



The 69 cent problem

Why better targeting is making exploration's real problem worse, and what a board can do about it

The number

Between 2012 and 2021, the world spent US\$196 billion looking for mineral deposits. In return, it generated an estimated US\$135 billion of discovered value. That's 69 cents back for every dollar in. The figure isn't mine. It belongs to Richard Schodde of MinEx Consulting, from his decade-scale accounting of global exploration performance (IMARC, Sydney, October 2023). I think it is the most important number in this industry that almost nobody talks about.

Here's one way to calibrate it. A poker machine in an Australian pub is required by law to return between 85 and 92 cents in the dollar, depending on the state. So, over the last decade global mineral exploration, tens of thousands of technically trained people spending two hundred billion dollars of other people's money, returned materially less than a machine that is legally engineered to take money off you.

I want to be careful with that comparison, because it is easy to make it glib. Exploration isn't gambling. The gambler buys entertainment, and the machine's odds are fixed. The explorer buys information, and can, at least in principle, change the odds. But that is exactly what makes the number damning rather than just unflattering. The machine cannot beat its own program. We can. But over a full decade, in aggregate, we didn't.

It isn't geology's fault either. Over the same ten years, on the same accounting, Australia and Africa each returned \$1.19 per dollar spent. Latin America returned \$0.42, and Canada \$0.78. The rocks didn't get three times better crossing an ocean. Some of that spread is endowment and data maturity, fair enough. But a gap that wide, inside one decade and one industry, isn't a geology story. It's a story about where the money went, and how the decision to send it was made.



The bottleneck has moved

For most of this industry's history, the binding constraint was finding. Targets were scarce, data was thin, and a geologist who could generate a credible drill target was the most valuable person a company had. Every system we run on was built in that world. The budgets, the promotions, the conference programs, our very self-image.

That world is ending. Machine learning and AI have made target generation cheap, and it is getting cheaper. Prospectivity models rank ground at continental scale. Agentic systems mine a century of legacy reports for the drill holes everyone forgot about. Every year the tools get better at producing exactly what the old world prized, the credible, drill-ready prospect. That is genuine progress, and I'm not knocking it.

But watch what it does to a junior explorer's quarter. Better targeting produces more drill-ready prospects than any exploration budget can test. The constraint quietly moves from finding targets to choosing between them, and here is the asymmetry we haven't faced. We have spent decades and serious money getting better at finding, and we have no shared discipline at all for choosing.

Prospectivity ranking isn't choosing. A prospectivity score tells you which target is most likely to host mineralisation. It doesn't tell you which program creates the most shareholder value per dollar spent. That second question depends on cost, on what success would actually be worth, on what else you could do with the money, and on what you'd do with the answer. They're different questions. The industry answers the first with growing sophistication and, in most boardrooms, doesn't ask the second at all.



How programs actually get approved

I should say plainly where I'm writing this from. I'm a geologist who has spent thirty years in exploration, made a number of discoveries, including one that is now a multi-million-ounce deposit, and ran two ASX-listed explorers as managing director. I have written the board papers, argued for the programs, faced the shareholders and lived with the results. So, when I describe how programs get approved, it isn't an observation from the outside. I was in the room.

The quarterly reality of a junior explorer looks like this. More projects than budget, and a board that wants progress on all of them. Programs get approved on geological merit, market appetite and momentum. Somewhere in the paper, the phrase "highly prospective" does an enormous amount of load-bearing work. The question that rarely appears, and almost never as a number, is the only one that actually matters to the people whose money it is. If we spend this, what do we expect back?

There are three honest reasons the question goes unasked. The first is training. Geologists are taught to find, not to price. The toolkit for valuing a decision under uncertainty just isn't in the curriculum, and most of us met it late and by accident, if we met it at all. The second is incentive. For a company whose real business model is capital raising, hope is easier to sell unquantified, and a stated expected value can feel like a hostage handed to the market. The third is difficulty. The probability of success and the value of success are genuinely hard to estimate, and easy to dispute.

All three reasons are real. None of them survives a look at the industries next door.



The tools exist. Mining never imported them

The oil and gas sector has priced its prospects since the 1960s. No serious operator commits a rig without a probability of success built from the chance of charge, reservoir, seal and trap, a risked volume, and an expected monetary value that has to beat the alternatives. Venture capital sizes every cheque against a power-law distribution of outcomes in which most investments fail, a return profile that looks uncannily like greenfields exploration. Private equity won't hold an asset without an exit thesis stated in advance.

All three of those industries operate under uncertainty entirely comparable to ours. All three treat the difficulty of the estimate as the reason to make it carefully, not the excuse to skip it, and the toolkit they run on is unglamorous and decades old. Expected value analysis, real options, Bayesian updating. None of it is beyond a profession that inverts geophysics before breakfast. Mining just never imported it. What I know from the coal face is the decision-making this toolkit needs to replace. Hunch and experience overriding formal assessment, and the in-vogue commodity chased because it raises well, not because it returns well. If any of what follows is wrong, I'd rather be corrected than agreed with. This has to work as science, not salesmanship. It should survive peer review, and I expect it to get some.

The discipline, in three questions

Strip the toolkit down to what a board actually needs, and it reduces to three questions that should be answerable, in numbers, in the board paper, before any exploration program is approved.

Question one: If we spend \$500,000 on this program, how much shareholder value do we expect to create?

Take a drill-ready target and call it Prospect A. Write down the probability the program succeeds. Not "good". A number, with its basis stated. Say 15 per cent, benchmarked against what programs of this type, in this terrane, at this stage, have historically delivered. Then value what success would be worth before the rig arrives. Say \$6 million of added project value. The expected value of the program is 15 per cent of \$6 million, less the \$500,000 cost. That comes to positive \$400,000, an expected \$1.80 in value for every dollar spent.



