

The Anatomy of an Optimistic Forecast

Behavioural bias in the ALTO project — a diagnostic reading of the Flyvbjerg framework

June 2026 · citizenresearch.ca

ABSTRACT

Public debate about ALTO has so far been conducted largely in the currency of individual numbers — a cost estimate here, a ridership projection there — contested one at a time. This paper argues that the more revealing question is not whether any single figure is wrong, but whether ALTO's figures are wrong in a *patterned* way, and what that pattern signifies. Drawing on Bent Flyvbjerg's behavioural account of megaproject planning, it treats ALTO's forecasts as a case to be diagnosed rather than merely audited. The central instrument is Flyvbjerg's distinction between cognitive bias (innocent optimism) and political bias (deliberate strategic misrepresentation), together with his demonstration that the two are separable by the *direction and consistency* of forecasting error rather than by any claim about the inner states of forecasters. On that test, the paper sets out why ALTO's profile places it where the theory predicts strategic distortion will dominate, and why the appropriate response is not the imputation of motive but the substitution of an outside view for the proponent's inside one. The analysis is diagnostic only; the design of an alternative framework is reserved for other Initiative work.

1 The wrong argument to be having

With the public consultation now closed, the contest over ALTO has settled into a familiar shape: the proponent advances a figure, critics advance a rival figure, and the exchange proceeds number by number. This is an argument the proponent is structurally well placed to win, because it concedes the most important point before the first number is spoken — the premise that each estimate is an independent technical product to be checked on its own terms. Bent Flyvbjerg's body of work, accumulated over three decades and the largest project database of its kind, exists to deny exactly that premise. His finding is that the estimates are neither independent nor merely technical: across project types, eras, and continents, they err in the same direction, by large margins, with no improvement over time.¹

That regularity changes the nature of the inquiry. If forecasting error were technical noise, it would be scattered — sometimes high, sometimes low — and the right response would be a better model. Because it is instead systematic and directional, the right response is to ask what produces a *bias* rather than an error. This paper pursues that question for ALTO. It asks what kind of distortion is in play, how an observer could tell one kind from another without reading minds, and what follows for how the project should be appraised. It is a diagnosis, not a verdict on any person, and it stops short of proposing what should be built instead.

2 Two theories of a bad forecast

Flyvbjerg's decisive move is to refuse the assumption, common in behavioural economics, that all behavioural distortion reduces to cognition. Cognitive bias, he argues, is only half the story; political bias is the other half, and behavioural economics commits a "psychology bias" of its own when it overlooks the second.² The two halves correspond to two competing explanations of the same observed outcome — costs that come in high and benefits that come in low. The first is *optimism bias*: a genuine cognitive failing, non-deliberate, in whose grip planners are unaware they are being optimistic. The second is *strategic misrepresentation*: the deliberate distortion of information to secure a desired end, which by the definitions Flyvbjerg borrows from the philosophical literature on deception is, plainly, lying.³

The distinction is not academic. The two diagnoses share a symptom and an outcome but differ in everything that matters for response. Optimism is a defect of method, curable by better technique — forcing distributional information into the estimate. Misrepresentation is a defect of incentive, curable only by changing what forecasters are rewarded and held accountable for. Confuse the two and the prescribed remedy will miss. Flyvbjerg's now-settled position, reached through a long exchange with Daniel Kahneman, is that real decisions involve *both*, with the mix shifting along a scale of political-organisational pressure: where stakes and pressure are low, optimism dominates;

¹Bent Flyvbjerg, "What You Should Know about Megaprojects and Why: An Overview," *Project Management Journal* 45, no. 2 (2014): 6–19.

²Bent Flyvbjerg, "Top-Ten Behavioral Biases in Project Management: An Overview," *Project Management Journal* 52, no. 6 (2021): 531–546.

