ALTO HSR Citizen Research Initiative · Methodology Paper

A Framework for Independent Evaluation of the ALTO HSR Project

An annual fiscal ledger framework and a seven-stage analytical pipeline

Purpose

This paper sets out the analytical methodology the ALTO HSR Citizen Research Initiative (CRI) has developed to evaluate major federal infrastructure projects. The methodology is built around a single accounting identity — the annual fiscal ledger every operating rail corridor must balance — and a seven-stage analytical pipeline that produces a defensible value for each of its terms. The pipeline is replicable, transparent, and grounded in established academic and federal-policy literature. It has been applied in full to the ALTO high-speed rail proposal; the supporting CRI research notes listed in the Sources section at the end of this paper provide the detailed methodology, results, and underlying data for each stage. The paper is intended to support independent review by the Parliamentary Budget Officer or another federal analytical body, with the supporting documents as appendices.

The annual fiscal ledger framework

Every operational rail corridor must balance an annual fiscal ledger. The costs of running the corridor in a given year — debt service on the capital outlay, plus operations and maintenance expenditure, plus fleet capital renewal — must equal the revenue collected from those who use the corridor and from those who subsidise it. The identity is:

Capex × CRF + O&M and fleet capital = Ridership × Fare + Public subsidy + Land value capture

(annual debt service on capital) + (annual operating cost) = (annual farebox revenue) + (annual public subsidy) + (annual LVC contribution)

The identity is an accounting truism; what makes it analytically useful is that each of its terms is independently anchored. None can be set at will; each has a defensible value that emerges from a specific empirical or engineering methodology, rather than from political assertion. The seven-stage pipeline produces the value for each term.

Ledger term	What anchors its value	Pipeline stage that produces it
Capex × CRF (annual debt service)	Capital outlay times capital recovery factor; CRF reflects cost of capital and debt term. Capex itself is calibrated against the international cost-overrun reference class,	Stages 1 and 2 (ECI, CFI) → Stage 3 (capex RCF). CRF is a financing parameter.

Ledger term	What anchors its value	Pipeline stage that produces it
	with engineering and community complexity as predictors.	
O&M and fleet capital (annual operating cost)	Estimated bottom-up from corridor asset inventory and service-level inputs across three streams (maintenance and renewals, operating categories, fleet capital recapitalisation). Cold-climate and route-specific factors raise the benchmark; international comparators used for cross-validation only.	Stage 4 (engineering ground-up O&M cost estimation).
Ridership × Fare (annual farebox revenue)	Modal-shift S-curves applied to corridor demographics fix the ridership achievable at any given fare. The (ridership, fare) pair sits on the corridor's modal-shift frontier; the two cannot be chosen independently.	Stage 5 (reference-class modal-shift ridership analysis).
Public subsidy (federal + provincial)	Residual: whatever closes the gap between annual costs and the sum of farebox plus LVC. Bounded below by zero (the corridor cannot pay passengers to board) and above by total cost.	Output of the ledger once the other four terms are set; reported as part of Stage 5 (subsidy frontier).
Land value capture (LVC)	Supplementary revenue from land value uplift attributable to the corridor — station-area development levies, business-rate supplements, betterment charges, joint development. Bounded above by the actual attributable land value increment.	Where a project discloses an LVC mechanism, the methodology evaluates its bounding values against international comparators.

Each piece is independently anchored. Debt service is set by capex (Stages 1–3) and the cost of capital. O&M is set by engineering benchmarks against international HSR operators (Stage 4). Farebox revenue is set by ridership times fare on the modal-shift frontier (Stage 5). The proponent's choice to claim a particular capex, ridership, and fare combination does not satisfy the ledger unless the corresponding subsidy figure is also disclosed and the fare-and-ridership pair sits on the frontier. A claim that is not internally consistent — for example, the ridership of one operating regime combined with the fare of another — does not balance the ledger; it only appears to.

Subsidy and LVC together close the gap. Public subsidy is the dominant gap-closer in every operational HSR network in the world: every HSR system except the four highest-density Japanese and Chinese trunks operates with a structural annual operating subsidy on top of capital service support, and even those four required the full capital outlay from public funding. Land value capture is the only large-scale supplementary mechanism with empirical track record — HS1 in the United Kingdom, Crossrail's Business Rate Supplement, Hong Kong's MTR Rail+Property model, and several Japanese private-railway joint development arrangements.

Across these comparators LVC typically funds five to fifteen per cent of capital service requirements; the remainder closes through public subsidy. ALTO has not published a business case. The Government's written reply to Q-923 (April 22, 2026)[1] confirms that the analysis to date is preliminary and that the High-Speed Rail initiative's business case "will continue to be updated alongside its development" through co-development. The materials ALTO has published — the *Fast Forward* explanatory document, the consultation materials, the FAQ, and the CEO's parliamentary and public testimony — disclose no land value capture mechanism. The forecast 60,000 to 63,000 new residential units around stations is invoked as a downstream property-tax benefit accruing to municipalities, not as a financing source for the corridor. Bill C-15 (the High-Speed Rail Network Act) provides streamlined expropriation and right-of-first-refusal authority but no betterment levy, tax-increment financing district, special assessment district, joint development framework, or air-rights regime; the Senior Director, Commercial and First Nations Financial Participation role addresses Indigenous equity in Alto itself, not station-area land value capture. Under the current published scope, therefore, the entire gap closes through public subsidy.

