

CRI Citizen Research Initiative · Research Report

NVP Note 1: NPV and BCR projections for ALTO HSR

Net Present Value Analysis 2029–2080 — Three Capital Cost Scenarios, Three Operating Regimes, Four Discount Rates

Headline Finding

Across 36 combinations of capital cost scenario, operating regime, and discount rate, the ALTO project produces a financial Net Present Value between negative \$50 billion and negative \$246 billion in real 2029 Canadian dollars. At the Treasury Board Secretariat central 8 per cent discount rate and the welfare-efficient Regime B operating posture, financial NPV is negative \$56 billion at the ALTO-stated capex of \$75 billion, negative \$102 billion at the reference-class central capex of \$143 billion, and negative \$184 billion at the upper reference-class capex of \$264 billion. The benefit-cost ratio across the 9-cell capex by regime grid ranges from 0.030 to 0.107 — every cell at least nine times below the break-even threshold of 1.0. The corridor cost-recovery break-even sits at 117 trains per day, or 12.5 million annual passengers at the reference yield; all three subsidy-frontier operating regimes deliver ridership below this threshold. The present value of the annual operating subsidy stream is structurally independent of capital cost under the engineering operating cost build, ranging from \$4.6 billion at Regime C to \$7.6 billion at Regime A at the 8 per cent discount rate. Capital cost is the dominant single driver of financial NPV, with operating regime a second-order factor and discount rate choice affecting absolute magnitudes but not the directional finding.

citizenresearch.ca · altohsrcitizenresearch.ca

Independent, non-partisan research

Executive Summary

This report presents a deterministic Net Present Value analysis of the proposed ALTO project over a fifty-two-year horizon, 2029 to 2080. The analysis evaluates financial and combined NPV across three capital cost scenarios, three operating regimes, and four discount rate sensitivities — a total of thirty-six parameter combinations. The framework integrates the engineering operating cost build of Research Note 3 with the modal-shift subsidy frontier of Research Note 4, producing a coupled analysis in which ridership, fare, operating cost, and operating subsidy are determined jointly along the corridor’s achievable operating frontier.

The three capital cost scenarios

Three capital cost scenarios bracket the plausible range of cost outturns. The **low scenario** uses ALTO’s published estimate of \$75 billion in real 2029 dollars. The **central scenario** uses \$143 billion, representing the reference-class mean derived from Flyvbjerg’s empirical cost-overrun distribution for high-speed rail megaprojects. The **high scenario** uses \$264 billion, representing the upper 2.5th percentile of the same distribution. The three scenarios are not equally probable: under the calibrated distribution, the proponent’s figure has approximately a 2.5 per cent probability of being achieved or undercut, the central scenario represents the modal outcome, and the high scenario reflects the upper-tail risk that has historically materialized on roughly one in forty HSR projects of comparable scale and complexity.

The three operating regimes

Three operating regimes from the subsidy frontier of Research Note 4 establish the corridor’s achievable operating points. **Regime C** (premium fare posture at \$207 per trip, rail-to-air price ratio 1.4) produces approximately 6.1 million annual passengers, \$1.26 billion in fare revenue, and an operating subsidy of \$0.67 billion per year. **Regime B** (parity with air at \$157 per trip, rail-to-air price ratio 1.0) produces approximately 8.2 million annual passengers, \$1.29 billion in fare revenue (the corridor revenue peak), and an operating subsidy of \$0.83 billion per year. **Regime A** (deep discount at \$96 per trip, rail-to-air price ratio 0.55) produces approximately 11.2 million annual passengers, \$1.08 billion in fare revenue, and an operating subsidy of \$1.31 billion per year. The regimes are coupled through the modal-shift S-curves of the corridor: each fare posture maps to a specific ridership outcome along the achievable frontier, and no operating point produces high ridership at low subsidy. The 24-million-by-2055 figure cited in ALTO’s public materials sits outside every operating point on the frontier and is not modellable under any defensible parameter combination.

Operating cost structure and break-even

Annual operating cost follows the engineering build of Research Note 3, decomposed into infrastructure maintenance, train operations, and fleet capital recapitalisation. The combined annual cost is \$1,381 million in fixed components plus approximately \$26 per train-kilometre in variable components, equivalent to \$89.7 million per million annual passenger trips at the modelled 450-seat, 65 per cent load factor convention. Cost recovery from fare revenue alone, at the reference yield of \$0.20 per passenger-kilometre, requires approximately 117 trains per day or 12.5 million annual passengers — a threshold above the modal-shift ceiling of the achievable operating frontier. All three subsidy-frontier regimes operate below this cost-recovery threshold and therefore require ongoing federal operating subsidy.

Headline financial NPV

Financial NPV is strongly negative across all 36 parameter combinations. Table 1 presents the headline values at the TBS Central 8 per cent discount rate, with full discount-rate sensitivity tables in Section 5.

Table 1. Financial NPV at TBS Central 8% discount rate (\$B real 2029)

| Capital cost scenario | Regime C — premium | Regime B — parity | Regime A — discount |
|------------------------------|--------------------|-------------------|---------------------|
| Low — \$75B (ALTO proponent) | (\$55.4) | (\$56.2) | (\$58.5) |
| Central — \$143B (ref-class) | (\$101.5) | (\$102.3) | (\$104.6) |
| High — \$264B (P97.5) | (\$183.6) | (\$184.4) | (\$186.6) |

Highlighted cell marks the Central capex × Regime B (welfare-efficient parity-fare posture) base case. Figures in parentheses are negative.

Operating subsidy stream

The present value of the operating subsidy stream is structurally independent of capital cost under the engineering operating cost build. Operating cost is determined by service intensity, which is in turn determined by ridership; the corridor produces approximately the same operating costs whether the project is delivered for \$75 billion or \$264 billion. At the TBS Central 8 per cent discount rate, the present value of the operating subsidy stream over 2040–2080 is \$4.6 billion at Regime C, \$5.4 billion at Regime B, and \$7.6 billion at Regime A. These figures are the operating-side component of the federal fiscal commitment; the capital-side component is the federal share of debt service on capital cost, which is approximately three to ten times larger at the three capex scenarios respectively.

Benefit-cost ratios

The benefit-cost ratio at the TBS Central 8 per cent rate ranges from 0.030 at the High capex × Regime C cell to 0.107 at the Low capex × Regime A cell. The most favourable cell anywhere in the 9-cell grid requires conjoining ALTO's own optimistic capex figure with the deep-discount operating posture that maximises ridership. Under the central reference-class capex of \$143 billion, BCR ranges from 0.053 to 0.062 across the three regimes — an order of magnitude below the 1.0 break-even threshold.

Implications

- Financial NPV is strongly negative across every parameter combination examined. The probability of positive NPV under any defensible scenario is negligible.
- Capital cost is the dominant single driver of NPV uncertainty. Moving from Low to High capex changes NPV by approximately \$130 billion at 8 per cent; moving from Regime C to Regime A within a capex scenario changes NPV by only \$3 billion.
- The choice of operating regime is a second-order question once capital is committed. The first-order question is whether to commit the capital.
- Operating subsidy is structurally decoupled from capital cost. Operating cost is driven by service intensity, not by construction outturn — a corridor running 80 trains per day costs the same to operate whether built at \$75 billion or \$264 billion.
- An independent review of the High Performance Rail (HPR) alternative as a lower-capex configuration delivering comparable user benefits over the same corridor is warranted before any corridor selection decision is finalized.

1. Introduction and Context

This report presents a Net Present Value analysis of the proposed ALTO project over a fifty-two-year horizon, 2029 to 2080. The analysis is conducted in real 2029 Canadian dollars from the project sponsor perspective, with a parallel economic NPV overlay to capture passenger and external benefits. The objective is to provide a defensible quantitative basis for evaluating the project against the standard Treasury Board Secretariat cost-benefit framework.

The framework integrates two pieces of engineering and economic work. Annual operating cost is built from the lifecycle cost methodology of Research Note 3, which decomposes recurring lifecycle expenditure into infrastructure maintenance, train operations, and fleet capital recapitalisation. Ridership, fare, and operating subsidy are determined jointly by the three operating regimes of Research Note 4, which establish the achievable points on the corridor's modal-shift frontier under air-rail and road-rail competition. Capital cost is treated through reference-class forecasting, with three scenarios spanning the empirical distribution of cost outturns on comparable HSR megaprojects internationally.