³Bent Flyvbjerg, Mette K. Skamris Holm, and Søren L. Buhl, "Underestimating Costs in Public Works Projects: Error or Lie?" *Journal of the American Planning Association* 68, no. 3 (2002): 279–295.

where a minister or chief executive must have a particular project, strategic misrepresentation dominates, with optimism present and reinforcing rather than absent.⁴

The mechanism Flyvbjerg names is brutally simple, and worth stating in its bare form because it is the engine of everything that follows: underestimated costs plus overestimated benefits equals funding. A low cost estimate is more easily approved, and so produces overrun; a high benefit estimate is more easily approved, and so produces shortfall. The bias is therefore not random but *functional* — it points in the direction that wins the competition for scarce capital. Flyvbjerg has called the resulting practice design by deception, and it is the practice, not the individual, that the framework indicts.⁵

3 Where ALTO sits on the scale

If the balance between innocent optimism and deliberate misrepresentation depends on the degree of political-organisational pressure, then locating a project on that scale is the first analytical task, and Flyvbjerg gives it a structure. His Proposition 1 holds that for small projects with low strategic import and little top-management attention, bias, if present, originates mainly in cognition. His Proposition 2 holds that for large projects with high strategic import and ample top-management attention, bias originates mainly in politics — in strategic misrepresentation — though cognitive bias remains present.⁶

ALTO sits at the upper extreme of every variable in that proposition. It is delivered through a Crown corporation carrying a multi-billion-dollar mandate; it enjoys explicit ministerial sponsorship; and it competes with every other federal priority for a finite pool of capital. By the framework's own logic, this is precisely the configuration in which strategic misrepresentation should be expected to be the dominant bias, with optimism layered on top. This conclusion is worth stating carefully: it is not yet a finding that ALTO's numbers are distorted. It is a prediction, derived from the project's structural profile, about *where to look* and what kind of distortion to expect if distortion is present. The remaining sections test that prediction against the evidence the framework makes available.

4 Uniqueness, the inside view, and the reference-class problem

The deepest thread runs through three of Flyvbjerg's biases that are really one problem under three names: uniqueness bias, the inside view, and base-rate neglect. His own illustration is almost too apt. A jurisdiction that has never built high-speed rail treats the undertaking as unique; uniqueness licenses the "inside view," in which the estimate is built bottom-up from the specifics of this project; and the inside view licenses ignoring the base rate of comparable projects elsewhere. The promotional framing of ALTO — the first true high-speed rail in Canada, a singular corridor, distinctive Shield geology — is structurally identical to that pattern.

⁴Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), 250–251.

⁵Bent Flyvbjerg, "Design by Deception: The Politics of Megaproject Approval," *Harvard Design Magazine*, no. 22 (2005): 50–59.

What the appeal to uniqueness accomplishes is epistemological, and it is the crux of the whole dispute. To call a project unique is to set the size of its reference class to one, and a reference class of one renders the proponent's bottom-up estimate the only admissible evidence. This is, at bottom, the reference-class problem from the philosophy of probability: any individual case belongs to indefinitely many classes, and the probability one assigns depends entirely on which class is chosen. The proponent wants the operative class to be "this project." The Initiative's instruments — the dual-regime cost model, the load-factor analysis measured against the international viable band, the comparative quality-assurance assessment — are, in this light, a single sustained argument that ALTO is a member of the class "high-speed and intercity rail megaprojects," for which abundant outcome data exist. The disagreement is not, at root, about any one number. It is about the unit of analysis.

That reframing matters because the outside view carries decisive quantitative content. On the largest dataset of its kind, rail projects deliver, on average, costs about 1.40 times their estimate and benefits about two-thirds of their forecast.⁷ Demand forecasts are worse still: for nine of ten rail projects passenger forecasts are overestimated, by an average of roughly 106 per cent, and for the high-speed subclass specifically the average cost escalation is higher than for rail as a whole.⁸⁹ A single illustrative operation follows. Apply the generic rail correction to any proponent's own benefit-cost ratio — multiply by roughly 0.66 divided by 1.40 — and the realised ratio falls to about 0.47 of what was promised, before a single ALTO-specific complication is added. Where the Initiative's appraisal already places ALTO's social benefit-cost ratio far below break-even on its own terms, the outside-view correction compounds on top of it. The philosophical point is that this correction imputes no motive whatsoever. It is simply what the base rate is.

