough to require professional medical attention; it does *not* mean fatality (fatalities are a subset).

Equation 2.1 · Mission Expected Casualties

$$E_c = \Sigma_i \; P_i \cdot A_{c,i} \cdot \rho_{p,i}$$

where i indexes failure-mode/phase combinations; P_i is the probability of impact in cell i; $A_{c,i}$ is the casualty area; $\rho_{p,i}$ is the population density there.

Ec is the most-used number in launch range safety. It is the metric the FAA's Part 450 §450.101 sets at $\mathbf{1\times10^{-4}}$ casualties per launch for the collective public, and the metric Transport Canada has aligned its LSO threshold to at the same level. It is intuitive: an Ec of 1×10^{-4} means "if you flew this exact mission ten thousand times, you would expect one casualty in total."

Individual risk: Probability of Casualty (Pc)

Pc answers a different question: "for any one specific person — say, an unlucky resident of Canso — what is the probability they personally are a casualty of this launch?" Ec is a sum across people; Pc is a per-person number. A launch can have low Ec because almost nobody lives downrange, while still having unacceptable Pc for the few who do.

Equation 2.2 · Individual Probability of Casualty

$$P_c(x, y) = \Sigma_i \; P_i \cdot f_i(x, y) \cdot A_{c,i}$$

where $f_i(x, y)$ is the probability density of fragment impact at location (x, y), in $1/m^2$, for failure mode i.

FAA Part 450 §450.101 sets the individual Pc limit at $\mathbf{1\times10^{-6}}$ per launch for any member of the public, and 1×10^{-5} for neighbouring operations personnel (people working nearby but not on the launch team). Visualised across a map, Pc forms *risk contours* (Figure 2.1).

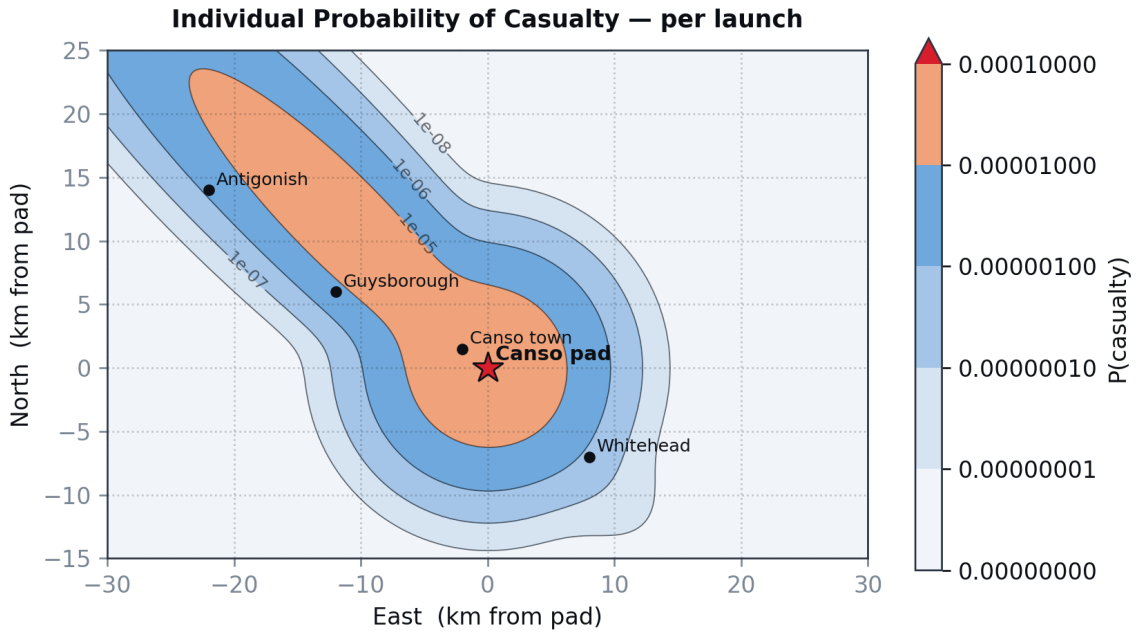

Figure 2.1 — Individual Pc as a function of position around the Canso pad, for the worked-example mission. Contour lines mark powers of ten. The 10^{-6} contour is the regulatory boundary: no member of the public should sit inside it without justification. (Illustrative; population labels are placeholders for the worked example.)

Societal risk: F–N curves

Ec and Pc miss something important. A mission with Ec = 10^{-4} might represent ten thousand events with one casualty each, or a single event with ten thousand casualties — the same expectation, very different social acceptability. Society is risk-averse to large-consequence events. The F–N curve captures this by plotting, on log axes, the frequency F of events causing N or more casualties, against N.

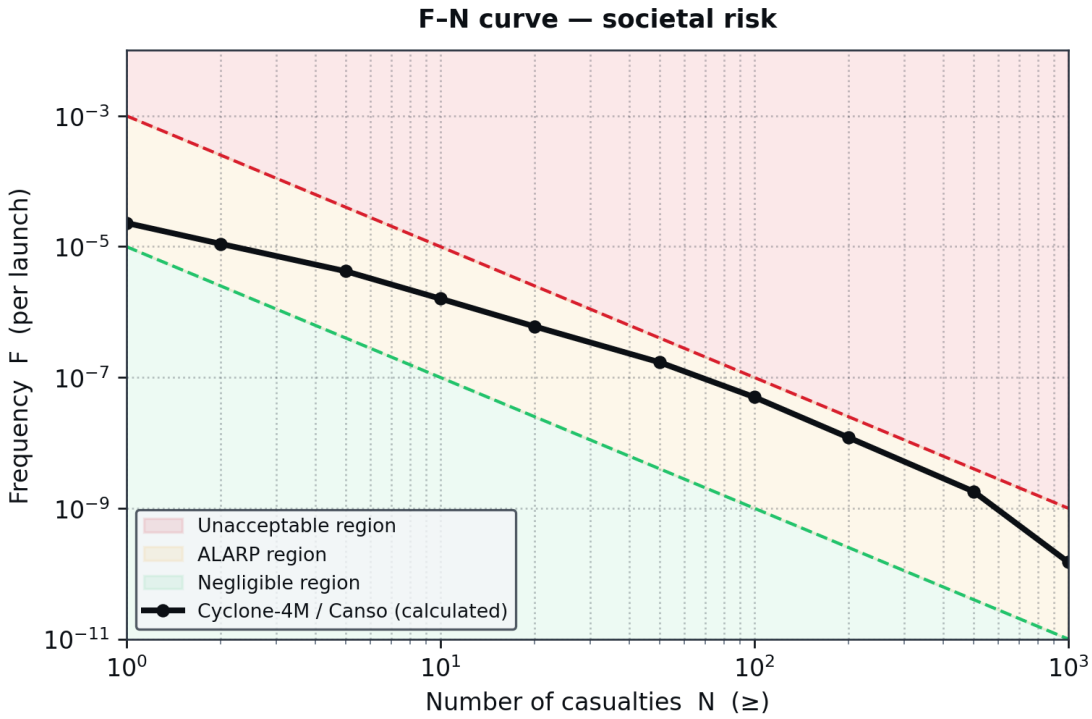

Figure 2.2 — F–N curve for the worked-example Cyclone-4M mission. The shaded bands show the conventional ALARP framework: the upper (red dashed) line is the unacceptable region; the lower (green dashed) is the broadly-acceptable region; between them is the ALARP region where additional mitigation must be justified.

2.2 The six-step QRA process

Throughout this book we use a standard six-step process, drawn from NASA-STD-8719.13 practice, RCC Standard 321-07, and the DNV process-safety QRA literature, and adapted for the launch context.

Step	Activity	Output
1	Hazard identification — enumerate failure modes and what they produce (debris, blast, toxic plume).	Hazard register
2	Failure probability — assign P(fail) per phase of flight, with uncertainty bounds.	Reliability table
3	Consequence modelling — fragment catalogues, ballistic trajectories, casualty area Ac.	Debris models, Ac per fragment
4	Population exposure — static (residents) plus dynamic (aircraft, vessels) populations, with sheltering.	Exposure grid ρp(x,y)
5	Risk integration — compute Ec, Pc(x,y), and the F–N curve from the products of steps 1–4.	Ec, Pc, F–N
6	Acceptance and mitigation — compare to thresholds; if over, modify the mission and iterate.	Go / mitigate / no-go

Table 2.1 — The six-step QRA process. Each step is the subject of one of Chapters 3–7 in this volume; Chapter 8 is the assembly.