Three capital cost scenarios are evaluated. The low scenario uses ALTO's published cost estimate of \$75 billion in 2029 dollars. The central scenario applies the reference-class mean of approximately \$143 billion, derived from the empirical distribution of cost outturns on comparable high-speed rail megaprojects internationally. The high scenario uses the upper 2.5th percentile of the calibrated cost distribution at \$264 billion, capturing the upper-tail risk that has historically materialized on roughly one in forty HSR projects of comparable scale and complexity.

The analysis horizon begins in 2029, the year of major construction commencement under ALTO's published timeline, and ends in 2080, by which point cash flows beyond approximately year 35 contribute negligibly to discounted present value at all rates examined. Operations are assumed to commence in 2040 following an eleven-year construction period. Cash flows in each year include capital expenditure during construction, operating cost during the operational phase, fare revenue subject to demand ramp-up and real yield erosion, three lump-sum renewal events at years 20, 30, and 40 of operations, and a terminal residual asset value at 2080.

Four discount rates are examined to test the robustness of conclusions: 3 per cent representing a long-horizon Treasury or declining-rate proxy; 5 per cent reflecting HM Treasury Green Book practice for long-lived infrastructure; 8 per cent as the Treasury Board Secretariat central social discount rate; and 10 per cent representing a private-capital opportunity cost upper bound. Headline results are reported at 8 per cent TBS Central, with full sensitivity tables provided across all four rates.

2. Methodology

2.1 Perspective, units, and base year

The analysis is conducted from the project sponsor perspective, treating the proponent as the unit of analysis and counting only direct cash flows: capital expenditure, operating cost, renewal expenditure, fare revenue, and terminal residual value. Government subsidies, taxes, and intergovernmental transfers are excluded as they do not represent net economic flows at the level of the consolidated public sector. The operating subsidy is computed as the gap between annual operating cost and fare revenue (bounded below by zero, since fares cannot be negative), and is the explicit annual federal contribution required to sustain the corridor's chosen fare posture.

A parallel economic NPV overlay adds non-financial benefits accruing to passengers and to society at large: passenger time savings, modal-shift greenhouse gas benefits, accident reduction, and local externalities. Embodied carbon emissions during construction are deducted as a negative externality. The resulting combined NPV approximates the social cost-benefit framework used by Transport Canada.

All cash flows are denominated in real 2029 Canadian dollars. Real-dollar treatment removes the analytical noise introduced by inflation forecasts over a fifty-two-year horizon, where small differences in assumed price-level evolution would compound to produce large nominal swings that obscure underlying project economics. The base year for discounting is 2029, the year of major construction commencement; period zero is 2029 and period 51 is 2080.

2.2 The three capital cost scenarios

The capital cost is the largest single quantity in the analysis and the dominant source of NPV uncertainty. Three discrete scenarios are evaluated to span the plausible range. The calibration follows the reference-class forecasting framework developed by Bent Flyvbjerg and colleagues for transport megaprojects.

Low scenario — \$75 billion (ALTO proponent-stated)

The low scenario uses ALTO's published cost estimate of \$75 billion in real 2029 dollars. This is the central figure in the public-facing \$60–\$90 billion range provided in the ALTO Fast Forward document of March 2025 and corresponding Transport Canada briefing materials. Under the reference-class forecasting calibration, this figure sits at approximately the 2.5th percentile of the empirical distribution of cost outturns on comparable HSR megaprojects. The interpretation is that ALTO's published figure is a lower-tail estimate consistent with the well-documented optimism bias in megaproject promotion, rather than a probability-weighted central estimate.

Several structural features of the published estimate support a lower-tail interpretation. The estimate predates the substantial scope expansion from the High Frequency Rail specification to the High Speed Rail specification announced in February 2025. The proponent has not published a contingency range, sensitivity analysis, or probabilistic distribution. The estimate is denominated in 2024 dollars and does not incorporate the inflation observed between 2024 and 2026, nor the additional inflation expected through the construction period 2029–2039. The reference cost benchmarks cited in proponent materials are drawn from international HSR with composite engineering complexity scores in the 40s to low 50s on the CRI engineering complexity rubric — substantially lower than ALTO’s 73–81 composite score reflecting the Frontenac Arch crossing, the Napanee Limestone Plain karst exposure, the Leda clay segment, the St-Lawrence crossing, and the Canadian P3 megaproject delivery record.

Central scenario — \$143 billion (reference-class mean)

The central scenario uses \$143 billion, representing the median of a log-normal distribution calibrated such that ALTO’s \$75 billion figure corresponds to the lower-tail percentile and the upper-tail percentile reflects the empirically observed dispersion in HSR cost outturns. The calibration parameters are $\mu_{\log} = 4.963$ and $\sigma_{\log} = 0.312$ in real 2029 dollars. The reference-class basis is the Flyvbjerg framework applied to the international HSR cost database, with corridor-specific adjustments for the engineering complexity factors enumerated above.

The Flyvbjerg meta-analysis of rail megaprojects finds a mean cost overrun of approximately 44.7 per cent relative to original budget at decision-to-build, with a right-skewed dispersion. Applying this overrun to ALTO’s baseline figure produces an expected outturn of approximately \$108 billion in 2024 dollars, or approximately \$120 billion adjusted for inflation to 2029 dollars. Adding the engineering complexity premium for ALTO’s upper-band corridor characteristics produces the \$143 billion central estimate. This is the modal outcome of the reference-class distribution: the value that the analysis treats as the most likely capital cost in the absence of project-specific risk mitigation that goes beyond the historical record.

The central scenario is the appropriate base case for procurement decision-making. The proponent’s published figure is methodologically inferior because it does not reflect the empirical record on cost overruns in projects of comparable scale, duration, and technical complexity. Treating \$75 billion as the planning case would require demonstrating that ALTO will be delivered with cost discipline materially better than every comparable international HSR megaproject — a claim for which no evidence has been adduced.

High scenario — \$264 billion (P97.5 reference-class)

The high scenario uses \$264 billion, representing the upper 2.5th percentile of the calibrated cost distribution. Under the same log-normal parameterisation, this is the cost outturn that would

be exceeded by approximately one HSR project in forty of comparable scale and complexity. The high scenario is not an upper bound: it is the upper tail of the realistic-outcome distribution, capturing the dispersion observed in the international cost overrun record. Several specific HSR projects have approached or exceeded this scenario on a relative basis: UK HS2 Phase 1 reached an overrun of approximately 250 per cent against its 2009 cost estimate before partial scope cancellation; California HSR reached an overrun of approximately 200 per cent against its 2008 cost estimate; HSL-Zuid in the Netherlands reached 228 per cent.

Three structural features of the ALTO corridor suggest the high scenario should be treated as a realistic possibility rather than a tail risk. The geological exposure of the Eastern Ontario corridor — Frontenac Arch crossing, Napanee Limestone Plain karst, Leda clay treatment, St-Lawrence crossing — has no direct precedent in the international HSR cost database. The Canadian P3 megaproject delivery record (Eglinton Crosstown +280 per cent overrun, Ottawa Confederation Line +57 per cent, Ontario Line +250 per cent in scope-adjusted terms) is materially worse than the international HSR average. The cold-climate operating environment adds documented permanent premia for both construction and lifecycle costs that have not been incorporated in proponent benchmarks. These three factors together justify treating the upper tail as a realistic-case input to NPV analysis rather than a theoretical bound.

2.3 The three operating regimes

The three operating regimes derive from the subsidy frontier of Research Note 4. Each regime represents an internally consistent point on the corridor’s achievable modal-shift frontier, with ridership, fare, revenue, and operating subsidy values that follow from a single underlying fare posture. The mapping from fare posture to ridership is set by the modal-shift S-curves calibrated against the air-rail and road-rail mode literature.