No fare posture eliminates the bill. Raising fares reduces ridership along the air-rail and road-rail S-curves; lowering fares reduces revenue per rider. The frontier of achievable (subsidy, ridership, fare, revenue) combinations sits within a single-degree-of-freedom space — choosing any one variable determines the other three. No point on the frontier produces a balanced ledger without public subsidy at the corridor scale ALTO proposes. The minimum-subsidy posture (premium fare, low ridership) still requires approximately one billion dollars per year in operating subsidy at central capex; the welfare-efficient posture (fare parity with air) requires approximately two billion; the high-ridership posture (deep fare discount) requires approximately three and a half billion. The choice between operating regimes is therefore a political-economy choice about the distribution of fiscal cost between farebox and treasury, not a choice between subsidised and unsubsidised operation. The Government's written reply to Q-923 (April 22, 2026)[2] asserts that ALTO operations will be "financially self-sustaining, with revenues covering operations and maintenance costs and eliminating the need for ongoing operating subsidies." This claim is inconsistent with the empirical reference class of operational HSR networks worldwide and with the Stage 4 and Stage 5 outputs of this methodology: no point on the achievable modal-shift frontier, at any defensible operating posture, produces fare revenue sufficient to cover the $2.15 billion annual recurring cost identified in Stage 4. Verifying or rejecting the Government's self-sustaining-operations claim against the empirical reference class is one of the specific analytical tasks this methodology is designed to support.

Implications for the analytical pipeline. Stages 1 through 5 produce defensible values for the four independently-anchored ledger terms. Stage 6 (NPV / BCR) takes the populated ledger as the annual cash flow input and computes its present value over the analysis horizon; the NPV is

---

[1] Canada, Parliament, House of Commons, Q-923, written response to Order Paper question asked by Philip Lawrence, MP (Northumberland—Clarke), 45th Parliament, 1st Session, asked 5 March 2026, answered 22 April 2026; reply signed by Mike Kelloway, Parliamentary Secretary to the Minister of Transport.

[2] Q-923, op. cit. (22 April 2026).

the discounted lifetime sum of the annual gaps. Stage 7 (QA1 / QA2) reviews the project against the institutional criteria, with the populated ledger as part of the evidence base. The pipeline therefore does not produce an NPV in isolation — it produces a balanced annual fiscal ledger, whose present value is the NPV, whose subsidy term is the federal commitment the project actually implies, and whose ridership-and-fare pair is the operating posture the corridor would actually have to adopt.

The seven-stage analytical pipeline

The pipeline is designed around the principle that megaproject evaluation is not a single analysis but a sequence of analyses that must interlock to populate the annual fiscal ledger above. A capital cost estimate is meaningless without a complexity assessment that calibrates the reference class. An operating cost estimate is meaningless without an engineering ground-up build from asset inventory and service-level inputs. A ridership forecast is meaningless without a modal-shift framework that grounds the corridor's achievable demand in empirical S-curves and corridor demographics. A benefit-cost ratio is meaningless without a probabilistic cost distribution and a defensible ridership envelope. An institutional review verdict is meaningless without the underlying CBA and risk analysis. Each of the seven stages contributes a distinct output to the integrated assessment.

Stage	Component	What it produces	ALTO finding
1	**Engineering Complexity Index (ECI)**	Technical risk score 0–100, against international reference class of comparable projects	Peak severity 82/100; exposure-adjusted 73/100 (upper High band)
2	**Community Friction Index (CFI)**	Governance, consultation, and consent risk score 0–100	Current 43; projected mid-energy 65 across construction window
3	**Reference-class capital cost estimation**	Probabilistic capex distribution calibrated by ECI and CFI, against the empirical cost-overrun literature	$75B proponent-stated; $143B central (RCF); $264B P97.5
4	**Engineering ground-up O&M cost estimation**	Annual recurring cost built bottom-up from asset inventory and service-level inputs across three streams: infrastructure maintenance and renewals, operating categories, and fleet capital recapitalisation. Canadian climate, terrain, and traffic uplifts applied; international comparators used for cross-validation only	$2.15B/yr combined at 80 trains/day MID for 1,000 km corridor: maintenance $1.27B + operations $700M + fleet capital annuity $180M; cost-recovery ratio 0.80 at $0.20/pkm baseline; break-even at ≈12.5M pax/yr
5	**Reference-class modal-shift ridership analysis**	Independent ridership envelope, operating-regime structure, and subsidy frontier, derived from air-rail and road-rail S-curves applied to corridor demographics	Achievable frontier: 5–12M annual riders across the subsidy spectrum; welfare-efficient point Regime B (~8M riders at ~$2B annual

