

Fundamental Analysis Course

- Module 1: The Lasso Curve & The 10 Stages of a Mining Company
- Module 2: JORC Code & The Resource / Reserve Hierarchy
- Module 3: Ore Grades, Cut-Offs & What "Good" Actually Looks Like
- Module 4: Reading Drill Result Announcements
- Module 5: Economic Studies (Scoping, PFS, DFS)
- Module 6: Capital Structure, Dilution & The Share Register
- Module 7: Catalysts & The Catalyst Calendar
- Module 8: Red Flags, Green Flags & Sneaky Tactics
- Module 9: Macro Overlay — Commodity Cycles & The ASX Small-Cap Resource Cycle
- Module 10: Putting It Together — The "FA Story" and TA Confluence

Module 1: The Lassonde Curve & The 10 Stages of a Mining Company

Why this is the most important framework

The **Lassonde Curve** — named after Pierre Lassonde, founder of Franco-Nevada — plots a typical mining company's market cap over its lifecycle, from grassroots exploration to production decline. It looks like two peaks separated by a deep valley.

- **First peak:** discovery / maiden resource excitement
- **Valley:** the "orphan period" or "valley of death" during studies, permitting, and financing
- **Second peak:** production ramp-up and cash flow realisation
- **Decline:** as reserves deplete

Most retail investors buy the first peak (discovery hype), get smashed in the valley, and capitulate at the bottom. Sophisticated capital buys the valley and rides the second peak. The framework's value is letting you recognise which peak/valley you are standing in *right now*.

Every other module in this course feeds back into this curve.

The 10 stages

Stage 1 — Concept / Grassroots

- **Typical MC:** under \$5m, often \$2-3m shells
- **What's happening:** tenement pegging, geophysics, soil/stream sediment sampling, mapping. No drilling yet.

- **Catalysts:** essentially none. Tenement grants, MOUs, "JV signed" announcements.
- **SP behaviour:** drifts; pumped briefly on cheap PR.
- **Reality:** you are buying a story, a postcode, and management. 95%+ never advance past this stage.

Stage 2 — Drilling Begins

- **Typical MC:** \$5-20m
- **What's happening:** maiden drill program. Usually RC (reverse circulation) for cost reasons before diamond core.
- **Catalysts:** "rig mobilised", "drilling commenced", "first hole completed", "assays pending".
- **SP behaviour:** the **drill speculation ramp** — SP often runs *before* assays return as punters position for a hit.
- **Reality:** ~90% of maiden programs return nothing economically meaningful. The pre-assay ramp is often the entire trade.
- **ASX example: BSR (Bison Resources)** post-IPO in late 2025 — early drilling on its WA gold tenements with the SP running into anticipated first assays. **TOR (Torque Metals)** is similar — early-stage drilling with drill-for-equity contractor arrangements (a green flag, see Module 6/8).

Stage 3 — Discovery

- **Typical MC:** \$20-100m+, sometimes much more in hot commodity cycles
- **What's happening:** high-grade hits in multiple holes, geological continuity emerging, step-out drilling extending the mineralised footprint.
- **Catalysts:** every assay batch is market-moving. Comparisons to nearby tier-1 deposits.
- **SP behaviour: THE FIRST PEAK.** 5x to 50x moves possible in weeks. This is the dopamine spike.
- **Trap:** many discoveries don't substantiate at scale. Step-out drilling reveals grade falls off, mineralisation pinches, or it's a one-pod wonder.
- **ASX example: PC2 (PC Gold)** — Stage 3/4 NT gold discovery with high-grade hits driving repeated re-rates through 2025 into early 2026. The hallmarks were there: multiple consistent intercepts, geological continuity, step-outs extending the footprint, and a tight cap structure (Module 6) that translated grade hits into market cap moves.

Stage 4 — Resource Definition

- **Typical MC:** \$50-200m
- **What's happening:** closely-spaced drilling to define a JORC-compliant resource. Metallurgical testwork begins.

- **Catalysts: maiden JORC resource** is a major re-rate event. Subsequent upgrades (Inferred → Indicated → Measured) are progressively less impactful.
- **SP behaviour:** can hold or drift higher into maiden resource, then often sells the news.
- **ASX example: MI6 (Minerals 260)** — defining resources at its WA gold and base-metal projects. **PC2** transitions through this stage as it works toward a maiden JORC resource (planned for mid-2026).

Stage 5 — Scoping Study

- **Typical MC:** flat or starting to soften
- **What's happening:** conceptual economic study, $\pm 35\text{--}50\%$ accuracy. Tests whether the project is *worth* progressing.
- **Catalysts:** scoping study results.
- **SP behaviour:** muted. Market starts thinking about capex, dilution, permitting timelines. **The valley begins.**

Stage 6 — PFS (Pre-Feasibility Study)

- **Typical MC:** often **down 30–60%** from discovery peak
- **What's happening:** $\pm 20\text{--}25\%$ accuracy. Mining method selected, processing flowsheet designed, infrastructure costed.
- **Catalysts:** PFS results. Capex sticker shock is real and frequent.
- **SP behaviour:** deep valley. This is where capital raises hurt most because the SP is depressed.

Stage 7 — DFS (Definitive / Bankable Feasibility Study)

- **Typical MC:** starting to recover *if* economics are robust
- **What's happening:** $\pm 10\text{--}15\%$ accuracy. Lender-ready engineering. Reserves declared from Indicated/Measured resources after applying Modifying Factors (Module 2).
- **Catalysts:** DFS results, **Final Investment Decision (FID)**, financing close, offtake agreements.
- **SP behaviour:** bottom of the valley typically here or just before. **This is where smart money accumulates** — the project is de-risked but the market hasn't re-rated yet.
- **ASX example: SLS (Solstice Minerals)** sits in the studies/DFS region for parts of its portfolio — the kind of stock that grinds in the valley while the work gets done. The opportunity is recognising stocks where the work *is* progressing (catalysts being delivered, no red flags from Module 8) but the SP hasn't responded yet.

Stage 8 — Financing & Construction

- **Typical MC:** re-rating
- **What's happening:** debt + equity + offtake/streaming/royalty package finalised. Construction begins. EPC contractor mobilised.
- **Catalysts:** financing close, ground-breaking, construction milestones, first ore on ROM pad.
- **SP behaviour:** grinds higher with milestones. Setbacks are punished hard.
- **Watch out: 30-50% capex blowouts are common, not exceptional.** Schedule slippage is the norm.
- **ASX example: GLN (Galan Lithium)** — Phase 1 construction completed 31 March 2026 on time and on budget (the exception, not the rule). Note that GLN's recent re-rate from ~\$0.09 in mid-2025 to ~\$0.47 in January 2026 was the Lassonde **second peak build** — sophisticated capital (Clean Elements) bought the valley at \$0.11 and the same fund paid 3.7x more per share five months later (covered in Module 6).

Stage 9 — Commissioning & Ramp-Up

- **Typical MC:** significant re-rate
- **What's happening:** plant commissioning, first concentrate / doré / cathode, ramp toward nameplate capacity.
- **Catalysts:** first product, first revenue, first cash flow, achievement of nameplate.
- **SP behaviour: THE SECOND PEAK** builds here. But ramp problems are normal — recovery shortfalls, throughput issues, grade reconciliation problems — and can crater the SP temporarily.
- **ASX example: GLN** crosses from Stage 8 into Stage 9 in Q2 2026 as it processes first brine and produces first lithium chloride concentrate. This transition is the most volatile point of the curve — get commissioning right and the second peak builds quickly; commissioning problems can crater the SP and force dilutive raises at compressed prices.

Stage 10 — Production & Depletion

- **Typical MC:** tied to commodity price × production × remaining reserves
- **What's happening:** steady-state production, reserve replacement battle, expansion studies, M&A.
- **Catalysts: quarterlies** (production, AISC, cash position, hedge book), reserve updates, exploration around mine, M&A.
- **SP behaviour:** tracks commodity price + operational delivery. Depletion without replacement = terminal decline.
- **ASX example: ELV (Elevra Lithium)** — Quebec hard-rock producer. Through 2024–early 2025, ELV traded in the typical Stage 10 producer pattern: SP tracking spodumene price, operational delivery measured quarterly, reserves and grade updates

as recurring catalysts. This is where commodity-cycle plays compound — a producer can 5-10x on the underlying commodity moving without doing anything new operationally.

A reference map of typical Lasso

nde positions

This is a snapshot — stocks move along the curve over time, sometimes forward, sometimes backward, occasionally sideways for years. Use it as a way of orienting any new ticker quickly:

Stage	Typical MC	Risk profile	Position sizing default	ASX examples (April 2026)
1 — Concept	<\$5m	Lottery	0.5-2%	Most \$2-3m shells
2 — Drilling	\$5-20m	Lottery	0.5-2%	BSR, TOR (early drilling)
3 — Discovery	\$20-100m+	High volatility, asymmetric	2-4%	PC2
4 — Resource def	\$50-200m	High volatility	2-4%	MI6, PC2 (transitioning)
5 — Scoping	flat / softening	Orphan period	2-5%	varies
6 — PFS	-30 to -60% from peak	Deep valley	2-5%	varies
7 — DFS	recovering	Smart money zone	2-5%	SLS (parts of portfolio)
8 — Construction	re-rating	Execution risk	3-6%	GLN (entering Stage 9)
9 — Commissioning	significant re-rate	Ramp risk	3-6%	GLN (Q2 2026 onwards)
10 — Production	commodity × tonnes	Macro-driven	5-10%	ELV

(Position sizes are defaults — Module 10 covers sizing in detail. They depend on conviction, total mining allocation, and your risk tolerance.)

Why the valley exists (and why it's the opportunity)

1. **No catalysts.** Discovery dopamine has worn off. Months of nothing while consultants run studies.
2. **Capital raises.** PFS/DFS work costs millions. Issued at deep discounts because SP is weak. Existing holders get diluted.
3. **Capex sticker shock.** \$200m-\$2bn+ capex numbers scare retail.
4. **Time.** 3-7 years from discovery to first production is typical. Retail attention does not last that long.
5. **Permitting.** Indigenous Land Use Agreements, environmental approvals, water licences, native title — glacial pace, especially in WA and NT.

The valley is where 10m → 100m+ MC stories actually get *built*. Not at the discovery pop (gambling), but at the bottom of the valley when a real project with real economics is being de-risked and the market hasn't woken up yet.

The GLN case study is worth internalising here. In mid-2025, lithium was deeply out of favour, GLN's SP had collapsed from \$1.70+ in 2022 to ~\$0.09. Most retail had capitulated. That August, a specialist lithium fund (Clean Elements) did 77 days of due diligence and put \$20m in at \$0.11 — a 21% **premium** to market. Five months later, after Phase 1 construction had visibly progressed, the same fund paid \$0.41 in another placement at premium-to-VWAP pricing. The valley was where the asymmetric trade existed; by the time mainstream coverage and broker upgrades arrived, the easy money had been made.

Where 10m → 100m MCs come from (honestly)

Three legitimate paths:

1. **Discovery (Stages 2 → 3).** Highest-conviction trade if you're early. Lowest probability per attempt. Lottery economics — most go nowhere; the rare ones go 20-50x.
2. **Re-rating from valley (Stages 7 → 9).** Lower upside per stage but much higher probability if you've done the work. The "boring" path that compounds.
3. **Commodity cycle.** A producer (Stage 10) can 5-10x just because the underlying commodity moved. Macro overlay — covered in Module 9.

A fourth path that exists but is harder to time: **takeover**. Mid-tier or major acquires the developer at a 20-80% premium. Often happens at Stage 6-8. (GLN rejected a US\$150m takeover bid from Huayou + Renault in late 2024 — sometimes management sees more value standalone, sometimes they're wrong.)

Using AI tools to research mining stocks — a 2-stage workflow

A practical note that's worth flagging early because most retail are now using AI tools (ChatGPT, Claude, Gemini, Grok, etc.) to research stocks and getting burned by it.

What AI is good at

- Explaining frameworks and concepts (the Lassonde Curve, JORC categories, capex intensities, what AISC means)
- Summarising publicly disclosed information *if* it has been fed the actual document
- Comparing similar concepts side-by-side (hard rock vs brine, Inferred vs Indicated, scoping vs PFS)
- Generating questions you should be asking about a stock
- Pattern recognition once you describe what you're seeing

What AI is bad at and routinely gets wrong

- **Specific drill intercepts and assay results** — these get hallucinated frequently
- **Current cap structure** — SOI counts, recent placements, options outstanding are often stale or invented
- **Recent share prices and market caps** — without a search tool, models work from training data that's months to years out of date
- **Identity of management, directors, major holders** — gets people confused, mixes up similarly-named companies
- **Specific tenement boundaries, geology, drill hole IDs** — these are often invented if the model is asked confidently
- **Recent ASX announcements** — not in training data; fabricated when extrapolated from older info

The 2-stage workflow

Stage 1 — Frame the question with AI. Use the AI to:

- Explain the framework relevant to the stock (Lassonde stage, study type, deposit type, jurisdiction)
- Generate a list of specific questions you need answered
- Identify what good and bad would look like for that specific situation
- Compare it conceptually to similar companies you already understand

Stage 2 — Verify every specific fact against primary sources. That means:

- The company's ASX announcements page directly (not summarised by AI)
- The most recent quarterly Appendix 5B (Section 8 cash burn, Section 1.2 cash position)
- The most recent annual report and any maiden resource / scoping / PFS / DFS document
- Director announcements (Appendix 3Y) for on-market buying / selling
- The Top 20 holders list at the back of the annual report

The AI-research failure mode is doing Stage 1 and skipping Stage 2 — letting the AI confidently tell you that a company has X tonnes at Y g/t when it has neither. The number sounds plausible. It isn't. **Every specific number, name, or date that matters to your decision must be verified from the company's own ASX disclosures.** No exceptions.

A simple rule: if you're about to put real money behind a fact, that fact has to be linked back to a specific ASX document you've actually read. If the AI gave it to you and you can't find it on the ASX announcements page, treat it as not real until proven otherwise.

Practical exercise

For every ASX mining stock you're tracking, write down:

1. What stage is it at?
2. What was its peak MC and when?
3. Where on the curve is it sitting *now* relative to its peak?
4. What's the next genuine catalyst that could move it?
5. What ASX document did you verify each of those answers from?

If you can't answer all five, you don't actually know the company well enough to size a position.

What I'm uncertain about

- MC ranges by stage vary massively by commodity, jurisdiction, and macro cycle. The numbers above are typical Australian junior patterns over the last decade — treat as rough guide, not law.
- The Lasso Curve is a heuristic. Plenty of companies skip stages, fail at any stage, or backslide. A discovery can de-rate to below pre-discovery levels if step-out drilling disappoints.
- Timeframes have been getting *longer* in WA due to permitting bottlenecks. 3–7 years discovery-to-production is increasingly optimistic for new projects.
- The ASX examples cited above are illustrative of stage positions as at April 2026 — they will move along the curve over time. Re-classify them yourself when you re-read this in 6 months.

Module 2: JORC Code & The Resource / Reserve Hierarchy

Why this matters


The JORC Code is **the rulebook** for how Australian-listed mining companies report exploration results, mineral resources, and ore reserves. It is the single biggest source of retail confusion and the single biggest place companies legally mislead unsophisticated investors.

If you can't read a JORC table and explain what each category means, you are gambling — not investing — in mining.

The basics

- **JORC** = Joint Ore Reserves Committee
 - **Current code:** JORC Code 2012 (released December 2012). A revised version has been in consultation for several years; check the JORC website for current status.
 - **Mandatory** for all ASX-listed companies reporting exploration results, resources, or reserves.
 - Aligned to the international **CRIRSCO** template, so JORC, NI 43-101 (Canada), SAMREC (South Africa), PERC (Europe) and SK-1300 (US SEC) are broadly equivalent though not identical. (More on the differences later in this module.)
-

The hierarchy (most important diagram in mining)



Two key rules baked into this diagram:

1. **Inferred resources cannot be converted directly to reserves.** They must first be upgraded to Indicated or Measured via additional drilling. This is a hard rule and one of the most common things companies blur in their wording.
2. **Reserves require Modifying Factors to be applied.** A resource is a geological estimate; a reserve is an *economic* statement that the rock can be mined and processed profitably under defined assumptions.

Resource categories

Inferred Resource

- Lowest geological confidence.
- Drill spacing wide (often hundreds of metres).
- Tonnage and grade are estimated based on limited sampling.
- **Cannot be used in reserve calculations.**
- **Cannot be used in PFS or DFS economics** (with very narrow exceptions).
- *Can* be used in Scoping Studies — which is one reason scoping studies are not bankable.

Indicated Resource

- Moderate confidence.
- Drill spacing tight enough that geological and grade continuity can reasonably be assumed.
- **Can be converted to Probable Reserves** after Modifying Factors.
- Required basis for PFS economics.

Measured Resource

- Highest geological confidence.
- Very tight drill spacing, often supported by trenching, channel sampling, or underground development.
- **Can be converted to Proved Reserves** (or Probable, depending on confidence in Modifying Factors).

Sub-categories you'll see

Modifiers like "in-situ" vs "extractable", "open-pittable" vs "underground", or constraints by depth or pit shell. Read carefully — companies sometimes report a big global resource then quietly disclose only a fraction is within an economic pit shell.

Reserve categories

Probable Reserve

- Derived from Indicated (or sometimes Measured) Resources after applying Modifying Factors.
- Lower confidence than Proved.

Proved Reserve

- Derived from Measured Resources after applying Modifying Factors.
 - Highest confidence economic mining inventory.
-

The Modifying Factors

To convert a resource to a reserve, the Competent Person must apply realistic assumptions about:

- **Mining** (method, dilution, recovery, geotech)
- **Processing** (metallurgical recovery, throughput)
- **Metallurgical** (test work supporting flowsheet)
- **Infrastructure** (power, water, roads, port access)

- **Economic** (commodity price, opex, capex, royalties, taxes)
- **Marketing** (offtake potential, product specifications)
- **Legal** (tenements, mining leases)
- **Environmental** (approvals, rehabilitation)
- **Social** (community agreements, ILUAs)
- **Governmental** (royalty regimes, foreign investment rules)

If any of these don't stack up, the resource cannot become a reserve. Many "huge resources" never become reserves because of metallurgy, infrastructure, or social licence.

Competent Person (CP) requirement

Every public report of exploration results, resources, or reserves must:

1. Be signed off by a **Competent Person** — a member of AusIMM, AIG, or a recognised overseas equivalent (RPEQ, etc.).
2. Have at least 5 years' relevant experience in the style of mineralisation/deposit type and the activity being reported.
3. Include a **named consent** statement that the CP has reviewed and approved the form and context of the announcement.

If the CP statement is missing, vague about credentials, or the CP is the company's CEO with no relevant background — that's a flag.

Reading a JORC resource table

A standard table looks like:

Category	Tonnes (Mt)	Grade (g/t Au)	Contained Au (Moz)
Measured	5.0	2.5	0.40
Indicated	12.0	1.8	0.69
Inferred	25.0	1.4	1.13
Total	42.0	1.65	2.22

What to actually look at:

1. **% of resource that is Inferred.** If it's 60%+, the headline ounces are largely speculative. Many junior "1Moz+" announcements are 70-90% Inferred.
 2. **Cut-off grade applied.** Always disclosed in the announcement footnotes. A resource at 0.3 g/t cut-off vs 0.8 g/t cut-off describes a fundamentally different deposit.
 3. **Top-cut applied.** Were extreme high-grade assays capped (e.g., capped at 30 g/t)? If not, the average grade may be skewed by a few bonanza intercepts.
 4. **Density assumption.** Tonnage = volume × density. Wrong density = wrong tonnage.
 5. **Constraints.** Was the resource constrained inside a pit shell at a given commodity price? At what price?
-

Common gotchas (where companies legally mislead)

"Exploration Target"

This is **NOT** a JORC resource. It's a conceptual range (e.g., "10–25 Mt at 1.5–2.5 g/t Au"). The JORC code allows it but requires explicit wording that it is conceptual, not a resource estimate, and based on insufficient drilling. Companies often headline this and bury the disclaimer.

Historical Estimates

Often quoted by companies that have just acquired old projects. **Not JORC-compliant** unless re-validated by a Competent Person under the 2012 Code. Headlines like "historical resource of 500,000 oz" are essentially marketing until JORC-validated. Covered in detail in the foreign-and-historical-estimates section below.

"Ore"

Strictly, "ore" means rock that is *economically* extractable — i.e., reserve-quality. Companies use "ore" loosely for inferred resources, which is technically non-compliant. Watch the language.

Resource on resource on resource

Multiple updates in a year, each headlined as if it's a fresh discovery, when it's the same deposit being re-drilled and re-categorised. Look at how the *Indicated + Measured tonnes* are growing, not just the headline total.

"No new drilling, just changed parameters" — the parameter-shuffle MRE

This is one of the more important patterns to recognise because it looks like progress and isn't.

A Mineral Resource Estimate (MRE) update can grow the headline resource without any new drilling at all, simply by changing the input parameters:

- **Lowering the cut-off grade** — sweeps in more low-grade material that was previously below cut-off
- **Raising the assumed commodity price** — extends the economic pit shell, capturing more material at depth or at the edges
- **Loosening the geological continuity assumption** — links blocks of mineralisation that were previously separate
- **Changing the top-cut** — letting bonanza intercepts have more influence on the average grade
- **Re-classifying based on revised drill spacing rules** — moving material from Inferred to Indicated without new drilling, on the basis of the geologist's revised judgment

None of this is necessarily wrong. Cut-offs should reflect current commodity prices. Geological models do get refined. But when an MRE update grows materially **without new drilling**, the right read is "the same rock has been re-described" — not "we've discovered more ore."

How to spot it

When an MRE update is announced:

1. **Read the methodology section.** It will state how many new holes were drilled since the last update. If the answer is "none" or "very few", the change is parametric.
2. **Compare the cut-off used.** A drop from 0.5 g/t to 0.3 g/t can grow tonnage 30–50% on the same drilling.
3. **Compare the commodity price assumption.** If the price assumption moved from \$1500/oz to \$2500/oz, the pit shell and economics expand without any geological change.
4. **Look at the Indicated + Measured tonnage in particular.** Adding Inferred tonnes is easier than adding I+M tonnes. The latter requires real new drill data.
5. **Read the Competent Person statement.** It will (or should) note any methodology changes.

