

Modal Shift Note 1:

Modal shift between high-speed rail and air on the ALTO corridor

An assessment of the international evidence on when rail substitutes for air, applied to the Toronto–Ottawa and Toronto–Montréal corridor pairs that account for the bulk of the corridor's air-substitutable demand. Compares ALTO, current VIA Rail service, and a High Performance Rail alternative on both travel time and price.

SUMMARY

The international literature converges on a logistic S-curve in which rail captures the majority of the combined rail+air market on city pairs with station-to-station travel times of two to four hours, and rail's share collapses rapidly above five hours.

The corridor's modal-shift opportunity is concentrated in two city pairs — Toronto–Ottawa and Toronto–Montréal — that together account for the dominant share of air-substitutable demand. Both fall inside the competitive zone under any modern dedicated-track scenario.

The majority of the achievable modal shift on each principal pair is captured by moving from VIA's current shared-track service to a dedicated, electrified High Performance Rail corridor at conventional 200 km/h speeds. ALTO's additional 300+ km/h capability captures a further 19 to 20 percentage points at price parity — a real but residual gain.

Once price is added to the analysis, the picture shifts. Under canonical price assumptions — VIA at $r \approx 0.5$, HPR at $r \approx 0.7$, ALTO at $r \approx 1.0$ — ALTO's apparent 19–20 percentage-point time-only advantage shrinks to 11–13 points on the principal Toronto pairs. The cost-per-point for that incremental modal shift is several billion dollars; the cost-per-point of the larger HPR step that precedes it is much lower.

1. Travel time: the competitive zone

The empirical literature on rail–air substitution converges on a remarkably consistent set of distance and travel-time thresholds. Studies in Europe, China, Japan, and elsewhere identify a competitive break-even threshold of roughly 400 to 600 kilometres (about 2 to 3 hours door-to-door) for short-haul routes, beyond which aviation begins to regain a time advantage. Medium-distance corridors of 600 to 1,100 kilometres (about 3 to 5 hours) show the greatest demand elasticity, where both carbon and time savings favour rail when supported by high service frequency. Long-haul segments above 1,400 kilometres show minimal substitution potential under current infrastructure, with substitution typically below 10 per cent.

The International Council on Clean Transportation summarises the pattern more bluntly: high-speed rail can typically substitute for flights up to 1,000 km in distance, and is most competitive for trips under 700 to 800 km. The mechanism behind these thresholds is the door-to-door time calculation. Below roughly 700 km, the time spent reaching the airport, checking in, clearing security, boarding, taxiing, and reaching the destination city centre adds enough overhead that the total air journey time matches or exceeds high-speed rail. Above this distance, air's faster line-haul speed begins to dominate, and rail's share falls steeply once journeys exceed about 4.5 hours.

This relationship is conventionally modelled as a logistic S-curve, sometimes referred to in the transport literature as the Jorritsma curve. The shape is characteristic: under two hours rail captures essentially the entire air market; between three and four hours rail typically captures 60 to 80 per cent; between four and five hours rail's share collapses; above five hours rail captures only a residual share of price-sensitive or rail-loyal travellers. Frequency, station centrality, fare structure, and reliability shift the curve up or down by several percentage points but do not change its overall shape.

Empirical anchors

Three European routes anchor the empirical baseline. On Paris–Lyon, the TGV Sud-Est cut travel time from almost four hours to about two hours; rail's market share in the rail+air market rose from 40 per cent to 72 per cent, while air collapsed from 31 per cent to 7 per cent. On Madrid–Seville, with a 471 km line completed in 1992, rail share rose from 16 to 52 per cent of all modes and air shrank from 40 to 13 per cent. The Madrid–Barcelona AVE — at 621 kilometres and 2 hours 30 minutes the cleanest modern parallel to ALTO's longer pairs — has resulted in roughly 75 per cent of travellers between the two cities now choosing the train over the plane on the rail-vs-air pair, with the AVE taking around four years to decisively overtake the air bridge after its 2008 opening.

Asian comparators reach further. The 2019 World Bank review of Chinese high-speed rail found that service is very competitive with bus and aircraft for distances between 150 and 800 kilometres, and that 350 km/h services remain competitive with air for journeys up to about

1,200 kilometres. The Beijing–Shanghai HSR (1,318 km, 4 hours 18 minutes) is the canonical case where high frequency and 350 km/h operating speed maintain rail dominance at distances that would normally favour air. Tokyo–Osaka on the Tokaido Shinkansen (552 km, 2 hours 22 minutes) is another textbook 80+ per cent rail-dominant pair.




Figure 1. Modal-shift S-curve showing rail's share of the combined rail+air market as a function of station-to-station rail journey time. Logistic curve fitted with inflection at 3.5 hours and steepness parameter $k = 1.3$. The curve divides naturally into three zones: a short-haul band below 2 hours where rail dominates the rail+air market; a competitive zone between 2 and 4 hours where rail's share moves through the steep middle of the curve and infrastructure investment can decisively shift modal share; and a long-haul band above 4 hours where air's faster line-haul speed begins to dominate and rail's share collapses. The international comparator data discussed in this section is consistent with this shape, with all major HSR services in the competitive zone achieving rail shares of 70 to 85 per cent on the rail-vs-air pair.

Rail wins decisively under three hours, competes strongly at three to four hours, and degrades rapidly after that — with high frequency and central-station access being decisive variables alongside line-haul time.

2. Price: the elasticity factor

The S-curve in Figure 1 represents rail's share of the rail+air market as a function of travel time, with prices held implicitly at parity. Real modal choice is two-dimensional: passengers weigh both time and price, and the relative price of rail to air shifts the entire S-curve up or down. A logit choice model with a price utility term captures this directly: each doubling of the rail-to-air price ratio shifts the curve's inflection point earlier by an amount that depends on the price

coefficient. Figure 2 shows what this looks like in practice — a family of curves at six representative price ratios.