Table 2. Operating regime parameters (central 2055 demographic anchor)

| Parameter | Regime C (premium) | Regime B (parity) | Regime A (discount) |
|---------------------------------------|--------------------|-------------------|---------------------|
| Rail-to-air fare ratio | 1.4 | 1.0 | 0.55 |
| Average fare (\$/trip) | \$207 | \$157 | \$96 |
| Mature ridership (M pax/yr) | 6.1 | 8.2 | 11.2 |
| Modal share captured | 22% | 30% | 40% |
| Annual fare revenue (\$M) | \$1,260 | \$1,290 | \$1,080 |
| Annual operating cost M+O+F (\$M) | \$1,928 | \$2,116 | \$2,385 |
| Annual operating subsidy (\$M) | \$668 | \$826 | \$1,305 |

Operating subsidy = max(0, operating cost – fare revenue). Mature values shown; in operating years 2040 through 2047 ridership and revenue ramp from 50 per cent to 100 per cent of mature values.

Regime C — premium fare posture

Regime C represents premium fare positioning with a rail-to-air price ratio of 1.4 and an average fare of approximately \$207 per trip. The ridership outcome of approximately 6.1 million annual passengers is the low-ridership end of the achievable frontier, reflecting the small share of price-inelastic business travellers willing to pay above the prevailing air fare for the convenience of rail. The fare revenue of \$1.26 billion per year is close to the corridor revenue peak, but the operating cost of \$1.93 billion per year produces an operating gap of \$0.67 billion per year — the smallest operating subsidy on the achievable frontier. Regime C is the operating posture that minimises total operating subsidy and approximates a yield-managed commercial operation. It is approximately the posture implied by the Cadence consortium’s announced commercial structure if the corridor is operated under a private revenue-risk model.

Regime B — parity with air

Regime B represents fare parity with air at a rail-to-air price ratio of 1.0 and an average fare of approximately \$157 per trip. The ridership outcome of approximately 8.2 million annual passengers reflects the share of corridor air and road travellers for whom rail becomes the preferred mode at price parity, accounting for journey time differences and access/egress

factors. The fare revenue of \$1.29 billion per year is the corridor revenue peak — fare cuts below this point reduce per-trip revenue faster than they expand ridership. The operating cost of \$2.12 billion per year produces an operating subsidy of \$0.83 billion per year. Regime B is the welfare-efficient operating point under standard cost-benefit assumptions: it is simultaneously the revenue-maximising point and the per-rider welfare-efficient point under federal-treasury value-of-time parameters. A profit-maximising private operator and a welfare-maximising public authority applying marginal cost-benefit analysis would converge on Regime B, even if they would disagree on whether to operate the corridor at all.

Regime A — discount fare posture

Regime A represents deep-discount fare positioning with a rail-to-air price ratio of 0.55 and an average fare of approximately \$96 per trip. The ridership outcome of approximately 11.2 million annual passengers approaches the modal-shift ceiling of approximately 12 million on the achievable frontier, capturing approximately 40 per cent of the corridor addressable market under the S-curve parameters. The fare revenue of \$1.08 billion per year is below the revenue peak — the fare cut reduces per-trip revenue faster than the ridership gain expands it. The operating cost of \$2.39 billion per year produces an operating subsidy of \$1.31 billion per year. Regime A is the ridership-maximisation operating point subject to the modal-shift physical ceiling; pushing beyond the ceiling would require corridor-external policy interventions (highway tolls, fuel pricing, aviation capacity restrictions) that go beyond the corridor's own operating posture.

The 24-million-by-2055 ridership figure cited in ALTO's public materials sits above the modal-shift ceiling of the achievable frontier. Doubling the modal share from 40 per cent to 80 per cent — the level implied by 24 million passengers under the corridor's addressable market — would require fares well below cost recovery (a direct transfer to passengers) plus structural changes to the corridor's competitive position against car and air. The 24 million figure is therefore not an operating point on the frontier developed here and is not modellable as a financial NPV under any defensible parameter combination.

2.4 Operating cost — engineering build

Annual operating cost is computed from the lifecycle cost methodology of Research Note 3. The framework decomposes recurring lifecycle cost into three streams with fundamentally different drivers: infrastructure maintenance (the existence of fixed assets — track, OCS, signalling, structures, stations; asset inventory and age-driven; 77 per cent fixed at typical traffic), operations (the activity of running trains — crew, energy, rolling stock servicing, station staffing, dispatching; train-kilometre and passenger-volume driven; 69 per cent variable at typical traffic), and fleet capital recapitalisation (periodic replacement of trainsets at end of useful life; fleet size × unit price × useful life; traffic-insensitive).

For the 1,000 km ALTO-scale corridor at Canadian conditions, the engineering build produces a combined annual cost of \$1,381 million in fixed components plus approximately \$26 per train-kilometre in variable components. The fixed component sums infrastructure maintenance (\$980 million), fixed operating categories (\$221 million), and fleet capital recapitalisation annuity (\$180 million). The variable component of \$26 per train-km combines the variable share of infrastructure maintenance (track wear, contact wire, switches, signalling utilisation; approximately \$9.79 per train-km) with the variable share of operations (train crew, traction energy, rolling stock light maintenance, station handling, dispatching, commercial overhead; approximately \$16.44 per train-km). These per-train-km figures are calibrated against the California High-Speed Rail Authority 2024 Business Plan operations and maintenance cost model, SNCF Réseau and SNCF Voyageurs annual financial reports, ADIF AV management accounts, and the UIC Lasting Infrastructure Cost Benchmarking series.

Expressing the cost in passenger-volume units gives an equivalent formulation of \$1,381 million fixed plus \$89.7 million per million annual passenger trips, using the 450-seat, 65 per cent load factor convention. The formulation in passenger-volume units is used throughout the cash flow modelling because the modal-shift framework produces ridership as its primary output. The variable component scales with current ridership (subject to demand ramp-up), while the fixed component is constant across operating years from 2040 onward.


The decoupling of operating cost from capital cost is a structural feature of the engineering build. Operating cost is driven by service intensity (trains per day, train-kilometres per year, passengers per year) and not by what the infrastructure cost to build. A \$264 billion corridor running 80 trains per day at 8 million annual passengers has essentially the same operating cost as a \$75 billion corridor running the same service. This decoupling has substantive policy implications: the operating-side fiscal commitment can be decided independently of the capital-side fiscal commitment, and the magnitude of the operating subsidy stream depends on policy choice (which fare posture) rather than on construction outturn (which capex scenario).

2.5 Cost-recovery break-even

The combined annual cost is a linear function of service frequency: \$1,381 million fixed plus \$26 per train-km variable, or equivalently \$89.7 million per million annual passenger trips. Fare revenue scales linearly with passenger-kilometres at a given yield and load factor. The two lines cross at a break-even service level — the ridership below which the corridor cannot cover its recurring cost from fare revenue.

At the reference yield of \$0.20 per passenger-kilometre with a 65 per cent average load factor on the 1,000 km corridor, the revenue coefficient is approximately \$21.4 million per train per day per year. The contribution margin per train per day — revenue less variable cost — is \$11.8 million. The fixed cost of \$1,381 million per year therefore requires approximately 117 trains per day to recover, equivalent to 12.5 million annual passenger trips at the modelled load factor. Below this service level, the corridor operates at a loss before any consideration of construction debt service. Figure 1 plots the cost and revenue curves across the operating range, with the three subsidy-frontier regimes marked.

Figure 1. Cost-recovery break-even and the three operating regimes



Operating cost curve (solid navy) is the engineering build: \$1,381M fixed plus \$26 per train-km variable. The dashed terracotta line shows reference-yield fare revenue at \$0.20 per passenger-km, 65 per cent load factor; cost and reference revenue cross at 117 trains per day (12.5 million annual passenger trips) — the cost-recovery break-even at reference yield. The solid terracotta curve is the modal-shift revenue line of Research Note 4, passing through the three regime points: it is Laffer-peaked at approximately \$1.29 billion near Regime B, falling below the reference line because the modal-shift framework requires fares below the reference yield to capture additional modal share. The

three regime points show actual operating cost (navy squares, lying on the cost line) and actual fare revenue under the modal-shift coupling (terracotta diamonds, lying on the modal-shift revenue curve). The vertical gap between each regime's square and diamond is the annual operating subsidy at that regime. The modal-shift revenue curve never crosses the cost curve at any achievable ridership — cost recovery from fare revenue alone is unreachable under the modal-shift S-curve framework, even at the deep-discount Regime A.