Stage	Component	What it produces	ALTO finding
			subsidy); ALTO 24M target outside the frontier
6	**NPV / Benefit-Cost Analysis**	BCR at three discount rates; iso-BCR analysis identifying the (ridership, fare) parameter space for break-even, using both the Stage 3 cost distribution and the Stage 5 ridership envelope	Central BCR ≈ 0.11; required fare $381–$1,596 at ALTO's 24M forecast; 24M itself unreachable on the Stage 5 frontier
7	**Institutional review (QA1 / QA2)**	Structured pass/fail assessment against Norwegian QA1 / QA2 criteria	Fails 5 of 8 QA1 criteria; forward-fails 5 of 6 QA2 criteria

Stages 1 and 2 (engineering and community) are inputs to Stage 3 (capital cost estimation). Stages 1 and 2 also inform Stage 5 (modal-shift ridership analysis), where the operational corridor experience reflected in the CFI bears on phase maturity and the consent environment in which the corridor would operate. Stage 4 (O&M cost estimation) is built from an independent engineering ground-up specification of corridor assets and service level — it is not derived as a fraction of the Stage 3 capex figure, which is a structural advantage over the standard "percentage of capex" benchmark approach. Stages 3, 4, and 5 are the **three parallel reference-class and engineering inputs** — capital cost, recurring cost, and ridership-and-fare — that together feed Stage 6 (NPV / BCR). Stages 1 through 6 are inputs to Stage 7 (institutional review). Stage 7 produces the recommendation for political decision-making. The pipeline does not stop at any one stage — a project that fails one stage will, in the typical case, fail the subsequent stages, and the structural finding is the cumulative result.

Stage 1 — Engineering Complexity Index (ECI)

The ECI scores a candidate project against a defined set of physical engineering dimensions, weighted to one hundred. The current rubric uses ten dimensions covering subgrade conditions, bedrock and foundation requirements, hydrology and drainage, climate exposure, topography and grade, seismic risk, ecological sensitivity, heritage and cultural overlay, corridor land-use friction, and urban engineering content. Each dimension is granularly scored against published anchor descriptions; the integrated score is then exposure-adjusted using a parameter that accounts for the proportion of the corridor exposed to each condition. The output is a single composite index on a 100-point scale, banded as Low (20–40), Moderate (41–60), High (61–80), or Extreme (81–100).

The ECI sits in the international reference class of comparable HSR projects, allowing direct comparison against thirteen completed or in-construction projects. ALTO's peak severity score of 82/100 places it at the boundary between High and Extreme; the exposure-adjusted score of 73/100 places it in the upper High band. The ECI is methodologically a generalisation of the engineering-complexity component of cost-overrun models in the academic literature (Flyvbjerg, Welde and Odeck), made transparent and replicable through the published rubric and scorecard.

Methodology and ALTO scoring: CAPEX Note 1 (engineering complexity rubric); CAPEX Note 2 (ALTO scorecard application).

Stage 2 — Community Friction Index (CFI)

The CFI scores a candidate project against a parallel set of governance, consultation, and consent dimensions, weighted to one hundred. The current rubric uses five components: public opposition intensity, legal challenges filed, forced route changes, environmental and NGO opposition, and political controversy. Each component scores 0–20, producing a composite 0–100 score. The CFI is also constructed as a trajectory — each component evolves through the project lifecycle, and the four-year forward projection through the construction window is part of the methodology, not a separate analysis.

ALTO's CFI assessment yields a current score of 43 and a projected mid-energy trajectory peak of 65 by the mid-construction window. Like the ECI, the CFI is calibrated against the international reference class of comparable projects: California HSR, HS2, Stuttgart 21, and others have CFI scores in the 60–85 range during their construction windows, and these are the projects whose cost overruns and schedule delays have been most extensively documented. The CFI is the governance analogue of the ECI — it captures the institutional and political risks that are systematically missing from purely technical cost-overrun models.

Methodology and ALTO scoring: CAPEX Note 3 (community friction and HSR cost, international comparative analysis).

Stage 3 — Reference-class capital cost estimation

The reference-class capital cost estimation methodology applies the empirical cost-overrun pattern documented across 258 transport infrastructure projects in 20 countries (Flyvbjerg, Skamris Holm and Buhl, 2003) to the proponent's point estimate, calibrated by the ECI and CFI scores from Stages 1 and 2. The output is a probability distribution rather than a single figure, reported at the central reference-class adjustment (RCF) and at upper-tail percentiles (P85 and P97.5). For ALTO, applied to the proponent-stated $75 billion concept-stage figure, the methodology produces a central estimate of $143 billion and a P97.5 estimate of $264 billion. The implied P85 figure — the parliamentary cost frame under Norwegian QA2 rules — is approximately $185 billion. The capex distribution feeds Stage 6 (NPV / BCR) directly; Stage 4 (O&M cost) is built independently from an engineering ground-up specification, and Stage 5 (ridership) supplies the complementary demand-side input.