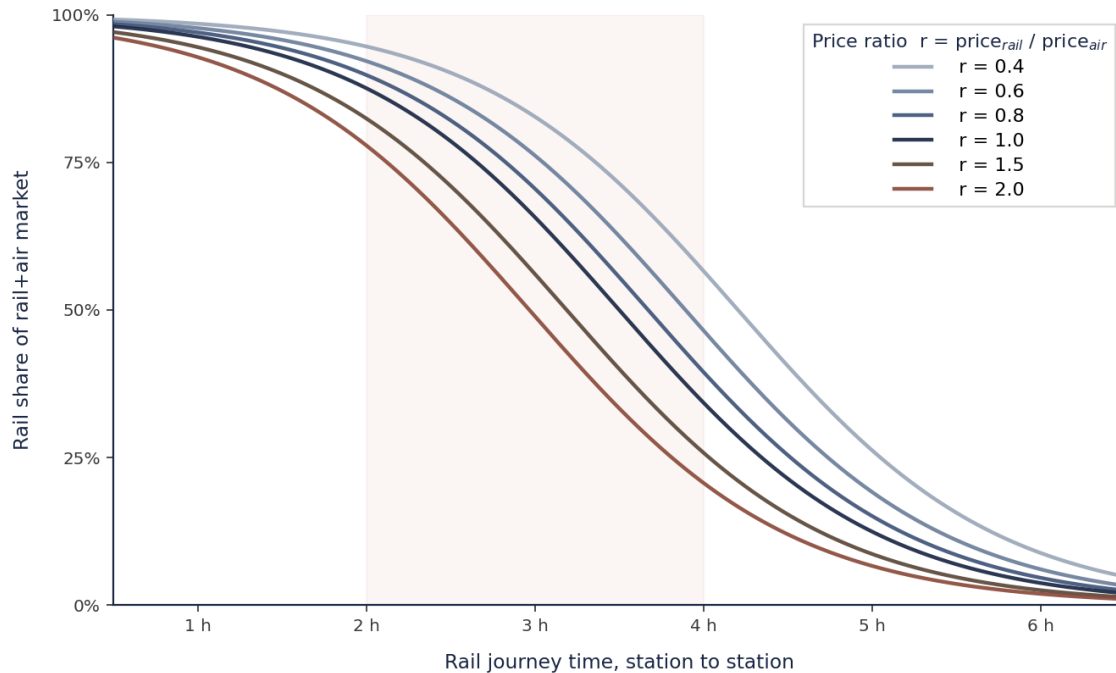


Figure 2. Family of modal-shift S-curves at six rail-to-air price ratios ($r = \text{rail price} \div \text{air price}$). The middle navy curve corresponds to $r = 1.0$ — the price-parity case from Figure 1. Curves to its upper right (the cooler tones) show what happens when rail is priced below air: the whole curve lifts, and rail captures a larger share at any given travel time. Curves to its lower left (the warmer tones) show the opposite — rail priced above air, and a corresponding loss of share.

How to read the chart

The simplest way to use Figure 2 is as a lookup. Pick a travel time on the horizontal axis, then pick the curve that matches the rail-to-air price ratio for that route, and read off the predicted rail share. A 3-hour rail journey at price parity ($r = 1.0$) sits at roughly 60 per cent rail share. The same 3-hour journey priced at half the air fare ($r = 0.5$) sits closer to 75 per cent. The same 3-hour journey priced at 1.5 times the air fare ($r = 1.5$) drops to around 45 per cent. The shift is symmetric in log-price: $r = 0.5$ lifts the curve by the same amount that $r = 2.0$ drops it.

What makes Figure 2 useful is that it lets time and price be assessed together. A faster service at a higher price can deliver lower modal share than a slower service at a lower price; the family of curves shows how the two effects combine.

Empirical anchors on price

International HSR routes that achieve high modal shares typically combine competitive travel times with rail fares well below air. On Madrid–Barcelona, AVE Básico fares routinely sit in the €40–70 range against air fares of €100–200, putting the rail-to-air price ratio in the 0.4 to 0.6

band; this favourable price position contributes meaningfully to the 75 per cent rail share at a 2 hour 30 minute travel time. Madrid–Seville and Paris–Lyon show similar patterns. Tokyo–Osaka on the Tokaido Shinkansen is the contrasting case: rail and air prices are roughly comparable (a price ratio in the 0.7–0.9 band), but central-station access and reliability sustain rail dominance even without a price advantage.

Price sensitivity differs by traveller type

Figure 2 shows a single family of curves with one price coefficient, but the price coefficient itself is not the same for every traveller. The international literature consistently finds that business travellers — particularly those whose fares are paid by employers and whose time is heavily valued — show much lower price sensitivity than leisure travellers. Typical price elasticities sit in the -0.4 to -0.7 range for business demand against -1.0 to -1.6 for leisure demand: a 10 per cent fare increase reduces business volumes by 4 to 7 per cent and leisure volumes by 10 to 16 per cent.

The implication for Figure 2 is that each curve is really a weighted average of two underlying curves — a flatter business curve that responds modestly to price changes, and a steeper leisure curve that responds more sharply. On routes where business travel dominates demand (typical of weekday peak intercity), modal share is less responsive to fare changes; on routes where leisure travel dominates (off-peak, weekends, holidays), fare changes produce larger swings in share. A pricing strategy that targets one segment can therefore deliver materially different modal outcomes than the aggregate family of curves suggests.

Air's connecting-flight advantage

Beyond time and price, air retains a structural advantage that the simple time-and-price model does not capture: the connecting-flight network. Travellers continuing beyond the corridor endpoint to international or transcontinental destinations face significant friction in switching modes — they must transfer between station and airport, re-check baggage, clear security separately, and lose the through-fare and frequent-flyer benefits that come with a single-carrier itinerary. Hub airports retain a portion of the air market that even a faster, cheaper rail service would not capture, because the overall journey ends at a long-haul gate, not a downtown station.

This effect is largest for trips originating or terminating outside the route in question — for example, a traveller flying short-haul to a hub and connecting onto a long-haul flight — and smallest for downtown-to-downtown trips entirely within a corridor. It interacts with traveller type as well: international business travellers in particular tend to be heavily concentrated on connecting itineraries, which insulates the short-haul air segment from rail competition. The modal-share envelope in Figure 2 should therefore be read as a ceiling for the rail-substitutable portion of the market, not for the air market as a whole.

