CITIZEN RESEARCH INITIATIVE · MODAL SHIFT ANALSIS NOTE 4

Modal Shift Note 4:

Subsidy frontier and optimisation

A continuous-spectrum framework relating operating subsidy, fare revenue, ridership, and net public cost for the ALTO corridor under the modal-shift and ridership envelope developed in Notes 1, 2, and 3. The framework identifies the welfare-efficient and revenue-maximising operating points on the corridor's subsidy frontier.

BOTTOM LINE

High ridership and low subsidy are mutually exclusive on the ALTO corridor. The modal-shift framework developed in Notes 1 and 2, combined with the corridor demographics of Note 3, produces a fixed frontier of (subsidy, ridership) combinations: each point on the frontier reflects a different operating posture, and the choice between them is a single-degree-of-freedom decision. The corridor cannot simultaneously deliver Regime A ridership (11–12 million annual passengers) at Regime C subsidy levels ($0.5–1.5 billion per year). Any public communication implying otherwise is selecting figures from different points on the frontier and presenting them as one operating outcome.

Ridership rises concavely with subsidy from approximately 5 million at $0.3 billion per year to approximately 12 million at $5 billion — diminishing returns as the corridor approaches the modal-shift ceiling. Revenue is hump-shaped, peaking at approximately $1.29 billion at $1.9 billion subsidy and declining beyond that as fare cuts overwhelm ridership gains. The marginal net public cost per added rider has a U-shaped minimum at approximately $400 per rider near the Regime B operating point.

The optimal operating posture depends on the optimisation objective. Maximising fare revenue and minimising per-rider net public cost both point to Regime B. Minimising total net public cost points to Regime C. Maximising ridership subject to a fiscal cap points to Regime A. The choice between them is a political-economy decision about how to weight commercial yield, public welfare, ridership, and fiscal restraint. ALTO's published 24-million-by-2055 target sits outside every point on the frontier developed here and is incompatible with any defensible operating-regime choice.