A LIVE DATUM FROM THE REFERENCE CLASS — HS2, JUNE 2026

The abstraction acquires a face in Britain's High Speed Two, the nearest contemporary member of the class. A National Audit Office report published on 29 June 2026 records that the cost of the London–Birmingham programme has roughly doubled since 2020 — an increase of some £36 billion excluding inflation — and that the full railway is now expected between three and thirteen years later than first planned. Most telling for the present argument is the fate of the project's benefit–cost ratio. At the 2020 decision to proceed it stood at 1.2, or "low value for money." Recomputed with the costs now known — had those costs been visible in 2020 — the auditor puts it at 0.3 to 0.4: "poor value for money."

The operation is not identical to the reference-class correction above; it substitutes the realised cost while holding benefits roughly fixed, rather than adjusting the two together. But its direction and magnitude corroborate the same claim, and do so from an independent auditor's evidence rather than a critic's model: the approval-stage ratio was an artefact of underestimation, and on realised costs the case for the project fell below viability from the outset.

⁷Bent Flyvbjerg and Dirk W. Bester, "The Cost-Benefit Fallacy: Why Cost-Benefit Analysis Is Broken and How to Fix It," *Journal of Benefit-Cost Analysis* 12, no. 3 (2021): 395–419. Rail base rates (cost A/E $\approx$ 1.40; benefit A/E $\approx$ 0.66) are drawn from this dataset and reproduced in Flyvbjerg (2021), Table 2.

⁸Bent Flyvbjerg, Mette K. Skamris Holm, and Søren L. Buhl, "How (In)accurate Are Demand Forecasts in Public Works Projects? The Case of Transportation," *Journal of the American Planning Association* 71, no. 2 (2005): 131–146.

⁹Bent Flyvbjerg, Nils Bruzelius, and Werner Rothengatter, *Megaprojects and Risk: An Anatomy of Ambition* (Cambridge: Cambridge University Press, 2003). The high-speed subdivision (average cost escalation $\approx$ 52% for high-speed rail, 45% urban, 30% conventional) derives from this body of work.

A qualification sharpens rather than softens the point. The proponent’s own benefit case does reach for the outside view — but only for the half of the ledger that flatters it. The two studies ALTO released in June 2026 build their benefit magnitudes almost entirely from the international high-speed-rail literature, the same European and Chinese reference class the Initiative invokes. What they import from that class is the size of the upside; what they decline to import is its base rate for realisation — that rail benefits arrive at about two-thirds of forecast and passenger numbers are overstated by roughly a hundred per cent. The class is admitted where it raises the estimate and excused where it would discipline it: base-rate neglect not as an oversight but as a selection rule.

5 The evidential signature: deception versus error

Here the analysis must be most disciplined, because here it is most tempting to overreach. Intent cannot be observed, and a non-partisan research initiative should not pretend otherwise. The framework, read carefully, does not ask it to. What it supplies instead is a distributional signature. Genuine technical error would scatter symmetrically around zero — a roughly normal distribution of overshoots and undershoots, centred near accuracy. What the data actually show is error that is consistently directional: costs under, benefits over, stable across decades and continents, with no improvement as techniques supposedly advance. That asymmetry is the tell. Innocent error is not supposed to know which way to point. When error reliably points in the funding-favourable direction across an entire population of projects, the hypothesis that cognition alone is responsible is the hypothesis that gets falsified.

This is also how Flyvbjerg reads the verdict of Martin Wachs, who after decades studying transportation forecasting concluded that the persistent gaps between forecast and outcome amount not to a technical failing but to a collective failure of professional ethics.¹⁰ For ALTO, the methodologically honest claim is therefore not “the proponent is lying,” which cannot be established and which would forfeit the Initiative’s standing, but something more precise and more durable: that ALTO’s estimates exhibit the canonical directional signature — every adjustable assumption resolved in the direction that favours viability — and that this signature is, on the largest body of evidence in the field, the fingerprint of strategic distortion rather than honest error. The structure of the error carries the inference; the reader is left to draw the conclusion about agency. That is both the more rigorous posture and the more defensible one.

A LIVE DATUM FROM THE PROPONENT’S SIDE — ALTO’S BENEFIT CASE, JUNE 2026

If High Speed Two shows the cost half of the mechanism coming true after the fact, two studies ALTO released the same month — June 2026, two months after the consultation had closed — show the benefit half being assembled before it. A computable-general-equilibrium assessment of structural economic impacts reports a national real-GDP gain of about \$24.4 billion a year; a corridor tourism study adds up to \$3.9 billion in GDP and 43,000 jobs. Neither nets a cost. The macro study excludes construction and operating expenditure by design; the tourism study has no cost side to exclude. What both offer is a benefit total unaccompanied by the outlay required to obtain it.