Modal share depends on time, price, traveller type, and itinerary structure. The family of S-curves in Figure 2 captures the first two; the third and fourth shift the realistic modal-share envelope further.

3. Where the corridor sits on the curve for travel time

The corridor is not a single market. It is a sequence of overlapping city pairs whose distances place each segment in a different position on the modal-shift curve. The bulk of the corridor's air-substitutable demand is concentrated in two pairs: Toronto–Ottawa and Toronto–Montréal. The Toronto–Montréal air market alone runs roughly 900,000+ annual seats, with substantially more once Toronto–Ottawa connecting flows are added. These two pairs are the focus of the analysis that follows.

ALTO's published targets call for approximately 1,000 km of dedicated, electrified track at speeds of 300 km/h or more, with stops at Toronto, Peterborough, Ottawa, Montréal, Laval, Trois-Rivières, and Québec City. Published target travel times for the two principal pairs are about 2 hours Toronto–Ottawa and just over 3 hours Toronto–Montréal — both inside the 1.5–3.5 hour zone where international comparators consistently capture 70 to 90 per cent of the rail+air market.

Current VIA Rail baseline

VIA Rail's existing Corridor service is well outside the air-competitive zone on the principal Toronto pairs. Toronto–Montréal currently averages 5 hours 13 minutes over 538 km, with about 40 weekly departures in each direction; the fastest scheduled service runs about 4 hours 37 minutes. Toronto–Ottawa typically runs 4 to 4.5 hours. Current trains have a design top speed of 200 km/h but are limited to a maximum operating speed of 160 km/h due to infrastructure constraints, and the network is dominated by shared track with CN Rail freight, which is the principal cause of both the slow line-haul speed and the poor reliability — VIA's on-time performance was around 67 per cent as of 2021.

Despite this, the Corridor is VIA's commercial backbone. Corridor service contributes 81 per cent of VIA's revenue and 95 per cent of its ridership; the Montréal–Ottawa–Toronto route alone moved over 2.1 million passengers in 2023. Against the Toronto–Montréal air market of roughly 900,000+ annual seats — and the broader corridor air market once Toronto–Ottawa is included — VIA's current rail share of the rail+air market on these pairs is structurally low. At 4 hour 30 minute and 5 hour 13 minute journey times, this is precisely what the literature predicts.

Segment-by-segment travel times under each scenario

City pair	Distance	VIA current	HPR (200 km/h)	ALTO (300+ km/h)
Toronto–Ottawa	~450 km	~4 h 30 min	~2 h 55 min	~2 h
Toronto–Montréal	~540 km	5 h 13 min*	~3 h 38 min	~3 h
Ottawa–Montréal	~190 km	~1 h 55 min	~1 h 30 min	~1 h

Table 1. Indicative travel times for the principal corridor city pairs under each scenario. VIA current values are scheduled times on existing shared-track infrastructure. HPR values are Express journey times published in the CRI HPR Strategy (Chapter 4): a dedicated, electrified 401-corridor mainline at 200 km/h with downtown access via crossover spurs, and Ottawa via the Smiths Falls Subdivision from a Brockville crossover at 140 km/h. ALTO values are the published targets for the 300+ km/h network. *Toronto–Montréal under current VIA service runs 5 h 13 min on the 538 km direct routing via Kingston.

Plotted onto the modal-shift S-curve, these times produce the three pictures in Figures 3a–3c. Each panel shows the same time-only curve with the two principal Toronto pairs highlighted under one scenario; the contrast between panels traces the modal-shift trajectory at price parity as corridor infrastructure improves. The price-shifted reading — locating each scenario on its appropriate curve in the family from Figure 2 — is taken up immediately afterwards.

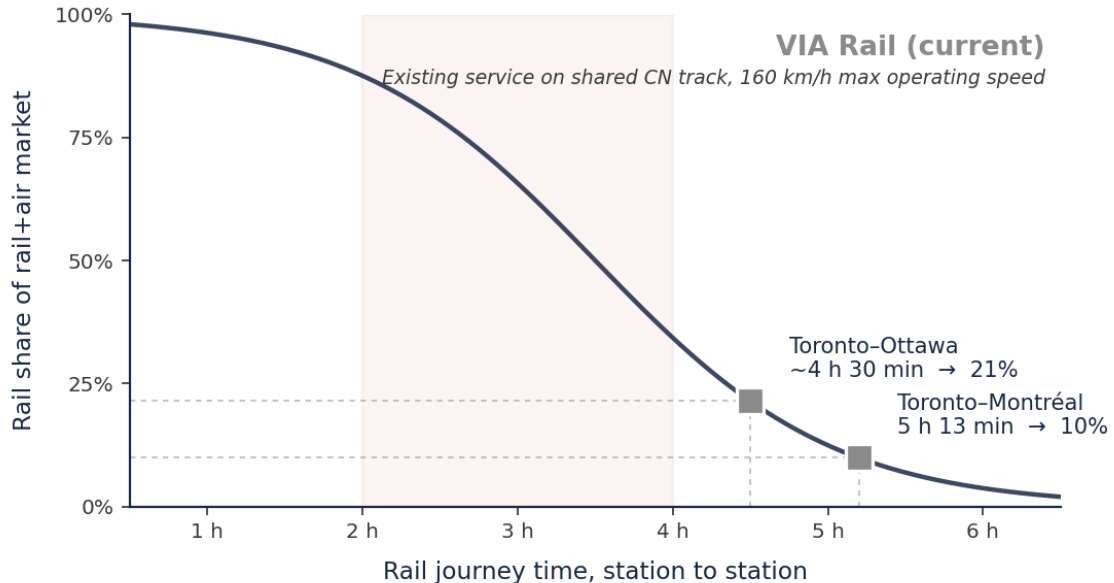


Figure 3a. Current VIA Rail service. Both principal Toronto pairs sit well below the inflection point: Toronto–Ottawa at roughly 4 h 30 min captures around 21% of the rail+air market, and Toronto–Montréal at 5 h 13 min captures around 10%. The corridor’s air-substitutable demand is structurally outside the competitive zone.