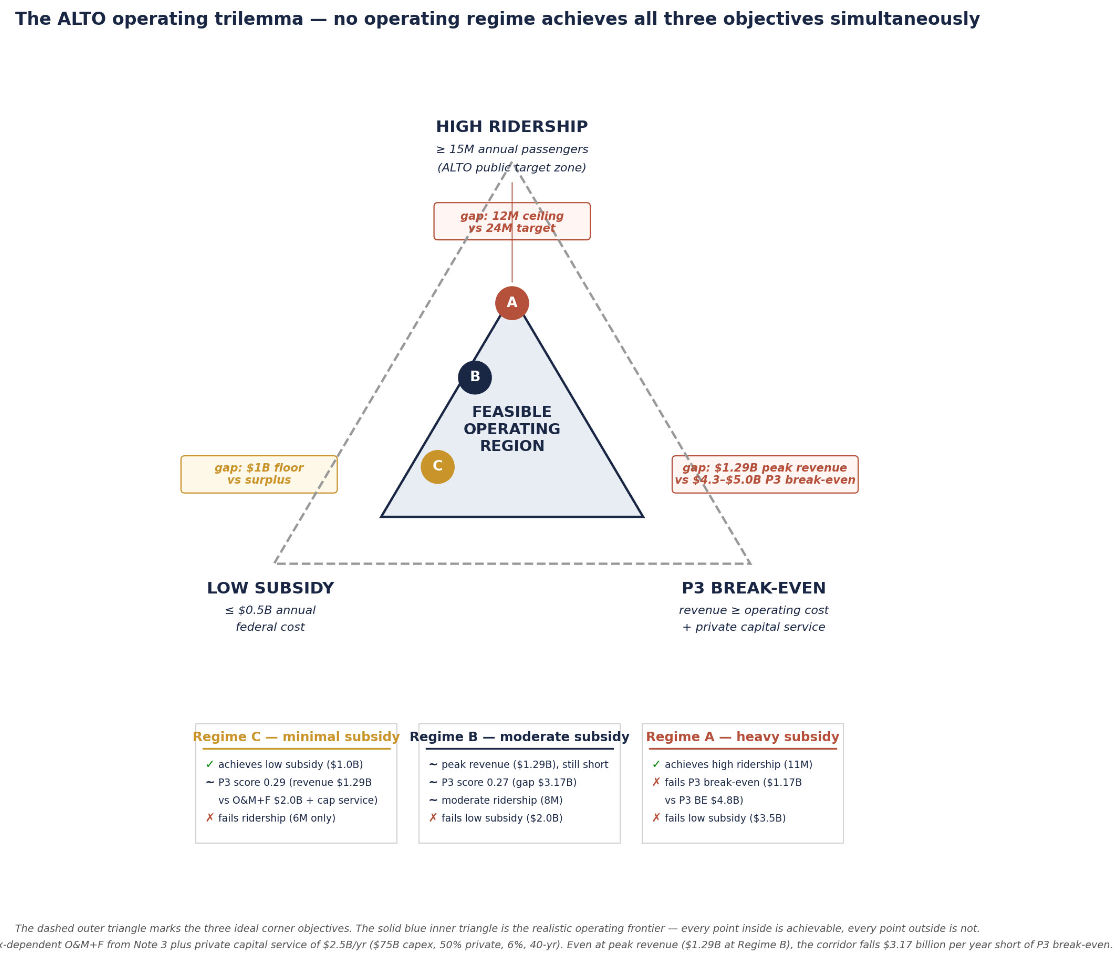

Figure 1. *The ALTO operating trilemma. The dashed outer triangle marks the three ideal corner objectives — high ridership (at the level of ALTO's public targets), low subsidy (operating surplus), and P3 break-even (revenue covering operating cost plus private capital service). The solid blue inner triangle is the realistic operating frontier: every point inside is achievable under some combination of fare, subsidy, and modal-shift parameters consistent with the analysis in Notes 1, 2, and 3; every point outside (the three corner wedges between the dashed and solid triangles) is structurally infeasible. Regimes A (heavy subsidy) and C (minimal subsidy) approach their respective ideal corners but cannot reach them; Regime B sits on the frontier edge between them, achieving the revenue peak and the per-rider welfare-efficient operating point but not approaching any corner. The P3 break-even corner is the structurally unreachable extreme: with operating cost (O&M plus fleet capital) varying from approximately $1.8 to $2.5 billion per year across the ridership spectrum (Note 3, pax-dependent formula) plus private capital service of $2.49 billion per year (computed at $75 billion total project cost, 50% private share, 6% blended cost of capital, 40-year amortisation — the ALTO proponent-stated base case), P3 break-even revenue is approximately $4.3 to $5.0 billion per year, against an achievable peak fare revenue of $1.29 billion at Regime B. The gap of $3.17 billion per year at peak revenue — under the base-case capital structure — cannot be closed by operating-posture choice alone. The gap widens further under the CRI reference-class capital scenarios analysed in Section 5. The bottom panels summarise each regime's check-mark, partial, and cross-mark profile across the three objectives.*

Figure 2 presents the same three objectives in ternary-plot form. The advantage of the ternary view is that it shows the actual operating locus — the path that an operator traces through objective space as the subsidy level varies continuously — rather than just three discrete regime points. Each operating point on the subsidy frontier is mapped to barycentric coordinates representing its normalised achievement of the three corner objectives: ridership score = N / 15M target; low-subsidy score = 1 − S / $5B; P3 break-even score = R / P3 BE(pax), where P3

BE(pax) = O&M+F(pax) + private capital service and O&M+F(pax) is the pax-dependent recurring lifecycle cost from Note 3 ($1,381M + $89.7M × pax/yr in millions). Private capital service is $2.49 billion per year under the harmonized base case ($75 billion total capex, 50 per cent private share, 6 per cent blended cost of capital, 40-year amortisation). The three scores are then normalised to sum to 1, giving each point a unique location in the triangle. This formalisation has the property that distance from each corner is a direct measure of how far short of that objective the operator falls.