2.3 Acceptance criteria — TC LSO and FAA Part 450 compared

It is worth lining up Canadian and US thresholds side by side, because they are similar but not identical, and operators who are used to the FAA framework sometimes assume equivalence where there is none.

Metric	FAA Part 450 §450.101	TC LSO Framework 2024 (illustrative)	Comment
Collective Ec, public	$\leq 1\times10^{-4}$ per launch	$\leq 1\times10^{-4}$ per launch	Aligned
Collective Ec, NOPS personnel	$\leq 2\times10^{-4}$ per launch	$\leq 2\times10^{-4}$ per launch	Aligned
Individual Pc, public	$\leq 1\times10^{-6}$ per launch	$\leq 1\times10^{-6}$ per launch	Aligned
Individual Pc, NOPS personnel	$\leq 1\times10^{-5}$ per launch	$\leq 1\times10^{-5}$ per launch	Aligned
Aircraft Pi (debris ≥ casualty)	$\leq 1\times10^{-6}$ per launch	$\leq 1\times10^{-6}$ per launch	Coordinated via Nav Canada
Conditional Ec (CEC), high-consequence	$\leq 1\times10^{-3}$ per phase	Case-by-case	TC retains discretion
Inert debris kinetic energy threshold	11 ft·lbf (≈ 15 J)	15 J (SI direct)	Equivalent
Blast overpressure casualty threshold	1.0 psi (≈ 6.9 kPa)	6.9 kPa	Equivalent

Table 2.2 — Numerical acceptance criteria. The Canadian framework as of 2024 mirrors FAA Part 450 on collective and individual numerical limits but reserves more discretion on conditional cases. The implication: anywhere a Canadian filing has slack relative to FAA, an operator should not assume that slack is policy.

Worked Example · setting up the problem

Throughout this book, we evaluate a single proposed mission:

Vehicle: Cyclone-4M, two-stage to LEO, GLOW ≈ 287 t, kerolox first stage, kerolox upper.
Site: Spaceport Nova Scotia, Canso, NS (≈ 45.30°N, 60.95°W).
Trajectory: Sun-synchronous, azimuth 135°, target 600 km circular at ≈ 97° inclination. Atlantic flyout.
Payload: 3,350 kg into target SSO (vehicle nameplate).
Launch window: Daytime, summer, prevailing south-westerly wind 10–15 kn.

Acceptance criteria (TC LSO illustrative):
Ec $\leq 1\times10^{-4}$ per launch (collective public)
Pc $\leq 1\times10^{-6}$ per launch (any individual public)
Aircraft Pi $\leq 1\times10^{-6}$ (any aircraft in active corridor)

Our task across the next six chapters: produce defensible numbers for each of these, and compare them to the thresholds.

Chapter 3

Hazard identification and the running example

Step 1 of QRA is to enumerate, exhaustively, the bad things that could happen and what each one produces. This is engineering work, not probability work — at this stage we are not assigning likelihoods, we are building a complete list. Skip a hazard here and the entire downstream calculation is invalidated.

3.1 Vehicle: Cyclone-4M from Canso

The Cyclone-4M is a two-stage liquid-propellant orbital launch vehicle developed by Yuzhnoye SDO of Ukraine. Its first stage is a Zenit-derived kerolox core powered by four RD-870 engines (a variant of the RD-120) with gimballed nozzles for thrust vectoring. The upper stage carries the RD-861K, also kerolox, for orbital insertion. From Canso, with a south-easterly trajectory, the vehicle is rated to deliver about **3,350 kg to sun-synchronous orbit** at 600 km altitude. Total mass at lift-off (GLOW) is approximately 287 tonnes, of which about 245 tonnes is propellant.

For QRA the relevant features of the vehicle are not the orbital ones, but the destructive ones: how much propellant, of what kind, in what tank arrangement; what the dry-mass distribution looks like; what kind of flight termination system is fitted; and what the staging events are, because each event is a hazard transition.

3.2 Mission profile and phases of flight

Phase	T+ (s)	Altitude	Downrange	Dominant hazards
Pre-ignition (holdown)	–10 to 0	0	0	Pad fire, vehicle anomaly on the pad, NOPS personnel only
Lift-off and tower clear	0–10	0–0.3 km	0–0.5 km	Pad-area fragments, blast, SRB exhaust (N/A here — kerolox)
Stage-1 ascent (max-Q)	10–135	0.3–55 km	0.5–95 km	FTS-driven breakup, propellant fragmentation, drift to populated areas
Stage separation	135–145	≈ 55 km	≈ 95 km	Failure to separate; residual propellant ignition; jettisoned hardware
Stage-2 ascent	145–540	55–180 km	95–1,500 km	Mostly suborbital over open ocean; reentering debris if failure
Fairing jettison	≈ 250	≈ 110 km	≈ 600 km	Inert fairings — must clear shipping lanes
Orbital insertion	540	600 km	Orbital	Negligible ground risk; on-orbit collision risk handled separately

Table 3.1 — Cyclone-4M mission phases for the worked example. T+ values are nominal and the timeline shifts modestly under wind and performance dispersions.

Phases matter because failure probability, hazard type, and population exposure all change rapidly along the trajectory. A failure at T+5 s puts debris on the pad and possibly into Canso town; the same failure at T+250 s drops fragments into the broad Atlantic. We will assign different P(fail), different fragment models, and different exposure calculations to each phase.

3.3 Hazard inventory

From the phases we derive a structured hazard register. Every entry links a *cause* (an initiating failure) to an *effect* (a release of harmful energy).

Initiating event	Phase	Energy released	Population at risk
Engine combustion instability / catastrophic failure	Lift-off → stage-1	Propellant fireball, fragments, blast	On-pad NOPS, Canso town (3 km N)
Loss of thrust vector control / guidance failure	Stage-1 ascent	Off-trajectory flight ending in FTS	Atlantic shipping, Sable Island sector
FTS premature trigger	Lift-off	Vehicle destruction over land	Canso town, Whitehead, Hazel Hill
FTS fails to trigger	Stage-1	Uncontrolled overflight, eventual ground impact	Wide footprint, downrange to St. Pierre
Stage separation failure	Stage sep	Stuck stack, propellant venting	Open Atlantic, low risk
Upper-stage failure	Stage-2 ascent	Suborbital reentry, debris field	Broad ocean area
Toxic plume release (kerolox: minimal)	Lift-off → ascent	Combustion CO/CO_2; HC unburned	Air-quality exposure, low concern
Far-field blast overpressure	Lift-off	Acoustic / shock to structures	Buildings within ≈ 4 km

Table 3.2 — Hazard register for the worked example. This list is complete in the sense that every reasonably-foreseeable initiating failure leads to one or more rows; nothing without a row contributes to the eventual Ec. Auditors check this list closely.

Why hazard identification cannot be skipped

A common QRA failure mode is a beautifully-quantified analysis built on an incomplete hazard register. If FTS-fails-to-trigger is missing from the register, the calculated Ec will be artificially low, the operator will look in good shape, and the eventual incident — when it happens — will reveal that the model never knew the failure mode existed. RCC 321-07 §2.5 and FAA AC 450.115-1 both treat hazard completeness as a precondition for everything else.

Worked Example · hazard register summary

From Tables 3.1 and 3.2, the worked-example mission has **seven** categories of initiating events spread across **seven** phases of flight. We will treat all 49 cells as candidates, but in practice many combinations have negligible contribution (e.g. "toxic plume during orbital insertion") and are screened out at Step 2.

The four cells that will dominate the eventual Ec, by inspection, are: **(a)** stage-1 failure with FTS response, debris over Atlantic; **(b)** stage-1 failure, FTS fails, ground impact on land; **(c)** guidance loss with FTS response, debris drift to land; **(d)** stage-2 reentering debris over shipping lanes.

Chapter 4

Failure probability — building P(fail)

Step 2 of QRA assigns a probability to each hazard cell. This is the most contested input in launch QRA: small launchers fly rarely, so frequentist estimates have wide bounds; new vehicles have no flight heritage at all. The discipline here is honest treatment of that uncertainty rather than borrowing comforting numbers from a different vehicle.