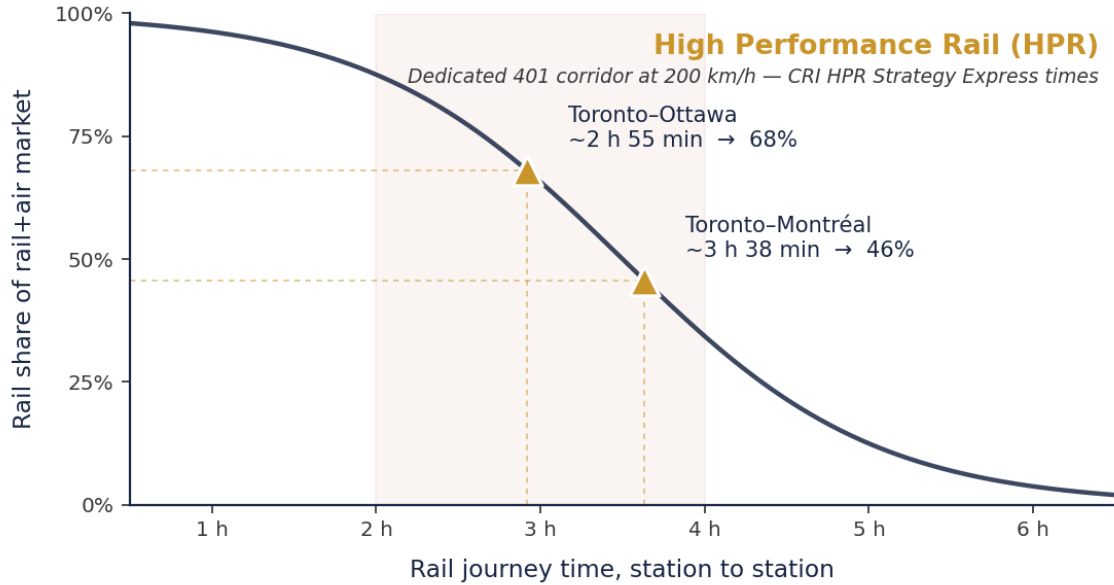


Figure 3b. High Performance Rail at 200 km/h on a dedicated, electrified 401-corridor mainline, with Express journey times of 2 h 55 min Toronto–Ottawa and 3 h 38 min Toronto–Montréal published in the CRI HPR Strategy. Toronto–Ottawa moves onto the upper portion of the curve at roughly 68% rail share at price parity. Toronto–Montréal moves to roughly 46% — across the inflection but still in the steeper portion of the curve.




Figure 3c. ALTO at 300+ km/h on a dedicated 1,000 km HSR network, with published targets of 2 hours Toronto–Ottawa and approximately 3 hours Toronto–Montréal. Toronto–Ottawa moves further onto the upper plateau at roughly 88% rail share at price parity. Toronto–Montréal moves to roughly 66% — still on the steeper portion of the curve, where additional time savings continue to produce meaningful share gains.

The implied rail-vs-air market shares for each scenario are summarised in Table 2.

City pair	VIA current	HPR (200 km/h)	ALTO (300+ km/h)
Toronto–Ottawa	~21%	~68%	~88%
Toronto–Montréal	~10%	~46%	~66%

Table 2. Predicted rail share of the combined rail+air market on each principal pair under each scenario, derived from the logistic curve in Figure 1 with prices held at parity. Values should be read as order-of-magnitude estimates; actual shares would also depend on fare structure, frequency, reliability, station accessibility, and traveller mix. The price dimension is taken up explicitly in Section 4.

These are the time-only readings — what each scenario would deliver if its fares matched air. In practice, fares depend on capital structure, and the corridor's three scenarios sit at quite different points on the price axis. Section 4 takes up that dimension.

4. Where the corridor sits on the curve for price

Section 2 introduced the family of curves showing how rail's modal share responds to the rail-to-air price ratio. This section maps the corridor's three scenarios onto that price axis, then tracks how modal share would respond as the price ratio varies across the realistic 0.4–2.0 range.

Capital cost, fares, and the corridor's price ratios

Current VIA Rail Toronto–Montréal Economy fares sit in the \$80–120 range for advance bookings against Air Canada Economy fares of \$200–400 — a price ratio in roughly the 0.4–0.6 band, the same band as Madrid–Barcelona and Paris–Lyon. The structural fare advantage that international HSR uses to lift modal share is already in place on the corridor; the binding constraint on current rail share is travel time, not price. For the analysis that follows, VIA is taken to operate at a rail-to-air price ratio of approximately 0.5.

Whether each new-build scenario can preserve a fare advantage against air depends on capital-cost recovery and the subsidy structure. The CRI HPR Strategy estimates corridor capital costs in the order of \$19 million per kilometre — roughly \$19–25 billion for the full Windsor–Montréal programme — producing annual debt service of \$1.0–1.3 billion at a 5 per cent rate. Spread across plausible corridor ridership of 6 to 10 million passengers per year, the per-trip debt-service burden is roughly \$100–200. Under a public-infrastructure subsidy model that absorbs most of that debt service into the public capital account — the standard model for European HSR and the model that VIA already operates under — HPR fares could plausibly sit at a modest premium over current VIA levels, in the order of \$130–170 against air fares of \$250–300. This places HPR at a rail-to-air price ratio of approximately 0.7: still meaningfully cheaper than air, but somewhat above current VIA fares to recognise the new-build capital.

ALTO's \$60–90 billion capital envelope produces annual debt service of roughly \$3.0–4.5 billion — three to four times higher. Spread across ALTO's projected 24 million annual passengers (a figure with no North American precedent), the per-trip debt-service burden is \$125–190 per

passenger; spread across more conservative 10–15 million ridership, \$200–450. Even with the same proportional subsidy treatment as HPR, ALTO's residual fare burden would push fares to or beyond air parity. Under a fare cap that prevents the price ratio from moving above 1.0 — a defensible operating constraint, since rail share collapses sharply once rail is priced above air — ALTO would settle at a rail-to-air price ratio of approximately 1.0 (parity with air), with the operator or government absorbing the remaining capital-cost gap.