¹⁰Martin Wachs, “The Past, Present, and Future of Professional Ethics in Planning,” in *Policy, Planning, and People*, ed. N. Carmon and S. S. Fainstein (Philadelphia: University of Pennsylvania Press, 2013), 101–119.

Their internal architecture is the directional signature in miniature. Each is built as a fan of scenarios — pessimistic to optimistic, low to high coordination — and in each the entire fan sits above zero. Every table of the macro study prints the same line, that welfare increases in every scenario; the tourism study's weakest case is still \$177 million and two thousand jobs. The scenario space has a floor at the baseline and no downside tail: the modelled question is only ever how large the gain is, never whether there is a loss. Even the reports' own adverse mechanisms are kept from reaching the total — the tourism study concedes that faster trains shorten stays and turn overnight visits into day trips, and shows length of stay going negative in several cities, yet the aggregate is arranged to rise regardless.

The sharpest tell is where the two documents contradict each other. The macro study omits domestic tourism on the ground that it is largely substitution from other household spending, with little net effect on national output; the tourism study builds most of its \$33.7-billion base, and most of its headline uplift, from precisely that in-corridor domestic travel, counted through gross input–output multipliers that assume no such displacement. Where the promoter's two reports disagree, each resolves the disagreement toward its own larger number. Both, to their credit, label their outputs illustrative, order-of-magnitude, and not forecasts, and make the largest figures conditional on tourism policy the railway itself does not deliver — but the numbers that leave the page are round and unconditional. The caveats stay in the prose; the figures travel. As with HS2, no claim about anyone's honesty is required: it is enough that every adjustable assumption has resolved in the direction that favours the project.

6 Survival of the unfittest

The most consequential idea in the framework, for understanding how a project like ALTO comes to exist at all, is Flyvbjerg's inverted Darwinism. It is not the best projects that get built, he argues, but the projects that look best on paper — and the projects that look best on paper are precisely those with the largest cost underestimates and benefit overestimates, which makes them, in reality, among the worst.¹¹ The approval process thus operates as an adverse-selection mechanism, a Gresham's law for infrastructure in which optimistic estimates drive out honest ones, because the candid project that books realistic costs and realistic ridership loses the funding contest to the one that does not.

This reframes the central question. The issue is not merely whether ALTO is a sound project that may encounter difficulties. It is what it signifies that this project, rather than a more modest alternative, is the one that cleared the hurdles. On the selection logic, a project may clear those hurdles partly *because* it presented numbers a more candid competitor could not match and still survive. The very fact of approval, in an environment that rewards optimism, is therefore itself a piece of evidence — not proof of bad faith, but a structural reason to distrust the survivor's own paperwork.

7 Power, convexity, and the exclusion of the outside view

Flyvbjerg's claim that power amplifies cognitive bias — that powerful decision-makers are, in his phrase, convexity generators, more swayed by what comes readily to mind and more optimistic

¹¹Bent Flyvbjerg, "Survival of the Unfittest: Why the Worst Infrastructure Gets Built—and What We Can Do about It," *Oxford Review of Economic Policy* 25, no. 3 (2009): 344–367.

about risk — connects this framework to his earlier study of rationality and power. The mechanism that should most interest an observer of ALTO is institutional rather than psychological: he documents that those in power tend to exclude experts and deliberative scrutiny when the stakes are highest, precisely because deliberation threatens to disturb a decision already taken.¹²

A public consultation is, in principle, the institutional site at which the outside view ought to enter — the moment when base rates, comparator projects, and independent reference-class evidence acquire standing against the proponent's inside view. The outside view is, after all, the established corrective: quality control by way of comparison with completed projects.¹³¹⁴ The philosophical question a consultation poses, then, is whether it is genuinely structured to admit that evidence, or whether it functions to ratify a conclusion reached in advance. The Initiative's critique of the consultation's adequacy can be restated in exactly these terms: it is the claim that the outside view is being structurally excluded — which is what the theory predicts will happen at the high-pressure end of the scale, where ALTO sits. Exclusion, it should be said, is not always outright refusal; as the next section shows, the outside view can also be admitted so late that it can no longer change the answer, which is exclusion by another clock.