4.1 Historical reliability data

As of 2026, the orbital launch industry has flown roughly 7,500 orbital missions over 65 years, with a long-run success rate near 92%. But lumping all vehicles together is misleading. Maiden flights of new vehicles fail roughly 50% of the time; vehicles with ten or more successful flights typically run above 95%. The Cyclone family — Tsyklon-2, Tsyklon-3, Tsyklon-4 ancestry — has a long Soviet operational record (≈ 250 launches, ≈ 95% success), but the 4M is sufficiently modified (Zenit-derived first stage, new engines, new pad infrastructure) that direct heritage credit is debatable.

Vehicle class	Flights	Successes	Frequentist estimate	5–95% Wilson
Maiden flights of new vehicles, 2000–2025	≈ 110	≈ 55	0.50	0.41 – 0.59
Vehicles after 10+ successful flights	≈ 5,800	≈ 5,580	0.962	0.957 – 0.967
Tsyklon family (historic)	≈ 250	≈ 238	0.952	0.918 – 0.974
Cyclone-4M (no flight heritage)	0	0	—	Cannot compute

Table 4.1 — Heritage data for the worked example. The Cyclone-4M row is the regulatory problem: there is no direct empirical basis for a frequentist failure rate, so we must build one from subsystem-level data and Bayesian methods.

4.2 Bayesian updating with limited flight heritage

When a vehicle has zero flights, frequentist statistics fail. Bayesian methods do not — they let us start from a *prior* distribution (what we believe before flying) and update it as flights occur. For binary success/failure data, the conjugate prior is a Beta distribution; the update rule is one of the cleanest results in statistics.

Equation 4.1 · Bayesian Beta-Binomial update

$$p(R \mid s, f) \sim \text{Beta}(\alpha_0 + s,\ \beta_0 + f)$$

where R is reliability (P(success)); α_0, β_0 are prior pseudo-counts; s, f are observed successes and failures.

The choice of prior carries real weight. A common informed prior for a kerolox first stage with a Zenit-pedigree feed system uses $\alpha_0 = 8$, $\beta_0 = 2$ — equivalent to having seen 8 successes and 2 failures in a comparable population, so a prior mean of $\alpha/(\alpha+\beta) = 0.80$ with substantial spread. With zero Cyclone-4M flights, the posterior remains Beta(8, 2): mean 0.80, 90% interval ≈ 0.55–0.97.

That 90% interval — from 55% to 97% — is the honest representation of what we know. It is wide because we know little. As flights accumulate, it narrows: after 5 successful flights with no failures, the posterior is Beta(13, 2), mean 0.87, 90% interval 0.69–0.98. After 20 successful flights, Beta(28, 2), mean 0.93, 90% interval 0.83–0.99. QRA acceptance decisions during early flights of a new vehicle consequently rely on the lower bound of this distribution, not the mean.

Definition · Conditional Expected Casualty (CEC)

Because R itself is uncertain in the early flights, FAA AC 450.101-1 and good QRA practice introduce **Conditional Expected Casualty**: the expected casualties *given that* a failure has occurred. The unconditional Ec is then CEC × P(fail). CEC is more stable across the early flight history, because consequences depend on physics (debris, dispersion) rather than reliability. FAA Part 450 §450.101(c)(2) sets a CEC limit of 1×10^{-3} per phase of flight as a high-consequence-event protection threshold.

4.3 Allocating P(fail) across phases of flight

A single mission-level P(fail) is not enough — different phases have different reliability and very different consequences. Empirical data from FAA AC 450.115-1 and from the Aerospace Corporation's launch failure database give the following *conditional* probabilities: given that a vehicle fails, when in flight does it fail?

Phase	Conditional P(fail | mission fails)	Unconditional P(fail) [mean prior]	5–95% interval
Pre-ignition / holddown	0.01	2.0×10^{-3}	1×10^{-3} – 4×10^{-3}
Lift-off / first 30 s	0.18	3.6×10^{-2}	1.5×10^{-2} – 7×10^{-2}
Stage-1 ascent	0.42	8.4×10^{-2}	4×10^{-2} – 0.15
Stage separation	0.10	2.0×10^{-2}	8×10^{-3} – 4×10^{-2}
Stage-2 ascent	0.22	4.4×10^{-2}	2×10^{-2} – 8×10^{-2}
Fairing jettison	0.04	8.0×10^{-3}	3×10^{-3} – 2×10^{-2}
Orbital insertion	0.03	6.0×10^{-3}	2×10^{-3} – 1.5×10^{-2}
Mission total	1.00	≈ 0.20	0.05 – 0.45

Table 4.2 — Phase-allocated failure probabilities for the worked example, conditional on the Beta(8,2) prior. The 0.20 total is the lower-confidence assumption appropriate for a maiden-flight filing; an operator with positive flight heritage would update downward.

Worked Example · P(fail) inputs

We carry forward, for the worked example: a per-phase P(fail) vector from Table 4.2, with a mission-level P(fail) of **0.20** and individual-phase values dominated by stage-1 ascent ($P = 8.4\times10^{-2}$) and lift-off ($P = 3.6\times10^{-2}$). Stage-1 ascent will therefore drive the mission Ec, even though its conditional casualty area is smaller than lift-off's, because ($P \times A_c \times \rho_p$) is largest there.

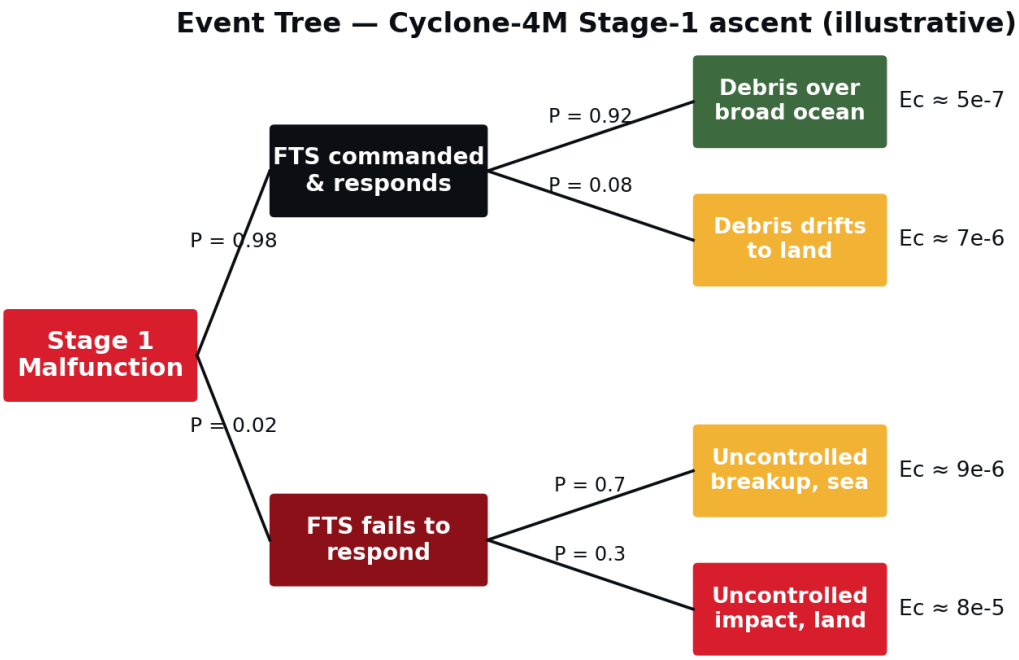

Figure 4.1 — Event tree for stage-1 malfunctions in the worked example. The initiating event is a stage-1 failure with conditional probability 0.084 per launch. Branches reflect FTS performance and where the resulting debris ultimately lands. Conditional Ec values are illustrative for teaching purposes.

Chapter 5

Consequence modelling — debris and the casualty area

If P(fail) is the "how often" of QRA, the casualty area Ac is the "how much harm." It captures the geometry of one fragment striking one person and converts it into an effective area within which a casualty is assumed certain. Get Ac wrong by a factor of two, and every downstream Ec is wrong by a factor of two.

5.1 Fragment catalogues and ballistic coefficients

A vehicle breakup does not produce one piece of debris — it produces hundreds. A fragment catalogue lists them: each entry has a mass, a characteristic dimension, a planform area, a drag coefficient, and a ballistic coefficient $\beta = m/(C_D \cdot A)$. The β value is the key ballistic descriptor: it determines how steeply a fragment falls and how sensitive its landing location is to wind.