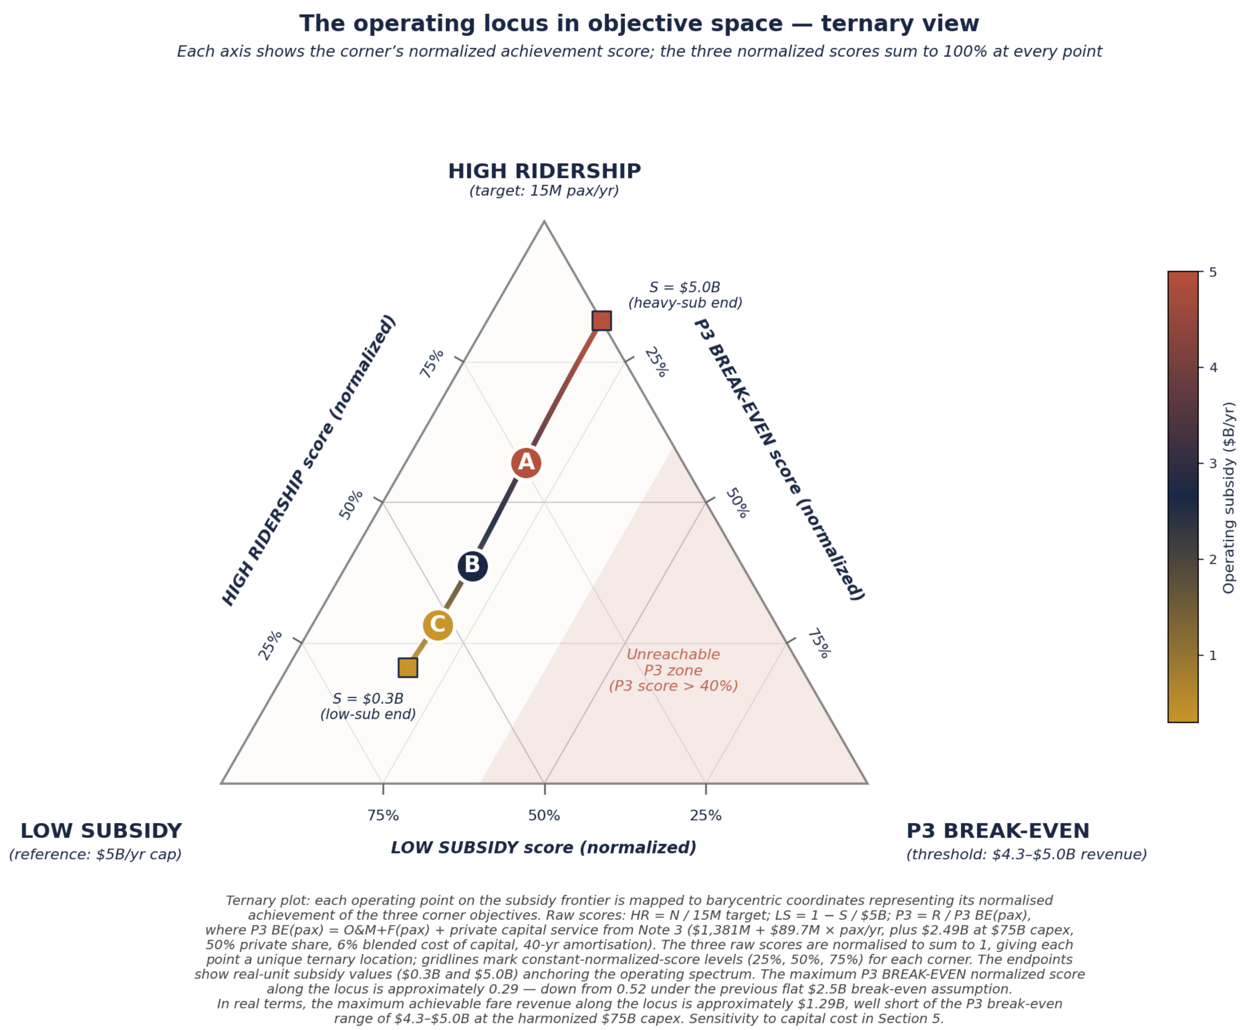

Figure 2. *The operating locus in objective space, ternary view. Each operating point on the subsidy frontier (Figure 3) is plotted as a point in the triangle whose distances to the three corners reflect normalised achievement of the three ideal objectives. The locus is the path through this space as operating subsidy varies from $0.3 billion (low-sub end, gold square at lower-left) to $5.0 billion (heavy-sub end, terracotta square at upper-right), passing through Regime C, Regime B, and Regime A in sequence. The colour gradient along the locus encodes subsidy level (gold at low, navy at moderate, terracotta at high). Two structural features stand out. First, the locus is a one-dimensional curve, not a two-dimensional region: the corridor has only one operational degree of freedom (the subsidy level), so the achievable space is a curve through the triangle rather than a filled region. Second, the locus tracks the LOW SUBSIDY ↔ HIGH RIDERSHIP edge closely; it never enters the right wedge near the P3 BREAK-EVEN corner (shaded faintly in red). The maximum P3 score achieved along the locus is approximately 0.30 under the harmonized $75B capex base case — down from approximately 0.52 under the previous flat $2.5B break-even assumption — because the pax-dependent O&M+F values from Note 3 plus the harmonized $2.49B private capital service produce*

a higher and ridership-rising P3 break-even revenue ($4.3 to $5.0 billion). The maximum P3 score occurs around Regime C, where revenue is close to its peak but operating cost is at its lowest. Sensitivity to the capital cost assumption is analysed in Section 5. This is the structural meaning of the trilemma: the three objectives are not symmetric tradeoffs but a single dominant tradeoff (ridership ↔ subsidy) with P3 break-even as a structurally unreachable third axis.

The ternary view also confirms the regime-positioning claim from the trilemma diagram (Figure 1). Regime C sits near the LOW SUBSIDY corner with modest P3 score, Regime A sits between centre and the HIGH RIDERSHIP corner, and Regime B sits on the locus between them — not near any corner — confirming that B achieves the revenue peak (an interior optimum) rather than approaching any corner objective. The empty right wedge confirms that no operating regime achieves the P3 corner objective. The single curve of the locus through the triangle is the geometric statement of the headline finding: the corridor cannot simultaneously achieve high ridership, low subsidy, and P3 break-even — and the limiting form of the tradeoff is the LOW SUBSIDY ↔ HIGH RIDERSHIP edge that the locus tracks.

1. Framework

Note 3 developed three discrete fare and subsidy regimes — A (heavy subsidy), B (moderate subsidy at parity with air), and C (minimal subsidy with P3 yield management) — that bracket the realistic policy space. The three regimes produce aggregate corridor modal shares of approximately 40, 30, and 22 per cent of the addressable market respectively, and require annual operating subsidies of approximately $3.5 billion, $2.0 billion, and $1.0 billion. Each regime corresponds to a different rail-to-air price ratio (approximately 0.55, 1.0, and 1.4 respectively) and a different per-mode capture rate along the modal-shift S-curves derived in Notes 1 and 2. This note extends that analysis from three discrete regimes to a continuous subsidy spectrum, in order to identify the optimisation properties of the corridor's operating posture.

The framework relates four quantities along the subsidy spectrum: annual subsidy (the federal operating contribution required to support the chosen fare posture), annual ridership (the resulting modal shift across the corridor's three sub-markets: air, road, and existing rail), annual fare revenue (riders times average fare), and net public cost (subsidy minus revenue, with negative values indicating a self-financing operation). Each quantity is anchored on the central demographic 2055 scenario from Note 3 (corridor population 20.1 million, addressable trips 34.2 million) with phase maturity reflecting the regime-coupled schedule developed in Note 3 Section 5. The mapping from subsidy to fare ratio is a smooth logistic that reproduces the three regime anchor points: at $1.0B annual subsidy the fare ratio is approximately 1.3 (deep premium to air), at $2.0B it is approximately 1.0 (parity with air), and at $3.5B it is approximately 0.6 (deep discount to air). The mapping from fare ratio to per-mode capture rate is derived directly from the Note 1 air-rail S-curves and Note 2 road-rail S-curves.