One objection presents itself immediately, and it is worth meeting head-on. It might be said that ALTO did not exclude the outside view at all — that it went out and bought it. In 2024 the proponent issued an advance contract award notice, PAS240625-002-00, for reference-class forecasting, should-cost and should-schedule modelling, and a series of Challenge Boards, and named a single pre-identified supplier on the ground that only one firm was capable of the work. That firm is Oxford Global Projects, the consultancy founded by Bent Flyvbjerg and Alexander Budzier — the commercial vehicle of the very framework this paper applies, retained to take the outside view on ALTO's own numbers. On its face this cuts against any claim of exclusion: the proponent engaged the outside view's own author's firm.

But procuring the instrument is not the same as letting it bind, and that distinction is the whole of the matter. Reference-class forecasting debiases only when its outside-view figure is permitted to move the decision; a should-cost that is commissioned, filed, and left beside an unchanged inside-view estimate is not a corrective but a credential. The framework is explicit that the failure mode is not the absence of the outside view but its subordination — the number produced and then declined. The decisive record, accordingly, is not the existence of the forecast but the comparison: does ALTO's published capital cost and benefit-cost ratio reflect its own reference-class should-cost, or diverge from it? That single document — the inside view and the outside view set side by side — would settle more than any figure the Initiative could model, because it would be the proponent's own instrument speaking. This yields a falsifiable prediction rather than an accusation: if the commissioned reference-class numbers are more conservative than the figures ALTO has advanced in public, the outside view was procured and parked; if they match, the cost critique weakens accordingly. The test is available, and it is coming due.

¹²Bent Flyvbjerg, *Rationality and Power: Democracy in Practice* (Chicago: University of Chicago Press, 1998).

¹³Bent Flyvbjerg, "Quality Control and Due Diligence in Project Management: Getting Decisions Right by Taking the Outside View," *International Journal of Project Management* 31, no. 5 (2013): 760–774.

¹⁴Bent Flyvbjerg, "From Nobel Prize to Project Management: Getting Risks Right," *Project Management Journal* 37, no. 3 (2006): 5–15.

8 Escalation, lock-in, and the manufactured point of no return

Escalation of commitment enters this analysis chiefly as a prospective warning rather than a present diagnosis. Flyvbjerg ties it to preferential attachment: the projects that look best on paper attract the initial funding; initial funding creates lock-in; and once a point of no return is passed, further funds flow to close the gap between the original underestimate and the real cost — good money thrown after bad. Early disbursement is not incidental to this process. It is frequently the instrument by which the point of no return is engineered, so that cancellation comes to entail an irretrievable loss of money and of face.

Read in this frame, the contract-commitment data emerging through the Initiative's access-to-information work is significant less as a record of spending than as a measure of how far the lock-in mechanism has already advanced. The more that is committed before the numbers are independently tested, the harder it becomes for any future government to halt the project, whatever the evidence then shows. The implication is about timing, not motive: the window in which an outside view can still alter the decision is open now and closing, which is the strongest available argument for the urgency of independent appraisal before commitment hardens into inevitability.

The same access-to-information channel now supplies a timing datum of its own. The Initiative's request for the reference-class records described above — the workbook, the should-cost and should-schedule outputs, and above all any document setting the inside view beside the outside view — was met in June 2026 with a ninety-day extension carrying the response to 18 September 2026, and with a notice invoking third-party consultation under section 27. That combination foreshadows a commercial-confidence claim over precisely the should-cost and should-schedule figures that would make the comparison legible. The mechanism is the one this section describes, observed in real time: the record capable of disciplining the decision is scheduled to arrive, if at all, in redacted form and only after further commitment has hardened. Whether it plays out that way is, again, a matter the disclosure itself will settle — but the sequence is the point, and the sequence is the framework's.