Example of what this looks like in announcements

A company reports: "Updated MRE delivers 35% increase in contained gold, project now hosts 2.7 Moz Au." Reading the body:

- Same number of drill holes as the prior MRE
- Cut-off lowered from 0.5 g/t to 0.3 g/t
- Gold price assumption lifted from US\$1800/oz to US\$2400/oz
- Top-cut raised from 25 g/t to 50 g/t
- Pit shell expanded as a result of the higher price

The deposit is exactly the same. The headline 35% growth is parametric. The Indicated + Measured fraction may even be lower than before (because the new Inferred tonnes added at lower cut-off are likely Inferred-only).

A genuine resource upgrade would include a meaningful new drilling campaign, with hole counts and metres drilled since the prior estimate disclosed in the methodology section.

When parameter changes are legitimate

They aren't always misleading. If gold prices have genuinely doubled, lowering the cut-off and re-running the pit shell is the correct technical response — the prior MRE was conservative for current conditions. The flag isn't "they changed parameters." The flag is "they changed parameters and headlined it as a resource growth event without making clear what drove the change."

The honest version of the same announcement would be: "MRE re-stated under updated price and cut-off assumptions; headline tonnage increased 35% with no change to the underlying geology or drilling."

The promotional version is: "Resource grows 35%."

Same numbers, different read. The body of the announcement is where the truth lives.

Foreign and historical estimates — what they are and how to read

them

A common situation: an ASX company acquires a project that already has a resource estimate calculated under a non-JORC code (Canada's NI 43-101, the US SEC's SK-1300, South Africa's SAMREC, etc.) or under no modern code at all (a "historical estimate" from the 1980s, 1990s, or earlier).

These cannot be reported as a JORC resource, but ASX **does** allow them to be disclosed under specific rules.

The CRIRSCO family

The Committee for Mineral Reserves International Reporting Standards (CRIRSCO) sets a template that JORC, NI 43-101, SK-1300, SAMREC, PERC, and others all derive from. They share the same hierarchy (Inferred → Indicated → Measured for resources, Probable → Proved for reserves) and the same conceptual framework around Modifying Factors and Competent Persons (or "Qualified Persons" in the North American codes).

The codes are conceptually equivalent but **not identical**. Differences exist in:

- Specific drill spacing requirements for each classification
- The level of detail required in technical reports
- Filing and disclosure mechanisms
- The application of classification judgment in marginal cases

NI 43-101 (Canada)

National Instrument 43-101 is Canada's mandatory mineral disclosure standard, in force since 2001. Reports are signed off by a "Qualified Person" (the Canadian equivalent of a Competent Person), filed publicly on SEDAR+, and tend to be more detailed than the equivalent ASX disclosures.

Key practical points:

- NI 43-101 technical reports are typically more comprehensive — they include detailed sections on geology, drill data, QA/QC, metallurgical testwork, and economic analysis
- They're publicly downloadable from SEDAR+ ([sedarplus.ca](https://www.sedarplus.ca)) at no charge
- The Indicated → Measured threshold under NI 43-101 is broadly aligned with JORC, though application can differ
- A NI 43-101 resource is the closest international cousin to a JORC resource

SK-1300 (US SEC)

The US SEC modernised its mineral disclosure rules in 2018 (effective 2021), replacing the old "Industry Guide 7" — which was famously restrictive (it didn't allow companies to disclose Resources at all, only Reserves). SK-1300 brought the US into rough alignment with JORC and NI 43-101.

- Uses "Qualified Person" terminology
- Allows Resource and Reserve reporting in the same hierarchy
- Filed via the SEC's EDGAR system
- Slightly more conservative in some areas than JORC/NI 43-101

SAMREC (South Africa) and PERC (Europe)

Functionally similar to JORC for African and European projects respectively. PERC tends to be more common for European projects (especially Sweden, Finland, Spain, Portugal). SAMREC is the standard for South African and many sub-Saharan African operations.

ASX Listing Rule 5.12 — how foreign estimates get reported

ASX Listing Rule 5.12 allows a company to disclose a foreign estimate (an estimate calculated under NI 43-101, SK-1300, SAMREC, or another recognised foreign code) but with strict requirements:

1. **The estimate is not classified as a JORC resource** until a Competent Person reviews and signs off
2. The company must clearly state that it is **not reported in accordance with JORC Code 2012**
3. The company must disclose **whether a Competent Person has done sufficient work to consider the estimate JORC-compliant** — and if not, what work is required to bring it up to JORC standard
4. The company must state that the estimate **must not be relied upon** in the same way as a JORC-compliant resource

This is why ASX announcements about foreign-acquired projects often include a long disclaimer paragraph alongside the foreign estimate numbers.

Why this matters legitimately

The use case is straightforward: a company acquires a project that already has, say, an NI 43-101 resource from a prior Canadian or TSX listing. On day one of the acquisition, the company can disclose that the project hosts X tonnes at Y g/t under NI 43-101. They then commission a Competent Person to do the work — usually 6 to 18 months — to convert the estimate to JORC. This work involves reviewing the underlying drill data, validating QA/QC, sometimes drilling twin holes, and producing a fresh CP report.

Letting the market price in the underlying value during this transition is reasonable. Forcing the company to keep a known resource secret for 18 months would be the wrong outcome.

Where to be careful

Pre-2014 NI 43-101 reports used slightly different rules. Always check the date of the foreign estimate report. Reports from the early 2000s may have been calculated under standards that are now considered loose by current practice.

Historical estimates are *different* from foreign estimates. A historical estimate is a pre-modern-code calculation — a 1980s calculation, a 1990s calculation, a Soviet GKZ resource — that doesn't conform to *any* current code. ASX allows these to be disclosed under Listing Rule 5.12 too, but with even stricter disclaimers. The qualification typically reads something like: "the historical estimate is not reported in accordance with the JORC Code 2012 and a Competent Person has not done sufficient work to classify the historical estimate as a Mineral Resource in accordance with the JORC Code 2012."

Treat historical estimates with significantly more scepticism than NI 43-101 or SK-1300 estimates. The methodology, sampling, and assays underpinning them often don't meet modern standards. Re-validation frequently results in significant downgrades.

What to expect when foreign estimates get converted to JORC

A few rules of thumb:

- The JORC version is usually the **same size or smaller** than the foreign estimate, rarely larger
- Material classified as **Indicated under NI 43-101 may be reclassified as Inferred under JORC** because of differing application of confidence judgment
- Old historical estimates frequently lose 20–40% of their headline tonnage on JORC re-validation, sometimes more
- **Grades may be revised downward** if top-cuts get applied where they weren't before, or if old assays from inadequate QA/QC procedures get excluded

- The re-stated JORC number is the one you should use for valuation; the foreign estimate is for orientation

Do not extrapolate from headline foreign estimates. Wait for the JORC re-statement, then re-anchor your valuation.

The perpetual-foreign-estimate pattern (covered in Module 8)

Some companies disclose a foreign estimate at acquisition and never seem to deliver the JORC version. 12 months pass, then 18, then 24, with quarterly reports continuing to mention "the JORC re-statement is in progress." This is a behavioral red flag covered in detail in Module 8 — at minimum it raises the question of why the JORC work is taking so long, and at worst it suggests the company knows the JORC version will be substantially smaller than the foreign estimate they keep referring to.

Why "10Moz" doesn't mean what you think

A 10Moz Au resource sounds enormous. But:

- If 70% is Inferred, you have 3Moz of bankable rock until further drilling.
- If grade is 0.6 g/t, the project is marginal at best (gold).
- If cut-off was set at 0.3 g/t, lifting the cut-off to 0.8 g/t might leave you with 2Moz.
- If the deposit is at 800m depth and the company is a \$20m junior, no one is mining it for 15 years.
- If it's in a jurisdiction with permitting/sovereign risk, it might never be mined.
- If it's reported as a NI 43-101 estimate that hasn't been JORC-validated, the JORC version may be smaller.
- If the headline grew through parameter shuffle rather than new drilling, the deposit hasn't changed.

Resource size **without context** is meaningless. This is the most common retail mistake in mining.

Practical exercise

Pull up the latest resource update from any ASX gold or copper junior you follow. Answer:

1. What % of the contained metal is Inferred vs Indicated vs Measured?
2. What cut-off grade was used? What was used in the prior estimate?
3. Was a top-cut applied? At what value? Same as prior estimate?
4. Is the resource constrained inside a pit shell? At what commodity price?
5. Who is the Competent Person and are they independent of the company?
6. **How many new drill holes / metres have been drilled since the prior estimate? If "none", the headline change is parametric.**
7. **Is any part of the disclosed resource a foreign estimate or historical estimate, not yet JORC-validated?**

If you can answer these for every position you hold, you're ahead of 95% of retail.

What I'm uncertain about

- Whether the JORC 2012 Code revision (long in consultation) has been finalised at the time you're reading this. Worth checking the JORC website (jorc.org) for current status.
- Specific recent ASX enforcement actions — there have been periodic crackdowns on misleading resource statements but I'd want to verify recent specifics rather than name them from memory.
- Specific application differences between JORC and NI 43-101 in marginal cases. The high-level framework above is solid; for specific edge cases, refer to the actual code documents.
- The 20-40% downgrade rule of thumb for historical estimate re-validation is anecdotal and varies enormously by deposit. Some are fine; some collapse to nothing.

Module 3: Ore Grades, Cut-Offs & What "Good" Actually Looks Like

Why this matters

A resource is just tonnage $\times$ grade. The grade tells you whether the project is economic. Without grade context, you can't tell a tier-1 deposit from a moose pasture.

This module gives you benchmarks by commodity so when you see a drill hit or resource update, you immediately know whether it's interesting or marketing.

Key concepts first

Cut-off grade

The minimum grade above which rock is classified as economic to mine. Below cut-off = waste.

- Lower cut-off = more tonnes, lower average grade
- Higher cut-off = fewer tonnes, higher average grade
- Companies can manipulate the headline by adjusting cut-off. Always check the footnote.

Strip ratio

Open-pit mining only. Tonnes of waste that must be moved per tonne of ore.

- 1:1 = excellent (ore right at surface)
- 3:1 to 5:1 = typical for most metals
- 8:1+ = high cost; only justified for high-grade or high-value ore

- 15:1+ = usually marginal at best

Recovery rate

% of contained metal that can actually be extracted through processing.

- 90-95% = excellent (typical for clean sulphides)
- 70-85% = average
- <60% = problematic ("refractory" ore — needs expensive processing like roasting, pressure oxidation, or fine grinding)

Head grade vs in-situ grade

- **In-situ grade** = grade of rock in the ground (the resource grade)
- **Head grade** = grade of rock fed to the processing plant
- They're often different because of dilution from waste rock during mining and selective mining (taking only the best stuff first in early years).

Mining dilution

Waste rock that gets mixed with ore during extraction. Typically 5-15% in open pit, sometimes more in narrow-vein underground. Reduces head grade vs in-situ.

Grade benchmarks by commodity

Gold (g/t Au)

Context	Grade
Bulk-tonnage open pit	0.5-1.5 g/t
Average open pit	1-3 g/t
High-grade open pit	3-5 g/t
Underground average	4-8 g/t
High-grade underground	8-15 g/t
Bonanza	30+ g/t

Context	Grade
World-class deposit	usually >5 g/t with scale, or 1–2 g/t at very large scale (Cadia, Boddington)

Cut-off grades typically:

- Open pit: 0.3–0.8 g/t
- Underground: 2–4 g/t

ASX framing: PC2 drill intercepts in the 5–15 g/t Au range over meaningful widths place it in the high-grade open pit / underground crossover zone — well above marginal, well above the cut-off.

MI6 discoveries in similar high-grade ranges. The grade range matters more than any single eye-catching intercept — Module 4 covers how to read drill announcements properly.

Copper (% Cu)

Deposit type	Typical grade
Porphyry (large open pit)	0.3–0.8% Cu
World-class porphyry	>0.6% Cu with billion-tonne scale
IOCG (Olympic Dam style)	variable; 0.5–2% Cu often with Au, U, REE credits
Sediment-hosted (Zambian/Congo style)	1–5% Cu
Underground vein	1–5% Cu typical

Often quoted as **CuEq** (copper equivalent) — combines Cu with by-product credits (Au, Ag, Mo). Always check the assumptions: which prices, which recoveries, are by-product recoveries realistic? See the dedicated section below — equivalent grades are the single most common misread in mining announcements.

Lithium

Spodumene (hard rock): measured in % Li₂O

- Greenbushes (WA, world's best): ~2.0%+ Li₂O
- Pilgangoora, Mt Marion, Wodgina: 1.0–1.5% Li₂O typical
- Cut-off: 0.5–0.7% Li₂O
- Below 1% is generally marginal in the current market

Brine (Salar): measured in mg/L Li or ppm

- Atacama (Chile, world's best): 1500–2700 mg/L
- Argentine salars: typically 200–800 mg/L

- Mg:Li ratio matters — high magnesium = expensive processing

ASX framing: GLN's Hombre Muerto West sits at 859 mg/L average (with recent tests hitting 981 mg/L) — upper tier for Argentine brines, but well below Atacama. Critically, the Mg:Li ratio is low at HMW, which is why the project is economic where higher-grade salts with worse impurity profiles aren't. **ELV's** spodumene operation in Quebec operates in the typical 1.0–1.5% Li₂O range; the economics there depend on hard-rock conversion costs and spodumene pricing more than on grade differentiation.

Sedimentary / clay: newer, no proven commercial production at scale (yet). Treat with skepticism.

Nickel

Sulphide nickel: % Ni

- 0.5–1% = marginal
- 1–2% = average
- 2%+ = high grade
- World-class: Voisey's Bay, Norilsk, Mt Keith

Laterite nickel: % Ni (different processing — HPAL or ferronickel)

- 0.8–1.5% typical
- Capex-intensive; HPAL plants have a brutal history of cost overruns and operational issues

Iron Ore (% Fe)

- DSO (Direct Shipping Ore): >58% Fe is the rough benchmark for sale
- Premium: >62% Fe (the Pilbara benchmark)
- Magnetite concentrates: 65–70% Fe after beneficiation
- Hematite from raw mining: 55–62% typical
- Below 50%: marginal without significant beneficiation

Impurities matter as much as grade — silica, alumina, phosphorus all carry penalties. Discounts and premiums are real money.

Rare Earths (REE)

Reported as **TREO%** (Total Rare Earth Oxide). But TREO alone is misleading.

- TREO grade: 1–5% typical for hard rock; 0.05–0.2% for ionic clay deposits (much lower but easier to process)

- **The split matters more than total grade.** What % of the TREO is the magnet rare earths — Nd, Pr, Dy, Tb? Cerium and Lanthanum are oversupplied and worth little.
- "NdPr%" of TREO is the key second number. >20% is good.
- Mt Weld and Mountain Pass are world-class because of grade *and* split.

Uranium (U_3O_8 or % e U_3O_8)

- 0.05–0.1% U_3O_8 (500–1000 ppm): low grade but mineable in Australia (e.g., in-situ leach amenable)
- 0.2–1% U_3O_8 : average to good
- 1%+ U_3O_8 : world-class (Athabasca Basin, Canada)
- Cigar Lake / McArthur River: 15%+ U_3O_8 (the only deposits in this league globally)

Silver (g/t Ag)

Often a by-product. As primary commodity:

- 50–150 g/t Ag = average
- 150–400 g/t = high grade
- 500+ g/t = bonanza

Silver-equivalent (AgEq) calculations are common — same caveats as CuEq below.

Zinc / Lead (% Zn, % Pb)

Usually reported together as Pb+Zn% or as Zn-equivalent:

- 5% Pb+Zn = marginal
- 8–12% Pb+Zn = average
- 15%+ Pb+Zn = high grade
- Often by-product silver credits

Equivalent grades — the recurring trap

This is the single most common way mining announcements mislead retail. It deserves a dedicated section because it sits at the intersection of every other concept in this module.

What equivalent grades are

When a deposit contains multiple commodities, companies often report a single "equivalent" grade by converting all by-products into units of the primary commodity at assumed prices and recoveries.

So a deposit with 1.0 g/t Au, 5 g/t Ag, and 0.5% Cu might be reported as "1.8 g/t AuEq" or "8.5 g/t AgEq" or "1.1% CuEq" depending on which metal the company chooses as the primary.

Why this is a trap

The headline equivalent grade can be misleading in three different directions, and you need to check all three:

- 1. Price assumptions.** The conversion uses commodity prices set by the company. If gold is priced at US\$3000/oz and copper at US\$10,000/t in the calculation, the equivalent figure looks great. If gold drops to US\$2000/oz, the "AuEq" calculation no longer reflects reality. Always check the price deck in the footnote and ask: do these prices look conservative, current, or aggressive?
- 2. Recovery assumptions.** Different metals recover at different rates from the same ore. Gold might recover at 92%, copper at 88%, silver at 75%. If the company assumes 90% recovery across all metals to simplify the equivalent calculation, the number is inflated. Real plants don't recover all metals equally.
- 3. By-product reality.** Some by-product credits never get realised because:
 - The off-take agreement doesn't pay for them (gold credits in copper concentrate are paid; rare earth credits often aren't)
 - The processing flowsheet doesn't actually recover them at scale
 - Smelter terms charge penalties that consume part of the credit value
 - The market for the by-product is too thin or volatile to monetise reliably

The "low primary grade hidden by big credits" pattern

When you see a low equivalent grade for the primary metal, it usually means the actual primary grade is much worse than what the equivalent suggests.

Example pattern: a company announces "1.5 g/t AuEq" sounds reasonable for an open pit. But when you read the footnotes:

- Actual gold grade: 0.4 g/t

- The remaining 1.1 g/t AuEq comes from silver, lead, and zinc credits
- The economics are now entirely dependent on silver/lead/zinc prices, recoveries, and concentrate marketability — not gold

This is structurally different from a project that's actually 1.5 g/t Au with no by-products. The risk profile is different. The investor case is different. But the headline reads identically.

How to read equivalent-grade announcements

Always do this short check:

1. **What's the primary metal grade alone?** The footnotes will state it. If they don't, that itself is a flag.
2. **What's the by-product split?** What % of the equivalent is coming from each metal?
3. **What price deck is used?** Are those prices current, lagged, or forward?
4. **What recoveries are assumed?** Are they consistent with the metallurgical testwork actually done?
5. **Does the off-take or smelter pay for those by-products?** Polymetallic concentrates have notoriously complex terms.

If the actual primary grade is below the marginal threshold for that commodity (Module 3 benchmarks above), the project's economics live or die on the by-products — even if the equivalent number looks healthy.

When equivalent grades are legitimate

They are not always misleading. A genuine polymetallic deposit (think Olympic Dam — copper, gold, uranium, REE, silver) genuinely produces all those metals and the equivalent is a useful summary measure. A genuine VMS deposit (zinc, lead, copper, silver, gold) is sold as a polymetallic concentrate and the equivalent reflects realised value.

The trap is when equivalent grades get used to dress up a project where the primary metal isn't economic on its own and the by-product credits are speculative.

One-line rule

Always check the primary grade alone. If the primary alone is below the marginal threshold, the equivalent grade is doing real work in the announcement — find out exactly what work, before the equivalent goes into your thesis.

How to spot a tier-1 deposit in 30 seconds

A tier-1 deposit generally combines:

1. **Scale** — billion-tonne+ for porphyry, 100Mt+ for sulphide-hosted base metals, multi-Moz for gold
2. **Grade** — well above the marginal threshold for that deposit type (and on the *primary* metal, not equivalent)
3. **Strip ratio** — favourable, especially in early years
4. **Metallurgy** — clean, high recovery, no refractory issues
5. **Infrastructure** — close to power, water, roads, port
6. **Jurisdiction** — Australia, Canada, USA, Scandinavia, parts of South America are tier-1; many parts of Africa/Asia carry sovereign risk discounts

If a deposit checks 5/6 of those, it'll get built. If it checks 2/6, it'll bounce around as a "story" stock for years.

How to spot a marginal deposit

- Primary metal grade close to or below the cut-off for that commodity
- Headline grade is an "equivalent" grade that's doing heavy lifting from speculative by-products
- Strip ratio above 6:1 for non-precious metals
- Metallurgical recovery below 75%
- Remote with no infrastructure
- Sovereign risk jurisdiction
- High capex relative to projected NPV (capex > NPV is a red flag)

These projects often "become economic" only at peak commodity prices, then disappear when prices normalise.

The grade-tonnage trade-off

Big and low-grade vs small and high-grade — both can work, but not always for the same investor.

- **Large/low-grade** (e.g., porphyry copper at 0.5%): high capex, long mine life, leverage to commodity price, harder to permit. Good for majors, harder for juniors.
- **Small/high-grade** (e.g., narrow-vein gold at 15 g/t): lower capex, shorter mine life, faster payback, easier to permit. Good for juniors and mid-tiers.

A junior pursuing a 1bn tonne porphyry needs \$2bn capex — they will get diluted to oblivion or sell to a major. A junior with a 500koz high-grade UG project at 8 g/t can self-fund construction.

Match the deposit to the company's ability to develop it.

Practical exercise

For every deposit you're tracking, write down:

1. Commodity and deposit type
2. Resource grade and where it sits on the benchmark above
3. **If reported as an equivalent grade — what's the primary metal grade alone?**
4. Strip ratio (if open pit)
5. Recovery rate from met testwork
6. Cut-off applied to the resource
7. Is this a tier-1, tier-2, or marginal deposit?

If grade context is missing from your thesis, you are speculating on the *story*, not the *rock*.

What I'm uncertain about

- "Marginal" thresholds shift with commodity prices. A 1 g/t Au open pit was uneconomic at \$400/oz gold; it's clearly economic at \$3000+/oz. Always think in current/forward price context.
- REE economics are unusually complex because pricing varies enormously across the 17 elements and end-market demand is policy-driven (magnet demand for EVs, wind turbines, defence). Specific cut-offs change with NdPr pricing.
- Lithium grade thresholds have moved a lot through the 2022–2025 cycle; some marginal projects that were viable at \$80k/t spodumene are deeply uneconomic at \$700/t. Verify against current commodity prices.
- Equivalent grade conventions vary by company and broker. The math is simple in principle but the assumptions are everything. When in doubt, ignore the equivalent and look at the primary metal grade alone.

Module 4: Reading Drill Result Announcements

Why this matters

Drill result announcements are where companies most actively shape perception. The headlines are often technically accurate but designed to obscure what an experienced reader would immediately spot.

Learning to read drill results properly is the single highest-leverage skill in junior mining investing. Get this right and you avoid 80% of the traps.

