QUANTUM COMPUTING AND CATALYST DESIGN: WHAT INVESTORS NEED TO KNOW ABOUT NITROGEN FIXATION

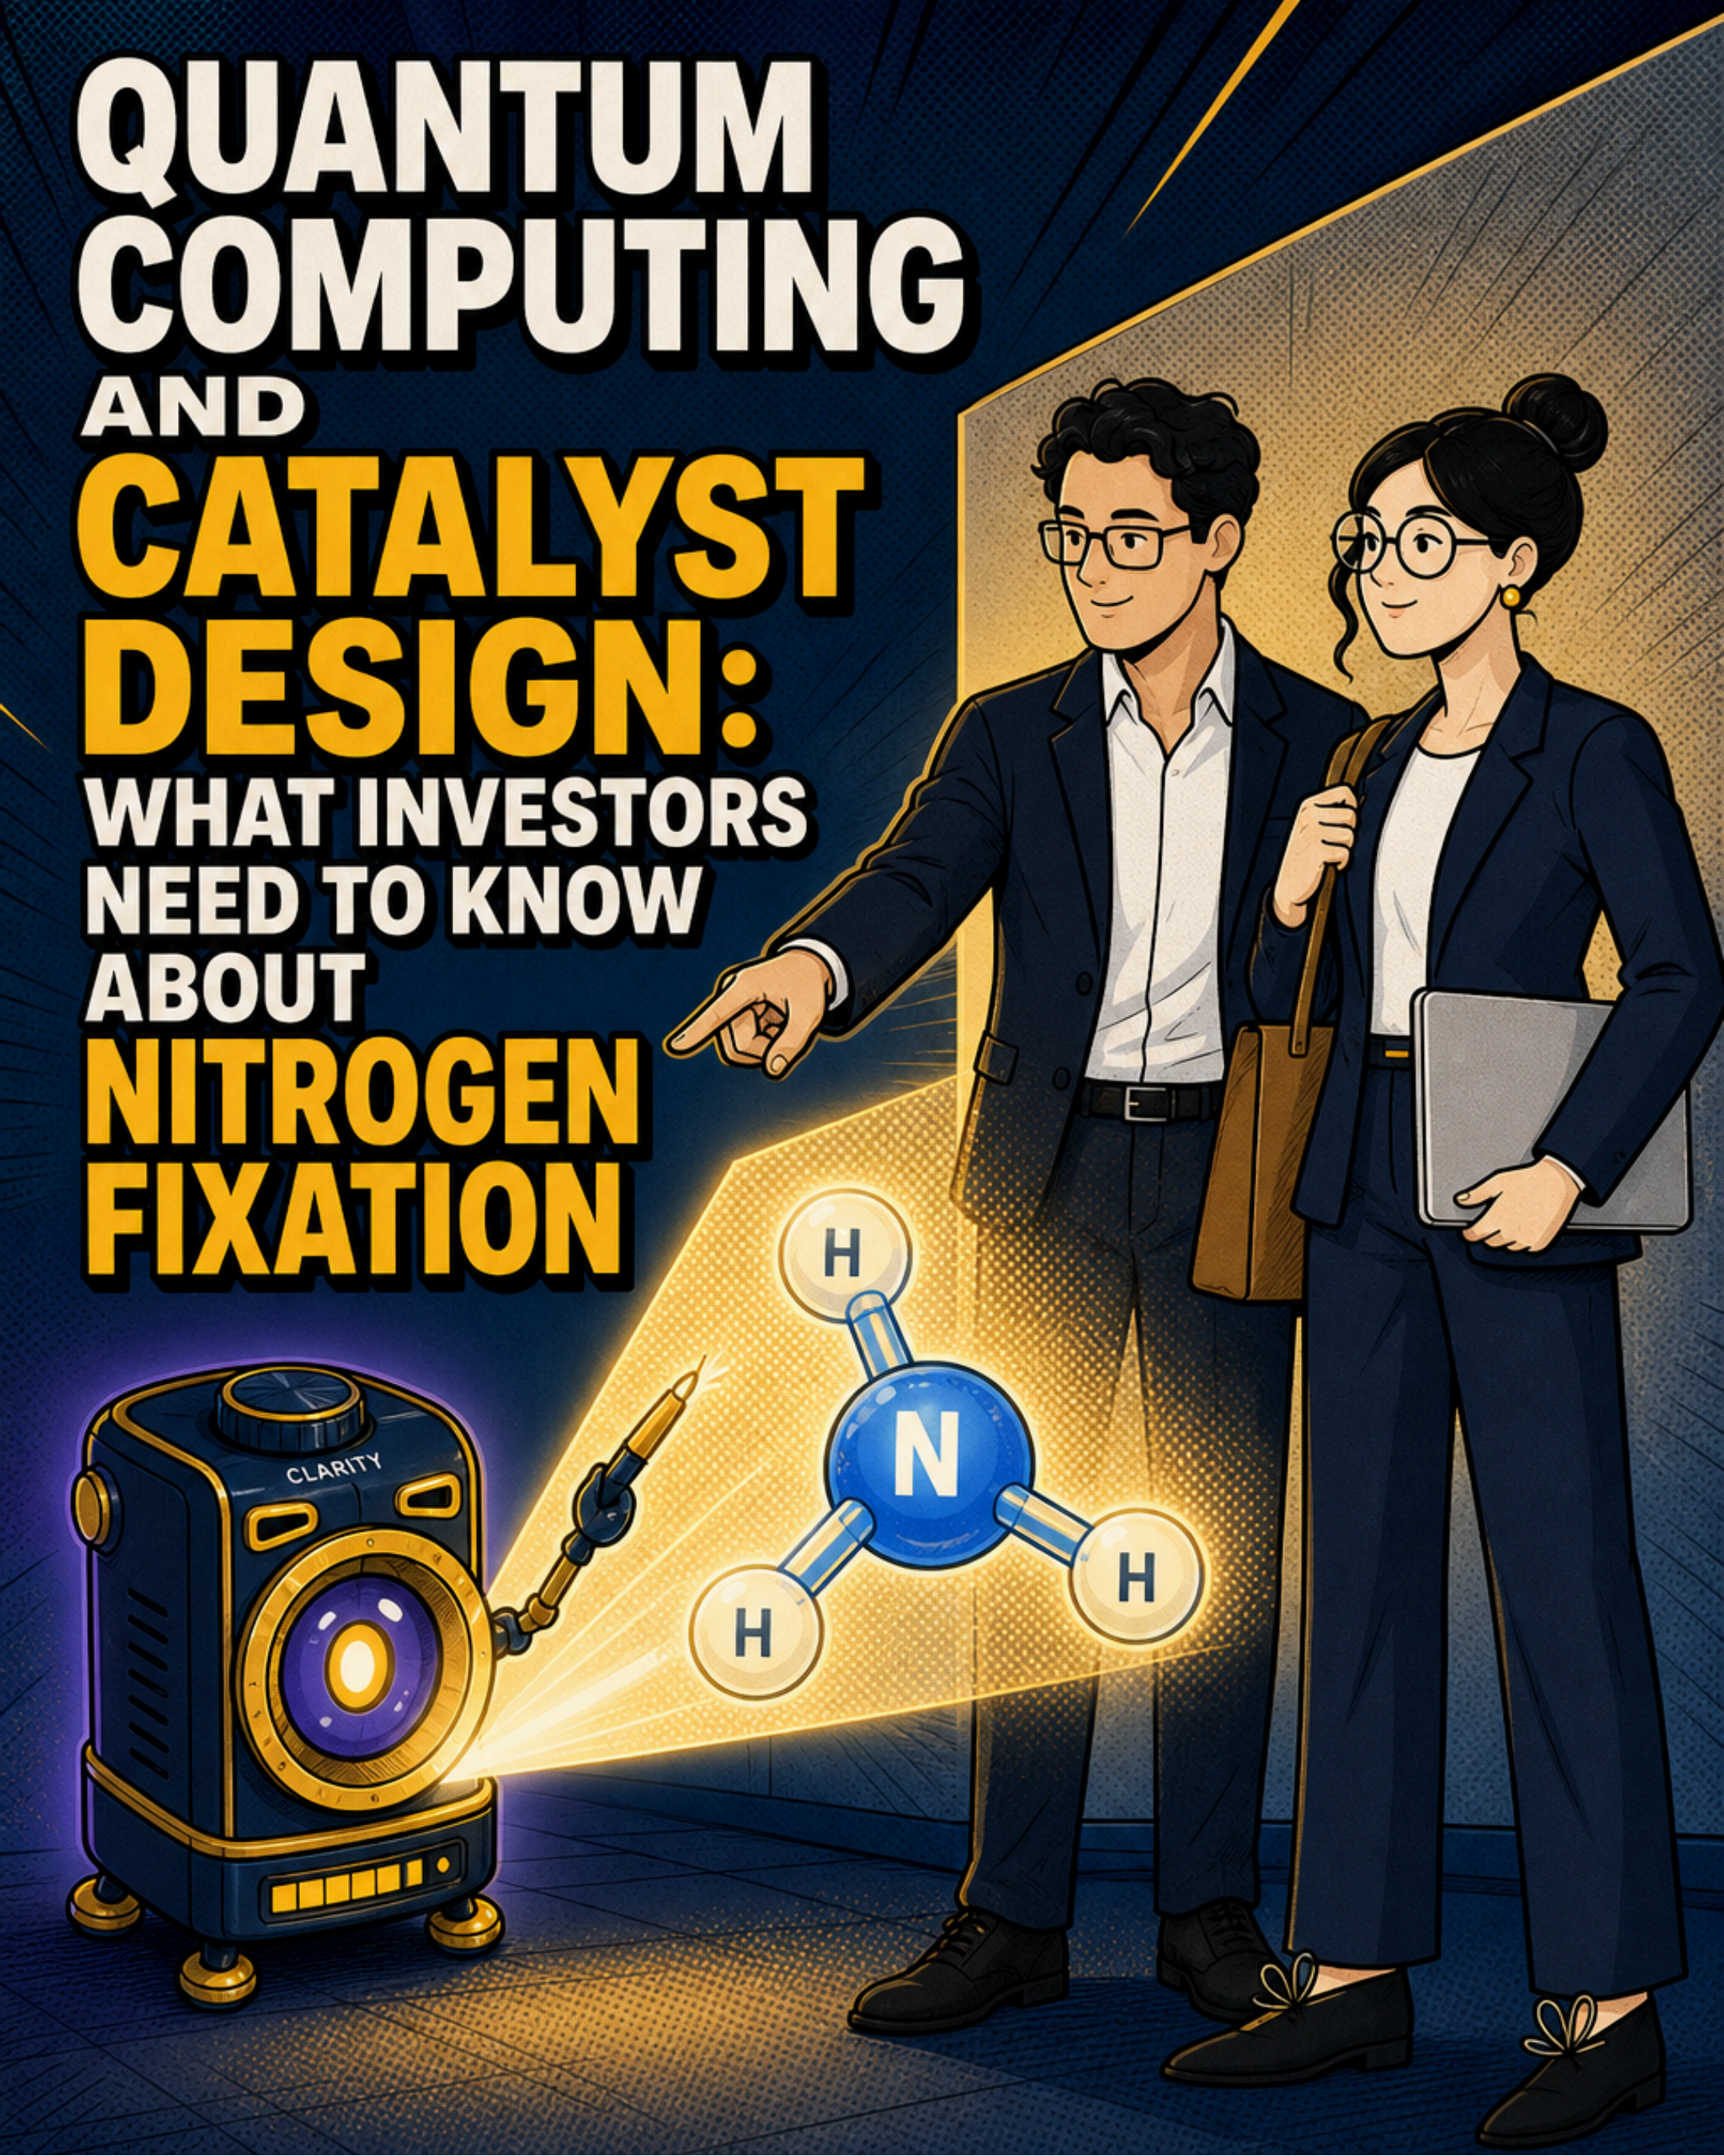

The molecule that feeds the world, and why it costs so much to make

About half the nitrogen in the human body passed through an industrial reactor at some point. That reactor runs the Haber-Bosch process, the century-old reaction that combines nitrogen and hydrogen to make ammonia, the foundation of synthetic fertiliser. It is, by most measures, the chemical process most responsible for the current scale of human civilisation.