High Speed Two shows the lock-in mechanism operating in plain sight, and in a form more counter-intuitive than the theory usually advertises. By 2026, with some £47 billion already spent, the National Audit Office found that the benefit–cost ratio for completing the programme had risen to a range of 1.5 to 6.4 even as the programme grew more expensive — because the estimated cost of cancelling it had more than quadrupled, to a figure comparable with the cost of finishing, and that avoided cost is subtracted from the remaining bill. This is escalation of commitment rendered as arithmetic: once enough is sunk, the books can show that continuing is “value for money” precisely because so much would be forfeit by stopping. The decision to proceed, the auditor records, rested on advice that the ratio merely exceeded 1.5 rather than on the full range. It is worth adding that the independent scrutiny the programme now receives — mega-project assurance panels, a central decision panel — was largely imported after that lock-in rather than before it. The outside view was not so much refused as deferred until it could no longer change the answer. For ALTO the lesson is about sequence: the cheapest moment to apply the test is now, before the commitment that will later make the same test read the other way.

9 The causal relocation: bias as root cause, complexity as alibi

The framework's most important metaphysical claim is that bias is the root cause of overrun, while scope changes, geology, weather, and complexity are merely proximate causes — the visible forms through which the underlying underestimation manifests. Behavioural science, in Flyvbjerg's summary, tells the planner: your biggest risk is you. The Shield was always there to be reckoned with; the expropriation friction and the input-cost inflation were always foreseeable as a class. What is typically missing is not information about them but an honest reckoning with them at the planning stage.¹⁵

This pre-empts the alibi ALTO can be expected to offer when overruns arrive — that they were caused by unforeseeable geological, legal, or market conditions. On the framework's account these are not exogenous shocks but the predicted shape of upstream underestimation: the causal chain runs from bias, to underestimation of scope during planning, to unaccounted-for scope changes during delivery, to overrun. This is also why two of the Initiative's instruments are the most Flyvbjergian in its arsenal, whether or not they were built with him in view. An engineering-complexity scorecard and a community-friction index are attempts to quantify, *in advance*, the magnitude of precisely what the inside view suppresses — to put a number on the complexity and social resistance that will later be offered as an excuse, while that number can still discipline the decision. That is the de-biasing operation the framework prescribes.

High Speed Two supplies an unusually candid illustration of the root-versus-proximate distinction — from the proponent's own hand. Asked to account for the doubling of costs, the programme's delivery body attributed the increase not principally to external shocks but to its own estimates: roughly a third to underestimation, a further quarter to inefficient delivery, and a further tenth to scope change, with inflation making up the balance. Its working definition of scope change is the decisive tell — “the addition of necessary works that were missed from the original scope.” That is not an exogenous event befalling the plan; it is the plan's original incompleteness surfacing during delivery, which is exactly the causal order the framework asserts. When even the builder's own decomposition places underestimation ahead of every other single non-inflationary factor, the alibi of unforeseeable complexity is hard to sustain.

10 A caution, in the service of rigour

One critical qualification protects the credibility of the entire exercise. The vocabulary of bias has a self-sealing tendency that the framework only half-acknowledges. Symmetric error can be relabelled noise; directional error, bias or lying; almost any outcome can be folded back into the scheme after the fact. Gigerenzer has pressed this point as a “bias bias,” and even sympathetic practitioners concede it is often impossible to identify which specific bias is operating or to exclude alternative explanations.¹⁶ Wielded loosely, the bias lexicon becomes unfalsifiable and reads as motive-imputation dressed up as analysis — which is the fastest route by which a non-partisan initiative is recast as a partisan one.

¹⁶Gerd Gigerenzer, “The Bias Bias in Behavioral Economics,” *Review of Behavioral Economics* 5 (2018): 303–336.

The discipline that keeps the analysis honest is to rest the weight on the parts that are empirical and falsifiable — the reference-class comparison, the directional signature, the base-rate correction — and to treat the attribution of *deliberate* deception as an inference the reader is invited to draw from structure, never as a claim asserted about named persons. That line is not merely ethical caution. It is, conveniently, the same line that separates an argument which survives hostile scrutiny from one that does not.

