AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

TECHNICAL PAPER
NO. 02 · BUILD

02

# The Discipline of Confidence

AACE cost-estimate maturity from concept screening to a bankable decision, and why the class of an estimate is earned by project definition, not conferred by its title.

TECHNICAL PAPER · BUILD

Aurus Technical Committee
July 2026 · ~55-min read

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

# Evidence before assertion.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT

CONTENTS

# What this paper covers



HOW TO READ THIS PAPER

The argument runs from problem to practice. Chapters 1 and 2 establish why cost estimates mislead and how the AACE classification framework corrects the error. Chapters 3 and 4 examine the basis of estimate and the contingency that make a class honest. Chapters 5 and 6 follow the ladder to a bankable decision and set out what makes mineral estimates harder. Chapter 7 offers a maturity model an owner or lender can apply to any estimate. Every figure is traced to a published standard or a peer-reviewed source in the reference chapter, and framework exhibits are labelled as capability instruments carrying no project data.

# A class number is a statement about knowledge, not a promise about price

Every capital decision in mining rests on a cost estimate, and every cost estimate carries an accuracy that the project has earned through the work done to define it. The industry keeps confusing the label on that estimate with the confidence behind it. This paper sets out the AACE International classification framework, its application to mineral projects, and the practice that turns a class number from a marketing claim into a defensible statement.

The record is unforgiving. A review of 63 international mining and smelting projects found that as-built capital costs averaged about 14 percent above the bankable-feasibility-study estimate, that roughly half of projects landed outside a band of plus or minus 15 percent around the study figure, and that overruns of 100 percent or more struck about one project in thirteen (Bertisen and Davis, The Engineering Economist, 2008). These are not rounding errors. They are the cost of treating an estimate as a number rather than as the product of a defined body of engineering.

AACE International resolves the confusion with a single governing rule: the class of an estimate is determined by the maturity of project definition, and by nothing else (Recommended Practice 18R-97). Five classes run from Class 5, a concept screen built on 0 to 2 percent of full definition, to Class 1, a check estimate built on 50 to 100 percent. Each class carries an expected accuracy range, and those ranges are wide at the top of the ladder and narrow at the bottom: a Class 5 estimate can sit anywhere from 50 percent low to 100 percent high, while a Class 1 estimate is held to within roughly 10 percent (18R-97, Table 1).

For mineral projects the framework maps cleanly onto the study ladder that lenders and securities codes already recognise (Recommended Practice 47R-11). A Class 5 estimate is the cost engine of a Preliminary Economic Assessment, generally built on Inferred Resources from early drilling. A Class 4 estimate belongs to a pre-feasibility study. A Class 3 estimate is what a bankable or definitive feasibility study requires, prepared on Probable and Proven reserves defined to the confidence the codes demand. The nouns are not interchangeable, and the discipline of using them correctly is where credibility begins.

That discipline is the spine of this paper. A cost estimate is only as good as its basis of estimate, the documented scope, assumptions, exclusions and pricing that let a reviewer test it (Recommended Practice 34R-05). Its accuracy range is only honest once an appropriate contingency has been applied, sized so that overrun and underrun are equally likely rather than chosen to flatter the economics (Recommended Practices 40R-08 and 47R-11). And its class is only real when the project definition behind it actually exists.

Aurus works across this full ladder. Our consultants have prepared estimating mandates spanning AACE Class 5 to Class 1: delivered at Class 5 on scoping-level studies, and contracted to Class 3 and Class 1 accuracy on definitive studies for developers and lenders. We treat the basis of estimate as the deliverable that matters, and we say what class an estimate is, and what it is not.

The closing chapter offers a maturity model any owner or lender can apply: a way to read where an estimate sits, what its number can and cannot support, and what definition work stands between today's figure and a decision that will hold. The premise throughout is the one printed on the cover of every paper in this series. Evidence before assertion.

AT A GLANCE

# Five figures that frame the estimating problem

DEFINITION, NOT LABEL

**5**

AACE ESTIMATE CLASSES, CLASS 5 TO CLASS 1

Set by level of project definition alone. AACE RP 18R-97.

CLASS 5 ACCURACY

**+100%**

HIGH SIDE OF A CONCEPT-SCREEN ESTIMATE

Range -20% to -50% low, +30% to +100% high. AACE RP 18R-97, Table 1.

CLASS 1 ACCURACY

**10%**

ORDER OF A CHECK-ESTIMATE ACCURACY BAND

Range -3% to -10% low, +3% to +15% high. AACE RP 18R-97, Table 1.

THE MATURITY GAP

**14%**

MEAN AS-BUILT OVERRUN VS FEASIBILITY ESTIMATE

63 mining and smelting projects. Bertisen and Davis, 2008.

CONTINGENCY, DONE RIGHT

**50%**

TARGET PROBABILITY OF OVERRUN EQUALS UNDERRUN

The P50 test for a Class 3 mining estimate. AACE RP 47R-11.

HOW TO READ THIS PAGE

EACH FIGURE IS A PUBLISHED STANDARD VALUE OR A PEER-REVIEWED FINDING, CITED ON THE TILE. NO NUMBER HERE IS DERIVED, AVERAGED, OR ATTRIBUTED TO AURUS. THE PAPER CARRIES THE FULL SOURCES IN ITS REFERENCE CHAPTER.

> An estimate does not become more accurate because someone calls it a feasibility study. It becomes more accurate because the project is better known, and the class number is simply the honest record of how well.

Each figure returns later in the argument: the full five-class matrix in Exhibit 2, the way accuracy narrows and effort climbs in Exhibit 9, and the whole framework gathered into one shape in the signature confidence funnel, Exhibit 18.

DEFINITION SETS THE CLASS · THE CLASS SETS THE ACCURACY · THE ACCURACY IS THE HONEST WIDTH AROUND THE NUMBER

1

BUILD · THE MATURITY GAP

# Why estimates miss before they are wrong

Capital projects rarely fail at the moment of overrun. They fail earlier, when a number is asked to carry more confidence than the work behind it can support.

**63**
MINING AND SMELTING PROJECTS REVIEWED

**~14%**
MEAN AS-BUILT CAPITAL OVERRUN VS FEASIBILITY ESTIMATE

**1 in 13**
PROJECTS OVERRAN BY 100% OR MORE

There is a familiar sequence to a troubled mine build. A study is published with a headline capital figure. Financing is arranged against it. Construction begins, quantities grow, schedules slip, and a second number appears that bears little resemblance to the first. The post-mortem asks what went wrong during execution. The more useful question is what the first number was ever entitled to claim.

The evidence that this is a systemic problem, not a run of bad luck, is now decades old and well documented. In a study of 63 international mining and smelting projects, as-built capital costs averaged about 14 percent above the estimate carried in the bankable feasibility study. That average conceals a wider truth: roughly half of the projects finished outside a band of plus or minus 15 percent around the feasibility figure, and overruns of 100 percent or more occurred in about one project in thirteen. The authors found little sign that the bias shrank over time (Bertisen and Davis, The Engineering Economist, Vol. 53, No. 2, 2008).

A cost overrun of that scale is not a modelling error. It is the visible symptom of an estimate that was classified above the maturity of the project it described. The feasibility label promised control-grade confidence; the underlying definition supported something looser. The gap between the two is where value is destroyed, and it opens long before the first earthworks contract is let.

What makes the finding so useful is its persistence. The same study found little sign that the bias shrank across the two decades it covered, which tells us the problem is not a shortage of skill that better software or more experience will cure. It is structural. It lives in the decision to classify an estimate above the maturity of the project, and it will recur on the next project and the one after unless the classification itself is disciplined. That is an encouraging conclusion, because a structural problem has a structural remedy, and the remedy is already written down.

## The principle that closes the gap

AACE International, the professional body for cost engineering, resolves the problem with a rule that sounds almost too simple. In its Recommended Practice 18R-97, the class of a cost estimate is determined by one thing: the maturity of project definition. End usage, methodology, preparation effort and expected accuracy all follow from definition; they do not set the class themselves. An estimate is Class 3 because the project is defined to a Class 3 level, not because a study cover says feasibility.

Definition is expressed as a percentage of the engineering and project deliverables that a complete project would carry. A concept screen sits at 0 to 2 percent of full definition. A bankable feasibility study sits somewhere between 10 and 40 percent. A control estimate ready to govern construction sits above 50 percent. The percentages are indicative rather than measured to a decimal, but the ordering is absolute, and it is the ordering that matters: accuracy is bought with definition, and it cannot be bought any other way.

The consequence of ignoring that ordering is not abstract; it is written into the historical record. When a project is priced at a definition it has not reached, the estimate is not merely optimistic. It is uninformed about its own scope, and the definition it never carried returns later as cost the number could not have anticipated. The distribution in Exhibit 1 shows how wide that gap runs across real projects: the spread is not a tail of unlucky builds but the ordinary shape of estimates classified above the maturity that earned them, and reading it as bad luck is the first mistake the framework is built to prevent.

EXHIBIT 1

### Half of mine projects finish outside a plus or minus 15 percent band

Distribution of as-built capital cost against the bankable-feasibility estimate, 63 mining and smelting projects.

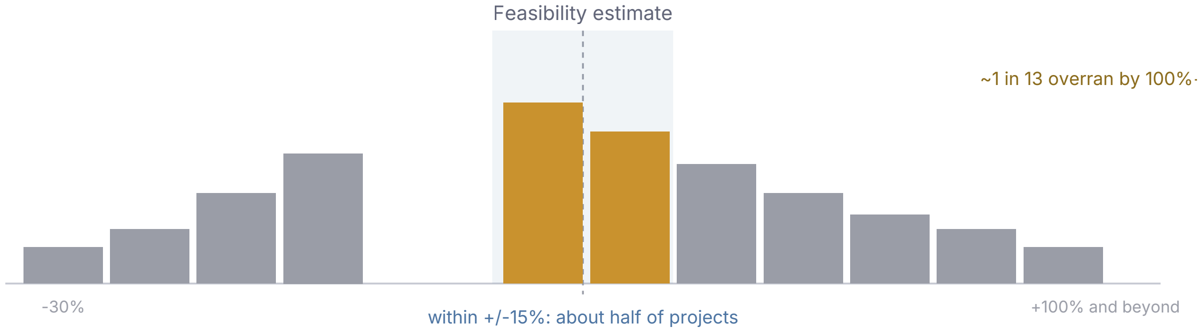


Shape illustrative of the cited findings; category counts not published as a table. Source: Bertisen and Davis, The Engineering Economist, 2008 (63 projects). Framework exhibit, no Aurus project data.

Two forces drive the overrun bias in the same direction. The first is technical: early estimates rest on limited engineering, and the unknowns that remain almost always add scope rather than remove it. The second is behavioural: a project seeking finance has an incentive to present its economics favourably, and an estimate is the easiest place for optimism to hide. A mature estimating practice cannot remove the second force, but it can make it visible, by stating the class honestly and by pricing the uncertainty that the class implies.

## Accuracy is not precision

A first source of confusion is the difference between accuracy and precision. Precision is how finely a number is expressed; accuracy is how close it sits to the eventual truth. An estimate quoted to the nearest thousand dollars looks precise, and that precision is routinely mistaken for confidence. It is nothing of the kind. A Class 5 concept screen can be carried to eight figures and still be entitled to a range of minus 50 to plus 100 percent, because its accuracy is set by the definition behind it, not by the number of digits on the page. The AACE ranges are a corrective to that illusion: they attach an honest width to every number, and they widen it precisely where a project is least known. A reader who takes one idea from this paper should take this one. The decimal places are not the confidence.