For the corridor's three scenarios, plausible operating price ratios are: VIA at $r \approx 0.5$ (current subsidised rail), HPR at $r \approx 0.7$ (modest premium over current VIA, partial capital recovery), ALTO at $r \approx 1.0$ (price parity with air, with subsidy absorbing the larger debt-service gap from a higher capital base).

Modal share as a function of price ratio

Holding each scenario's travel time fixed at its published value, modal share becomes a one-dimensional function of price ratio. Figure 4 plots this function for each scenario on each principal Toronto pair, with markers placed at the canonical price ratio assumed for each.

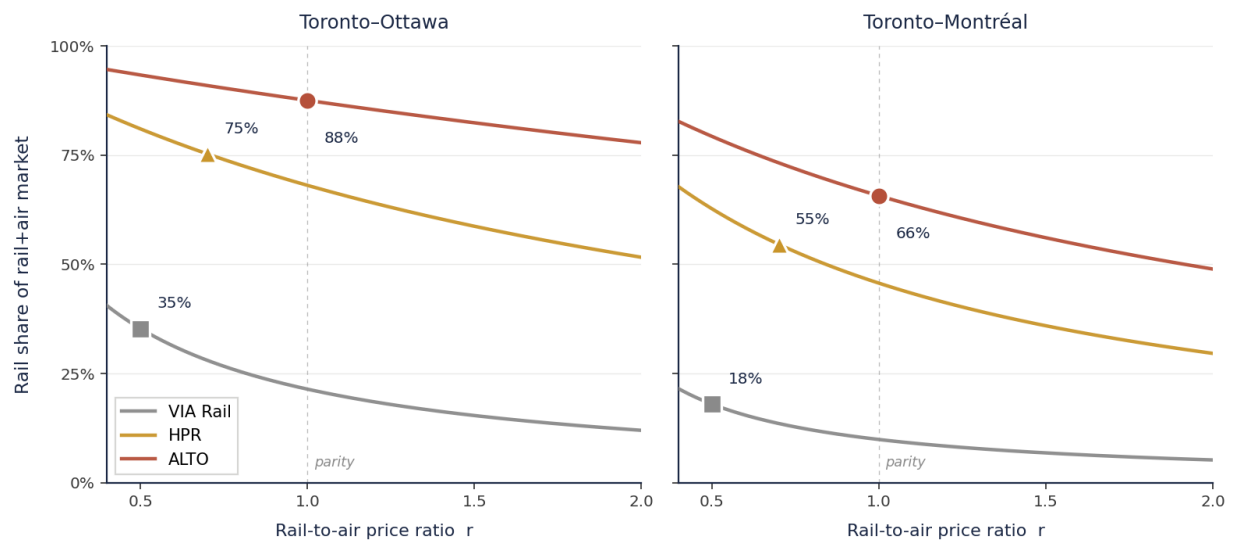


Figure 4. Modal share as a function of rail-to-air price ratio, with each scenario's travel time held fixed at its published value (VIA 4 h 30 min / 5 h 13 min, HPR 2 h 55 min / 3 h 38 min, ALTO 2 h / 3 h on Toronto–Ottawa and Toronto–Montréal respectively). Markers indicate the canonical operating price ratio for each scenario: VIA at $r = 0.5$, HPR at $r = 0.7$, ALTO at $r = 1.0$. The dashed vertical line marks rail-air price parity. The vertical separation between lines at any given price ratio shows how much modal share is driven by infrastructure (i.e., by travel time); the slope of each line shows how price-sensitive that scenario is at its operating point.

Reading the corridor pairs against the price family

On Toronto–Montréal, the three scenarios at their canonical price ratios deliver predicted rail shares of 18 per cent (VIA at $r \approx 0.5$), 55 per cent (HPR at $r \approx 0.7$), and 66 per cent (ALTO at $r \approx$

1.0). ALTO retains an 11-percentage-point modal-shift advantage over HPR — meaningful, but markedly smaller than the 20-percentage-point gap suggested by the price-parity readings in Table 2. The gap closes because ALTO's higher capital cost flows through to a higher price ratio, dragging its share down the curve, while HPR retains a price advantage that lifts its share back up. If HPR were operated at the current VIA price ratio ($r \approx 0.5$), its share would rise to roughly 63 per cent — narrowing the gap to just 3 percentage points.

On Toronto–Ottawa, both new-build scenarios sit well into the upper portion of the curve, where price effects on share are smaller. ALTO at $r \approx 1.0$ captures roughly 88 per cent; HPR at $r \approx 0.7$ captures roughly 75 per cent — a 13-percentage-point ALTO advantage. With HPR priced at $r \approx 0.5$, its share rises to about 81 per cent, narrowing the gap to 7 percentage points.

Two patterns are visible across the chart. First, on both pairs ALTO retains a real but reduced modal-shift advantage over HPR under the canonical price assumptions — between 11 and 13 percentage points, against the 20-percentage-point time-only gap from Section 3. Second, on both pairs the principal lever for closing the modal-shift gap between HPR and ALTO is the operator's pricing model, not its line-haul speed. A subsidy structure that holds HPR fares closer to current VIA levels could close most of the remaining gap; a P3 cost-recovery structure that pushed HPR fares above $r \approx 0.8$ would widen it.

The HPR pricing lever, with ALTO held at parity

A more direct way to put the HPR pricing decision in front of the reader is to fix ALTO at its operating constraint of price parity with air ($r = 1.0$) and ask how the HPR vs ALTO comparison varies as HPR's fare moves relative to ALTO's. The X-axis becomes the HPR-to-ALTO fare ratio directly. Figure 5 plots HPR's modal share alongside ALTO's fixed share for each principal pair; Figure 6 plots the resulting differential (ALTO minus HPR) on a single axis.