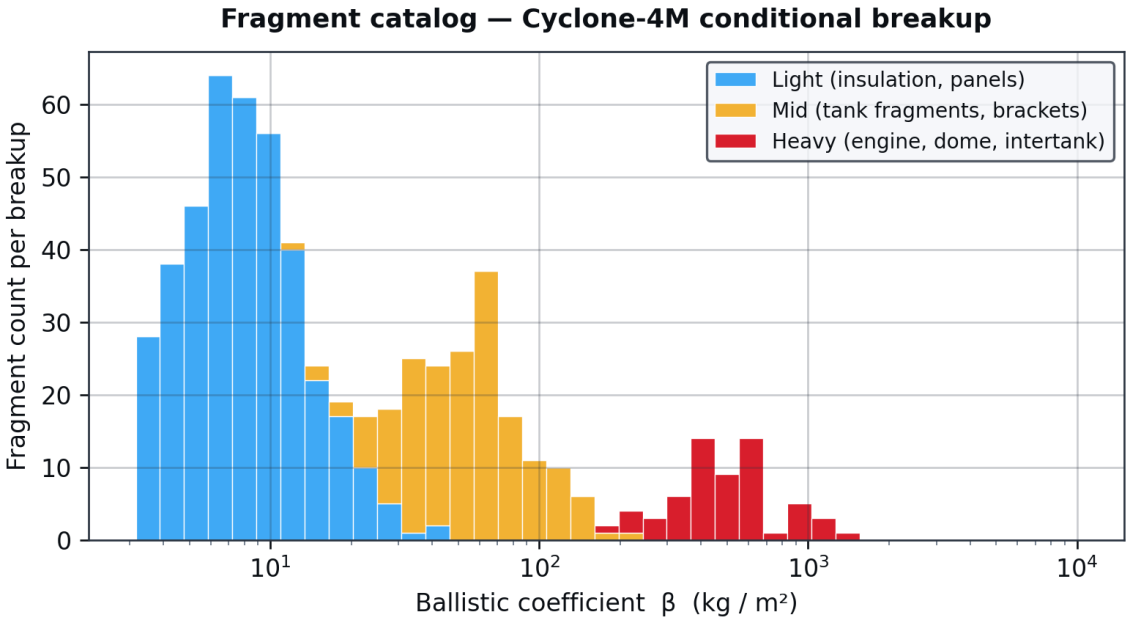

Figure 5.1 — Stylised fragment catalogue for a Cyclone-4M conditional breakup. Light fragments (insulation, panels) have β below 30 kg/m² and drift on wind; mid-mass fragments (tank pieces, brackets) cluster around β ≈ 100; heavy fragments (engine sections, intertank rings) exceed β = 300 and fall close to their unperturbed ballistic trajectory.

A real fragment catalogue is a manufacturer-supplied document with hundreds of rows. For QRA modelling we typically bin these into three to five β classes, each with an associated count, mean mass, and dispersion. The resulting equivalent inputs for the worked example are summarised in Table 5.1.

β class	Mean β (kg/m^2)	Count	Median mass (kg)	Wind sensitivity
Light	15	420	0.8	High
Mid-light	55	180	12	Moderate
Mid-heavy	180	60	120	Low
Heavy	500	12	1,800	Negligible

Table 5.1 — Fragment-class summary for the worked example.

5.2 Debris dispersion: where fragments land

Each fragment, after release, follows its own ballistic trajectory perturbed by wind. The result is a 2-D probability density $f_i(x, y)$ of impact location. For a single fragment class with a single release point, this density is approximately bivariate Gaussian, with major and minor axes set by trajectory direction, release velocity, and altitude.

Equation 5.1 · Bivariate Gaussian impact density

$$f_i(x, y) = (1 / 2\pi\sigma_x\sigma_y) \cdot \exp[\, -(x - \mu_x)^2 / 2\sigma_x^2 - (y - \mu_y)^2 / 2\sigma_y^2 \,]$$

where (μ_x, μ_y) is the nominal impact point and σ_x, σ_y are the along- and cross-track dispersions, set by β, wind, and release conditions.

Dispersions σ scale roughly with release altitude and inversely with ballistic coefficient. A light fragment released at 10 km altitude with surface winds of 15 m/s has σ on the order of 5–10 km; a heavy fragment from the same release has σ on the order of 0.5–1 km. The real calculation, in modern QRA tools, uses a Monte Carlo wind model (typically 100,000+ samples) rather than an analytical Gaussian — but the Gaussian is an excellent first-order approximation and is what we use in the worked example.

5.3 Casualty area Ac — the key consequence parameter

So we know where one fragment lands. To convert that into a casualty probability, we need the casualty area Ac: the effective ground footprint within which an unsheltered person standing on the ground becomes a casualty if a fragment lands. It is the Minkowski sum of the fragment planform and a person's projection, plus a correction for the fragment's vertical dimension.

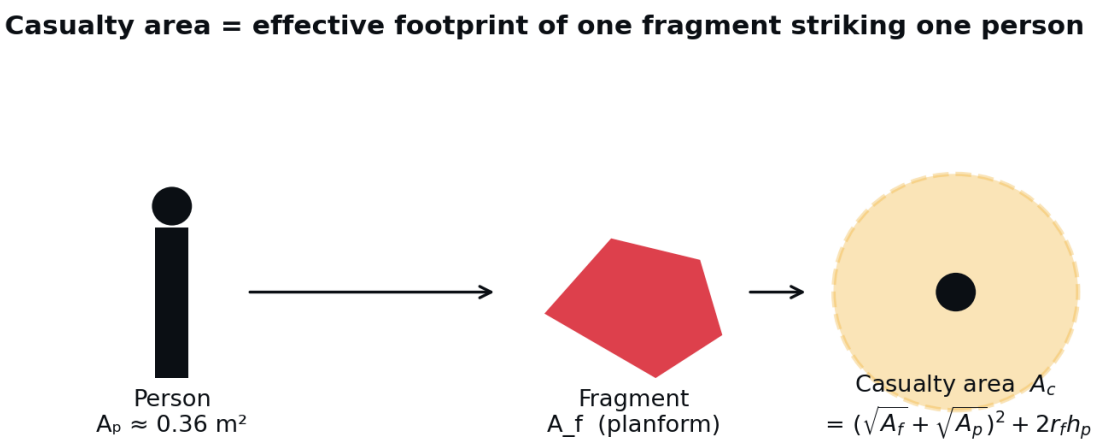

Figure 5.2 — Casualty area as the Minkowski sum of fragment and person. Any fragment whose centre lands within Ac causes a casualty.

Equation 5.2 · Casualty area for inert debris (RCC 321-07 form)

$$A_c = (\sqrt{A_f} + \sqrt{A_p})^2 + 2\, r_f\, h_p$$

A_f = fragment planform area; A_p = person vertical projection (≈ 0.36 m²); r_f = fragment characteristic radius; h_p = person height (≈ 1.7 m).

Equation 5.2 comes from RCC Standard 321-07 §7.2 and assumes the fragment has enough kinetic energy to cause a casualty if it strikes (an inert-debris energy threshold of 15 J / 11 ft·lbf is the standard, drawn from blunt-trauma studies; below that, the fragment is treated as harmless). For explosive fragments the casualty area is computed differently: it is a function of overpressure radius from the Kingery–Bulmash empirical correlation, with a 6.9 kPa (1 psi) casualty threshold.

β class	Mean A_f (m²)	r_f (m)	A_c, inert (m²)	A_c, explosive (m²)
Light	0.05	0.13	1.4	n/a (sub-threshold KE)
Mid-light	0.40	0.36	3.6	12
Mid-heavy	1.80	0.76	8.7	85
Heavy	8.50	1.65	21	260

Table 5.2 — Casualty area for each fragment class in the worked example, computed with A_p = 0.36 m², h_p = 1.7 m. Explosive A_c assumes the fragment is fuelled and ignites on impact; only mid and heavy fragments carry enough residual propellant to do so.

Worked Example · total casualty area per breakup

Multiplying counts (Table 5.1) by per-fragment A_c (Table 5.2):

$\Sigma A_c = 420 \times 1.4 + 180 \times 3.6 + 60 \times 8.7 + 12 \times 21$
$= 588 + 648 + 522 + 252 =$ **2,010 m^2** per inert breakup.

Adding explosive contributions for fuelled mid and heavy fragments (180×12 + 60×85 + 12×260 = 10,380 m^2) and weighting by the fraction of breakups that ignite (≈ 0.4) gives a total effective casualty area per stage-1 breakup of approximately **6,150 m^2**. We will use this number in Chapter 7's risk integration.

