



The Covenant Sizes the Works

How a PPP's capital structure, not its business case, decides how much infrastructure gets built — and what the guarantees, grants and sharing clauses a ministry of finance is asked for actually buy

Author: David Duarte Arancibia, Austral Fiscal & Infrastructure Intelligence

Date: October 2026

Audience: ministries of finance and budget offices; PPP units and infrastructure ministries; fiscal councils and supreme audit institutions; lenders and development banks that structure or guarantee PPP debt; academics in public finance and project finance

Series: paper 27 of the Austral series on the fiscal risks of public-private partnerships. It is the interaction paper the series was missing. *Foundations for Deciding When to Use a Third-Party Guarantee*, *The Optimal Risk-Retention Frontier* and *The Viability Gap Is an Output, Not an Input* each price one instrument at a time; this paper switches them all on inside a single capital structure, which is where the results that contradict how instruments are requested and evaluated appear.

Disclosure and method. The author was head of the public-private partnership unit at Chile's Ministry of Finance and later head of the Public Finance Division at DIPRES, Chile's budget office; he is a co-developer of PFRAM, the IMF/World Bank PPP Fiscal Risk Assessment Model, and was the World Bank's lead on PPP fiscal-risk assessment. He is the founder of Austral, which sells fiscal-risk tools; readers should weigh that interest.

Provenance. The questions in this paper arose in advisory work. **No figure, sector, design element or document from any mandate appears here, and none was consulted in writing it.** The paper is built on a synthetic toll road in Costa Andina, the synthetic country of *The PPP Budget Nobody Publishes*, and every number in it is generated by a script published alongside the paper, whose parameters are all declared (Appendix A). The results are structural: the road makes them visible and measurable, and §11 says what it cannot prove.

Abstract

A ministry of finance is routinely asked for a guarantee, a grant or a waiver of its revenue share "to improve the bankability" of a public-private partnership. This paper starts from an identity. When something in the structure is free to adjust — the scope, the subsidy bid, the toll — lenders size the debt against a coverage covenant, and the structure grows until the covenant binds again. Wherever the covenant binds, coverage is the constraint, met with equality. **What an instrument buys is works, or a smaller subsidy bid. The substance is in the price.**

On a synthetic toll road with 2,000 traffic paths, financed on a bid case 30% above the State's own traffic study — the optimism Bain found across more than a hundred toll roads — raising the minimum revenue guarantee from 85% to 100% lifts the works from USD 358 million to 385 million, with coverage pinned at 1.30×. From a floor of about 115%, the works stop depending on the forecast. **The lender is lending against the guarantee.**

What that costs depends on whose traffic study it is costed on. **On the bidder's case, the floor buys 3.7 dollars of works per dollar of expected cost — the cheapest of the State's priced instruments. On the State's own study, it buys 0.68, less than a grant,** and 0.18 per dollar of 95th-percentile cost. Across calibrations the level on the State's study straddles a grant; the gap with the bidder's case does not close. The same optimism makes the State's revenue share look four times larger than it is. Waiving that share buys no works while the covenant binds, and almost two dollars of works per dollar where the sponsor's return binds instead.

Where the covenant binds, it fixes the debt in dollars, so leverage decides only how much equity sits on top: at an imposed 70%, next to the sponsor's own peak, the road is USD 196 million smaller than at the leverage where the covenant and the return bind together. Sizing on the closing margin rather than the covenant gives 39% more works and coverage below 1.0× in one path in seven.

The recommendations follow. Cost every guarantee and every sharing income on the State's own study, never on the bidder's. Declare instruments by the works they buy per dollar of expected and of tail cost, never by the coverage they are said to improve. Publish the distribution of coverage in the year that decides alongside the closing margin. And in negotiated works, ask for the capital structure.

Keywords: public-private partnerships · project finance · debt sizing · coverage covenants · minimum revenue guarantees · revenue sharing · viability gap funding · contingent liabilities · fiscal risk · capital structure · optimism bias.

Feedback & discussion. Published openly for challenge. Comment at <https://austral-intelligence.com/research> ; email dduarte@austral-intelligence.com.

Contents

1. The request a ministry of finance receives
2. A toll road in Costa Andina
3. Coverage is the constraint; works are the output
4. The covenant reads the tail; the return reads the mean
5. Leverage: the covenant fixes the debt in dollars
6. The year that decides
7. When the State's share of the upside costs works
8. The number the bank reads first
9. Optimism does not raise the State's bill. It hides it.
10. The price list
11. Limitations, and what this paper does not prove
12. The policy lead

Appendix A — the model and how to reproduce it · Appendix B — robustness · Appendix C — glossary
· References

1. The request a ministry of finance receives

Every PPP unit has heard the sentence, and every budget office has received it in writing: *the project needs a guarantee to be bankable*. Sometimes it is a minimum revenue guarantee. Sometimes it is a capital grant, a longer term, a waiver of the State's share of revenue above a band, or a lighter equity requirement. The justification is always the same — the lenders need more coverage — and the request is usually evaluated on that footing: does the instrument raise the debt service coverage ratio enough to get the project financed, and what will it cost?

This paper argues that the question is malformed. In any structure where something is free to adjust — the scope of the works, the subsidy the bidder asks for, the toll it offers, the payment it makes to the State — the lenders' coverage requirement is not something an instrument improves. It is the constraint against which that free variable is sized. The sponsor and its lenders push the structure to the point where the covenant binds, and no further. Give them an instrument that improves the project's risk profile and they will not sit on extra coverage. They will build more, bid less or offer more, until the covenant binds again.

So wherever the covenant binds, coverage at the margin is exactly the covenant. That much is an identity, and the paper treats it as one. **The substance is in what the instrument buys and at what price** — and that is invisible as long as instruments are studied one at a time, because it only appears when debt, the waterfall, the support instruments and the size of the project are all live in the same model. The series has priced each instrument on its own: *Foundations for Deciding When to Use a Third-Party Guarantee* asks when a guarantee is the right tool, *The Optimal Risk-Retention Frontier* how much risk a ministry should keep, *The Viability Gap Is an Output, Not an Input* why the subsidy is a result of the structure rather than a parameter of it. This is the paper where they are all switched on at once.

None of the mechanics is new to a project-finance banker. Sizing debt by a coverage ratio, subject to a gearing limit, is textbook (Yescombe 2014; Gatti 2018), and so is the observation that leverage is an output of that sizing. What has not been drawn out is the consequence for the State's side of the table: what its instruments buy when the structure they enter is sized against a covenant, why the answer depends on whose traffic study the ministry costs them on, and why that makes its usual vocabulary — "bankability", "coverage", "closing margin" — the wrong one.