2. The subsidy frontier

Figure 3 shows the four relationships across the continuous subsidy spectrum, with the three regime anchor points (C, B, A) marked. The four panels disaggregate the relationships that are folded together in Note 3's regime-based summary.

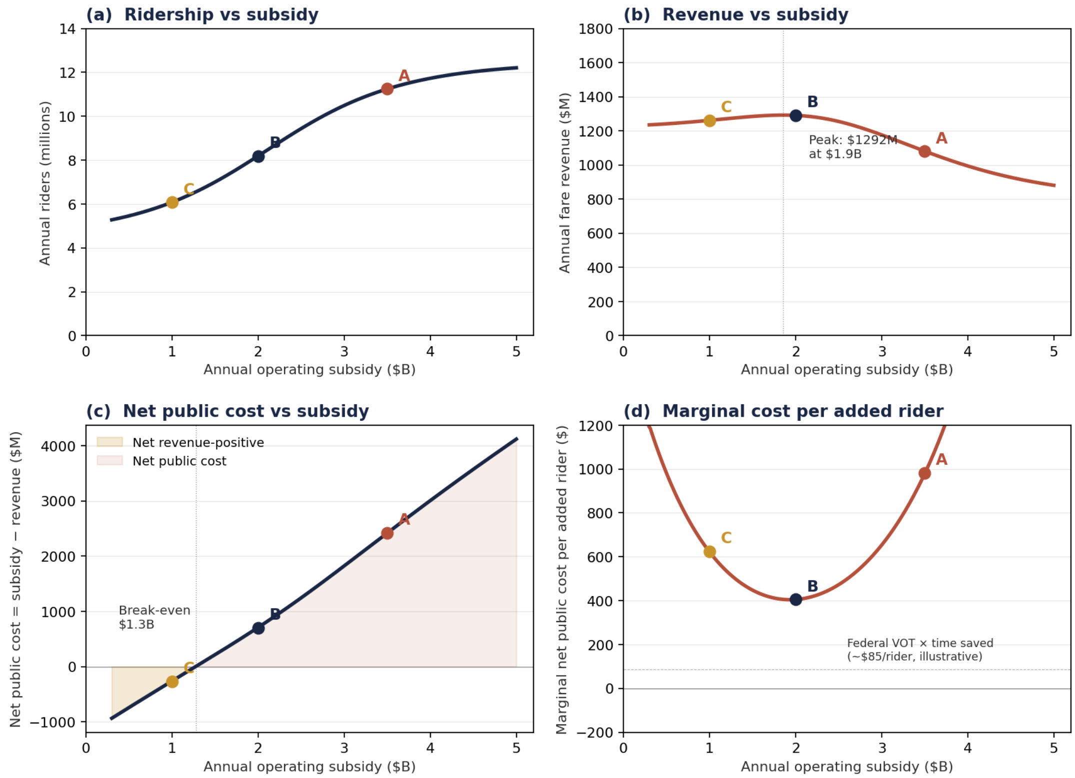

Figure 3. *The subsidy frontier for the ALTO corridor at the central demographic 2055 anchor. Panel (a) shows annual ridership rising concavely from approximately 5 million at $0.3 billion annual subsidy to approximately 12 million at $5 billion — diminishing returns as the corridor approaches the modal-shift ceiling. Panel (b) shows annual fare revenue rising to a peak near $1.9 billion subsidy at approximately $1.29 billion, then declining as additional fare cuts overwhelm the additional ridership they buy — the Laffer-like structure of corridor revenue. Panel (c) shows the net public cost (subsidy minus revenue) crossing zero near $1.3 billion subsidy: below that level the corridor operates at a net public revenue surplus, above it the corridor requires net public outlay rising to approximately $4 billion at $5 billion subsidy. Panel (d) shows the marginal net public cost per added rider, which has a smooth U-shaped minimum at approximately $400 per rider near the Regime B operating point and rises to approximately $1,000 per rider at Regime A. Reference horizontal line at $85 per rider indicates a rough illustrative federal value-of-time figure for the time savings each additional rider would experience switching from car to rail.*

The four panels together describe the corridor's subsidy frontier. Three structural features are visible. First, ridership is genuinely concave in subsidy: the first dollars of subsidy buy lots of additional riders because they move the corridor up the steeper part of the modal-shift S-curves; the last dollars buy progressively fewer because they push the corridor into the saturating top of the curves. The slope of panel (a) is approximately 2.5 million additional riders per billion dollars of subsidy at the low end and approximately 0.4 million per billion at the high end — a sixfold reduction in marginal effectiveness across the spectrum. Second, revenue is hump-shaped: at low subsidy the corridor is in the premium-fare zone where each additional rider pays more than the average rider in the higher-subsidy zone, so revenue rises as ridership grows; past the peak, the fare reduction overwhelms the ridership gain. Third, net public cost transitions cleanly from negative (revenue exceeds subsidy) to positive (subsidy exceeds revenue) at around $1.3 billion of annual subsidy — a point that sits between the Regime C anchor of $1.0 billion and the Regime B anchor of $2.0 billion.

3. Revenue and subsidy versus ridership

Figure 4 presents the same data with annual ridership on the horizontal axis, showing how subsidy and revenue diverge as the corridor moves up the ridership scale. The chart also overlays the federal capital service ($2.49 billion per year at the harmonized $75 billion base case, 50 per cent federal share, 6 per cent blended cost of capital, 40-year amortisation), so each regime now shows three quantities: annual operating subsidy (navy circle), annual fare revenue (terracotta circle), and full annual federal cost (slate diamond — operating subsidy plus federal capital service). This is the structurally informative view for the full fiscal commitment: the question is not simply how much operating subsidy is needed, but how the full federal cost stacks up at any given ridership target. Sensitivity to higher capital cost scenarios ($143B reference-class central, $264B P95) is analysed in Section 5.