## The anatomy of an overrun

When a project overruns, the growth almost never comes from a single dramatic surprise. It accumulates from four ordinary sources, and each maps to a weakness in the early estimate. Quantities grow as design matures and the drawings reveal the steel, concrete and cable that an early layout only implied. Scope grows as items carried as vague allowances become defined items with real cost. Prices move as the pricing basis ages and escalation is under-provided. And productivity disappoints when the labour and site assumptions prove optimistic. A mature estimate anticipates all four, not by predicting them one by one, but by holding a contingency sized to the class and by naming its allowances honestly. An immature estimate meets all four for the first time during construction, when the only remaining variable is the overrun.

The value at stake is not abstract. A mineral project sanctioned on a feasibility figure it cannot honour will draw down its financing, exhaust its contingency, and return to its lenders for more, at a moment when its negotiating position is weakest and its options are fewest. The correction costs far more than the diligence that would have caught the mislabelled class at study stage. This is why the classification framework is not a documentation exercise for the estimating team but a governance instrument for the board and the lender: it is the cheapest control available against the most expensive mistake a resource project can make.

THEMATIC ASIDE

### Why the bias runs one way

If early estimates were merely imprecise, overruns and underruns would be equally common and the average error would be near zero. It is not. The record shows a persistent tendency to underestimate, which points to something beyond noise. Two mechanisms explain it. Unknowns at low definition are asymmetric: an undrawn system can only add cost once it is drawn, never remove it. And a project competing for scarce capital faces a standing incentive to present favourable economics, which makes the estimate the natural place for optimism to settle. Neither mechanism is dishonest by necessity, and both are answered by the same control: state the class truthfully, and price the uncertainty it implies.

> **Accuracy is bought with definition, and it cannot be bought any other way. A class number is a statement about how much of the project is known, not a promise about what it will cost.**
>
> AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

## What this paper does

The chapters that follow build from the framework to the practice. Chapter 2 sets out the five AACE classes and maps them onto the mineral study ladder that securities codes and lenders already use, and it defends the study nouns that the industry routinely abuses. Chapter 3 examines the basis of estimate, the document that makes an estimate testable, and the deliverables that move a project up the definition scale. Chapter 4 addresses contingency and the accuracy range, the mechanism by which an estimate becomes honest about its own uncertainty. Chapter 5 follows the ladder to a bankable decision, through the stage gates and validation steps that a lender expects. Chapter 6 sets out why mineral estimates are harder than industrial ones, where the resource, the corridor, and the long tail of capital compound the uncertainty an estimate must carry. Chapter 7 offers a maturity model that an owner or a financier can apply to any estimate placed in front of them.

The paper is written for two readers in particular. The first is the owner or sponsor who receives an estimate and must decide how much weight it can bear, and who is rarely told in plain terms what the class number on the cover actually promises. The second is the lender or investment committee that must test an estimate written by someone else, against a headline figure that has already been socialised. For both, the argument is the same: the class is a claim about definition, the basis of estimate is where that claim is checked, and the accuracy range is the honest width around a number no discipline can make narrower than the project's own maturity allows.

Throughout, the discipline is the same one Aurus applies on mandate: name the class, document the basis, price the risk, and never let a noun outrun the definition beneath it.

**The question is never what does it cost. It is how much of the project does this number already know.**

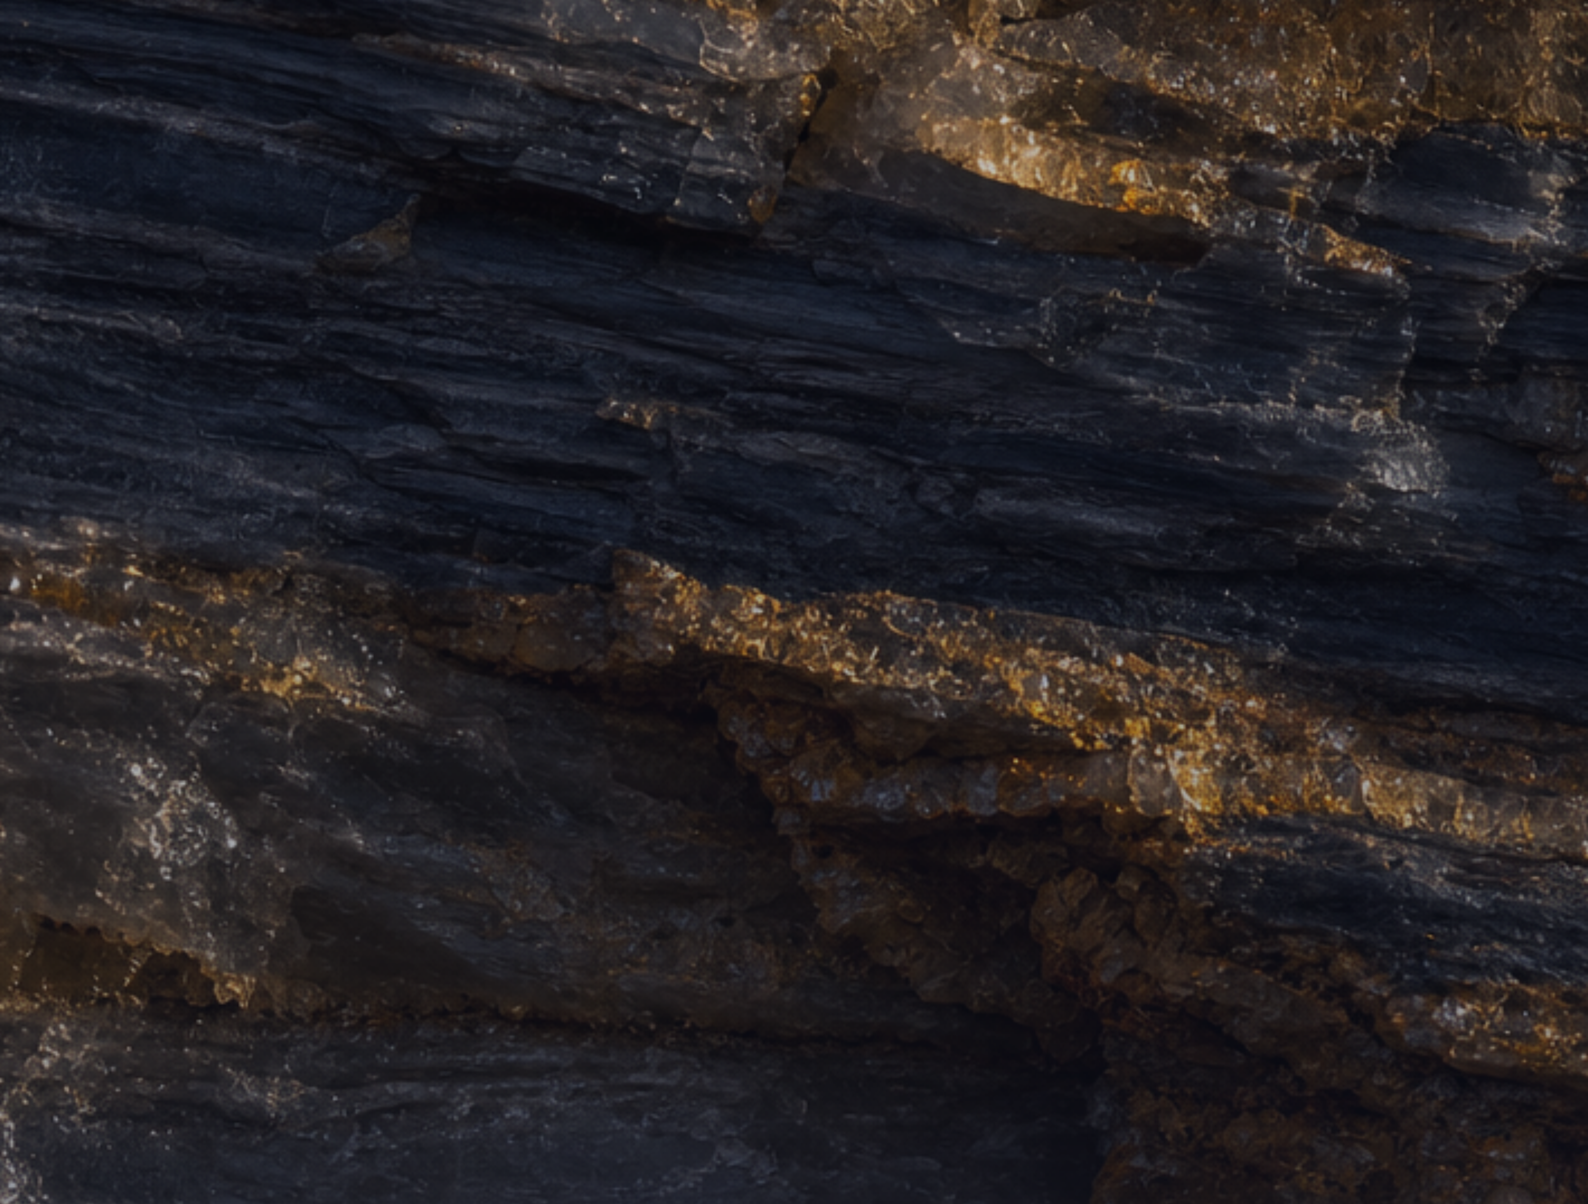

# 2

BUILD · THE CLASSIFICATION FRAMEWORK

## Five classes and the mineral study ladder

The AACE system is not a filing scheme. It is a contract between a level of knowledge and a permitted claim, and it fits the mineral study ladder exactly.

---

**5 to 1**
ESTIMATE CLASSES, SCREENING TO CHECK ESTIMATE

**PEA**
THE MINING STUDY A CLASS 5 ESTIMATE SUPPORTS

**Class 3**
THE ESTIMATE A BANKABLE STUDY REQUIRES

AACE International publishes its classification in two registers relevant here: a generic matrix in Recommended Practice 17R-97, and industry-specific versions that apply the same logic to particular sectors. Recommended Practice 18R-97 covers process-industry engineering, procurement and construction, and Recommended Practice 47R-11 applies the system to mining and mineral processing. The structure is identical across all three; only the deliverables used to judge definition change.

## The five classes

Each class is defined by a band of project definition, a typical purpose, a characteristic estimating methodology, and an expected accuracy range. The accuracy range is stated as a low and a high value, and it describes the typical variation of the actual cost from the estimate after an appropriate contingency has been applied, quoted at an 80 percent confidence interval. Exhibit 2 sets out the full matrix. It is the reference the rest of this paper returns to.