The standard format

A drill result announcement typically headlines an intercept like this:

“**PRG02 returned 24m @ 3.5 g/t Au from 85m, including 6m @ 11.2 g/t Au from 92m**”

You need to decode every part of that:

- **PRG02** — hole identifier
- **24m** — down-hole length of the intercept
- **3.5 g/t Au** — length-weighted average grade across that 24m
- **from 85m** — depth at which the intercept starts
- **including** — a sub-interval within the parent intercept (almost always the high-grade core)
- **6m @ 11.2 g/t Au from 92m** — that high-grade core

What's *not* in the headline:

- True width
- Hole orientation and dip

- Rock type / mineralisation style
- Whether top-cut was applied
- Recovery / RQD (rock quality)
- Assay method

These are usually in the JORC Table 1 appendix, which retail rarely reads. Read it.

Length-weighted vs grade-weighted

Headline grade is **length-weighted**:

$$\text{weighted grade} = \Sigma(\text{interval}_i \times \text{grade}_i) / \Sigma(\text{interval}_i)$$

Example: 24m intercept made of:

- 18m @ 1.0 g/t (low-grade halo)
- 6m @ 11.0 g/t (high-grade core)

$$\text{Length-weighted} = (18 \times 1.0 + 6 \times 11.0) / 24 = (18 + 66) / 24 = \mathbf{3.5 \text{ g/t}}$$

Looks like a uniform 24m of 3.5 g/t mineralisation. Reality: 6m of bonanza surrounded by marginal halo. Whether you can mine that depends entirely on geometry, mining method, and selectivity.

Always read the "including" sub-intervals to see what's really there.

Down-hole length vs true width

Drilling is often angled to intersect mineralisation perpendicular (or as close as possible). The reported intercept is **down-hole length** — the actual length of drilled core.

True width is the actual thickness of the mineralised body measured perpendicular to its strike and dip.

If the hole is drilled at an angle to the mineralisation:

- Hole drilled exactly perpendicular: true width = down-hole length
- Hole drilled at 45° to mineralisation: true width $\approx$ down-hole length $\times$ 0.7

- Hole drilled subparallel (worst case): down-hole length wildly exaggerates true width

Companies are required to disclose true width or state it cannot be determined. When you see "true width estimated at ~70% of reported length", apply that mentally to all intercepts.

When the announcement says nothing about true width, check the long-section diagram (if provided) and figure it out yourself.

Gram-metres (g·m) — the universal currency

Professional analysts rank holes by **gram-metres** (or %·m for base metals):

$$\text{g}\cdot\text{m} = \text{grade} \times \text{down-hole length}$$

So 24m @ 3.5 g/t = **84 g·m**

This single number lets you compare holes across different deposits, different widths, different grades.

Rough benchmarks for gold drill holes:

- <20 g·m = barely interesting
- 20–50 g·m = solid intercept
- 50–100 g·m = strong hole
- 100–500 g·m = exceptional
- 500+ g·m = world-class / company-making

For copper porphyry: %·m equivalent (e.g., 200m @ 0.6% Cu = 120 %·m).

Use g·m to cut through the marketing. A "100m @ 0.5 g/t" intercept (50 g·m) is being sold as more impressive than "5m @ 12 g/t" (60 g·m), but the latter is actually a better hit — and far easier to mine.

Drilling-type taxonomy — what kind of hole are you looking at?

Not all drill holes have the same purpose, and the same intercept means very different things depending on what kind of drilling produced it. This is the most under-appreciated part of reading announcements properly.

Grade control drilling

- **Purpose:** confirm the grade and geometry of ore *immediately about to be mined* — typically 1–6 months ahead of mining
- **Spacing:** very tight, often 5–10m
- **Where you see it:** producers (Stage 10) and projects in late commissioning (Stage 9)
- **What it tells the market:** very little new about the deposit. The ore was already in the reserve. Grade control just confirms the model.
- **Headline grade is usually high** because the drilling is happening inside high-grade ore that's already classified as Indicated/Measured. That's by design — they're not going to grade-control drill barren rock.
- **Read:** unless the grade control results materially differ (positively or negatively) from the reserve model, this is a routine operational update, not a discovery event.

Infill drilling

- **Purpose:** raise the geological confidence of an existing resource — typically converting Inferred to Indicated, or Indicated to Measured
- **Spacing:** moderate, 25–50m typically
- **Where you see it:** Stage 4–7 (resource definition through DFS)
- **What it tells the market:** confirms the deposit. Strike rate is typically very high (you're drilling inside known mineralisation).
- **Headline is usually a high hit rate** — "12 of 12 holes intersected mineralisation" is normal for infill, not exceptional
- **Read:** valuable for resource confidence and converting categories, but doesn't grow the deposit. The catalyst is the resource update that follows, not the individual hole results.

Step-out drilling (sometimes called "extension drilling")

- **Purpose:** test whether mineralisation continues *beyond* the known boundary of the resource — along strike, at depth, or in a parallel structure
- **Spacing:** wider, 50–200m typically, depending on the deposit type
- **Where you see it:** Stage 3–4 (discovery through resource definition)
- **What it tells the market:** real new information. A successful step-out grows the deposit. A failed step-out defines the edge.

- **Strike rate is much lower** — geology pinches, faults intervene, grade falls off — so successful step-outs deserve attention
- **Read:** the most informative single category of drilling. A 50 g·m step-out hole 200m beyond the known boundary is more valuable than a 100 g·m infill hole inside the known zone.

Exploration drilling (greenfield / regional)

- **Purpose:** test a new prospect that isn't part of an existing resource — geophysical anomaly, surface geochemistry, structural target, satellite deposit
- **Spacing:** very wide, often 100–500m
- **Where you see it:** Stage 1–3 (concept through discovery), and around producing operations seeking new ore
- **Strike rate:** low. Most exploration holes return nothing.
- **Read:** when an exploration hole hits, the asymmetric return is enormous. A discovery hole on a greenfield target is the biggest single catalyst in mining (Stage 3 from Module 1). When it misses, the SP can decay fast as the prospect is de-risked downward.

Why the distinction matters

The same announcement reading "drill program returns 30m @ 4 g/t Au" tells you completely different things depending on which type of drilling produced it:

- **Grade control:** routine, no SP impact (already in the reserve)
- **Infill:** confirms the model, modest SP impact (resource confidence improvement priced separately)
- **Step-out:** material SP impact (deposit is bigger than the prior boundary)
- **Exploration:** potentially huge SP impact (new discovery)

The announcement headline rarely tells you which kind of drilling produced the intercept. The body and the supporting plan-view diagrams will. **Read them.**

A useful framing: ask "is this hole inside or outside the existing resource boundary?" If inside, treat it as routine confirmation. If outside, treat it as new information, and ask how far outside and in what direction.

Common cherry-picking and obfuscation patterns

1. "Up to" reporting in headlines

“Up to 45 g/t Au returned from drilling”

That's a single peak assay over a 1m sample interval. The actual intercept might be 1m @ 45 g/t with 19m of barren rock either side. Always look for the **weighted intercept**, not the peak assay.

2. Loose "includes" wrapping

A 50m parent intercept at 0.4 g/t with a 1m core at 12 g/t lets the company headline either:

- "50m @ 0.7 g/t including 1m @ 12 g/t" (sounds bulk-mineable)
- Or just "1m @ 12 g/t" (sounds high-grade)

Same hole, two stories. Both true. Neither tells you if the deposit has scale *and* grade together.

3. Composite reporting hiding internal waste

A reported "20m @ 2.0 g/t" intercept might actually be:

- 5m @ 6 g/t
- 10m @ 0.1 g/t (essentially waste)
- 5m @ 4 g/t

Composited, it averages 2.0 g/t and looks like a single zone. In reality, two separate thin lodes with waste between them. Mining selectively to avoid the waste is much more expensive.

Look at the **assay table** in the appendix. The granular metre-by-metre assays tell the truth.

4. No top-cut applied

A rogue 200 g/t assay in a single sample can lift a 30m intercept from 1.5 g/t to 8 g/t on paper. Best-practice resource estimation applies a top-cut (e.g., cap all assays at 30 g/t). Drill result announcements *don't* usually apply top-cuts because doing so kills the headline. Mentally cap any single-metre assay above ~5x the bulk grade.

5. Hole orientation games

A hole drilled subparallel to a vein system will produce huge down-hole intercepts that have small true widths. "100m @ 5 g/t" sounds amazing but if true width is 8m, it's a far less impressive hole.

6. Selective announcement timing

The good holes get headlined immediately. The poor holes from the same program get bundled into a quarterly months later, or never reported individually. Cross-reference the program's planned holes (in earlier announcements) against the holes that were actually reported on. If they drilled 30 holes and you've only seen results from 12, the other 18 weren't great.

7. Mislabelling the drilling type

A subset of the prior issue. Companies sometimes describe routine infill drilling as if it were step-out or discovery, hoping the reader doesn't check the plan view. The plan-view diagram will show whether holes are inside or outside the existing resource shell. If the diagram is omitted or unclear, that's a flag.

Headline vs body — the discipline

Most retail read the headline of a drill announcement and stop. That's the trap. The body of the announcement contains the qualifications, the what-isn't-said, and frequently the actual story.

What to read in the body, in order

1. **The intercept table** — every reported intercept with hole ID, from-to depths, length, grade, and (ideally) true width. Compare reported holes against the program plan announced earlier — what holes are missing?
2. **The methodology / collar table** — hole positions, dips, azimuths. Lets you reconstruct geometry vs true width.
3. **The geological description** — does the company describe the rock type, structural setting, alteration? Vague descriptions are a flag.
4. **The plan view and long-section diagrams** — where are the holes in space? Are they inside or outside prior drilling? In what direction does the resource grow?
5. **The forward plan section** — what does the company say about next steps, next drilling, next assays? This is where catalysts get telegraphed.

6. **The JORC Table 1 appendix** — sampling, sample preparation, QA/QC, lab procedures, top-cut treatment, density assumptions.

The "no headline" tell

Sometimes a company drills a program, waits, and then publishes results in a quarterly activities report rather than a dedicated drill-results announcement. **The absence of a headline announcement is itself information.** If the same company would have headlined a strong intercept, the lack of a headline often means the results were below expectations.

The contrast pattern looks like this. Two companies drill maiden holes in the same week:

- **Company A** publishes a dedicated ASX announcement titled "Strong Discovery Hole Confirms High-Grade Mineralisation" with a plan view, long section, and 5 holes with grades, prefacing future drilling
- **Company B** mentions the program in passing in their quarterly: "drilling has been completed at site X with results currently being interpreted"

A is making the most of strong results. B is downplaying weak ones. The body of the announcements (or the absence of an announcement) tells you which is which.

The "body matches headline" check

A good announcement is one where the body confirms and elaborates on what the headline claims. The headline says "exceptional intercept"; the body shows hole-by-hole assays that include exceptional intercepts. The headline says "step-out drilling extends mineralisation 200m"; the body shows a plan view with the new holes plotted 200m beyond the prior boundary.

A bad announcement is one where the body softens the headline:

- Headline: "Significant new gold zone discovered"
- Body: 1m @ 4 g/t in a single hole, no follow-up drilled, geology speculative, true width unclear

If you can't reconcile the headline to specific data points in the body, the announcement is doing more work as marketing than as information.

Cross-reference against other companies' announcements on the same district

When a company in a known mineral district announces results, look at what the neighbours are saying. Often other companies on adjacent tenements publish results that contextualise — sometimes they show the structure pinches out, or grades fall off in a particular direction, or a metallurgical issue is more widespread than originally implied. The full picture rarely lives in one company's announcement.

What a *good* drill announcement looks like

Green flags when reading:

- True width is disclosed for each intercept
- Long sections and plan view diagrams included
- Multiple holes, in step-out configuration, all showing economic intercepts
- Grade continuity along strike and at depth, not just one bonanza pod
- Geological description matches a known deposit model
- JORC Table 1 is detailed and consistent with prior announcements
- Assay table provided as an appendix
- Independent CP signs off
- Drilling type explicitly stated (step-out vs infill vs exploration)
- Forward plan section identifies the next batch of catalysts

When you see all of that, you're looking at a serious result. When you see "up to" headlines, no diagrams, and one big intercept with everything else "pending", treat the announcement as marketing.

A practical reading checklist

For every drill result announcement, answer:

1. **What type of drilling produced this intercept — grade control, infill, step-out, or exploration?**
2. **What's the headline intercept in g·m?**
3. **What's the true width vs down-hole length?**
4. **Is the headline a single hole or multiple holes?**
5. **Where does the hole sit — inside the existing resource boundary or outside?**
6. **What's the high-grade core stripped out, vs the parent interval?**
7. **Are there any single-metre assays driving the average?**

8. **How many holes from the program are still pending? Why?**
9. **Does the body of the announcement substantively support the headline, or does it walk it back?**
10. **Does the geology match the deposit model the company is selling?**

You will be ahead of 95% of retail if you do this for every announcement on every stock you hold.

Pre-event positioning patterns

The drill result trade is well-known and well-played:

- **Pre-drilling ramp:** SP runs as the market positions for assays, often 20–100% above pre-news levels.
- **Assay window:** depending on lab queues and program size, results can take weeks. SP often peaks 1–5 days *before* the actual announcement as insiders/anticipators take profit.
- **Sell the news:** even good results often see a SP fade because the "best case" was already priced in.
- **Surprise upside:** genuinely exceptional holes (200+ g·m for gold) trigger fresh leg up because they exceeded the priced-in best case.

Don't confuse "good drill result" with "good trade". The trade depends on what was already priced in.

What I'm uncertain about

- Lab turnaround times have varied considerably with commodity cycle activity (gold/lithium booms cause queue delays). A "results expected in 4 weeks" guidance can stretch to 8+ weeks during peak cycles.
- ASX continuous disclosure thresholds for what constitutes price-sensitive vs not have been refined a few times — companies retain meaningful discretion in batching results.
- The drilling-type taxonomy above is conventional but the boundaries are fuzzy. A "step-out" hole at 30m beyond the boundary is still partially confirming continuity rather than testing entirely new ground. Use the framework as a guide to reading announcements, not as a hard classification.

Module 5: Economic Studies (Scoping, PFS, DFS)

Why this matters

Economic studies are the company's formal pitch that a deposit can be turned into a profitable mine. They are also where the most consequential lies-by-omission happen — base case prices set 30% above spot, capex estimates that consistently understate by 30–50%, NPV calculations that ignore real-world ramp problems.

Read every study with the assumption that it represents the company's most optimistic internally-defensible case. Then mentally degrade it.

The three studies, in order

Scoping Study (in Canada: PEA — Preliminary Economic Assessment)

- **Accuracy:** $\pm 35\text{--}50\%$
- **Cost:** \$200k – \$1m
- **Timeline:** 3–6 months
- **Resource categories allowed:** Inferred, Indicated, Measured (yes, Inferred is allowed at this stage — important caveat)
- **Purpose:** "Is this worth taking further?"
- **Output:** rough capex, opex, NPV, IRR, mine plan concept
- **NOT bankable.** Cannot be used to raise project debt.

PFS (Pre-Feasibility Study)

- **Accuracy:** $\pm 20\text{--}25\%$

- **Cost:** \$2m – \$10m
- **Timeline:** 9–18 months
- **Resource categories allowed:** Indicated and Measured only (Inferred excluded from reserves and economics)
- **Purpose:** "Is this technically and economically robust?"
- **Output:** maiden Ore Reserves declared, refined capex/opex, NPV, IRR with sensitivities
- **Bankable for some lenders, but most still want DFS.**

DFS / BFS (Definitive / Bankable Feasibility Study)

- **Accuracy:** ±10–15%
 - **Cost:** \$10m – \$50m+
 - **Timeline:** 12–24 months
 - **Resource categories allowed:** Indicated and Measured only
 - **Purpose:** "This is the project we will build."
 - **Output:** detailed engineering, procurement-ready BOMs, refined reserves, lender-grade economics
 - **Bankable for project finance.** The basis for FID (Final Investment Decision).
-

The scoping → PFS resource confidence step-up — why it matters

The progression from Scoping to PFS is not just an accuracy tightening. It's a fundamental change in what categories of resource can carry the economics.

Scoping studies are allowed to use Inferred resources in the production schedule and NPV calculation. This is one reason scoping economics often look better than the equivalent PFS economics on the same deposit — Inferred resources get included in the mine plan that won't survive into the PFS.

PFS and DFS are not allowed to use Inferred resources in reserves or in the economic case (with extremely narrow exceptions). When the project moves from scoping to PFS, all the Inferred tonnage that was carrying production in years 6–15 gets removed, replaced only by what's been upgraded to Indicated or Measured by additional drilling.

The practical consequence is that NPV often falls between scoping and PFS not because the project is worse, but because the rules tightened. A scoping study showing a 15-year mine life on a deposit that's 70% Inferred will become a PFS showing a 6-year mine life on the same deposit unless the company does the infill drilling required to upgrade Inferred → Indicated.

The "PFS-downgraded-to-scoping" tactic

A pattern worth flagging: a company has previously released a scoping study with strong numbers, then quietly re-classifies subsequent work as a "scoping update" rather than progressing to PFS. The reason is usually that the PFS-allowable resource (Indicated + Measured only) wouldn't support the economics the scoping showed.

Another version: a project that has been at PFS-stage in prior announcements gets walked back into a "revised scoping study" without explanation. That's a signal the PFS-grade economics didn't work, and the company is trying to reset to a less rigorous standard.

When you see a study labelled in unusual terms — "preliminary economic update", "scoping refresh", "concept study" — read carefully. The labelling matters because the rules differ. A scoping study with 70% Inferred in the production schedule is a different document from a PFS based only on Indicated and Measured.

Key outputs in every study

NPV (Net Present Value)

The discounted value of all future cash flows from the project, minus initial capex.

- Usually presented as **post-tax NPV at 5-8% discount rate**
- Pre-tax NPV is also often shown — pre-tax NPV is always higher; companies sometimes lead with pre-tax for the bigger number
- Sensitive to commodity price assumption and discount rate

Rule of thumb: project NPV should be at least 2-3x capex for a developer-stage company. NPV = capex means barely worth building.

IRR (Internal Rate of Return)

The discount rate at which NPV = 0. Measures capital efficiency.

- <15% IRR: marginal, hard to finance
- 15–25%: solid project
- 25–40%: strong project
- 40%+: excellent (or assumptions are aggressive)
- 60%+: assumptions are almost certainly aggressive — verify the price deck and capex

Capex

Initial capital expenditure to build the project.

- Pre-production capex: actual build cost
- Sustaining capex: ongoing capital to maintain production through mine life (added separately to AISC)

Always check what's included. Some studies exclude:

- Owner's costs (the company's own salaries, financing fees)
- Working capital
- Contingency below industry-standard levels
- Land acquisition / royalty buyouts
- Power/water infrastructure if "to be provided by third party"

Opex (Operating Cost)

Cost per tonne of ore processed or per unit of metal produced. Usually presented as:

- \$/t ore (mining + processing)
- C1 cash cost (direct production cost per unit)
- AISC (all-in sustaining cost — see below)

AISC (All-In Sustaining Cost)

Originally a World Gold Council metric (2013), now used loosely across commodities. Includes:

- Cash production costs
- Royalties
- Sustaining capex
- Site G&A
- Reclamation accruals

It does NOT include:

- Initial capex

- Exploration capex (unless sustaining)
- Corporate G&A (sometimes)
- Income tax
- Interest expense

Different companies calculate AISC differently. **Compare AISC across companies with caution.**

For a project to be robust, AISC should sit comfortably below the long-term commodity price. AISC at 80% of spot = healthy margin. AISC at 95% of spot = no margin, no resilience.

Payback Period

Years until cumulative cash flow recovers initial capex.

- <2 years: excellent
- 2-4 years: solid
- 4-6 years: marginal
- 6+ years: high risk in a cyclical commodity

Mine Life (LOM — Life of Mine)

Years of production at planned throughput based on current reserves.

- <5 years: short, hard to attract debt
- 7-15 years: typical for mid-tier projects
- 20+ years: tier-1 reliability

LOM extends if exploration adds resources and conversion to reserves continues.

Natural run rate — what plants actually do

A concept that doesn't appear in most studies but matters enormously for how you read them.

Natural run rate is the throughput a plant actually achieves in steady-state operation, independent of nameplate. It's almost always different from nameplate, and the direction of the difference depends on the plant's history.

Below nameplate (most plants in the first 1-3 years):

- Recovery shortfalls require slower throughput
- Bottlenecks in specific circuit components
- Ore hardness exceeds design assumptions
- Mining can't deliver enough ore at design grade

At nameplate (well-engineered plants reaching design after 12–24 months):

- The original design parameters were realistic
- Operations team has worked through commissioning issues

Above nameplate (mature operations with optimisation):

- Bottleneck debottlenecking
- Operator skill has improved beyond design assumptions
- Mining is delivering above-grade ore consistently

The Bellevue Gold case is the canonical example of a plant that operated materially below nameplate for an extended period — design said one tonnes-per-year, actual was meaningfully lower. This delta drives everything: AISC rises (fixed costs across less production), revenue falls, payback extends, dilution risk increases.

How to read announcements about run rate

When a producer or commissioning project reports actual throughput vs nameplate, what you're looking for is:

1. **The trajectory.** Is throughput improving quarter-on-quarter, or stuck below nameplate?
2. **The reason.** "Plant down for scheduled maintenance" vs "ongoing recovery optimisation" vs "ore characteristics different from design" — each implies a different fix and a different timeline.
3. **The forward guidance.** Has management revised guidance down? Multiple guidance revisions in the same direction is a tell.
4. **The cash flow consequence.** AISC at actual throughput vs design is the real margin you can model.

Studies typically don't quantify natural run rate explicitly. You back it out from the production guidance and the historical achievement vs guidance. A project guiding 5 ktpa nameplate but historically achieving 4 ktpa has a natural run rate around 4 ktpa, and the economics need to be re-cut at that level.

Sensitivity tables — where to focus

Every study includes a sensitivity table showing how NPV/IRR changes with input changes:

Sensitivity	-20%	Base	+20%
Commodity price	\$X	\$Y	\$Z
Capex			
Opex			
FX rate			
Discount rate			

Where to focus:

1. **Commodity price assumption.** What price did they use? If the base case price is *above* current spot, the project is more sensitive than it appears at first glance. Recalculate NPV at current spot prices.
2. **Capex sensitivity.** A project with NPV that flips negative on +20% capex is a very fragile project, given that 30-50% capex blowouts are normal.
3. **Discount rate.** Most companies use 5-8%. Use 10% as a stress test for junior/risky projects.