One case is outside the argument. **When a project is already built and its debt drawn, nothing is free to adjust**, and an instrument granted then does buy coverage: it buys the avoided default. Such instruments are rescues, and should be priced as rescues. Everything that follows concerns instruments granted while the structure is still being sized — at tender, in an unsolicited proposal, or when additional works are negotiated into an existing contract.

The paper reaches eight results, each closing in a rule:

- **Instruments buy works, never coverage (§3).** Raising the floor from 85% to 100% of reference revenue lifts the works from USD 358 million to 385 million with coverage pinned at 1.30×. From about 115%, the floor replaces the traffic forecast as what the lender is lending against.
- **What an instrument buys depends on which condition binds (§4).** The covenant reads the tail of the distribution, the equity return reads its mean. The same waiver of revenue sharing buys nothing at 70% debt and USD 75 million of works at 30%.
- **Leverage matters mainly through the covenant (§5).** Where the covenant binds, the debt is about USD 250 million whatever the leverage ratio. The sponsor's value peaks at about 70% debt; at 70% the road is USD 196 million smaller than where the two conditions bind together.
- **What moves the year that decides depends on how lenders test the ramp-up (§6).** A shorter grace period or tenor always costs works; the amortisation profile and a longer tenor matter only if coverage is tested from the first principal year.
- **The State's upside is free in works only while the covenant binds (§7).** Where the return binds, a waiver buys works; and the sharing band costs works when it sits below the revenue the lender reads.
- **The closing margin hides the risk (§8).** Sizing by the margin rather than the covenant gives 39% more works and coverage below 1.0× in one path in seven.
- **Optimism does not raise the State's bill; it hides it (§9).** On the bidder's case the base floor costs a tenth of what it costs on the State's study, and the State's revenue share looks four times larger.
- **The price list (§10).** Costed on the State's own study, the floor buys 0.68 dollars of works per dollar of expected cost and 0.18 per dollar of tail cost; a grant buys one; a waiver of revenue sharing buys zero or nearly two, depending on which condition binds.

2. A toll road in Costa Andina

Why a synthetic road

The paper needs a structure in which every piece can be switched on and off, and in which every number can be regenerated by a reader. A real project offers neither. The paper therefore builds a **synthetic toll road** in Costa Andina, the synthetic country the series introduced in *The PPP Budget Nobody Publishes*. A toll road is chosen deliberately. Its mechanics — demand risk, a traffic ramp-up, a minimum revenue guarantee, revenue sharing above a band, a debt waterfall with a coverage covenant — are the common case of demand-risk PPPs, and the evidence on traffic forecasts is unusually well documented: forecast error in transport projects in general (Flyvbjerg, Skamris Holm and Buhl 2005) and optimism in toll-road forecasts specifically (Bain 2009). None of the results depends on the revenue coming from tolls.

The road

The concession runs for **30 operating years**. Traffic ramps up over three years — to 55%, 78% and 93% of steady state — and then grows at 2% a year in real terms. The **reference case**, the State's own traffic study, gives revenue of USD 34.7 million in the first year and USD 64.3 million in the first steady-state year.

The bidder presents a more optimistic case, and the structure is financed on it. **The bid case runs 30% above the reference**. The number follows Bain (2009): across more than a hundred privately financed toll roads, actual traffic averaged 0.77 of the forecast in the first year and 0.79 by the third — a shift in level that does not fade, and whose inverse is about 1.3. The State's reference study is treated as unbiased on average, which is consistent with Flyvbjerg et al.'s (2005) finding of no average overestimate for roads in general, alongside errors larger than $\pm 20\%$ in half of them.

Traffic uncertainty has two parts, because forecast error and year-to-year volatility are different things. Each of 2,000 paths draws **one forecast-error shock to the level of traffic** (log standard deviation 20%) and then **annual shocks of 5%** that accumulate as a random walk. The dispersions those two studies report across projects are closer to 30%, but they mix differences in bias between projects with uncertainty within one; a lender's uncertainty about a single project's level, after its own due diligence, is smaller. The paper uses 20% and reports the main results at 30% in Appendix B, where several magnitudes move a long way. Every comparison uses the same 2,000 paths, so differences between cases are differences in structure, not in luck.

Two traffic worlds share those paths. The structure — the sponsor's expected return and the lender's covenant — is read on the bid case, because that is the case it is financed on. **The State's flows — what it pays under the floor and collects from revenue sharing — are costed on the reference case**, its own study, because that is the case the State believes. Bid-case costs are reported alongside,

because they are what a ministry sees if it costs its instruments on the bidder's numbers, and §9 and §10 turn on the difference.

The contract

Two support instruments are written into the concession, both anchored on the State's reference case, as many programmes anchor them.

- A **minimum revenue guarantee** (the floor) tops revenue up to **85% of reference steady-state revenue**, following the ramp in the first three years and flat in real terms thereafter. What the State pays under it is a fiscal cost.
- A **revenue-sharing band**: above **115% of the reference path**, the State takes **50%** of the excess. What it collects is fiscal income.

The finance

The sponsor finances **70% of the works with debt**, at 6.2% real. Principal has a **three-year grace period** matching the ramp-up, then amortises in equal principal repayments over **22 years**, finishing five years before the contract ends. Below a coverage of $1.20\times$ the waterfall **locks cash up**; locked cash is released when coverage recovers or the debt is repaid, and shortfalls draw on it before the sponsor injects equity. Operating costs are 22% of toll revenue plus USD 6 million a year. Taxes are paid at 25%, with interest deductible and losses carried forward. The cost of equity is the CAPM rate relevered for leverage by Hamada's (1972) formula, from an unlevered beta of 0.65, a real risk-free rate of 3.5% and an equity risk premium of 5.5%: **13.3% at 70% debt**.

Two conditions, one output

The structure has to satisfy two conditions at once.

- **The return condition.** The internal rate of return on the sponsor's expected distributions must reach the cost of equity at the leverage chosen.
- **The covenant.** In nine paths out of ten, coverage must never fall below $1.30\times$: formally, the 10th percentile across paths of the life-minimum annual debt service coverage ratio (ADSCR) must be at least 1.30. This is a stochastic version of what lenders do when they size debt on a downside case with a minimum coverage.