EXHIBIT 2

### The AACE five-class estimate classification, applied to mineral projects

Level of project definition governs class; accuracy ranges are after contingency, at 80 percent confidence. Mining study-stage names per RP 47R-11.

| CLASS | PROJECT DEFINITION | MINING STUDY STAGE | METHODOLOGY | ACCURACY LOW | ACCURACY HIGH |
|---|---|---|---|---|---|
| Class 5 | 0% to 2% | Concept screen; PEA | Parametric, judgment | -20% to -50% | +30% to +100% |
| Class 4 | 1% to 15% | Pre-feasibility | Factored, parametric | -15% to -30% | +20% to +50% |
| Class 3 | 10% to 40% | Feasibility, bankable | Semi-detailed unit costs | -10% to -20% | +10% to +30% |
| Class 2 | 30% to 70% | Control, bid or tender | Detailed unit cost, forced take-off | -5% to -15% | +5% to +20% |
| Class 1 | 50% to 100% | Check estimate, bid or tender | Detailed unit cost, detailed take-off | -3% to -10% | +3% to +15% |

Source: AACE International, Recommended Practice 18R-97, Table 1 (definition, methodology, accuracy); Recommended Practice 47R-11 (mining study-stage mapping). The two Recommended Practices carry one accuracy structure and differ only in the deliverables used to judge definition (see the sidebar, one system, three registers). Archived: AACE 18R-97 process-industries and generic captures, 2026-07-09.

Two features of the matrix repay attention. The accuracy ranges are themselves ranges, not single numbers, and they overlap between adjacent classes. A Class 5 estimate for a repeat project with good cost history can be as accurate as a Class 3 estimate for a project using unproven technology (18R-97, Figure 1). The class is a statement about definition; the accuracy that definition delivers still depends on how novel and how well benchmarked the project is. The second feature is the effort column, examined in Chapter 4: the work required to move down the classes rises far faster than the definition percentage suggests.

## How each class is actually built

The methodology column of the matrix is not incidental. It describes how the number is produced, and the method is what the definition allows rather than what the estimator prefers. At Class 5 the estimate is stochastic: cost-capacity curves, ratios to analogous plants, and parametric models that turn a capacity figure and a location into an order-of-magnitude total in hours rather than weeks. At Class 4 the method becomes semi-stochastic: equipment is factored, with installed cost built up from major-equipment cost through Lang or Guthrie-type factors, and gross unit rates applied to preliminary quantities. At Class 3 the estimate turns deterministic in its spine: semi-detailed unit costs are applied to quantities taken off a developing design, assembly by assembly, and major equipment is priced from budget quotations rather than factored. At Class 2 and Class 1 the estimate is built from a forced detailed take-off against near-complete drawings, priced from firm quotations and committed rates. Each step down the ladder replaces an assumption with a measurement, and that substitution, repeated across thousands of line items, is what narrows the accuracy range.

> THEMATIC ASIDE
>
> **One system, three registers**
>
> AACE publishes the classification in a generic register and in industry-specific ones. Recommended Practice 17R-97 is the non-industry matrix. Recommended Practice 18R-97 applies the same logic to engineering, procurement and construction in the process industries, the register that governs a mineral-processing plant. Recommended Practice 47R-11 applies it to mining and mineral processing as a whole and names the study stages a geologist and a lender already use. The five classes and their accuracy ranges are identical across the three; only the deliverables used to judge definition change from a chemical plant to an open pit. A paper on mineral projects can therefore cite 18R-97 for the accuracy structure and 47R-11 for the study-stage names without contradiction, because they are the same framework seen from two industries.

# What moves an estimate down a class

Because the class is set by definition, moving an estimate to a tighter class is a matter of producing specific deliverables, not of reworking the arithmetic. The step from Class 4 to Class 3 is the clearest example. It requires the process design to firm up from block flows to developed flowsheets and preliminary piping and instrumentation diagrams; the plant layout to advance from a conceptual arrangement to a preliminary general arrangement with defined footprints; major equipment to move from factored costs to budget quotations; civil and structural scope to be quantified rather than allowed for; and the mine plan to rest on Indicated moving to Probable and Proven reserves rather than Inferred material. Each deliverable replaces a stochastic estimate with a measured one, and the accuracy range narrows as a consequence. A project that wants a Class 3 number without funding these deliverables is asking for a class it has not earned, and no amount of estimating effort on the existing definition will produce one.

# The mineral study ladder

The value of Recommended Practice 47R-11 is that it speaks the language of the securities codes. Its mapping is direct. A Class 5 estimate is the cost engine of a Preliminary Economic Assessment, generally prepared on Inferred Resources delineated by early drilling. A Class 4 estimate belongs to a pre-feasibility study. A Class 3 estimate is what a feasibility study, including a bankable or definitive feasibility study, requires, and it must be built on Probable and Proven reserves defined to the confidence the codes demand. The estimate class and the resource category advance together, because neither can lead the other honestly.

This coupling is what the codes enforce and what a disciplined lender reads first. An estimate cannot claim a class the underlying resource has not earned, and a resource defined only to Inferred confidence cannot be governed by a Class 3 cost estimate whatever the study cover says. Exhibit 3 sets the three ladders side by side, in the order code-compliant reporting already assumes they climb.

EXHIBIT 3

**Estimate class, study stage and resource confidence advance as one**

The crosswalk lenders and code-compliant reporting already assume.

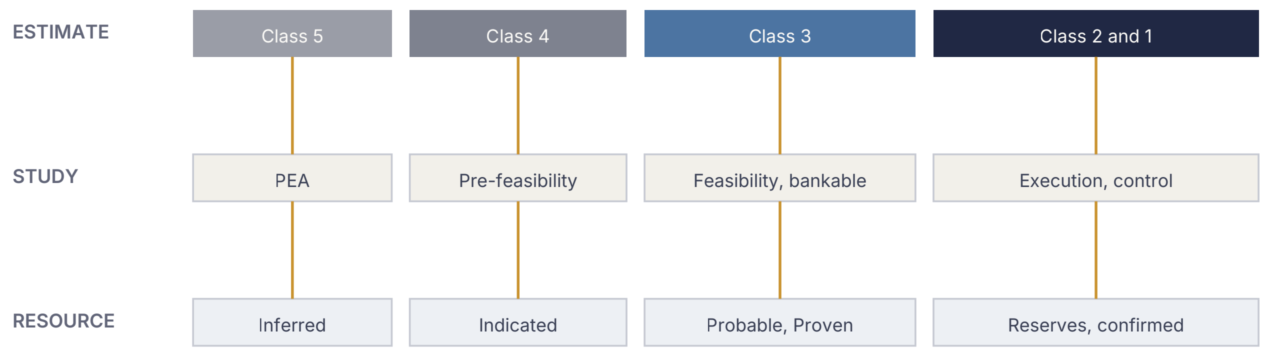


Source: AACE International, RP 47R-11 (class to study-stage mapping); CRIRSCO-family resource and reserve categories (JORC 2012, NI 43-101 CIM Definition Standards). Framework exhibit, no Aurus project data.

# The study-noun discipline

Because the class and the study stage are locked together, the words used to name a study carry technical weight. The nouns feasibility study, definitive feasibility study and bankable feasibility study are reserved for Class 3-or-better estimating on declared Mineral Resources. Work performed at Class 5 accuracy over an exploration target is a scoping-level techno-economic study, whatever its cover page claims. A study that attaches a feasibility label to a concept-grade estimate is not being ambitious; it is misrepresenting its own confidence, and an experienced technical reviewer reads the mismatch immediately.

The same discipline governs the resource base. An exploration target is not a Mineral Resource, and under the CRIRSCO-family codes it must be expressed as a range with a cautionary statement, never as a point estimate, and it must never carry project economics. A Class 5 study can and should model the economics of a concept, but it must state plainly that the tonnage it prices is a target, that there has been insufficient work to estimate a Mineral Resource, and that it is uncertain whether further work will define one. Honesty about the resource is inseparable from honesty about the estimate.

The securities codes reinforce the discipline from their own side. Under NI 43-101 a Preliminary Economic Assessment may consider Inferred Resources but must carry the cautionary language that its economics are preliminary, while a feasibility study must rest on Mineral Reserves and cannot be supported by Inferred material at all. The estimate class and the permitted study noun are therefore governed twice over, once by AACE from the cost side and once by the reporting codes from the resource side, and the two point in the same direction. A study that satisfies one but not the other is not a borderline case; it is a study whose authors have not reconciled what they are allowed to claim with what they have actually defined.

# Modifying factors: why the estimate and the reserve mature together

The link between estimate class and resource confidence is not a coincidence of naming. Under the CRIRSCO-family codes, a Mineral Resource becomes an Ore Reserve only after the application of the modifying factors: the mining, processing, metallurgical, infrastructure, economic, marketing, legal, environmental, social and governmental considerations that decide whether the material can actually be extracted and sold at a profit. Almost every one of those factors is also a cost input. Infrastructure and processing definition drive capital cost; the economic and marketing factors set the prices and escalation the estimate assumes; the environmental and social factors carry closure, resettlement and compliance cost. A project cannot hold Proven and Probable reserves without having defined the very things a Class 3 estimate must price. The reserve and the estimate mature on the same evidence, which is why a bankable study needs both at once and why a Class 5 estimate, resting on Inferred material and an exploration target, cannot honestly be called bankable no matter how the number is dressed.

THEMATIC ASIDE

### The Competent Person and the estimate

The reporting codes require that a public resource or reserve statement be signed by a Competent Person under JORC 2012 and SAMREC, or a Qualified Person under the CIM Definition Standards incorporated by NI 43-101. That named professional takes responsibility for the confidence category, and by extension for the modifying factors that a reserve depends on. The cost estimate sits inside that responsibility, not beside it: a reserve declared on a capital and operating estimate that the definition cannot support is a reserve declared on sand. The estimate class and the reporting sign-off are two views of one judgment about how well the project is known.

AURUS PRACTICE NOTE

Aurus prepared the scoping-level techno-economic study, at AACE Class 5 accuracy, for a US$2.2-billion potash and magnesium complex concept on the Central African coast: 2 million tonnes a year of muriate of potash with salt, magnesia and magnesium-metal co-products, and integrated mine-to-port infrastructure, over an exploration target defined from six drill holes and extensive 2D and 3D seismic surveys. The target is conceptual: there has been insufficient exploration to estimate a Mineral Resource, and it is uncertain whether further exploration will result in one. The study was scoped, presented and signalled as Class 5 throughout, which is precisely what a concept of that scale is entitled to claim.

The discipline this chapter argues for is finally a discipline of naming. A study earns its noun from the definition beneath it, not from the ambition of the party that commissioned it, and the same holds for the resource category a number is allowed to assume. When the noun and the definition agree, the estimate can be read at face value by anyone who knows the framework. When they diverge, the gap does not disappear; it is deferred to the party least able to absorb it, usually the lender asked to finance the difference or the owner asked to explain it after the fact. Naming the class and the noun to the same fact is the cheapest honesty a resource project has, and it is available before a single quantity is priced.

> The class number and the study noun are two names for the same fact: how much of the project is actually known.

# 3

BUILD · THE BASIS OF ESTIMATE

## Definition is the input, the basis is the proof

An estimate a reviewer cannot test is not an estimate, it is an assertion. The basis of estimate is the document that makes a number answerable.

**0 to 100%**
THE DEFINITION SCALE THAT SETS THE CLASS

**34R-05**
THE AACE BASIS-OF-ESTIMATE RECOMMENDED PRACTICE

**1 doc**
THAT DECIDES WHETHER AN ESTIMATE CAN BE AUDITED

If the class of an estimate is set by the maturity of project definition, then the honest first question about any estimate is not what does it cost, but how much of the project does this number actually know. That question is answered by two things: the deliverables that define the project, and the basis of estimate that records how the number was built from them.