Capex blowout — the pattern that always repeats

Capex estimates almost always rise from study to study, then again from DFS to actual build:

Stage	Typical capex vs final actual
Scoping	40-70% of final actual
PFS	60-85% of final actual
DFS	75-95% of final actual
Final actual	100% (often more)

This is not because companies are stupid. It's because:

- Engineering detail increases over time, finding scope that wasn't visible earlier
- Inflation in steel, copper, labour, EPC contracts

- Permitting conditions add infrastructure (water treatment, dust suppression, road upgrades)
- Currency moves
- Site conditions different from assumed (geotech, hydrology)

Stress-test every DFS by adding 25-35% to capex and see if the project still works at current spot commodity prices.

Capex tactics worth recognising

Plant relocation / refurb capex savings — usually less than advertised

A company acquires a second-hand processing plant from a closed operation, plans to disassemble, transport, and re-erect it on their site, and headlines a major capex saving versus building new.

The reality is usually:

- **Disassembly and transport costs** are larger than headline implies. Cranes, trucking, port logistics, storage at destination — all expensive
- **Component condition** varies. Mills, crushers, and tanks may need refurbishment, replacement liners, new motors. The "refurbished" plant frequently needs \$20-50m of work to be operational
- **Site-specific modifications** — every site has different ore characteristics, climate, layout. The plant has to be re-engineered to match. This is engineering work that wasn't in the headline saving
- **Permitting** for a re-located plant can take longer than for a new build, because regulators look at it as a new installation
- **Lead time** to get it operational is often longer than building new — disassembly takes time, components arrive in the wrong order, re-engineering reveals issues during reassembly

The BML (Bellevue Gold's predecessor approach for some assets) and similar refurb projects in the gold and base metals sector have had mixed track records. Some saved money. Several didn't. The pattern is that the headline saving is usually 60-70% of the eventual reality after refurbishment.

If a company headlines significant capex savings via a refurb, ask:

1. Where was the plant before? What condition?
2. What's the refurbishment scope and budget? Is it included in the headline capex?

3. What's the lead time from acquisition to first ore — and is that realistic?
4. Has the company done this before? What was their track record?

Staged capex deferral

A company splits the project into Phase 1 (low capex) and Phase 2 (higher capex, deferred to later years and supposedly funded from Phase 1 cash flow). The headline NPV often combines both phases.

This looks fundable. In reality:

- **Phase 1 cash flow rarely materialises at the predicted level** because of the natural-run-rate / ramp-up issues described above
- **Phase 2 capex is rarely funded purely from Phase 1 cash flow** — requires another raise or debt facility
- **The headline NPV depends on Phase 2 actually happening** — but Phase 2 frequently gets indefinitely deferred when commodity prices fall, when Phase 1 underperforms, or when management changes
- **Sub-scale operations result** when Phase 2 is permanently deferred — Phase 1 alone rarely has the economics to justify the overheads

The Santana Minerals (SMI) staged approach to its Bendigo-Ophir gold project is a recent example of staged development being announced with the intent to fund expansion from operating cash flow. The honest read on staged capex deferral is: model Phase 1 standalone and treat Phase 2 as optionality, not part of the base case. If Phase 1 NPV alone doesn't justify the build, the headline phased NPV is doing more work than it should be.

Stockpile inclusion in mine plan

Some studies fold low-grade stockpile (rock that's been mined but not processed because it's below cut-off grade for current operations) into the production schedule for late mine life. The rationale is that as commodity prices rise or processing costs fall, the stockpile becomes economic to process.

This can be legitimate. It can also be padding to extend headline mine life and total contained metal. Tells:

- Stockpile economics are based on a future commodity price assumption that's higher than current spot
- The metallurgical recovery on stockpiled material is assumed to be the same as fresh ore (often it's lower because the rock has weathered)
- The stockpile contributes meaningfully to LOM tonnes but minimally to NPV (because it's late in life and discounted)
- The stockpile inclusion is in the body of the announcement but not in the Reserves table

The LGM (an example flagged in industry commentary) included stockpile material in their mine plan in a way that materially extended LOM headlines. When the stockpile was excluded, the project mine life was meaningfully shorter and the economics weaker.

If a study includes a stockpile in the production schedule, ask:

1. What % of total contained metal comes from stockpile vs fresh ore?
 2. What recovery is assumed on the stockpile?
 3. What commodity price is required for stockpile processing to be economic?
 4. Is the stockpile in the Reserve, or is it sitting outside the Reserve as a notional add-on?
-

Permitting and environmental — what timelines are normal

A common misread: a "two-year environmental permitting process" gets framed as a delay or a problem.

A two-year environmental baseline study and approval process is normal in Australia for any new mining project of meaningful scale. It is not a delay caused by activism, government incompetence, or company failure. It's the standard regulatory timeline for:

- 12 months of seasonal environmental baseline data (flora, fauna, water, air, noise) — minimum required
- Government agency review of the EIS (Environmental Impact Statement)
- Public consultation periods
- Indigenous Land Use Agreement (ILUA) negotiation if not already in place
- Conditions of approval being negotiated and finalised

Examples like ALM (Alma Metals) and JJ-style projects through the standard permitting cycle take 18–30 months as a matter of course, not as exceptions. The right question isn't "why is permitting taking so long" — it's "is permitting on schedule for a normal Australian permitting timeline?"

When a company guides "first production 2027" for a project that hasn't started baseline environmental work yet in 2026, that timeline is unrealistic regardless of how confident the company sounds. The regulator dictates the floor on this timeline, not the company. Be sceptical of any project guidance that compresses the environmental timeline below 18 months from baseline start to approval.

Hidden games to watch for

1. Optimistic commodity price deck

The most common manipulation. Industry-standard practice is to use the consensus 5-year forward curve or current spot, whichever is more conservative. Some studies use:

- Long-term incentive prices (the price required to justify new supply) — fine in theory, but optimistic
- "Analyst consensus" cherry-picked from bullish forecasts
- Spot price during a temporary price spike

Always recalculate NPV at current spot. If the project doesn't work, the study is selling you a price view, not a project.

2. Aggressive recovery rates

Met testwork on a few drill samples extrapolated to 90%+ recovery. Real plants typically achieve 5–10% below pilot test rates in early years, sometimes never reach pilot rates.

3. Optimistic ramp profile

Assuming nameplate capacity in Year 1 or even Year 2. Most plants ramp slower — especially complex flowsheets. Stress test with a 12–24 month ramp.

4. Excluding pre-production capitalised costs

Owner's team costs during construction, working capital, financing fees can add 10–20% to total capital required to first cash flow.

5. Cherry-picking the best mine plan years for production guidance

"Average annual production: 200 koz Au" might be true over 10 years but Year 1–3 might be 280 koz (high-grade core mined first) and Year 8–10 might be 130 koz. Most analyst valuations rely on the headline number.

6. Single-product NPV when by-products are required

A copper project with by-product gold credits often shows two NPVs — copper-only and copper+gold. The copper+gold NPV is usually headlined. If the gold credit doesn't materialise (lower grade than predicted, lower recovery), the project economics collapse.

7. "Phased development"

Splitting the project into Phase 1 (small, low-capex) and Phase 2 (larger, capex deferred to later years from operating cash flow). Looks fundable. In reality, Phase 2 often gets indefinitely deferred when commodity prices fall, leaving a sub-scale operation. (Covered in detail above.)

What a *good* study looks like

Green flags:

- Base case commodity prices at or below current spot
- Conservative ramp assumption (12+ months to nameplate)
- Adequate contingency (15%+ on capex)
- Met recoveries based on locked-cycle pilot tests, not just bench-scale
- Reasonable strip ratio assumptions through life of mine, not just early years
- Sensitivity to -20% commodity price still shows positive NPV at 8% discount
- AISC well below long-term commodity price
- Published assumptions fully transparent
- Stockpile contribution disclosed clearly if any
- Permitting timeline realistic vs jurisdictional norms
- If refurb plant: refurbishment scope and budget clearly itemised within capex

When you see all of that, the study is being run as engineering, not marketing.

How to read a study in 30 minutes

1. **Read the announcement headline numbers** (NPV, IRR, capex, AISC, LOM)
 2. **Find the commodity price assumption.** Compare to current spot.
 3. **Find the discount rate.** Re-mentally apply 10% if 5–8% used.
 4. **Read the capex breakdown.** Note contingency level and what's excluded.
 5. **Read the sensitivity table.** Stress to -20% commodity price and +25% capex simultaneously. Does NPV survive?
 6. **Read the production schedule.** Is Year 1 production materially higher than LOM average? (High-grading early years.)
 7. **Compare PFS → DFS capex if both are out.** What was the blowout?
 8. **Check who did the study.** Tier-1 consultancy (SRK, AMC, Lycopodium, GR Engineering, etc.) carries more weight than in-house or unknown consultants.
 9. **Check the resource categories used.** Scoping using Inferred — flag. PFS/DFS using only Indicated/Measured — correct.
 10. **Check for stockpile contribution and refurb plant elements** in the mine plan and capex.
-

Practical exercise

For any study published on a stock you follow, answer:

1. What commodity price was used vs current spot?
2. What's the NPV/capex ratio?
3. What's the IRR pre-tax and post-tax?
4. What contingency was applied to capex?
5. If you stress capex +25% and price -20%, does NPV stay positive?
6. What's the AISC vs current commodity price?
7. **What % of contained metal in the mine plan comes from Inferred (scoping only) or stockpile?**
8. **What's the natural run rate assumption — does the production schedule assume nameplate in Year 1, or a graduated ramp?**

If those answers don't make the project look investable, the company will need either a commodity price tailwind or an M&A bid to make money for shareholders.

What I'm uncertain about

- "Industry-standard" capex contingency varies (10–25%) and there's genuine disagreement among consultants about the right level.
- AISC methodology variations between companies are significant enough that direct comparisons require careful normalisation. World Gold Council guidance helps for gold;

less standardised for other commodities.

- Specific consultant reputations shift over time. A few prominent names have had project blowouts that hurt their credibility — worth checking recent track record on similar projects.
- Refurb plant track records vary enormously; the 60–70% headline savings rule of thumb is a generalisation, not a hard rule. Some refurbishments have come in close to budget and ramped well.
- Permitting timelines vary by jurisdiction within Australia (WA vs Qld vs NT vs SA) and have been getting longer in recent years. The 18–30 month rule of thumb is a starting point, not a guarantee.

Module 6: Capital Structure, Dilution & The Share Register

Why this matters

A great deposit owned by a company with a broken capital structure will not make you money. A mediocre deposit owned by a company with a tight register, aligned management, and disciplined capital management can.

Most retail focus on the project. The professionals focus on the cap table first, then the project. **The cap table tells you how the company has been valued and treated by people with more information than you.** Read it that way and capital structure becomes the most information-dense part of any FA Story.

This module rewrites the Module 6 framework around a single organising idea: **every component of capital structure is a signal about how informed money has chosen to participate in the company.** The mechanics matter, but the signal-reading is what generates edge.

What the cap table tells you that the project doesn't

The project description is the company's pitch. It's the story management wants you to hear. The cap table is the record of every decision informed money has made about the project — at what price, on what terms, with what conviction.

Three contrasts illustrate why this matters more than the project page on the corporate website:

Contrast 1. Two gold juniors with identical-looking projects. Junior A's most recent placement was at a 20% discount to VWAP with one free option for every two shares. Junior B's most recent placement was at a 2% premium to VWAP with no options attached, and the same specialist

cornerstone fund participated three times in 18 months at progressively higher prices. The projects look identical. The signal from the cap tables is opposite.

Contrast 2. Two lithium developers. Developer A has 1.5bn fully-diluted shares after raising through the bear market at progressively lower prices, with major dilution overhang from cheap options. Developer B has 200m shares with directors having materially added to their positions on-market with their own cash over the past 12 months. The deposit could be the same; the per-share economics on a re-rate are dramatically different.

Contrast 3. Two copper explorers, both Stage 3 with strong drill results. Explorer A has a tight register dominated by long-only specialist resources funds. Explorer B has a top-20 register dominated by trading desks and short-term promotional money. A's SP holds gains; B's SP collapses on every up-day as the trading money rotates out.

The project doesn't tell you any of this. The cap table does.

The discipline this module is teaching: read every component of capital structure as a signal, not just as a number. Who participated, at what price, on what terms, in what direction over time? Those four questions give you more about a company's quality than any drill announcement.

The components, briefly

You need to know the mechanics, but the mechanics are the easy part. We'll cover them once and then spend the rest of the module on what they signal.

Shares on issue (SOI)

Total ordinary shares currently issued. Multiplied by share price gives market cap.

Options

Right to buy a share at a set strike price by a set expiry date. Listed options trade on ASX (e.g., XYZO); unlisted options are held by directors, employees, advisors, seed investors. If options are deep in the money, assume they will be exercised. That's pending dilution.

Performance rights / performance shares

Zero-strike (or near-zero) options that vest on milestones — maiden resource, completing DFS, first production, share price targets. Pure dilution when they vest, no cash inflow.

Convertible notes / loans

Debt that converts to equity at a set price (often a discount to market). Common in distressed financing situations. Watch for floor conversion prices and "ratchet" or "death spiral" structures with no floor.

Escrowed shares

Shares held in escrow (cannot be sold) for a defined period. Required by ASX for certain seed and promoter shares post-IPO — typically 12–24 months. Watch escrow expiry dates.

Fully diluted SOI

SOI + all unexercised options + performance rights + convertible note conversion shares. **This is the number that matters.** Use fully diluted MC as your real valuation metric, not headline MC.

A company with 200m SOI at 10c (MC = \$20m) but 100m unexercised options at 5c, 50m performance rights, and a \$5m convertible at 8c is actually:

- $200\text{m} + 100\text{m} + 50\text{m} + 62.5\text{m} = 412.5\text{m}$ fully diluted
- Fully diluted MC at 10c = \$41.25m
- Real ownership share is half what the headline suggests

The mechanical job is to compute this number for every position. The interesting job is reading what the structure tells you.

Reading the share register as a signal

Pull the **Top 20 shareholders** list (in the annual report, the half-yearly report, and quarterly Computershare/NSX updates).

What concentration tells you

- **Top 20 holding %**

- **Above 80%:** very concentrated. Concentration can be a strength (aligned long-term holders) or a weakness (one big seller can crater the SP). Read it in context.
- **60-80%:** concentrated. Look at who the holders are.
- **40-60%:** broadly held. Liquidity is usually adequate. Lower risk of single-holder selling crashes.
- **Below 40%:** retail-heavy. Volatile. Subject to sentiment swings and sector rotation.

Who is on the register

The names on the register tell you what kind of money has done DD on the company. Categorise as:

Specialist resources funds. Institutional money that exclusively or primarily invests in mining. Names you'll see on Australian small-cap registers include Acadian, Regal, Tribeca, Lowell, Paradise, L1 Capital, Collins St, RCF, Resource Capital Funds, Sprott. Specialist money on the register is a strong signal — these funds have analysts who do real DD on resource quality, capital structure, and management. They don't take small positions in companies they haven't validated.

Generalist institutional money. Pension funds, sovereign wealth funds, generalist long-only funds. Less rigorous on resource-specific factors but signals that the company has reached scale. Generalist money tends to follow specialist money into a name; rarely leads.

Strategic investors / corporates. A major mining company taking a 5-15% stake (e.g., a tier-1 producer buying into a junior in their commodity). The strongest possible signal — a strategic has more information than anyone, and a position taken at premium to market is a near-direct vote of asset quality.

Trading desks. Macquarie, UBS, JP Morgan, Goldman Sachs nominees. These are usually short-term holders, often arbitrage or facilitation positions. Their presence isn't a signal of conviction.

Nominee accounts. HSBC Nominees, Citicorp Nominees, JP Morgan Nominees, BNP Paribas Nominees — these obscure real ownership. Substantial holder notices (below) are how you find out who's really behind them.

Vendor / promoter shares. Original founders or vendors who took shares as consideration for the asset. Often in escrow at IPO; watch for escrow expiry dates. Vendor shares are a neutral signal — sometimes vendors are aligned long-term holders, sometimes they're sellers waiting for escrow to lift.

Director holdings — the alignment signal

From the annual report, directors' shareholdings are disclosed. What to look for:

- **Meaningful absolute holding.** Directors with millions of shares acquired with their own money are aligned. Directors with grant-only holdings (i.e., shares received as compensation, not bought) are less aligned.
- **Recent on-market buying.** Disclosed via Appendix 3Y filings. Directors buying shares on-market at current prices, with their own cash, is the strongest single alignment signal in the register.
- **Director participation in placements is *not* the same as on-market buying.** Directors putting their own money into a placement at the placement price (often a discount to market) is positive — but it's not as strong as buying on-market at full price. Many directors participate in placements only to maintain their percentage holding; that's a defensive posture, not a conviction position.
- **Director selling.** Disclosed via Appendix 3Y. Directors selling material amounts on-market is a tell, especially if it occurs after a positive announcement. Always check the explanation.

Substantial holder notices — the dynamic signal

Required disclosure when a holder reaches 5% of SOI, and again on every 1% change up or down (Chapter 6C of Corporations Act).

These are gold for tracking what informed money is doing in real time:

- **Initial substantial holder notices** (a fund crossing 5%) tell you who has just built a meaningful position. Match the name against your specialist/generalist/strategic categorisation above.
- **Subsequent change notices** show you whether the holder is accumulating further or trimming. A specialist fund that goes from 5% to 7% over six months is increasing conviction. The same fund going from 7% to 5% is reducing conviction.
- **Ceasing to be a substantial holder** notices (dropping below 5%) tell you when a meaningful holder has exited. Worth checking against the SP action that followed.

The 3% creep provision

A specific mechanism worth knowing about because it shows up periodically in active register stocks:

Holders above 19.9% of a company are normally subject to takeover provisions and can't accumulate freely. But the "**creep**" exception under Section 611 of the Corporations Act allows a holder to acquire up to **3% of the company every 6 months** without triggering takeover provisions, provided they've been above 19.9% for at least the prior 6 months.

So a holder at, say, 22% can buy up to 25% in the next 6 months under the creep provision, then up to 28% in the 6 months after that, and so on.

Why this matters:

- A persistent ~3% accumulation pattern by a major holder over consecutive 6-month windows is the creep provision in action
- This is a signal of intent — the holder is positioning for influence or potential control
- It often precedes a formal takeover offer, a board push, or a strategic move
- The CAY (Cazaly Resources) creep accumulation pattern observed in some periods is a worked example of how the mechanic plays out

Track the substantial holder notices on stocks with one or more major holders above 19.9%. The 3% creep is one of the cleanest tells of strategic intent in the entire ASX register.

How juniors raise capital — and what each method signals

The mechanism matters less than what the choice of mechanism tells you about the state of the company and the demand for its paper.

Placement (most common)

Issue of shares to "sophisticated investors" under section 708 of the Corporations Act. No prospectus required. No shareholder approval required if within 7.1 capacity (15%) or 7.1A capacity (additional 10% for eligible smaller companies).

Typical placement structure:

- 10-25% discount to recent VWAP
- Often includes free options ("sweeteners")
- Funds typically arrive within days of completion
- Brokers take 4-6% fee

The signal in placement pricing. A placement priced at a steep discount with free options attached is a placement where demand was weak — the broker had to discount and sweeten to fill the book. A placement priced at a small discount with no options is a placement where demand was strong. A placement priced **at a premium to VWAP** is the rarest and strongest signal — demand exceeded supply at full pricing.

The premium-to-VWAP placement — the strongest demand signal

Premium-to-VWAP placements are unusual because the natural state of a placement is a discount (broker needs to incentivise the placee to buy in size). When a placement gets done at a premium, it tells you:

- A specialist or strategic investor has done substantial DD and concluded the asset is worth more than the current SP
- The investor was prepared to compete for allocation rather than negotiate for a discount
- Other investors in the placement got the same premium pricing — the broker didn't have to discount to fill the book
- Often associated with cornerstone arrangements where the lead investor has effectively pre-committed at a premium and the rest of the book follows

The Galan Lithium (GLN) story is the cleanest worked example available in 2025–2026. Two placements in five months:

- **August 2025:** Clean Elements Fund (specialist lithium fund), \$20m at A\$0.11 per share, 21% premium to last close, after a 77-day technical and legal DD period
- **January 2026:** \$40m institutional placement plus \$1m director placement at A\$0.41 per share, 2% premium to 5-day VWAP and premium to 10 and 15-day VWAPs, Clean Elements participating again, directors committing own cash, Canaccord Genuity as lead manager

The same cornerstone investor paid 3.7x more per share over five months at premium pricing on both occasions. That's a cap-structure signal you don't get from drill results, broker reports, or chart analysis. It's a specialist fund with deeper information than the public market putting real money behind a specific operational view.

What to read in any placement announcement

Every placement announcement lets you extract these signals:

1. **Pricing vs VWAP.** Discount or premium? Compared to which VWAP window (5-day, 10-day, 15-day)?
2. **Discount or premium to last close?** Standard placements run 13–20% below last close even when they're at premium to VWAP, because of the typical SP run before placement.
3. **Who's the lead manager?** Tier-1 brokers (Canaccord Genuity, Bell Potter, Petra Capital, Morgans, Macquarie, Euroz Hartleys) carry more weight than unknown brokers.

4. **Which institutions participated, and were any of them existing holders?** Re-participation by a known cornerstone is a stronger signal than first-time generalist participation.
5. **Did directors participate, and at what scale?** Director participation with own cash, especially material amounts, is positive.
6. **What attaching options (if any)?** Heavy options attached suggests weak demand. No options or modest options at high strike suggests strong demand.
7. **What's the use of funds?** Specific milestones (Phase 1 expansion, infill drilling for resource update) are positive. Generic "working capital and general corporate purposes" is a flag — usually means the company doesn't have specific milestones to fund.

ASX Listing Rule 7.1 and 7.1A

Companies can issue up to 15% of their SOI without shareholder approval in any rolling 12-month period (Rule 7.1). Smaller companies (broadly, MC below ~\$300m and not in S&P/ASX 300) can seek shareholder approval at AGM for an additional 10% capacity, taking total to 25% in 12 months (Rule 7.1A).

If you see 7.1A approval on the AGM agenda, the company is signalling "we expect to raise capital aggressively over the next 12 months." It's not necessarily bad — it might just be prudent capacity for known capex needs — but it tells you to expect placements within the year.

SPP (Share Purchase Plan)

A pro-rata offer to existing retail shareholders. Capped at \$30,000 per holder. Often attached to a placement (institutions get the placement, retail gets the SPP at the same price) to soften the optics of dilution.

SPP take-up rates are a tell. Heavy retail take-up (oversubscribed SPPs) suggests retail enthusiasm; light take-up suggests retail isn't interested at the offer price.