The works the structure carries — written W throughout, in USD million — is the largest investment that satisfies both. In the model the works generate no revenue of their own: more works means more scope for the same traffic, or the same road for a smaller subsidy. W is a measure of financing capacity. Every result in the paper is a statement about how W responds when an instrument, a financing term or a sizing rule changes.

What is solved and what is imposed

One distinction organises the paper. Some quantities are **unknowns of the model**: the works, and the coverage and returns they produce. Others are **decisions of the parties**: the leverage, the floor, the band, the tenor. Moving a quantity from one group to the other changes the question the model answers. With leverage solved for, the model asks how much debt the road can carry; with leverage imposed — the realistic case, since sponsors arrive with a capital structure — it asks what the road can carry *given* that structure. The base case imposes 70% and §5 lets it go.

Table 1 — Costa Andina toll road: declared parameters

Block	Parameter	Value
Road	operating years · ramp-up	30 · 3 years (55%, 78%, 93% of steady state)
	reference revenue, year 1 · first steady-state year	USD 34.7m · USD 64.3m; real growth 2%
	bid case (financing)	$1.30 \times$ reference (Bain 2009: 0.77 in year 1)
	traffic uncertainty	level shock 20% (log s.d.) + annual shocks 5%; 2,000 paths
Contract	floor (minimum revenue guarantee)	85% of reference steady-state revenue, following the ramp, flat thereafter
	revenue sharing	State takes 50% of revenue above 115% of the reference path
Costs	operating costs	22% of toll revenue + USD 6.0m a year; tax 25%, losses carried forward
Finance	debt · cost · grace · tenor	70% of works · 6.2% real · 3 years · 22 years, linear
	waterfall	lock-up below $1.20\times$, released on recovery or repayment
	cost of equity	$r_f 3.5\% + \beta_L \times 5.5\%$, $\beta_U 0.65$ relevered by Hamada → 13.3% at 70% debt
Lender	covenant	10th percentile of life-minimum ADSCR $\geq 1.30\times$, read on the bid case
State	costing basis · discount rate	reference case · 3.5% real, its borrowing cost

What the road confirmed and what it killed

The paper was outlined around six results expected from practice. Before it was written, the synthetic road was built to check whether they were results of the mechanics or of particular circumstances. **Four reproduced. One did not:** the expectation that revenue sharing collected during the ramp-up breaks the structure survives only as a condition, and §7 reports it as such. **One turned out to be arithmetic rather than a finding** — that the gap between the project's value at the social discount rate and at the cost of equity is the effect of the cost of capital, which is true by construction — and it is dropped.

3. Coverage is the constraint; works are the output

The base case

With 70% debt imposed, the road carries **USD 358 million of works**. The binding condition is the covenant. The return condition is slack: at USD 358 million the sponsor's expected return is **21.2%**, against a cost of equity of 13.3%, and the return condition alone would allow USD 498 million. The lenders, not the sponsor's hurdle, decide the size of the road. Debt is USD 250 million.

Raise the floor

Now give the project what a ministry is most often asked for: a higher floor.

Table 2 — The floor buys works, not coverage (70% debt)

Floor (% of reference steady-state revenue)	Works carried, USD m	Coverage at the 10th percentile	Expected PV of floor payments, State's study	95th percentile, State's study	Probability the floor ever pays, State's study	Expected PV, bidder's case
70%	355.1	1.300	2.6	11.0	12%	0.1
85% (base)	357.6	1.300	16.9	100.9	38%	1.6
92.5%	362.5	1.300	33.0	169.4	53%	4.1
100%	384.6	1.300	56.8	244.3	66%	8.9
115%	450.1	1.300	128.0	411.0	85%	30.4

USD million. Works and coverage on the bid case; fiscal costs discounted at 3.5%.

Coverage does not move by a thousandth. Raising the floor from 85% to 100% adds USD 27 million of works.


The flat column is not a coincidence of calibration. It is an identity. If the works are sized until the covenant binds, then coverage at the margin is the covenant, whatever the instrument: an instrument that raises cash flow in the states the covenant reads lets the lender lend more, the sponsor builds more against the extra debt, and coverage returns to 1.30×. Stated generally:

In a PPP whose size is free and whose debt is sized against a coverage covenant, any instrument that improves the risk profile is converted entirely into works and never into coverage.

That much follows from the definition. What does not follow from the definition, and is the substance of the rest of the paper, is **how much** each instrument buys, at **what cost** to the State, and **under which conditions** it buys anything at all.

The consequence for policy is immediate. **Asking for a guarantee "to improve bankability" is a category error.** Wherever the covenant binds, bankability is fixed by it and met with equality. What is being asked for is more works — or, in a fixed-scope tender, a lower subsidy bid, which is the same thing seen from the other side — and the request should be stated, evaluated and published as such.

Figure 1 — The floor buys works, never coverage



When the floor replaces the forecast

Figure 1 shows a second result, and for a ministry of finance it is the more important one.

At the base floor of 85%, the works depend on how uncertain the traffic is: with a 10% forecast error the road carries USD 422 million, with 20% it carries 358, with 30% it carries 317. **At a floor of 115%, all three carry between USD 448 and 452 million** — and so does a case with annual volatility of 8% instead of 5%.

The reason is mechanical, and it has a precise trigger. The covenant reads the 10th percentile of revenue in the year that decides (§6) — here, the first operating year, where that percentile of the bid case sits at 98% of the reference revenue. **Once the floor rises above that level, revenue in the states the covenant reads is the floor**, and the forecast no longer enters the sizing. The lender has stopped lending against traffic and started lending against the State's guarantee. A demand-risk concession has become, in the eyes of the only party that prices its debt, a loan secured on the State with a toll road attached.

That has a direct reading in the language of fiscal commitments. A floor set above the lender's downside revenue is contingent in law and firm in substance: it is what the debt is sized on. At 115% it pays in 85% of paths on the State's own study, with an expected cost of USD 128 million and a 95th percentile of 411. And the trigger is not exotic. With a forecast error of 30% — the dispersion the literature reports — the lender's downside revenue falls to 84% of the reference, below the base floor of 85%: **in the base contract the debt would already be sized on the guarantee.** *The Recognition Rule* asks when a PPP belongs on the government balance sheet. A floor above the lender's downside revenue should at least be disclosed in the fiscal risk statement as a guarantee the debt is sized on, with its cost on the State's own study.

Robustness

The identity holds in every variant in Appendix B — a pure random walk for traffic with 14% annual volatility, forecast errors of 10% and 30%, annual volatility of 8%, covenants of 1.20× and 1.40×, a covenant read at the 5th rather than the 10th percentile, and a covenant tested only from the first principal year. In each, coverage at the works carried equals the covenant to three decimals, and raising the floor moves only the works.