It is also extraordinarily expensive to run.

The Haber-Bosch process is responsible for 1–2% of total global energy consumption and 1–3% of global CO_2 emissions, consuming roughly 3–5% of the world's natural gas production each year.

To put a number on the scale:

around 170 million metric tonnes of ammonia are produced globally each year, with approximately 80% going into fertilisers.

THE HABER–BOSCH PROCESS

$$N_2 + 3H_2 \rightleftarrows 2NH_3$$

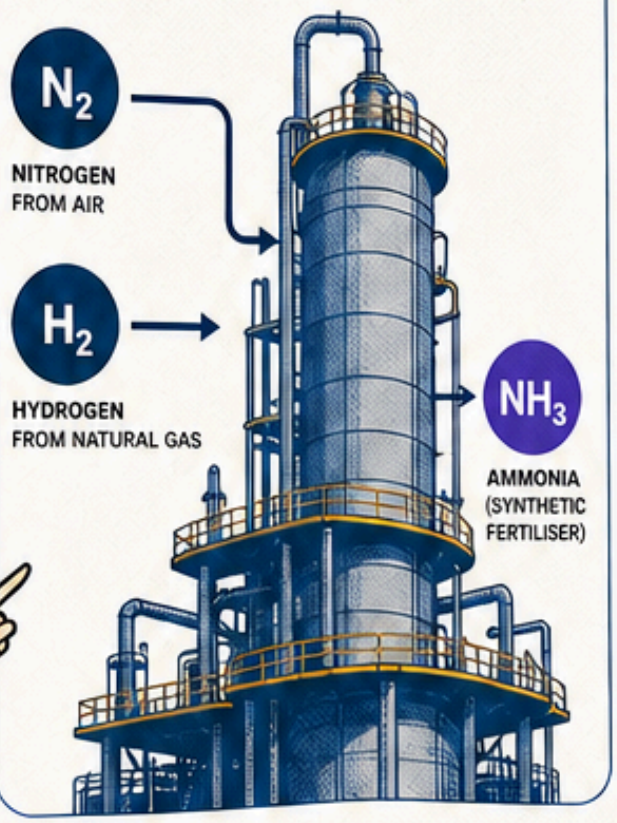

AMMONIA PRODUCTION BY THE NUMBERS

~170 MILLION METRIC TONNES PER YEAR GLOBALLY

~80% GOES INTO FERTILISERS

1–2% OF TOTAL GLOBAL ENERGY CONSUMPTION

1–3% OF GLOBAL CO_2 EMISSIONS

3–5% OF THE WORLD'S NATURAL GAS PRODUCTION

Those figures are not a chemistry trivia problem. They are an investment signal. A meaningful improvement in catalyst efficiency for nitrogen fixation would be one of the largest single decarbonisation and cost-reduction events in the global chemical industry. The question is what it would actually take to find a better catalyst. **And that is where quantum computing enters the story.**