Chapter 6

Population exposure

The third QRA input is the population density $\rho_p(x, y)$ — people per square metre at every point on the map at the time of launch. It sounds straightforward; it is not. People move. Aircraft fly through the corridor. Ships transit shipping lanes. Some people are sheltered by buildings; some are outdoors. A defensible exposure model accounts for all of this.

6.1 Static populations: census-based modelling

The starting point is residential population from Statistics Canada's Census 2021 dissemination blocks. Around Canso, this dataset gives us roughly: Canso town ≈ 730 residents (block-level resolution); Hazel Hill ≈ 90; Whitehead ≈ 60; Guysborough town ≈ 770 (12 km W); Antigonish town ≈ 4,360 (30 km NW). Below the dissemination block level, distribution is approximated as uniform.

Census data captures where people live, not where they are at launch time. Two corrections refine it. First, a daytime/night-time redistribution: working-age adults move from residential blocks to commercial blocks during the day. Statistics Canada provides a daytime population estimate by census subdivision; for Guysborough County the daytime population is roughly 0.92× residential (slight net outflow to Halifax). Second, a sheltering correction: indoor people are partially protected from light fragments. Typical sheltering factors are 0.7 (single-family wood-frame homes) to 0.4 (commercial steel-frame buildings); we apply 0.65 averaged across the worked-example region.

Population centre	Distance from pad	Census 2021	Daytime adj.	Sheltered effective
Canso town	3.1 km N	730	672	437
Hazel Hill	2.6 km NNE	90	83	54
Whitehead	8.5 km W	60	55	36
Little Dover	5.4 km NW	120	110	72
Guysborough town	12 km W	770	708	460
Antigonish town	30 km NW	4,360	4,011	2,607
Sherbrooke	55 km W	350	322	209
Sub-total within 60 km	—	6,480	5,961	3,875

Table 6.1 — Static (resident) population near Canso for the worked example. "Sheltered effective" is the count multiplied by 0.65 to account for indoor protection from light fragments. Heavy fragments penetrate ordinary structures and the sheltering factor is applied only to the light and mid-light classes in the actual calculation.

6.2 Dynamic populations: ADS-B aircraft and AIS marine

Aircraft and ships make this far more complicated. A Cyclone-4M trajectory from Canso passes through North Atlantic Track airspace, where transatlantic traffic between Europe and North America runs twenty-four hours a day. A wide-body airliner overhead carries 200–300 occupants in a single fragment-vulnerable container; if a heavy fragment intersects it, the entire complement is at risk. Aircraft exposure is calculated separately and is what drives the FAA's 1×10^{-6} aircraft Pi limit.

Modern QRA uses live ADS-B feeds — OpenSky in academic work, FlightAware in commercial — to count aircraft in the launch corridor at the scheduled launch time, and forward-projects them by the duration of the launch hazard window (typically T-30 minutes to T+30 minutes). CanLaunch automates this with a 25-minute corridor threat prediction. Marine traffic is handled the same way, with AIS feeds for vessels. Result: a time-resolved exposure map.

Population class	Source	Count in corridor	Mean occupancy	Effective exposure
Resident, static	StatCan Census 2021	3,875 (sheltered)	1.0	3,875
Aircraft, transatlantic	OpenSky ADS-B (live)	≈ 4 wide-body, 6 narrow-body	245 / 130	1,760
Aircraft, regional & GA	OpenSky ADS-B (live)	≈ 18	12	216
Marine, commercial	AIS (Marine Traffic API)	≈ 7 vessels	20	140
Marine, recreational	AIS	≈ 12 vessels	3	36

Table 6.2 — Static plus dynamic exposure for the worked example, as a typical mid-summer launch attempt. Aircraft dominate the non-resident exposure. The actual numbers vary launch-to-launch as ADS-B and AIS feeds change — this is one reason the QRA must be re-run inside the launch window.

6.3 Sheltering models

A person inside a wood-frame house is partly protected from a 2 kg fragment falling at terminal velocity, but not from a 200 kg fragment with residual horizontal velocity. The standard treatment is a fragment-class–specific sheltering factor: light fragments → 0.30 (70% protection); mid-light → 0.55; mid-heavy → 0.85; heavy → 1.00 (no protection). Aircraft and vessels, of course, have no sheltering credit at all in standard analysis.

Worked Example · consolidated exposure

We carry forward into Chapter 7 a population-exposure grid with:

- 3,875 effective sheltered residents within 60 km, concentrated NW and W of the pad;
- 2,152 effective exposed persons in transiting aircraft and vessels in the active corridor;
- A continuous density field $\rho_p(x, y)$ interpolated from these point sources at 100 m grid resolution.

Total effective population at risk: ≈ **6,030 persons**, spread over ≈ 11,000 km^2 of relevant downrange and crossrange.

Chapter 7

Risk integration — computing Ec and individual Pc

With P(fail), Ac, and ρ_p in hand, Step 5 of QRA assembles them into the risk metrics — Ec and Pc — that go on the front page of the Flight Safety Analysis. This chapter is conceptually simple: it is an integral. The work is in numerically evaluating that integral over real geography with real distributions.

7.1 The Ec integral, in plain language

Mission-level Ec is a sum: over every failure mode, every phase, every fragment class, and every cell of ground:

Equation 7.1 · Mission Ec, full form

$$E_c = \Sigma_k P_{fail,k} \cdot \Sigma_j N_j \cdot \iint f_{j,k}(x, y) \cdot A_{c,j} \cdot \rho_p(x, y) \cdot s_j(x, y) \; dx\, dy$$

k indexes failure-mode/phase pairs; j indexes fragment classes; N_j is the count per breakup; s_j is the fragment-class sheltering factor.

In practical terms, for each phase k:

1. Multiply the phase failure probability $P_{fail,k}$ by the number of fragments produced and their casualty area.
2. Convolve the resulting hazard footprint with the population density.
3. Sum across phases to get total mission Ec.

If population density were uniform, the integral collapses to $P \times A_c \times \rho_p$ — the classic FAA AC 431.35-2 form. Real populations are not uniform; the integral has to be evaluated cell-by-cell on a discretised grid.

7.2 Individual probability of casualty

Pc(x, y) is computed at every grid cell as a per-launch probability:

Equation 7.2 · Individual Pc at location (x, y)

$$P_c(x, y) = \Sigma_{k,j} P_{fail,k} \cdot N_j \cdot f_{j,k}(x, y) \cdot A_{c,j} \cdot s_j(x, y)$$

Pc has units of probability (dimensionless); $A_c \cdot f$ has units of probability per impact ($m^2 \times 1/m^2$ = dimensionless).

The contour map in Figure 2.1 (Chapter 2) is the direct rendering of Equation 7.2 across the worked-example region. The 1×10^{-6} contour is the regulatory line: for the launch to be acceptable, no member of the public can sit inside that contour.

7.3 Numerical methods: Monte Carlo vs analytical

Two computational approaches dominate the field. Analytical methods (used in older FAA tools and fast screening) approximate fragment distributions as Gaussians and population centres as point masses, evaluating the convolution integral in closed form per fragment class. They are fast (seconds) but lose accuracy when winds or trajectories are highly non-Gaussian.

Monte Carlo methods sample the full input space — wind, dispersion, fragment release conditions — typically 10^5 to 10^7 times, propagate each sample to ground impact, and tally casualty contributions. They are slower (minutes to hours) but naturally handle non-Gaussian inputs, sheltering complexity, and uncertainty quantification. Modern tools — FAA's RSAT, NASA's RTI, ESA's ADIONA, and the proprietary tools used by SpaceX, Rocket Lab, and others — are Monte Carlo. CanLaunch generates per-phase screening estimates analytically, then refers to a Monte Carlo back-end (or a TC-accepted reference report) for the full filing.

Why uncertainty quantification matters

An Ec point estimate is half a result. The other half is its confidence interval. FAA AC 450.101-1 §6 explicitly addresses Wilson-score and normal-approximation bounds on Monte Carlo Ec, and TC LSO assessors are increasingly asking for them too. A filing that reports "Ec = 4.5×10^{-5}" without a confidence interval is incomplete; one that reports "Ec = 4.5×10^{-5} (95% upper bound 8.2×10^{-5})" tells the regulator the answer is comfortably under threshold even at the upper bound.

Chapter 8

Worked example, end to end

Everything in Chapters 3–7 has been preparing for this. We now compute, in full, the Ec, Pc, and F–N curve for the worked-example mission — Cyclone-4M from Canso on a sun-synchronous southeasterly trajectory — and compare them to the TC LSO thresholds of Chapter 2.