Rule 1. Never request, grant or evaluate an instrument "to improve bankability". State what it buys — works, scope or a lower subsidy bid — and price it in those units. When a floor sits above the revenue the lender reads in the year that decides, disclose it as a guarantee the debt is sized on, with its cost on the State's own study.

4. The covenant reads the tail; the return reads the mean

The identity of §3 holds when the covenant is the binding condition. It is not always. At low leverage the equity return binds instead, and then the question becomes which instruments the return condition values.

Table 3 — What an instrument buys depends on which condition binds

	30% debt: the return binds	70% debt: the covenant binds
Works carried, base	574.6	357.6
Raise the floor 85% → 100%	+4.5	+27.0
Waive the State's revenue share	+75.0	0.0
Capital grant of 50	+50.0	+50.0

USD million. Same road, same 2,000 paths, works sized on the bid case.

The pattern is the whole mechanism in one table. **The covenant reads the tail of the distribution; the equity return reads its mean.** The floor pays only in bad states, so it is valuable to the lender — whose covenant looks at the 10th percentile — and worth little to the sponsor's expected return. Revenue sharing is collected only in good states, so waiving it is worth a great deal to the sponsor's expected return and nothing to the lender, who never looks at those states. A grant pays in every state and buys exactly its face value in either case; that one-for-one is true by construction, since grant-funded works carry no debt and no equity.


So an instrument buys works in proportion to **what it pays where the binding condition looks**. The same instrument can be decisive in one structure and worthless in another, and a ministry that evaluates it without asking which condition binds is evaluating it blind.

Rule 2. Before pricing an instrument, establish which condition binds the structure — the lender's covenant or the sponsor's return — and in which states the instrument pays. Tail instruments buy works mainly against a binding covenant; mean instruments mainly against a binding return.

5. Leverage: the covenant fixes the debt in dollars

Debt is a number, not a ratio

Let the leverage go. Figure 2 sizes the road at every debt share from 20% to 80%.



Where the covenant binds — from about 46% to 80% debt here — the **debt it allows is almost constant, at USD 249–253 million**. The covenant sizes the debt in dollars, from the cash flow available in the states it reads; the leverage ratio only decides how much equity sits on top of that debt, and therefore how large the project is. At 80% the road carries USD 312 million. At the crossing, 45.7% debt, it carries USD 553 million — **77% more works on essentially the same debt**, with almost five times the equity.

The return condition moves far less

Below the crossing the return binds, and there the works respond only weakly to leverage: from 25% to 80% debt, the return condition falls by about a fifth (–20%), against more than two-thirds (–69%) for the covenant condition. Part of that slope comes from the cost of risky debt: with debt priced at the risk-free rate it shrinks to –8%. The rest reflects other features of the model, among them that the cost of equity is held at its leverage at financial close while the debt amortises. **The covenant is the channel through which leverage matters most.**

The interior maximum belongs to fixed hurdles

One might expect an *interior* maximum: a leverage at which the works peak, with too little debt on one side and too much on the other. Under a leveraged cost of equity there is none — below the crossing the works keep rising slowly as leverage falls, to USD 585 million at 20%. An interior maximum does appear, at **53.2% debt and USD 475 million of works**, when the sponsor applies a **fixed hurdle rate** that does not rise with leverage. Then the return condition grows with debt, the covenant condition falls, and they cross. Fixed hurdles are common in practice, so the interior maximum is real in many negotiations; but it is a property of how sponsors set their hurdle, not of the mechanics, and it should be reported that way.

Whose optimum

The more consequential result is in the returns. At the imposed 70%, the sponsor's expected return is 21.2% against a cost of equity of 13.3%, and the net present value of its equity at that cost is **USD 83 million**. That value is not an accident of the imposed number: across the leverage grid, **the sponsor's equity value peaks at about 70%** — USD 83 million at 68% and at 70%, falling to 69 at 80% and to zero at the crossing.

Table 4 — Two optima: the sponsor's and the road's

Debt share	Works carried, USD m	Binding condition	Sponsor's expected return	Cost of equity	Equity NPV at the cost of equity, USD m
45%	554.4	return (crossing at 45.7%)	9.3%	9.3%	0.0
55%	459.4	covenant	12.7%	10.4%	55.0
70%	357.6	covenant	21.2%	13.3%	83.2
80%	311.8	covenant	32.4%	17.8%	69.3

At the imposed 70%, next to the sponsor's peak, the road is USD 196 million smaller than at the crossing. The gap is a rent created mechanically by a capital structure above the crossing: the covenant caps the debt, the equity is a fixed share of a smaller project, and the return is left slack. Competition on the free variable — bidders competing on works, scope or subsidy — erodes that rent by pushing leverage down until the return binds. **Without competition, the rent stays with the sponsor.** And the places where competition is weakest are precisely the ones a ministry of finance reviews most often: unsolicited proposals, additional works negotiated inside existing contracts, and renegotiations.

The crossing is not a number the ministry can simply demand: it depends on the cost of equity, which the sponsor states, and at the crossing the sponsor earns no more than that cost. What the ministry can


do is compute the crossing with an independent cost of equity, measure the rent the proposed structure leaves, and negotiate part of it back — through a minimum equity share, or a cap on the equity return.

Rule 3. Treat leverage as an output, not an assumption. In negotiated works — unsolicited proposals, additional works, renegotiations — compute the crossing with an independent cost of equity, measure the rent the proposed structure leaves with the sponsor, and negotiate part of it back through a minimum equity share or a cap on the equity return.

6. The year that decides

The covenant is read on the *minimum* coverage over the life of the debt. Which year delivers that minimum matters, because only instruments and terms that move that year move the works — and which year it is depends on how the lenders test the ramp-up.

Figure 3 — The year that decides



Source: Austral, synthetic toll road in Costa Andina, 2,000 traffic paths; covenant_model.py. USD million, real. Debt 70%; works 358 in the left panel.

If the lenders test coverage in every year with debt service, **the year that decides is the first operating year**: the 10th-percentile coverage is $1.31\times$ in year 1 against $1.42\times$ in year 4, and year 1 holds the life-minimum in 80% of paths (10th pct. $1.31\times$). Year 1 is the ramp-up year, revenue is at 55% of steady state and — because of the grace period — debt pays interest only. The amortisation profile cannot move a year in which no principal is paid: switching from linear to annuity amortisation lifts the works by about 1%.