WHY CLASSICAL COMPUTERS HIT A WALL

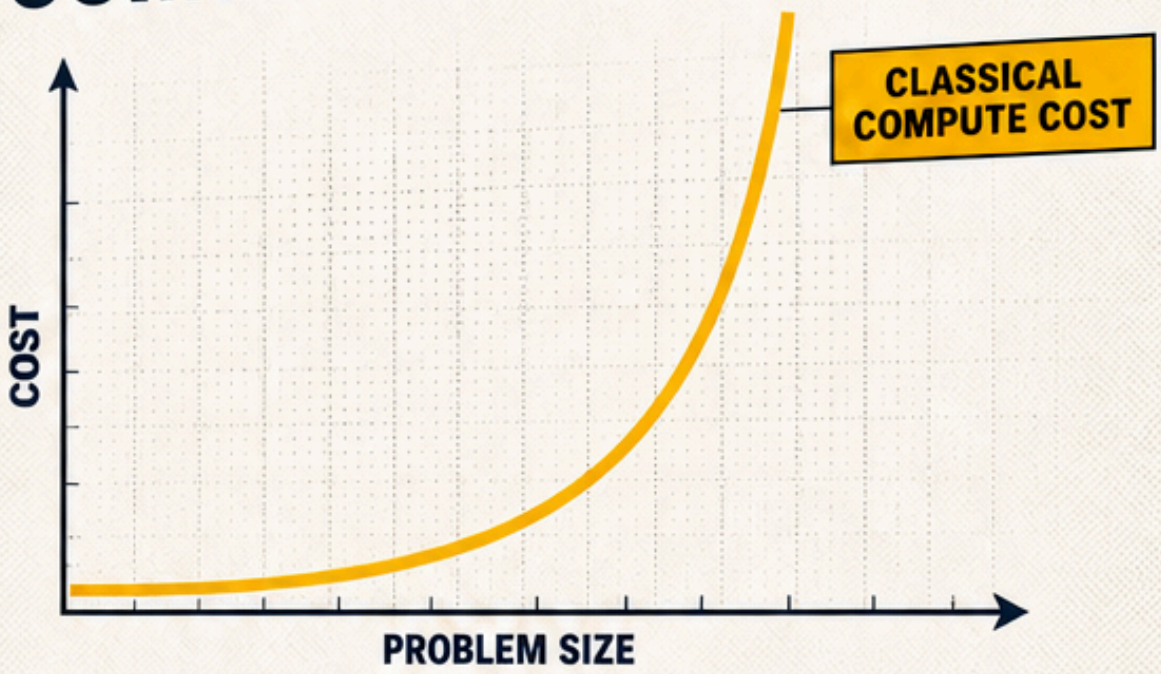

DFT has a hard ceiling.
It faces limitations in quantitatively describing strong electronic correlations, electron spin states in magnetic catalysts, and subtle energy differences — the very things that matter most.

The root problem is computational scaling.
Classical computers face an exponential growth in complexity when modelling strongly correlated electron systems.

This is a structural limitation.
Not a temporary gap that better supercomputers will close.

Quantum computers are built differently.
They represent information using qubits in superposition and exploit quantum correlations directly.

Electronic structure simulations represent one of the most promising application areas where quantum computers have the potential to outperform classical methods.

The molecule that quantum researchers keep coming back to

If you want to understand why the quantum chemistry community is so focused on nitrogen fixation specifically, you need to know about nitrogenase and its FeMo-cofactor.

Nitrogenase is an enzyme found in certain soil bacteria. It does something the Haber-Bosch process requires 400–500°C and 150–300 atmospheres of pressure to replicate: it splits the triple bond in atmospheric nitrogen and converts it into ammonia.

It does this under ambient conditions, at room temperature and standard pressure.

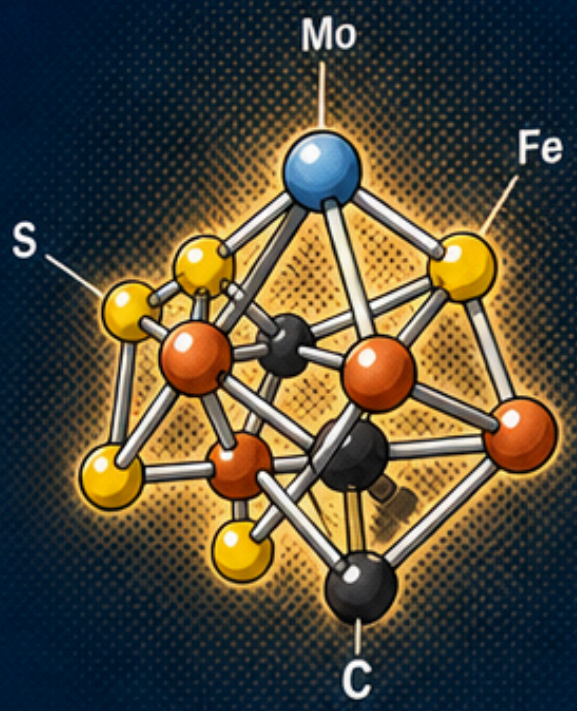

FeMoco

The problem is that FeMoco is precisely the kind of system classical computers **cannot model accurately.**

FeMoco contains multiple metal centres (iron and molybdenum atoms), making the spin dynamics particularly complex and difficult to simulate.

The canonical resource estimate for FeMoco, established by Reiher et al. in a widely cited 2017 PNAS paper, places the required active space at approximately (54 electrons, 54 orbitals) — a configuration that exceeds what is tractable via exact classical methods such as full configuration interaction, and which remains deeply challenging even for approximate classical approaches such as DMRG and FCIQMC.

That observation is worth sitting with. The most important catalyst design problem in industrial chemistry is, in a precise technical sense, intractable for classical supercomputers using exact methods.

It isn't a matter of waiting for a faster processor.

What VQE actually does

The Variational Quantum Eigensolver (VQE) finds the lowest energy (ground-state energy) of a molecule.

1. Molecule

The system we want to simulate.

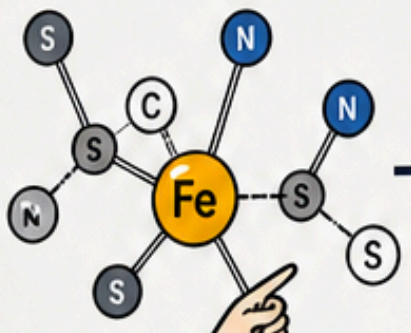

2. Quantum circuit

Encodes a trial wavefunction with adjustable parameters.

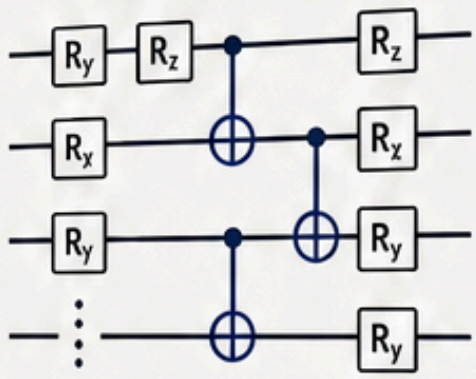

3. Measure energy

Run the circuit and measure the energy.