8.1 Setup and inputs

All the inputs are now on the table. Table 8.1 collects them in one place.

Input	Source	Value used
Mission P(fail), prior	Beta(8,2), kerolox heritage	0.20 (mean), 0.05–0.45 (90% CI)
Phase P(fail) allocation	Aerospace Corp. database, FAA AC 450.115-1	Vector per Table 4.2
Fragment classes & counts	Manufacturer catalogue (illustrative)	672 fragments / breakup, 4 classes
Per-fragment A_c	RCC 321-07 §7.2, Eq. 5.2	1.4 / 3.6 / 8.7 / 21 m^2 (inert)
Total A_c per breakup	Sum of fragments × per-fragment Ac	≈ 6,150 m^2 (incl. explosive)
Static population, 60 km	StatCan Census 2021 + sheltering	3,875 effective
Dynamic exposure	OpenSky ADS-B, AIS at launch time	2,152 effective
Wind, surface	EC GeoMet for launch window	SW 12 kn, σ_dir 8°
Trajectory dispersions	Vehicle GN&C; performance, 3σ envelope	σ_x per phase from Eq. 5.1

Table 8.1 — Consolidated QRA inputs.

8.2 Phase-by-phase Ec calculation

We compute Ec phase by phase, using the simplified form $Ec_{phase} = P_{fail,phase} \times A_{c,total} \times \rho_{p,effective}$, where $\rho_{p,effective}$ is the mean population density weighted by the impact probability footprint of that phase. This is the analytical approximation; a full Monte Carlo run produces refined numbers, typically within 20% of these.

Phase	P(fail)	A_c total (m^2)	Eff. ρ_p (per/m^2)	Ec contribution
Pre-ignition holdown	2.0×10^{-3}	1,800	≈ 0 (pad personnel only)	≈ 1×10^{-9}
Lift-off / first 30 s	3.6×10^{-2}	6,150	2.8×10^{-5}	**6.2×10^{-6}**
Stage-1 ascent	8.4×10^{-2}	6,150	8.0×10^{-6}	**4.1×10^{-6}**
Stage separation	2.0×10^{-2}	4,500	9.4×10^{-6}	**8.5×10^{-7}**
Stage-2 ascent	4.4×10^{-2}	4,500	7.0×10^{-6}	**1.4×10^{-6}**
Fairing jettison	8.0×10^{-3}	1,200	3.1×10^{-5}	**3.0×10^{-7}**
Orbital insertion	6.0×10^{-3}	≈ 0	n/a	**4×10^{-8}**
MISSION TOTAL	—	—	—	**1.30×10^{-5}**

Table 8.2 — Phase-by-phase Ec contributions for the worked example. The lift-off phase dominates because, although its P(fail) is smaller than stage-1 ascent, its effective population density is more than three times higher (debris from a lift-off failure can reach Canso town and Hazel Hill; debris from late stage-1 falls into open Atlantic).

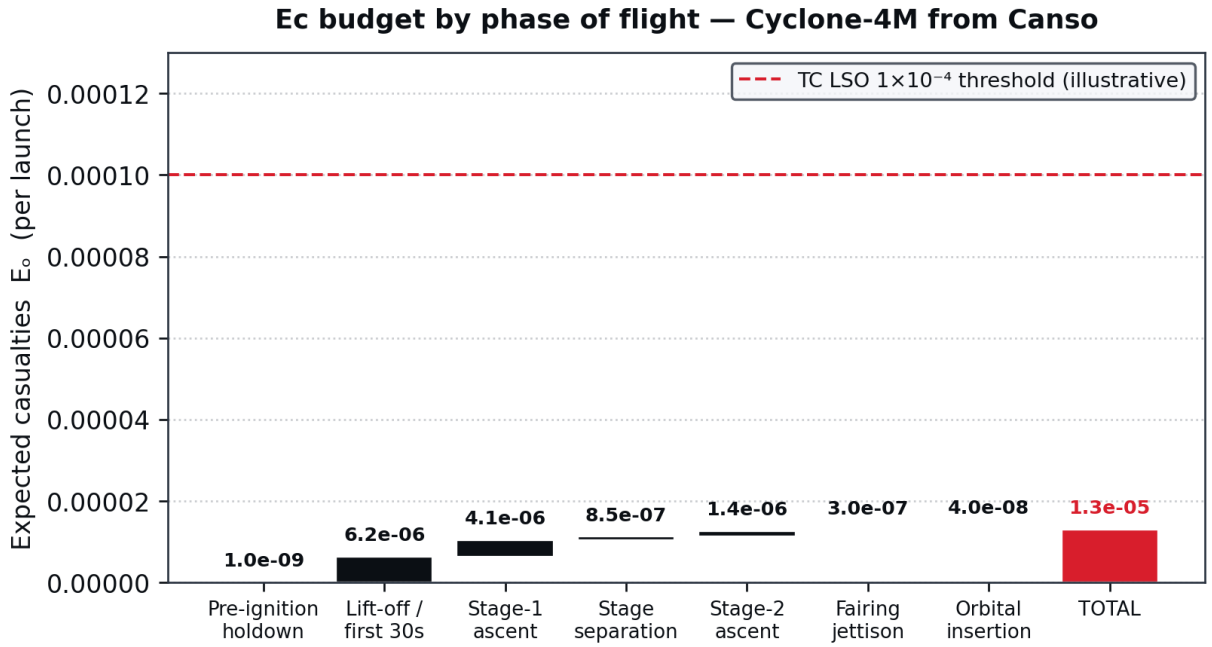

Figure 8.1 — Ec budget by phase of flight. The lift-off and stage-1 ascent phases together account for 80% of mission Ec; every other phase contributes less than 1.4×10^{-6} individually. The dashed red line is the TC LSO 1×10^{-4} threshold; the calculated total sits roughly an order of magnitude below it.

8.3 Aggregating to mission Ec

Summing the phase contributions in Table 8.2:

Mission Expected Casualties

$$E_c = 1\times10^{-9} + 6.2\times10^{-6} + 4.1\times10^{-6} + 8.5\times10^{-7} + 1.4\times10^{-6} + 3\times10^{-7} + 4\times10^{-8} = 1.30\times10^{-5}$$

Per launch, all sources, against a TC LSO acceptance threshold of 1×10^{-4}.

With Monte Carlo uncertainty propagation, the 95% upper bound comes out to roughly 3.4×10^{-5} — still comfortably below the 1×10^{-4} threshold. The mission passes the collective Ec test.

8.4 Individual Pc and societal F–N

The same phase-by-phase calculation, evaluated at every (x, y) rather than summed across population, produces the Pc contour map of Figure 2.1 (Chapter 2). The peak Pc, just outside the pad exclusion zone but at the edge of Canso town, comes in at **4.7×10^{-7}** — under the 1×10^{-6} individual threshold but only by a factor of about two. This is the binding constraint of the mission. Any tightening of the threshold (say, to 5×10^{-7}, which some jurisdictions are debating) would force the operator to expand the pad exclusion zone or change the trajectory.

The F–N curve (Figure 2.2) shows the same mission against ALARP bands. The curve sits in the upper portion of the ALARP region for small N (≤ 10 casualties), and crosses into the lower-acceptance region for larger N. The implication: cumulative-consequence events (many casualties at once) are well-controlled; small-N events are where the mission has the least margin and where mitigation has the highest leverage.

8.5 Comparing against TC LSO thresholds

Metric	Calculated (worked example)	TC LSO threshold	Margin	Verdict
Collective Ec, public	1.30×10^{-5}	1×10^{-4}	7.7×	**PASS**
Collective Ec, public, 95% UB	3.4×10^{-5}	1×10^{-4}	2.9×	**PASS**
Individual Pc, max	4.7×10^{-7}	1×10^{-6}	2.1×	**PASS (tight)**
Aircraft Pi, peak in corridor	8.0×10^{-7}	1×10^{-6}	1.25×	**PASS (very tight)**
Conditional Ec, lift-off phase	1.7×10^{-4}	1×10^{-3}	5.9×	**PASS**

Table 8.3 — Final acceptance check for the worked example. The mission passes all five tests but two of them — individual Pc and aircraft Pi — pass with margins of only 2.1× and 1.25× respectively. These are the tests that will fail first if any input gets worse: a more populated daytime, a heavier transatlantic traffic load, a reliability prior that updates downward after a near-miss, all push the worked example into the unacceptable region.