Many term sheets do not test the ramp-up years. Tested from the first principal year, the picture changes, and the profile becomes a lever.

Table 5 — What moves the works depends on how the covenant is tested (70% debt)

Term	Tested in every year with debt service	Change	Tested from the first principal year	Change
Base: 3 years of grace, 22-year linear amortisation	357.6	—	398.9	—
Annuity amortisation instead of linear	361.0	+1.0%	498.9	+25%
1 year of grace	308.4	−14%	308.4	−23%
No grace	194.3	−46%	194.3	−51%
Tenor 25 years	361.0	+1.0%	432.4	+8%
Tenor 18 years	344.1	−4%	353.3	−11%
Tenor 15 years	315.7	−12%	316.1	−21%
Tenor 10 years	241.9	−32%	241.9	−39%

USD million. The repeated 361.0 under the first convention is the ceiling set by year 1 alone: no term that leaves year-1 debt service unchanged can lift the works above it.

Two terms cost works under either convention when they are shortened. **The grace period is the largest lever:** removing it puts principal into the ramp-up year and cuts the works roughly in half. And **the tenor can carry the whole closing.** The base case amortises over 22 years and repays five years before the contract ends, which is long for commercial bank debt in most emerging markets. At 15 years the road loses 12–21% of its works; at 10 years, about a third or more. A structure that closes only on a tenor no lender offers has not closed. It depends on someone lending that way, and that dependence belongs in the file. The amortisation profile, by contrast — and a longer tenor — is worth negotiating in this structure only if the covenant skips the interest-only years.

Rule 4. The debt tenor, the grace period and the way lenders test coverage during the ramp-up are market assumptions, not modelling choices. Declare all three, with the sensitivity of the works to each.

7. When the State's share of the upside costs works

Free while the covenant binds

At 70% debt, waiving the State's revenue share would not add a single dollar of works (Table 3). Deferring its start to the first steady-state year, a common remedy, changes nothing either. **Revenue sharing is collected in good states; the covenant is read in bad ones.** The band sits at 115% of the reference path; the revenue the covenant reads in the year that decides lies below it. In the states the lender looks at, the State collects nothing, so there is nothing for a waiver or a deferral to give back. The income the State would forgo goes instead to the sponsor's return, whose equity value rises from USD 83 million to 121 million.

A currency where the return binds

Where the return binds instead, the same waiver buys USD 75 million of works at 30% debt. The share the State gives up is worth **USD 40 million** on its own study. On the bid case — the case the sponsor finances on — the same share, at the same discount rate, is worth **USD 153 million**. A waiver therefore buys almost two dollars of works per dollar the State believes it is giving up: the optimism that makes the guarantee look cheap (§9) makes the State's upside look valuable to the bidder, and in a structure bound by the return the State can trade on that.

The condition under which the band bites

Even with the covenant binding, sharing starts to cost works when its band falls below the revenue the covenant reads. That happens when the bid case runs far enough above the reference path on which the band is anchored.

Table 6 — The anchoring condition (band fixed at 115% of the reference path, 70% debt)

Bid case over reference	Band ÷ revenue the covenant reads in year 1	Works: sharing from year 1	Sharing deferred to year 4	No sharing
1.00×	1.53	307.0	307.0	307.0
1.30× (base)	1.18	357.6	357.6	357.6
1.50×	1.02	425.8	425.8	425.8
1.70×	0.90	465.6	495.1	495.9
2.00×	0.77	504.9	545.4	600.4

USD million. "Revenue the covenant reads" is the 10th percentile of first-year revenue in the bid case.

While the band sits above the revenue the lender reads (ratio above one), sharing costs nothing and deferring it buys nothing. Once the band falls below it — here, when the bid runs more than about 50% above the State's own study — the State collects in the very states the covenant reads, and the works fall. At $1.70\times$ the band takes USD 30 million of works, and **deferring sharing to the first steady-state year recovers almost all of it**. At $2.00\times$ the band bites in later years too and deferral recovers less than half.

So the expected result survives, but not as a rule. **The error is not to collect early. It is to anchor the band on one path and finance the project on another, far above it** — at which point the State shares in revenue the project needs to pay its debt.

Rule 5. Where the covenant binds, keep revenue sharing: it costs no works. Where the return binds, treat it as a currency, and price any waiver on the State's own study, not on the bidder's. If the bid case runs so far above the reference that the band falls below the revenue the lender reads, raise the band or accept fewer works.

8. The number the bank reads first


Project files published by PPP units typically report a net present value, an internal rate of return and a *closing margin*: the gap between the project's return and its cost of capital, or the subsidy that closes it. **The financier reads something else first: the distribution of coverage in the year that decides.** A structure can pass the margin test and fail the covenant in a large share of paths.

The road makes the point sharply. Size it by the closing margin — the largest works at which the sponsor's expected return equals its cost of equity, ignoring the covenant — and compare with sizing by the covenant.

Table 7 — Same road, same traffic, two sizing rules (70% debt)

Sized by	Works, USD m	Coverage at the 10th percentile	Paths that lock cash up at least once	Paths with coverage below 1.0x at least once
the covenant	357.6	1.30	5%	0%
the closing margin (return = cost of equity)	498.3	0.94	45%	15%

Bid case. Below 1.0x: debt service exceeds the cash flow available for it; in the model it is met from locked-up cash or sponsor equity, so no path defaults outright.



Sizing by the margin delivers **39% more works** — and a structure that locks cash up in almost half the paths and cannot cover its debt service from operating cash in one path in seven. Both pass the test a PPP file applies — a non-negative margin — and nothing in that test separates them.

No lender would lend at a 10th-percentile coverage of 0.94, so the margin-sized road does not reach financial close as drawn. The danger lies elsewhere. **The State sizes on the margin** whenever its own feasibility study or viability-gap calculation sets the subsidy so that the project "just closes", and whenever it evaluates an unsolicited proposal on its stated return. Then the subsidy is set too low for the lenders, and the gap reappears later — as a failed close, a renegotiation, or a request for a guarantee.

The number that separates the two structures is one every financial model is a short step from computing and almost no public file reports: **the distribution of coverage in the year that decides**, read on the State's own study. For the covenant-sized road it is sobering enough: on the reference case, coverage falls below the 1.30× covenant at some point over the life of the debt in 48% of paths, and below the 1.20× lock-up in 34%.

Rule 6. Alongside the closing margin, publish the distribution of coverage — in the year that decides and over the life of the debt — on the State's own traffic study: at minimum the probability that coverage falls below the covenant, below the lock-up and below 1.0×. A file that reports only the margin cannot tell a bankable structure from one that is short of its debt service in one path in seven on the bidder's own case, and in more than half on the State's.