Methodology and detail: CAPEX Note 4 (engineering complexity and community friction, joint predictors of HSR cost).

Stage 4 — Engineering ground-up O&M cost estimation

Stage 4 estimates the corridor's annual recurring cost from the bottom up. The methodology builds each cost line item from a defensible engineering specification — asset inventory, service level, unit costs, useful asset lives, and Canadian operating-condition adjustments — rather than from a percentage-of-capex benchmark calibrated against international comparators. The international comparator set (Spanish ADIF, California HSRA, UK Network Rail HS1, SNCF Réseau, Network Rail CP7, DB InfraGO) is used at the end of the build for cross-validation against the assembled bottom-up figure, not as the primary estimating method. This is a structural advantage over the standard "percentage of capex" operating cost approach: the ground-up estimate is independent of whatever capex figure the proponent adopts, and therefore independent of the optimism bias documented in the cost-overrun reference class.

The recurring annual cost decomposes into three streams with fundamentally different cost drivers, sensitivities to traffic, and underlying methodologies. Conflating them — as is common in summary business cases — hides the structure that determines whether the corridor can recover its costs at any achievable level of demand.

- **Infrastructure maintenance and renewals.** Built from a ten-asset-class decomposition of the corridor (plain line track, switches and crossings, overhead contact system, signalling, bridges, tunnels, drainage, fencing, stations, communications). For each asset class, annual maintenance is the sum of a fixed component (independent of traffic, driven by inspection and age-based degradation) and a variable component (driven by gross-tonne-km and a speed-fatigue term). Renewal capital expenditure on long-life components is annuitised over each asset's useful life using the standard capital recovery factor. Multiplicative Canadian uplifts are applied for climate (1.25), mixed traffic (1.00 for dedicated HSR), and terrain (1.10), combined factor 1.375.
- **Operating cost.** Built from an eight-category decomposition (train operations / traincrew; traction energy; rolling stock light and intermediate maintenance; station operations; network control / OCC; commercial; insurance; general and administrative overhead). Three of the eight scale with train-km, three are fixed, one scales with revenue, and one is a markup on direct costs. Each is calibrated against international unit-cost benchmarks (California HSRA 2020 Plan, SNCF Réseau, FRA HSIPR), with a separate Canadian climate uplift of 1.10 for operations.
- **Fleet capital.** Trainsets have useful lives of 25 to 35 years and require periodic replacement. The acquisition capex is annuitised over the assumed useful life at the social discount rate, producing a recurring annual fleet capital line. This is methodologically distinct from operating-stream rolling-stock light and intermediate maintenance, and it is the cost stream most often omitted from summary business cases.

For ALTO at MID baseline assumptions — 1,000 km dedicated double-track corridor, 300 km/h design speed, 80 trains per day, 450-seat trainsets at 65 per cent load factor, seven staffed stations, Canadian climate and terrain uplifts applied — the three streams produce combined recurring cost of **approximately $2.15 billion per year:** $1.27 billion for infrastructure maintenance and renewals (Note 1), $700 million for operating cost (Note 2), and $180 million for fleet capital annuity (30 trainsets at $70 million each, annuitised over 25 years at 7 per cent real). The LOW–MID–HIGH envelope on the maintenance stream alone is $1.08–$1.27–$1.52 billion per year; the full envelope on the combined figure is approximately $1.8–$2.15–$2.5 billion per year. Forty-year present value at 7 per cent real is approximately $28.6 billion (MID).

The combined cost has a clean fixed-plus-variable structure: **approximately $1.38 billion per year fixed, plus $9.58 million per year per train per day variable**. The fixed component is 77 per cent of infrastructure maintenance, 32 per cent of operations, and 100 per cent of fleet capital (fleet size is set by peak service requirement, not utilisation). This linearity allows Stage 4 to be evaluated at any service level — including the service levels implied by each of the Stage 5 modal-shift operating regimes — without re-running the ground-up build.

The substantive finding for ALTO is that the corridor's break-even service level — at which fare revenue exactly equals recurring annual cost — is approximately **117 trains per day, corresponding to approximately 12.5 million annual full-corridor-equivalent passengers** at a $0.20-per-pkm yield and a 65 per cent load factor. This figure sits at the upper edge of the Stage 5 modal-shift frontier ceiling (approximately 12 million annual riders at $5 billion annual subsidy). At the modal-shift welfare-efficient operating point (Regime B, 8 million annual riders), the corridor recovers approximately 80 per cent of recurring cost from farebox at the MID assumptions, leaving an annual deficit of approximately $440 million — and this is the deficit before any allowance for the capital service term in the annual fiscal ledger. The cost recovery ratio approaches but does not exceed 1.0 even at the upper modal-shift envelope; recurring cost recovery from farebox alone is structurally infeasible across the achievable operating space.