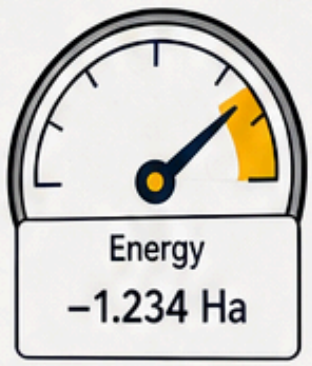

4. Optimise

Classical optimiser updates the parameters to lower the energy.

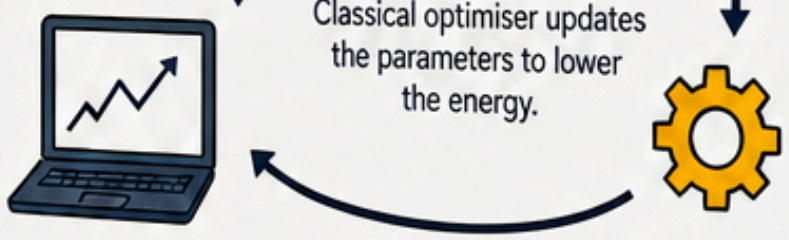

5. Converge

Repeat until the lowest energy is found.

VQE = Quantum processor for the hard part. Classical computer for the optimisation loop.

COMPRESSING THE DISCOVERY CYCLE

Traditional trial-and-error vs. In-silico pre-screening

TRADITIONAL TRIAL-AND-ERROR

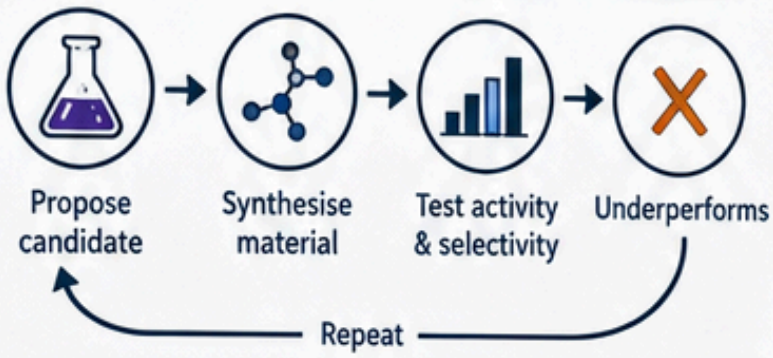

- TIME TO VALIDATED MATERIAL: LONG
- ITERATION COUNT: HIGH
- SUCCESS PROBABILITY PER CAMPAIGN: LOW
- COST (TIME & MATERIALS): SUBSTANTIAL

High cost. Low odds.
Weak competitive position.

VS.

IN-SILICO PRE-SCREENING

- TIME TO VALIDATED MATERIAL: SHORTER
- ITERATION COUNT: REDUCED
- SUCCESS PROBABILITY PER CAMPAIGN: HIGHER
- COST (TIME & MATERIALS): LOWER

Faster discovery. Higher odds.
Stronger competitive position.

THE PLATFORM EFFECT:

WHY NITROGEN FIXATION IS THE BEACHHEAD, NOT THE DESTINATION

The simulation stack that solves FeMoco applies to a lot of other problems too.

The same underlying capability — accurate ground-state energy calculations for strongly correlated molecular systems — applies across several categories of commercially important chemistry:

CO_2 REDUCTION CATALYSTS.

Turning CO_2 into useful chemical feedstocks requires catalysts with similarly complex transition-metal active sites. The quantum advantage argument is structurally identical to the nitrogen fixation case.

BATTERY ELECTROLYTES.

The behaviour of lithium and sodium salts in next-generation battery electrolytes depends on precisely the kind of ion-solvent interactions that DFT struggles to model accurately.

Quantum simulations of battery electrolytes using VQE-based methods have been actively developed, targeting the accurate modelling of $LiPF_6$, $NaPF_6$, and related salts critical to next-generation batteries.

PHARMACEUTICAL CATALYSIS.

Drug metabolism in the body depends heavily on cytochrome P450 enzymes, which contain iron centres with the same kind of complex spin dynamics as FeMoco. AstraZeneca has engaged in quantum computing collaborations for computational chemistry in drug synthesis — though readers should consult AstraZeneca's own published work and press releases for the precise scope and partners of specific programmes, as secondary summaries of these collaborations vary in detail.

McKinsey has estimated the broader quantum computing opportunity in life sciences could reach hundreds of billions of dollars by the mid-2030s, though such projections cover all quantum computing applications in the sector, not quantum chemistry alone, and carry significant methodological uncertainty inherent to long-range technology forecasts.