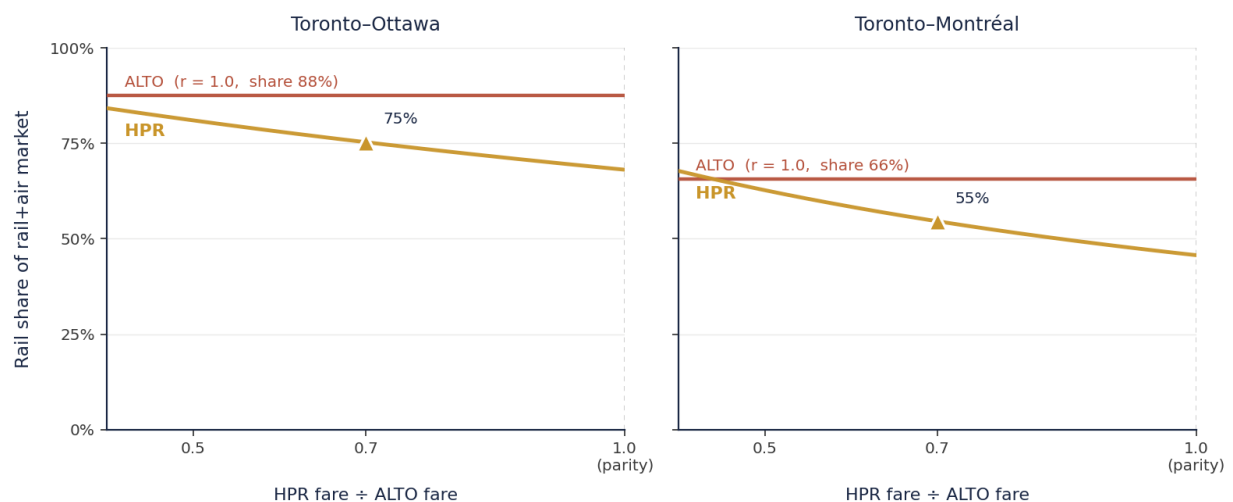


Figure 5. HPR and ALTO modal share as a function of the HPR-to-ALTO fare ratio, with ALTO held fixed at rail-air price parity ($r = 1.0$) and HPR's fare expressed as a fraction of ALTO's fare. ALTO's share appears as a flat reference

at the top of each panel; HPR's varies along the gold curve. Markers show HPR at the canonical operating point (HPR/ALTO = 0.7), placing HPR fares modestly above current VIA levels but well below ALTO and air.

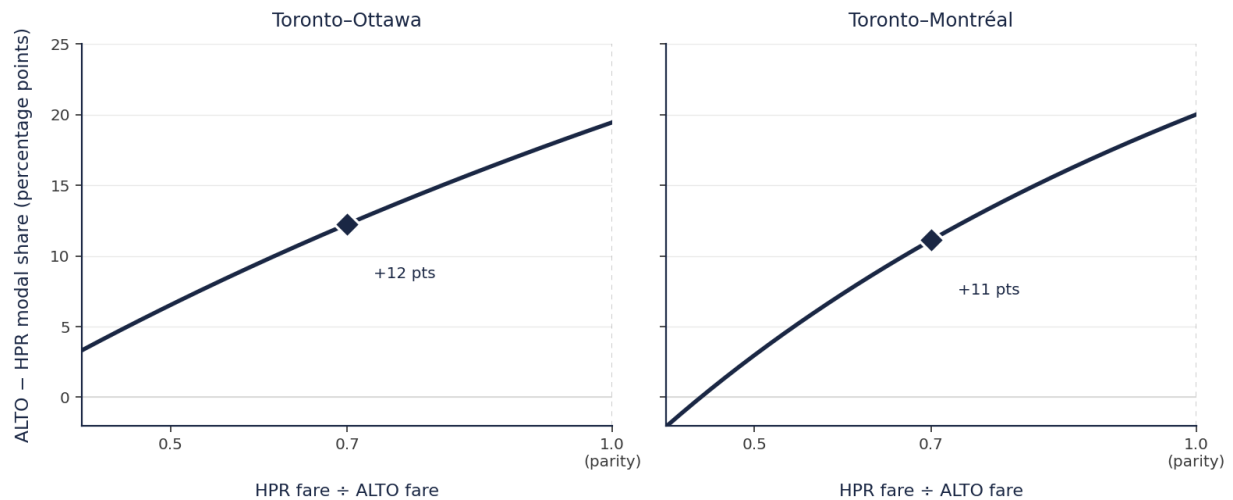


Figure 6. ALTO – HPR modal-share differential as a function of the HPR-to-ALTO fare ratio, on the same horizontal axis as Figure 5. The differential rises monotonically as HPR fares move toward ALTO's: from approximately 7 percentage points on Toronto–Ottawa and 3 on Toronto–Montréal at HPR/ALTO = 0.5, to 19 and 20 percentage points respectively at HPR/ALTO = 1.0 (rail-fare parity between the two scenarios). The diamond marker indicates the canonical HPR/ALTO = 0.7 operating point: 12 percentage points on Toronto–Ottawa and 11 on Toronto–Montréal.

The two figures make explicit what is implicit in the canonical-readings analysis above: ALTO's modal-shift advantage is highly contingent on HPR's pricing model. If HPR fares can be held to roughly half of ALTO's fares (HPR/ALTO \approx 0.5) — broadly equivalent to maintaining current VIA fare levels in real terms — the modal-shift gap between the two scenarios is in the order of 3 to 7 percentage points across the principal pairs. If HPR fares drift up to 70 per cent of ALTO's (the canonical position), the gap is 11 to 13 percentage points. If HPR fares converge with ALTO's altogether (HPR/ALTO = 1.0), the gap returns to the full 19–20-percentage-point time-only advantage that ALTO's headline case rests on. The corridor decision is therefore as much a question about HPR's intended subsidy structure as it is about the choice of infrastructure type — a question that sits inside the operator's hands rather than the engineer's.

The traveller-type and connecting-flight effects from Section 2 sharpen this further on the corridor's actual demand profile. Both Toronto–Ottawa and Toronto–Montréal carry significant business volumes, particularly at peak; lower price sensitivity in that segment means HPR's price position matters somewhat less for business travellers and somewhat more for leisure travellers. At the same time, the connecting-flight advantage that air retains is concentrated on the international business segment connecting through Pearson and Trudeau — a portion of the air market that neither HPR nor ALTO would capture in any scenario. The rail-substitutable market on the principal Toronto pairs is therefore the leisure and downtown-to-downtown business demand, which is the segment where price sensitivity is highest.