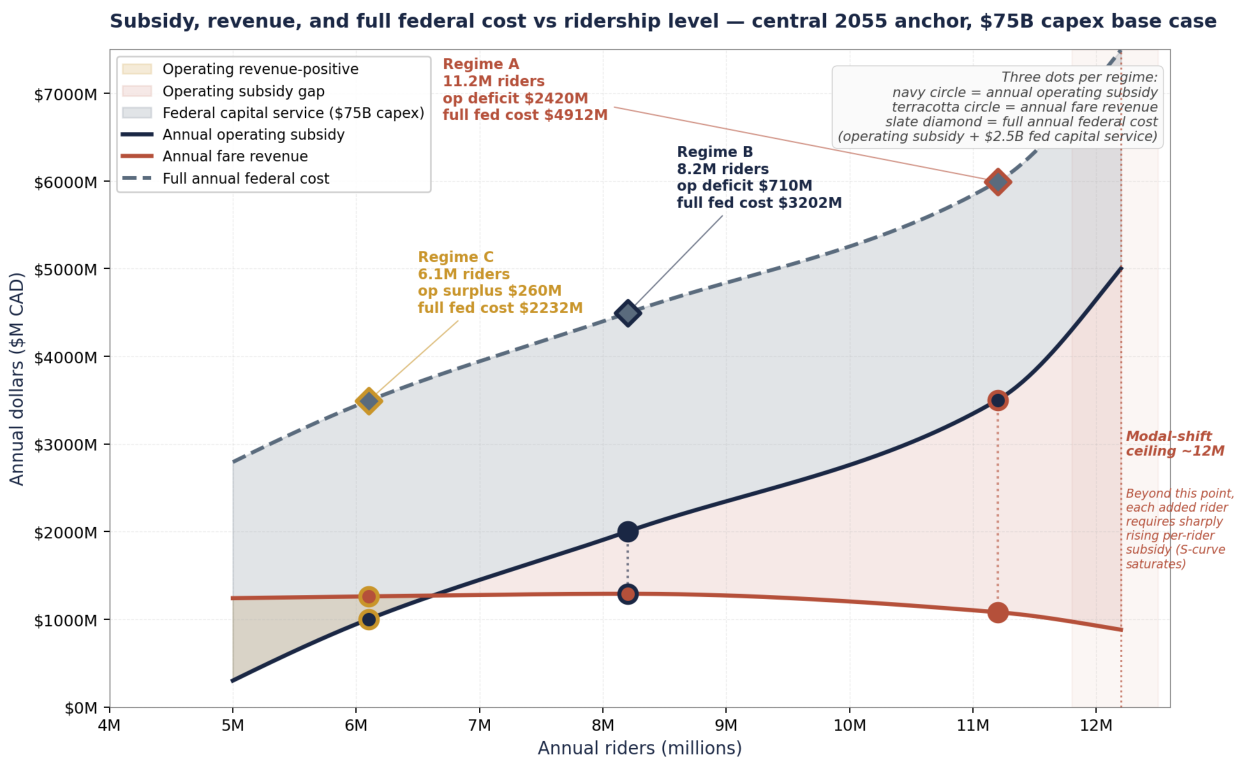

Figure 4. *Subsidy and revenue plotted against annual ridership level, central 2055 anchor. The two curves form a scissors structure: subsidy (navy) rises convexly with ridership, while revenue (terracotta) is essentially flat between 6 and 10 million riders and declines mildly above 10 million. The gap between the two curves is the net public cost, shaded gold (net revenue-positive, riders below approximately 6.5 million) or terracotta (net public outlay, riders above approximately 6.5 million). At Regime C (6.1 million riders), the corridor returns a net operating surplus of approximately $260 million per year; with the $75B base-case federal capital service added, the full federal cost is*

approximately $2.23 billion per year. At Regime B (8.2 million riders), the corridor requires approximately $710 million in net operating outlay; with capital service added, the full federal cost is approximately $3.20 billion. At Regime A (11.2 million riders), the net operating outlay is approximately $2.42 billion; full federal cost approximately $4.91 billion. Capital service exceeds operating subsidy at every regime, even under the proponent-stated base case. The chart is capped at the modal-shift ceiling of approximately 12 million annual riders, where the rising operating subsidy reflects deeply discounted fares pushed below operating-cost recovery: the subsidy here is not simply operating cost minus revenue (operating cost in the central case is approximately $1.8 to $2.5 billion per year across the ridership range, per Note 3), but the federal contribution required to support the chosen fare posture under the modal-shift framework. At high subsidy levels, the public pays both the operating gap and a fare buy-down that makes per-rider fares fall below cost recovery in order to attract additional modal-shift ridership. Beyond approximately 12 million riders, each additional rider requires sharply rising per-rider subsidy as the modal-shift S-curve saturates.