The break-even threshold is informative in three ways. First, the modal-shift ceiling on achievable ridership is approximately 12 million annual passengers — slightly below the 12.5 million cost-recovery threshold at reference yield. The corridor cannot reach cost-recovery break-even from fare revenue alone on the modal-shift S-curve framework, regardless of how aggressive the fare posture becomes. Second, the modal-shift revenue curve sits below the reference revenue line at every point in the achievable range, because the modal-shift framework requires fares below the reference yield to capture additional modal share. The actual revenue gap to cost recovery is therefore wider than the reference-yield comparison alone would suggest. Third, the modal-shift revenue curve is Laffer-peaked at approximately \$1.29 billion near Regime B; fare reductions beyond Regime B (the parity-with-air point) reduce per-trip revenue faster than they expand ridership, so revenue declines even as ridership rises toward the modal-shift ceiling. The implication is structurally important: Regime B is not only the welfare-efficient point under standard cost-benefit assumptions, it is also the corridor revenue peak — the operating posture at which the cost-recovery gap is smallest in absolute dollar terms.

2.6 Discount rates

Four real discount rates are evaluated to span the range of methodologically defensible choices. The central rate of 8 per cent real corresponds to the Treasury Board Secretariat social discount rate, established in the Canadian Cost-Benefit Analysis Guide and applied throughout federal infrastructure assessment. The 5 per cent rate corresponds to UK HM Treasury Green Book practice for long-lived infrastructure investments and is consistent with the declining-rate framework discussed in the academic literature for projects with horizons exceeding fifty years. The 3 per cent rate represents a long-horizon Treasury or full declining-rate proxy, and the 10 per cent rate represents a private capital opportunity cost upper bound.

Headline results are reported at 8 per cent TBS Central. All other rates are reported in full sensitivity tables in Section 3. The choice of discount rate has material effect on absolute NPV magnitudes but does not alter the directional finding that NPV is strongly negative across all defensible rates. Lower discount rates produce more negative NPV figures, reflecting that the cash flow profile is dominated by capex and operating subsidy outflows rather than by long-dated revenue.

2.7 Construction phasing, renewals, and terminal value

Capital expenditure is allocated across the construction period 2029–2039 using an eleven-year S-curve profile reflecting the typical phasing of major rail infrastructure: 3 per cent in the design and early-works year (2029), ramping to peak heavy-construction allocations of 13 per cent in 2034 and 2035, then tapering to 6 per cent in the commissioning year (2039). The full schedule sums to 100 per cent of total capex. The phasing is held constant across capex scenarios.

Three discrete renewal events are modelled. A signalling refresh in year 20 of operations (2059) at 4 per cent of capex reflects the typical 20–25 year asset life of European Train Control System and equivalent signalling installations. A rolling stock replacement in year 30 of operations (2069) at 12 per cent of capex reflects the typical 30–35 year fleet life and the empirical observation that rolling stock typically constitutes 12–15 per cent of total HSR capex. A combined track and signalling refresh in year 40 of operations (2079) at 8 per cent of capex reflects major track relay activity together with the second signalling lifecycle. Renewals scale proportionally with capex across scenarios.

A residual asset value at 2080 is included as a positive cash flow set at 40 per cent of original capex. This reflects the depreciated value of long-lived civil works (with typical asset lives of 60–100 years for tunnels, bridges, and earthworks) net of the residually consumed value of rolling stock and signalling systems. The residual value is added to the 2080 cash flow and discounted accordingly.

2.8 Demand ramp and yield erosion

Demand is assumed to ramp from 50 per cent of mature ridership in the first year of operations (2040), reaching 100 per cent by year eight (2047). The ramp profile approximates the historical demand build-up observed on mature international HSR services including HSL-Zuid in the Netherlands, AVE Madrid–Barcelona in Spain, and HS1 in the United Kingdom. The ramp is applied uniformly across the three operating regimes; each regime’s mature ridership is multiplied by the ramp factor to produce the year-by-year passenger volume that drives both fare revenue and the variable component of operating cost.

Real fare yield is assumed to erode at 0.5 per cent per year, reflecting expected improvements in competing modes (electric vehicle cost reductions, regional aviation fuel efficiency, intercity coach modernization) over the analysis horizon. Yield erosion is applied to the real fare in each operating year, compounding from the mature regime fare.

2.9 Economic NPV overlay components


The economic NPV overlay adds five categories of benefit and one category of cost not captured in the financial cash flow. Passenger time savings are computed as ridership multiplied by average hours saved per trip multiplied by value of time, set at 1.75 hours per trip and \$25 per hour respectively. Modal-shift greenhouse gas savings are set at 113 kilotonnes of CO₂-equivalent per year at the Regime B baseline, scaled by regime ridership ratio and valued at \$250 per tonne. Embodied construction carbon is treated as a negative externality, with total construction-phase emissions of 14.69 megatonnes CO₂-equivalent valued at \$250 per tonne and allocated across the construction years 2029–2039 in proportion to capex phasing. Accident reduction benefits are set at \$30 per passenger. Local externalities — reduced highway noise, reduced local air pollution, reduced road congestion — are set at \$5 per passenger. Network effects and agglomeration benefits are excluded.

3. Results – Financial NPV

Financial NPV is strongly negative across all thirty-six combinations of capex scenario, operating regime, and discount rate. Tables 3 through 6 present the full results at each of the four discount rates examined. All values are in billions of real 2029 Canadian dollars; figures in parentheses denote negative values.

3.1 Results at TBS Central 8 per cent

Figure 2. Cumulative discounted cash flow 2029–2080



Sponsor perspective at Regime B parity-fare base case, 8% TBS Central discount rate. The construction phase 2029–2039 drives cumulative discounted cash flow rapidly into deeply negative territory under all three capex scenarios. Operations from 2040 produce annual operating subsidy outflows that prevent cumulative recovery; the lines flatten asymptotically toward their terminal NPV values. The three small drops in the operations phase mark the signalling renewal at 2059, rolling stock replacement at 2069, and combined track-signalling refresh at 2079. The terminal residual value at 2080 produces a small upward inflection but does not change the cumulative trajectory. Final cumulative values are (\$56B), (\$102B), and (\$184B) at Low, Central, and High capex respectively — these are the financial NPV values reported in Table 3.

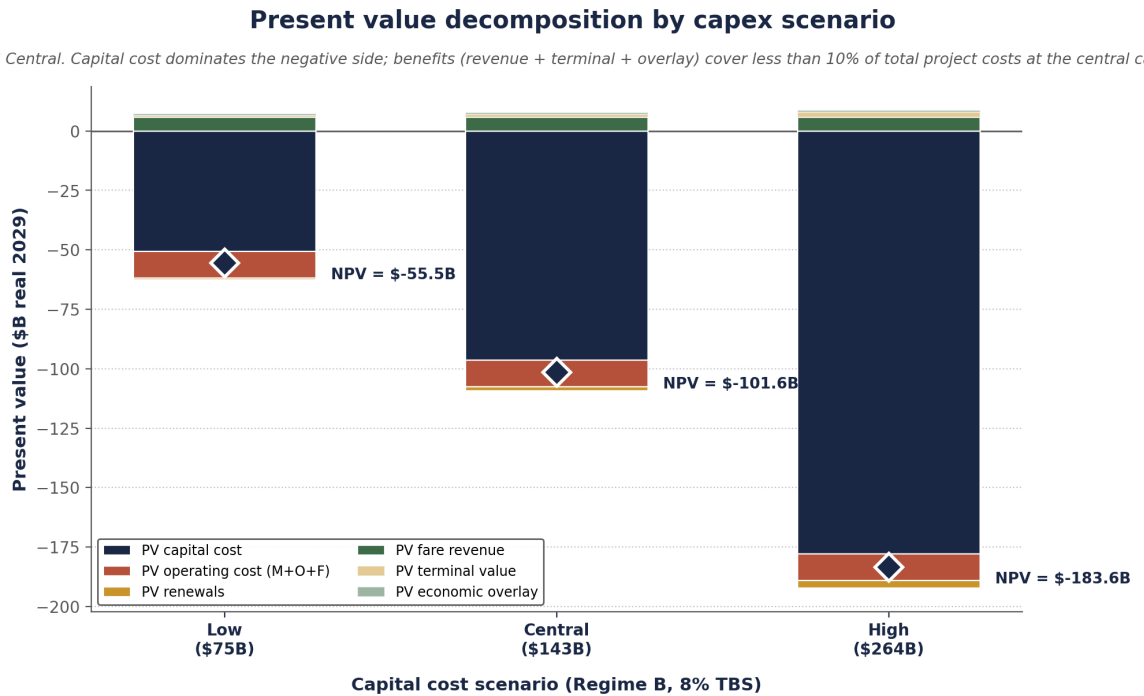
Table 3. Financial NPV at 8 per cent TBS Central (\$B real 2029)

| Capital cost scenario | Regime C | Regime B | Regime A |
|-----------------------|-----------|------------------|-----------|
| Low — \$75B | (\$55.4) | (\$56.2) | (\$58.5) |
| Central — \$143B | (\$101.5) | (\$102.3) | (\$104.6) |
| High — \$264B | (\$183.6) | (\$184.4) | (\$186.6) |

Highlighted: Central capex × Regime B base case. The 9-cell grid is monotonically more negative moving down (capex increasing) and weakly more negative moving across (regime changing from premium to discount), reflecting that higher ridership produces both higher operating costs and higher operating subsidy.