> ***For a company building quantum chemistry simulation capability, nitrogen fixation is a credible, high-profile demonstration of the technology — but the addressable market extends well beyond ammonia.***

WHAT THE TIMELINE ACTUALLY LOOKS LIKE

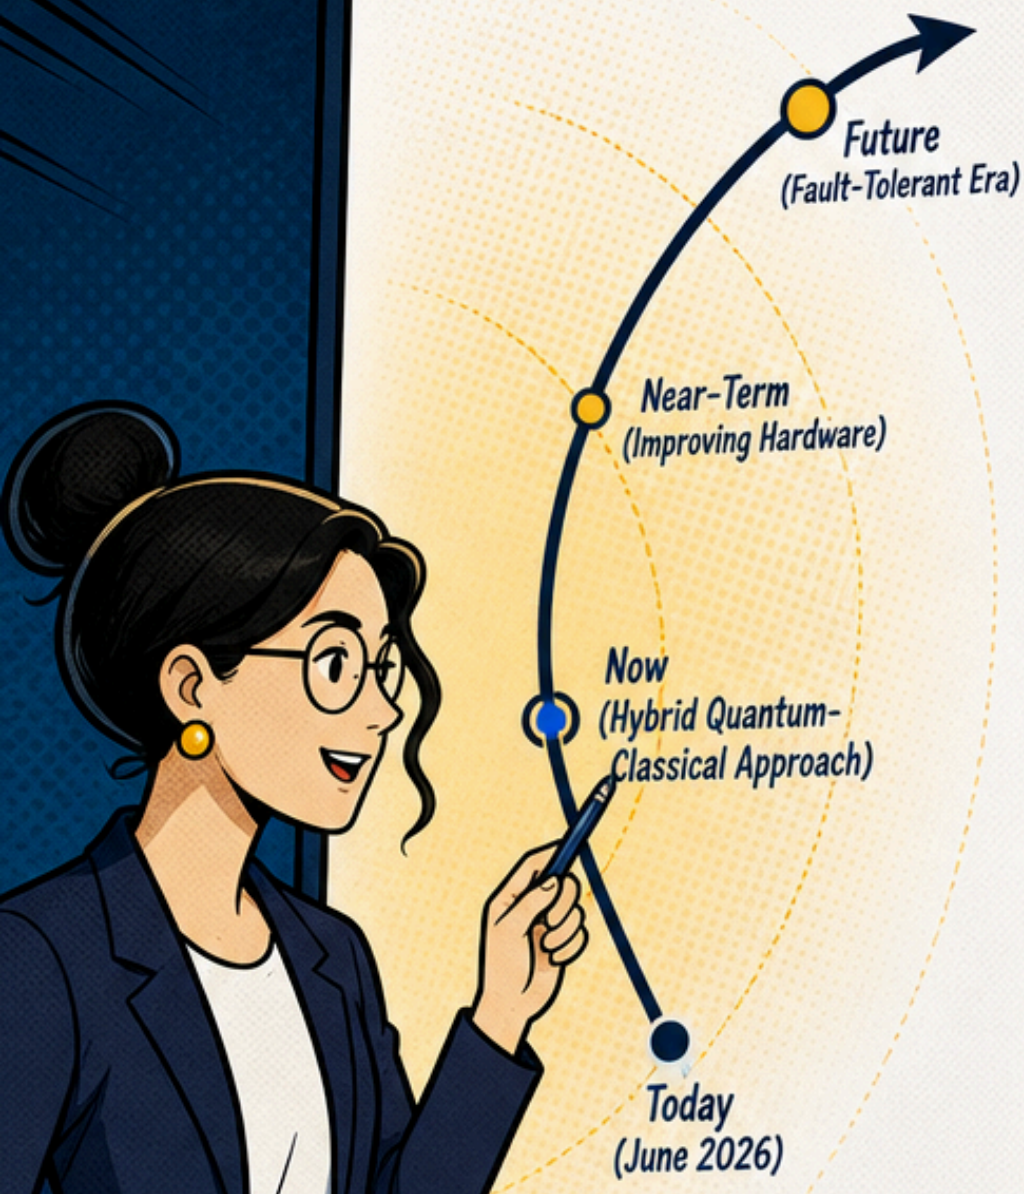

It's worth being honest about where the technology sits in June 2026.

Fully fault-tolerant quantum computers capable of simulating FeMoco at the precision needed for definitive catalyst design don't yet exist in commercially accessible form. The algorithms are well-defined. The hardware roadmaps from major players are credible. The resource estimates have improved substantially over the past three years. But near-term quantum hardware still operates with error rates that limit circuit depth.

What does exist, and what is genuinely useful now, is the hybrid quantum-classical approach: **using quantum processors to handle the strongly correlated regions of a molecular simulation while classical algorithms address the remainder.**

This division of labour is already producing results on real molecules. Recent research has demonstrated hybrid methods — combining techniques such as Density Matrix Embedding Theory with quantum diagonalization — achieving agreement with classical reference benchmarks within chemically relevant thresholds for small test systems on current-generation hardware. It is important to note that agreement with a classical benchmark is not the same as agreement with experimental truth; classical methods themselves carry errors that can compound. Nonetheless, demonstrating that quantum-classical hybrid workflows can reproduce established reference results is a meaningful step toward practical utility.

For investors evaluating companies in this space, the relevant questions are: which organisations have built simulation software stacks that will run on fault-tolerant hardware when it becomes available? Which teams understand the chemistry well enough to know which molecular targets are worth solving? And which have built the hybrid workflows that deliver value on today's hardware while the full capability matures?