The scissors structure has direct policy implications. At any ridership target below approximately 6.5 million annual passengers (the right edge of the gold-shaded zone), the corridor is operating at a net public revenue surplus: fare revenue exceeds the subsidy needed to sustain that ridership level. This is consistent with a posture where the corridor functions as a self-financing commercial enterprise, similar to a yield-managed P3. Above 6.5 million riders, the corridor crosses into net-public-cost territory, and the public cost rises convexly with the ridership target. By 11 million annual passengers (close to the Regime A operating point), the corridor requires approximately $2.4 billion annually in net public outlay above and beyond the fare revenue it generates. Beyond 11.5 million the curve steepens sharply, indicating that pushing the corridor toward the ALTO public targets of 24 million by 2055 would require an entirely different operating regime than any of the three considered here.

4. Optimisation objectives

The subsidy frontier supports several distinct optimisation objectives that produce different optimal operating postures. Table 1 summarises five candidate objectives commonly applied to public transport investments — revenue maximisation, net-public-cost-per-rider minimisation, total-net-public-cost minimisation, fiscally-constrained ridership maximisation, and total welfare maximisation — each of which selects a different point on the subsidy frontier developed above.

Objective	Optimal regime	Riders 2055	Subsidy	Revenue	Net public cost
Maximise fare revenue	Regime B (parity with air)	~8M	$1.9–2.0B	$1.29B (peak)	+$0.7B
Minimise net public cost per rider	Regime B (parity with air)	~8M	$1.9–2.0B	$1.29B	$400 marginal
Minimise total net public cost	Regime C (yield mgmt) or below	~6M	$0.5–1.5B	$1.26B	+$0.2B or surplus
Maximise ridership s.t. fiscal cap	Regime A (heavy subsidy)	~11M+	$3.5B+	$1.08B	+$2.4B
Maximise total welfare[1]	Between B and A	~9M	$2.5B	$1.2B	+$1.3B

Table 1. Optimal operating posture under different objective functions, central 2055 anchor. Note [1]: "total welfare" includes ridership × value-of-time savings × emissions avoided − net public cost. The optimum under this composite objective depends on assumed parameter values; for illustrative federal-treasury parameters (value of time ≈ \$25/h, average time saved ≈ 1.5 h per rider, emissions avoided ≈ 30 kg CO_2 per rider at \$170/tCO_2), the social value per rider is approximately \$42 in time savings plus \$5 in emissions = \$47 per rider, which falls short of the \$400 marginal net public cost at Regime B. Under more favourable assumptions (value of time ≈ \$40/h, network effects included, induced demand counted at full value), the social value can reach \$150–250 per rider, which still sits below the marginal net public cost across most of the subsidy spectrum but closer to break-even at lower subsidy levels. A formal cost-benefit analysis would need to specify these parameters explicitly.

Four observations follow. First, the revenue-maximisation and the per-rider welfare-efficiency objectives converge on Regime B. This is not coincidental: the same condition (marginal revenue equals marginal cost) that defines the Laffer peak in panel (b) of Figure 3 also defines the minimum of the marginal-cost-per-rider curve in panel (d). A profit-maximising private operator and a welfare-maximising public authority — under the assumption that they apply marginal rather than total welfare as the criterion — would, under simplifying assumptions, choose the same fare posture. They would differ on whether the corridor should operate at that posture at all (the profit-maximiser would require revenue to exceed full cost including capital service; the welfare-maximiser would require social value to exceed net public cost), but conditional on operating, they would choose similar fare structures.

Second, the minimum-total-net-public-cost objective points toward Regime C or below. At Regime C the corridor generates a small operating surplus (revenue exceeds subsidy by approximately \$200 million annually). Below Regime C — at subsidy levels under \$0.8 billion — the surplus grows to approximately \$400 million annually. The cost to the public treasury is minimised, but at the cost of ridership: only about 5 to 6 million annual passengers, which is below the McGill TRAM stated-preference projection and well below the corridor's modal-shift potential. This is approximately the operating posture implied by the Cadence consortium's announced commercial structure.

Third, the ridership-maximisation-subject-to-fiscal-cap objective points toward Regime A or beyond. Regime A produces approximately 11 to 12 million annual riders at a net public cost of approximately \$2.4 billion. To approach ALTO's published 24-million-by-2055 target would require pushing past Regime A into a posture with subsidy substantially above \$5 billion annually and modal share above the 40 per cent ceiling — not a feasible operating posture under the modal-shift framework developed in Notes 1 and 2. The 24-million target is therefore not the welfare-efficient operating point under any reasonable parameter choice; it is achievable, if at all, only under heroic assumptions about every operating, demographic, and modal-shift variable simultaneously.

Fourth, the total-welfare-maximisation objective — which weighs total ridership times per-rider social value against total net public cost rather than the marginal versions — produces an optimum between Regime B and Regime A, contingent on the assumed social value per rider. The total welfare function $W(S) = N(S) \times v - [S \times 1000 - R(S)]$ (where N is annual ridership, v is social value per rider in dollars, S is operating subsidy in billions, and R is fare revenue in millions) has its maximum where the marginal-cost-per-rider curve in Figure 3 panel (d) crosses the assumed value v. At the illustrative federal-treasury value of approximately $85 per rider, the maximum is at or below Regime C (the marginal cost only falls below $85 at very low subsidy levels). At a moderately higher value of $200 per rider — appropriate if induced demand and emissions externalities are counted at full value — the maximum sits near Regime B. At a high value of $400 per rider — appropriate only if network effects, large emissions externalities, and corridor agglomeration benefits are all credited — the maximum sits between Regime B and Regime A. Total welfare maximisation is therefore strongly sensitive to the social value parameter, while marginal-cost minimisation is robust to it.