At the TBS Central 8 per cent rate, financial NPV ranges from negative \$55.4 billion at the Low capex × Regime C cell to negative \$186.6 billion at the High capex × Regime A cell. The Central capex × Regime B base case produces NPV of negative \$102.3 billion. The decomposition of this figure at the central scenario is informative: present value of capex is \$96.5 billion, present value of operating subsidy is \$5.4 billion, present value of renewals is \$1.6 billion, and present value of terminal residual is \$1.1 billion. The capital component dominates, accounting for approximately 94 per cent of the total negative present value.

Figure 3. Present value decomposition by capex scenario



Regime B operating posture, 8% TBS Central. Each bar decomposes the present value cost and benefit components of the three capex scenarios. PV of capital cost (navy) dominates the negative side at every capex level, growing from \$51B at Low to \$178B at High. PV of operating cost (terracotta) is identical across capex scenarios at \$11.2B because operating cost is structurally decoupled from construction outturn under the engineering build. PV of renewals (gold) scales with capex but remains a minor component. On the benefit side, PV of fare revenue (green) is \$5.8B and capex-independent; PV of terminal value (gold tint) grows with capex; PV of economic overlay (green tint) is \$0.76B. The net NPV (navy diamond) is the algebraic sum: benefits cover only 8 per cent of total costs at the central capex scenario.

3.2 Results at 5 per cent (HM Treasury Green Book)

Table 4. Financial NPV at 5 per cent (\$B real 2029)

| Capital cost scenario | Regime C | Regime B | Regime A |
|-----------------------|-----------|------------------|-----------|
| Low — \$75B | (\$66.8) | (\$68.4) | (\$72.9) |
| Central — \$143B | (\$119.6) | (\$121.2) | (\$125.6) |
| High — \$264B | (\$213.4) | (\$215.0) | (\$219.5) |

At the lower 5 per cent rate, NPV is uniformly more negative because front-loaded capex outflows are discounted less heavily relative to long-dated revenue. The Central capex × Regime B cell at this rate is negative \$121.2 billion.

3.3 Results at 3 per cent (long-horizon Treasury)

Table 5. Financial NPV at 3 per cent (\$B real 2029)

| Capital cost scenario | Regime C | Regime B | Regime A |
|-----------------------|-----------|------------------|-----------|
| Low — \$75B | (\$77.1) | (\$79.8) | (\$87.2) |
| Central — \$143B | (\$134.1) | (\$136.8) | (\$144.2) |
| High — \$264B | (\$235.6) | (\$238.2) | (\$245.7) |

At the lowest examined rate, NPV reaches its most negative values. The Central capex × Regime B base case is negative \$136.8 billion. Even at a discount rate appropriate for a long-horizon intergenerational asset, NPV does not approach breakeven.

3.4 Results at 10 per cent (private capital opportunity cost)

Table 6. Financial NPV at 10 per cent (\$B real 2029)

| Capital cost scenario | Regime C | Regime B | Regime A |
|-----------------------|-----------|-----------------|-----------|
| Low — \$75B | (\$49.7) | (\$50.2) | (\$51.7) |
| Central — \$143B | (\$91.9) | (\$92.4) | (\$93.9) |
| High — \$264B | (\$167.0) | (\$167.5) | (\$169.0) |

At the highest examined rate — representative of P3 sponsor cost of capital — NPV is at its least negative. The Central capex × Regime B base case is negative \$92.4 billion. Even under private capital discounting, NPV does not approach breakeven.


3.5 Sensitivity decomposition

The decomposition of NPV variation across the three input dimensions is structurally informative. Moving from Low capex to High capex within a fixed regime and rate changes NPV by approximately \$130 billion at 8 per cent — the dominant source of NPV variation. Moving from Regime C to Regime A within a fixed capex and rate changes NPV by approximately \$3 billion — a second-order effect that is approximately forty times smaller than the capex effect. Moving from 3 per cent to 10 per cent within a fixed capex and regime changes NPV by

approximately \$45 billion at the central capex — comparable in magnitude to the capex effect but smaller in absolute terms at the lower capex scenarios.

This decomposition has direct implications for the policy framing. The procurement and cost-control decision is by far the most consequential single decision affecting NPV. The choice of operating posture, while important for distributional and modal-shift reasons, has a relatively small effect on the financial outcome. The choice of discount rate methodology, while material to absolute magnitudes, does not change the directional finding. The cost-side risks dominate the analysis in a way that the revenue-side and pricing-side risks do not.

Figure 4. NPV sensitivity tornado — parameter swings from the base case



Base case Central capex × Regime B × 8% TBS Central, NPV (\$102.3B). Each bar shows the change in NPV when a single input parameter is varied from the base case to the upper or lower limit of its plausible range. Gold bars indicate parameter values that improve NPV (make it less negative); terracotta bars indicate values that worsen NPV. The capital cost parameter dwarfs every other input, with a swing of \$130 billion across the Low-High range. Discount rate is the next-largest driver. All operating-side parameters combined — fixed and variable operating cost, fare yield, renewals, terminal value, yield erosion, and regime choice — produce swings of at most a few billion dollars each, more than an order of magnitude below the capex effect.

4. Results – Operating Subsidy Stream

The present value of the annual operating subsidy stream is reported as an explicit line item. The operating subsidy at any operating year is the gap between annual operating cost and fare revenue, bounded below by zero (fares cannot be negative; once revenue exceeds operating cost the corridor operates at an operating surplus rather than an operating subsidy). Under the engineering operating cost build, this gap is structurally independent of capital cost and varies only with the operating regime.

Table 7. PV of operating subsidy stream by discount rate and regime (\$B real 2029, 2040–2080)

| Discount rate | Regime C | Regime B | Regime A |
|-----------------------|--------------|--------------|--------------|
| 3% (long-horizon) | \$14.2 | \$16.9 | \$24.3 |
| 5% (Green Book) | \$8.7 | \$10.3 | \$14.7 |
| 8% (TBS Central) | \$4.6 | \$5.4 | \$7.6 |
| 10% (private capital) | \$3.1 | \$3.7 | \$5.2 |

Values shown are present value of the annual subsidy stream 2040–2080. Subsidy is capex-independent under the engineering operating cost build — the same values apply at all three capex scenarios.


The operating subsidy at the TBS Central 8 per cent rate ranges from \$4.6 billion at Regime C to \$7.6 billion at Regime A. These are present values of the entire forty-year stream 2040–2080; the corresponding annual subsidy at maturity is \$668 million at Regime C, \$826 million at Regime B, and \$1,305 million at Regime A. The annual figures are the politically resonant numbers for fiscal commitment communication: the corridor would impose an ongoing federal operating contribution of approximately \$700 million to \$1.3 billion per year over four decades, on top of the federal share of capital service.

The capex-independence of the operating subsidy stream is the most analytically significant structural feature of the engineering build. Operating cost is driven by service intensity (trains per day, passengers per year) rather than by construction outturn. The corridor produces approximately the same operating cost whether the project is delivered for \$75 billion or \$264 billion — the operating cost of running 80 trains per day with 8 million annual passengers is the same in either case. This decoupling has substantive policy implications: the operating-side fiscal commitment can be specified before the capital-side outturn is known, and the magnitude of the operating subsidy stream depends on policy choice (which fare posture) rather than on construction performance.