## Definition is a set of deliverables, not a feeling

AACE frames project definition as a maturity matrix of specific engineering and project deliverables. The degree to which each deliverable exists, from a preliminary block flow diagram through to issued-for-construction drawings, is what places a project on the 0 to 100 percent definition scale and therefore fixes the estimate class. Definition is countable. A reviewer can ask which deliverables are complete, which are preliminary, and which are still assumed, and the answer places the estimate on the ladder without reference to what anyone wishes it to be.

EXHIBIT 4

**Definition matures deliverable by deliverable, and the class follows**

Illustrative maturity of representative deliverables against estimate class.

| DEFINITION DELIVERABLE | CLASS 5 | CLASS 4 | CLASS 3 | CLASS 2 AND 1 |
|---|---|---|---|---|
| Process and mass balance | Assumed | Preliminary | Developed | Complete |
| Plant layout and general arrangements | None | Conceptual | Preliminary | Detailed |
| Equipment list and datasheets | Analogue | Factored | Priced quotes | Purchase orders |
| Civil, structural and infrastructure design | Allowance | Sized | Quantified | Take-off |
| Mine plan and reserve basis | Inferred, target | Indicated | Probable, Proven | Confirmed |
| Execution plan and schedule | Notional | Outline | Integrated | Resource-loaded |

Structure per the AACE project-definition maturity matrix, RP 18R-97 and RP 47R-11. Framework exhibit, no project data.

## The basis of estimate: the estimate's own defence

The basis of estimate is the companion document that lets a reader reconstruct the number. AACE defines it as a recommended practice in its own right (Recommended Practice 34R-05). A complete basis states the scope covered and the scope excluded, the source of every major quantity and unit price, the currency and exchange assumptions, the escalation basis and the pricing date, the productivity and labour assumptions, the allowances carried for undefined scope, and the contingency method. It names what is quoted, what is factored, and what is assumed. Where a headline figure is a single line, the basis is where its integrity lives.

A basis of estimate is also what makes an estimate portable between the people who must rely on it. The estimator who built the number will not be in the room when the board sanctions it or when the lender's engineer tests it, and a headline figure carries none of the reasoning that produced it. The basis is how that reasoning travels: it lets a reviewer who never met the team confirm what was priced from a firm quotation, what was factored from a database, and what was still an assumption on the day the number was struck. Exhibit 5 sets out the sections that basis should contain.

EXHIBIT 5

### The anatomy of a defensible basis of estimate

The eight sections a reviewer should be able to find and test.

| | | | |
|---|---|---|---|
| Scope and battery limits<br>included and excluded | Estimate class and purpose<br>definition basis stated | Quantity basis<br>take-off, factored, assumed | Pricing basis<br>quotes, databases, date |
| Currency and exchange<br>rates and reference date | Escalation<br>basis and to-date | Allowances<br>for undefined scope | Contingency and risk<br>method and confidence |

Every headline number resolves down through these sections to a quote, a database, or a stated assumption.
An estimate with no traceable basis cannot be classified, because its definition cannot be examined.

Structure per AACE International, RP 34R-05, Basis of Estimate. Framework exhibit, no project data.

## From deliverables to a number

A well-built estimate assembles the cost the way the project will be built: a work breakdown structure that mirrors the execution plan, quantities drawn from the design at the maturity the class allows, unit prices from quotes or maintained cost databases, and factored allowances only where definition genuinely does not yet exist. The build-up in Exhibit 6 is deliberately ordinary. Its discipline is that each layer is sourced, and the sourcing is what the basis of estimate records. An estimate that jumps from a capacity figure to a total, with the middle layers absent, is a Class 5 number wearing a Class 3 label.

EXHIBIT 6

### A number is only as sound as the layers beneath it

The build-up from scope to a sanctioned total.

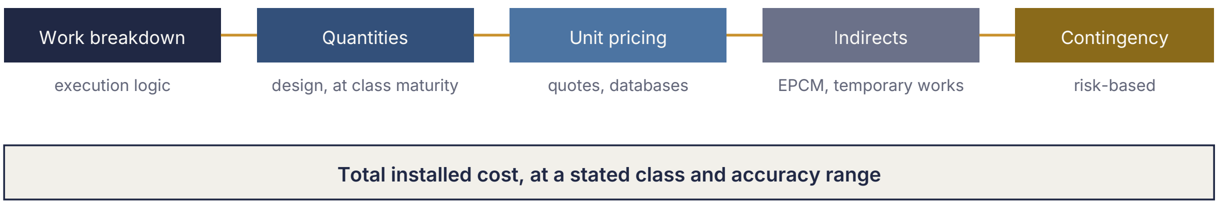


Base estimate (direct plus indirect) carries the deterministic cost; contingency carries the residual uncertainty (Chapter 4).

Structure per AACE estimating practice, RP 18R-97 methodology column and RP 34R-05. Framework exhibit, no project data.

## The costs the headline forgets

Direct costs, the equipment, materials and installation labour of the plant and mine, are the part of an estimate that receives the most attention and causes the fewest surprises. The overruns tend to live in the categories a thin estimate treats lightly. Indirect costs cover engineering, procurement and construction management, temporary works, construction power and water, camps and logistics, commissioning and spares. Owner's costs cover the project team, permitting, land, owner's insurance and financing during construction. Escalation carries the estimate from its pricing date to the midpoint of expenditure, and in a long build it is a material line in its own right. A mature Class 3 estimate carries all of these explicitly, each with its own basis; an immature one buries them in a single percentage and discovers their true size during execution.

Currency deserves the same discipline. A mineral project in Central Africa typically prices equipment in one currency, labour in another and fuel in a third, and it earns revenue in a fourth. The basis of estimate must state the reference rates and the reference date, because an estimate that is silent on currency is silently exposed to it. None of this is exotic. It is simply the difference between an estimate that has been thought through and one that has been rounded up.

## Structure is what makes an estimate reviewable

An estimate is only testable if it is organised so that a reviewer can find things. A code of accounts gives every cost a consistent address, so that direct and indirect costs, owner's costs, escalation and contingency can be read at any level of the work breakdown structure without double counting or gaps. The work breakdown structure itself should mirror the execution plan, so that the estimate, the schedule and the eventual contracts all speak the same language. When the estimate structure and the schedule structure diverge, the reconciliation between cost and time becomes guesswork, and guesswork is where contingency quietly disappears. A Class 3 estimate built on a disciplined code of accounts can be interrogated line by line; a lump-sum estimate with no structure can only be believed or disbelieved as a whole, which is another way of saying it cannot be reviewed at all.

## The estimate is reviewed before it is issued

A mature estimate is tested by its own authors before it ever reaches a decision-maker. An internal estimate review checks that the class claimed matches the definition held, that the basis of estimate is complete, that quantities reconcile to the current design, that pricing is current and sourced, and that contingency has been set by a method fit for the class rather than by habit. It reconciles the estimate against the previous class estimate for the same project and explains the movement, because an unexplained jump between a pre-feasibility and a feasibility number is itself a finding. This internal discipline is what a later independent review, described in Chapter 5, repeats from the outside. An estimator who has run the review honestly has little to fear from the lender's engineer; an estimator who has skipped it will meet the same questions at a far higher cost. The review is not bureaucracy. It is the moment the estimate is made to defend its own class before anyone acts on it.

The review is also the estimate's last private moment. After it, the number belongs to the decision it was built to support, and every weakness left in it becomes someone else's to find. A practice that treats the internal review as the real test, rather than a formality cleared before issue, is the one whose numbers survive contact with a lender's engineer.

> THEMATIC ASIDE
>
> **Allowance, contingency, reserve: three different things**
>
> Three terms are routinely blurred, and the blur is expensive. An allowance is money in the base estimate for scope that is real but not yet detailed, for example a lump sum for cabling before the cable schedule exists. Contingency is money for the cost growth of scope that is defined but uncertain, sized by risk analysis and held at P50 (Chapter 4). A management reserve sits outside the estimate entirely, held by the owner for changes in scope that have not yet occurred. AACE keeps them distinct because they answer different questions and are drawn down under different authority. An estimate that folds all three into one number cannot say what its contingency is actually covering, and a contingency that cannot be explained cannot be defended.

**The basis of estimate is where a number stops being a claim and becomes a document: the place a reviewer who never met the project can trace every figure back to a quote, a database, or an assumption stated in the open.**

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

Contingency is not optimism removed. It is uncertainty priced, so that the accuracy range can be believed.


4

BUILD · CONTINGENCY AND THE ACCURACY RANGE

# The number that makes a class honest

Contingency is not padding and not optimism. It is the price of the uncertainty the class implies, and it is what lets an accuracy range be believed.

**P50**
CONTINGENCY TARGET: OVERRUN AS LIKELY AS UNDERRUN

**80%**
CONFIDENCE INTERVAL BEHIND THE ACCURACY RANGE

**5 to 100**
RELATIVE EFFORT INDEX, CLASS 5 TO CLASS 1

The accuracy range attached to each class is not a decoration. It is a probabilistic claim about how far the actual cost may sit from the estimate, and that claim is only true once an appropriate contingency has been added. Contingency is where an estimate stops pretending that everything not yet known will cost nothing.

## What contingency is, and what it is not

AACE treats contingency as the funds added to the base estimate to cover cost growth from within-scope uncertainty, as part of the overall risk-management process (Recommended Practice 40R-08). It is not an allowance for scope that is simply not yet designed, which belongs in the base estimate as a named allowance. It is not a management reserve for changes in scope, which sits outside the estimate. And it is not a lever for making the economics look better. Contingency covers the variance that remains after the base estimate has done its honest best, and the accuracy range describes the distribution that variance follows.

## The P50 test

A contingency is set correctly when it is sized so that the project is as likely to come in under the estimate as over it. In mining practice, Recommended Practice 47R-11 states the convention plainly: contingency on a Class 3 estimate is typically set to achieve roughly a 50 percent probability of overrun versus underrun, with the plus or minus accuracy expressed at an 80 percent confidence interval. The moment contingency is trimmed to protect a headline number, the estimate ceases to sit at P50, the accuracy range becomes a fiction, and the overrun bias in Chapter 1 is being manufactured on purpose.

EXHIBIT 7

**Contingency to P50: overrun and underrun equally likely**

Cumulative probability of final cost; base estimate plus contingency set at the median.

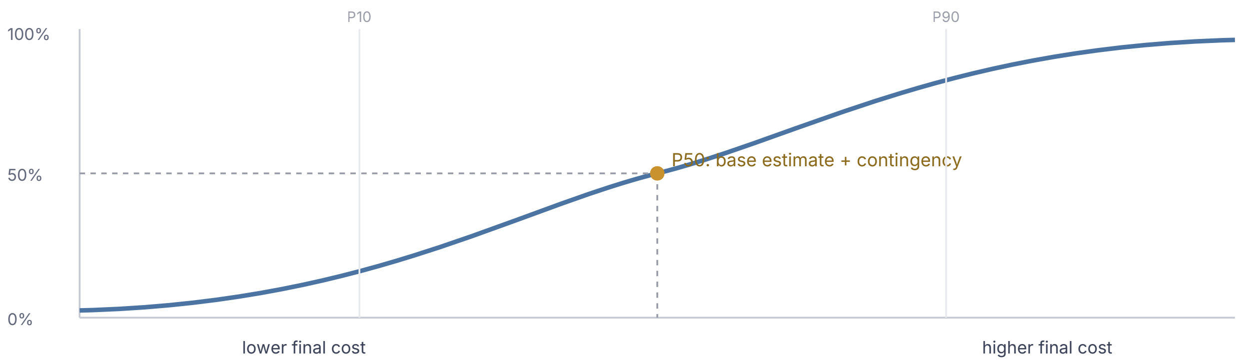


Convention per AACE International, RP 47R-11 (P50 target) and RP 18R-97 (80 percent confidence). Framework exhibit, no project data.