Rights Issue

A pro-rata offer to all shareholders.

- **Renounceable:** rights can be sold on-market if you don't want to participate
- **Non-renounceable:** rights expire worthless if not exercised — forces you to participate or accept dilution

Rights issues are usually deeply discounted (30%+ below SP) to ensure take-up. They're often a sign of difficulty raising capital through institutional placement at acceptable terms.

Convertible note / loan

Debt-style instrument that converts to equity. Used when the company can't easily place at acceptable terms. Often a sign of distress.

The danger structure is the **death spiral convertible**: a convertible note where the conversion price is set at a discount to the prevailing share price *at the time of conversion*. If the SP falls, the conversion price falls, more shares are issued for the same dollar amount of debt, the SP falls further, and so on. Avoid companies that have these structures in place.

Streaming / royalty financing

Sell future production at a discount, or sell a royalty on revenue, in exchange for upfront cash. Non-dilutive but encumbers the project. Wheaton Precious Metals, Franco-Nevada, Royal Gold are the major streamers.

The signal from streaming is mixed:

- **Positive read:** non-dilutive financing, often at attractive terms relative to equity dilution
- **Negative read:** the streamer takes a permanent slice of project economics; some forms of streaming are effectively perpetual

For developers, a streaming arrangement is often the right call — preserving equity in exchange for a fixed cost of capital. For producers, streaming is rarely necessary and is sometimes a sign that conventional debt or equity wasn't available.

Drill-for-equity contractor arrangements

A specific arrangement worth flagging because it's a green flag pattern often missed by retail.

The structure: a drilling contractor accepts payment in shares (or a mix of cash and shares) rather than full cash. The contractor effectively bets their own labour cost on the success of the program — they get paid more if the drilling produces results that lift the SP, and less if it doesn't.

Why it's a green flag:

- The contractor has done their own DD on the project and concluded the asset is worth taking a position in
- The company is demonstrating capital discipline by stretching cash through equity instead of burning it on contractor invoices
- The arrangement aligns the contractor's incentive with the shareholder's interest
- Contractor share-and-cash arrangements are most common at small junior explorers (e.g., TOR-style structures, PGO-style structures) where every dollar of cash matters

Caveats: small drill-for-equity arrangements (a few hundred thousand dollars worth) are positive signals. Large-scale ones can become problematic if the contractor needs to liquidate large parcels at times that don't suit the company.

Flow-through shares (Canadian projects)

If a company has a Canadian project (or an Australian-Canadian dual-listing), you may see **flow-through shares** in the capital structure.

The mechanic: under Canadian tax law, exploration expenditure on Canadian projects can be "flowed through" from the company to investors as a tax deduction. Investors in flow-through shares effectively pay a premium for the shares because they get a tax benefit equivalent to a deduction equal to the share purchase price. The company gets the cash to spend on Canadian exploration; the investor gets the tax benefit.

Why this matters for ASX investors:

- Flow-through shares are typically issued at a premium to standard placement prices
- They lock the funds to specific Canadian exploration expenditure (regulated, must be spent within a defined window)
- The premium pricing improves the per-share dilution math vs equivalent standard placements
- SMM (Summit Minerals) and similar Canadian-asset ASX juniors have used flow-through structures to fund exploration without the dilution penalty of equivalent standard placements

The flow-through premium is a real cap-structure positive when used appropriately — it's effectively non-Australian-investor money paying a premium for tax benefits the Australian holder doesn't capture, with the entire value of the deal accruing to the project rather than to dilution.

Cash position and runway — the Section 8 calculation

This is the single most useful 2-minute calculation in mining. Every quarterly Appendix 5B has the data; most retail don't read it.

The Section 8 method

Open the Appendix 5B for the most recent quarter. Look at:

- **Section 1.2 — Cash and cash equivalents at quarter end.** This is current cash on hand.
- **Section 1.5 — Net cash from / (used in) operating activities.** This is operating cash burn for the quarter (negative for nearly all explorers and developers).
- **Section 2.6 — Net cash from / (used in) investing activities.** This is exploration spend, equipment purchases, and project capex (negative for nearly all explorers and developers).
- **Section 8 — Estimated cash available for future operating activities.** This explicitly states the company's calculation of how many quarters of runway they have at current burn rates.

Or do it yourself:

Quarterly burn = absolute value of (Section 1.5 + Section 2.6)

Runway in quarters = Section 1.2 / Quarterly burn

What runway tells you

- **Below 2 quarters:** capital raise is imminent. Markets price this in well before the announcement. SP weakness in this regime is partly anticipatory dilution.
- **2-4 quarters:** capital raise is on the horizon. Watch for the typical pattern: positive announcement → SP run → trading halt → placement.
- **4-8 quarters:** company has flexibility to time a raise. They can wait for a positive catalyst to lift the SP before raising.
- **Above 8 quarters:** company is well-funded relative to current burn. Unlikely to need to raise in the near term unless burn rate accelerates (e.g., entering DFS or construction).

Why this beats almost any other dilution prediction

Every other framework for predicting capital raises is fuzzier. Some companies telegraph in their corporate presentations; some don't. Some companies are opportunistic; some are predictable. But the cash-and-burn calculation is mechanical — once you have it, the timing window for the next raise is constrained to a 1-2 quarter range.

The same calculation applied to your entire mining portfolio every quarter tells you which positions are about to face dilution and which aren't. Position sizing should reflect that.

What's not in Section 8

A few caveats:

- **Operating activities don't include construction capex.** A developer entering construction will have a step-change in cash burn that won't be reflected in the prior quarter's Section 8 number. Read the company's construction guidance separately.
 - **Some companies use "non-recurring" framings** to make recurring spend look temporary. If a company says quarterly G&A is "elevated due to one-off costs" for three consecutive quarters, the costs aren't one-off.
 - **Drawdown of facilities and prepayments** can extend cash without a placement. Authium-style prepayments (the GLN US\$6m facility example) and similar arrangements can defer dilution.
-

Capital raises as entry timing — the contrarian angle

Most retail treat capital raises as bad news. The discounted SPP, the dilution, the SP weakness around the announcement — all feel negative. But for the contrarian investor, capital raise windows are often the best entry points in the entire mining cycle.

The pattern

A junior runs a capital raise. Process:

1. Trading halt called
2. Placement is bookbuilt with institutions over 24-48 hours
3. Trading halt lifts; placement at, say, 15% discount to last close
4. SPP follows for retail at the same price
5. SP trades around the placement price for several weeks as placement stock unlocks and gets sold by short-term holders
6. Eventually the placement supply gets absorbed and the SP recovers

For a contrarian investor, **the period between placement and supply absorption is often the best entry point**. The dilution has happened, the discounted shares are in weak hands rotating out, the structural overhang is visible to everyone, and there's no near-term need for further dilution. You're buying after the bad news has already been priced in.

When this works and when it doesn't

The pattern works when:

- The capital raise was for a clear milestone-driven purpose
- The company is fundamentally sound (FA Story still holds)
- The placement was supported by quality money (specialist funds, strategics, directors)
- The macro for the commodity hasn't deteriorated

The pattern fails when:

- The capital raise was distress-driven (avoiding insolvency rather than funding milestones)
- The placement was supported only by retail-style investors and short-term funds
- The macro is deteriorating (more dilution likely to come)
- The fundamentals were the reason for the raise (project doesn't work as advertised)

The signal in placement quality

A key tell that the entry-timing trade is set up properly: the placement was at premium to VWAP and supported by specialist money. Those are the placements that get absorbed quickly because the holders aren't selling at a loss. The SP recovers within weeks, not months.

A placement at heavy discount with retail-heavy support and free attaching options is a placement where the supply pressure persists for months. Don't enter on the assumption that supply will absorb quickly.

How to use this in your process

Three rules:

1. **Don't be afraid of capital raise weakness on quality companies.** It's often the best entry window in the cycle.
2. **Don't buy capital raise weakness on weak companies.** The dilution is a symptom, not an opportunity.
3. **Wait until the trading halt has lifted and the SPP period is closing before deciding.** That's typically 2–4 weeks after the placement. By then, the supply dynamics are visible.

Patterns of capital raising by stage

Stage	Typical raise pattern	What to watch
1 — Concept	\$1–3m placements every 6–12 months	Dilution accumulating without milestones

Stage	Typical raise pattern	What to watch
2 — Drilling	\$3-8m to fund drilling, often before each program	Pre-drilling raises locked in adequate cash
3 — Discovery	Larger raise (\$10-30m) immediately after the discovery hole, at premium to pre-discovery price	Quality of cornerstone support
4 — Resource definition	\$5-15m to fund infill drilling and metallurgy	Whether the resource update justifies the dilution
5-6 — Studies	\$5-20m to fund PFS/DFS work; SP often weak so dilution is heavy	Premium-to-VWAP raises here are exceptionally strong signals
7 — DFS complete	The big one: \$50-500m+ project finance package (debt + equity + offtake/streaming)	Equity component vs debt vs streaming mix
8-9 — Construction	Cost overrun raises, often deeply dilutive at depressed SPs	Watch for staged drawdown facilities
10 — Production	Reserve replacement exploration funded from cash flow; minimal further dilution if managed well	Any equity raise at this stage is a flag

The capital raise immediately after a positive announcement is a recurring pattern. The good news pumps the SP, the placement gets done at the higher price, dilution is "less bad" than it would have been. Watch for ASX trading halts within 1-3 days of major positive announcements.

The signals from how a company raises capital — consolidated

Strong positive signals:

- Premium-to-VWAP pricing on a placement
- Specialist resources fund as cornerstone, especially repeat participation at higher prices
- Director participation with own cash on top of placement participation
- Strategic investor (major mining company) at premium to market
- Funds raised align with stated next 12-18 month milestones
- Tight pricing (small discount), no free attaching options
- Streaming/royalty financing for project finance instead of equity dilution
- Drill-for-equity contractor arrangements at small explorers
- Flow-through shares for Canadian exploration (premium pricing, locked-in spend)

Yellow flags:

- Heavy free options attached (signals demand was weak)
- Placement at significant discount with no specialist cornerstone
- Vague use-of-funds (general working capital)
- 7.1A capacity sought when there's no obvious milestone-driven need

Red flags:

- "Death spiral" convertibles with floating conversion prices and no floor
 - Repeated raises with no stated milestone progress between them
 - Insider participation suspiciously timed (raise immediately *before* a major positive announcement)
 - Capital raised vastly exceeds disclosed funding requirements (overcapitalisation = SP overhang)
 - Placements consistently priced below previous placements (the value is being destroyed structurally)
-

Practical exercise

For every junior you hold or are considering, build out:

1. Current SOI
2. Total options + performance rights = pending dilution
3. Fully diluted SOI and fully diluted MC
4. Top 20 holding % and identifiable institutions on register, categorised (specialist / generalist / strategic / trading / nominee / vendor)
5. Director holdings and recent on-market buying/selling activity (Appendix 3Y filings)
6. Substantial holder change pattern over the past 12 months — who's accumulating, who's exiting?
7. Section 8 cash burn calculation — current runway in quarters
8. Last 3 capital raises: dates, amounts, prices vs VWAP at the time, dilution %, who participated
9. **Were any of those raises at premium to VWAP? Who were the cornerstone investors? Have they participated again at higher prices?**

If you can't reproduce this from memory for any holding, you don't actually understand what you own. The cap structure is half the FA Story.

What I'm uncertain about

- ASX Listing Rules 7.1 and 7.1A have been amended over time; verify current % thresholds and eligibility criteria on the ASX Listing Rules site if you need precision.
- The retail SPP cap was increased from \$15k to \$30k some years ago; \$30k is the current figure but worth confirming.
- Substantial holder threshold is 5% under Corporations Act Chapter 6C; this is stable but technical disclosure timing rules have been refined periodically.
- The 3% creep provision technical mechanics around the qualifying period (the 6-month "above 19.9%" requirement) have edge cases worth verifying for any specific situation. The framework is broadly correct but precision matters in active register stocks.
- Flow-through share mechanics are Canadian tax law specifics; verify with reference to Canadian Income Tax Act provisions if you're modelling a specific structure precisely. The framework above is the simplified view.
- "Specialist resources fund" categorisation is a judgment call. The named funds are well-known examples but individual fund mandates change over time. Verify each fund's current strategy if it's material to your thesis.

Module 7: Catalysts & The Catalyst Calendar

Why this matters

Mining stocks don't move on time. They move on **events**. Between events, they drift, decay, and bleed sentiment. Map the events ahead of time and you know when to be in, when to be out, and what to expect.

The retail mistake is owning a stock through the dead time between catalysts. The professional approach is to enter ahead of high-conviction catalysts and exit after them — or to enter the orphan period and wait for the next catalyst cycle to lift the SP.

What is a catalyst?

A **catalyst** is a discrete event that changes the market's information set about a company's value or risk. It must be:

1. **Identifiable** — you know it's coming
2. **Time-bound** — there's a window when it'll arrive
3. **Material** — capable of moving the SP meaningfully

Vague things like "exploration upside" or "macro tailwinds" are not catalysts. They are themes.

The catalyst hierarchy by stage

Stage 1–2 (Concept → Drilling)

- Tenement grants
- Geophysical survey results

- Drilling commencement
- First-hole results
- "Visual indicators" announcements (often fluff but can move SP)

Catalyst potency: low to moderate. Drill commencement runs are predictable; visual results are unreliable as a fundamental signal but can drive 20–50% SP moves.

ASX framing: **BSR** post-IPO 2025 is a classic Stage 2 catalyst pattern — drill commencement, first hole, assay results in sequence over 3–6 months. **TOR** runs the same playbook with the additional drill-for-equity contractor structure as a green flag.

Stage 3 (Discovery)

- **Discovery hole assays** — the most asymmetric catalyst in mining
- Step-out hole assays confirming continuity
- Geophysics extending the prospective footprint

Catalyst potency: very high. A genuine discovery hole can 5–10x a junior in days.

ASX framing: **PC2** through 2025 — successive high-grade Au intercepts, each batch market-moving, step-out drilling extending the footprint. The Stage 3 catalyst flow is the most concentrated information rate in mining.

Stage 4 (Resource Definition)

- **Maiden JORC resource** — major re-rate event
- Resource upgrades (Inferred → Indicated)
- Metallurgical testwork results

Catalyst potency: high for maiden, declining for subsequent upgrades.

ASX framing: **PC2's** maiden JORC resource (planned mid-2026) is the next major catalyst. **MI6** has been delivering resource upgrades and metallurgical results across its portfolio.

Stage 5–6 (Scoping → PFS)

- Scoping study results
- PFS results
- Permitting application lodgement
- Strategic investor announcement
- Offtake MOUs

Catalyst potency: mixed. Study results often disappoint vs hyped expectations (capex sticker shock). Offtake MOUs are usually overhyped relative to their actual binding force — see Module 4 on reading the body of announcements vs the headline.

Stage 7 (DFS)

- DFS results
- **Final Investment Decision (FID)**
- Permitting approvals
- Financing milestones (mandate signed, term sheet, drawdown)
- Binding offtakes

Catalyst potency: high. This is where the orphan period ends if the project is real.

Stage 8 (Construction)

- Ground-breaking
- Major equipment orders / arrivals
- First ore on ROM pad
- Construction milestones (mill installation, tailings dam complete, etc.)
- Schedule/budget updates (often negative)

Catalyst potency: moderate. Each milestone is a small re-rate. Cost overrun announcements are heavily punished.

ASX framing: GLN delivered the rare "construction completed on time and on budget" milestone on 31 March 2026. The default expectation should be the opposite — 30-50% capex blowouts and schedule slippage are the norm — so the absence of bad news at this stage is itself information.

Stage 9 (Commissioning)

- First concentrate / doré / cathode produced
- First sale / first revenue
- First positive cash flow quarter
- Achievement of nameplate capacity

Catalyst potency: very high. The second peak of the Lassonde curve builds here.

ASX framing: GLN's Q2 2026 first lithium chloride concentrate production is a company-defining milestone. The Stage 9 transition is when the market starts pricing the asset on cash flow rather than NPV — a different valuation regime entirely.

Stage 10 (Production)

- **Quarterly activities reports** (the heartbeat of producer SPs)
- Production guidance updates
- AISC reporting
- Reserve updates
- Dividend announcements
- Exploration results around mine
- M&A activity

Catalyst potency: moderate per event, but cumulative and consistent.

ASX framing: ELV through 2025–2026 — quarterly production and AISC become the core read on the stock. Macro overlay (Module 9) drives most of the SP movement; quarterly delivery either confirms or undermines the thesis the macro is implying.

Recurring scheduled catalysts (every ASX-listed company)

Build these into your calendar for every position:

Quarterly Activities Report (Appendix 5B)

Due dates: within 1 month of quarter end:

- Q1 (Jan–Mar): due 30 April
- Q2 (Apr–Jun): due 31 July
- Q3 (Jul–Sep): due 31 October
- Q4 (Oct–Dec): due 31 January

Contains: cash position, cash flow, quarterly activities, project updates, exploration spend.

For producers: production volume, AISC, sales, hedging position. The most important regular catalyst for any producer.

The Section 8 read — every Appendix 5B has a section showing actual quarterly cash burn (operating + investing). Divide cash on hand by that burn rate to get quarters of runway. This single calculation predicts cap raise timing better than almost anything else (covered in Module 6).

Half-yearly and Full-year financial reports

- Half-year (HY): typically March (for Dec year-end) or September (for Jun year-end)
- Full-year: typically August (Jun year-end) or February (Dec year-end)

AGM

Usually November for Dec year-end companies. May for Jun year-end. Sets up:

- 7.1A capacity vote
- Director re-election
- Performance rights grants
- Often a forward strategy presentation

Annual Mineral Resource and Ore Reserve Statement

ASX Listing Rule 5.21 requires an annual update. Often released alongside the annual report. Track resource/reserve growth year-over-year — this is where reserve depletion vs replacement becomes visible.

How to build a catalyst calendar

For each stock, track in a single document or spreadsheet:

Date / Window	Event	Type	Conviction	Expected SP impact
2026-Q2	First lithium chloride production (GLN example)	One-off	Very high	Major re-rate if successful; major fall if delayed
2026-mid	Maiden JORC resource (PC2 example)	One-off	High	+50-100%
2026-Q4	Scoping/PFS study results	One-off	Medium	+20-40% (if positive)
Quarterly	Drill assay batches	Recurring	Variable	-10 to +30%
2026-10-31	Sept quarterly	Scheduled	Low base case	Cash position read

Update monthly as company guidance shifts. Most companies telegraph their next 6 months in their quarterly investor presentations — read those.

Pre-event positioning patterns

The catalyst trade is well-understood and well-played. Recurring patterns:

Run-up into catalyst

SP often rises in the 2-8 weeks before a known catalyst as positioning builds. This is the cleanest part of the trade.

Sell-the-news

On the day of (or 1-2 days after) the catalyst, even good results often see SP fade as positioned holders take profit. Standard pattern especially for known/expected results.

Surprise upside

When results materially exceed the priced-in best case, you get a fresh leg up after the initial sell-the-news fade. Discovery holes far above expectations, DFS economics significantly above PFS, met recoveries 5%+ above pilot — these can re-rate.

Surprise downside

Cap-ex blowouts, met recovery shortfalls, permitting denials, delayed financing — punished hard, often -30% to -60% in a session. Recovery can take years.

How to use catalysts in position sizing

A position with a known catalyst within 8 weeks should be sized differently from a position with no catalyst for 12 months.

Crude framework:

- **Pre-catalyst position (≤ 8 weeks out):** can be sized to your full conviction allocation
- **Mid-cycle position (catalyst 3-6 months away):** half of full allocation; add into the catalyst window
- **Orphan period position (no catalyst within 6 months):** quarter allocation max; you're being paid for time risk via lower entry price, not catalyst risk

This naturally aligns position size with the rate at which information arrives.

What kills a catalyst trade

Slippage

Companies miss guided dates. A "Q2 2026 PFS" routinely becomes Q3, then Q4. Every quarter of slippage is a quarter of dilution risk and SP decay.

Pre-leak

Information sometimes leaks. The SP runs ahead of the announcement, and by the time results are public, the trade is over. Watch for unexplained volume / SP moves in the days before scheduled events. ASX speeding tickets ("Aware?" letters) are a tell.

Better-priced alternatives

Even a positive result can fail to move the SP if the market has already moved past it. A good DFS in a hot lithium market in 2022 moved stocks 50%+. The same DFS in 2025's bear market barely moved the SP.

Macro override

A perfect drill result during a commodity crash can be ignored entirely. Sentiment and macro determine whether the catalyst actually translates to SP movement.

When the absence of a catalyst is itself the catalyst

Some companies announce nothing because there is nothing to announce. This is structurally different from a company in the orphan period waiting for a known event. Tells:

- **Quarterly activities reports that copy-paste content from previous quarters.** Phrases like "drilling continues at site X" or "metallurgical testwork progressing" recurring verbatim across multiple quarters indicate the project is dormant. Change the date and the announcement could be from 18 months ago.
- **No drill rig on site for an extended period** despite quarterly references implying work is ongoing.
- **Cash burn that doesn't match the activities described** — if the quarterly activity report describes active drilling but the Section 6 spend on exploration is below what a rig program would cost, the activity is overstated.
- **Management remuneration high relative to actual progress** — covered in Module 6/8.

The dormant company is the worst trade in mining. The thesis dies of slow attrition rather than a single bad event. By the time you realise nothing has happened in 18 months, the SP has bled half its value and the cap structure has dilated through funding rounds.

Practical exercise

For every position in your portfolio:

1. List the next 3 catalysts in chronological order
2. Estimate the date window for each
3. Assign a conviction level (high/medium/low)
4. Estimate plausible SP impact ranges (positive and negative)
5. Identify which catalyst you're being paid to wait for
6. **Re-read the last 4 quarterly activities reports — does each one describe substantively new progress, or are phrases recurring verbatim?**

If you can't list 3 catalysts within the next 12 months for a stock you hold, you're holding a story with no scheduled events to validate or invalidate it. That's not investing — that's a time-decay trade against your own capital.

Cross-reference: ASX continuous disclosure

Under ASX Listing Rule 3.1, listed companies must immediately disclose information that a reasonable person would expect to materially affect the SP, subject to certain carve-outs (incomplete proposals, confidential negotiations, etc.).

This means catalysts cannot be "saved up" for AGM season. When something material happens, it must be announced. The exceptions are where companies wedge new "material" news into already-scheduled releases (like quarterlies) to make the news look less significant.

Watch for material content buried in the body of quarterlies that wasn't pre-announced — that's a deliberate framing choice. (Module 4 covers reading announcement bodies vs headlines as a discipline.)