Cross-validation against the international comparator set confirms the figures. The MID figure of $1.27 million per route-kilometre per year for infrastructure maintenance is approximately 1.4 times the SNCF Réseau LGV figure (approximately $900 thousand per route-km per year), consistent with the 1.375 Canadian uplift applied to a European-baseline unit cost. Stripping the uplift recovers an underlying figure ($920 thousand per route-km per year) that sits at the top of the European HSR range — the appropriate position given Canadian labour rates and the absence of a domestic HSR supply chain. Operating cost at $24 per train-km sits in the mid-range of the international benchmark band ($20–$30 per train-km for dedicated HSR at mature service).

Methodology and detail: O&M Note 1 (infrastructure maintenance cost methodology and worked example); O&M Note 2 (operating cost methodology and worked example); O&M Note 3 (combined recurring cost, fleet capital recapitalisation, and break-even cost-recovery analysis).

Stage 5 — Reference-class modal-shift ridership analysis

Stage 5 produces an independent ridership envelope for the corridor by applying empirical modal-shift S-curves to corridor demographics and operating-posture choices, in place of accepting a proponent-stated headline figure as an analytical input. It is the ridership-side analogue of Stages 3 and 4: just as those stages derive probabilistic cost distributions from international cost reference classes rather than from the proponent's point estimates, Stage 5 derives a continuous (subsidy, fare, ridership, revenue, net public cost) frontier from international and Canadian modal-shift evidence rather than from the proponent's headline forecast.

The stage has four components. The air-rail S-curve relates the rail-to-air price ratio and the rail journey time to the share of air travellers captured by rail (Appendix J). The road-rail S-curve relates the rail-to-car generalised-time ratio and the relative price to the share of car travellers captured by rail (Appendix K). The demographic ridership envelope applies the two S-curves to the corridor's addressable population and trip pool, producing a 2035–2080 ridership envelope under three discrete operating regimes — A (heavy subsidy, deep fare discount to air), B (moderate subsidy at fare parity with air), and C (minimal subsidy with yield-managed pricing) (Appendix L). The subsidy frontier then extends the three discrete regimes to a continuous subsidy spectrum, producing the relationships between annual subsidy, annual ridership, annual fare revenue, and net public cost across the operating space (Appendix M).

The output of Stage 5 is a one-dimensional frontier through a four-variable space: each point on the frontier represents a defensible operating posture for the corridor, with the choice between points being a single-degree-of-freedom political-economy decision. The frontier supports several optimisation objectives — revenue maximisation, per-rider net-public-cost minimisation, total net-public-cost minimisation, fiscally-constrained ridership maximisation, and total welfare maximisation — each of which selects a different point. The first two converge on Regime B (approximately 8 million annual riders at approximately $2.0 billion annual operating subsidy, with peak fare revenue of approximately $1.29 billion and a marginal net-public-cost minimum of approximately $400 per added rider). Total-cost minimisation points toward Regime C or below; ridership maximisation points toward Regime A.

For ALTO, the structural finding is that the corridor's published 24-million-by-2055 ridership target **sits outside every point on the achievable frontier**. The frontier tops out at approximately 12 million annual riders at $5 billion annual subsidy — beyond which the corridor would need to push past the empirical modal-shift ceiling (approximately 40 per cent of the addressable market) under conditions none of the international modal-shift literature supports. The headline target is therefore not a defensible operating point under any of the candidate optimisation objectives. The fare and subsidy assumptions implicit in any operating regime that approaches the published target are not internally consistent with the modal-shift framework that the corridor's demand must come from. The Government's written reply to Q-923[3] cites additional ridership figures — 1.21 billion trips in the first forty years of full operations (averaging approximately 30 million annually) and 43 million annually by 2084 — that are not internally consistent with the 24-million headline target and that each remain outside the Stage 5 achievable frontier. The reply's $100 billion fare-revenue projection over the same forty-year window implies an average fare of approximately $83 per trip, a (fare, ridership) pair the modal-shift framework does not produce.

[3] Q-923, op. cit. (22 April 2026): "Class 5 capital costs are estimated at $60–90 billion (2024 Canadian dollars). Expected ridership is up to 1.21 billion trips in the first 40 years of full operations, with fare revenues projected to exceed $100 billion (2024 Canadian dollars) over the same period. Ridership is expected to grow to 43 million annually by 2084."

Stage 5 also produces a full-cost view in which the operating subsidy from the frontier is combined with capital service across the three Stage 3 capital cost scenarios and the Stage 4 engineering O&M figure. At the welfare-efficient Regime B operating point, the full annual federal cost ranges from approximately $4.3 billion (proponent-stated capex) to approximately $9.9 billion (P97.5 worst-case), with full-cost-per-rider of $540 to $1,240 — six to fourteen times the federal value-of-time per rider under standard cost-benefit parameters. This is the input to Stage 6's NPV / BCR.