9. Optimism does not raise the State's bill. It hides it.

The road is financed on a bid case 30% above the State's reference — Bain's ratio. Suppose the reference is right: traffic comes in on the State's study, with the same shocks around it.

Table 8 — Financed on the bid case, living on the reference case (works fixed at USD 358 million, 70% debt)

	If the bid case were right	When the reference case is right
Coverage at the 10th percentile	1.30	1.12
Paths that lock cash up at least once	5%	34%
Paths with coverage below 1.0×	0%	0% (16% without the floor)
Sponsor's expected return (cost of equity 13.3%)	21.2%	16.6%
Sponsor's equity NPV at its cost of equity, USD m	83.2	34.2
Expected PV of floor payments, USD m	1.6	16.9
95th percentile of floor payments, USD m	1.9	100.9
Expected PV of the State's revenue share, USD m	152.6	39.5
State's expected net position (share less floor), USD m	+151.0	+22.6

Read the floor rows carefully. What the State pays under the floor depends only on actual traffic, because the floor is anchored on the State's study: it pays the same whether the road was sized at 307 or 358, and whatever the bid said. **Under the contract as written, optimism does not raise the State's bill. It hides it.** On the bidder's case the floor is a rounding error — USD 1.6 million. On the State's own study it is USD 16.9 million expected and 101 million at the 95th percentile: more than ten times larger. The same optimism inflates the other side of the ledger: the State's revenue share, worth USD 153 million on the bid case, is worth 40 on its own study. A ministry that books both on the bidder's numbers sees a net position of +151; on its own, +23.

Who carries the optimism itself? Had the road been sized on the reference case from the start, it would have carried USD 307 million, not 358. The extra 51 million of works are financed by optimism and carried by the lenders' headroom — lock-ups rise from 5% to 34% of paths — and by the sponsor's return, whose equity value falls from 83 to 34 million and still clears its cost of capital. **The floor is what keeps those paths whole:** without it, the same road would fall below 1.0× in 16% of paths. The guarantee does not pay more because of the optimism; it is what lets the optimistic structure survive the realistic traffic.

A ministry that costs its instruments on the bidder's case will approve guarantees at between a tenth and a fifth of their cost — a tenth for the base floor, a fifth for raising it to 100% — and count sharing income at four times its value. Nothing in the structure will tell it so, because the structure is sized on the same case.

Rule 7. Cost every guarantee and every sharing income on the State's own traffic study, never on the bidder's. Publish the bid case as a multiple of that study and of what comparable assets have achieved.

10. The price list


Every instrument has a price in works

If every instrument converts into works, every instrument has a price: so many dollars of works per dollar of fiscal cost. Because the State's cost has a distribution, the price has two parts — per dollar of expected cost and per dollar of tail cost. And because §9 showed that the cost depends on whose study it is measured on, each price has two readings.

Table 9 — The price list (works sized on the bid case; State's flows discounted at 3.5%)

Instrument	Works bought, USD m	Expected cost, State's study	Works per dollar of expected cost, State's study	Works per dollar of 95th-percentile cost, State's study	Works per dollar of expected cost, bidder's case
<i>70% debt — the covenant binds</i>					
Raise the floor 85% → 100%	+27.0	39.9	0.68	0.18	3.73
Capital grant of 50	+50.0	50.0	1.00	1.00	1.00
Waive the State's revenue share	0.0	39.5 forgone	0.00	0.00	0.00
Tenor 25 years instead of 22	+3.4	—	free, if a lender offers it		
Cap leverage at the crossing (negotiated works)	+195.8	—	free; paid from the sponsor's rent		
<i>30% debt — the return binds</i>					
Raise the floor 85% → 100%	+4.5	39.9	0.11	0.03	0.63
Capital grant of 50	+50.0	50.0	1.00	1.00	1.00
Waive the State's revenue share	+75.0	39.5 forgone	1.90	0.33	0.49

USD million. Expected and 95th-percentile costs are the path-by-path change in the State's net position (floor payments less sharing receipts) against the base contract; for the waiver, the "tail" is income forgone in the best states, not a payment in the worst.

Figure 5 — The price list depends on whose traffic study you cost it on

Source: Austral, synthetic toll road in Costa Andina, 2,000 traffic paths; covenant_model.py. USD million, real. Debt 70%, covenant binding; works sized on the bid case. State's flows discounted at 3.5%.

The same guarantee, two prices

On the bidder's case, raising the floor buys 3.7 dollars of works per dollar of expected cost — the cheapest of the State's priced instruments. On the State's own study, in the base calibration, it buys 0.68: less than a grant. At the 95th percentile it buys 0.18. The gap between those numbers is the whole of \$9 in one row: the covenant is read on the bid case, so the floor buys works as if traffic will be high; it pays on actual traffic, so it costs as if traffic will be what the State expects.

Irwin (2007) argues that a guarantee should be valued with its risk, not at its expected cost. The price list adds the other side of the ledger — what the guarantee buys — and shows that the costing basis matters as much as the risk. The point is robust where the level is not. Across the variants in Appendix B, the floor's price on the State's study ranges from 0.15 to 1.30 dollars of works per dollar — sometimes above a grant, more often below — while on the bidder's case it is always flattered, by a factor of between two and about twenty.

That is the quantitative content behind a pattern the series has documented from the other side. *The Useful Silence* argues that governments prefer liabilities their fiscal rules do not see. The price list shows how easy the preference is to rationalise case by case: costed on the numbers in the file, a guarantee buys several times more infrastructure per dollar than a grant — and its cost does not appear in the budget at signature.

The other instruments

A capital grant is the honest benchmark: one dollar buys one dollar of works, expected or in the tail, and it is recorded when it is paid. A waiver of revenue sharing is either worthless or a bargain, depending on the regime: nothing while the covenant binds, almost two dollars of works per dollar where the return binds, because the bidder values the upside on its own case. And the two free levers

on the list are not the State's to give: a longer tenor is a market term, and a leverage cap is a negotiation in which the price is paid out of the sponsor's rent rather than the budget.

Three cautions

Prices are not constant. The floor's price changes as it rises: on the State's study, 0.31 dollars of works per dollar from 85% to 92.5%, then 0.93 from 92.5% to 100% and 0.92 from 100% to 115% — it gets cheaper exactly as it crosses the lender's downside revenue, which is also where it becomes the collateral (§3). A ministry that simply buys the cheapest lever will be drawn towards the guarantee that should be disclosed as firm.