What I'm uncertain about

- Specific quarterly due dates can be extended in unusual circumstances (e.g., COVID-related extensions in 2020). Default to the dates above.
- The threshold for what triggers continuous disclosure is a judgment call and ASIC has periodically issued enforcement actions where they disagreed with companies' interpretations. Recent specifics worth verifying if you need precision.
- Catalyst potency varies enormously with the broader macro cycle. The same announcement that drives a 50% re-rate in a bull market may drive a 10% bounce in a bear market. Module 9 covers how to read the cycle backdrop.

Module 8: Red Flags, Green Flags & Sneaky Tactics

Why this matters

Junior mining is the most retail-hostile sector on the ASX. The structures are complex, the disclosure is technical, the time horizons are long, and the promoters are professional. Most of what's done is technically legal. Some of it is genuinely fraudulent. All of it is designed to make money for the company, the promoters, the directors, and the broker pipeline — sometimes alongside shareholders, sometimes at their expense.

Your job is to learn the patterns. Once you can spot them, you can quickly disqualify 80% of the junior universe and concentrate work on the 20% worth analysing.

This module rewrites the original red-flag list around a different organising principle. The original was a checklist of items. The rewrite is a framework of **behavioural patterns** built around a single meta-question: *why does this company exist, and who benefits from its current structure?* Once you can answer that question for any junior, the specific flags become consequences of the answer rather than a list to memorise.

The meta-question: why does this company exist?

Every ASX-listed junior exists for a reason. The reason is not always "to develop a profitable mining project for shareholders." Many companies exist because:

- The shell was acquired by promoters as a pre-listed vehicle for a story they wanted to tell
- The directors are receiving meaningful remuneration regardless of project progress
- The capital structure provides ongoing liquidity to founders and seed investors
- The company is a regulatory shelter for prior-stage investments that need exit liquidity
- The company exists to recycle capital through repeat raises with each cycle's hot commodity

These motivations are not necessarily fraudulent. Some are explicitly disclosed in the prospectus. But they shape every decision the company makes about announcements, capital raises, and disclosure framing — and **none of them are aligned with a long-term shareholder buying at current prices.**

The single most useful question to ask of any junior is:

“If shareholders fully understood what this company is and isn't, would the share price be higher or lower?”

If the answer is **higher**, you've found a fundamental-price gap candidate (Module 10). The information disadvantage runs in your favour — you know more than the average market participant.

If the answer is **lower**, you've found a junior whose current price depends on shareholders not paying attention. The information disadvantage runs against you — the company is selling a story that doesn't survive scrutiny.

Most retail mistakes in junior mining come from owning the second category while believing it's the first. Every behavioural pattern below is a way to identify which category you're actually looking at.

Behavioural pattern 1: The dormant-company tell

This is the most expensive pattern in retail mining because it kills capital through slow attrition rather than visible loss. Years can pass with nothing happening; the SP bleeds 5-10% per quarter on dilution and decay; by the time the shareholder realises, the cumulative loss is 50-70%.

What it looks like

The company files quarterly activities reports that read like prior quarters' reports. The same project descriptions recur with minor wording variations. The same forward-looking statements ("further drilling to be conducted in coming quarters", "metallurgical testwork progressing", "permitting discussions ongoing", "the asset remains in our secure facility") appear verbatim across multiple quarters.

A worked example pattern: a company holds a uranium concentrate inventory in a secure facility while waiting for prices to recover. Quarter after quarter, the announcement reads "the [material]

remains in our secure storage facility, awaiting market conditions for sale." Change the date and the announcement could be from 18 months ago.

Why this pattern persists

The company has no material new activity. Genuine progress would generate dedicated announcements (drill commencement, assay results, study commencement, study results, permitting milestones). Copy-paste quarterly reporting is a fallback when the activity to report doesn't exist.

The company keeps the lights on because the shell has value. Directors continue to be paid. The cap structure can be tapped for further raises when needed. The story can be relabelled if the commodity goes out of favour. None of these require actual project progress.

How to detect it

- **Re-read the last 4-6 quarterly activities reports side by side.** Open them in tabs. If the same paragraphs recur across multiple quarters, the company is dormant.
- **Look for drill rig activity.** A company with active drilling will name the rig contractor, the metres drilled, the campaign budget, the assay turnaround timeline. A company with no drilling will use vague language: "exploration activities are progressing", "drilling is being planned", "rig mobilisation is anticipated."
- **Check Section 6 of the Appendix 5B against described activities.** If the announcement claims active drilling but the exploration spend is below what a single-rig program costs (\$150-400k per quarter at minimum for an active program), the activity is overstated.
- **Compare director remuneration to project progress.** A board collecting \$300k+ each in a \$20m MC company that hasn't produced a substantive announcement in 12 months is a signal that the company exists primarily to pay the board.
- **Check ASX announcements page.** A 12-month period with only quarterly reports and AGM notices, no operational announcements, is dormant.

What this means for your portfolio

The dormant company is the worst trade in mining because it dies slowly. There's no single bad announcement to react to. The thesis dies of attrition while the SP bleeds and the cap structure dilutes. By the time you've waited "just one more quarter" three times, two years have passed and the loss is structural.

Run the dormant-company check on every junior in your portfolio quarterly. If it triggers, exit before you've sunk another year's wait into a position that isn't going anywhere.

Behavioural pattern 2: The announcement-headline mismatch

The body of an announcement either supports its headline or quietly walks it away. Reading both — and comparing how different companies in the same district handle similar news — tells you what kind of operator you're dealing with.

What it looks like

Two companies drill comparable programs in the same week. Their announcements differ in revealing ways:

Company A (strong-results posture): dedicated ASX announcement titled "Significant high-grade intercepts confirm continuity of XYZ deposit." The body includes hole-by-hole assays, a plan view showing the new holes plotted relative to prior drilling, a long section, true widths disclosed, geological description, comparison to the deposit model, and a forward plan section identifying the next batch of holes and assays.

Company B (weak-results posture): a passing mention in their quarterly report — "drilling has been completed at site X, with results currently being interpreted." No dedicated announcement. No plan view. No hole-by-hole data. The forward plan is vague.

The two announcements are about similar work. The difference is what each company is willing to put a spotlight on. A is making the most of strong results. B is downplaying weak ones.

Why the pattern works

Announcement style is a high-fidelity signal of underlying results because it costs nothing to make. A company with strong results has every incentive to lead with them — bigger headlines, more diagrams, more data, more forward guidance. A company with weak results has every incentive to bury them — less prominent placement, less data, vague language, deferred follow-up.

The asymmetry between the two postures is much larger than the underlying difference in geological results. A 30m @ 2 g/t hole gets a dedicated announcement at one company and a one-line mention at another, depending on whether the hole was the best or the worst of the program.

How to detect it

- **For every operational announcement in your portfolio, do the body-vs-headline check from Module 4.** Does the body confirm and elaborate on the headline, or soften it?
- **Track which companies in your sector consistently make the most of their announcements vs which consistently bury content.** Over multiple announcements, the pattern becomes obvious.
- **Watch for the "no headline" tell.** When a company drills a program and reports it only inside a quarterly, that's typically because the results didn't justify a standalone announcement. The absence is itself information.
- **Cross-reference against neighbours.** When a company in a known mineral district publishes underwhelming results while neighbouring companies on adjacent tenements publish stronger ones, that's a signal about the specific tenement, not the district.

What this means for your portfolio

You can build a quality-of-disclosure ranking across your watchlist. Companies that consistently lead with strong, data-rich announcements are doing the work. Companies that consistently bury content are managing perception around results that don't speak for themselves. Position sizing should reflect the difference.

Behavioural pattern 3: The repeat-IPO cluster

Some promoters and management teams recur across multiple junior mining IPOs over time, often in clusters, often with similar capital structures, similar story arcs, and similar outcomes.

What it looks like

Three or more new ASX listings appear within a 12-month window with overlapping management, similar promoter networks, similar broker syndicates, comparable capital structures (heavy seed and promoter shares, generous escrow lifting on similar timelines), and stories pitched to whatever commodity is currently in vogue.

The pattern is observable in real time during late-cycle phases. When uranium is hot, you'll see clusters of uranium-themed listings. When gold is hot, gold-themed listings. The same promoter networks reactivate around whatever has retail attention.

Why the pattern works for the promoters

The ASX listing process is the most expensive part of getting a small mining company to market. Once a shell exists with a story attached, the marginal cost of promoting it to retail is comparatively small. Promoters who've successfully listed multiple juniors have process efficiencies — broker relationships, IPO investor lists, marketing templates, story frameworks — that make repeated listings cheaper and more reliable than building a single project to a meaningful outcome.

For the promoter, the IPO itself is often the trade. Seed shares acquired at fractions of a cent are listed at \$0.20 with escrow expiring 12–24 months later. The successful exit is the escrow lift, not the project's eventual mine development.

How to detect it

- **Look up the last 5 directorships of every board member.** ASIC's company register will show. If a director has been on the board of 4+ junior mining IPOs in the past decade, all of which are now sub-\$10m MC, you're looking at a repeat-IPO operator.
- **Check the lead manager.** The same broker syndicates appear repeatedly across these clusters. Some brokers specialise in this kind of listing — that's a signal in itself.
- **Look at the seed investor list in the prospectus.** Many of these clusters share seed investors across companies. The same names recur across multiple prospectuses.
- **Compare prospectus wording across recent IPOs.** Story templates often recur with the commodity name swapped out.
- **Watch for unusual listing volume in a single commodity over a short window.** A cluster of 5 lithium IPOs in the same 6-month window of 2022 was a tell. The same applies to other commodities at different points in the cycle.

What this means for your portfolio

The repeat-IPO cluster isn't always fatal — sometimes good projects come through these structures. But the base rate of success is low, and the structural incentives are misaligned with shareholders. Default to scepticism. The single project that proves viable is the exception, not the rule. If you're going to participate, do so with the working assumption that the IPO itself is the promoter's trade and your edge is finding the rare cluster member where the underlying project is actually viable.

Behavioural pattern 4: The PFS-to-scoping downgrade

A specific tactic that signals the project's real economics didn't survive the discipline of a PFS framework.

What it looks like

A company has previously released a scoping study with strong headline numbers. The expected next step is a PFS. Instead, subsequent announcements label the work as a "scoping update", "revised scoping study", "preliminary economic update", or "concept study refresh". The PFS doesn't materialise. The labelling stays at scoping-grade indefinitely.

Why the pattern works for the company

Scoping studies are allowed to use Inferred resources in the production schedule and economic case. PFS studies are not (Module 5). When a project's headline scoping economics depend on Inferred tonnes carrying production years 6-15, the PFS-grade economics on the same deposit will be materially weaker — shorter mine life, lower NPV, less attractive financing case.

Rather than publish a PFS that walks back the scoping headlines, the company stays at scoping-grade and continues to refer to the original numbers. The market remembers the strong scoping numbers; the absence of a worse PFS is invisible.

The market mostly forgives the lack of progress. The company keeps fundability optionality without testing whether the project actually withstands PFS-grade scrutiny.

How to detect it

- **Track the labelling of consecutive economic studies on the same project.** Scoping → PFS is the normal progression. Scoping → Scoping update → Scoping refresh is the pattern.
- **Compare Inferred percentage in the resource that drives the economics.** If a meaningful percentage of contained metal in the production schedule is Inferred, the scoping economics depend on tonnes that won't survive into a PFS.
- **Check the timeline.** A company that completed scoping 24+ months ago with no PFS in sight is either dormant on study work or unable to PFS-validate the project economics.
- **Look at the management commentary.** Honest companies will explicitly say "we are progressing infill drilling to upgrade Inferred to Indicated for PFS purposes." Companies

running the downgrade tactic will avoid mentioning the resource confidence step-up.

What this means for your portfolio

A scoping study with strong headlines that hasn't progressed to PFS within 18 months should be re-modelled at PFS-allowable resource categories. If the project doesn't work at Indicated + Measured only, the scoping headlines are overstating the project's real value.

Behavioural pattern 5: The perpetual-foreign-estimate

The Module 2 cross-reference. A company acquires a project that already has a resource calculated under NI 43-101, SK-1300, SAMREC, or another non-JORC code, and discloses it on ASX under Listing Rule 5.12. The expected next step is a JORC re-statement, typically 6–18 months. The pattern is when this re-statement never quite arrives.

What it looks like

Quarter after quarter, the company refers to the foreign estimate in investor presentations, marketing material, and broker briefings. Each quarterly mentions that "JORC re-classification work is progressing" or "Competent Person review is ongoing." The timeline keeps slipping. 12 months pass, then 18, then 24. The foreign estimate stays the only resource statement on the project.

Why the pattern works for the company

The foreign estimate is usually the most favourable version of the deposit's tonnes and grade. If the JORC re-statement comes in materially smaller (which, per Module 2, is the typical pattern — historical estimates often lose 20–40% on JORC re-validation), the company has to publish that smaller number and the SP re-rates downward.

By staying perpetually mid-process, the company keeps the favourable foreign estimate as the live reference point in retail's mind without ever taking the SP hit of the JORC version.

How to detect it

- **Track the time elapsed since acquisition.** A foreign estimate that hasn't been JORC-restated 18+ months after acquisition warrants direct questioning.
- **Read the Competent Person's commentary in successive quarterlies.** Honest companies will quantify the work remaining (additional twin-hole drilling, density work, QA/QC review). Pattern-running companies will use vague language without specific milestones.
- **Compare the foreign estimate scale to the project's drill budget.** If the foreign estimate is substantial but the JORC re-statement work has had minimal funding allocated, the company isn't actually trying to deliver the JORC version.
- **Look for partial JORC restatements.** A company that JORC-restates only the highest-grade portion of the deposit (cherry-picking the part most likely to validate at scale) while leaving the rest as foreign estimate is running a sub-pattern of the same tactic.

What this means for your portfolio

A perpetual foreign estimate should be valued at zero in your model until JORC-restated. Any premium the SP carries on the foreign estimate is risk you're absorbing without compensation. If the company can't or won't deliver the JORC re-statement within a reasonable window, the most likely reason is that the JORC version doesn't support the foreign estimate's scale.

Behavioural pattern 6: The capital-raise-after-positive-news

Pattern: positive announcement → SP runs 30–80% on the day → trading halt the next morning → capital raise priced at small discount to the elevated SP. Common, legal, but tells you the capital raise was already prepared and the announcement was timed to optimise the placement price.

How to detect it

- **Time between positive announcement and trading halt.** If less than 5 trading days, the raise was almost certainly preparation-complete before the announcement.
- **Cross-reference against the prior 6 months of capital raises.** If the company has used this pattern repeatedly, it's structural.
- **Look at the cash position before the raise.** A company with 2–3 quarters of runway is in normal-pre-raise territory; a positive announcement timed to land just before the raise is the classic optimisation. A company with 6+ quarters of runway who still raises after positive news is using the news to capture higher pricing on dilution that wasn't urgent.

- **Watch the placement participants.** A genuine cornerstone-led raise after positive news has a specialist fund taking the lead at the elevated price. A weaker pattern has the raise dominated by short-term placement money rotating out within weeks.

What this means for your portfolio

The pattern itself isn't a hard red flag — it's standard operating procedure across the junior space. But the *quality* of the pattern matters. A company doing this consistently with quality cornerstone support, alongside genuine project milestones, is just optimising treasury management. A company doing this with poor placement support and no underlying project progress is using the announcement cycle to manage their cap structure rather than to build the project.

(Cross-reference: the entry-timing trade discussed in Module 6 — the period after the placement when the dilution has been absorbed and the SP recovers is often the best entry window for quality companies.)

Behavioural pattern 7: The transformational acquisition pivot

A company in a declining sector announces a "transformational" acquisition into a hot sector. Lithium 2022. Uranium 2024. Gold 2025. New tickers, new logos, new website — same management, new dilution. The acquired asset is usually low-quality (the good ones don't need rescue listings).

How to detect it

- **Check the company history.** Has the commodity focus changed multiple times in the past decade?
 - **Check who the asset was acquired from.** Often a private vendor associated with directors or seed investors — a related-party transaction dressed up as a strategic pivot.
 - **Check the consideration paid.** Often a large parcel of shares at a price that makes the vendors substantial holders of the post-acquisition company. Look at the dilution caused by the acquisition vs the value being acquired.
 - **Check the technical work supporting the acquired asset.** Often a thin technical report, sometimes a foreign estimate (cross-ref Pattern 5), often a prior-stage project that didn't progress under previous owners for reasons that haven't been addressed.
-

Behavioural pattern 8: Shell recycling

Companies that have been listed for 15+ years, changed names 3+ times, pivoted from gold to lithium to uranium to whatever's hot. Same management, same shell, new story. The shell was kept alive specifically because being already-listed is valuable for promoters, not because the underlying business has continuity.

How to detect it

- **Look at the company history on ASX.** Frequent name/ticker changes are the tell.
 - **Check directorship continuity.** If the same director has been on the board through 3 commodity pivots and 2 name changes, the shell has been passed through cycles.
 - **Look at IPO date vs current commodity focus.** A company listed in 2007 as a uranium explorer, that became a graphite play in 2014, a cobalt play in 2017, a lithium play in 2021, and is now a critical-minerals play in 2025, is a shell that's been re-themed 4 times.
-

Other red flags worth knowing

These are the patterns that don't quite warrant the deep-dive treatment above but matter as part of the disqualification process. Most are technically covered in other modules; flagged here for completeness.

Death spiral convertibles

Convertible notes where the conversion price floats at a discount to the prevailing SP, with no floor. If the SP falls, more shares are issued for the same dollar of debt, the SP falls further, and so on. (Module 6.) Avoid companies that have these structures in place.

Going concern emphasis in the auditor's report

The auditor flags "material uncertainty regarding the company's ability to continue as a going concern" — i.e., they may run out of money. Read the audit report on page 1 of the financial statements. Any going concern emphasis is a serious flag.

Auditor changes mid-year

Auditors don't usually walk away from clients without reason.

Heavy free options attached to placements

Cross-reference Module 6. Multiple free options per share suggest weak demand at the placement and persistent dilution overhang afterward.

Related-party tenement / asset deals

The company acquires an asset from an entity associated with a director, often for shares + cash. The vendor is a separate legal entity but the beneficial owner is a director or their associate. Sometimes legitimate; often grossly overpriced. Check the Annexure to the deal in the announcement.

"Up to" reporting in headlines

Cross-reference Module 4. "Up to 35 g/t Au returned from drilling" always means a single peak assay, not a weighted intercept.

Promotional language without backing data

"World-class", "company-making", "transformational", "tier-1 potential", "elephant in the room" — language that belongs in marketing decks, not technical announcements. Legitimate technical announcements describe what was found in measured terms.

Excessive director remuneration

Directors paid \$400k+ each in a \$20m MC company that hasn't drilled a hole in 12 months. Total board + KMP comp above 5% of MC for a non-producer is excessive.

Selective announcement timing

Bad news released on Friday afternoon, on a major news day, or after market close. Good news released first thing in the morning, sometimes simultaneous with an investor roadshow.

Buried material content in quarterlies

Material new information disclosed within a 30-page quarterly rather than as a standalone announcement. Less likely to attract market attention or regulatory scrutiny than a dedicated release.

"Non-binding MOU" headlines

Headlined as offtake or partnership. Buried in body: "non-binding", "subject to negotiation", "conditional on completion of feasibility study, financing, and regulatory approvals". Often nothing comes of it.

Performance rights vesting on share price targets

Performance rights that vest when SP reaches specified levels align insiders with hitting *short-term* SP marks, which can be done via promotion rather than fundamental delivery.

"Strategic review" language

"The Company is undertaking a strategic review of its asset portfolio." Often code for "we're trying to sell this and haven't found a buyer." Sometimes preludes a major impairment.

Hidden royalty stack

The project has multiple royalties already encumbering it: government royalty, vendor royalty from previous owner, streaming agreement, native title benefits. A 10% combined royalty burden destroys margin.

Foreign jurisdiction asset with onshore-only listing

The asset is in a high-risk jurisdiction but the company is ASX-listed because that's where retail capital is easiest. Sovereign risk, expropriation risk, and security risks may not be priced in.

ASX queries / speeding tickets

ASX issues "please explain" notices when SP or volume moves unusually. Frequent queries signal either market manipulation, leaks, or careless disclosure. Search for "ASX query", "Aware Letter", "Price and Volume Query" on the company's announcements page.

Resource update with shifted parameters

Cross-reference Module 2 — "no new drilling, just changed parameters" pattern. Resource grows year-over-year because the cut-off was lowered, not because new mineralisation was found.

Stockpile inclusion in mine plan

Cross-reference Module 5. Late-life stockpile processing folded into headline LOM, often economic only at favourable price assumptions, sometimes with recovery assumptions higher than realistic for weathered material.

Selective benchmarking

Comparing your project to globally famous deposits with very different geology, jurisdiction, or scale. Always read which deposits are being used as comparators and check whether the geology is genuinely analogous.

Backdated or revised guidance

Production guidance revised down repeatedly. Cost guidance revised up. Each revision presented as "due to one-off factors". After three revisions, the issues aren't one-off — they're structural.

What a green-flag company actually looks like

The inverse of the patterns above. A green-flag company exhibits multiple of these characteristics simultaneously:

On the cap structure (Module 6)

- Tight share register with specialist resources funds as identifiable holders
- Director on-market buying with own cash (Appendix 3Y)
- Strategic / cornerstone investor at premium-to-VWAP placements
- Disciplined capital raises tied to specific milestones (no vague "general working capital")
- Streaming/royalty financing instead of equity dilution at FID
- Drill-for-equity contractor arrangements at small explorers

On the project (Modules 2–5)

- Conservative assumptions in studies (commodity price at or below current spot, met recoveries supported by locked-cycle pilot tests, contingency 15%+ on capex, ramp profile of 12+ months)
- Resource and reserve growth year-over-year through actual drilling, not parameter changes
- Independent technical reports signed off by reputable consultants (SRK, AMC, Mining Plus, Snowden, CSA Global, Lycopodium, GR Engineering, Wood)
- Binding offtake with credible counterparty (not "non-binding MOU")

On management and disclosure

- Strong technical board with executives who've been through full Stage 1→10 cycles
- Stable senior management — same MD, CFO, Exploration Manager for 3+ years
- No-drama announcement style — technical announcements written by geologists, not marketers, with specific data, sober tone, full JORC Table 1
- Forward plan sections that identify next milestones with realistic timeframes
- No frequent ASX queries

On corporate behaviour

- Project progresses through Lasso stages on a defensible timeline

- Capital raises align with milestone delivery, not announcement-cycle optimisation
- Directors don't sell into positive news
- Quarterly reports describe substantively new progress quarter on quarter, not copy-paste content

A junior that exhibits 8+ of these characteristics simultaneously is rare. When you find one, the FA Story (Module 10) is much easier to write because the cap structure and the disclosure pattern are aligned with shareholder interests rather than working against them.