## How contingency is determined

The number itself can be reached in more than one way. A deterministic method applies risk-based percentages to estimate line items using judgment and precedent. A parametric method models contingency as a function of systemic drivers such as project definition, technology novelty and team experience (Recommended Practice 42R-08). An expected-value method builds a risk register and computes contingency from the probability and cost impact of identified risk events (Recommended Practice 44R-08), and a full quantitative risk analysis runs the estimate through a probabilistic model to produce the cumulative distribution shown in Exhibit 7 (Professional Guidance Document 02). Current AACE guidance treats the accuracy range and the contingency as a single question rather than two (Recommended Practice 119R-21). The right method depends on the class: a Class 5 estimate is served by a parametric or judgment approach, while a Class 3 estimate for a lender warrants a quantitative risk analysis.

EXHIBIT 8

### Contingency methods scale with the class they serve

Matching the rigour of the contingency method to the maturity of the estimate.

| METHOD | HOW IT WORKS | BEST FIT | AACE REFERENCE |
|---|---|---|---|
| Deterministic percentages | Risk-based percentage on line items, by judgment | Class 5 to 4 | 40R-08 |
| Parametric model | Contingency as a function of systemic risk drivers | Class 5 to 3 | 42R-08 |
| Expected value | Risk register, probability times cost impact | Class 3 to 2 | 44R-08 |
| Quantitative risk analysis | Probabilistic model of the full estimate | Class 3 to 1, lender grade | PGD-02, 119R-21 |

Source: AACE International, RP 40R-08, 42R-08, 44R-08, 119R-21; PGD-02, Guide to Quantitative Risk Analysis.

THEMATIC ASIDE

### What a risk register actually contains

An expected-value contingency is only as good as the register beneath it. A usable register names each risk, assigns it a probability and a cost impact with a stated basis, records whether it is a threat or an opportunity, and identifies an owner and a response. Summed probabilistically, that register produces the cumulative distribution in Exhibit 7 and the contingency that sits at its median. The value is more than the number it yields. It is also the conversation it forces: a risk that no one will price is usually a risk no one has thought about, and a register with three lines on a billion-dollar project is telling a reviewer more about the estimate's maturity than the headline ever will.

## The risks contingency does not cover

Contingency covers within-scope uncertainty. It does not, and should not, cover everything that can move a final cost. Escalation is not contingency: it is the expected movement of prices between the pricing date and the point of expenditure, and it belongs in the base estimate as a computed line, not in the risk allowance. Currency movement is not contingency: it is an exposure to be stated and, where the owner chooses, hedged, not a variance to be buried. Scope change is not contingency: it is the domain of the management reserve, drawn under the owner's authority when the project is deliberately altered. Force-majeure and truly exceptional events sit outside the estimate altogether. Keeping these categories separate is what lets a reviewer see whether a single contingency percentage is doing honest work or quietly absorbing four different risks it was never sized for. An estimate that cannot say which risk its contingency covers has not analysed its risk; it has averaged its anxiety.

## Sensitivity is not risk

A last distinction keeps the accuracy range honest. A sensitivity analysis asks how the outcome moves when one input is flexed, holding the rest still: what happens to project value if the metal price falls ten percent, or if capital is ten percent higher. It is a useful communication tool, and it is not a risk analysis. A risk analysis asks how all the uncertain inputs behave together, with their probabilities and their correlations, to produce the distribution of outcomes from which contingency and the accuracy range are read. Presenting a tornado of single-variable sensitivities as if it sized the contingency is a common error, because it ignores the joint behaviour that actually drives cost growth. A mature study uses sensitivities to explain what matters and a probabilistic model to size what is set aside, and it never confuses the one for the other.

# Accuracy ranges overlap, and effort compounds

Two properties of the framework decide how it should be read. First, the accuracy ranges overlap between classes, so a class label alone does not fix an accuracy: a well-benchmarked Class 5 estimate can outperform a Class 3 estimate for a first-of-a-kind process, which is why Table 1 publishes ranges rather than points. Second, the effort to prepare an estimate rises far faster than its definition percentage. On the AACE relative index, effort runs at 1 for a Class 5 estimate and reaches 5 to 100 for a Class 1, because detailed take-off, quoted pricing and quantitative risk analysis are expensive to produce. A sponsor who wants Class 3 confidence must fund Class 3 work; the accuracy cannot be conjured from a Class 5 budget.

EXHIBIT 9

### Ranges narrow and effort climbs as definition matures

Expected accuracy band (bars, drawn to the Table 1 values) and relative preparation effort (line), Class 5 to Class 1.

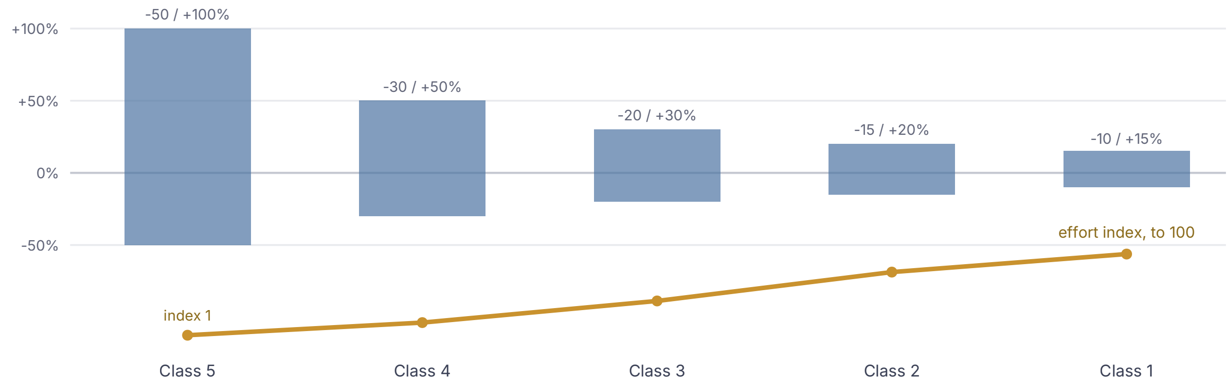


Bars drawn to the AACE RP 18R-97, Table 1 low and high accuracy values at a common vertical scale (0.8 px per percent), so the low side and high side of each class track the published figures. Gold line: the relative preparation-effort index from the same table (1 at Class 5, rising to 5 to 100 at Class 1), shown as a rising index rather than on the linear scale of the bars. Archived: AACE 18R-97 captures, 2026-07-09. Framework exhibit, no Aurus project data.

# Reading the range honestly

The accuracy range is a probabilistic statement, and it is worth reading as one. When AACE quotes a Class 3 range of minus 10 to plus 30 percent at an 80 percent confidence interval, it means that the eventual cost is expected to fall inside that band roughly four times in five, after the P50 contingency has been added. It does not mean the cost cannot fall outside it; one project in five is expected to. That is why a lender does not fund a project at its P50 estimate alone but sizes headroom against a higher percentile, and why a sponsor who reports only the P50 figure without its range has told half the story. The range is not a hedge or a disclaimer. It is the most honest single thing an estimate can say about itself.

THEMATIC ASIDE

### Why P50 for the estimate and a higher percentile for the funding

Setting contingency to P50 makes the estimate an unbiased best guess: the number the project is equally likely to beat or exceed. That is the right basis for comparing options and for reporting an expected cost. Funding is a different question. A lender or board wants confidence that money will not run out, so it provides against a higher percentile, often P70 to P90, through a funding buffer or a cost-overrun facility held above the estimate. The two numbers serve two purposes and should never be conflated: an estimate quietly set to P80 to look prudent is no longer an expected cost, and a funding plan built only to P50 is under-provided by design.

> A trimmed contingency does not make a project cheaper. It only moves the overrun from the estimate, where it can be managed, to the build, where it cannot.

5

BUILD · THE LADDER TO A DECISION

# From concept to a bankable decision

A financing decision does not need certainty. It needs an estimate whose class matches the decision being taken, and a basis a lender can independently test.

**Class 3**
THE LENDER-GRADE ESTIMATE A BANKABLE STUDY RESTS ON

**Gates**
DEFINITION THRESHOLDS, NOT CALENDAR DATES

**Class 5 to 1**
THE ESTIMATING SPAN AURUS WORKS ACROSS

A cost estimate exists to support a decision, and different decisions need different classes. A board screening a portfolio needs Class 5. A board choosing between development options needs Class 4. A board and its lenders sanctioning construction need Class 3, and increasingly Class 2 for the packages that will be committed first. The maturity ladder is the sequence by which a project earns the right to each decision in turn.

## Stage gates are definition thresholds

The discipline that keeps a project on the ladder is the stage gate. A gate is not a date on a schedule; it is a test of whether the definition deliverables that a class requires actually exist, applied before the project is allowed to claim that class and take the decision that depends on it. A project that passes a feasibility gate on the strength of a Class 4 definition has not accelerated; it has simply moved its overrun forward to a point where it is harder to correct. The gate exists to stop exactly that.

EXHIBIT 10

### Each decision is gated on the class it requires

Estimate class, the decision it supports, and the gate that guards it.

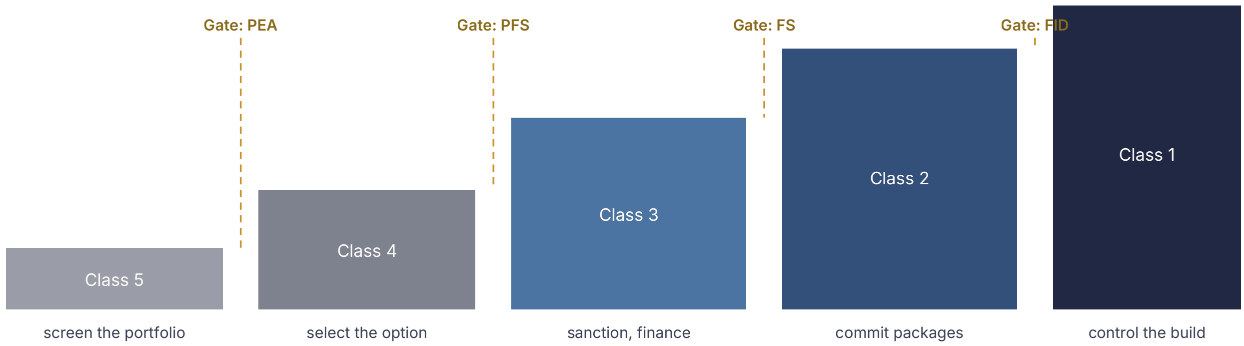


Class to decision mapping per AACE RP 47R-11; gate names per the mineral study ladder. Framework exhibit, no project data.