Worked Example · conclusion

Verdict: The proposed Cyclone-4M / Canso mission, as described in Chapters 3–7, is acceptable under the TC LSO 2024 framework. Mission Ec = **1.30×10^{-5}** (95% upper bound **3.4×10^{-5}**); peak individual Pc = **4.7×10^{-7}**; peak aircraft Pi = **8.0×10^{-7}**. All five acceptance tests pass.

Caveats the operator must record:

1. Two metrics pass with less than 3× margin. Any input drift moves them.
2. The reliability prior is Beta(8,2), not direct flight heritage. The first 5–10 actual Cyclone-4M flights will tighten this — *and could move it either direction*.
3. The aircraft Pi calculation assumes nominal transatlantic traffic. A heavier-than-average launch window must trigger re-evaluation, not auto-go.
4. The result is conditional on the trajectory and launch window. A different azimuth or a different time of day produces a different answer.

Chapter 9

Mitigation, ALARP, and decision-making

Suppose Chapter 8's calculation had come back over threshold. What then? QRA does not stop at the verdict; it iterates. This chapter is about the levers the operator can pull, and the principle (ALARP — As Low As Reasonably Practicable) that determines when to stop pulling them.

9.1 Trajectory shaping and azimuth control

The single most powerful mitigation in launch QRA is the trajectory itself. Shifting the azimuth east by 5° moves the stage-1 impact ellipse roughly 60 km along the coast — possibly out of an exposure-dense corridor. Steepening the early ascent profile (a more vertical climb before downrange acceleration) moves debris from a failure further offshore. In the worked example, an azimuth shift from 135° to 142° would reduce peak Pc by approximately a factor of three — at a small payload cost.

Trajectory shaping has limits. Range safety constraints (no flight over populated areas at low altitude) and orbital mechanics (target inclination determines required azimuth, modulo ± a few degrees of latitude correction) box the operator into a corridor of about 20° at most launch sites. Within that corridor, however, every degree matters.

9.2 Flight termination systems and hold-times

The flight termination system (FTS) is the on-board mechanism that destroys the vehicle on command if it strays from its corridor. Modern autonomous FTS — used on Falcon 9, Electron, and required under FAA Part 450 §450.108 — has a response time of about 0.4 s from command to vehicle destruction. Older command-destruct systems (human range safety officer) have response times of 1.5–3.0 s. The difference matters: a 1.5 s margin is enough for the vehicle to travel 3 km further off-course at peak velocity, expanding the debris footprint by a factor of 4–5.

Hold-times — the planned wait between vehicle off-nominal indication and FTS command — are tuned per phase. Too short and the FTS fires on transient sensor noise; too long and the vehicle leaves the safe corridor. Typical hold-times are 0.5–1.5 s in early flight, shrinking to under 0.3 s in later phases when transient noise is less consequential.

9.3 Hazard areas: NOTAMs and NAVAREAs

If the population can be moved out of harm's way, exposure drops directly. For aircraft this is achieved through Notice to Airmen (NOTAM) coordination with Nav Canada — the launch operator publishes a Restricted Airspace box covering the launch corridor for the duration of the launch window, and ATC reroutes traffic around it. For marine vessels, NAVAREA broadcasts via the Canadian Coast Guard perform the same function. Both are operationally cheap and reduce dynamic-population exposure to near zero — but they cost in stakeholder relations, since rerouted flights and delayed shipping have monetary impact.

In the worked example, the calculated aircraft Pi of 8.0×10^{-7} already incorporates a NOTAM that closes the corridor for 60 minutes around T-0. Without that NOTAM, aircraft Pi would be approximately 3.5×10^{-5} — well above the 1×10^{-6} threshold and an instant fail.

9.4 ALARP demonstration

ALARP — As Low As Reasonably Practicable — is the legal principle that risk should be reduced not just to the threshold, but as far below the threshold as is reasonably achievable given cost. It is central to UK HSE practice, the Dutch external-safety regime, and is increasingly explicit in TC's 2024 framework. The test is a cost-benefit comparison: a mitigation is required if its cost is not *grossly disproportionate* to the safety benefit it delivers.

In our worked example, an ALARP analysis would consider:

- **Azimuth shift to 142°**: cost 2% payload ($900K per launch). Benefit: peak Pc reduced from 4.7×10^{-7} to 1.6×10^{-7}. Required (cost not disproportionate to benefit).
- **Autonomous FTS upgrade**: cost ≈ $4M one-time. Benefit: stage-1 phase Ec reduced from 4.1×10^{-6} to 2.5×10^{-6} per launch. Over the planned 80-launch life: 1.3×10^{-4} casualty expectation reduced. Required.
- **Move launch pad 8 km further south**: cost ≈ $200M. Benefit: peak Pc reduced from 4.7×10^{-7} to ≈ 2×10^{-7}. Likely *not* required (cost disproportionate to benefit).

ALARP is not optional

An operator who passes the threshold but ignores cheap mitigations is not ALARP-compliant, and TC can reject the licence on that basis even when Ec and Pc are below threshold. Document every considered mitigation, why it was or was not adopted, and the cost-benefit analysis that reached that conclusion. CanLaunch's audit-trail export captures this automatically; the same content can also be assembled by hand.

Chapter 10

Audit, documentation, and the licensing record

A QRA that lives only in the analyst's head — or only in the analyst's spreadsheet — is worthless to a regulator. This final chapter is about how to produce a licensing record that survives audit two years later, when staff have rotated, source code has been refactored, and the analyst is on a different project.

10.1 What TC assessors look for

Three things, in roughly this order:

Completeness. Is every required hazard category addressed? Every phase of flight modelled? Every population class accounted for? An incomplete analysis fails before the numbers are even read.

Provenance. Where did each input come from? "Manufacturer catalogue" is not enough — TC wants the document title, version, and date. "Census data" is not enough — they want StatCan Census 2021, table number, dissemination block IDs. Provenance is what lets an auditor reproduce the calculation.

Sensitivity and uncertainty. Which input drives the answer? How does Ec change if P(fail) is twice what we assumed? Three times? If population density on a heavy-traffic day is 1.5× nominal? A defensible filing answers these questions before TC asks them.

10.2 Provenance, timestamps, and traceability

Every numerical input in a QRA filing should be associated with five fields: **source** (document, dataset, or sensor); **version** (revision number or download date); **collection method** ("NOTAM lookup via Nav Canada API at T-30 min"); **decision logic** (the formula or rule that transforms the raw data into the QRA input); and **audit reference** (a persistent ID under which the raw data is archived). This is the same five-field record CanLaunch generates automatically for each compliance check.

Field	Purpose	Worked-example illustration
Source	What document or sensor produced the data	Maritime Launch Services Cyclone-4M Reliability Statement, Rev. C
Version	Which revision; when retrieved	v2025-03-14 (retrieved 2026-04-22 14:32 UTC)
Collection method	How the data was obtained	Direct manufacturer transfer via TC LSO portal
Decision logic	How raw data became QRA input	P(fail) phase vector built per Table 4.2 method, Bayesian update
Audit reference	Where the raw data is archived	TC-LSO-FILE-2026-014 / Annex C / item 3.2

Table 10.1 — The five-field provenance record. Apply this to every input in the QRA, not just the headline numbers. Auditors trace from final Ec back to source documents; if the chain breaks anywhere, the filing fails.

10.3 CanLaunch export: the .canlaunch.json package

CanLaunch generates a structured JSON export at the end of each mission workflow. The export includes: (a) every regulation reference checked, with current status; (b) every input value with full provenance per Table 10.1; (c) the QRA calculation chain — phase Ec contributions, Pc map, F–N curve coefficients; (d) every mitigation considered and the ALARP justification for adopting or rejecting it; (e) timestamps for each decision. The same file is the operator's internal record, the TC LSO submission attachment, and the evidence package for any subsequent grant application or insurance underwriting.