The five objectives together describe the structurally informative point: there is no single "optimal" operating posture without specifying the optimisation criterion. The five candidate optima span Regime C (revenue-positive operation, minimum total public cost), Regime B (revenue peak, per-rider welfare efficiency), an intermediate position (total welfare under moderate social-value assumptions), and Regime A (maximum ridership). The corridor decision is therefore not a single quantitative question but a sequence of three sequential decisions: whether to build the corridor at all (a comparison of full social value against full federal cost, addressed in Section 5), what fare posture to operate it under once built (a comparison of operating subsidy against operating revenue and marginal ridership benefits, addressed in this section), and how to communicate the chosen posture transparently in the public business case (addressed in Section 6).

5. Capital cost and full-cost accounting

The subsidy frontier developed above considers operating subsidy only — the annual federal contribution required to sustain the corridor's chosen fare posture. It does not include capital cost service, which dominates the corridor's total fiscal commitment. The full-cost picture of the corridor's public funding requirement is therefore the operating subsidy from Figure 3 plus the capital service requirement, and the capital cost itself is subject to substantial uncertainty that must be made explicit before any full-cost conclusion can be drawn.

ALTO's public materials cite a total project capital cost of approximately $60 to $90 billion. This is the proponent-stated range, prepared without reference-class adjustment for cost overruns. The Citizen Research Initiative's reference-class analysis (Flyvbjerg methodology applied to the international HSR cost database, with corridor-specific adjustments for the Frontenac Arch, Napanee Limestone Plain, Leda clay segment, and Canadian P3 megaproject delivery record) produces three scenario points that bracket the realistic capital cost range. The proponent-stated value of $75 billion is retained as the P50 central estimate under the as-stated specification, $143 billion is the P50 central estimate after reference-class adjustment for the 44.7 per cent average rail-project overrun and the corridor's specific complexity premia, and $264 billion is the P95 worst-case estimate accounting for the upper-tail risk that the Flyvbjerg distribution shows for projects in the engineering complexity range that ALTO occupies (composite score 73–81 on the CRI's engineering complexity rubric, the upper part of the High band). The three scenarios should be understood as a probability-weighted range, not as alternative point estimates: the realistic expected value of the corridor's capital cost is somewhere between $143 billion and $264 billion, with the proponent's $75 billion sitting at approximately the 25th percentile of the reference-class distribution.

The capital service implications of the three scenarios, under a representative procurement structure (federal share 50 per cent, blended cost of capital 6 per cent, amortised over 40 years), are shown in Table 2. Combined with the Regime B operating subsidy of $2 billion per year, the full annual federal cost ranges from $4.3 billion at the proponent-stated capital cost to $9.9 billion at the upper reference-class scenario. The corresponding full-cost-per-rider at the central 2055 ridership of approximately 8 million annual passengers ranges from $540 to $1,240.

Capital cost scenario	Total capital	Annual debt service (total)	Federal share (50%)	Full annual federal cost[1]	Full cost / rider[2]
ALTO proponent-stated	$75B	$4.5B	$2.3B	$4.3B	$540
CRI reference-class central	$143B	$8.6B	$4.3B	$6.3B	$790
CRI P95 worst-case	$264B	$15.8B	$7.9B	$9.9B	$1,240

Table 2. Full federal cost implications across three capital cost scenarios. The $75 billion ALTO proponent estimate, the $143 billion CRI reference-class central, and the $264 billion CRI P95 worst-case are derived from Flyvbjerg-methodology reference-class forecasting applied to the international HSR cost database, with corridor-specific complexity adjustments. [1] Full annual federal cost = federal share of capital debt service + Regime B operating subsidy of $2.0 billion per year (the welfare-efficient operating point from Section 4). [2] Full cost per rider = full annual federal cost ÷ 8 million annual riders (Regime B central 2055 ridership). Debt service computed at 6 per cent blended cost of capital, 40-year amortisation.

Three observations follow. First, capital service dominates operating subsidy at every scenario. Even at the proponent's $75 billion, federal capital service ($2.3 billion per year) exceeds Regime B's operating subsidy ($2.0 billion). At the CRI reference-class central of $143 billion, capital service is more than twice the operating subsidy. At $264 billion, capital service is approximately four times the operating subsidy. The full-cost optimisation is therefore dominated by the capital cost assumption, not by the operating regime choice. This shifts the framing of the corridor decision: choosing Regime A versus Regime B versus Regime C is a second-order question once the capital is committed, while choosing whether to commit the capital at all is the first-order question.

Second, the full-cost-per-rider across the three scenarios spans $540 to $1,240. Against an illustrative federal value-of-time per rider of approximately $85 (1.5 hours saved per trip at $25 per hour plus emissions value, the Treasury Board reference parameters), the corridor is approximately 6 to 14 times more expensive than the public benefit it delivers under standard cost-benefit assumptions. Under more generous social-value parameters that credit network effects, full induced-demand valuation, and large emissions externalities, the social value per rider can be estimated at $200 to $250 — still 2 to 6 times below the full cost per rider across the three capital scenarios. The corridor cannot be made welfare-positive at central ridership levels under any of the three capital cost scenarios using standard cost-benefit parameters.