5. Where the modal-shift returns sit on the curve

The modal-shift S-curve has an important policy implication that is easily missed when ALTO is described in headline terms. Because the curve is logistic — flat at the top, steep in the middle, flat at the bottom — the value of additional travel-time savings depends critically on where on the curve a route sits to begin with. On Toronto–Montréal, moving from VIA's current 5 hour 13 minute service to HPR's published Express time of 3 hours 38 minutes crosses much of the curve's steep middle and delivers a large modal-shift gain; the further move to ALTO's 3-hour service stays in the steeper portion of the curve and adds a meaningful further increment. On Toronto–Ottawa, HPR's 2 hour 55 minute Express service places the route on the upper portion of the curve, and ALTO's 2-hour service moves it onto the upper plateau, where additional time savings produce smaller share gains.

Figure 7 decomposes the total modal-shift opportunity on each pair into the share captured by moving from VIA's current service to HPR, and the additional share captured by moving from HPR to ALTO. The HPR step does the larger part of the work on both pairs, but the size of the residual ALTO increment varies.

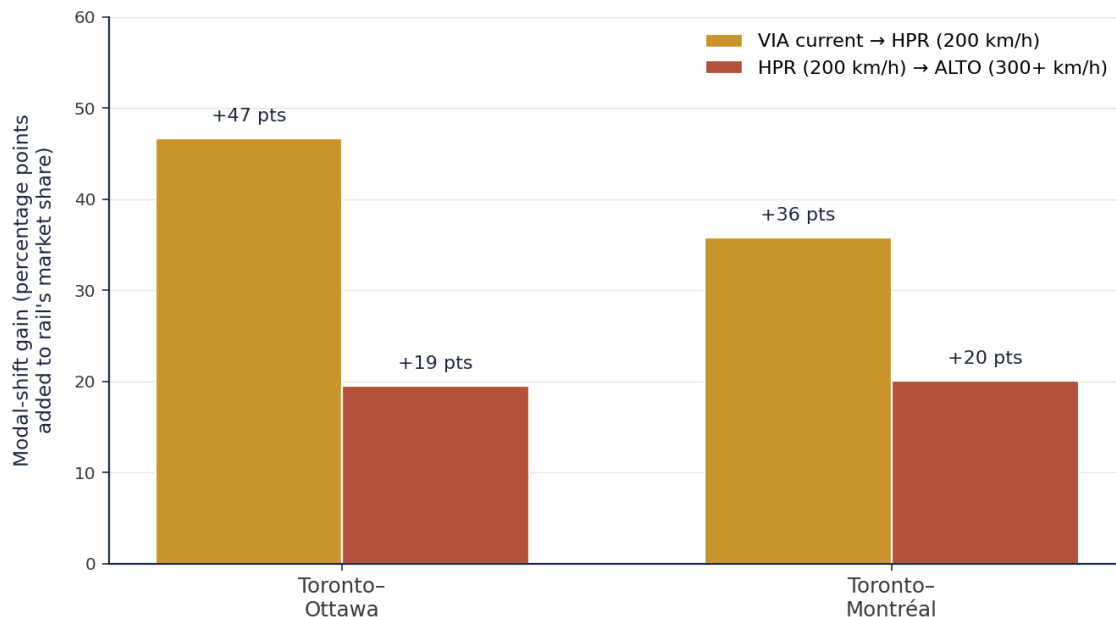


Figure 7. Decomposition of modal-shift gain by investment step on the two principal Toronto pairs. Gold bars show the percentage-point gain in rail's market share from moving from current VIA service to HPR (200 km/h dedicated corridor, with Express times from the CRI HPR Strategy: Toronto–Ottawa 2 h 55 min, Toronto–Montréal 3 h 38 min). Terracotta bars show the additional gain from moving from HPR to ALTO (300+ km/h: Toronto–Ottawa 2 h, Toronto–Montréal 3 h). At price parity, the HPR step delivers between 36 and 47 percentage points across the two pairs; the additional ALTO step delivers between 19 and 20 percentage points.

On Toronto–Ottawa, the move from VIA's current 4 hour 30 minute service to HPR's 2 hour 55 minute Express service captures an estimated 47 percentage points of modal shift; the further

move from HPR to ALTO's 2 hour service adds 19 more. On Toronto–Montréal, HPR captures 36 percentage points and ALTO adds 20 — a slightly larger residual contribution than on Toronto–Ottawa, because HPR's 3 hour 38 minute time still sits on the steeper portion of the curve. The HPR step delivers the majority of the achievable modal shift on both pairs (roughly 65 to 70 per cent of the total), but the residual available from ALTO's higher speeds is real at price parity.

HPR delivers the majority of the achievable modal shift on both Toronto pairs at price parity. ALTO's additional speed adds 19 to 20 percentage points — a residual gain that further shrinks to 11 to 13 percentage points once the canonical price assumptions from Section 4 are applied.

The cost-effectiveness comparison sharpens this picture. The CRI HPR Strategy estimates corridor capital costs in the order of \$19 million per kilometre — roughly \$19–25 billion for the full Windsor–Montréal HPR programme. ALTO's published \$60–90 billion capital envelope, before Canadian P3 cost escalation, represents an incremental investment of \$40–70 billion above the HPR option. Spread across the additional 11 to 13 percentage points of modal shift that ALTO captures over HPR under the canonical price assumptions in Section 4, this works out to roughly \$3 billion to \$6 billion per percentage point of incremental modal shift — a cost-per-point several times worse than the HPR step that precedes it. The headline 19-to-20-point time-only advantage that motivates ALTO's case shrinks substantially once realistic pricing is admitted; the cost-per-point of the remainder is the figure that the corridor decision turns on.

6. Implications for the corridor decision

Four conclusions follow from putting the international modal-shift literature, segment-level travel times, and the price dimension alongside one another.