Percentiles do not add. A tail budget has to be computed on the whole package of instruments, not instrument by instrument.

The regime decides the ranking. The same list read at 30% debt reverses the order of the floor and the waiver. In a fixed-scope tender the list translates into subsidy rather than works — raising the floor from 85% to 100% lets a bidder at 70% debt ask for USD 27 million less grant, at an expected cost to the State of USD 40 million on its own study — but that translation assumes the bidder keeps its leverage.

Rule 8. Publish the price list, costed on the State's own study: for every instrument requested or granted, the works it buys per dollar of expected cost and per dollar of tail cost. Set a tail budget for the package as a whole, and never compare a guarantee with a grant on expected cost alone.

11. Limitations, and what this paper does not prove

A synthetic road is not evidence of frequency. The results in §§3–10 are structural: they follow from the mechanics of the covenant and the waterfall, and the road makes them visible and measurable. They do not show how often, in real portfolios, the covenant rather than the return binds, how far bid cases run above reference cases, or how much rent negotiated structures leave with sponsors. Testing those propositions requires project-level financing data — leverage, tenors, covenants and realised coverage — that few PPP programmes publish, which is itself one of the paper's points.

The identity is robust; the magnitudes are not. Appendix B shows the identity of §3 holding in every variant tried, the bidder's case flattering the floor in every one, and the headline results keeping their direction at a 30% forecast error, with one exception: the floor's price on the State's study rises above a grant. The size of each effect depends on the calibration: the works a higher floor buys depend on where it sits relative to the revenue the covenant reads, the floor's price on the State's study straddles a grant across variants, and the region in which revenue sharing bites depends on the anchoring. The paper reports these as conditions, with the tests that locate them.

The model is deliberately simple, and two simplifications matter. Construction is instantaneous — the works are paid at financial close and the road opens the following year — so there is no construction risk, no interest during construction and no completion support; together with interest-only ramp-up years, that is what makes year 1 decide when every year is tested, and §6 reports the alternative convention for that reason. And the cost of equity is held at its leverage at close while the debt amortises, which is part of the slope of the return condition in §5. Beyond those: there is no refinancing, no cash sweep, no debt service reserve account beyond the lock-up, no inflation and no exchange-rate risk; the sponsor's return is measured on expected distributions, with shortfalls met by the sponsor rather than by default; and the covenant is a single stylised criterion, where real lenders combine a minimum coverage on a base case, a downside case and a gearing limit.

Competition is not modelled. The paper treats the works as the free variable and reports the largest works a structure can carry. Whether a real bidder offers those works, or keeps part of the difference, depends on the competition for the contract. §5 and §10 say where that matters most.

What would test it. A portfolio test would need, for each project, the financing case, the State's reference case, the covenant, the support instruments and the realised coverage. The rules in §12 are written so that the data a programme would publish to follow them is the data a test would need.

12. The policy lead

The paper's starting point is an identity: in a PPP whose size is free, coverage is the constraint and works are the output. What makes it useful is the price list that follows, and the finding that the price depends on whose traffic study it is measured on. Eight rules follow, and they are deliberately operational.

1. **Never evaluate an instrument by the coverage it improves.** State what it buys — works, scope or a lower subsidy bid. When a floor sits above the revenue the lender reads in the year that decides, disclose it as a guarantee the debt is sized on.
2. **Establish which condition binds before pricing anything.** Tail instruments buy works against a binding covenant; mean instruments against a binding return.
3. **Treat leverage as an output.** In negotiated works, compute the crossing with an independent cost of equity, measure the sponsor's rent, and negotiate part of it back through a minimum equity share or a cap on the equity return.
4. **Declare the tenor, the grace period and how coverage is tested in the ramp-up as market assumptions,** with the sensitivity of the works to each.
5. **Keep revenue sharing where the covenant binds; price any waiver on the State's own study where the return binds;** raise the band if the bid case runs far above the reference.
6. **Publish the distribution of coverage in the year that decides,** on the State's own study, alongside the closing margin.
7. **Cost every guarantee and every sharing income on the State's own study, never on the bidder's,** and publish the bid case as a multiple of it.
8. **Publish the price list:** works bought per dollar of expected and of tail cost, with a tail budget for the package.

None of this requires new institutions. Most of the inputs already sit in the financial model behind any PPP that reaches financial close; the distributions need a simulation layer of the kind PPP units already run to value revenue guarantees. The reform is to stop arguing in the vocabulary of bankability, coverage and margins, and to report the prices — on the State's own numbers — alongside the coverage distribution.

Appendix A — The model and how to reproduce it

Script. `covenant_model.py` (Python 3, NumPy), published alongside the paper at austral-intelligence.com/research/covenant-sizes-the-works/covenant_model.py, generates every number: `python3 covenant_model.py` prints every table in the paper and writes `out/results.json`; `--figures` also draws the five figures through `covenant_figures.py` (matplotlib). Parameters are declared in a single frozen dataclass (`Road`); random numbers are drawn once from a fixed seed and shared by every case.

Revenue. For path i and operating year t , gross toll revenue is

$$R_{it} = m \cdot r_t \cdot \exp\left(\varepsilon_i + \sum_{s \leq t} \eta_{is}\right),$$

where r_t is the reference path (ramp-up then 2% growth), $\varepsilon_i \sim N(-\sigma_L^2/2, \sigma_L^2)$ the forecast-error level shock ($\sigma_L = 0.20$) and $\eta_{is} \sim N(-\sigma^2/2, \sigma^2)$ the annual shocks ($\sigma = 0.05$). The **bid case** sets $m = 1.30$; the **reference case** sets $m = 1$, with the same ε_i and η_{is} .

Support. Revenue to the project is $\tilde{R}_{it} = \max(R_{it}, F_t) - s \cdot \max(\max(R_{it}, F_t) - b r_t, 0)$, where F_t is the floor — 85% of reference steady-state revenue, following the ramp in years 1–3 and flat in real terms thereafter, so that by year 30 it is about half the reference path — $b = 1.15$ the band and $s = 0.5$ the State's share, from the year set by `SHARE_FROM`. Floor payments are $\max(F_t - R_{it}, 0)$; sharing receipts are the second term.