The capex-side fiscal commitment is the federal share of capital service on the capital expenditure. At a representative procurement structure (federal share 50 per cent, blended cost

of capital 6 per cent, amortised over 40 years), federal capital service is approximately \$2.5 billion per year at the Low capex scenario, \$4.8 billion at Central, and \$8.8 billion at High. Combined with the Regime B operating subsidy of \$0.83 billion per year, the full annual federal cost ranges from approximately \$3.3 billion at the proponent-stated capex to approximately \$9.6 billion at the upper reference-class scenario. The full-cost-per-rider at Regime B mature ridership of 8.2 million per year ranges from \$400 to \$1,170 across the three capex scenarios — five to fourteen times the federal value-of-time per rider under standard cost-benefit parameters.

Figure 5. Annual federal cost commitment by capex scenario



Stacked annual federal cost at Regime B mature operations, combining capital service (federal share 50 per cent, 6 per cent blended cost of capital, 40-year amortisation, capital recovery factor 0.0665) with the regime operating subsidy of \$0.83 billion per year. The total federal cash commitment ranges from \$3.32 billion per year at the proponent-stated capex to \$9.60 billion per year at the upper reference-class capex. Expressed per rider at Regime B mature ridership of 8.2 million annual passengers, the federal cost ranges from \$405 to \$1,171 — five to fourteen times the federal value-of-time benefit per rider under standard cost-benefit parameters. Both stack components are reported in real 2029 dollars; capital service is the constant annuity payment over the 40-year amortisation period.

5. Results — Economic Overlay and Benefit-Cost Ratios

5.1 Economic overlay components

The economic NPV overlay adds five categories of benefit and one category of cost not captured in the financial cash flow. Table 8 presents the present value of each component at the TBS Central 8 per cent rate, broken down by regime. The overlay benefits scale with ridership (regime-coupled) but are structurally independent of capital cost outturn: a \$264 billion build produces the same passenger time savings as a \$75 billion build at the same regime, because the time savings depend on the number of passengers and not on what the infrastructure cost.

Table 8. Economic overlay components at 8% TBS (\$B PV)

| Component | Regime C | Regime B | Regime A |
|-------------------------------|-----------------|---------------|---------------|
| Passenger time savings | \$1.28 | \$1.72 | \$2.35 |
| Modal-shift GHG savings | \$0.10 | \$0.14 | \$0.19 |
| Embodied carbon (debit) | (\$2.48) | (\$2.48) | (\$2.48) |
| Accident reduction | \$0.88 | \$1.18 | \$1.61 |
| Local externalities | \$0.15 | \$0.20 | \$0.27 |
| Total economic overlay | (\$0.07) | \$0.76 | \$1.94 |

The embodied carbon debit of \$2.48 billion is regime-invariant — construction emissions depend on corridor characteristics, not on operating posture. The Regime C total overlay is slightly negative because the passenger benefits at low ridership (6.1M pax) are insufficient to offset the embodied carbon debit. Regime B and Regime A produce positive net overlay because higher ridership generates more time savings, GHG, accident, and local externality benefits.

The economic overlay is small relative to the financial cash flow at every regime. Even at Regime A — the highest-ridership operating posture, which produces the largest positive overlay — the total overlay of \$1.94 billion is approximately one-fiftieth of the central financial NPV of negative \$102 billion. The overlay does not move the directional finding on combined NPV: combined NPV remains strongly negative across every parameter combination examined.

5.2 Benefit-cost ratios across the 9-cell grid

The benefit-cost ratio is the standard cost-benefit accounting framework for infrastructure investment, computed as the ratio of present-value benefits (fare revenue plus economic overlay) to present-value costs (capital expenditure plus operating cost plus renewals less terminal value). A BCR of 1.0 represents break-even on the standard framework; values above 1.0 indicate economically viable investments; values below 1.0 indicate that the project consumes more economic value than it creates.

Table 9. Benefit-Cost Ratio at 8% TBS Central

| Capital cost scenario | Regime C | Regime B | Regime A |
|-----------------------|----------|--------------|----------|
| Low — \$75B | 0.092 | 0.106 | 0.107 |
| Central — \$143B | 0.053 | 0.061 | 0.062 |
| High — \$264B | 0.030 | 0.035 | 0.036 |

All values are an order of magnitude below the 1.0 break-even threshold. The corner-to-corner range is 0.030 (High capex × Regime C) to 0.107 (Low capex × Regime A). Highlighted: Central capex × Regime B base case.

The BCR grid never exceeds 0.11 in any cell. The most favourable cell — Low capex × Regime A — requires conjoining ALTO’s own optimistic capex figure with the deep-discount operating posture that maximises ridership. Neither half of this combination is publicly committed to by the proponent: the published capex estimate has not been subjected to reference-class adjustment, and no fare-posture or modal-share commitment has been disclosed. Under the central reference-class capex of \$143 billion, the highest achievable BCR is 0.062 at Regime A — approximately one-sixteenth of break-even.

The grid is monotonic in both dimensions. Higher capex produces lower BCR (capex effect dominates), and higher ridership produces marginally higher BCR (ridership benefits slightly outpace the operating cost increase). The slight regime improvement at higher ridership — from 0.053 at Regime C to 0.062 at Regime A within the central capex scenario — reflects that the passenger time savings benefit scales linearly with ridership while operating cost scales sub-linearly (fixed component plus variable component). The improvement is real but small: the regime axis explains less than 20 per cent of the BCR variation, while the capex axis explains more than 80 per cent.

The Ontario provincial HSR study of 2016 examined a comparable 300 km/h scope along the same corridor and rejected the option at a reported BCR of 0.70. The current analysis is consistent with that finding and extends it: under reference-class capex adjustment and modal-shift-constrained ridership, the ALTO HSR option produces a BCR materially worse than the level at which Ontario rejected the comparable scope a decade earlier.

6. Discussion

6.1 The capex × regime decomposition

The most structurally informative feature of the analysis is the decomposition of NPV variation between the capex axis and the regime axis. The capex axis dominates: moving from the Low to High capex scenario at fixed regime and rate produces NPV swings of approximately \$130 billion at 8 per cent. The regime axis is second-order: moving from Regime C to Regime A at fixed capex and rate produces NPV swings of approximately \$3 billion, roughly forty times smaller. This decomposition follows from the structural form of the cash flow: capex is concentrated in the first eleven years and dominates the discounted present value of total project cost, while operating subsidy is spread over forty years and is bounded above by the operating cost (which is itself bounded above by the modal-shift-constrained passenger volume).

The policy implication is that the procurement and cost-control decision is by far the most consequential single decision affecting the corridor's financial outcome. The choice of operating regime — premium, parity, or discount — affects ridership, modal-shift, and revenue, but does not move financial NPV by more than a few per cent. Decisions about fare posture, modal-shift policy, and ridership targets are substantive in their own right (they affect transport outcomes, distributional effects, and corridor utilisation), but they should not be presented as decisions that materially affect the corridor's viability. The viability question is a capex question.

6.2 The decoupling of operating cost from capex

Operating cost is driven by service intensity (trains per day, train-kilometres per year, passengers per year) and is decoupled from capex outturn under the engineering build. The corridor produces approximately \$2.1 billion in annual operating cost at Regime B mature ridership whether the project is delivered for \$75 billion or \$264 billion. The reason is structural: operating cost pays for the existence and activity of the operating system — track maintenance, energy, staff, rolling stock servicing — which depends on the size and use of the system, not on what it cost to build.