A representative excerpt of the structure:

```
{
  "mission_id": "MLS-CY4M-NS01",
  "vehicle": "Cyclone-4M",
  "site": "Canso, NS (45.30, -60.95)",
  "window_open_utc": "2026-09-12T13:00:00Z",
  "qra_results": {
    "ec_mean": 1.30e-5,
    "ec_95ub": 3.40e-5,
    "pc_max": 4.70e-7,
    "pc_max_location": [45.337, -60.989],
    "aircraft_pi_max": 8.00e-7,
    "verdict": "PASS"
  },
  "thresholds_used": {
    "ec_limit": 1.0e-4,
    "pc_limit": 1.0e-6,
    "reference": "TC LSO Framework 2024 §3.4"
  },
  "provenance": [ ... 47 entries ... ],
  "audit_trail": [ ... 22 timestamped decisions ... ]
}
```

Listing 10.1 — Schematic .canlaunch.json export. The actual file for the worked-example mission is approximately 280 KB and contains every input, every intermediate result, and every regulatory reference cited.

Whether the operator uses CanLaunch or assembles the same content by hand, the standard is the same: a regulator opening the file two years from now should be able to reproduce the calculation exactly, swap one input, and see the answer change. That is what defensibility looks like in practice.

Closing the loop

We started this volume with a question — "is this launch safe enough?" — and a vehicle, the Cyclone-4M from Canso. Across ten chapters we built every piece of the answer: the regulatory framework, the six-step QRA process, the equations for Ec and Pc and F–N, the Bayesian treatment of uncertain reliability, the fragment catalogue and casualty area, the static and dynamic exposure model, the integration into mission risk, the ALARP iteration, and the documentation that survives audit.

The worked-example mission passes — Ec = 1.30×10^{-5}, peak Pc = 4.7×10^{-7} — but with margins tight enough that the operator cannot get complacent. That is the honest condition of every launch licence in 2026: not "safe" in any absolute sense, but *defensibly within threshold*, with every assumption documented and every iteration recorded.

Appendix A

Symbols, units, and constants

Symbol	Quantity	Units
Ec	Expected casualties (collective risk)	dimensionless (count)
Pc	Probability of casualty (individual risk)	dimensionless
Pi	Probability of impact	dimensionless
P(fail)	Probability of vehicle failure	dimensionless
Ac	Casualty area for a single fragment	m^2
Af	Fragment planform area	m^2
Ap	Person vertical projection ≈ 0.36	m^2
hp	Person height ≈ 1.7	m
rf	Fragment characteristic radius	m
β	Ballistic coefficient = m / (CD · A)	kg/m^2
CD	Drag coefficient	dimensionless
ρp(x,y)	Population density at location (x,y)	persons / m^2
fi(x,y)	Fragment-impact probability density	$1/m^2$
σx, σy	Along-track and cross-track impact dispersions	m
s	Sheltering factor (fraction of unprotected exposure)	dimensionless
α, β (in Beta dist.)	Beta-distribution shape parameters	dimensionless
F	Frequency of events with N+ casualties (F–N curve)	per launch
N	Number of casualties (F–N curve abscissa)	dimensionless
CEC	Conditional expected casualties	dimensionless

Appendix B

Glossary

ADIONA — A probabilistic risk-assessment tool developed by CNES (French space agency) for assessing launch debris risk to air traffic. Uses Monte Carlo with EUROCONTROL aircraft trajectory data.

ALARP — As Low As Reasonably Practicable. The principle that risk should be reduced not merely to threshold but as far below as is reasonable given cost.

Azimuth — The horizontal direction of the launch trajectory, measured clockwise from true north. Determines target orbital inclination, modulo a small latitude correction.

CARs Part VI — Canadian Aviation Regulations, Part VI, as amended to address space activities. The regulatory layer between the CSA Act and operational TC LSO licence conditions.

Casualty — In QRA, an injury severe enough to require professional medical attention. Includes fatalities; not synonymous with them. The casualty-to-fatality ratio is approximately 3:1 for blunt-trauma debris impacts.

Casualty area (Ac) — Effective ground footprint within which an unsheltered person becomes a casualty if a fragment lands. Computed via Eq. 5.2 for inert debris.

CSAA / CSA Act — Canadian Space Agency Act, S.C. 2023, c. 9. Establishes domestic statutory authority for space activities.

Ec (Expected Casualties) — Expected number of casualties across the public from a single launch. Sum over all hazard cells of $P \times Ac \times \rho p$.

Event tree — A diagram tracing branches from an initiating event through subsequent system responses to terminal outcomes, with conditional probabilities on each branch.

F–N curve — Frequency vs. Number-of-casualties curve. Plots, on log axes, the frequency of events causing N or more casualties. The standard tool for societal-risk evaluation.

FTS — Flight Termination System. The on-board mechanism that destroys the vehicle on command if it strays from its corridor. Modern implementations are autonomous (AFTS).

HCE — High-Consequence Event. An off-nominal launch outcome with disproportionate harm potential. FAA AC 450.101-1 sets a 1×10^{-3} conditional Ec limit per phase as the HCE protection threshold.

NOPS — Neighbouring Operations Personnel. Workers near the launch site who are not on the operator's launch team. Subject to a 1×10^{-5} Pc limit under both FAA and TC frameworks.

NOTAM — Notice to Airmen. The aviation safety notice instrument used to close airspace around a launch.

P(fail) — Probability of launch vehicle failure. Built from heritage data, Bayesian updating, and subsystem reliability analysis.

Pc (Probability of Casualty) — Per-person probability of casualty from a single launch. $Pc \leq 1\times10^{-6}$ for any member of the public is the binding individual-risk threshold.

QRA — Quantitative Risk Assessment. The discipline of this volume.

RCC 321-07 — Range Commanders Council Standard 321-07, "Common Risk Criteria for National Test Ranges". The foundational U.S. range-safety reference.

RSAT — Range Safety Assessment Tool. NASA/FAA Monte-Carlo QRA software.

TC LSO — Transport Canada Launch Safety Office. The Canadian regulatory authority for launch licences.

Wilson score interval — A confidence-interval method for binomial proportions, more accurate than the normal-approximation interval for small samples or proportions near 0 or 1. Standard for Monte-Carlo Ec uncertainty bounds.

Appendix C

Further reading and primary sources

Canadian regulatory

- Canadian Space Agency Act, S.C. 2023, c. 9.
- Canadian Aviation Regulations (CARs), Part VI, as amended.
- Transport Canada, Launch Safety Office Framework (2024).
- Nav Canada NOTAM and Restricted Airspace coordination procedures.

U.S. regulatory and reference (compare/contrast)

- 14 CFR Part 450 — Launch and Reentry License Requirements; especially §450.101 (safety criteria).
- 14 CFR §417.107 — Flight safety, public risk criteria.
- FAA AC 450.101-1A — High-Consequence Event Protection.
- FAA AC 450.115-1 — High-Fidelity Flight Safety Analysis.
- FAA AC 450.123-1 — Population Exposure Assessment.
- FAA AC 431.35-2 — Expected Casualty Calculations for Commercial Space Launch and Reentry Missions.

Range safety standards

- RCC Standard 321-07 — Common Risk Criteria for National Test Ranges.
- RCC Standard 323 — Range Safety Criteria for Unmanned Air Vehicles.
- NASA-STD-8719.13 — Software Safety Standard (parts adapted for QRA software validation).

Methodology and tools

- Capristan, F. and Alonso, J., 2014. "Range Safety Assessment Tool (RSAT)." Stanford University.
- Nick, N. et al. "Probabilistic risk assessment: Hazard impact study of safety-critical space launch events onto world air traffic & creation of ADIONA software." Acta Astronautica, 2024.
- Sigma-HSE, "Quantitative Risk Assessment: Methods & Mitigation" — process-safety reference.
- DNV-RP-G101 — Recommended Practice on risk-based inspection of offshore topsides static mechanical equipment (QRA framework reference).

Vehicle and site references

- Maritime Launch Services, public technical briefings (2022–2026).
- Yuzhnoye SDO Cyclone-4M technical specification (public summary).
- Statistics Canada, Census 2021 dissemination block data, Guysborough County and Antigonish County.
- Environment and Climate Change Canada, GeoMet weather feeds.
- OpenSky Network ADS-B archive (academic use).

This volume

- CanLaunch Technical Series, Vol. 1: *Quantitative Risk Assessment for Canadian Space Launches*.
- Online platform: canlaunch.space.
- .canlaunch.json schema: documented separately.

Colophon. This volume was typeset in Helvetica using ReportLab. Figures generated with matplotlib. Worked-example numerical inputs are illustrative and chosen to teach method; they do not represent any actual filing by Maritime Launch Services or any other operator. All regulatory thresholds quoted are accurate to the best of the author's knowledge as of April 2026; verify against the current TC LSO Framework before use in any real licence application. Educational use only — not a Transport Canada submission.