Sneaky tactics — the consolidated taxonomy

Beyond the major behavioural patterns above, smaller tactics are worth recognising:

1. **The "transformational acquisition" pivot** (covered above as Pattern 7)
2. **Reverse takeover (RTO) shell games** — live ASX shells sold to vendors who reverse-list private assets into them. Sometimes legitimate (faster than IPO) but often involve poor-quality assets that couldn't pass IPO due diligence.
3. **The 7.1A overhang** — companies seek 7.1A approval at AGM "for flexibility" then quietly use it months later. The capacity itself signals dilution coming.
4. **Capital raise straddled across positive announcements** (covered above as Pattern 6)
5. **Selective announcement timing** (Friday afternoons for bad news, Monday mornings for good)
6. **Buried material content in quarterlies**
7. **Non-binding MOU headlines**
8. **Performance rights vesting on SP targets**
9. **Resource update with shifted parameters** (Module 2 cross-ref)
10. **"Strategic review" language**
11. **Internal name changes hiding history** (Pattern 8)
12. **"Spin-out" of unwanted asset** — parent company spins off a non-core asset to existing shareholders as a separate listing. Often the spin-out has poor economics; the parent kept the good stuff.
13. **Paid promotion / sponsored research** — large portion of "research" coverage on small-cap miners is paid for by the company. Some research houses disclose this clearly; many bury it. Any "Buy" recommendation with a 12-month price target 3-5x current SP on a Stage 1-3 company with no revenue is almost certainly paid coverage.
14. **Nearology plays** — pegging ground next to a major's discovery and running on "in the same belt as [tier-1 deposit]" marketing. The geology is rarely the same — the major's geology team had pick of the best ground.

ASIC and ASX enforcement context

The line between aggressive promotion and securities fraud is jurisdictionally specific. ASX has historically been more permissive than the US SEC, though this is changing slowly.

What's worth knowing about the regulatory backdrop:

- **ASIC** (Australian Securities and Investments Commission) is the corporate regulator. They handle continuous disclosure breaches, misleading statements, market manipulation cases, and director duty breaches.
- **ASX** runs the listing rules and conducts initial review of disclosure issues, often via "speeding tickets" (Aware Letters / Price and Volume queries) before escalating to ASIC.
- **The continuous disclosure regime** (Listing Rule 3.1) requires immediate disclosure of price-sensitive information, with carve-outs for incomplete proposals and confidential negotiations. The threshold for what qualifies as price-sensitive is judgment-based and ASIC has periodically issued enforcement actions where they disagreed with companies' interpretations.
- **Resource and reserve reporting** has been the subject of multiple ASIC enforcement campaigns over the years, including specific actions around misleading exploration target disclosures, foreign estimate disclaimers, and inferred-resource economic studies. The pattern is periodic crackdowns followed by lighter touch periods.
- **Class actions** by shareholders are an alternative enforcement mechanism — the major plaintiff law firms (Slater and Gordon, Maurice Blackburn, Shine Lawyers) regularly bring continuous disclosure class actions against ASX-listed companies that have re-rated sharply downward after disclosure events.

For any specific borderline case, check the ASIC media releases page for recent enforcement actions involving similar conduct. This is the most reliable way to gauge current regulatory tolerance.

The 10-question disqualification checklist

Before doing any deeper work on a junior, run these 10 quick checks. **If 3+ red flags appear, move on.** The whole point of the checklist is efficiency — most of the junior universe should be disqualified before you spend serious analytical time on it.

1. ☐ Has the company changed names / pivoted commodities in the last 5 years?

2. ☐ Are there ASX queries / speeding tickets in the last 12 months?
3. ☐ Going concern emphasis in latest auditor's report?
4. ☐ Director / CFO turnover in the last 12 months?
5. ☐ Headline grades reported as "up to" rather than weighted intercepts?
6. ☐ Reliance on non-JORC (foreign or historical) estimates in marketing material that hasn't been re-validated within 18 months of acquisition?
7. ☐ Any related-party asset transactions in the last 24 months?
8. ☐ Capital raise within 5 trading days of major positive announcement, recurring as a pattern?
9. ☐ Director remuneration exceeds 5% of MC for a non-producer?
10. ☐ Recent quarterly activity reports use copy-paste language across multiple quarters?

A junior that triggers on 5+ of these is almost certainly not investable. A junior that triggers on 3–4 needs a much higher conviction reason to proceed. A junior that triggers on 0–2 is worth deeper analytical work.

Practical exercise

For every position currently in your portfolio, run the 10-question checklist. Be honest about the answers — if you can't immediately answer one, that's the question to research first.

For each position that triggers 3+ red flags:

1. Decide whether you have a specific reason that overrides the flags, or whether the position should be exited
2. If exiting, set a specific timeline (next 30 days; out by end of quarter)
3. If holding, document the override reason in your FA Story so you can re-test it next quarter

The discipline isn't to never own flagged stocks — sometimes the asymmetric upside justifies the risk. The discipline is to **know which positions have flags** and to size accordingly. A position with 4 red flags should be a 1% allocation, not a 5% one.

Final framing

The single best question to ask of any junior is the meta-question from the start of this module: *"If shareholders fully understood what this company is and isn't, would the share price be higher or lower?"*

Every behavioural pattern, every red flag, every sneaky tactic in this module is a way to identify cases where the answer is "lower" — where the current SP depends on shareholders not paying attention to specific things.

The practical edge in junior mining isn't picking winners. It's avoiding losers. The 80% of the universe you disqualify with the framework above is where most retail capital goes to die. The 20% you spend real analytical time on is where the work pays off.

If you can do the disqualification quickly and focus your hours on the smaller set of genuine candidates, you'll be ahead of almost every retail investor in the sector.

What I'm uncertain about

- ASIC and ASX enforcement intensity varies over time. Some periods see active crackdowns on misleading resource statements; others see lighter touch. Check ASIC media releases for recent examples if you're researching a borderline case.
- Specific consultant reputations shift over time. The list of "reputable independents" earlier in the module is current as I understand it but worth verifying recent project track records.
- The line between aggressive promotion and securities fraud is jurisdictionally specific. ASX has historically been more permissive than the US SEC, though this is changing slowly.
- The behavioural patterns above are observational generalisations from market history. They don't apply uniformly — some companies that exhibit elements of these patterns are still legitimate. The framework is a screening tool, not a categorical rule. Use it to allocate analytical time, not to make final investment decisions in isolation.
- Repeat-IPO clusters and shell recycling are particularly subjective to identify. The framework above is conservative — when in doubt, treat the situation as a flag and require a higher conviction reason to proceed.

Module 9: Macro Overlay — Commodity Cycles & The ASX Small-Cap Resource Cycle

Why this matters

The same project gets a 5x valuation in a bull cycle and a 0.2x valuation in a bear cycle.

Commodity cycles dominate everything else. A great project at the wrong point in the cycle will lose you money. A mediocre project at the right point in the cycle can multi-bag.

Most retail focus only on company-specific factors. The professionals build their thesis around the cycle first, then pick the best vehicles within it.

The two cycles you need to understand

1. The commodity cycle (10–15 years)

The underlying physical commodity moves through long cycles of:

- Underinvestment → supply tightness → price spike → over-investment → oversupply → price crash → underinvestment

These are driven by mine development lead times (5–15 years from discovery to production) and capital allocation that lags price signals by years.

2. The ASX small-cap resource cycle (3–7 years)

Driven by retail capital flows, broker IPO pipelines, and sector sentiment. Tends to lag the commodity cycle by 12–24 months at the start of a bull, then leads at the top by 6–12 months.

When both cycles align bullishly, juniors print money for everyone. When they diverge, only the best projects survive.

Why mining is so cyclical

Demand inelasticity in the short run

Industrial users can't quickly substitute most metals. Copper users need copper. Lithium users need lithium. So short-run demand barely responds to price.

Supply inelasticity in the short run

You can't build a copper mine in 2 years. From discovery to first production: typically 7–15 years for a major. From production decision to first metal: 3–6 years for a development-ready project. Supply responds to price with massive lag.

Result: prices overshoot in both directions

- Tight markets → spike to multiples of long-run average
- Surplus markets → crash below all-in costs of the marginal producer

This is structural, not cyclical noise.

Structural vs narrative-driven supply and demand

This is the single most useful distinction in macro mining analysis. Most commodities are at any moment driven by some mix of structural factors (real, long-cycle) and narrative factors (sentiment, headlines, positioning). Knowing which is which prevents trading the wrong thesis.

Structural drivers

These are the long-cycle forces that determine the actual physics and economics of supply and demand:

- **Mine development lead times** (5–15 years from discovery to production)
- **Mine depletion rates** (existing mines depleting their reserves at known rates)
- **Capex required to bring on new supply** (incentive prices)
- **Geological constraints** (where economic deposits actually exist; how grades have been trending; how much exploration has been funded over the past decade)
- **Long-term demand drivers** (EV adoption, grid build, demographic-driven consumption, urbanisation)
- **Substitution economics** (when does aluminium replace copper; when does sodium-ion replace lithium-ion)

These move slowly. They don't show up in headlines. They drive 5–10 year price trends.

Narrative drivers

These are the short-cycle forces that move prices in the next quarters and shape sentiment:

- **Specific announcements** ("Country X bans exports of commodity Y", "Producer Z guides production lower")
- **Geopolitical events** (sanctions, conflicts, election outcomes)
- **Speculative positioning** (CFTC reports showing extreme long or short concentration)
- **Specific stimulus programs** ("China announces \$X billion infrastructure program")
- **Regulatory shifts** (US Inflation Reduction Act, EU Critical Raw Materials Act)
- **Broker upgrades and downgrades**
- **Investor letters and conferences**

These move fast. They're loud. They drive 1–6 month price action.

Why the distinction matters

A move that's structural is usually durable. A move that's purely narrative is usually mean-reverting.

Consider two cases of a 30% commodity price rise:

Case A — narrative-driven: "Country exports of metal X get banned for 3 months due to political dispute. Spot price rises 30%." Reading: the underlying production capacity exists; once the dispute resolves, supply returns. Mean-reverting move; don't chase juniors at the top of the spike.

Case B — structural: "Major copper mine in Chile faces 5-year permitting delay; Indonesian nickel grades have been declining for 18 quarters; lithium hard rock incentive price is now above spot for sustained periods." Reading: the supply curve has shifted, not just temporarily disrupted. Higher prices are more likely to persist; juniors with fundable projects can re-rate durably.

The mistake is treating a narrative move as if it were structural and piling into late-stage juniors who only work at peak prices. The reverse mistake is treating a structural move as if it were narrative and selling into the early innings of a cycle.

How to read the difference

Three questions clarify which kind of move you're looking at:

1. **Has the supply curve actually shifted, or is supply just temporarily disrupted?** Mine closures vs export bans vs strikes are different. A mine closure that requires 3 years to restart is structural. A strike that resolves in 6 weeks is narrative.
2. **Has demand been re-rated based on durable factors, or based on near-term sentiment?** A change in EV adoption forecasts based on 5-year manufacturing build-out is structural. A change based on this quarter's Tesla sales is narrative.
3. **Are the major producers responding with capex commitments?** Producers don't move on narrative. When they announce major capex programs to expand or open new mines, they're confirming a structural read. When they keep capex flat through a price rise, they're treating the rise as narrative.

The commodities most prone to dangerous narrative-led moves are the ones that get hot on social media and in financial press: rare earths, uranium, lithium during peaks, gold during panics. The commodities most consistently structural in their moves are bulk industrial metals (copper, iron ore) where the supply economics dominate over headlines.

Reading the underlying commodity

You cannot invest in mining without watching the commodity itself. The minimum:

Where to track prices

- **Gold/Silver:** LBMA, COMEX (CME futures)

- **Copper, Aluminium, Zinc, Nickel, Lead, Tin:** LME (London Metal Exchange)
- **Iron Ore:** Singapore Exchange (SGX) 62% Fe futures, Platts IODEX
- **Lithium:** Fastmarkets, Benchmark Mineral Intelligence (subscriber data, but quoted in trade press); spodumene CIF China is the key Australian benchmark
- **Rare Earths:** Shanghai Metal Markets, Argus, Asian Metal — the market is opaque and Chinese-dominated
- **Uranium:** UxC, TradeTech (spot and term prices reported weekly)
- **Coking coal / Thermal coal:** Platts, Argus, IHS

What to watch

- **Spot price vs forward curve** — contango (forwards above spot) suggests storage costs/expected supply growth; backwardation (forwards below spot) suggests tightness
- **Inventory levels** — LME warehouse stocks, SHFE stocks, major producer inventories
- **CFTC Commitment of Traders reports** (for futures) — speculative positioning
- **Major producer guidance** — quarterly updates from the supermajors give early supply signals
- **Chinese demand data** — especially copper, iron ore, steel, lithium — China is 50%+ of most base metal demand

Incentive price

The commodity price required to justify building the next marginal mine. If spot is above incentive price, new supply is coming. If spot is below, supply tightens.

Rough current incentive prices (these change with input costs and grade benchmarks — verify if precision matters):

- Copper: ~\$10,000–12,000/t for new tier-2 projects
- Gold: ~\$1,800–2,200/oz for new mid-tier projects
- Lithium spodumene: ~\$1,500–2,500/t for new hard rock projects
- Nickel: ~\$22,000–28,000/t for new sulphide projects

When spot trades persistently below incentive, the cycle's bottom is forming. When spot trades far above, the top is approaching.

The China commodity manipulation playbook

China has been running variations of the same playbook on multiple commodities for the better part of 40 years. Recognising it is essential for any commodity where China is the dominant producer or consumer. The pattern works because it relies on differential time horizons — China can play decades; Western juniors and even majors can't.

The playbook in stages

Stage 1 — Build dominance. China invests heavily in domestic production capacity (rare earths, lithium chemicals, gallium, germanium, graphite, vanadium, magnesium, tungsten, antimony) or in offshore acquisitions (copper in Africa, lithium in South America, nickel in Indonesia). The investments are often loss-making at the time on a stand-alone basis but build production share.

Stage 2 — Suppress prices to crush competition. Once dominant, China expands supply or tolerates oversupply, pushing prices below the all-in cost of Western producers. Western projects close, mothball, or fail to be financed. The marginal cost curve effectively becomes determined by Chinese producers.

Stage 3 — Wait out the lag. Mining is a slow industry. New projects take 5–15 years to bring online. Once Western capacity is depleted and the next development cycle hasn't started yet, China's dominance becomes structural rather than just market-share-based.

Stage 4 — Tighten supply selectively when leverage is needed. Through export bans, quotas, environmental crackdowns, or "operational issues" at strategic facilities. Prices spike. Western governments and companies panic. Funding flows to Western alternatives, but at a 5-15 year lead time before any meaningful new capacity comes online.

Stage 5 — Allow Western capacity to be partially built, then crush again. Once Western governments have funded some alternatives and incentive prices are high enough that projects look fundable, China can release capacity, allow exports to resume, or expand its own production to undercut the new Western entrants before they reach steady state.

Where you can see this pattern

- **Rare earths:** dominant since the 1990s; periodic export quota cuts; the 2010 export restriction to Japan; the 2024 export controls on certain processing technologies
- **Tungsten, antimony, gallium, germanium, graphite:** dominance + selective export controls in 2023–2024
- **Lithium chemicals:** rapid expansion of Chinese refining capacity 2020–2023, contributing to the 2023–2025 lithium price collapse that killed many Western projects
- **Magnesium, silicon metal:** periodic supply tightening linked to electricity rationing in Chinese provinces
- **Steel-grade iron ore:** China's pricing leverage as the dominant buyer

What this means for ASX investors

A few practical implications:

1. **Be very cautious about the back end of the cycle in any commodity China dominates.** The price spike phase looks great for juniors but Chinese capacity can be released faster than Western projects can come online. The window between Western project FID and Western project commissioning is exactly when China's response can crush the trade.
2. **Distinguish between commodities China can crush and commodities it can't.** China can crush rare earths, lithium chemicals, graphite, antimony — they have or can build the supply. China cannot crush, e.g., Athabasca-grade uranium (geology dominates), Chilean copper (geography dominates), Western Australian iron ore (China is the buyer, not the seller).
3. **Watch Western policy responses.** Critical mineral lists, defence stockpiling, IRA-style subsidies, EU Critical Raw Materials Act — these are the response that creates the next investable cycle. But the policy response leads and the actual incentive prices can take years to translate to project economics.
4. **The "China is going to be cut off" thesis is a real structural change but a slow one.** Western re-shoring of critical minerals supply is a multi-decade project, not a 12-month trade. The juniors that benefit are the ones still standing in 2030, not necessarily the ones up 200% in 2026.

Why the playbook works

China's central planning gives them a longer time horizon than Western capital markets. A Chinese state-owned enterprise can absorb 5–10 years of losses to build dominance; a Western junior cannot. A Western board cannot defend a project that's losing money for three years to a stockholder base looking for quarterly results.

This is the structural reality of the resource sector. The ASX investor who ignores it will repeatedly buy Western critical mineral juniors at the top of supply-driven price spikes and watch them get crushed when the supply dynamic flips.

The thin-thematic premium

Some commodities are so small in dollar terms that even modest capital inflows produce extreme price moves and extreme junior valuations. This phenomenon — the **thin-thematic premium** — is worth understanding because the same projects that look "cheap" can become expensive in a hurry, and the same projects that look "expensive" can crash hard when the thematic capital

rotates out.

Examples of thin-thematic markets

- **Tin:** annual market ~\$8-10bn USD; a few hundred million dollars of fund flow can move the entire space. ASX names like MLX (Metals X), ELT (Elementos), SRZ (Stellar Resources) and others have moved 5-10x in tin price spikes.
- **Antimony:** even smaller; <\$3bn annual market. China dominates; supply concerns drive episodic price moves.
- **Tungsten:** small market; defence-driven buying creates premiums.
- **Vanadium, magnesium, germanium, gallium:** very small markets; thematic capital can create dramatic valuation swings.
- **Smaller rare earths:** specific elements like dysprosium, terbium have tiny markets and extreme price volatility.
- **Index inclusion effects:** when a uranium ETF gets bigger or a thematic ETF rebalances, individual junior stocks can see step-function volume changes. DYL (Deep Yellow) is an example where index/ETF flows have driven significant moves alongside the underlying uranium thesis.

How the thin-thematic premium works

When thematic capital decides a particular small commodity is "the next thing":

- A small dollar inflow chases a small number of listed names
- Liquidity-constrained trading drives extreme price moves
- Junior valuations re-rate to multiples of underlying NPV
- Capital raises become easy and oversubscribed
- IPO pipeline floods in to capture the premium

When the thematic capital rotates out:

- Same dynamics in reverse
- Even structurally good projects see their valuations collapse
- Capital raises become impossible or punitively dilutive
- The IPO pipeline dries up overnight

The trade implication

The thin-thematic trade is real, but it's a trading position, not an investment thesis. Treat thin commodities as cyclical thematic plays where:

- Position sizing is smaller than for major commodity exposure
- Entries are during the early thematic phase (before the IPO pipeline floods)

- Exits are when the thematic capital is clearly broad-based and IPOs are pricing at premium multiples
- Position duration is months to a couple of years, not 5+ year holds

The juniors that survive the thematic bust are usually the ones that were going to be built anyway based on the underlying economics. The ones that needed peak-of-cycle pricing to be funded usually don't make it through to the next cycle.

A useful test: at the long-term mean price for the commodity (not the spike price), is the project still economic? If yes, the thin-thematic move is just timing. If no, you're trading the thematic and need an exit plan.

The ASX small-cap resource cycle phases

Phase 1 — Capitulation (after a bust)

- Junior valuations 70–95% off prior peaks
- Capital raises difficult, deeply discounted, often involve free options
- Brokers reduce coverage
- Companies hibernate; no drilling, minimum activity to maintain tenements
- Acquirers pick up assets at distressed prices
- **What to do:** accumulate quality survivors at the bottom; ignore broken companies

Phase 2 — Stealth recovery

- Commodity prices begin rising
- Producers re-rate first (operating leverage to commodity price)
- Junior capital raises easier but still at premium to lows
- Brokers re-introduce coverage
- M&A activity picks up
- **What to do:** focus on near-production developers (Stage 7–8); they convert commodity move directly into equity value

Phase 3 — Broad bull

- Producers fully re-rated; multiples expanding

- Developers re-rated as projects look bankable at higher prices
- Explorers begin to move on hopes of being next discovery
- IPOs accelerate
- **What to do:** sell into developers, rotate into earlier-stage explorers selectively

Phase 4 — Mania

- Anything with the right commodity in its name moves
- Quality differentiation collapses; trash trades alongside good companies
- IPO pipeline floods the market
- Capital raises oversubscribed at premium pricing with no options attached
- **What to do:** reduce exposure across the board; raise cash. This is when the best risk/reward is on selling.

Phase 5 — Bust

- Commodity peak passes; spot starts falling
 - Broker enthusiasm fades
 - Mania-priced IPOs collapse first; then Stage 1-3 explorers; then developers
 - Dilutive raises return; SOIs balloon
 - Capitulation resets the cycle
 - **What to do:** preserve capital; if you're heavy, you've already failed
-

Macro factors that drive commodities

USD strength

Most commodities priced in USD. Strong USD = headwind for commodity prices (dollar buys more of the same physical good). Weak USD = tailwind.

Watch DXY (US Dollar Index) — sustained above 105 is generally bearish for commodities; sustained below 95 is bullish.

Real interest rates

Real rates = nominal rates - inflation. Particularly important for **gold**:

- Rising real rates → headwind for gold (it pays no yield, opportunity cost rises)
- Falling real rates → tailwind for gold

Real rate proxy: 10-year US TIPS yield.

Inflation

Generally supportive of commodities — they're real assets. But the relationship is messier than the bumper-sticker version: high inflation can also crush demand, which hurts industrial metals.

Chinese demand

~50%+ of global copper, iron ore, aluminium, nickel demand. Chinese property starts, infrastructure spend, and EV production are direct demand drivers.

China stimulus = bullish for industrial metals. China credit tightening = bearish.

Energy transition supply story

Long-cycle bullish thesis: copper, lithium, nickel, cobalt, rare earths, uranium, silver all have demand growth from EV / battery / grid / nuclear / solar / wind buildout. The demand story is more reliable than the supply response timing.