First, the corridor's modal-shift potential is real and well-supported by international evidence. It is concentrated in two principal city pairs — Toronto–Ottawa and Toronto–Montréal — that account for the bulk of the corridor's air-substitutable demand. Modelling the corridor as a single 1,000 km market obscures this. The real question is segment-level travel time and segment-level price, not headline line-haul speed.

Second, on travel time alone — that is, holding price at parity — HPR's published Express times of 2 hours 55 minutes Toronto–Ottawa and 3 hours 38 minutes Toronto–Montréal place both principal pairs into the upper portion of the modal-shift curve. ALTO's faster published targets capture a real incremental gain of 19 to 20 percentage points on top of HPR — meaningful, but residual relative to the larger HPR step that precedes it. On a capital-cost-per-percentage-point basis, the incremental modal shift purchased by ALTO's 300+ km/h infrastructure is several times more expensive than the modal shift purchased by the HPR step.

Third, the price-shifted reading reduces ALTO's apparent advantage substantially. Each scenario operates under a different cost structure: VIA today at a rail-to-air price ratio of approximately 0.5, HPR plausibly at $r \approx 0.7$ (a modest fare premium over current VIA, with most of the new-build capital absorbed into the public account), and ALTO at $r \approx 1.0$ (price parity with air, with subsidy absorbing the larger debt-service gap from a substantially higher capital base). Under these canonical price assumptions, ALTO's modal-shift advantage on Toronto–Montréal narrows from 20 percentage points at price parity to 11 percentage points (66 per cent vs 55 per cent), and on Toronto–Ottawa from 20 percentage points to 13 (88 per cent vs 75 per cent). If HPR were operated at the current VIA price ratio ($r \approx 0.5$), the gap would close further still — to 3 percentage points on Toronto–Montréal and 7 on Toronto–Ottawa.

Fourth, this is precisely the regime in which the international literature finds frequency, reliability, station-centrality, and price to matter more than headline line-haul speed — which is the core rationale for the High Performance Rail framework. Capturing the bulk of the modal-shift opportunity does not require operating at the global frontier of high-speed rail technology. It requires getting passenger trains off shared freight track, electrifying the corridor, providing high-frequency service into central stations, delivering reliability that air carriers cannot match in winter conditions, and pricing the service competitively against air. The CRI HPR Strategy delivers all of these at roughly a quarter of ALTO's capital cost, and the lower capital cost is itself what enables the more favourable pricing position that closes most of the remaining modal-shift gap.

The corridor is thus a textbook case of why high-speed-rail claims need to be unbundled. The modal-shift opportunity is genuine. The majority of it is captured by conventional high-performance rail speeds on a dedicated, electrified, reliable corridor priced competitively against air. ALTO's additional 300+ km/h capability buys a real but reduced gain in modal shift once realistic pricing is admitted — between 11 and 13 percentage points on the principal Toronto pairs, at an incremental capital cost of \$40–70 billion. Whether the corridor decision turns on the right framework — segment-level, two-dimensional analysis of time and price — is what determines whether the public investment achieves the modal-shift outcome it is intended to produce.

§ Methodology and sources

Modelling approach

The S-curve in Figures 1, 3a–3c and 4 is a standard logistic of the form $S(t) = 1 / (1 + \exp(k \cdot (t - t_0)))$, where $S(t)$ is rail's share of the combined rail+air market as a function of station-to-station rail journey time t . The parameters used here are $k = 1.3$ and $t_0 = 3.5$ hours, calibrated by visual fit to the published international comparator data. The family of curves in Figure 2 extends this to two dimensions by adding a price-utility term: $S(t, r) = 1 / (1 + \exp(k \cdot (t - t_0) + \gamma \cdot \ln r))$, where r is the rail-to-air price ratio (rail price divided by air price) and $\gamma = 1.0$ is the price coefficient. This

binary-logit specification is the simplest defensible form of the time–price modal-choice model used routinely in transport demand work (Wardman, Train, Ben-Akiva and Lerman). More elaborate discrete-choice and stated-preference models (multinomial logit, mixed logit, nested logit) include additional regressors for frequency, station access time, reliability, and traveller demographics; these models produce more precise predictions for individual routes but tend to confirm the same overall S-shaped relationship and the same direction of the price effect. The parameters used here should be treated as illustrative rather than predictive; sensitivity analysis at k between 1.0 and 1.6, t_0 between 3.0 and 4.0 hours, and γ between 0.6 and 1.4 produces the same qualitative conclusions about HPR's modal-shift performance and ALTO's price-driven degradation of the time advantage.

Principal sources

- ALTO HSR Citizen Research Initiative (2026). HPR Strategy, Chapter 4 — High Performance Passenger Rail (Express journey times). citizenresearch.ca.
- International Council on Clean Transportation (2022). The bullet train to lower-carbon travel.
- Mineta Transportation Institute (2017). Modal Shift and High-Speed Rail: A Review of the Current Literature. P. Haas.
- World Bank Group (2019). China's High-Speed Rail Development.
- Bergantino, A. and Madio, L. (2020). Intermodal competition and substitution: HSR versus air transport. *Research in Transportation Economics*, 79.
- AECOM (2011). High-Speed Rail Overseas Experience Report. C. Nash.
- Sun, X. et al. (2024). A review on research regarding HSR interactions with air transport. *Transport Policy*, 157.
- Wardman, M. (2014). Price Elasticities of Surface Travel Demand: A Meta-analysis of UK Evidence. *Journal of Transport Economics and Policy*, 48.
- Ben-Akiva, M. and Lerman, S. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Train, K. (2009). *Discrete Choice Methods with Simulation*, 2nd ed. Cambridge University Press.
- Comisión Nacional de los Mercados y la Competencia (CNMC), annual rail market reports for Spain.
- VIA Rail Canada Annual Report 2023; VIA Rail published timetables, station-pair travel times, and Economy fare ranges.
- Alto Inc. published travel-time targets and corridor descriptions, February 2025 announcement; ALTO communications regarding self-funded P3 model and zero subsidy.
- Energies (2025). *Emission Reductions in the Aviation Sector: A Systematic Review of the Sustainability Impacts of Modal Shifts*.