## Validation is what makes an estimate bankable

A lender does not accept an estimate because it is labelled bankable. It accepts an estimate because the estimate survives independent scrutiny. Lender-grade validation tests the basis of estimate against the class claimed, reconciles quantities to the design, benchmarks unit rates and productivity against comparable projects, confirms that escalation and currency assumptions are current, and checks that contingency has been set at P50 by a method appropriate to a Class 3 estimate rather than trimmed to protect the return. Where a parameter is genuinely open, a mature estimate names it as open and prices an allowance for it, rather than assuming it away. The checklist in Exhibit 11 is the shape of that scrutiny.

What makes the exercise lender-grade is not its length but its independence. A review run by the team that built the estimate can confirm the arithmetic but cannot see the assumption it shares with the number, and the assumptions that move a project are precisely the ones its authors no longer notice. An estimate becomes bankable when someone who did not write it, and whose interest is in finding the weak point rather than defending it, applies each of these tests and the number still holds.

EXHIBIT 11

### What lender-grade validation actually tests

The questions an independent review puts to a Class 3 estimate.

| TEST | QUESTION THE REVIEWER ASKS | PASS CONDITION |
|---|---|---|
| Class integrity | Does the definition support the class claimed? | Deliverables present at Class 3 maturity |
| Basis completeness | Can every major number be traced to a source? | Basis of estimate covers all eight sections |
| Quantity reconciliation | Do quantities match the current design? | Take-off reconciles to drawings and models |
| Pricing currency | Are rates and escalation current and sourced? | Quotes and databases within pricing window |
| Contingency adequacy | Is contingency at P50 by a fit-for-class method? | Quantitative or expected-value basis documented |
| Open parameters | Are unknowns named and allowed for? | No silent assumptions on material scope |

Structure per AACE RP 18R-97, 34R-05 and 47R-11 review practice. Framework exhibit, no project data.

This is the discipline Aurus brings to the infrastructure that carries a mineral project to market. The firm applies bankable feasibility and Category A environmental and social assessment practice to deep-water port and terminal infrastructure, prepared to AACE Class 3 accuracy for international lenders, where the estimating standard is set by what a lender must be able to test rather than by what a sponsor would prefer to show.

## Benchmarking: the outside view

The single most effective guard against the optimism in Chapter 1 is to check an estimate against what comparable projects actually cost, rather than only against its own build-up. An estimate assembled purely from the inside, line by line, inherits every optimistic assumption its authors hold. Benchmarking unit rates, productivity, indirect-cost ratios and total installed cost against a set of delivered projects gives an outside view that the inside view cannot supply, and it is the discipline most likely to catch a Class 4 estimate wearing a Class 3 label before a lender does. The overrun record exists precisely because the outside view is so often skipped: a project that benchmarked its feasibility estimate against the distribution in Exhibit 1 would rarely present a number at the optimistic tail without saying so. A mature practice runs both views and reconciles them, and it treats a large gap between the two as a finding to be explained rather than a discrepancy to be smoothed.

## The decision data pack

A final investment decision is taken on a defined body of evidence, not on a headline. The capital and operating estimates sit at the centre of that pack, but they are read alongside the basis of estimate, the risk register and quantitative risk analysis that justify the contingency, the execution schedule the estimate was built against, the reserve statement and its Competent Person sign-off, and the independent technical review that tested the whole. The maturity of the estimate is legible from the pack itself: a Class 3 decision pack is thick with traceable definition, while a pack that offers a number and little else is announcing its own class regardless of the label on the cover. The discipline of assembling the pack is the discipline of earning the decision, and it is the point at which every earlier shortcut becomes visible.

## Beyond first capital

A bankable decision rests on more than the initial capital estimate. The same maturity discipline extends to the numbers that decide whether a project pays back. Sustaining capital, the ongoing investment a mine needs to hold production through its life, is routinely under-estimated at feasibility because it is easy to defer on paper. Ramp-up cost and the shape of the ramp, the months between first production and steadystate output, sit at the join between the capital and operating estimates and are a frequent source of disappointment. Operating cost carries its own accuracy question: while AACE classes formally describe capital estimates, a mature study holds its operating and energy estimates to a comparable standard of definition, because a bankable return is as sensitive to unit operating cost as to first capital. A study that presents a precise capital figure alongside a thin operating estimate has simply moved the uncertainty from one line to another.

THEMATIC ASIDE

### The lender's own engineer

On a project financing, the sponsor's estimate is rarely the last word. The lenders appoint an independent technical adviser, often called the lender's or owner's engineer, to test the estimate against the same questions in Exhibit 11: is the class supported by the definition, is the basis complete and current, is contingency set to P50 by a fit-for-class method, and are open parameters named rather than assumed. A sponsor who has prepared the estimate to survive that review has little to fear from it. A sponsor who has trimmed contingency or labelled a Class 4 study bankable discovers the gap at precisely the moment it is most expensive to close. The discipline this paper describes is, in the end, the discipline of being ready for a reader who did not write the number.

## The estimate does not end at sanction

A control estimate is the baseline the project is then measured against, so the discipline continues into execution. Trends capture cost movements as they emerge, change control governs scope so that the management reserve is drawn deliberately rather than eroded by drift, and a regular forecast-at-completion reprices the remaining work against actual performance. Contingency is drawn down against the risks it was provided for, and its remaining balance is a live measure of how much uncertainty is left in the project. An estimate treated as a one-time number, filed at sanction and never reconciled, tells the owner nothing about where the build is heading until the overrun is already booked. An estimate maintained as a control baseline turns the same discipline that classified it into a means of steering, which is the whole point of having earned the class in the first place.

## Estimating across the full ladder

An estimating practice earns credibility by working honestly at every rung, not by claiming the bottom rung on every project. Aurus consultants have prepared estimating mandates spanning AACE Class 5 to Class 1: delivered at Class 5 on scoping-level studies, and contracted to Class 3 and Class 1 accuracy on definitive studies. The practice note below illustrates the top of that span, a mandate contracted to Class 1 accuracy, and it also illustrates the honesty the framework demands, because Class 1 is stated as what the contract requires, not as what a definitive feasibility study will typically achieve.

AURUS PRACTICE NOTE

Aurus was engaged under a multi-million-euro mandate to deliver the Definitive Feasibility Study for a potash (carnallite) solution-mining development in Central Africa. The mandate was contracted to AACE Class 1 accuracy on capital cost, with operating and energy cost estimates held to the same contractual accuracy requirement, alongside well-field and cavern design, long-term creep, triaxial and shear geomechanics with a full rock-mechanics model, complete plant engineering including HAZOP, P&IDs and single-line diagrams, and investment-grade economic evaluation. Class 1 accuracy is named here as the contractual estimating obligation the mandate set. Definitive feasibility practice more typically lands at Class 3, and a mature practice says so rather than presenting the tightest class as an industry norm.

**A lender does not accept an estimate because it is called bankable. It accepts an estimate because the estimate survives being tested by someone who did not write it.**

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

# 6

BUILD · THE MINERAL CONTEXT

## Why mineral estimates are harder

A mine is not a process plant on a serviced site. The framework is the same, but the definition a mineral project must reach is deeper, and more of it sits underground.

**2 unknowns**
THE OREBODY AND THE COST, EACH ESTIMATED

**mine to port**
INFRASTRUCTURE OFTEN THE LARGEST CAPITAL BLOCK

**decades**
OF SUSTAINING CAPITAL AND CLOSURE TO PRICE

The overrun record in Chapter 1 is a mining record, and that is not an accident. The classification framework holds for any capital project, but a mineral project must assemble a larger and less visible definition to earn a class, which makes the discipline more demanding rather than less. A mine is estimated twice over: once for what it will cost, and once for the orebody that is itself only an estimate, and some of that ground can never be fully seen.

## The resource is itself an estimate

A process plant on a serviced industrial site begins from a known feedstock. A mine begins from an orebody that is itself an estimate, expressed in confidence categories that carry their own uncertainty. A cost estimate built on Inferred Resources inherits that uncertainty and compounds it: the tonnes, the grades, the metallurgical recovery and the mine plan are all provisional, and every one of them drives cost. This is the deeper reason the estimate class and the resource category advance together. A Class 3 capital estimate on a Probable and Proven reserve is standing on defined ground; a Class 3 label on Inferred material is standing on a guess wearing a number. The compounding of two uncertainties, in the resource and in the cost, is what makes an honest mineral estimate harder to earn than an industrial one of the same nominal class.

EXHIBIT 12

### Two uncertainties compound in a mineral estimate

Resource confidence and cost definition advance on the same evidence.

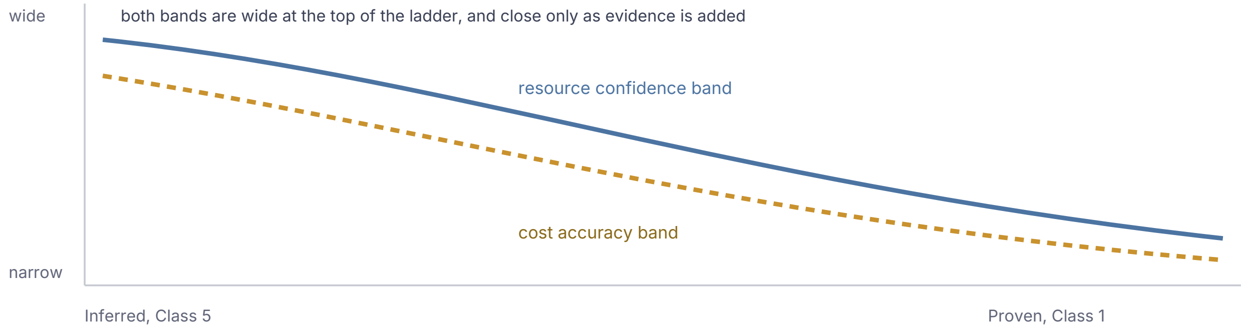


Schematic of the joint maturing of resource category and estimate class per AACE RP 47R-11 and the CRIRSCO codes. Framework exhibit, no project data.

## Greenfield ground and remote sites

A brownfield expansion inherits infrastructure, cost history and a workforce, so its definition matures quickly and its estimates benchmark against real precedent. A greenfield mineral project in a frontier region inherits none of that. It must define, and price, the access roads, power, water, camps and logistics that an industrial estimator assumes as given, and it must do so where cost data is thin and productivity assumptions are fragile. Remoteness turns indirect costs and escalation from footnotes into major line items, and it widens the accuracy range at any given level of definition, because the benchmarks that would narrow it do not exist. The framework accounts for this: RP 18R-97 notes that a novel or unbenchmarked project sits toward the wide end of its class range, exactly where a frontier mine belongs.

EXHIBIT 13

**The cost drivers a mineral estimate must define, beyond the plant**

Where mineral definition reaches further than a serviced-site industrial estimate.

| DRIVER | WHY IT IS HARDER IN MINING | ESTIMATE CONSEQUENCE |
|---|---|---|
| Orebody definition | Grade and tonnage are themselves estimated | Cost inherits resource uncertainty |
| Mine-to-port infrastructure | Rail, port, power and roads over long distances | Often the largest single capital block |
| Remoteness and logistics | Frontier sites, thin cost data, fragile productivity | Large indirects, wide accuracy band |
| Sustaining capital | Decades of ongoing investment to hold output | Under-provided at feasibility |
| Closure and rehabilitation | Mandatory late-life cost, long liability tail | A provision, not an afterthought |
| Commodity-price exposure | Revenue set by a market, not a contract | Drives the economic modifying factor |

Framework exhibit, drawn from AACE RP 47R-11 scope and the CRIRSCO modifying factors. No project data.