Geopolitics / sovereign supply concentration

Some commodities are heavily concentrated in single jurisdictions:

- DRC: ~70% of cobalt
- China: ~85% of rare earth refining
- Indonesia: ~50% of nickel
- Russia: significant palladium, nickel, aluminium
- Niger / Kazakhstan: significant uranium

Disruption in a concentrated source = price spike. This is actively driving Western "friendshoring" capital — including ASX explorers in friendly jurisdictions.

How macro feeds back to junior valuations

The transmission goes:

Commodity price ↑



Producer cash flows ↑ → Producer SP re-rates



Producers initiate exploration / M&A → Developer SP re-rates



Higher developer multiples re-rate explorers in same commodity



Brokers initiate coverage on more juniors → IPO pipeline opens



Retail capital flows in → SP-and-volume rotation through the sector

The lag from each step to the next is months. By the time retail is buying juniors, the early money is already rotating to the next cycle's leaders.

Track the producers in your commodity of interest. They are the canary in the coal mine for the rest of the sector.

Sector rotation patterns

When one commodity tops, capital often rotates to whichever commodity is "next". Rough recent pattern:

- 2020–2022: Lithium mania
- 2023: Lithium top, capital rotates
- 2023–2024: Uranium re-rate
- 2024–2025: Gold re-rate, rare earths
- 2025–2026: Copper themes building (verify timing/state at point of reading)

Don't fight the rotation. If your commodity has had its move and capital is rotating elsewhere, even good projects will see SP decay until rotation returns to that commodity.

How to use macro in your process

Step 1 — Identify which commodities are in early-cycle vs late-cycle

Look at prices vs incentive levels, inventory trends, and producer behaviour (cap-ex announcements signal late-cycle; defensive cap-ex cuts signal late-bear).

Step 2 — Distinguish structural drivers from narrative drivers

For each commodity in the portfolio, ask: is the current price level supported by structural supply/demand factors, or by short-term narrative? Position sizing should reflect this — heavier in structural moves, lighter and shorter-duration in narrative moves.

Step 3 — Allocate weighting between commodities accordingly

Heavier in early-cycle commodities, lighter in late-cycle ones. Avoid commodities mid-bust unless you have very long horizon.

Step 4 — Pick stage within commodity to match cycle phase

- Early bull: producers and near-production developers
- Mid bull: developers and resource-stage explorers
- Late bull: rotate out of explorers into producers (defensive), or to cash

Step 5 — Apply the China-dominance check for relevant commodities

If China dominates supply or demand for your commodity, factor in the playbook risk. Be especially cautious about late-stage juniors where the thesis depends on sustained supply tightness.

Step 6 — Stress test individual stocks at lower commodity prices

Even in a bull cycle, your individual position thesis must work at -20% commodity price. Otherwise you're betting on the macro, not the company.

Practical exercise

For each commodity you have exposure to:

1. Where is the spot price vs incentive price for that commodity?
2. Where are inventory levels (rising, falling, neutral)?
3. What are the major producers saying about supply growth?
4. What's the demand growth narrative and is it credible?
5. **Is the current price level structural or narrative-driven?**
6. **Does China dominate supply, demand, or refining for this commodity?**
7. **Is this a thin-thematic market that's prone to extreme valuation swings?**
8. Which phase of the small-cap cycle is the ASX universe in for this commodity?
9. Does your individual stock match the cycle phase (producer vs developer vs explorer)?

If the answers don't align — e.g., you own a Stage 2 explorer in a commodity at peak mania, or a late-stage critical-mineral developer that needs sustained Chinese supply tightness to be economic — you're misaligned with the cycle.

What I'm uncertain about

- Specific incentive prices change with input cost inflation, grade trends, and discount rate environment. The ranges above are rough; use as direction, not precision.
- Sector rotation timing is impossible to call precisely. The pattern that "commodity X follows Y" is much weaker than narrative suggests.
- China's demand share for various commodities is gradually shifting as their economy rebalances. The 50%+ share for base metals is broadly current but trending differently for different commodities.

- The energy transition supply story has been narrative-rich for years but actual demand growth has often disappointed vs forecasts (especially EV adoption rates outside China). Be skeptical of long-dated demand projections.
- The China playbook framing above is a generalisation. Individual commodity dynamics differ. Use it as a lens, not as a predictive model.
- Thin-thematic premium dynamics are real but also genuinely opportunistic — sometimes the thematic move is also the start of a structural change. Tin in 2024–2025 had elements of both. The framework is most useful for distinguishing pure thematic moves from durable shifts.

Module 10: Putting It Together — The "FA Story" and TA Confluence

Why this matters

Modules 1-9 give you the components. This module shows you how to assemble them into a coherent thesis for any individual stock, and how to use technical analysis as the *timing* layer on top of the *fundamental* layer.

Without this synthesis, you have fragments of analysis. With it, you have a structured, repeatable process that scales across your whole portfolio.

The FA Story — what it is

An "FA Story" is a one-pager that summarises everything material about a stock into a coherent narrative. It forces you to:

1. State your thesis clearly
2. Identify what would invalidate it
3. Quantify upside, downside, and time horizon
4. Track catalysts that will validate or kill the thesis

If you can't write the FA Story for a stock, you don't understand it well enough to own it.

The FA Story template

For every position (or candidate position), fill out:

1. Company snapshot

- Ticker, name, current MC, fully diluted MC, cash position
- Project location, jurisdiction, commodity, deposit type
- Current Lasso Curve stage (Module 1)

2. Asset summary

- Resource: tonnage, grade, contained metal, % Inferred vs I+M (Module 2)
- Reserves (if any): grade, tonnes, mine life
- Where the deposit sits on the grade benchmark for this commodity (Module 3)
- Strip ratio, recovery, cut-off (Module 3)
- Met characteristics — clean / refractory / complex flowsheet
- Tier-1 / tier-2 / marginal classification

3. Economics (if past PFS)

- NPV, IRR, capex, AISC, payback, LOM (Module 5)
- Commodity price assumption vs current spot
- Sensitivity at -20% price and +25% capex
- Construction lead time

4. Capital structure

- SOI, options, performance rights, fully diluted SOI
- Top 20 holding %
- Identifiable institutions on register
- Director on-market activity
- Cash runway in quarters
- Likely next capital raise size and timing
- Estimated dilution from current to first cash flow

5. Catalysts

- Next 3 catalysts in chronological order with dates
- Conviction level and expected SP impact for each
- Current pricing-in level (is the catalyst already in the SP, or not?)

6. Macro positioning

- Commodity cycle phase for the relevant commodity
- Spot price vs incentive price
- Stock's stage matches cycle phase? (Module 9)

7. Red and green flags

- Any flags from Module 8 checklist
- Independent net assessment

8. Thesis statement (one paragraph)

"Stock X is a Stage [N] [commodity] [deposit type] in [jurisdiction] trading at [MC] vs estimated [intrinsic value/peer multiple]. The thesis is that [catalyst path] over the next [time horizon] will re-rate the SP because [specific drivers]. Downside is bounded by [floor scenario] because [reason]. The thesis is invalidated if [specific events]."

9. Position sizing

- Conviction level
- Stage-appropriate sizing (see below)
- Entry price plan and exit conditions

10. Risks

- Top 3 specific risks (not generic "commodity price risk")
- Hedging or position management response to each

If you can fill this out for every stock you own, your portfolio is built on actual analysis. If you can't, parts of your portfolio are built on vibes.

The fundamental-price gap pattern
— the highest-conviction
asymmetric setup

This is worth elevating to its own framework piece because it's the most reliable asymmetric setup in mining and most retail miss it.

What the pattern is

A stock is trading at a market cap that's significantly below the intrinsic value implied by its underlying fundamentals — but the market hasn't caught up because:

- The macro cycle for that commodity has been out of favour for an extended period
- The company has been quiet (no recent catalysts to drive coverage)
- The valley-of-death dynamics (Module 1) have crushed the SP while the project itself has been progressing
- Specialist money knows the asset is real, retail has moved on, and the broker pipeline hasn't woken up yet

The gap closes when one or more of: a commodity cycle turn, a specific catalyst, a strategic event (cornerstone investor, takeover bid), or a broker initiation reawakens market attention.

How to identify the pattern

Three things have to align:

1. **The fundamentals genuinely justify a higher valuation.** Resource quality, project economics, capital structure, jurisdiction — all have to stand up to scrutiny. The deposit has to be real, the economics have to be defensible, and the path to value has to be visible.
2. **The price has been compressed by something other than fundamental deterioration.** Cycle, time, neglect, or rotation — not the project breaking. If the project has actually deteriorated, the gap is justified, not an opportunity.
3. **A catalyst exists that could close the gap.** Something has to wake the market up. Without that, "cheap" can stay "cheap" for years.

Worked examples to recognise

The pattern recurs across cycles. A few archetypal cases that industry observers have flagged in 2025–2026 as displaying the fundamental-price gap:

- **PGO (Pacgold) and similar** — gold explorers/developers in good jurisdictions trading well below the implied value of their drilled-out resources, because the gold sector was out of favour earlier in the cycle and the broker pipeline hadn't reactivated
- **QML (Queensland Mining Corp) and similar** — small-cap producers/near-producers where the market cap had fallen below replacement value of the producing assets

- **CNB (Carnaby Resources) and similar** — copper and base metal developers where the macro re-rate hadn't yet flowed through to the SP
- **GRL (Godolphin Resources) and similar** — multi-commodity explorers in tier-1 jurisdictions trading at distressed multiples relative to their resource scale and grade

The specific tickers will rotate over time. The pattern doesn't.

Why the gap exists

Three structural reasons:

1. **Mining capital is highly cyclical.** Generalist money rotates in and out of the sector. When the sector is out of favour, even the highest-quality juniors trade at distressed multiples. When the sector is in favour, even mediocre juniors trade above intrinsic value.
2. **Information arrival in mining is lumpy.** A stock that's done good work on a project but hasn't had a major catalyst in 18 months effectively disappears from the market's mind. The work gets done, the value accrues, but no one is paying attention.
3. **Broker coverage follows price, not fundamentals.** Small-cap brokers initiate coverage when stocks are moving up and drop coverage when stocks are moving down. The result is that the lowest-attention period is often when the fundamentals are quietly improving the most.

The trade structure

The fundamental-price gap is the closest thing to free money in mining, but only if you do the work properly. Three rules:

Rule 1: Verify the fundamentals before getting interested in the price. A "cheap" stock is only interesting if the fundamentals are actually good. Run the FA Story. Check resource quality (Module 2/3), check capital structure (Module 6), check for red flags (Module 8). If the stock is cheap because it deserves to be, walk away.

Rule 2: Identify the specific catalyst that could close the gap. A cycle turn is a catalyst, but it's a slow one. A specific company catalyst (PFS, FID, takeover bid, cornerstone investor) is faster and more identifiable. The best fundamental-price gap setups have at least one identifiable catalyst within 6–12 months.

Rule 3: Size for time, not just outcome. Even with the right setup, the gap can take 12–24 months to close. Position sizing should reflect that you're being paid for time as well as for outcome. Don't go heavy and don't expect immediate validation.

When the pattern fails

Most often when:

- The fundamentals were actually weaker than they looked, and the cheap price was correct
- The catalyst gets pushed back repeatedly and the cycle turns against the commodity before the catalyst arrives
- A capital raise during the orphan period dilutes the per-share value and resets the math
- The stock gets de-listed or the company pivots into something unrelated (Module 8 transformational acquisition pattern)

A failed fundamental-price gap is usually an FA Story problem (you got the fundamentals wrong) or a timing problem (the catalyst arrived too late). Both are post-mortem-able.

Why it's the highest-conviction setup retail consistently miss

The pattern requires three uncomfortable things:

- **Buying in the absence of momentum** (the trend is sideways or down, not up)
- **Holding through quiet periods** (months between catalysts; SP barely moves)
- **Trusting your own work over market sentiment** (the market is telling you the stock is worth less than your analysis says it is)

Most retail can't do this. They want momentum, they want validation, they want short time horizons. The fundamental-price gap rewards the opposite — patience, conviction, and the discipline to ignore short-term price action when the underlying fundamentals are actually strengthening.

This is why the pattern persists. If most retail could execute it, the gap wouldn't exist.

Position sizing by stage

Risk is fundamentally different by stage. Sizing should reflect that.

Stage 1–2 (Concept / Early Drilling)

- **Max 0.5–2% per position** of total mining allocation
- Lottery economics: most go to zero; rare ones 20–50x

- Treat each position as a near-write-off; size to make the win meaningful but the loss survivable

Stage 3–4 (Discovery / Resource Definition)

- **2–4% per position**
- High volatility, high information arrival rate
- Trim into strength after major catalysts; redeploy on pullbacks

Stage 5–7 (Studies / DFS)

- **2–5% per position**
- The orphan period — long time horizons, limited catalysts
- Size for time as much as outcome
- **The fundamental-price gap pattern often appears here.** A stock that's done good work but hasn't had a re-rate event yet, in an out-of-favour sector, is the prototype.

Stage 8–9 (Construction / Commissioning)

- **3–6% per position**
- Highest reward-to-risk if execution is on track
- Cap-ex blowouts and ramp problems are the main risks; trim if they appear

Stage 10 (Production)

- **5–10% per position** if conviction is high
- Lower risk, lower per-stock upside, but compounds via dividends and reserves growth
- Suitable as portfolio anchors

These are **defaults**. Risk-tolerant investors can lean heavier in earlier stages; risk-averse investors should bias to producers.

Total mining sector exposure

A defensible cap on total mining sector exposure as % of total portfolio depends on your circumstances, but most professional resources investors keep total mining exposure under 30–50% of equity portfolio outside of bull-cycle peaks.

TA — what it actually does for you

Technical analysis tells you nothing about whether a stock is fundamentally cheap or expensive. It tells you about:

1. **Trend** — what direction is the SP moving in
2. **Momentum** — is the move accelerating or fading
3. **Support / resistance** — where price has historically reacted
4. **Volume** — is the move backed by participation or thin
5. **Sentiment / positioning** — through indicators like RSI, accumulation/distribution

TA is the **timing layer** on top of FA. FA tells you *what* to buy; TA tells you *when*.

Practical TA for mining FA investors

You don't need to be a chartist. The minimum useful kit:

1. Trend identification

- 50-day and 200-day moving averages
- "200dMA rising" = uptrend; "200dMA falling" = downtrend
- Don't fight the 200dMA

2. Support and resistance

- Prior swing highs and lows
- Round numbers (psychological levels)
- Levels of major capital raises (often act as resistance — placement shareholders often sell at break-even)

3. Volume

- Major moves on heavy volume = real

- Major moves on thin volume = suspect (often reverse)
- Volume on breakouts is critical confirmation

4. Relative strength

- Is the stock outperforming or underperforming peers and the commodity?
- Outperformance suggests something is being priced in that you don't yet see
- Underperformance suggests the opposite

5. Accumulation patterns

- Long sideways range with occasional volume spikes = often institutional accumulation
- Watch for breakouts above the range
- **The fundamental-price gap often shows this pattern just before the gap closes**
 - quiet accumulation by specialist money before the catalyst hits

6. Distribution patterns

- Failed rallies on heavy volume
- Lower highs after a sustained uptrend
- Often precede major down moves

What you don't need

Most chart patterns (head and shoulders, cup and handle, etc.), Elliott Waves, Fibonacci retracement debates, and the entire indicator soup. They add complexity without much edge for fundamentally-driven mining investing.

FA + TA confluence — the actual playbook

Setup A: FA bullish, TA bullish

- Strong fundamentals

- Uptrend, above 200dMA, recent breakout on volume
- **Action:** full position size, add on pullbacks to support

Setup B: FA bullish, TA bearish (the orphan trap / fundamental-price gap)

- Strong fundamentals
- Downtrend, below 200dMA, no buying interest
- **Action:** This is where the fundamental-price gap pattern lives. Wait, but watch. Even a great FA story can take 6-18 months to find a bid in a downtrend. Build a watchlist and start a small position when you see signs of accumulation (volume spikes on no news, sideways range starting to compress); don't go heavy until trend turns or a specific catalyst is imminent.

Setup C: FA neutral, TA bullish

- Mediocre fundamentals
- Stock is moving on momentum / sector rotation / promotion
- **Action:** if you're trading, take a small position with tight stop. If you're investing, ignore.

Setup D: FA bearish, TA bullish (the trap)

- Weak fundamentals (poor grade, poor jurisdiction, poor management)
- Stock running on hype
- **Action:** avoid. The eventual reckoning is brutal.

Setup E: FA bearish, TA bearish

- Avoid. Don't bottom-fish broken companies.

The best risk-adjusted setups come from Setup A and from Setup B *after* the trend turns or a catalyst is identified.

When to sell

Selling is the part most investors are worst at. Pre-define exit conditions for every position:

Sell conditions worth respecting

1. **Thesis invalidated** — the specific catalyst path you bought for has failed
2. **Position size too large** — winning positions can grow into oversized concentrations; trim back to plan
3. **Better opportunity** — capital is fungible; if you find a clearly better setup, rotate
4. **Macro cycle has turned** — trim into late-cycle mania regardless of how good a specific story still looks
5. **Capital raise dilution** — significant unexpected dilution may invalidate your math; recheck the FA Story
6. **The gap has closed** — for a fundamental-price gap trade, when the SP has re-rated to roughly fair value, the asymmetric edge is gone. Rotate to the next gap.

Sell conditions NOT worth respecting

1. "It's down 20%" — only sell on thesis change, not price action alone
2. "I want to lock in profits" — only relevant if you've hit your price target or thesis playing out
3. "Someone on Twitter said it's a fraud" — do your own work
4. Any single-day move

The best mining investments often look bad for periods between catalysts. Selling on price weakness alone is how retail consistently gives up the second peak of the Lassonde curve.

Building and maintaining the system

This is a practice, not a one-off exercise.

Weekly

- Skim ASX announcements for all holdings (15 min)
- Update catalyst calendar with new dates / slippage
- Note any unusual SP / volume action — particularly accumulation patterns on holdings with fundamental-price gap setups

Monthly

- Re-read FA Story for each holding; mark anything stale
- Update macro view on each commodity
- Review portfolio sector allocation vs cycle
- Scan for new fundamental-price gap candidates in commodities you understand

Quarterly

- Read the Appendix 5B for each holding in detail
- Update cash runway calculations
- Re-check fully diluted MC after any raises

Annually

- Read every annual report for every holding (yes, all of it — financial notes especially)
 - Re-test thesis from scratch
 - Cull positions you wouldn't initiate today
-

Final perspective

The course brief talked about creating a "well-rounded skill set" and an "FA Story" that integrates everything. Here's the honest version:

The framework is straightforward. The execution is a multi-year discipline. You will get things wrong — sometimes you'll lose 50%+ on a position despite doing the work. The point isn't to be right every time. The point is to:

1. Have a thesis that's structured, falsifiable, and defensible
2. Size positions so individual losses don't wipe you out
3. Hold long enough for catalysts to play out
4. Sell when thesis breaks, not when SP wobbles
5. Repeat across enough positions that variance averages out

Mining is a sector where the math punishes the lazy and rewards the patient. Most retail are neither. If you actually do the FA Story work for every position, hold size discipline, and respect cycles, you'll be ahead of the vast majority.

The fundamental-price gap is the highest-leverage edge in this entire framework — but it's only available if you've internalised everything else. Module 1's stage framework, Module 2's resource

literacy, Module 3's grade context, Module 4's drill reading, Module 5's economic study analysis, Module 6's capital structure reading, Module 7's catalyst calendaring, Module 8's flag detection, and Module 9's macro overlay all have to be working together for you to identify a fundamental-price gap correctly. Skip any one and you'll mistake a value trap for an opportunity.

Practical exercise — your first FA Story

Pick one stock from your current ASX mining holdings. Build the full FA Story above using only ASX-disclosed information (announcements, quarterlies, annual report, presentations). Spend at least 4 hours on it.

When done, ask yourself: *"If I had to defend this position to a sceptical professional investor in 30 minutes, do I have the answers?"*

If the answer is yes, you've internalised the framework. If no, the gaps tell you what to research next.

Practical exercise — finding a fundamental-price gap candidate

Separately, pick a commodity you understand and want to be long over the next cycle. Then:

1. List every ASX small-cap producer/developer/explorer in that commodity
2. Filter to those in tier-1 or tier-2 jurisdictions with no major Module 8 red flags
3. Filter to those whose SP is in the bottom third of its 2-year range
4. For the survivors, build a quick FA Story and identify whether the fundamentals justify a higher valuation
5. For the candidates that pass the fundamental check, identify the specific catalyst that could close the gap
6. Size positions to your conviction in the work, with the expectation that you may wait 12-18 months for validation

This exercise generates the actionable watchlist for the next cycle leg. Most retail will not do this work. That's exactly why the gap exists.

What I'm uncertain about

- Position sizing recommendations are conventional ranges; your circumstances (total wealth, time horizon, dependents, risk tolerance) should drive your specific allocations.
 - TA effectiveness in junior mining is debatable. The signal-to-noise is lower than in liquid large-caps because thin trading and concentrated holders can drive moves unrelated to fundamentals or technicals. Use TA as a tiebreaker, not a primary driver.
 - The "weekly/monthly/quarterly/annual" cadence above is a starting framework. Some periods of intense activity (live drill programs, study release windows) require daily attention; quiet periods need much less.
 - The fundamental-price gap pattern is real but the specific tickers cited above will rotate over time. The pattern doesn't generalise to "any cheap-looking stock"; the specific conditions (genuine fundamentals, identifiable catalyst, cycle-driven compression) all have to be present. Most candidates that look like the pattern are actually value traps.
-

Where to go from here

You now have:

- The Lassonde Curve framework (Module 1)
- JORC literacy + foreign estimates (Module 2)
- Grade context + equivalent grades (Module 3)
- Drill result decoding + drilling-type taxonomy (Module 4)
- Economic study analysis + operational realities (Module 5)
- Capital structure analysis (Module 6)
- Catalyst calendaring + dormant-company detection (Module 7)
- Red flag detection + behavioral patterns (Module 8)
- Macro overlay + structural-vs-narrative + China playbook (Module 9)
- Synthesis, position sizing, and the fundamental-price gap (Module 10)

The next step is **applying it to real stocks**. Not in theory — in your actual portfolio, with real money at stake. The framework only becomes valuable when you've used it enough that it becomes second nature.

When you're ready, bring me specific companies, specific announcements, or specific scenarios and we'll work through them together using the framework.