The comprehension gap in quantum chemistry investment

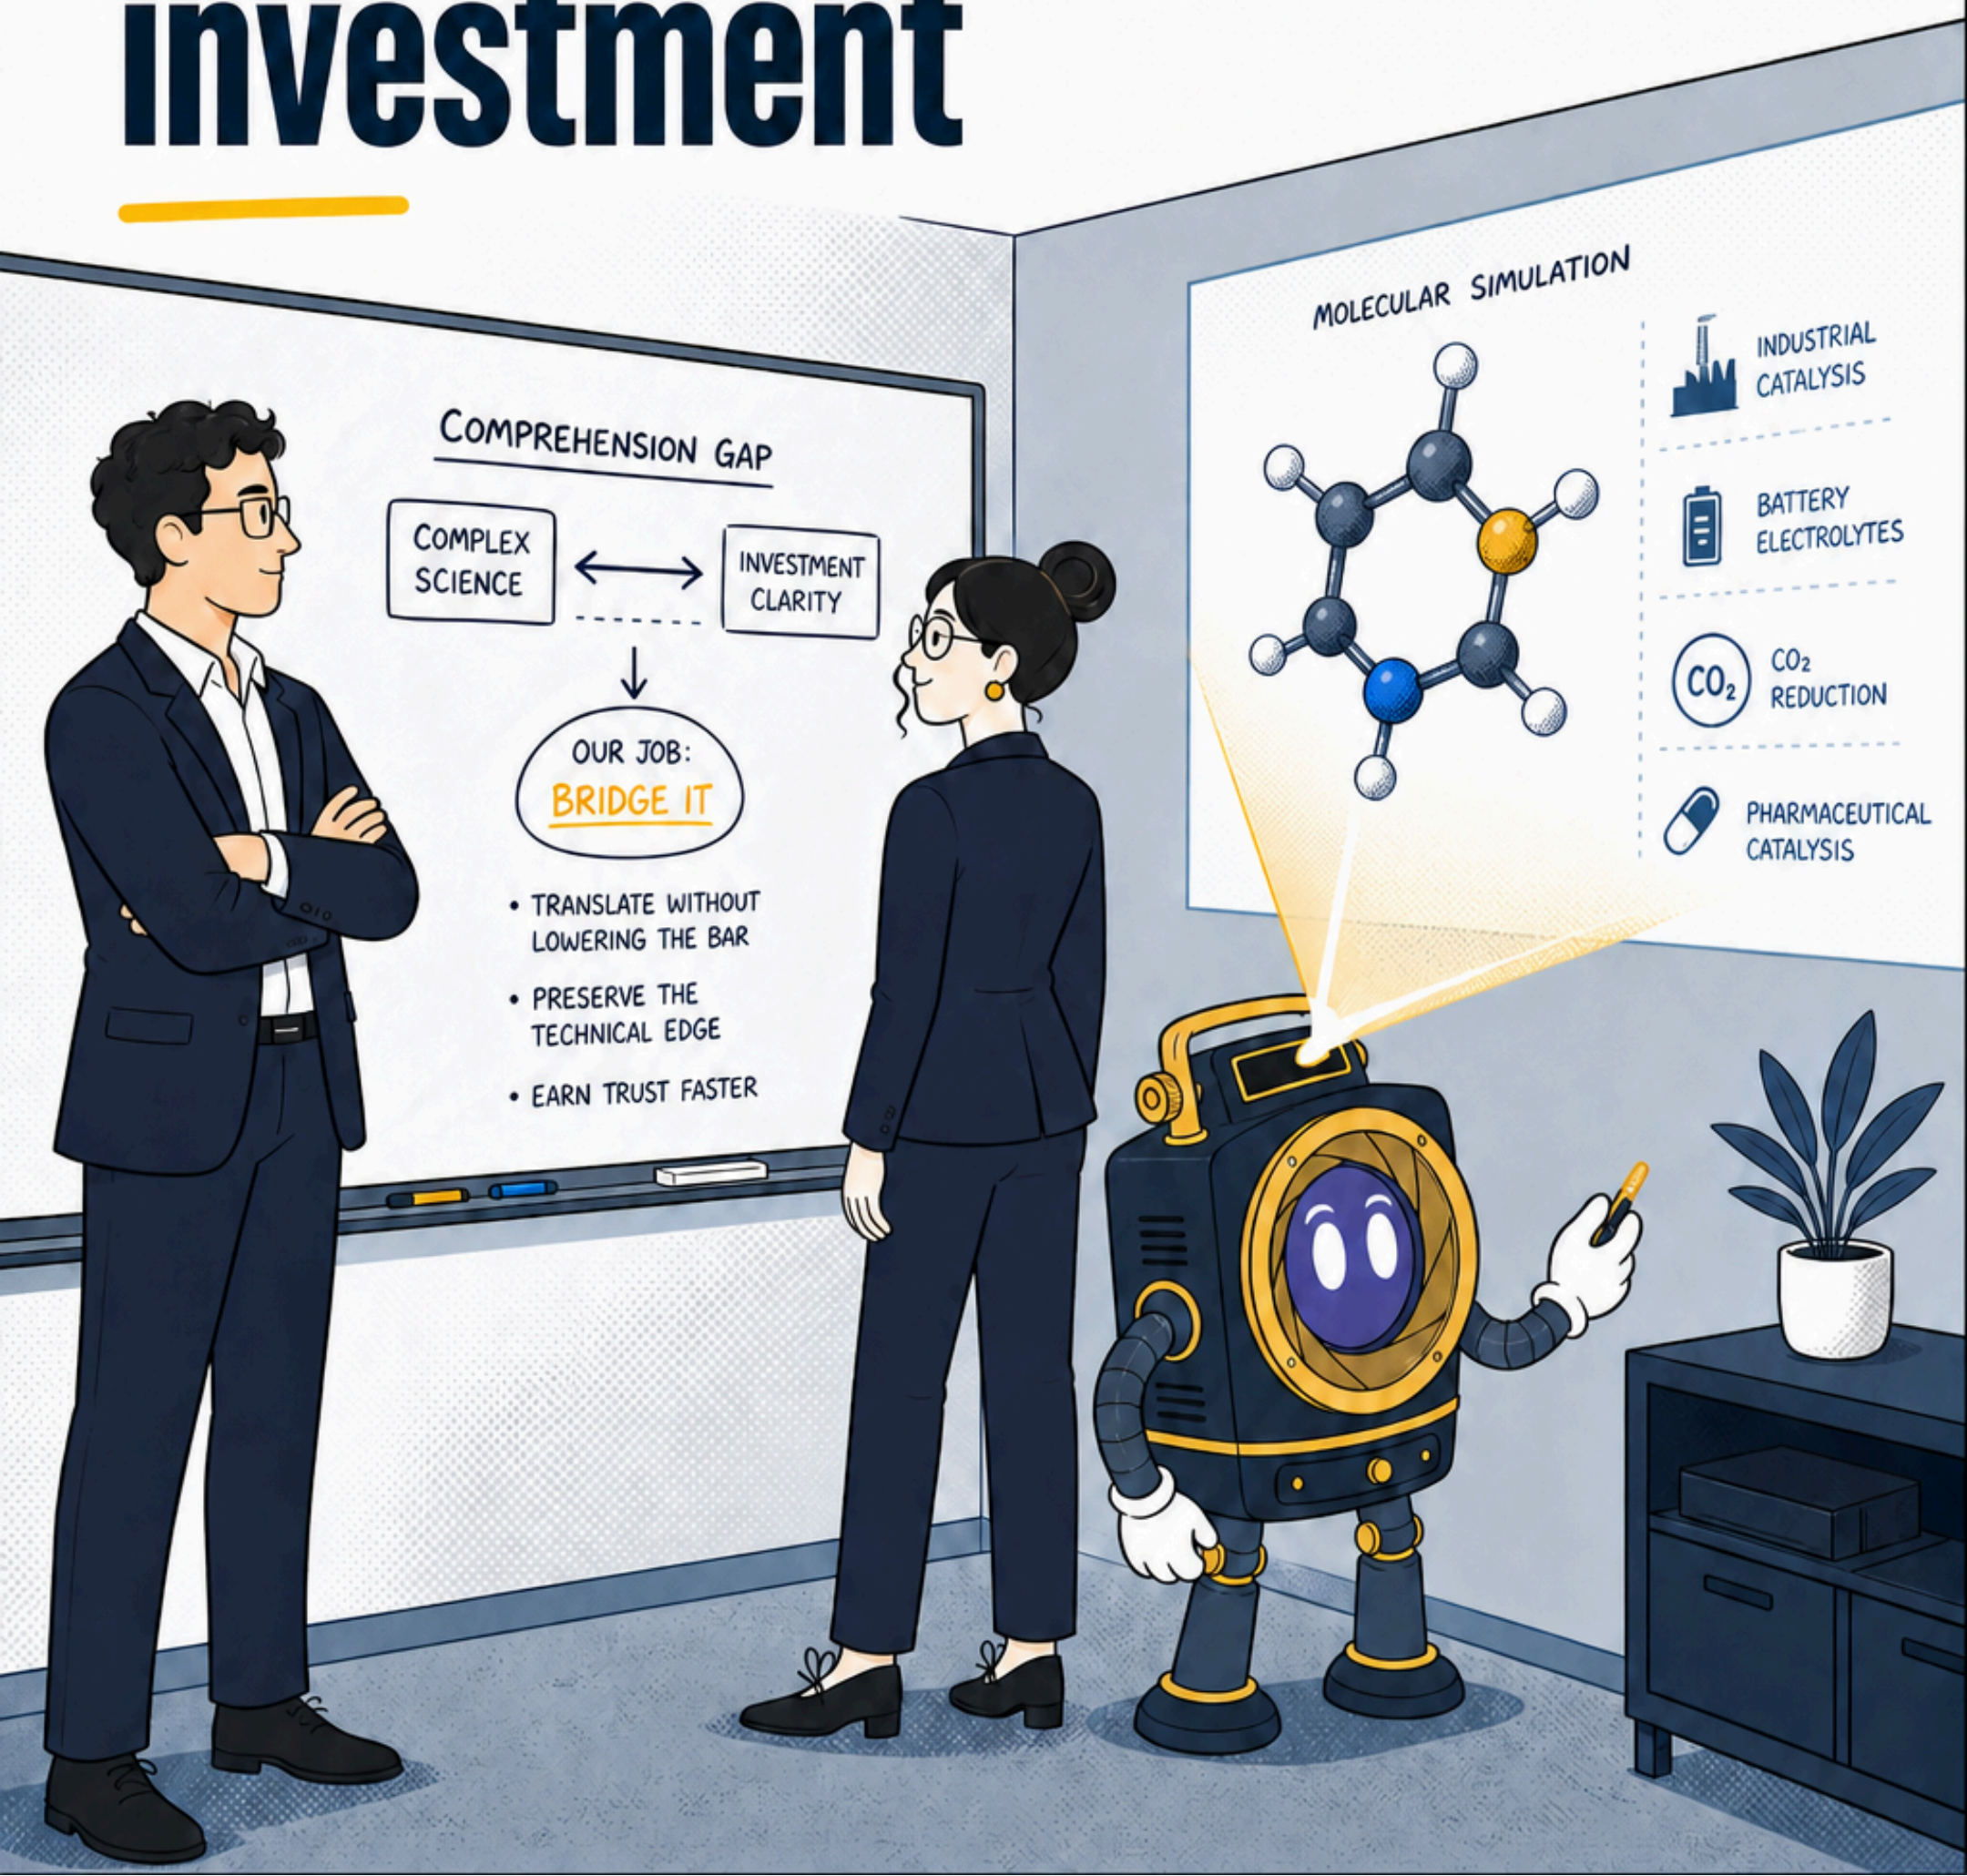

infrairis

https://startups.infrairis.com/