Third, the choice of which capital cost scenario applies is the dominant policy decision facing the corridor. The proponent-stated $75 billion would be defensible if the project were comparable to delivered international HSR (Madrid–Barcelona, Paris–Lyon, Tokyo–Osaka) — corridors with composite engineering complexity scores in the 40s to low 50s. ALTO's composite engineering complexity score is 73–81 (upper part of the High band, approaching Extreme), 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 (Eglinton Crosstown +280 per cent overrun, Ottawa Confederation Line +57 per cent overrun, Ontario Line +250 per cent overrun in scope-adjusted terms). Under Flyvbjerg reference-class forecasting, a corridor at this complexity level cannot be reliably costed from the lower-complexity international comparators that the proponent's estimate appears to draw on. The realistic expected capital cost is between the CRI reference-class central ($143B) and P95 worst-case ($264B), making the full-cost-per-rider at central 2055 ridership approximately $790 to $1,240 — well above any reasonable per-rider social value, and producing a benefit-cost ratio materially below 1.0 across the full plausible range.

The policy implication of full-cost accounting under the three CRI capital cost scenarios is that the corridor decision should be made before the capital is committed, not after. Once the capital outlay (between $75 billion and $264 billion) is sunk, the operating-regime choice between A, B, and C is between similarly-priced options, and the choice should default toward the welfare-efficient operating point (Regime B). But the prior question — whether to build the corridor at all — is sensitive to which capital cost scenario actually materialises. The published business case relies on the $75 billion proponent figure without reference-class adjustment, and does not address how the procurement structure changes if the realised cost is closer to the $143 billion CRI central or the $264 billion P95 worst-case. A defensible business case would publish a probability-weighted cost-benefit analysis across the three scenarios, with explicit treatment of the federal fiscal exposure under each.

6. Implications for the corridor decision

Three implications follow for independent review of the corridor's projected fiscal trajectory. First, the choice of operating subsidy 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 subsidy frontier is chosen. Operating at Regime C produces approximately 6 million riders at a small operating surplus; operating at Regime A produces 11 million riders at $2.4 billion annual net public cost. The choice between these is a political-economy choice about how to weight commercial yield, public welfare, ridership, and fiscal restraint, and that 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 near Regime B (parity with air, $1.9–2.0 billion annual operating subsidy, approximately 8 million annual riders, marginal net public cost approximately $400 per added rider). 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. The published ALTO business case does not specify which optimisation objective is being applied or what the implied operating posture will be.

Third, the public ridership targets cannot be reached from any operating point on the subsidy frontier developed here. The 24-million-by-2055 figure would require modal share above the 40 per cent ceiling under heavy subsidy, plus upper-case demographic growth, plus full-corridor mature operation in 2055 — three conditions that none of the modal-shift literature supports simultaneously. The frontier developed in this note brackets the realistic operating space; ALTO's published targets sit outside that space. An independent review should ask which point on the frontier the corridor is actually targeting, and what the implied fiscal commitment and modal-shift assumptions are.

Methodology and parameters

The framework anchors on the central demographic 2055 scenario from Note 3 (corridor population 20.1 million, addressable trips 34.2 million at 1.7 trips per capita) with the regime-coupled phase maturity factor (Regime C ≈ 0.80, Regime B ≈ 0.88, Regime A ≈ 0.94), where maturity follows a smooth logistic asymptoting to ≈ 0.96 at high subsidy. The market structure is air 15 per cent, existing rail 10 per cent, road 75 per cent of the addressable pool. The mapping from operating subsidy S (in $ billions) to fare ratio r is a logistic: $r(S) = 0.4 + 1.3 / (1 + \exp(S - 1.8))$, calibrated to reproduce the three regime anchor points. The mapping from fare ratio to per-mode capture is derived from the Note 1 air-rail S-curve at journey time 3.0 hours (mid-corridor) and the Note 2 road-rail S-curve at $\tau = 0.5$ (NA-calibrated, infrastructure-relevant time ratio under ALTO speed). The average air fare is $160 per one-way trip; rail revenue is computed as riders times average rail fare (air fare times r). Net public cost is subsidy minus revenue, with negative values indicating operating surplus. Capital cost scenarios in Section 5 ($75B, $143B, $264B) are derived from Flyvbjerg reference-class forecasting applied to the international HSR cost database with corridor-specific complexity adjustments for the Frontenac Arch, Napanee Limestone Plain, Leda clay segment, St-Lawrence crossing, and Canadian P3 megaproject delivery record (composite engineering complexity score 73–81 on the CRI rubric, upper part of the High band). Capital service computed at 6 per cent blended cost of capital (combining federal debt service and private equity return), 40-year amortisation, 50 per cent federal share. The CRI's full capital cost analysis is documented separately at citizenresearch.ca.

Sources

- ALTO HSR Citizen Research Initiative (2026). Modal shift between high-speed rail and air on the ALTO corridor (Note 1). citizenresearch.ca.
- ALTO HSR Citizen Research Initiative (2026). Modal shift between rail and car on the ALTO corridor (Note 2). citizenresearch.ca.
- ALTO HSR Citizen Research Initiative (2026). ALTO HSR ridership envelope, 2035–2080 (Note 3). citizenresearch.ca.
- Statistics Canada (2026). Population Projections for Canada (2025 to 2075). Catalogue 17-20-0003, released January 27, 2026.
- Transport Canada (2024). Guide to Benefit-Cost Analysis of Transportation Investments — including value-of-time and emissions valuation parameters.
- Treasury Board of Canada Secretariat (2007). Canadian Cost-Benefit Analysis Guide: Regulatory Proposals.