Methodology and detail: MS Note 1 (air-rail modal-shift S-curve); MS Note 2 (road-rail modal-shift S-curve); MS Note 3 (ALTO HSR ridership envelope 2035–2080); MS Note 4 (subsidy frontier and optimisation).

Stage 6 — NPV and Benefit-Cost Analysis

Stage 6 takes the capex distribution from Stage 3, the engineering ground-up O&M figure from Stage 4, and the ridership envelope from Stage 5 and integrates them with an externality framework (modal-shift greenhouse gas benefits, embodied-carbon construction debit, time-savings, accident reduction, local environmental costs) and the Treasury Board Secretariat 8% real central social discount rate. The output is a Net Present Value (NPV) and Benefit-Cost Ratio (BCR) at the three federal sensitivity rates (3% / 5% / 8% TBS), with Monte Carlo treatment for parameter uncertainty. A complementary iso-BCR analysis identifies the parameter space (ridership, fare, capex, discount rate) within which break-even (BCR=1) is achievable. For ALTO, the central reference-class BCR is approximately 0.11; iso-BCR analysis demonstrates that no defensible combination of (capex, ridership, fare, rate) within the empirically supported parameter space produces break-even at the central or high capex scenarios.

The iso-BCR framing produces an intuitive rendering of the analytical finding: at ALTO's published 2055 ridership forecast of 24 million passengers per year, the required fare for break-even is $381 per trip at the proponent-stated capex, $818 at the reference-class central, and $1,596 at the high-cost scenario. The Stage 5 ridership analysis sharpens this finding by establishing that the 24-million figure is itself outside the achievable modal-shift frontier — so the iso-BCR break-even fares should be read as the answer to a counterfactual the corridor cannot in fact reach. **Break-even is unreachable from both directions: the fares required at the stated ridership are economically implausible, and the stated ridership is not deliverable on any operating posture the modal-shift frontier supports.** The layperson's test — what would the ticket cost have to be? — is grounded in the underlying NPV mathematics, and is now paired with the operating-regime test from Stage 5 — what is the achievable ridership at any defensible operating posture?

Methodology and results: NPV Note 1 (financial analysis, NPV and BCR); ALTO Iso-BCR Research Note; ALTO NPV Analysis v3 Excel model.

Stage 7 — Institutional review (QA1 / QA2)

Stage 7 applies Norway's Quality Assurance scheme — the QA1 (concept-stage) and QA2 (pre-funding) review gates established under the Norwegian Ministry of Finance in 2000 — as a structured pass/fail framework. QA1 evaluates a project against eight substantive criteria covering needs, goals, requirements, alternatives, cost-benefit analysis, cost estimation methodology, risk and uncertainty disclosure, and distributional alignment. QA2 evaluates the project at pre-funding stage against six criteria covering stochastic cost estimation at P85, reduction list, management challenge assessment, updated CBA with benefit realisation plan, project management base, and QA1 compliance. The framework draws on twenty-five years of Norwegian operating experience across approximately 160 reviewed projects, with documented improvements in budget compliance from pre-QA cost overruns of 59–183 per cent to post-QA on-budget delivery of approximately 75 per cent.

ALTO fails five of eight QA1 criteria outright (alternatives, CBA, cost estimate, risk disclosure, distributional alignment), partially satisfies two, and conditionally passes one. The addition of Stages 4 and 5 reinforces the QA1 cost-benefit-analysis and risk-disclosure findings: the proponent's published materials rely on a single 24-million ridership figure that the Stage 5 modal-shift analysis shows to be outside the achievable frontier, treats O&M and fleet capital as small line-items rather than as a multi-billion-dollar annual ledger commitment, and does not specify which operating posture (fare, subsidy, modal share) the corridor would actually adopt to deliver that figure. A QA2 forward assessment confirms the project would also fail at pre-funding review. The structural finding is that ALTO is a paradigmatic case of the project category Norway's scheme was created to filter.

Methodology and assessment: Appendix F — ALTO Norway QA Research Note (institutional design and empirical track record); Appendix G — ALTO QA Assessment (structured pass/fail of ALTO).

Connecting the pipeline to political decision-making

The seven-stage pipeline produces a recommendation rather than a decision. Under Norwegian institutional rules, the QA1 reviewer's recommendation goes to the Cabinet, which retains decision authority. The Cabinet may accept the recommendation, request revised documentation, or reject the recommendation and proceed notwithstanding. The political optionality is preserved by the institutional separation of analysis from decision — but the analytical record is on the public record, reviewable by the equivalent of the Auditor General and parliamentary standing committees, and the political cost of proceeding against an unfavourable QA1 finding is non-trivial.

In the Canadian context, the institutional analogue would be the Cabinet receiving the analytical pipeline output as input to the corridor selection decision, with the Parliamentary Budget Officer or a comparable independent analytical body as the natural reviewer of the methodology and findings. The CRI methodology is structured to be replicable by such a reviewer: each stage publishes its rubric, its scoring, its data, and its supporting documentation. A reviewer disagreeing with any individual finding is invited to re-run the analysis under their own preferred assumptions. The methodology is contestable, but it is not opaque.