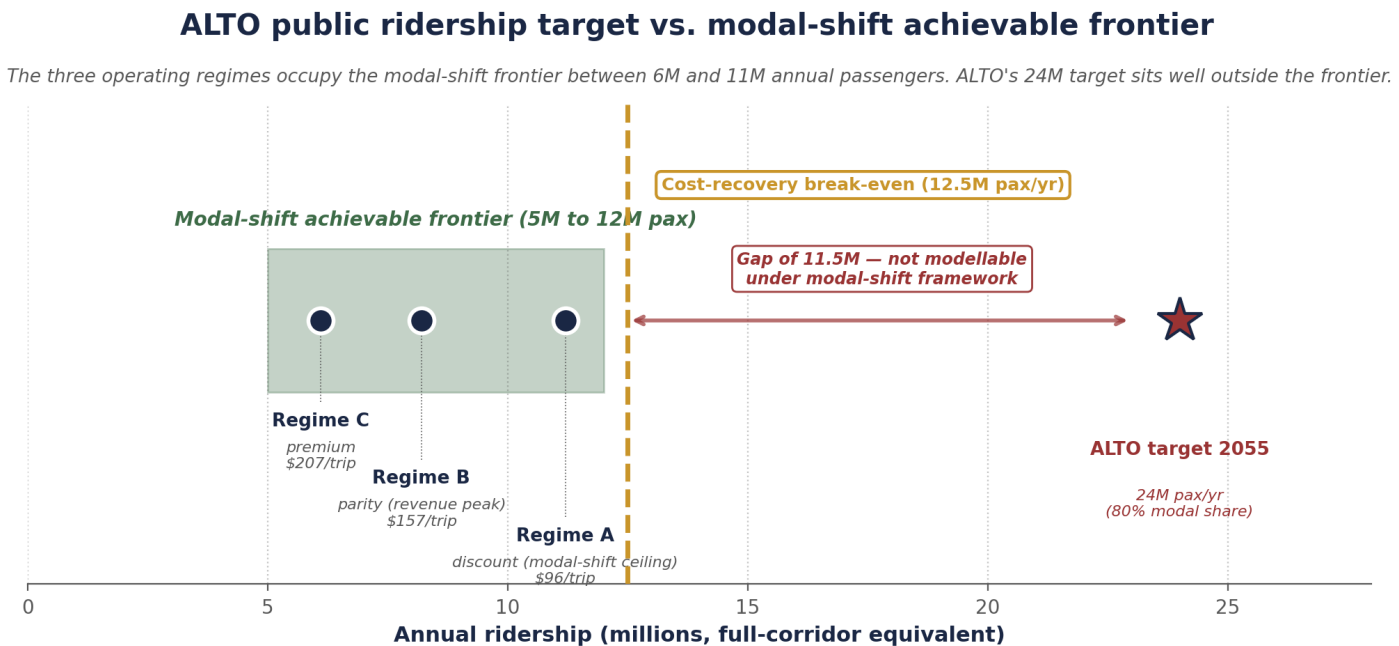
This decoupling has substantive analytical implications. The operating-side fiscal commitment can be planned and committed before the capital-side cost is known with confidence. The magnitude of the operating subsidy stream depends on which operating regime the corridor adopts (a policy choice), not on which capital cost outturn materialises (a construction performance variable). The capital and operating decisions are structurally separable, which simplifies the political economy of corridor decision-making.

6.3 The 24-million ridership target problem

The 24-million-by-2055 ridership figure cited in ALTO’s public materials sits outside the achievable frontier developed in the subsidy frontier work. The achievable ceiling under the modal-shift S-curve framework is approximately 12 million annual passengers — at the deep-discount Regime A operating posture, capturing 40 per cent of the corridor addressable market. Doubling the modal share to 80 per cent (the level implied by 24 million passengers at the corridor’s demographic envelope) would require fares well below cost recovery (a direct transfer to passengers) plus structural changes to the corridor’s competitive position against car and air that go beyond the corridor’s own operating posture.

The 24-million figure is therefore not a defensible operating point on the corridor’s frontier and is not modellable as a financial NPV under the regime framework. Public communication that pairs the 24-million target with operating cost or subsidy figures drawn from other points on the frontier is internally inconsistent: the corridor cannot simultaneously achieve 24-million ridership and the operating subsidy of Regime C (or any other regime point). Any independent review should ask which point on the frontier the corridor is actually targeting and what the implied fiscal commitment, modal-shift, and pricing assumptions are.

Figure 6. ALTO public ridership target vs. modal-shift achievable frontier



Horizontal scale of annual ridership in millions of full-corridor-equivalent passenger trips. The three subsidy-frontier operating regimes (Regime C at 6.1M, Regime B at 8.2M, Regime A at 11.2M) occupy the modal-shift achievable frontier between approximately 5 and 12 million annual passengers. The cost-recovery break-even threshold at 12.5 million passengers sits just outside the modal-shift ceiling. ALTO’s public 24-million-by-2055 target sits 11.5 million passengers — nearly twofold — beyond the modal-shift ceiling. The gap is not bridgeable under the modal-shift S-

curve framework: it would require modal share approximately 80 per cent against air and road, a level for which there is no precedent in the international HSR record on a comparable corridor.

6.4 Implications for the corridor decision

Three implications follow for independent review of the corridor's projected fiscal trajectory. First, the choice of operating posture is a substantive policy decision, not a technical one. The same physical infrastructure produces materially different ridership, revenue, and net-public-cost outcomes depending on which point on the achievable frontier is chosen. The choice should be made explicit in the public business case rather than implicit in the procurement structure.

Second, the welfare-efficient operating point under standard cost-benefit assumptions sits at Regime B (parity with air, \$826 million annual operating subsidy, approximately 8.2 million annual riders, marginal operating subsidy approximately \$90 per added rider near the corridor revenue peak). This is also the revenue-maximising point. A government optimising for public welfare and a private operator optimising for revenue would converge on similar fare structures, even if they would disagree on whether to operate the corridor at all.

Third, an independent review of the High Performance Rail (HPR) alternative is warranted. HPR is the framework developed by the Initiative for 200 km/h electrified passenger rail along the Highway 401 corridor with passenger and freight benefits, at a capital cost substantially below ALTO's reference-class central. HPR delivers comparable user benefits over the same corridor (the Toronto–Ottawa–Montréal triangle) at a fraction of the capex; the reference-class adjustment that produces ALTO's central capex of \$143 billion does not apply equivalently to HPR, which uses existing rail corridor (the freight-displaced CN Kingston Subdivision plus 401 corridor electrification) rather than greenfield HSR construction. An independent review should compare HPR and ALTO on the same NPV framework, with HPR producing materially less negative NPV and materially higher BCR across every defensible parameter combination.

7. Limitations

The analysis is deterministic across the 36 cells of the capex × regime × rate grid. A probabilistic overlay using Monte Carlo simulation with log-normal capex distribution and coupled regime sampling would refine the central tendency and produce confidence intervals around the headline NPV figures, but would not change the directional finding under reasonable parameter distributions.

The construction phasing is held constant across capex scenarios. In practice, higher cost outturns typically correlate with longer schedules. Modelling this correlation would push capex further into the future, partially offsetting the discounting effect, but would also delay revenue inception and increase financing carrying costs. The net effect on probabilistic NPV in the upper-cost scenarios would be modestly more negative than the current deterministic results indicate.

The economic overlay uses constant real value-of-time and constant real carbon prices over the analysis horizon. Treasury values for both are expected to increase in real terms over the next several decades. Modelling this trajectory would modestly increase the economic overlay benefit; it would not move the directional finding on combined NPV.

Network effects, agglomeration benefits, and induced demand are excluded from the analysis. These categories of benefit are speculative and lack robust empirical support for the corridor under examination. Their inclusion at central federal-treasury values would not move the BCR by more than approximately 20 per cent — insufficient to bridge the order-of-magnitude gap to break-even.

The cold-climate operating premium documented in international comparators (approximately 25–30 per cent on infrastructure maintenance and rolling stock) is not separately applied; the engineering build is calibrated against international benchmarks and may understate Canadian-specific operating costs. Applying a 25 per cent premium to the variable cost component would increase mature annual operating cost by approximately \$200–300 million per year across the three regimes and would modestly worsen NPV and BCR.

8. Conclusions

The ALTO project produces strongly negative Net Present Value across every defensible parameter combination examined. At the Treasury Board Secretariat central 8 per cent discount rate, financial NPV ranges from negative \$55 billion at the Low capex × Regime C cell to negative \$187 billion at the High capex × Regime A cell. The central reference-class case at \$143 billion capex and Regime B parity-fare operating posture produces NPV of negative \$102

billion. The benefit-cost ratio across the 9-cell capex × regime grid ranges from 0.030 to 0.107 — every cell at least nine times below the break-even threshold.