Cash flow and coverage. Operating costs are 22% of gross toll revenue R_{it} plus USD 6 million. Cash flow available for debt service is $\tilde{R}_{it}$ less operating costs and tax (25%, interest deductible, losses carried forward, straight-line depreciation of the privately financed works over 30 years). With a capital grant G , debt $D = d(W - G)$ and equity $(1 - d)(W - G)$ finance the rest; grant-funded works carry no debt, no equity and no depreciation. Debt is drawn at close, pays interest at 6.2% throughout and principal after a three-year grace, over 22 years (linear, or annuity where stated). ADSCR_{it} is cash flow available over debt service; the covenant requires

$$Q_{10}\left(\min_{t \in \mathcal{T}} \text{ADSCR}_{it}\right) \geq 1.30,$$

where $\mathcal{T}$ is every year with debt service (base) or every year that repays principal (`COV_FROM="principal"`), read on the bid case.

Waterfall. Cash after debt service is distributed when $\text{ADSCR} \geq 1.20$ and locked up otherwise; locked cash earns nothing and is released in full in the first year coverage recovers or when the debt is repaid,

and shortfalls draw on it before the sponsor injects equity. "Coverage below 1.0×" counts paths in which debt service exceeds cash flow available for it in at least one year; in the model no path defaults outright.

Return. The cost of equity is $k_e = r_f + \beta_U(1 + (1 - \tau)d/(1 - d)) \cdot \text{ERP}$ (Hamada 1972) at the leverage at close, or a fixed hurdle where stated. The return condition requires the IRR of the sponsor's expected distributions on the bid case, against the equity outlay, to reach k_e .

Sizing. W is the largest works satisfying both conditions, found by bisection on each and taking the lesser; the crossing is the debt share at which the two give the same works, found by bisection on the debt share.

State's flows. Floor payments and sharing receipts are discounted at 3.5%, the State's real borrowing cost, on the reference case unless stated. "Expected" is the mean across paths of their present value; "95th percentile" is the 95th percentile across paths. For an instrument in Table 9, both are taken on the path-by-path change in the State's net position against the base contract.

What is not in the repository. No file from any advisory mandate.

Appendix B — Robustness

Table B1 — The identity, and the floor's price, under alternative assumptions (70% debt)

Variant	Works, floor 85%	Works, floor 100%	Coverage at the percentile (both)	Floor 85→100%: works per dollar, State's study	Same, bidder's case
Base	357.6	384.6	1.300	0.68	3.73
Random walk, annual volatility 14%, no level shock	362.4	407.3	1.300	0.65	1.31
Forecast error 10%	421.6	425.6	1.300	0.15	2.86
Forecast error 30%	317.0	383.4	1.300	1.30	3.68
Annual volatility 8%	350.3	384.4	1.300	0.65	2.09
Covenant read at the 5th percentile	332.1	378.0	1.300	1.15	6.35
Covenant 1.20×	390.7	416.6	1.200	0.65	3.58
Covenant 1.40×	329.9	357.1	1.400	0.68	3.75
Covenant tested from the first principal year	398.9	441.7	1.300	1.07	5.91

USD million. In every variant the covenant binds and coverage at the works carried equals the covenant. The floor's price on the State's study straddles a grant (1.00); on the bidder's case it is higher in every variant, by a factor of 2 (random walk) to about 20 (forecast error 10%).

Table B2 — The headline results at a 30% forecast error

Result	Base (forecast error 20%)	Forecast error 30%
Works carried at 70% debt, USD m	357.6	317.0
Sizing by the closing margin: extra works · paths below 1.0×	+39% · 15%	+55% · 26%
Crossing · works lost at 70% against the crossing, USD m	45.7% · 196	40.1% · 237
Base floor: expected cost on the bidder's case ÷ on the State's study	1.6 ÷ 16.9	8.5 ÷ 37.5
Floor 85→100%: works per dollar, State's study · bidder's case	0.68 · 3.73	1.30 · 3.68
Waiver of revenue sharing at 30% debt: works per dollar, State's study	1.90	1.44

Every qualitative result holds at a 30% forecast error except one: the floor's price on the State's study rises above a grant, because the base floor already sits above the revenue the covenant reads (§3). The gap with the bidder's case persists, and several magnitudes move a long way.

Appendix C — Glossary

ADSCR — annual debt service coverage ratio: cash flow available for debt service in a year divided by the debt service due that year. **Bid case** — the traffic case on which the structure is financed; here 1.30 times the reference case. **Binding condition** — whichever of the covenant and the return condition limits the works first. **Closing margin** — the gap between a project's return and its cost of capital, or the subsidy that closes it; what PPP files usually publish. **Covenant** — the lenders' minimum coverage, here read on the 10th percentile of the life-minimum ADSCR. **Crossing** — the debt share at which the covenant and the return condition allow the same works. **Floor (minimum revenue guarantee)** — the State tops revenue up to a fixed level. **Lock-up** — the waterfall rule that stops distributions to the sponsor while coverage is below a threshold. **Reference case** — the State's own traffic study, on which the floor and the band are anchored and its flows are costed. **Revenue-sharing band** — the State takes a share of revenue above a threshold set relative to the reference case. **Works carried (*W*)** — the largest investment that satisfies both the covenant and the return condition. **Year that decides** — the operating year in which the covenant's percentile reaches its minimum.

References

Bain, R. (2009). "Error and optimism bias in toll road traffic forecasts." *Transportation* 36(5): 469–482.

Flyvbjerg, B., Skamris Holm, M. K. and Buhl, S. L. (2005). "How (in)accurate are demand forecasts in public works projects? The case of transportation." *Journal of the American Planning Association* 71(2): 131–146.

Gatti, S. (2018). *Project Finance in Theory and Practice*, 3rd ed. London: Academic Press.

Hamada, R. S. (1972). "The effect of the firm's capital structure on the systematic risk of common stocks." *Journal of Finance* 27(2): 435–452.

Irwin, T. C. (2007). *Government Guarantees: Allocating and Valuing Risk in Privately Financed Infrastructure Projects*. Washington, DC: World Bank.

Yescombe, E. R. (2014). *Principles of Project Finance*, 2nd ed. Oxford: Academic Press.

Austral series papers cited: *Foundations for Deciding When to Use a Third-Party Guarantee* · *The Optimal Risk-Retention Frontier — How Much Risk a Ministry of Finance Should Keep* · *The Viability Gap Is an Output, Not an Input* · *The Recognition Rule — When a PPP Lands on the Government Balance Sheet* · *The Useful Silence — Why Governments Choose the Liability They Cannot See* · *The PPP Budget Nobody Publishes — The Affordability Envelope and the Commitments It Pre-Spends*. All at austral-intelligence.com/research.