## The long tail of capital

A mine's cost does not end at first production. Sustaining capital holds output through a life measured in decades and is routinely thin in feasibility estimates because it is easy to defer on paper. Closure and rehabilitation are a mandatory late-life cost with a long liability tail, a provision the modifying factors require rather than an afterthought to be added later. And on the revenue side, a commodity price set by a market rather than a contract drives the economic modifying factor that decides whether a Resource is a Reserve at all. A mature mineral estimate prices the whole life, not just the build, because the decision it supports is a decision about the whole life. This is where integrated mine-to-market engineering earns its place: the corridor of rail, port and power that carries a mine to market is often the largest single block of capital, and it must be estimated to the same class as the plant it serves.

AURUS PRACTICE NOTE

Aurus designed the preliminary-engineering programme for an integrated mine-to-market logistics chain in Central Africa: mineral-port siting, railway alignment and high-voltage transmission routing structured as a single system, staged from field reconnaissance and multi-disciplinary site characterisation through the global development scheme, structural sizing, execution drawings and cost estimates. Treating the corridor as one system, and carrying its cost estimate to the same discipline as the process plant, is what keeps the largest block of mineral capital from being the least defined.

> The framework does not get easier for a mine. The definition simply reaches further, and more of it sits underground.

**Name the class. Document the basis. Price the risk. Never let a noun outrun the definition beneath it.**

THE DISCIPLINE OF CONFIDENCE · CHAPTER 7

# 7

BUILD · THE AURUS MATURITY MODEL

## Reading any estimate placed in front of you

A capability instrument, not project data: a way for an owner or a lender to locate an estimate on the ladder and decide what its number is allowed to support.

**4 lenses**
DEFINITION, BASIS, CONTINGENCY, VALIDATION

**1 question**
HOW MUCH OF THE PROJECT DOES THIS NUMBER KNOW

**2 readers**
THE OWNER WHO DECIDES, THE LENDER WHO FUNDS

The preceding chapters describe the framework and the practice. This one turns them into an instrument. The Aurus cost-estimate maturity model is a way to read any estimate, from any source, and answer three questions: what class is this estimate really, what decision can it honestly support, and what definition work stands between it and the next decision. It is a diagnostic lens, and it carries no project data.

## Four lenses

The model reads an estimate through four lenses, each drawn from the chapters above. Definition: which deliverables exist, and at what maturity. Basis: whether the estimate can be traced to sources a reviewer can test. Contingency: whether the risk has been priced at P50 by a method that fits the class. Validation: whether the estimate has survived scrutiny by someone independent of its authors. An estimate that scores well on all four is what its label says. An estimate that scores well on none is a headline in search of a project.

EXHIBIT 14

### The four-lens maturity read

What each lens asks, and what a mature answer looks like.

| LENS | DIAGNOSTIC QUESTION | IMMATURE SIGNAL | MATURE SIGNAL |
|---|---|---|---|
| Definition | Which deliverables exist, at what maturity? | Class claimed above deliverables held | Class matches the definition matrix |
| Basis | Can every major number be traced? | Headline with no basis of estimate | Basis covers scope, quantities, pricing, risk |
| Contingency | Is risk priced at P50, fit for class? | Contingency trimmed to protect returns | P50 by quantitative or expected-value method |
| Validation | Has an independent reviewer tested it? | Self-certified by the authoring team | Independent review against the class |

Aurus capability instrument, synthesising AACE RP 18R-97, 34R-05, 40R-08 and 47R-11. No project data.

## The four failure modes, and their controls

Read through the four lenses, the overrun record of Chapter 1 resolves into a small set of recurring failures, each with a control that a mature practice applies as a matter of course. Exhibit 15 names them. None is exotic. Each is the difference between an estimate that holds and one that does not.

What the four have in common is more telling than what separates them. Each is a place where the estimate is allowed to claim more certainty than the project has earned: a class asserted above its definition, a basis too thin to trace, a contingency trimmed to flatter the return, a number that no independent reviewer has tested. None is a failure of arithmetic, and none is cured by a larger spreadsheet. Each is a failure of discipline, which is why each has a control a mature practice can apply before the estimate leaves the room rather than after the overrun arrives.

EXHIBIT 15

### Four ways estimates fail, and the control for each

The maturity model read as a set of controls.

| | |
|---|---|
| **Class above definition**<br>Control: gate the class on the deliverables matrix | **No traceable basis**<br>Control: require a complete basis of estimate (34R-05) |
| **Contingency trimmed off P50**<br>Control: set contingency by risk method, hold at P50 | **No independent validation**<br>Control: independent review against the class claimed |

Every control is a definition question. The maturity model asks nothing about the number itself, only about the knowledge behind it, which is the one thing that decides whether the number can be trusted.

Aurus capability instrument, synthesising the AACE framework (RP 18R-97, 34R-05, 40R-08). No project data.

# A worked read

Applied to two estimates of the same notional project, the model separates them at once. The first carries a feasibility label, a single headline figure, a contingency stated as a round percentage with no method, and an internal sign-off. Read through the four lenses it is a Class 4 study wearing a Class 3 name: its definition does not reach feasibility maturity, its basis cannot be traced, its contingency is not at P50, and no independent reader has tested it. The second carries the same label but a complete basis of estimate, quantities reconciled to a developing design, a P50 contingency from a quantitative risk analysis, and an independent review on file. It is what it says. The two headline numbers may sit within a few percent of each other; the confidence behind them is a full class apart, and only the maturity read reveals it.

EXHIBIT 16

### The maturity read, applied

Two estimates carrying the same label, scored on the four lenses.

| LENS | ESTIMATE A, LABELLED FEASIBILITY | ESTIMATE B, LABELLED FEASIBILITY |
|---|---|---|
| Definition | Class 4 deliverables held | Class 3 deliverables complete |
| Basis | Headline only, no basis document | Full basis, all eight sections |
| Contingency | Round percentage, no method | P50, quantitative risk analysis |
| Validation | Internal sign-off | Independent technical review |
| **Read** | **Class 4 wearing a Class 3 name** | **What it says it is** |

Aurus capability instrument, applying the four-lens read of Exhibit 14. No project data.

# The cost of getting the class wrong

The maturity read matters because the two ways of getting the class wrong are both expensive, in opposite directions. Claiming a class above the definition, the failure behind the overrun record, destroys value by committing capital against a number the project cannot honour, and it surfaces at the worst possible time, during construction. Refusing to move up the ladder is the quieter error: a sponsor who insists on Class 1 certainty before any decision funds engineering the decision does not need and delays a project whose economics decay with time. The framework guards against both. It tells a screening decision that Class 5 is enough, and it tells a financing decision that nothing short of Class 3 will do. Maturity is not a synonym for precision; it is the match between the confidence of the estimate and the weight of the decision it is asked to carry. An estimate can be too green for its decision, and it can also be needlessly gold-plated for it, and a disciplined owner avoids both.

## What the model is not

The maturity model is a reading instrument, not a scoring gimmick. It does not replace a quantitative risk analysis, a benchmarking exercise or an independent technical review; it is the lens that tells a reader whether those things have been done and whether their results support the class claimed. It produces no number of its own, and it carries no project data. Used honestly it is uncomfortable, because it often reveals that a confident headline rests on Class 4 definition, and that the work needed to reach the claimed class has not yet been done. That discomfort is the point. An instrument that only ever confirmed the label on the cover would be worth nothing.

## For the owner, and for the lender

For a developer or owner, the model is a way to spend study money where it changes the decision: advance definition on the deliverables that gate the next class, rather than polishing a number whose class is already fixed by what is not yet known. The single next action is to take the estimate now in hand, place it against the definition matrix in Exhibit 4, and fund the one or two deliverables that stand between its current class and the next. For a lender or investor, it is a way to read a sponsor's estimate quickly and fairly: locate it on the ladder, confirm that its class matches its basis, and size the diligence to the gap between the two. The single next action there is to call for the basis of estimate and the contingency method before the headline number, and to treat any estimate that cannot produce them as a class below the one printed on its cover. Both readers are asking one question, and it is the question this paper began with. How much of the project does this number actually know.

Aurus brings that question to work across the study ladder and across a portfolio of mandates in mining, infrastructure, engineering and environment, including multiple definitive-feasibility class mandates and studies delivered at scoping grade. The published credentials in this series are selected mandates, a floor rather than a full account. The estimating principle behind all of them does not change with the project. Name the class, document the basis, price the risk, and never let a noun outrun the definition beneath it.

EXHIBIT 17

### The discipline, in one line per class

What each class is entitled to claim, and the decision it can carry.

| CLASS | WHAT IT HONESTLY IS | DECISION IT CAN CARRY |
|---|---|---|
| Class 5 | A screen on a concept, wide by design | Whether to study further |
| Class 4 | An option compared, not yet chosen | Which option to advance |
| Class 3 | A defined project, lender-testable | Sanction and finance |
| Class 2 | A project ready to commit in packages | Commit and contract |
| Class 1 | A project priced to control the build | Control against the estimate |

A summary of AACE RP 18R-97 end usage and RP 47R-11 study stages. Framework exhibit, no project data.

EXHIBIT 18 · SIGNATURE FIGURE

### The confidence funnel: accuracy narrows only as definition grows

Expected accuracy range against level of project definition, the whole framework in one figure. Bands per AACE RP 18R-97, Table 1 and Figure 1; mining study stages per RP 47R-11.