Every input will be disputed, and that is how it should be. The dispute is the discipline. A board arguing about whether the probability is 10 per cent or 20 is doing exactly its job. A board handed "highly prospective" has been given nothing to argue with. A crude, contested, written-down 15 per cent beats an uncontested adjective every time, because the written number can be checked against outcomes and improved. The adjective cannot.

Question two: With four projects and \$1 million to spend, what allocation produces the best return on shareholders' capital?

Same arithmetic, run comparatively, because programs approved in isolation are how value quietly dies. Consider a second target, Prospect B. Geologically credible, ranked "prospective" by any model you like, with a 10 per cent chance at \$3.5 million of value and the same \$500,000 cost. Its expected value is negative \$150,000. Seventy cents of value per dollar spent. That is the 69-cent problem, one cheque at a time. In a company that approves programs one by one, B gets drilled eventually. In a company that ranks every ready program by expected value per dollar, B never beats A for the money, and the board can say why.

Allocation also means staging. Split A's program in two, and agree, before hole one, what first-stage result would move the 15 per cent, and what result stops the spend. If encouraging results would lift the probability to 45 per cent, stage two is comfortably worth drilling. If discouraging results would drop it to 2, it isn't, and half the budget walks away to the next-best cheque. Run the numbers, and the pre-agreed right to stop adds roughly \$85,000 of expected value to the identical program on the identical rocks. That is a real option. Exploration is full of them. The industry just declines to price what it already owns.

Question three: If we ask new investors for \$1 million, what return should they expect once the program is run?

This is the first two questions, said out loud to the people paying. A board that has done the work can state a sentence almost no junior explorer can currently say. ***"Your \$500,000 funds Prospect A. The program is expected to return \$1.80 in value for every dollar spent, and we have already agreed the result that would stop us at half the spend."***



A board that cannot state it is pricing its raise on narrative alone. A decade of narrative-priced raises is exactly what 69 cents in the dollar looks like from above. I would add one more thing, wearing my shareholder hat rather than my geologist one. Every investor should be asking a CEO one question before every raise. What do you expect my \$1 million to be worth after you have completed the program you are raising it for? It is a fair question. It should be a routine one. That is exactly what this discipline produces.

How much of the answer can be stated publicly is a question for disclosure advice, and it will differ from raise to raise. What a board can always say, because it is simply a statement of fact, is that the work has been done, the figures exist, and the result that would stop the spend was agreed before the rig arrived. Investors should settle for nothing less.

The obvious objection

Can the odds actually be brought up? Not the way a gambler would hope, and it matters to be precise here. No framework makes an individual drill hole luckier. The orebody is there or it isn't, and expected value analysis will never put metal in barren ground.

What the discipline changes is the portfolio. It stops money flowing into programs whose expected value is negative before the rig arrives. It sizes each program to the value of the information it will produce. It kills fading projects quickly, by prior agreement, instead of funding them out of habit and hope. Same geology, same drill holes, fewer wasted dollars around them. That is how \$0.69 moves toward and past \$1.00. Not by changing the rocks, but by changing which cheques get written. The regional spread in Schodde's own numbers says the headroom is real. Parts of this industry already clear the bar by a wide margin.



Monday morning

None of this needs software or a data-science team. It needs a board to require it. Six checks, one page, before any program is approved. The probability is written down with its basis. Success is valued before drilling. The expected value is positive. The kill criterion is agreed in advance. The program is the best cheque available across the portfolio, not just a good one in isolation. If capital is being raised, the new dollar's expected return is a sentence the board can say out loud. The one-page version of that test lives at rocworth.com.au.

To be clear, requiring the six checks is the board's job. Producing them does not have to be. As Results pending sets out, many executives already hand their announcements to consultants who do them well and hand them back. The numbers behind the six checks can be handed over the same way. That is the work Rocworth does.

Start crude. The first probabilities will be wrong. The first valuations will be argued to a standstill. The first kill criterion will be honoured awkwardly. That is what the start of a discipline looks like, and it is the whole point of the Bayesian habit. Written-down estimates get checked against outcomes and improve with every program. Instinct doesn't. Instinct just gets older.

The blank line

Nearly every junior explorer can tell you, at length and with conviction, that its targets are compelling. Almost none can tell you the expected return of its own next program. The 69 cent problem lives in that blank line. Machine learning and AI, for all their genuine gifts, are going to keep making the finding easier and the blank line more expensive. The better the targeting gets, the more the choosing decides. The tools to fill the line in are sixty years old, proven next door, and within reach of any board willing to require them. What the last decade returned was sixty-nine cents. What the next one returns is, to a degree this industry hasn't yet taken seriously, a choice.



Benjamin Bell is an exploration geologist of thirty years whose record includes the discovery of the 2.44 Moz Katanning Gold Deposit and managing-director roles at two ASX-listed explorers. He writes on exploration economics at Rocworth (rocworth.com.au), where the companion essay “Results pending” sets out why Rocworth exists. He is presenting on this topic at the AIG ML/AI in Greenfields Exploration workshop, Perth, 29 July 2026. Figures in the worked examples are illustrative and are not a valuation of, or reference to, any actual project or company. Exploration performance data: R. Schodde, MinEx Consulting, IMARC Sydney, 31 October 2023, slides 22 & 30. All content is general information about decision methods. It is not financial product advice, takes no account of any person’s objectives, financial situation or needs, and is not a recommendation concerning any security. Companies and investors should obtain their own professional advice.