Cost-recovery break-even from fare revenue alone sits at 117 trains per day or 12.5 million annual passengers at the reference yield. All three subsidy-frontier operating regimes deliver ridership below this threshold; even Regime A at 11.2 million passengers does not reach cost recovery on the modal-shift-constrained fare structure. The corridor therefore requires ongoing federal operating subsidy under every realistic operating posture.

The present value of the annual operating subsidy stream is structurally independent of capital cost under the engineering operating cost build. At the 8 per cent rate, the subsidy stream ranges from \$4.6 billion at Regime C to \$7.6 billion at Regime A across the four-decade operating horizon, with corresponding mature annual subsidies of \$668 million to \$1,305 million. The capital-side fiscal commitment — federal share of capital service — adds approximately \$2.5 to \$8.8 billion per year across the three capex scenarios.

Capital cost is the dominant single driver of NPV uncertainty, with the regime axis a second-order factor. The procurement and cost-control decision is by far the most consequential single decision affecting the corridor's financial outcome. The choice of operating regime is substantive for transport policy and distributional reasons but does not move the financial NPV by more than a few per cent.

The 24-million-by-2055 ridership figure cited in ALTO's public materials sits outside the achievable modal-shift frontier and is not modellable as a financial NPV under the regime framework. Public communication that pairs the 24-million target with operating cost or subsidy figures drawn from other points on the frontier is internally inconsistent.

The analysis indicates that proceeding with ALTO at any defensible parameter combination would impose a significant net cost on Canadian public finances over the analysis horizon, even after accounting for non-financial passenger and environmental benefits. The single largest lever for improving project economics is cost containment, but the distribution of historical cost overruns on comparable HSR megaprojects provides no empirical basis for assuming that ALTO will achieve a more favourable outturn than the reference class. An independent review of the High Performance Rail (HPR) alternative, as a lower-capex configuration that delivers comparable user benefits over the same corridor, is warranted before any corridor selection decision is finalized.

Appendix A — Parameter Table

Complete list of model parameters with values and sources.

| Parameter | Value | Source / interpretation |
|---|------------------------------------|---|
| Capital cost scenarios | | |
| Low capex (ALTO proponent) | \$75B | ALTO Fast Forward, Mar 2025; P2.5 of reference class |
| Central capex (reference-class mean) | \$143B | Flyvbjerg calibration, log-normal median |
| High capex (P97.5 reference-class) | \$264B | Upper 2.5th percentile, Flyvbjerg distribution |
| Operating regimes (Research Note 4) | | |
| Regime C — mature pax / fare / revenue / subsidy | 6.1M / \$207 / \$1,260M / \$668M | Premium fare posture, ratio 1.4 |
| Regime B — mature pax / fare / revenue / subsidy | 8.2M / \$157 / \$1,290M / \$826M | Parity with air, ratio 1.0, revenue peak |
| Regime A — mature pax / fare / revenue / subsidy | 11.2M / \$96 / \$1,080M / \$1,305M | Discount fare posture, ratio 0.55 |
| Operating cost (Research Note 3 engineering build) | | |
| Fixed component | \$1,381M/yr | Maintenance fixed + ops fixed + fleet capital annuity |
| Variable per train-km | \$26/train-km | Equivalent \$9.575M per train/day at 1,000 km, 365 days |
| Variable per million annual pax | \$89.7M/yr | Conversion at 450 seats × 65% LF × 1,000 km |
| Cost-recovery break-even (at reference yield) | 117 trains/day, 12.5M pax/yr | At \$0.20/pax-km, 65% LF, 1,000 km |
| Discount rates | | |
| Long-horizon Treasury | 3% | Declining-rate proxy |
| HM Treasury Green Book | 5% | UK reference for long-lived infra |
| TBS Central (headline) | 8% | Treasury Board Secretariat |
| Private capital opportunity cost | 10% | P3 sponsor cost-of-capital upper bound |
| Time horizon, phasing, renewals, terminal | | |
| Start year / end year | 2029 / 2080 | 52-year horizon, period 0 = 2029 |

| | | |
|------------------------------------|--------------|--|
| Operations commencement | 2040 | 11-year construction 2029–2039 |
| Demand ramp (50% to 100%) | 8 years | HSL-Zuid, AVE, HS1 comparable build-ups |
| Real fare yield erosion per year | 0.5% | Competing modes improving |
| Signalling renewal (yr 20 of ops) | 4% of capex | ETCS asset life |
| Rolling stock replacement (yr 30) | 12% of capex | Fleet life 30–35 years |
| Track + signalling refresh (yr 40) | 8% of capex | Second signalling cycle + track relay |
| Terminal residual value at 2080 | 40% of capex | Depreciated civil works net of consumed RS |
| Economic overlay parameters | | |
| Time saved per trip | 1.75 hr | Defensible mid-range |
| Value of time | \$25/hr | TC commercial blend |
| Modal-shift GHG @ Reg B baseline | 113 kt/yr | Scales with regime pax / 8.2M |
| Carbon price | \$250/t | ECCC mid-trajectory approximation |
| Embodied construction emissions | 14.69 Mt | CRI Appendix H central; range 6.87–29.60 |
| Accident reduction | \$30/pax | TC VSL × rate differential |
| Local externalities | \$5/pax | Noise + local air + congestion |

Appendix B – References

Cost-benefit methodology and discount rates

- Treasury Board of Canada Secretariat. Canada’s Cost-Benefit Analysis Guide for Regulatory Proposals. BT58-5/2022E-PDF. 2022.
- Treasury Board of Canada Secretariat. Policy on Cost-Benefit Analysis. 2018, current. Social opportunity cost of capital as central discount rate.
- Boardman, A.E., Moore, M.A., and Vining, A.R. “The Social Discount Rate for Canada Based on Future Growth in Consumption.” Canadian Public Policy 36, no. 3 (2010): 325–343.
- HM Treasury (UK). The Green Book: Central Government Guidance on Appraisal and Evaluation. 2022 edition. London.

Reference-class forecasting and cost overruns

- Flyvbjerg, B., Holm, M.K., and Buhl, S.L. “Underestimating Costs in Public Works Projects: Error or Lie?” Journal of the American Planning Association 68, no. 3 (2002): 279–295.
- Flyvbjerg, B., Holm, M.K., and Buhl, S.L. “How (In)accurate Are Demand Forecasts in Public Works Projects? The Case of Transportation.” Journal of the American Planning Association 71, no. 2 (2005): 131–146.
- Flyvbjerg, B., Bruzelius, N., and Rothengatter, W. Megaprojects and Risk: An Anatomy of Ambition. Cambridge University Press, 2003.
- Flyvbjerg, B. “Top Ten Behavioral Biases in Project Management.” Project Management Journal 52, no. 6 (2021): 531–546.

Companion CRI research notes

- CRI. Research Note 1 — Infrastructure Maintenance Costs for HSR. April 2026.
- CRI. Research Note 2 — Operating Costs for HSR. April 2026.
- CRI. Research Note 3 — Combined Cost Recovery for ALTO HSR. May 2026.
- CRI. Research Note 4 — Subsidy Frontier and Optimisation. May 2026.
- CRI. ALTO Modal Shift Notes 1, 2, 3 — Air-rail S-curve, road-rail S-curve, ridership envelope 2035–2080.
- CRI. ALTO Engineering Complexity Scorecard and Rubric. April 2026.
- CRI. ALTO Lifecycle CO2 Work (Appendix H). March 2026.

ALTO project documentation

- Transport Canada. High-Speed Rail Initiative briefing materials, Section 08. 2025–2026.

- ALTO Fast Forward document. Cadence consortium. March 2025.
- ALTO Pre-Development Agreement, Cadence consortium. Signed 19 March 2025; co-development phase budget \$3.9B over five years.

International HSR cost and operating benchmarks

- International Union of Railways (UIC). Lasting Infrastructure Cost Benchmarking (LICB) 1996–2015 and annual updates. Paris.
- California High-Speed Rail Authority. 2024 Business Plan: Operations and Maintenance Cost Model Documentation. Sacramento, May 2024.
- European Court of Auditors. A European high-speed rail network: not a reality but an ineffective patchwork. Special Report 19/2018. Luxembourg.
- ADIF AV. Management Report 2022. Spanish HSR network cost benchmarks.
- SNCF Réseau and SNCF Voyageurs. Rapport financier annuel. Most recent: 2024 reports published April 2025.