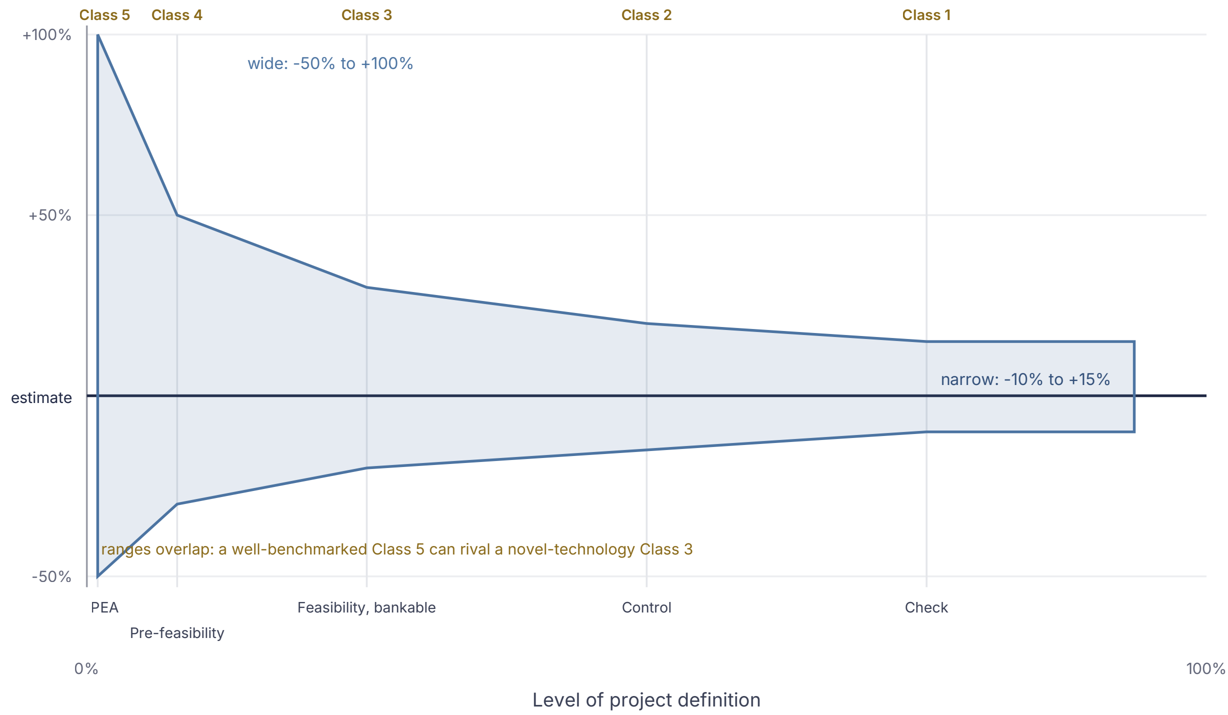


Source: AACE International, RP 18R-97, Table 1 and Figure 1 (accuracy structure and overlap); RP 47R-11 (mining study-stage mapping). The band edges are drawn to the Table 1 low and high values at a common vertical scale (2 px per percent) and set at each class's midpoint level of definition, so the funnel geometry tracks the published figures. Archived: AACE 18R-97 process-industries and generic captures, 2026-07-09. Framework exhibit, no project data.

# Reading the confidence funnel

The signature figure gathers the whole argument into one shape. Read from left to right, the funnel is the promise the framework makes and the limit it imposes at once. At the left edge, where almost nothing is defined, the band is enormous: a concept can be half its estimate or double it, and no estimating skill can close that width, because the width is a property of the ignorance, not of the estimator. The funnel narrows only as definition is added, deliverable by deliverable, and it reaches lender-grade width only at Class 3, where a feasibility study rests on Probable and Proven reserves and a developed design. The dashed line is the reminder that the bands overlap: a repeat project with good cost history can enter its class near the tight edge, while a first-of-a-kind process sits near the wide edge of the same class, which is why the framework publishes ranges and why benchmarking matters so much.

For the reader deciding whether to trust a number, the funnel is a single test. Find where the project sits on the definition axis, read the band the framework allows at that point, and compare it with the confidence the estimate claims. If the two agree, the estimate is mature. If the claim is tighter than the definition permits, the estimate has stepped outside the funnel, and the overrun record of Chapter 1 is what waits outside it. Everything else in this paper, the basis of estimate, the contingency at P50, the stage gates, the independent review, is machinery for keeping an estimate honestly inside the funnel its definition has earned.

The mature estimate is not the one with the smallest number. It is the one whose number knows exactly how much it does not yet know.

# Every figure, traced to its source

This paper follows the Aurus citation policy: every statistic, accuracy range and framework figure carries its institution, publication and year, ranges are preserved as published rather than reduced to points, and no number is presented as Aurus analysis. AACE International Recommended Practices are cited as published technical standards. Exhibits drawn from the published framework are labelled framework exhibits; the Aurus maturity-model instruments (Exhibits 14 to 16) are labelled Aurus capability instruments; both carry no project data. Aurus practice notes are anonymised credentials whose verb tense matches the execution status of the underlying mandate. The AACE and empirical figures were verified against archived source captures held in the Aurus evidence archive for this paper (the RP 18R-97 generic and process-industries documents, a transcription note carrying the 47R-11 mapping and the verified RP identifiers, and the Bertisen and Davis findings), and every exhibit and tile traces to a catalogued source through the paper's citation audit, confirmed orphan-free before publication.

## Standards and published sources


1. AACE International, Recommended Practice 17R-97, Cost Estimate Classification System (generic matrix).
2. AACE International, Recommended Practice 18R-97, Cost Estimate Classification System as Applied in Engineering, Procurement and Construction for the Process Industries (TCM Framework 7.3), Table 1 and Figure 1.
3. AACE International, Recommended Practice 47R-11, Cost Estimate Classification System as Applied in Engineering, Procurement and Construction for the Mining and Mineral Processing Industries.
4. AACE International, Recommended Practice 34R-05, Basis of Estimate.
5. AACE International, Recommended Practice 40R-08, Contingency Estimating: General Principles.
6. AACE International, Recommended Practice 42R-08, Risk Analysis and Contingency Determination Using Parametric Estimating.
7. AACE International, Recommended Practice 44R-08, Risk Analysis and Contingency Determination Using Expected Value.
8. AACE International, Recommended Practice 119R-21, Cost Estimate Accuracy Range and Contingency.
9. AACE International, Recommended Practice 10S-90, Cost Engineering Terminology.
10. AACE International, Professional Guidance Document 02, Guide to Quantitative Risk Analysis.
11. Bertisen, J. and Davis, G. A., Bias and Error in Mine Project Capital Cost Estimation, The Engineering Economist, Vol. 53, No. 2, 2008.
12. CRIRSCO-family reporting codes: JORC (2012) and the CIM Definition Standards incorporated by NI 43-101, for Mineral Resource and Ore Reserve categories and the exploration-target disclosure discipline.


## Exhibit index

| EXHIBIT | SUBJECT | SOURCES |
|---|---|---|
| 1 | Mine capital overrun distribution | Bertisen and Davis, 2008 (ref 11) |
| 2 | Five-class classification matrix | AACE 18R-97, 47R-11 (refs 2, 3) |
| 3 | Class, study stage, resource crosswalk | AACE 47R-11; CRIRSCO codes (refs 3, 12) |
| 4 | Definition deliverables maturity matrix | AACE 18R-97, 47R-11 (refs 2, 3) |
| 5 | Basis-of-estimate anatomy | AACE 34R-05 (ref 4) |
| 6 | Estimate build-up | AACE 18R-97, 34R-05 (refs 2, 4) |
| 7 | Contingency to P50 | AACE 47R-11, 18R-97 (refs 3, 2) |
| 8 | Contingency methods by class | AACE 40R-08, 42R-08, 44R-08, 119R-21, PGD-02 (refs 5 to 8, 10) |
| 9 | Accuracy ranges and effort | AACE 18R-97, Table 1 (ref 2) |
| 10 | Stage-gate ladder | AACE 47R-11 (ref 3) |
| 11 | Lender-grade validation tests | AACE 18R-97, 34R-05, 47R-11 (refs 2, 4, 3) |
| 12 | Compounding resource and cost uncertainty | AACE 47R-11; CRIRSCO codes (refs 3, 12) |
| 13 | Mineral cost drivers beyond the plant | AACE 47R-11; CRIRSCO modifying factors (refs 3, 12) |
| 14 | Four-lens maturity read | Aurus, synthesising AACE 18R-97, 34R-05, 40R-08, 47R-11 |
| 15 | Failure modes and controls | Aurus, synthesising AACE 18R-97, 34R-05, 40R-08 |
| 16 | The maturity read, applied | Aurus capability instrument (Exhibit 14 method) |
| 17 | The discipline, one line per class | AACE 18R-97, 47R-11 (refs 2, 3) |
| 18 | The confidence funnel (signature figure) | AACE 18R-97, Table 1 and Fig. 1; 47R-11 (refs 2, 3) |

GLOSSARY

# Terms used in this paper

**AACE International** The professional association for cost engineering that publishes the estimate-classification Recommended Practices used throughout this paper. Referenced in Chapters 1 to 6.

**Estimate class** One of five levels (Class 5 to Class 1) set by the maturity of project definition, each with an expected accuracy range. Chapters 1, 2.

**Level of project definition** The share of full engineering and project deliverables completed, expressed as a percentage; the sole determinant of estimate class. Chapters 2, 3.

**Basis of estimate** The document recording an estimate's scope, assumptions, exclusions, quantity and pricing sources, escalation and contingency method. Chapter 3.

**Contingency** Funds added to a base estimate to cover cost growth from within-scope uncertainty, distinct from allowances and management reserve. Chapter 4.

**P50** The point at which cost overrun and underrun are equally likely; the target for a correctly set contingency. Chapter 4.

**Accuracy range** The typical variation of actual cost from the estimate after contingency, quoted at an 80 percent confidence interval. Chapters 2, 4.

**PEA** Preliminary Economic Assessment; a scoping-level mineral study whose cost engine is typically a Class 5 estimate. Chapter 2.

**PFS** Pre-feasibility study; associated with a Class 4 estimate on Indicated Resources. Chapter 2.

**DFS, BFS** Definitive and bankable feasibility study; require Class 3 estimating on Probable and Proven reserves. Chapters 2, 5.

**Mineral Resource, Ore Reserve** CRIRSCO-family categories (Inferred, Indicated, Measured; Probable, Proven) reported under JORC 2012 and the CIM Definition Standards incorporated by NI 43-101. Chapters 2, 3.

**Exploration target** A conceptual tonnage and grade range that is not a Mineral Resource, disclosed with a cautionary statement and never carrying project economics. Chapter 2.

**Stage gate** A definition threshold that a project must pass before it may claim a class and take the decision that class supports. Chapter 5.

**FID** Final Investment Decision; the sanction to build, supported by Class 3 and control-grade estimating. Chapter 5.

**Allowance** Money carried in the base estimate for scope that is real but not yet detailed, distinct from contingency and from management reserve. Chapters 3, 4.

**Management reserve** Funds held by the owner outside the estimate for changes in scope that have not yet occurred, drawn under the owner's authority. Chapters 3, 4.

**Escalation** The expected movement of prices between the pricing date and the point of expenditure, computed in the base estimate rather than carried as contingency. Chapters 3, 4.

**Quantitative risk analysis** A probabilistic model of the full estimate that produces the cumulative cost distribution from which contingency and the accuracy range are read. Chapter 4.

**Sensitivity analysis** A test of how the outcome moves when one input is flexed while the rest are held still; a communication tool, not a substitute for a risk analysis. Chapter 4.

**Benchmarking** Checking an estimate against what comparable delivered projects actually cost, the outside view that an internal build-up cannot supply. Chapter 5.

**Modifying factors** The mining, processing, infrastructure, economic, marketing, legal, environmental, social and governmental considerations that convert a Mineral Resource into an Ore Reserve, almost all of which are also cost inputs. Chapters 2, 6.

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

The Discipline of Confidence: AACE Cost-Estimate Maturity in Mining and Mineral Infrastructure.

Aurus White Paper Series, No. 02 · Build. Technical Paper. By the Aurus Technical Committee, July 2026. Working draft 1: adversarial certification pending.

Citation policy: every statistic carries its institution, publication and year; AACE Recommended Practices are cited as published technical standards; ranges are preserved as published; framework exhibits and the Aurus maturity-model instruments both carry no project data; anonymised practice notes match the execution status of the underlying mandate.

Typefaces: Fraunces (display), Source Serif 4 (text), Inter (labels and exhibits).

Imagery generated and art-directed in-house.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT