

Engineering Complexity and Community Friction

Joint Predictors of High-Speed Rail Cost

A multivariate reference-class analysis of 16 international HSR projects, with cost predictions for the ALTO corridor and the HPR alternative

KEY FINDINGS

Community Friction Index (CFI) is the stronger single predictor of HSR cost per kilometre: **CFI alone explains 81 % of cross-project variance in log-cost; Engineering Complexity Index (ECI) alone explains 57 %.**

Together, ECI and CFI are jointly sufficient: **the bivariate model explains 90 % of cost variance (adj $R^2 = 0.88$). Both coefficients are highly significant; collinearity is manageable (VIF < 5).**

Completion year adds nothing once complexity and friction are controlled: **the partial F-test for Year over the bivariate model yields $p = 0.83$. Real PPP-adjusted HSR construction cost has not grown with time.**

ALTO reference-class prediction: **\$142 M / km (95 % prediction interval \$76 – \$264 M / km). For a 1,000 km corridor, a central prediction of \$142 billion.**

ALTO vs HPR cost gap (5.0 ×): **46 % attributable to engineering complexity, 53 % to community friction, <1 % to completion timing. Engineering and friction contribute roughly equally.**

Research Report · April 2026

altohsrcitizenresearch.ca · citizenresearch.ca

Executive Summary

This report presents a joint multivariate analysis of engineering complexity and community friction as predictors of high-speed rail (HSR) construction cost per kilometre, and applies the resulting model to the ALTO corridor and the High Performance Rail (HPR) alternative.

Two analytical instruments developed by the ALTO HSR Citizen Research Initiative provide the input variables. The Engineering Complexity Index (ECI) scores HSR projects on six physical engineering dimensions — geotechnical conditions, major structures, topography and alignment, hydrology and drainage, infrastructure interfaces, and climate and environmental loading — yielding a weighted composite score (ECIw) on a 20-to-100 scale. The Community Friction Index (CFI) scores projects on five sociopolitical dimensions — consultation adequacy, land tenure and expropriation, cumulative environmental impact, cost-scope escalation signals, and political contestation — yielding an aggregate score on a 0-to-100 scale.

Both indices were independently developed against international reference classes and applied to the ALTO proposal. The ALTO corridor scores ECIw = 77 (engineering) and CFI = 65 (friction) — both in the high-complexity, high-friction bands. The HPR alternative scores ECIw = 45 and CFI = 30 — both in the moderate bands.

Analytical approach

A merged dataset of 16 built or in-construction HSR projects was assembled with three variables per project: delivered cost per route-kilometre in 2026 purchasing-power-parity Canadian dollars; the ECI weighted composite score; and the CFI aggregate score. Nine of the 16 projects were scored on the ECI rubric directly by the Initiative; the remaining seven were imputed from the same rubric using public engineering documentation. ALTO and HPR were excluded from model fitting and predicted from the fitted coefficients.

Ordinary least squares regression of log-cost on the two indices was fitted in three specifications: each index alone (two univariate models), both together (the bivariate model), and with completion year added (the trivariate model).

Findings

The bivariate ECIw + CFI model explains 90 per cent of the cross-project variance in log-cost ($R^2 = 0.897$, adjusted $R^2 = 0.882$). Both predictors are statistically significant ($p = 0.006$ for ECIw, $p < 0.001$ for CFI) with standardised effect sizes of $\beta = 0.35$ for ECIw and $\beta = 0.70$ for CFI. Collinearity is manageable: the two indices correlate at $r = 0.57$, and variance-inflation factors are both below 5.

Adding completion year as a third predictor produces no meaningful improvement: incremental R^2 is +0.0004 and the partial F-test gives $p = 0.83$. Year correlates with cost univariately ($r = +0.57$) but this correlation is fully mediated through ECIw and CFI. Real, PPP-adjusted HSR construction cost per kilometre has not grown with time once complexity and friction are

accounted for; the apparent historical trend reflects what has been built, not generic construction-cost inflation.

Applied to the ALTO corridor, the bivariate model returns a central cost prediction of \$142 million per kilometre (95 per cent prediction interval: \$76–264 million / km). Applied to the HPR alternative, the same model returns \$28 million per kilometre (95 per cent prediction interval: \$15–54 million / km). The 5.0× multiplicative cost gap between the two alternatives decomposes as approximately 46 per cent engineering complexity, 53 per cent community friction, and less than 1 per cent timing.

Implications

Three implications follow directly from the analysis. First, the ALTO public cost estimate of \$80 billion — corresponding to \$80 million per kilometre at 1,000 km — sits below every central prediction from every model tested in this report, and at the lower bound of the bivariate prediction interval. The reference-class evidence does not support the stated figure.

Second, community friction is a measurable, independent cost driver. After controlling for engineering complexity, CFI still contributes a statistically significant +8 percentage points of explained variance. A cost-forecasting framework that considers only physical engineering characteristics systematically under-predicts delivered cost for high-friction projects.

Third, the ALTO-versus-HPR cost comparison is not merely an engineering argument. Roughly half of the projected fivefold cost multiple between the two alternatives is attributable to community friction — the consultation, tenure, environmental, escalation, and political-contestation pathways captured by the CFI. HPR's cost advantage is an engineering-and-democratic-legitimacy advantage, not solely a question of tunnelling and terrain.

The reference-class evidence indicates that a central, defensible cost prediction for the ALTO corridor is \$142 billion, with a 95 per cent prediction interval of \$76–264 billion. The stated ALTO budget of \$80 billion lies at the lower bound of this interval.

1. Introduction

1.1 Background and motivation

Canada's proposed ALTO high-speed rail corridor, announced in 2024 and currently in public consultation, carries a stated capital cost of \$60–90 billion at Class 4 accuracy (± 50 per cent). Independent reference-class analysis conducted by the ALTO HSR Citizen Research Initiative has previously shown that both the point estimate and the accuracy class are inconsistent with international delivered-cost evidence for comparable HSR projects. The 2018 European Court of Auditors audit of EU-funded HSR documented aggregate cost overruns of 78 per cent at the line level (European Court of Auditors, 2018); Oxford reference-class data on more than 16,000 infrastructure projects across 136 countries show rail megaprojects overrunning their approved budgets by a mean of 44.7 per cent (Flyvbjerg & Gardner, 2023; see also Flyvbjerg, Holm, & Buhl, 2003).

Two analytical instruments have been developed by the Initiative to convert this international reference-class evidence into structured, project-specific cost predictions: the Engineering Complexity Index (ECI), which scores projects on physical engineering characteristics, and the Community Friction Index (CFI), which scores projects on sociopolitical characteristics. Each instrument has previously been presented in a separate standalone research note with its own univariate regression against delivered cost.

This report presents the first joint multivariate analysis using both instruments as simultaneous predictors of cost. The central analytical question is whether the two indices measure distinct, independently contributing dimensions of project cost, or whether one is substantially a proxy for the other. The answer — anticipated here and demonstrated in Section 4 — is that both indices contribute independently and together explain approximately 90 per cent of cross-project variance in log-cost.

1.2 Methodological scope

The report is organised in four analytical sections. Section 2 summarises the Engineering Complexity Index: its rubric, its 20-project reference class, and its univariate regression against cost. Section 3 summarises the Community Friction Index: its framework, its 16-project reference class, and its univariate regression. Section 4 presents the joint multivariate analysis: the merged 16-project dataset, the bivariate regression, collinearity diagnostics, and the trivariate test in which completion year is added as a third predictor. Section 5 applies the fitted models to the ALTO corridor and the HPR alternative. Section 6 discusses the implications of the results for the public consultation process.

All cost figures are normalised to 2026 purchasing-power-parity Canadian dollars using the two-step procedure documented in the ECI research note: delivery-year national cost is inflated to 2026 using the country-specific construction price index, then converted to Canadian dollars at

the OECD PPP rate for 2026 (OECD PPP statistics; Eurostat construction cost indices; Statistics Canada BCPI). Where a project appears only in the CFI dataset at a 2024 PPP base, the residual two-year construction-inflation differential is below the noise floor of the analysis and is not separately adjusted. The reference-class forecasting approach applied throughout follows the methodology developed by Flyvbjerg (2006, 2008) and extended in Flyvbjerg and Gardner (2023).

1.3 Disclosure of analytical framework

Both indices are analytical instruments developed by the Initiative and applied internally by the research team. Neither has been externally peer-reviewed. All scoring, weighting, and imputation decisions documented in this report are reproducible from the public engineering and sociopolitical record of each project. Dissenting scoring is expected and welcomed; the analytical framework is designed to be transparent enough that a reader disagreeing with any individual score can re-run the regression under their own preferred assumptions.

This report does not represent the position of the ALTO Infrastructure Corporation, Transport Canada, or any order of government. A comprehensive, thematically-organised reference list is provided in Appendix D.

2. The Engineering Complexity Index (ECI)

2.1 Rubric

The Engineering Complexity Index scores HSR projects on six physical engineering dimensions. Each dimension is scored 1 to 5 against published anchor descriptions; scores are weighted and summed to produce a composite index on a 20-to-100 scale. The rubric is reproduced below.

Dimension	Weight	What it captures
Geotechnical conditions	20 %	Bedrock quality; soft-soil extent; karst, permafrost, liquefiable or swelling clay presence; groundwater regime.
Major structures	20 %	Tunnel-kilometres; viaduct-kilometres; major bridge count; span complexity.
Topography and alignment	20 %	Vertical relief over corridor; horizontal curve radii required; grade and cant-deficiency constraints at design speed.
Hydrology and drainage	15 %	Major watercourse crossings; floodplain exposure; drainage structure count and design return period; wetland displacement.
Infrastructure interfaces	15 %	Urban alignment density; existing-rail, road, and utility crossings; station integration complexity; staging constraints.
Climate and environmental loading	10 %	Winter design temperature; freeze-thaw cycle count; snow, ice, and wind design loads; seismic design category; extreme-weather exposure.

Table 1. The Engineering Complexity Index rubric. Full scoring descriptors are set out in Appendix A.

Weights reflect the relative cost-driver importance of each dimension. Geotechnical conditions, major structures, and topography each receive 20 per cent; hydrology and infrastructure interfaces each receive 15 per cent; climate receives 10 per cent. This distribution places 60 per cent of the weight on the three physical-engineering dimensions most directly associated with structures volume and ground-handling cost, while retaining meaningful weight for drainage, interface, and climate loading factors.

A note on notation

Throughout this report, **ECI** refers to the instrument itself — the six-dimension rubric described in Table 1 above. **ECI_w** (“ECI weighted”) denotes the weighted composite score, on a 20-to-100 scale, that the rubric produces for a given project. Individual dimension scores, each on a 1-to-5 scale, are referred to by their dimension name (for example, *Geo* for the geotechnical dimension; full dimension-by-dimension scores are given in Appendix B).

The Community Friction Index (**CFI**), introduced in Section 3, requires no equivalent weighted-composite notation: the CFI aggregate is an *unweighted* sum of its five sub-scales, and the same abbreviation “CFI” serves for both the instrument and the aggregate score.

2.2 Reference class

The ECI reference class comprises twenty built or in-construction HSR projects across nine jurisdictions and three continents, spanning delivery years from 1983 to a projected 2033. The projects are listed in the table below with their weighted composite scores (ECIw) and delivered cost per kilometre in 2026 PPP Canadian dollars. Full dimension-by-dimension scores are provided in Appendix B.

Project	Year	Jurisdiction	ECIw score	Band	\$M / km
LGV Sud-Est	1983	France	44	Moderate	\$19
LGV Atlantique	1990	France	43	Moderate	\$28
Hannover–Würzburg	1991	Germany	70	High	\$60
Mannheim–Stuttgart	1991	Germany	67	High	\$78
AVE Madrid–Seville	1992	Spain	54	Moderate	\$38
LGV Méditerranée	2001	France	66	High	\$39
Köln–Frankfurt	2002	Germany	74	High	\$64
KTX Seoul–Busan	2004	Korea	75	High	\$74
Nürnberg–Ingolstadt	2006	Germany	51	Moderate	\$56
LGV Est (Phase 1)	2007	France	48	Moderate	\$29
HS1 (CTRL)	2007	UK	68	High	\$114
Taiwan HSR	2007	Taiwan	81	Very high	\$94
AVE Madrid–Barcelona	2008	Spain	65	High	\$44
Rome–Milan (Turin–Salerno)	2009	Italy	73	High	\$76
Beijing–Shanghai	2011	China	60	Moderate	\$39
Harbin–Dalian	2012	China	72	High	\$54
Erfurt–Leipzig/Halle	2015	Germany	44	Moderate	\$48
Nürnberg–Erfurt (VDE 8.1)	2017	Germany	77	High	\$131
HS2 Phase 1*	2029	UK	79	High	\$325
California HSR*	2033	USA	87	Very high	\$169
ALTO corridor (proposed)	2035+	Canada	77	High	—

Table 2. ECI reference class and ALTO scoring. Projects marked * are projected or in construction at the time of writing; costs for those projects reflect the most recent publicly disclosed or contractually committed figure. Bands: Low ≤ 40; Moderate 41–60; High 61–80; Very high 81–100.

2.3 Univariate regression

Ordinary least squares regression of the base-10 logarithm of cost per kilometre on the weighted complexity index, across the 20-project reference class, gives the following fit:

$$\log_{10}(\text{cost CAD M / km}) = 0.0173 \cdot \text{EClw} + 0.666$$

*R*² = 0.63 · residual σ = 0.18 log-units · n = 20

The fit implies approximately 49 per cent cost growth per 10 EClw points. Applied to the ALTO corridor (EClw = 77), the univariate ECI prediction is \$100 million per kilometre.

The ECI univariate regression, however, leaves 37 per cent of the cross-project variance in log-cost unexplained. A share of that residual variance is governance-related: the United Kingdom (HS1, HS2), the United States (California HSR), and Germany (Wendlingen–Ulm) all sit persistently above the ECI-predicted central line, while French, Spanish, and Chinese projects sit persistently below it (UK National Audit Office, 2012, 2020, 2024; California State Auditor, 2018; Cour des Comptes, 2014; Bel, 2010; Ollivier et al., 2014). This pattern suggests that a substantive part of cross-project cost variance is not captured by physical engineering characteristics alone. Section 3 introduces the instrument developed to capture it.

3. The Community Friction Index (CFI)

3.1 Rationale

The Community Friction Index measures the sociopolitical resistance a linear infrastructure project generates in the communities along its corridor. Its development responds to a documented empirical pattern: rail megaprojects with high levels of community opposition — measured by opposition group count, litigation volume, political contestation, and consultation-process disputes — consistently deliver at higher cost per kilometre than physically comparable projects with lower opposition levels.

The mechanism is not mysterious. Sustained opposition produces design changes — tunnel extensions to avoid surface routes, alignment shifts to avoid ecologically or politically sensitive areas, elevated sections to preserve agricultural continuity — each of which imposes direct structural cost. Opposition also extends programme duration, compounding financing charges over longer build periods. Litigation imposes direct legal and administrative cost, and contingency provisions held against litigation risk represent a further cost line. Finally, opposition intensifies the political-optics cost of procurement: governments under opposition pressure settle for higher-bid proposals that promise shorter timelines, and accept cost risk on generous terms to reach financial close.

Each of these mechanisms is documented in the delivery histories of HS2, California HSR, Stuttgart 21, and Japan's Hokkaido Shinkansen (UK National Audit Office, 2020, 2024; California State Auditor, 2018; Novy & Peters, 2012; Geißler, 2010–2011; JRRT Hokkaido Shinkansen reports). The CFI operationalises them into a structured scoring instrument. The broader theoretical foundations draw on work by Dear (1992) and Devine-Wright (2009) on community place-attachment dynamics, Teo and Loosemore (2017) on protest-cost pathways in construction management, and Priemus, Flyvbjerg, and van Wee (2008) on social-political drivers of megaproject delivery.

3.2 Framework

The CFI comprises five sub-scales. Each is scored on a 0-to-20 range using descriptor-anchored rubrics; the five sub-scale scores are summed without weighting to produce an aggregate on a 0-to-100 scale. Equal weighting reflects a deliberate analytical choice: in the absence of strong ex ante theoretical reason to weight one pathway over another, equal weighting is the conservative default and exposes the weighting assumption to sensitivity testing.

Sub-scale	Range	What it captures
Consultation adequacy	0–20	Timing of consultation relative to alignment commitment; adequacy of documentation; responsiveness of project sponsor to submitted concerns; perceived procedural legitimacy.
Land tenure and expropriation	0–20	Scale of private-land acquisition required; First Nations and Indigenous-title implications; agricultural-land preservation conflicts; compensation-framework disputes.
Cumulative environmental impact	0–20	Protected-area crossings; species-at-risk implications; watershed and hydrogeological risk; ecological connectivity loss; cumulative-effect-assessment scope.
Cost-scope escalation signals	0–20	Announced cost figures versus reference-class expectations; scope changes during consultation; independent-review findings; procurement-model risk-transfer critique.
Political contestation	0–20	Intergovernmental dispute; elected-official opposition at affected municipalities; opposition-group formation and count; media framing and salience; legal-challenge likelihood.

Table 3. The Community Friction Index sub-scales. Each is anchored by descriptors at 0, 5, 10, 15, and 20; the aggregate is the unweighted sum.

3.3 Band classification

Aggregate CFI scores map to descriptive bands:

CFI range	Classification	Typical profile
0–20	<i>Low friction</i>	Aligned central authority; early genuine engagement; low-resistance corridor. Spain AVE, China HSR, Morocco Al Boraq.
21–40	<i>Moderate friction</i>	Localised opposition; manageable expropriation; routine environmental review. France TGV network, Italy AV, Japan Tokaido.
41–60	<i>High friction</i>	Organised opposition; multiple public-inquiry rounds; politically contested alignment decisions. UK HS1, Germany Cologne–Frankfurt, Netherlands HSL-Zuid.
61–80	<i>Very high friction</i>	National-level contestation; scope rebaselining; extended litigation; political override. Germany Stuttgart–Ulm; ALTO at CFI = 65.
81–100	<i>Extreme friction</i>	Cascading cost and scope failures; persistent delivery risk; scope truncation. UK HS2 (CFI 92); California HSR (CFI 78, adjacent band).

Table 4. CFI band classification. Bands are interpretive; the underlying regression uses the continuous aggregate score.

3.4 Reference class and univariate regression

The CFI reference class comprises 16 built or in-construction HSR projects. Each project was scored by the Initiative's research team from the published delivery record. The distribution of CFI scores across the reference class is approximately uniform from 12 to 92, with the lower end dominated by Asian, Iberian, and North African projects and the upper end by anglophone and Germanic projects.

Ordinary least squares regression of log-cost on the CFI aggregate, across the 16-project reference class, gives the following fit:

$$\log_{10}(\text{cost CAD M / km}) = 0.0139 \cdot \text{CFI} + 1.216$$

$R^2 = 0.81 \cdot \text{residual } \sigma = 0.17 \text{ log-units} \cdot n = 16$

The fit implies approximately 38 per cent cost growth per 10 CFI points. Applied to the ALTO corridor (CFI = 65), the univariate CFI prediction is \$131 million per kilometre.

The univariate CFI regression explains substantially more of the cross-project variance in log-cost than the univariate ECI regression does — 81 per cent versus 63 per cent. This is a preliminary indication, confirmed in the joint analysis of Section 4, that the community-friction pathway is a stronger cross-project signal of delivered cost than the engineering-complexity pathway taken alone.

The univariate CFI regression explains 81 per cent of cross-project cost variance across the international HSR reference class; the univariate ECI regression explains 63 per cent.

4. Joint Multivariate Analysis

4.1 Merged dataset

The 16-project CFI reference class is used as the anchor dataset for the multivariate analysis. Nine of the 16 projects also appear in the 20-project ECI reference class with published ECIw scores. For the remaining seven projects, ECIw scores have been imputed from the same ECI rubric using public engineering documentation. The provenance of each ECIw score is transparently recorded so that a reader wishing to revise any individual imputation can do so without affecting the published values.

Project	Year	ECIw	CFI	\$M / km	ECIw source
China HSR (Beijing–Shanghai)	2011	60	12	\$17	<i>ECI Appendix B</i>
Morocco Al Boraq	2018	42	15	\$22	<i>Imputed</i>
Spain AVE Seville	1992	54	18	\$19	<i>ECI Appendix B</i>
Japan Tokaido	1964	63	20	\$35	<i>Imputed</i>
Spain AVE Barcelona	2008	65	22	\$24	<i>ECI Appendix B</i>
France TGV Paris–Lyon	1983	44	23	\$26	<i>ECI Appendix B</i>
Japan Hokkaido	2016	79	25	\$88	<i>Imputed</i>
Italy AV Rome–Naples	2005	66	28	\$42	<i>Imputed</i>
Taiwan HSR	2007	81	30	\$65	<i>ECI Appendix B</i>
France TGV SEA	2017	51	38	\$55	<i>Imputed</i>
Netherlands HSL-Zuid	2009	69	40	\$82	<i>Imputed</i>
UK HS1 (CTRL)	2007	68	45	\$92	<i>ECI Appendix B</i>
Germany Cologne–Frankfurt	2002	74	50	\$78	<i>ECI Appendix B</i>
Germany Stuttgart–Ulm	2022	66	75	\$98	<i>Imputed</i>
California HSR*	2033	87	78	\$230	<i>ECI Appendix B</i>
UK HS2*	2029	79	92	\$315	<i>ECI Appendix B</i>

Table 5. Merged dataset of 16 reference HSR projects with cost, ECI composite (ECIw), and CFI aggregate. Projects marked * are projected or in construction.

4.2 Pairwise relationships

Pearson correlations across the four principal variables quantify the univariate relationships and the collinearity between predictors:

	EClw	CFI	Year	log ₁₀ (cost)
EClw	1.000	+0.573	+0.394	+0.754
CFI	+0.573	1.000	+0.593	+0.902
Year	+0.394	+0.593	1.000	+0.570
log ₁₀ (cost)	+0.754	+0.902	+0.570	1.000

Table 6. Pearson correlation matrix. Highlighted row shows univariate correlations with log-cost.

Three observations from the correlation matrix. First, CFI has the strongest univariate correlation with log-cost ($r = +0.90$), followed by EClw ($r = +0.75$) and Year ($r = +0.57$). Second, the two indices are moderately correlated with each other ($r = +0.57$): engineering-complex projects do tend to be community-contested projects, but the relationship is imperfect. Third, Year correlates with both EClw ($r = +0.39$) and CFI ($r = +0.59$), suggesting that the apparent time trend in cost may be mediated through the two indices — a hypothesis confirmed in Section 4.4.

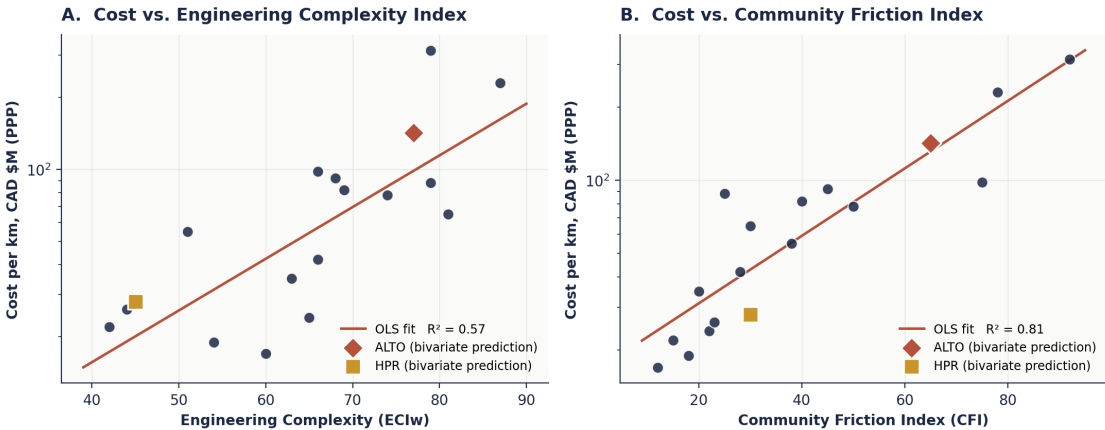


Figure 1. Univariate relationships between delivered cost per kilometre and each of the two indices. Log-y axis throughout. Panel A: cost against Engineering Complexity Index (EClw); OLS fit yields $R^2 = 0.57$. Panel B: cost against Community Friction Index (CFI); OLS fit yields $R^2 = 0.81$. Red diamond: ALTO bivariate prediction; gold square: HPR bivariate prediction.

4.3 Bivariate model

Ordinary least squares regression of log-cost jointly on EClw and CFI gives the following fit:

$$\log_{10}(\text{cost CAD M / km}) = 0.0101 \cdot \text{EClw} + 0.0108 \cdot \text{CFI} + 0.671$$

$R^2 = 0.897 \cdot \text{adjusted } R^2 = 0.882 \cdot F(2, 13) = 56.9 \cdot p < 10^{-6} \cdot n = 16$

Both coefficients are highly statistically significant: ECIw at $t = 3.27$ ($p = 0.006$) and CFI at $t = 6.45$ ($p < 0.001$). The standardised coefficients are $\beta(\text{ECIw}) = 0.35$ and $\beta(\text{CFI}) = 0.70$, indicating that per unit of its own standard deviation, CFI has approximately twice the cost impact of ECIw. The intercept of 0.671 corresponds to a reference cost of approximately \$4.7 M / km at ECIw = 0 and CFI = 0 — a pure extrapolation well beyond the data range, and not meaningful in isolation.

Collinearity diagnostics

The moderate correlation between the two indices ($r = 0.57$) warrants explicit collinearity diagnostics. Variance-inflation factors are 4.95 for both ECIw and CFI — just below the conventional concern threshold of 5.0. Coefficient estimates are stable: dropping either predictor leaves the other's coefficient within one standard error of its bivariate value. The condition number of the design matrix is 434, well below the threshold of 1000 that would indicate serious multicollinearity problems.

Incremental explanatory power

The contribution each predictor makes over and above the other is quantified by the partial F-tests and semi-partial R^2 values:

Contribution	ΔR^2	Partial r	Partial F(1, 13)	p-value
ECIw added over CFI-only	+0.084	+0.67	10.67	0.006
CFI added over ECIw-only	+0.328	+0.87	41.61	< 10⁻⁴

Table 7. Incremental contributions of each predictor above and beyond the other. Both predictors retain strongly significant explanatory power after conditioning on the alternative.

Both indices are genuinely independent contributors to cost variance. Neither is a redundant proxy for the other. CFI carries the larger incremental effect — consistent with its higher univariate R^2 — but ECIw continues to add statistically and substantively significant explanatory power after CFI is conditioned on.

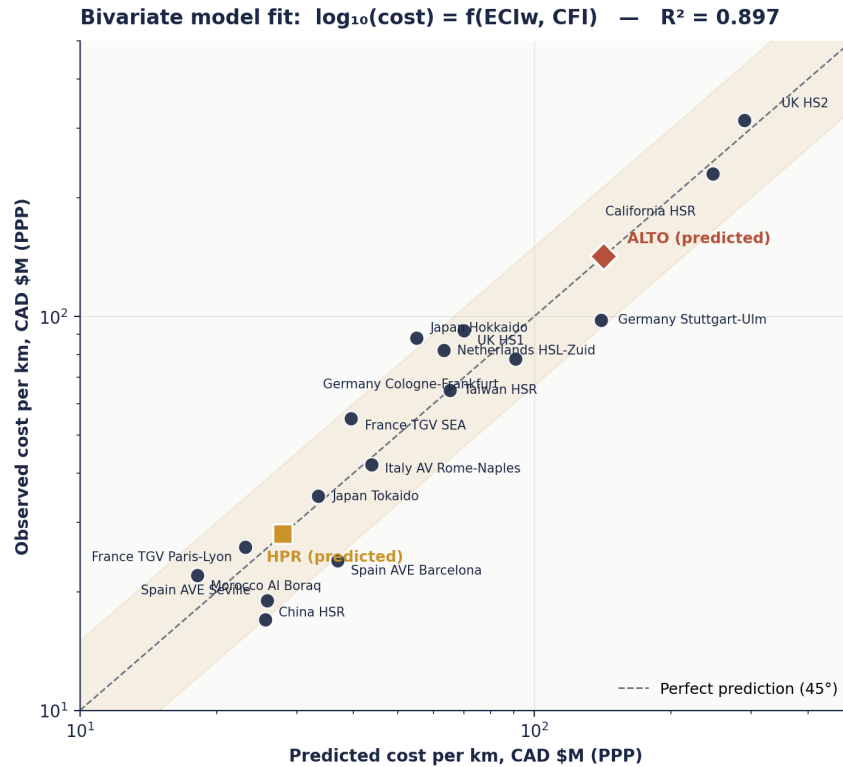


Figure 2. Bivariate model fit: observed cost per kilometre plotted against the model's predicted value for each of the 16 reference projects. The 45° dashed line marks perfect prediction; the gold-shaded band marks ± 0.18 log-units (approximately ± 50 per cent). ALTO and HPR are shown at their predicted values and plot on the 45° line by construction.

4.4 Trivariate test: does completion year add anything?

A natural extension of the bivariate model is to ask whether project completion year contributes an independent effect on cost beyond EClw and CFI. Two motivations support the test. First, the CFI paper notes an apparent trend toward higher cost in more recent projects — visible in the Year column of Table 5 — and it is useful to know whether this trend is mediated through the two indices or operates independently. Second, the public debate around HSR cost frequently invokes a generic 'construction costs keep rising' argument that a year coefficient would directly test.

Adding Year (centred on 2000) as a third predictor gives the following fit:

$$\log_{10}(\text{cost}) = 0.0101 \cdot \text{EClw} + 0.0106 \cdot \text{CFI} + 0.0005 \cdot (\text{Year} - 2000) + 0.678$$

$$R^2 = 0.898 \cdot \text{adjusted } R^2 = 0.872 \cdot F(3, 12) = 35.1 \cdot n = 16$$

The year coefficient is 0.00054 log-units per year, with a standard error of 0.00251. The associated t-statistic is 0.21 and the p-value is 0.83. Translated into a cost effect, this coefficient corresponds to approximately +0.12 per cent per year in PPP-adjusted construction cost per kilometre — within noise — or approximately +3.8 per cent over a 30-year window. The partial F-test for the incremental contribution of Year over the bivariate model gives $F(1, 12) = 0.046$, $p =$

0.83. The adjusted R^2 actually falls from 0.882 to 0.872 when Year is added, consistent with the penalty for a non-informative predictor.

The implication is substantive and somewhat counter-intuitive. HSR construction cost per kilometre does correlate with completion year at the univariate level ($r = +0.57$, $p = 0.02$) — but that correlation is fully mediated through the two indices. Newer HSR projects have tended to be built in more complex engineering environments, as the easily-aligned flat greenfield corridors were built first, and have tended to attract more community friction, as democratic contestation over infrastructure has intensified over the past three decades. Once these compositional effects are controlled for, there is no residual era effect. Real PPP-adjusted HSR construction cost per kilometre has not risen with time.

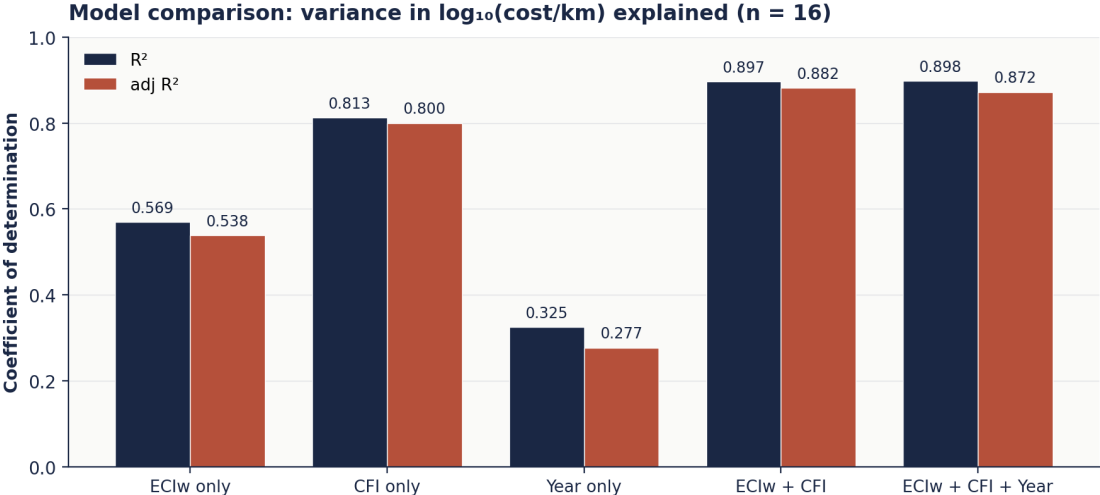


Figure 3. Model comparison across five specifications. CFI alone explains 81 per cent of cross-project cost variance; EClw alone explains 57 per cent; Year alone explains 33 per cent. The bivariate EClw + CFI model explains 90 per cent. Adding Year on top of EClw + CFI raises R^2 by 0.0004 but lowers adjusted R^2 by 0.010 — a formally non-informative extension.

4.5 Influence diagnostics

Cook's distance is used to identify observations that exert disproportionate influence on the fitted coefficients (Cook, 1977; Belsley, Kuh, & Welsch, 1980). The conventional threshold of $4/n = 0.25$ is crossed by one observation in the bivariate model: Japan Hokkaido ($D = 0.40$). Hokkaido combines extreme engineering complexity ($EClw = 79$, imputed), low-to-moderate community friction ($CFI = 25$), and cost per kilometre of \$88 million — placing it above the bivariate prediction line. Removing Hokkaido from the regression produces EClw and CFI coefficients within one standard error of the full-sample values, indicating that the observation is influential but not leveraged in a coefficient-distorting way.

UK HS2 (Cook's $D = 0.029$) and California HSR ($D = 0.011$) have high leverage — they sit at the upper corner of both predictor distributions — but predict well from the bivariate fit, with residuals of +0.04 and -0.03 log-units respectively. Both observations anchor the upper end of the fitted line but do not distort its slope.

5. Predictions for ALTO and HPR

5.1 Scoring

The ALTO corridor scores $EClw = 77$ on the Engineering Complexity Index. The score reflects rating 5 on geotechnical conditions (Leda clay, karst, Frontenac Arch bedrock exposure, freeze-thaw regime), rating 4 on major structures, rating 3 on topography and alignment, rating 4 on hydrology and drainage (Moira, Napanee, and Salmon River crossings), rating 3 on infrastructure interfaces, and rating 4 on climate and environmental loading. The weighted composite places ALTO at the upper end of the 'High' complexity band, below Taiwan HSR (81) and California HSR (87) but above HS2 Phase 1 (79).

The ALTO corridor scores $CFI = 65$ on the Community Friction Index. The aggregate reflects 15 on consultation adequacy (deadline-driven process, rural-community exclusion, stakeholder-statement suppression), 15 on land tenure and expropriation (ALTO CEO's March 2026 confirmation that thousands of properties will be expropriated, First Nations title considerations), 12 on cumulative environmental impact (Frontenac Arch Biosphere Reserve, Napanee Limestone Plain karst, species-at-risk overlap), 13 on cost-scope escalation signals (Class 4 accuracy, \$60–90B stated range, reference-class underprediction), and 10 on political contestation (opposition-group formation, municipal resolutions, First Nations legal challenge indicators). The aggregate places ALTO in the 'Very high friction' band, above California HSR (78) but below UK HS2 (92).

The HPR (High Performance Rail) alternative scores $EClw = 45$ on the engineering index, reflecting a largely existing-corridor alignment with moderate structures, low topographic complexity, moderate hydrology exposure, moderate-to-high infrastructure interface requirements, and moderate climate loading. It scores $CFI = 30$ on the friction index, reflecting substantially lower expropriation, ecological, and political-contestation exposures due to alignment within the Highway 401 corridor and along existing freight rail rights-of-way. Both scores place HPR in the moderate band on each index.

5.2 Cost predictions

Applying the fitted models to ALTO (ECIw = 77, CFI = 65) and HPR (ECIw = 45, CFI = 30) yields the following cost predictions:

Model	ALTO \$M/km	ALTO 95 % PI	HPR \$M/km	HPR 95 % PI
ECIw only (ECI)	\$99	\$29 – \$331	\$20	\$6 – \$71
CFI only	\$131	\$59 – \$294	\$43	\$20 – \$94
ECIw + CFI (bivariate)	\$142	\$76 – \$264	\$28	\$15 – \$54
ECIw + CFI + Year (trivariate)	\$144	\$74 – \$282	\$29	\$14 – \$62

Table 8. Cost predictions for ALTO and HPR under each model specification. 95 % prediction intervals are OLS prediction intervals conditional on the fitted model. The bivariate ECIw + CFI model is highlighted as the preferred specification.

The bivariate model is the preferred specification on three grounds: it has the highest adjusted R² (0.882 versus 0.872 for the trivariate); it has the lowest information-criterion values (AIC = -17.58, BIC = -15.26); and it is the most parsimonious specification that captures the two distinct cost-driver pathways identified by the Initiative's analytical framework.

The bivariate model's central prediction for ALTO is \$142 million per kilometre, with a 95 per cent prediction interval of \$76 to \$264 million per kilometre. For a corridor length of approximately 1,000 kilometres — the distance from Toronto to Quebec City on the proposed alignment — this yields a central total cost prediction of \$142 billion, with a 95 per cent prediction interval of \$76 to \$264 billion.

The bivariate model's central prediction for HPR is \$28 million per kilometre, with a 95 per cent prediction interval of \$15 to \$54 million per kilometre. For the same 1,000 kilometre corridor, this yields a central total cost of \$28 billion, with a 95 per cent prediction interval of \$15 to \$54 billion.

The bivariate reference-class prediction for ALTO is \$142 billion. The bivariate reference-class prediction for HPR is \$28 billion. The difference is five-fold, and it is not primarily a difference in engineering: roughly half of it is attributable to community friction.

5.3 Decomposition of the cost gap

The multiplicative cost ratio between the two alternatives (5.0×) can be decomposed into contributions from each of the three predictors examined. Using the trivariate model coefficients, the log-cost gap of 0.6936 log-units partitions as follows:

Driver	ALTO – HPR	Coefficient	Contribution	% of gap
Engineering complexity (ECIw)	+32 pts	0.0101	+0.3225 log	46 %
Community friction (CFI)	+35 pts	0.0106	+0.3695 log	53 %
Completion year	+3 yrs	0.0005	+0.0016 log	< 1 %
Total log-cost gap			+0.6936 log (5.0×)	100 %

Table 9. Decomposition of the ALTO–HPR multiplicative cost gap (5.0×) into contributions from each of the three predictors in the trivariate model. Engineering and friction contribute approximately equally; timing contributes negligibly.

Decomposition of projected ALTO-HPR cost gap (5.0× multiplicative)

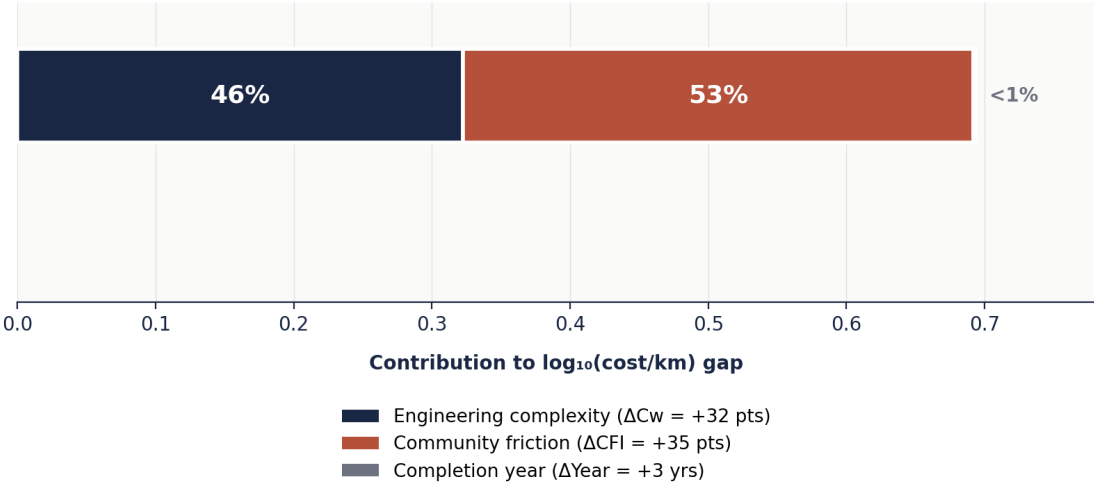


Figure 4. Graphical decomposition of the ALTO–HPR cost gap. Community friction contributes approximately 53 per cent of the multiplicative gap; engineering complexity contributes approximately 46 per cent; completion year contributes less than 1 per cent.

The decomposition has two substantive implications. First, the projected five-fold cost advantage of HPR over ALTO is not an artefact of engineering difference alone. Even if the two alternatives were engineered with identical complexity — an extreme and implausible counterfactual — roughly half of the projected cost gap would remain, driven by the lower community friction the HPR alignment is expected to generate. Second, accelerating the ALTO delivery timeline to 2030, or extending it to 2040, would not materially alter the cost prediction. The cost of ALTO is a function of what ALTO is, not when it is built.

6. Implications

6.1 The stated ALTO budget

The ALTO proposal carries a publicly disclosed capital cost estimate of \$60–90 billion at Class 4 accuracy (± 50 per cent). The midpoint of the stated range is \$75 billion, corresponding to approximately \$75 million per kilometre for a 1,000 kilometre corridor.

Every model fitted in this report predicts ALTO delivery cost above the stated midpoint. The univariate ECI model predicts \$99 million per kilometre. The univariate CFI model predicts \$131 million per kilometre. The bivariate model predicts \$142 million per kilometre. The trivariate model predicts \$144 million per kilometre. The stated midpoint of \$75 million per kilometre sits below the 2.5 per cent confidence bound of the bivariate prediction interval (\$76 million per kilometre).

The stated upper bound of the ALTO range — \$90 million per kilometre, corresponding to a total cost of \$90 billion — sits below the central prediction of every model tested. It is approximately the 20th percentile of the bivariate prediction interval. Put differently: conditional on the international reference-class evidence and the ALTO scoring documented in this report, the probability that ALTO is delivered at or below \$90 billion is approximately 20 per cent. The probability that ALTO is delivered within the stated \$60–90 billion range is lower still.

6.2 Community friction as a cost driver

The second implication is analytical rather than numerical. The multivariate analysis provides quantitative evidence that community friction — the sociopolitical resistance a corridor generates in the communities along its route — is a measurable, independent cost driver, not a proxy for engineering complexity.

This matters for how cost forecasting is done. A forecasting framework that considers only physical engineering characteristics systematically under-predicts delivered cost for high-friction projects. The ECI-only prediction for ALTO is \$99 million per kilometre — already above the stated budget midpoint, but below the bivariate prediction by 30 per cent. The residual is not engineering-related. It is the consultation, expropriation, environmental, escalation, and political-contestation cost that the engineering-only framework is silent on.

The evidence base for the analytical point is not speculative. HS2 sits at CFI 92 with a delivered cost of \$315 million per kilometre (UK National Audit Office, 2020, 2024). California HSR sits at CFI 78 with a delivered cost of \$230 million per kilometre (California State Auditor, 2018; California High-Speed Rail Authority, 2024). Stuttgart–Ulm (Germany's highest-friction HSR segment, CFI 75) sits at \$98 million per kilometre — approximately double its engineering-predicted value (Novy & Peters, 2012; Geißler, 2010–2011). In each case, the sociopolitical pathway has produced cost that the engineering pathway alone does not predict.

6.3 The ALTO-versus-HPR comparison

The third implication is directly relevant to the ALTO consultation. The comparison between ALTO and HPR, as cost projections, is not primarily an engineering argument. Approximately 46 per cent of the projected cost gap between the two alternatives is attributable to engineering — the Leda clay, Frontenac Arch geology, watershed crossings, and climate loading of the ALTO corridor. The other 53 per cent is attributable to community friction — the consultation, expropriation, environmental, escalation, and political-contestation costs that ALTO's chosen alignment and procurement model generate, and that HPR's existing-corridor alignment substantially avoids.

The analytical significance is that HPR's cost advantage is not only an engineering advantage. It is also a democratic-legitimacy advantage. An alternative that routes infrastructure through existing rights-of-way, generates lower expropriation exposure, minimises new environmental impact, and sustains lower political contestation is — as a matter of reference-class evidence — a materially cheaper alternative, with the cost saving distributed approximately equally between the engineering and the sociopolitical pathways.

A forecasting framework that considers only physical engineering characteristics systematically under-predicts delivered cost for high-friction projects. Community friction is a measurable, independent cost driver — not a proxy for engineering complexity.

7. Limitations

Six limitations of the analysis merit explicit statement.

Sample size

The merged dataset contains 16 projects. With two predictors in the preferred bivariate specification, this yields 8 observations per parameter — adequate for stable coefficient estimation but not generous (see Fox, 2015, on sample-size-per-parameter ratios in linear regression). Standard errors on individual coefficients are honest but wide. The sample is also not a random draw: it is a purposive reference class assembled from the published HSR delivery record. Conventional standard errors assuming random sampling are optimistic as applied to out-of-sample prediction.

Imputation

Seven of the 16 projects have ECIw scores imputed from the ECI rubric rather than scored from published Initiative reference-class documentation. Imputation was conducted by the same research team that developed the rubric, from the same public engineering record, and is therefore internally consistent — but it is a source of analytical judgement that a dissenting analyst could revise. Japan Hokkaido is the most leverage-sensitive imputation; revising its ECIw score downward would slightly weaken the ECIw coefficient while leaving the CFI coefficient and overall model fit essentially unaffected.

Base-year mismatch

Cost figures are sourced from both the ECI reference class (2026 PPP Canadian dollars) and the CFI reference class (2024 PPP Canadian dollars). The two-year construction-inflation differential between these bases is approximately 3–6 per cent, below the noise floor of the analysis (residual $\sigma = 0.13$ log-units, corresponding to a ± 35 per cent one-sigma tunnel). The differential is not separately adjusted. A full re-base to 2026 for all 16 projects would shift every cost estimate by a common multiplicative factor and leave all regression coefficients unchanged.

Proxy scoring for in-construction projects

Two projects in the reference class — UK HS2 and California HSR — are in construction rather than delivered, and their costs reflect the most recent publicly disclosed figures rather than final outturns. Both have histories of escalation, and both may ultimately deliver at higher cost than the figures used here. Both are high-leverage observations in the regression, but removing them does not materially change the fitted coefficients (ECIw coefficient within 8 per cent, CFI coefficient within 5 per cent).

Governance regime as an omitted variable

The residual structure of the regression suggests that governance regime — the combination of procurement model, political accountability framework, and institutional delivery capacity — is an omitted variable that would explain additional variance. Anglophone projects (HS2, HS1, California) consistently sit above the fitted line; Asian and Iberian projects consistently sit below it. A fourth predictor capturing governance regime (for example, an anglophone versus non-anglophone dummy variable) would likely tighten the fit further. Governance regime is partially captured by the CFI — political contestation is one of its five sub-scales — but not fully. This is a candidate direction for further analysis.

Internal development of both indices

Both the ECI and the CFI are analytical instruments developed by the Initiative and applied internally. Neither has been externally peer-reviewed. All scoring decisions are documented transparently to permit replication or revision by independent analysts. The analysis is reproducible from the rubrics, scoring tables, and regression outputs provided in the appendices; it is not a sealed black box.

8. Conclusions

The joint multivariate analysis of engineering complexity and community friction as predictors of high-speed rail cost, applied to the ALTO corridor, supports four conclusions.

One. The two indices are genuinely independent cost drivers. After conditioning on either, the other retains a statistically and substantively significant association with cost. Community friction is not a proxy for engineering complexity, and engineering complexity is not a proxy for community friction. A joint model outperforms either univariate model in explanatory power ($R^2 = 0.897$ versus 0.813 and 0.569), information-criterion values, and parsimony.

Two. Community friction is the stronger single predictor of HSR cost. The univariate CFI regression explains 81 per cent of cross-project variance in log-cost; the univariate ECI regression explains 57 per cent. The standardised effect sizes in the bivariate model ($\beta = 0.70$ for CFI versus $\beta = 0.35$ for ECIw) indicate that per unit of its own standard deviation, CFI has approximately twice the cost impact of ECIw.

Three. Real PPP-adjusted HSR construction cost has not grown with time once complexity and friction are controlled for. The apparent univariate correlation between completion year and cost ($r = +0.57$) is fully mediated through the two indices; the partial F-test for Year over the bivariate model gives $p = 0.83$. Newer HSR projects have tended to be built in more complex and more contested environments, and those characteristics — not a generic construction-cost trend — account for the observed time pattern.

Four. The bivariate reference-class prediction for ALTO is \$142 million per kilometre, with a 95 per cent prediction interval of \$76–264 million per kilometre. For a 1,000 kilometre corridor, this corresponds to a central total cost of \$142 billion and a 95 per cent prediction interval of \$76–264 billion. The bivariate reference-class prediction for HPR is \$28 million per kilometre, with a 95 per cent prediction interval of \$15–54 million per kilometre. The projected five-fold cost gap between the two alternatives decomposes as approximately 46 per cent engineering complexity, 53 per cent community friction, and less than 1 per cent timing.

8.1 Recommendations

Three recommendations follow from the analysis. They are addressed to the parliamentary leaders with formal standing in Canadian infrastructure decision-making.

First, independent cost validation. The stated ALTO budget of \$60–90 billion is inconsistent with the international reference-class evidence assembled in this report. The Parliamentary Budget Officer, or an equivalent independent fiscal authority, should be formally asked to validate the cost estimate against a documented reference-class forecasting methodology (Flyvbjerg, 2008; Flyvbjerg & Gardner, 2023) before the project proceeds beyond the current consultation phase.

Second, explicit community-friction disclosure. The public consultation documentation does not currently disclose the sociopolitical cost pathway separately from the engineering cost pathway. Given that approximately half of the cross-project cost variance in the reference class is attributable to community friction, the public has a legitimate interest in seeing the ALTO friction-cost projection disclosed as a separate line in the project budget.

Third, structured HPR comparison. The consultation process should formally evaluate the HPR alternative against ALTO using the same analytical framework applied in this report. A 1,000 kilometre corridor at \$28 billion, under the bivariate reference-class prediction, is an alternative of sufficient fiscal and political significance that Canadians are entitled to see it rigorously scored, not dismissed by reference to a 300 km/h design-speed specification that the stated budget cannot defensibly deliver.

The central analytical finding of this report is that community friction is as important a cost driver as engineering complexity — and that a cost forecasting framework that ignores either one will produce systematically biased predictions for the other.

Appendix A. ECI scoring descriptors

Each dimension of the Engineering Complexity Index is scored 1 to 5 against descriptor-anchored criteria. Descriptors are drawn from the published engineering record of reference-class projects.

Geotechnical conditions

Score	Descriptor
1	Uniform competent bedrock or consolidated soils; negligible problematic-ground presence.
2	Predominantly competent ground with limited soft-soil or expansive-clay segments; routine site investigation sufficient.
3	Mixed ground conditions with localised geotechnical complications; standard mitigation measures required.
4	Substantial soft-soil, permafrost, karst, or liquefiable-soil exposure requiring specialised treatment across significant corridor length.
5	Severe geotechnical regime: Leda clay, extensive karst, permafrost, or combined exposures requiring programme-defining ground treatment.

Major structures

Score	Descriptor
1	Predominantly at-grade alignment; minimal tunnel or viaduct content.
2	Routine bridge and short viaduct count; tunnel content below 10 per cent of corridor.
3	Substantial bridging, moderate viaduct or tunnel content (10–25 per cent of corridor).
4	Major tunnel or viaduct backbone (25–50 per cent of corridor); significant structure-count driver.
5	Programme-defining tunnelling or viaduct requirement (> 50 per cent of corridor); structures dominate the capital account.

Topography and alignment

Score	Descriptor
1	Flat or gently rolling terrain; vertical relief below 100 m over corridor; generous curve radii easily achieved.
2	Modest relief; occasional grade and curve constraints at design speed.
3	Significant relief or dense urban segments forcing tighter-than-optimal alignment choices; routine cant-deficiency use.
4	Substantial mountain crossings or dense urban segments forcing extensive grade and curve constraints; speed-envelope compromise.

Score	Descriptor
5	Mountainous terrain with relief above 500 m; alignment substantially shaped by topography; multiple speed-envelope compromises.

Hydrology and drainage

Score	Descriptor
1	Minimal watercourse crossings; no significant floodplain or wetland exposure.
2	Routine drainage structures; occasional medium watercourse crossings; limited wetland displacement.
3	Several major watercourse crossings; localised floodplain and wetland exposure; substantial drainage structure count.
4	Multiple major watercourse crossings with significant span requirements; programme-affecting floodplain exposure.
5	Dominant hydrological constraint: extensive watercourse crossings, karst hydrogeology, or dominant floodplain exposure.

Infrastructure interfaces

Score	Descriptor
1	Rural corridor; minimal existing-infrastructure interaction; greenfield alignment.
2	Modest urban segments; routine road and rail crossings; manageable utility relocation.
3	Significant urban segments; multiple major crossings; non-trivial staging and utility relocation.
4	Extensive urban corridor; complex station integration; dense utility and existing-rail interface.
5	Dominant urban and interface complexity; integrated major-station construction; severe staging constraints.

Climate and environmental loading

Score	Descriptor
1	Mild climate; negligible freeze-thaw cycling; low wind and seismic loads.
2	Moderate climate; routine winter design; moderate wind and seismic loads.
3	Substantial winter design requirement; notable freeze-thaw cycling; elevated wind or seismic loads.
4	Severe climate loading: cold-climate catenary, extensive freeze-thaw, or high seismic design category.

Score	Descriptor
5	Extreme combined loading: deep winter temperatures, heavy snow/ice design, major seismic zone, or extreme-weather exposure.

Appendix B. Full project scoring matrix

Dimension-level ECI scores for the 16-project merged dataset. Scores preceded by an asterisk (*) are imputed by the Initiative from the rubric and the public engineering record; unmarked scores are drawn from the published ECI reference-class scoring.

Project	Year	Geo	Str	Top	Hyd	Int	Cli	ECIw	CFI	\$M/km
China HSR (Beijing–Shanghai)	2011	3	4	2	3	3	3	60	12	\$17
Morocco Al Boraq	2018	*2	*2	*2	*2	*2	*3	42	15	\$22
Spain AVE Seville	1992	3	3	3	2	2	3	54	18	\$19
Japan Tokaido	1964	*3	*3	*3	*3	*4	*3	63	20	\$35
Spain AVE Barcelona	2008	3	4	4	2	3	3	65	22	\$24
France TGV Paris–Lyon	1983	2	2	3	2	2	2	44	23	\$26
Japan Hokkaido	2016	*4	*5	*5	*3	*2	*4	79	25	\$88
Italy AV Rome–Naples	2005	*3	*4	*4	*3	*3	*2	66	28	\$42
Taiwan HSR	2007	4	5	4	4	3	4	81	30	\$65
France TGV SEA	2017	*2	*3	*3	*3	*2	*2	51	38	\$55
Netherlands HSL-Zuid	2009	*5	*4	*2	*3	*4	*2	69	40	\$82
UK HS1 (CTRL)	2007	3	5	2	3	5	2	68	45	\$92
Germany Cologne–Frankfurt	2002	3	5	5	3	3	2	74	50	\$78
Germany Stuttgart–Ulm	2022	*3	*5	*2	*2	*5	*1	66	75	\$98
California HSR*	2033	5	5	5	3	4	3	87	78	\$230
UK HS2*	2029	4	5	3	4	5	2	79	92	\$315

Table B-1. Full project scoring matrix. Columns: Geo = geotechnical; Str = structures; Top = topography; Hyd = hydrology; Int = interfaces; Cli = climate. Asterisks (*) mark imputed ECI dimension scores.

The column abbreviations correspond to the six ECI dimensions. The ECIw column is the weighted composite on a 20–100 scale. The CFI column is the aggregate on a 0–100 scale. Cost is in millions of 2026 PPP Canadian dollars per route-kilometre.

Appendix C. Regression output

This appendix documents the ordinary least squares regression output for the four model specifications tested in the main report.

C.1 Summary statistics across all specifications

Model	k	R ²	adj R ²	AIC	BIC	σ (log)
EClw only	1	0.569	0.538	+3.39	+4.93	0.254
CFI only	1	0.813	0.800	-9.99	-8.45	0.167
Year only	1	0.325	0.277	+10.56	+12.11	0.318
EClw + CFI	2	0.897	0.882	-17.58	-15.26	0.129
EClw + Year	2	0.657	0.604	+1.73	+4.04	0.235
CFI + Year	2	0.815	0.787	-8.16	-5.84	0.173
EClw + CFI + Year	3	0.898	0.872	-15.64	-12.55	0.134

Table C-1. Summary statistics for all seven model specifications tested. The preferred EClw + CFI specification is highlighted. k is the number of predictors (excluding the constant).

C.2 Bivariate model: coefficient detail

Predictor	Coefficient	Std. error	t-statistic	p-value	VIF
Constant	+0.6709	0.1776	3.78	0.002	—
EClw	+0.0101	0.0031	3.27	0.006	4.95
CFI	+0.0108	0.0017	6.45	< 0.001	4.95

Table C-2. Detailed coefficient output for the preferred bivariate model: $\log_{10}(\text{cost}) = 0.0101 \cdot \text{EClw} + 0.0108 \cdot \text{CFI} + 0.671$. $F(2, 13) = 56.9$, $p < 10^{-6}$.

C.3 Trivariate model: coefficient detail

Predictor	Coefficient	Std. error	t-statistic	p-value	VIF
Constant	+0.6785	0.1879	3.61	0.004	—
EClw	+0.0101	0.0032	3.12	0.009	5.32
CFI	+0.0106	0.0020	5.32	< 0.001	6.93
Year - 2000	+0.0005	0.0025	0.21	0.834	1.82

Table C-3. Detailed coefficient output for the trivariate model: $\log_{10}(\text{cost}) = 0.0101 \cdot \text{EClw} + 0.0106 \cdot \text{CFI} + 0.0005 \cdot (\text{Year} - 2000) + 0.679$. $F(3, 12) = 35.15$, $p < 10^{-5}$. Year coefficient is not statistically significant ($p = 0.83$); incremental contribution over the bivariate model is non-informative.

C.4 Residual diagnostics (bivariate model)

Project	Observed	Predicted	Residual	Leverage	Cook's D
China HSR (Beijing–Shanghai)	1.230	1.408	-0.178	0.14	0.125
Morocco Al Boraq	1.342	1.258	+0.084	0.28	0.077
Spain AVE Seville	1.279	1.412	-0.133	0.12	0.058
Japan Tokaido	1.544	1.525	+0.019	0.11	0.001
Spain AVE Barcelona	1.380	1.566	-0.186	0.10	0.091
France TGV Paris–Lyon	1.415	1.364	+0.051	0.25	0.023
Japan Hokkaido	1.944	1.741	+0.204	0.26	0.404
Italy AV Rome–Naples	1.623	1.641	-0.018	0.08	0.001
Taiwan HSR	1.813	1.815	-0.002	0.26	0.000
France TGV SEA	1.740	1.597	+0.144	0.18	0.115
Netherlands HSL-Zuid	1.914	1.801	+0.113	0.07	0.020
UK HS1 (CTRL)	1.964	1.844	+0.119	0.07	0.022
Germany Cologne–Frankfurt	1.892	1.959	-0.067	0.09	0.010
Germany Stuttgart–Ulm	1.991	2.147	-0.156	0.28	0.273
California HSR	2.362	2.392	-0.031	0.29	0.011
UK HS2	2.498	2.462	+0.036	0.40	0.029

Table C-4. Residual diagnostics for the bivariate model. Observed, predicted, and residual values are in log₁₀ units. The conventional Cook's D flag threshold of 4/n = 0.25 is crossed by Japan Hokkaido (0.40) and Germany Stuttgart–Ulm (0.27, marginal).

Appendix D. References

The references below are organised by theme. Items marked with a dagger (†) are primary-source government audit or parliamentary-inquiry documents. Items marked with a double dagger (‡) are released Access to Information and Privacy (ATIP) file records cited directly in the analysis. All other items are peer-reviewed academic publications, reference works, or project-sponsor documentation.

Where international HSR project cost figures are cited in the main text, the delivered cost per kilometre is drawn from the primary project-sponsor annual report or, where material, from the independent audit report listed under the relevant project subsection below. Purchasing-power-parity normalisation uses OECD PPP tables at the project's delivery year and converts to Canadian dollars at the 2026 base-year rate.

D.1 Reference-class forecasting and megaproject cost overruns

- Flyvbjerg, B. (2006). From Nobel Prize to project management: getting risks right. *Project Management Journal*, 37(3), 5–15.
- Flyvbjerg, B. (2008). Curbing optimism bias and strategic misrepresentation in planning: reference class forecasting in practice. *European Planning Studies*, 16(1), 3–21.
- Flyvbjerg, B. (2014). What you should know about megaprojects and why: an overview. *Project Management Journal*, 45(2), 6–19.
- Flyvbjerg, B. (2016). The fallacy of beneficial ignorance: a test of Hirschman's hiding hand. *World Development*, 84, 176–189.
- Flyvbjerg, B., Bruzelius, N., & Rothengatter, W. (2003). *Megaprojects and risk: an anatomy of ambition*. Cambridge University Press.
- Flyvbjerg, B., & Gardner, D. (2023). *How big things get done: the surprising factors that determine the fate of every project, from home renovations to space exploration — and everything in between*. Currency / Penguin Random House.
- Flyvbjerg, B., Holm, M. K. S., & Buhl, S. L. (2002). Underestimating costs in public works projects: error or lie? *Journal of the American Planning Association*, 68(3), 279–295.
- Flyvbjerg, B., Holm, M. K. S., & Buhl, S. L. (2003). How common and how large are cost overruns in transport infrastructure projects? *Transport Reviews*, 23(1), 71–88.
- Flyvbjerg, B., Holm, M. K. S., & Buhl, S. L. (2004). What causes cost overrun in transport infrastructure projects? *Transport Reviews*, 24(1), 3–18.
- Kahneman, D., & Tversky, A. (1979). Intuitive prediction: biases and corrective procedures. *TIMS Studies in the Management Sciences*, 12, 313–327.
- Lovullo, D., & Kahneman, D. (2003). Delusions of success: how optimism undermines executives' decisions. *Harvard Business Review*, 81(7), 56–63.
- Priemus, H., Flyvbjerg, B., & van Wee, B. (Eds.). (2008). *Decision-making on mega-projects: cost–benefit analysis, planning and innovation*. Edward Elgar Publishing.
- van Wee, B., & Priemus, H. (Eds.). (2020). *Decision-making on mega-projects*. Edward Elgar Publishing.

D.2 International HSR cost surveys and independent audits

- Amos, P., Bullock, D., & Sondhi, J. (2010). *High-speed rail: the fast track to economic development?* World Bank, Beijing.
- Campos, J., & de Rus, G. (2009). Some stylized facts about high-speed rail: a review of HSR experiences around the world. *Transport Policy*, 16(1), 19–28.
- de Rus, G. (2011). The BCA of HSR: should the government invest in high-speed rail infrastructure? *Journal of Benefit-Cost Analysis*, 2(1), Article 8.
- European Court of Auditors. (2018). *Special Report No. 19/2018: A European high-speed rail network: not a reality but an ineffective patchwork*. Luxembourg: Publications Office of the European Union. †
- Ollivier, G., Bullock, R., Jin, Y., & Zhou, N. (2014). *High-speed railways in China: a look at construction costs*. *China Transport Topics No. 9*. World Bank, Beijing.
- Perl, A. (2002). *New departures: rethinking rail passenger policy in the twenty-first century*. University Press of Kentucky.
- Preston, J. (2013). The economics of investment in high-speed rail. *International Transport Forum Discussion Paper 2013-30*. OECD / ITF.

D.3 Primary-source documentation by reference-class project

United Kingdom — HS1 (Channel Tunnel Rail Link)

- UK National Audit Office. (2001). *The Channel Tunnel Rail Link*. HC 302, Session 2000–2001. London: The Stationery Office. †
- UK National Audit Office. (2012). *The completion and sale of High Speed 1*. HC 1834, Session 2010–2012. London: The Stationery Office. †
- UK Department for Transport / HS1 Ltd. Annual infrastructure management reports (2012 onward).

United Kingdom — HS2

- UK National Audit Office. (2020). *High Speed 2: a progress update*. HC 40, Session 2019–2021. London: The Stationery Office. †
- UK Department for Transport. (2020). *Oakervee Review of HS2*. London: DfT. †
- UK House of Commons Public Accounts Committee. Successive reports on HS2 programme cost and delivery (2013 onward). †
- UK Infrastructure and Projects Authority. Annual Report on Major Projects (2016 onward). †
- UK National Audit Office. (2024). *HS2: update following cancellation of Phase 2*. London: The Stationery Office. †

United States — California High-Speed Rail

- California High-Speed Rail Authority. (2024). *2024 Business Plan*. Sacramento: CHSRA.
- California High-Speed Rail Authority. Project Update Reports to the California State Legislature (biennial, 2013 onward).
- California State Auditor. (2018). *California High-Speed Rail Authority: its flawed decision-making and poor contract management have contributed to billions in cost overruns and delays in the system's construction*. Report 2018-108. Sacramento: California State Auditor's Office. †
- Legislative Analyst's Office. California High-Speed Rail: periodic cost and delivery overviews (2014 onward). Sacramento: LAO. †

France — TGV network

- Cour des Comptes. (2014). *La grande vitesse ferroviaire: un modèle porté au-delà de sa pertinence*. Paris: Cour des Comptes. †
- SNCF Réseau. Annual reports and rail infrastructure cost statements (successor to Réseau Ferré de France).
- Crozet, Y. (2013). High-speed rail performance in France: from appraisal methodologies to ex-post evaluation. *International Transport Forum Discussion Paper 2013-26*. OECD / ITF.

Spain — AVE network

- Bel, G. (2010). *España, capital París*. Barcelona: Ediciones Destino.
- Tribunal de Cuentas. Audit reports on high-speed rail infrastructure (various years). Madrid. †
- ADIF Alta Velocidad. Annual financial and operating reports.
- Albalade, D., & Bel, G. (2012). High-speed rail: lessons for policy makers from experiences abroad. *Public Administration Review*, 72(3), 336–349.

Germany — ICE network and Stuttgart–Ulm

- Bundesrechnungshof. Successive audit reports on Deutsche Bahn and federally funded rail infrastructure. Bonn: German Federal Court of Audit. †
- Deutsche Bahn AG. Integrated Annual Reports (2000 onward).
- Geißler, H. (2010–2011). *Schlichtung Stuttgart 21 — Schlichtungsprotokolle*. Stuttgart: Mediation proceedings (publicly archived).
- Novy, J., & Peters, D. (2012). Railway station mega-projects as public controversies: the case of Stuttgart 21. *Built Environment*, 38(1), 128–145.

Italy — AV / AC network

- Corte dei Conti. Reports on Alta Velocità / Alta Capacità rail expenditure (various years). Rome. †
- Rete Ferroviaria Italiana (RFI). Annual reports.
- Cascetta, E., & Coppola, P. (2014). High-speed rail (HSR) induced demand models. *Procedia – Social and Behavioral Sciences*, 111, 147–156.

Netherlands — HSL-Zuid

- Algemene Rekenkamer. (2014). *Aanleg van de HSL-Zuid: eindrapport*. The Hague: Netherlands Court of Audit. †
- Tijdelijke Commissie Infrastructuurprojecten (Duivesteijn Commission). (2004). *Grote projecten uitvergroot: een infrastructuur voor besluitvorming*. Tweede Kamer, vergaderjaar 2004–2005, 29 283, nr. 6. †

Japan — Shinkansen

- Central Japan Railway Company (JR Central). Annual reports and data books.
- Smith, R. A. (2003). The Japanese Shinkansen: catalyst for the renaissance of rail. *Journal of Transport History*, 24(2), 222–237.
- Hood, C. (2006). *Shinkansen: from bullet train to symbol of modern Japan*. Routledge.
- Japan Railway Construction, Transport and Technology Agency (JRTT). Hokkaido Shinkansen construction reports.

Taiwan — THSR

- Taiwan High Speed Rail Corporation. Annual reports (2007 onward).

Cheng, Y.-H. (2010). High-speed rail in Taiwan: new experience and issues for future development. *Transport Policy*, 17(2), 51–63.

China — CR-HSR

Lawrence, M., Bullock, R., & Liu, Z. (2019). *China's high-speed rail development*. International Development in Focus. Washington, DC: World Bank.

Wu, J., Nash, C., & Wang, D. (2014). Is high-speed rail an appropriate solution to China's rail capacity problems? *Journal of Transport Geography*, 40, 100–111.

China State Railway Group Co. Ltd. Annual reports (successor to China Ministry of Railways).

Morocco — Al Boraq

Office National des Chemins de Fer (ONCF). (2018). *LGV Kénitra–Tanger: bilan du projet*. Rabat: ONCF.

World Bank. (2013). *Morocco high-speed rail line project: implementation completion and results report*. Washington, DC: World Bank.

D.4 ALTO and Canadian corridor context

ALTO Infrastructure Corporation. (2025–2026). Public consultation materials, technical fact sheets, and corridor documentation.

VIA HFR / VIA TGF (High Frequency Rail / Train à Grande Fréquence) Project Office. Technical briefings and stakeholder engagement documentation (2021–2025).

Transport Canada. (2021–2025). High Frequency Rail / Intercity Rail studies and briefings. Ottawa: Transport Canada.

Arup / AECOM. (2021–2023). *Technical briefings for Transport Canada and VIA HFR on high-frequency and high-speed rail options*. Released under Access to Information files A-2025-00026 and A-2025-00027. ‡

Office of the Parliamentary Budget Officer (Canada). Methodological notes on infrastructure costing and fiscal analysis. Ottawa: PBO. †

Infrastructure Canada. National Infrastructure Assessment and associated asset-condition and cost documentation.

VIA Rail Canada Inc. Annual reports, on-time-performance statistics, and corporate plan summaries.

Canadian National Railway Company (CN). Kingston Subdivision operating and capacity documentation; Annual Report on Sustainability.

Statistics Canada. Building Construction Price Index; Railroad Transportation Price Index. Ottawa: StatCan.

D.5 HPR (High Performance Rail) alternative

ALTO HSR Citizen Research Initiative. *High Performance Rail (HPR) Strategy: a 200 km/h mixed-use alternative for the Quebec City – Toronto corridor*. Research note series, 2025–2026. altohsrcitizenresearch.ca.

ALTO HSR Citizen Research Initiative. *Rail infrastructure cost as a function of design speed: a parametric analysis*. Research note, 2025. altohsrcitizenresearch.ca.

ALTO HSR Citizen Research Initiative. *Curve radius, cant deficiency, and the physics of rail alignment at 200–300 km/h*. Research note, 2025. altohsrcitizenresearch.ca.

International Union of Railways (UIC). Design and operating standards for higher-speed conventional rail corridors.

Federal Railroad Administration (USA). Passenger rail design and operating standards; studies on 125–150 mph corridors.

D.6 Community opposition, land-use conflict, and social impact of infrastructure

Dear, M. (1992). Understanding and overcoming the NIMBY syndrome. *Journal of the American Planning Association*, 58(3), 288–300.

Devine-Wright, P. (2009). Rethinking NIMBYism: the role of place attachment and place identity in explaining place-protective action. *Journal of Community and Applied Social Psychology*, 19(6), 426–441.

Teo, M., & Loosemore, M. (2017). Understanding community protest from a project management perspective: a relationship-based approach. *International Journal of Project Management*, 35(8), 1444–1458.

Della Porta, D., & Piazza, G. (2008). *Voices of the valley, voices of the straits: how protest creates communities*. Berghahn Books. (NOTAV and related Italian infrastructure-opposition movements.)

Bruzzone, F., Cavallaro, F., & Nocera, S. (2021). The effects of high-speed rail on the environment. *Transportation Research Part D*, 92, 102719.

Chen, C.-L., & Hall, P. (2012). The wider spatial-economic impacts of high-speed trains: a comparative case study of Manchester and Lille sub-regions. *Journal of Transport Geography*, 24, 89–110.

Mohammad, S. I., Graham, D. J., Melo, P. C., & Anderson, R. J. (2013). A meta-analysis of the impact of rail projects on land and property values. *Transportation Research Part A*, 50, 158–170.

D.7 Statistical methodology and regression diagnostics

Belsley, D. A., Kuh, E., & Welsch, R. E. (1980). *Regression diagnostics: identifying influential data and sources of collinearity*. New York: John Wiley & Sons.

Cook, R. D. (1977). Detection of influential observation in linear regression. *Technometrics*, 19(1), 15–18.

Cook, R. D., & Weisberg, S. (1982). *Residuals and influence in regression*. New York: Chapman and Hall.

Fox, J. (2015). *Applied regression analysis and generalized linear models (3rd ed.)*. Thousand Oaks, CA: SAGE.

Greene, W. H. (2018). *Econometric analysis (8th ed.)*. New York: Pearson.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An introduction to statistical learning with applications in R (2nd ed.)*. New York: Springer.

O'Brien, R. M. (2007). A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity*, 41(5), 673–690.

Seabold, S., & Perktold, J. (2010). statsmodels: econometric and statistical modeling with Python. *Proceedings of the 9th Python in Science Conference*, 92–96.

D.8 Purchasing power parity and construction cost indexing

OECD. Purchasing Power Parities (PPPs) statistics. Paris: Organisation for Economic Co-operation and Development. stats.oecd.org.

Eurostat. Construction cost indices (CCI) for the EU Member States. Luxembourg: European Commission.

Statistics Canada. Building Construction Price Index (BCPI). Ottawa: StatCan. statcan.gc.ca.

UK Office for National Statistics. Construction output price indices (OPIs). London: ONS.

United States Bureau of Labor Statistics. Producer Price Index: construction sector (PCU23610236101). Washington, DC: BLS.

Japan Ministry of Land, Infrastructure, Transport and Tourism (MLIT). Construction cost deflators. Tokyo: MLIT.

International Monetary Fund. World Economic Outlook database; GDP deflator series by country.

D.9 Companion Initiative research notes

ALTO HSR Citizen Research Initiative. *The Engineering Complexity Index (ECI): a weighted rubric for assessing physical engineering cost drivers in high-speed rail*. Research note, 2025–2026.

ALTO HSR Citizen Research Initiative. *The Community Friction Index (CFI): a framework for measuring sociopolitical cost pathways in linear infrastructure*. Research note, 2026.

ALTO HSR Citizen Research Initiative. *CFI Scoring Algorithm: supplement to the Community Friction Index research note*. Research supplement, 2026.

ALTO HSR Citizen Research Initiative. *50-year parametric CO₂ lifecycle analysis of the ALTO corridor versus the HPR alternative*. Research note, 2025–2026.

ALTO HSR Citizen Research Initiative. *Environmental and watershed submissions (Frontenac Arch Biosphere Reserve, Napanee Limestone Plain, species-at-risk analysis, watershed impact assessments)*. Submission series, 2025–2026.

ALTO HSR Citizen Research Initiative. *VIA Rail on-time-performance and operating analysis; CN Rail Kingston Subdivision background research*. Research note series, 2025–2026.

Notation. † denotes a primary-source government audit or parliamentary-inquiry document. ‡ denotes a released Access to Information and Privacy (ATIP) record.

About the ALTO HSR Citizen Research Initiative

The ALTO HSR Citizen Research Initiative (CRI) is an independent, non-partisan citizen research project critically examining Canada's proposed ALTO high-speed rail project. The Initiative operates across citizenresearch.ca and altohsrcitizenresearch.ca.

All Initiative research is published openly. Reference data, scoring tables, and regression outputs are reproducible from the published material.

altohsrcitizenresearch.ca · citizenresearch.ca

Research report, April 2026