11 Conclusion

Read through Flyvbjerg, the scattered disputes over ALTO’s individual figures resolve into a single diagnosis. The project’s structural profile places it where strategic distortion is predicted to dominate; its forecasts display the directional signature that distinguishes such distortion from innocent error; its survival of the approval process is itself a mark of that selection pressure rather than a warrant of soundness; and the geological and social difficulties it will later cite are the anticipated form of an underestimation already present in the plan. None of this requires, or asserts, a claim about anyone’s honesty.

What the framework does assert is a relocation of the burden of proof. The proponent’s inside-view estimates carry a known, measurable, directional bias; the outside view is the established corrective; and the appropriate demand is therefore that the decision be tested against the base rate before lock-in forecloses the test. That demand is the whole of the present paper’s recommendation. What ought to be built instead, and on what evidence, is a separate question, reserved for other work of the Initiative. ■

WORKS CITED

- Alto (VIA HFR – VIA TGF Inc.). *Advance Contract Award Notice PAS240625-002-00* (project management and control expertise; pre-identified supplier Oxford Global Projects UK Limited). 2024.
- Alto (VIA HFR – VIA TGF Inc.). *Notice of extension, Access to Information request A-2026-0004*. June 2026. On file with the Initiative.
- Aviseo Consulting. *An Overview of the Structural Economic Impacts of Alto: Computable General Equilibrium Modelling Approach to Assessing High-Speed Rail in the Toronto–Québec City Corridor*. Prepared for Alto. June 2026.
- CPCS, in association with HDR. *Tourism in the Alto Corridor: Current Conditions and Potential Impacts*. Prepared for Alto. June 2026.
- Flyvbjerg, Bent. *Rationality and Power: Democracy in Practice*. Chicago: University of Chicago Press, 1998.
- Flyvbjerg, Bent. “Design by Deception: The Politics of Megaproject Approval.” *Harvard Design Magazine*, no. 22 (2005): 50–59.
- Flyvbjerg, Bent. “From Nobel Prize to Project Management: Getting Risks Right.” *Project Management Journal* 37, no. 3 (2006): 5–15.
- Flyvbjerg, Bent. “Survival of the Unfittest: Why the Worst Infrastructure Gets Built—and What We Can Do about It.” *Oxford Review of Economic Policy* 25, no. 3 (2009): 344–367.
- Flyvbjerg, Bent. “Quality Control and Due Diligence in Project Management: Getting Decisions Right by Taking the Outside View.” *International Journal of Project Management* 31, no. 5 (2013): 760–774.
- Flyvbjerg, Bent. “What You Should Know about Megaprojects and Why: An Overview.” *Project Management Journal* 45, no. 2 (2014): 6–19.
- Flyvbjerg, Bent. “Top-Ten Behavioral Biases in Project Management: An Overview.” *Project Management Journal* 52, no. 6 (2021): 531–546.
- Flyvbjerg, Bent, and Dirk W. Bester. “The Cost-Benefit Fallacy: Why Cost-Benefit Analysis Is Broken and How to Fix It.” *Journal of Benefit-Cost Analysis* 12, no. 3 (2021): 395–419.
- Flyvbjerg, Bent, Nils Bruzelius, and Werner Rothengatter. *Megaprojects and Risk: An Anatomy of Ambition*. Cambridge: Cambridge University Press, 2003.
- Flyvbjerg, Bent, Mette K. Skamris Holm, and Søren L. Buhl. “Underestimating Costs in Public Works Projects: Error or Lie?” *Journal of the American Planning Association* 68, no. 3 (2002): 279–295.
- Flyvbjerg, Bent, Mette K. Skamris Holm, and Søren L. Buhl. “How (In)accurate Are Demand Forecasts in Public Works Projects? The Case of Transportation.” *Journal of the American Planning Association* 71, no. 2 (2005): 131–146.
- Gigerenzer, Gerd. “The Bias Bias in Behavioral Economics.” *Review of Behavioral Economics* 5 (2018): 303–336.
- Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

- National Audit Office. *High Speed Two reset*. Report by the Comptroller and Auditor General, Session 2026-27, HC 52. London: National Audit Office, June 2026.
- Wachs, Martin. "The Past, Present, and Future of Professional Ethics in Planning." In *Policy, Planning, and People*, edited by Naomi Carmon and Susan S. Fainstein, 101–119. Philadelphia: University of Pennsylvania Press, 2013.