Possible methodological extensions

Three extensions to the current seven-stage pipeline are worth flagging as candidates for further development. Each addresses a category of analysis that the existing pipeline touches but does not yet treat as a standalone stage.

Lifecycle environmental assessment

The current pipeline includes embodied-carbon construction emissions as a fixed externality term in Stage 6 (NPV), drawing on the CRI's Appendix H lifecycle CO_2 work. A more comprehensive lifecycle environmental assessment — covering construction, operation, maintenance, and end-of-life — across air quality, water resources, biodiversity, and land use, with uncertainty bounds, would warrant a parallel stage feeding both the cost distribution and the externality framework. The Frontenac Arch ecological sensitivity, the Leda clay groundwater regime, and the cumulative impact of corridor land take across agricultural Eastern Ontario and Quebec are specific items that an expanded lifecycle assessment would treat formally.

Distributional and Indigenous consultation analysis

The current pipeline captures distributional issues partially through CFI and through QA1 criterion 8, but does not produce a standalone distributional analysis with quantified incidence by region, income decile, community type, and Indigenous nation. A formal additional stage, modelled on UK HM Treasury distributional impact assessment guidance and on UNDRIPA implementation requirements, would produce specific quantified findings on who benefits and who bears costs, treated as an output equal in standing to BCR.

Counterfactual / alternatives analysis

The current pipeline applies Stages 1 through 7 to the proponent's preferred concept. A complete evaluation requires the same pipeline applied to each conceptually different alternative — in the ALTO case, to HFR (the JPO 2021 high-frequency rail concept) and HPR (the high-performance rail alternative the Initiative has developed). The HPR Strategy work is the start of this; running the full pipeline on HPR with the same rigor as on ALTO — including, crucially, the Stage 5 modal-shift analysis under HPR's operating parameters — would produce a directly comparable BCR, ECI, CFI, ridership envelope, and QA verdict, completing the substantive content of QA1 criterion 4.

Suitability for review by the Parliamentary Budget Officer

The Parliamentary Budget Officer (PBO) has the explicit statutory mandate to provide independent analysis to Parliament on matters of fiscal and economic significance. PBO precedent on major federal projects includes independent cost analyses of the F-35 procurement, the Canadian Surface Combatant program, and the Trans Mountain pipeline expansion — in each case, accepting a methodology paper from a parliamentarian or committee, conducting independent verification, and reporting findings to Parliament. The CRI methodology paper, with its supporting documents as appendices, is structured to support exactly this institutional pathway.

The specific request that an MP could submit to the PBO, supported by this methodology paper, would be: *“Conduct an independent benefit-cost analysis of the ALTO high-speed rail project using the methodology summarised in this document, and report findings to Parliament before the federal corridor selection decision is finalised.”* The PBO has the analytical capacity, the institutional independence, and the parliamentary mandate to discharge such a request. The supporting CRI research notes listed in the Sources section at the end of this paper provide the underlying data, methodology, and worked-example findings the PBO's own analytical team would need as starting material.

The methodology is also available to inform PBO analysis of other major federal infrastructure projects, including those in the Canada Strong Fund pipeline. Each of the seven stages is generalisable beyond rail: the ECI rubric, the CFI rubric, the reference-class capital cost estimation methodology, the engineering ground-up O&M cost estimation methodology, the reference-class modal-shift ridership methodology (where the relevant demand-side framework can be substituted for non-passenger projects), the iso-BCR sensitivity framing, and the QA1/QA2 institutional review framework apply to any major federal investment of comparable scale and policy importance.

Sources — supporting CRI research notes

The supporting documents listed below are the research notes underpinning each stage of the analytical pipeline, and the cross-cutting notes that supply inputs to multiple stages. Each note is independently complete and replicable, with published methodology, data, and results. They are grouped here by pipeline stage.

Stage 1 — Engineering Complexity Index

- **CAPEX Note 1 — Engineering Complexity Rubric** — ten-dimension scoring framework for the worldwide database of high-speed rail projects; five descriptor levels (Minimal, Low, Moderate, High, Extreme) with weights summing to 100, producing both a Peak Severity and an Exposure-Adjusted composite on a 100-point scale with four banded interpretations.
- **CAPEX Note 2 — ALTO Engineering Complexity Scorecard** — application of the ten-dimension rubric to the proposed ALTO corridor, with dimension-by-dimension evidence and reference-class comparison against thirteen international HSR projects; produces ALTO Peak Severity 82/100 (Extreme band) and Exposure-Adjusted 73/100 (upper High band).

Stage 2 — Community Friction Index

- **CAPEX Note 3 — Community Friction and Cost in High-Speed Rail: An International Comparative Analysis** — international comparative analysis of 18 HSR projects relating Community Friction Index scores to cost per kilometre and cost overrun (R^2 = 0.80 and

0.66 respectively); includes the full five-component CFI methodology with 0–20 sub-scale anchors and ALTO's projected ~163% cost overrun at CFI 65.

Stage 3 — Reference-class capital cost estimation

- **CAPEX Note 4 — Engineering Complexity and Community Friction: Joint Predictors of High-Speed Rail Cost** — multivariate reference-class regression of log cost per kilometre on the Engineering Complexity Index and Community Friction Index across 16 international HSR projects (adj. $R^2 = 0.88$); produces the $142 M/km central prediction and $76–$264 M/km 95% prediction interval for the ALTO corridor, equivalent to $142 billion central capex for a 1,000 km line.

Stage 4 — Engineering ground-up O&M cost estimation

- **O&M Note 1 — Infrastructure Maintenance Costs for HSR** — ten-asset-class decomposition methodology with fixed-and-variable unit costs, capital recovery factor for renewals, and Canadian climate, mixed-traffic, and terrain uplifts; worked example produces $1.27B/yr at MID for a 1,000 km corridor.
- **O&M Note 2 — Operating Costs for HSR** — eight-category decomposition (traincrew, traction energy, rolling stock light/intermediate maintenance, station operations, network control, commercial, insurance, G&A overhead) with international unit-cost calibration; worked example produces $700M/yr at MID with $24/train-km.
- **O&M Note 3 — Combined Cost Recovery for ALTO HSR** — synthesis of Notes 1 and 2 with the addition of fleet capital recapitalisation as the third recurring-cost pillar; combined $2.15B/yr at MID, fixed-plus-variable form $1.38B + $9.58M per train-per-day, break-even at ≈12.5M annual passengers.

Stage 5 — Reference-class modal-shift ridership analysis

- **MS Note 1 — Air-rail modal-shift S-curve** — relating the rail-to-air price ratio and rail journey time to the share of air travellers captured by rail, calibrated against international HSR experience.
- **MS Note 2 — Road-rail modal-shift S-curve** — relating the rail-to-car generalised-time ratio and relative price to the share of car travellers captured by rail, calibrated against North American conditions.
- **MS Note 3 — ALTO HSR ridership envelope 2035–2080** — application of the air-rail and road-rail S-curves to corridor demographics under three operating regimes (A: heavy subsidy/deep discount; B: parity with air; C: minimal subsidy/premium yield).
- **MS Note 4 — Subsidy frontier and optimisation** — extension of the three discrete regimes to a continuous subsidy frontier across the four-variable (subsidy, ridership, fare, revenue) space, with the welfare-efficient operating point, revenue peak, and full-cost view across the Stage 3 capex scenarios; ALTO 24M target shown to sit outside the achievable frontier.
- **ALTO Modal-Shift Deck** — briefing deck integrating MS Notes 1–4 with the operating-regime framework and the subsidy frontier visualisations.

Land value capture (revenue-side ledger term)

- **LVC Note 1 — Land Value Capture: assessing the $12 billion claim in the McGill TRAM financial model** — analysis of the McGill TRAM December 2025 financial model's $12B LVC line; finds the figure is a reverse-engineered planning placeholder rather than a forecast, with a defensible figure for the present value of plausible station-area capture in the low single billions.

Stage 6 — NPV and Benefit-Cost Analysis

- **NPV Note 1 — Financial Analysis, NPV and BCR** — full NPV methodology, results across three discount rates (3%, 5%, 8% TBS Central), Monte Carlo treatment, JPO 2021 HFR comparison, and the central BCR finding of approximately 0.11.
- **ALTO Iso-BCR Research Note** — iso-BCR=1 analysis identifying the (ridership, fare) parameter space for break-even at three discount rates; required fares of $381 / $818 / $1,596 per trip at the proponent's 24M ridership target across the three capex scenarios.
- **ALTO NPV Analysis v3 (Excel model)** — NPV and BCR sheets with iso-BCR overlays, discount-rate comparison, Monte Carlo, and parameter overrides for replication.

Stage 7 — Institutional review (QA1 / QA2)

- **ALTO Norway QA Research Note** — institutional design of the Norwegian QA1/QA2 scheme, twenty-five-year empirical track record, and the policy case for a Canadian analogue.
- **ALTO QA Assessment** — structured pass/fail assessment of ALTO against the eight Norwegian QA1 criteria and the six QA2 forward-assessment criteria.

Supporting research notes

- **Lifecycle CO_2 work** — construction-phase embodied carbon central estimate of 14.69 Mt, range 6.87–29.60 Mt, with cold-climate construction premium and Leda clay treatment penalty; supplies the embodied-carbon externality term in Stage 6.
- **HPR Strategy** — detailed framework for High Performance Rail combining 200 km/h electrified passenger service along Highway 401 with freight displacement; supplies the alternative concept for QA1 alternatives analysis (criterion 4) and the candidate for parallel pipeline application identified in the counterfactual extension.