Why Your Gut Is Right About Chess and Wrong About Your Next Acquisition

The Opening Scenario

A chess player with twenty years at the board glances at a position for three seconds and knows, with something close to certainty, that the knight sacrifice works. He could not necessarily tell you why in real time. He just knows. And he is right — not occasionally, but at a rate that has been measured, tested, and confirmed across thousands of positions. That instinct is not luck. It is twenty years of fast, unambiguous feedback compressed into a single glance.

A CEO with twenty years running a company looks at an acquisition target for three weeks, feels the same kind of certainty, and approves the deal. He cannot fully articulate why either — the synergies feel real, the management team feels strong, the timing feels right. He is also drawing on twenty years of experience. But here the confidence is not a measured skill. It is a feeling with no track record behind it, because in capital allocation, nobody ever finds out whether their gut was right until years later, if ever.

Both instincts feel identical from the inside. One of them is reliable. The other is not. The difference has nothing to do with the intelligence, diligence, or seriousness of the person doing the judging. It has to do with the structure of the domain the judgment is being applied in.

Experience Is Not Expertise

These two words get used interchangeably, and the conflation is expensive. Expertise is a predictive model — built from repeated exposure to a problem with feedback fast enough and clear enough to tell you when the model is wrong, so the model gets corrected. Experience is simply repeated exposure. It produces familiarity, fluency, and confidence. It does not automatically produce a model that predicts outcomes, because correction requires knowing you were wrong, and knowing requires feedback that arrives in time to matter.

A CEO who has run twelve capital allocation decisions over a career has had twelve data points, most resolved years after the decision was made, most contaminated by everything else that happened to the business in the interim, and almost none accompanied by a clean counterfactual showing what would have happened with a different decision. Twelve data points spread across fifteen years, each one ambiguous, is not a track record. It is a set of stories the mind has organized into a feeling of competence. The feeling is genuine. The predictive power behind it is not established by the experience itself — it has to come from somewhere else, or it doesn't exist.

This is not a statement about any individual CEO's ability. It is a statement about what experience can and cannot manufacture on its own, in any domain, for anyone.

The Five Questions That Actually Separate Reliable Gut from Unreliable Gut

Some intuition is earned. Some only feels earned. Five questions reveal the difference, and they have nothing to do with how confident the feeling is or how long the person has been doing the work.

Does the environment hold still long enough to learn from it? Chess rules never change. A market a company competed in five years ago may not resemble the market it competes in today — new entrants, new technology, new regulation, new customer behavior. Skill built under yesterday's conditions does not transfer cleanly to today's, because the thing being learned keeps moving.

Do equally experienced people agree with each other? Put ten grandmasters in front of a position and most will converge on the same handful of strong moves. Put ten experienced CEOs in front of the same acquisition memo and you will get ten different recommendations, several of them confident and mutually exclusive. When credentialed experience produces convergence, that is a sign the domain has a learnable structure. When it produces scatter, the confidence each person feels is not evidence the domain is learnable — it is evidence of how the domain feels from inside, which is a different thing entirely.

Is the feedback fast and is it clean? A chess player loses a game in an hour and knows immediately, unambiguously, which move cost it. A CEO who approves an acquisition finds out whether it worked in three to five years, by which point the market shifted, two key executives left, a competitor made an unrelated move, and the original thesis is now buried under a decade of confounding variables. Fast, clean feedback builds a model. Slow, contaminated feedback builds a narrative.

Can a mistake be corrected, or is it final? A bad move in chess costs a game. A bad acquisition costs the capital, the management time, the credibility with the board, and frequently the years required to unwind it — and sometimes it cannot be unwound at all. Domains that tolerate mistakes let people calibrate through trial and error. Domains where mistakes are largely irreversible never let anyone build that calibration, no matter how many decisions they make.

Is there a structured reference point, or only judgment? A weather forecaster does not predict tomorrow from intuition alone — there is a model, a base rate, decades of pattern data behind every forecast, and the forecaster's judgment refines that structure rather than replacing it. Most capital allocation decisions have no equivalent. There is a financial model built on assumptions nobody stress-tested against history, and then there is judgment, unsupported by anything systematically extracted from how similar decisions have actually turned out elsewhere.

Run a major acquisition, an expansion bet, or a debt restructuring through these five questions and the pattern is not subtle. The environment keeps shifting under the decision. Equally experienced executives reach opposite conclusions on the same deal. Feedback arrives years late through a fog of confounding events. Mistakes are expensive and often permanent. And in most companies, the closest thing to a structured reference point is a spreadsheet built on assumptions, not on documented precedent.

Capital allocation does not score poorly on one or two of these conditions. It scores poorly on nearly all of them. That is not a coincidence, and it is not a comment on the people making the decisions. It is what the terrain looks like.

This Is Not a Competence Problem

It would be easy to read the above as a case for hiring smarter executives, or for the current CEO to trust their gut a little less and think a little harder. Neither response addresses the actual issue.

The chess player's intuition does not work because he is smarter than the CEO. It works because chess is a closed, stable, fast-feedback environment that rewards pattern accumulation with genuine predictive power. The CEO's intuition does not fail because he is less rigorous, less experienced, or less capable than the player. It fails because capital allocation is an open, shifting, slow-feedback environment that does not reward pattern accumulation the same way — no matter how rigorous, experienced, or capable the person applying it happens to be.

Put the most brilliant executive alive into this environment and the structural conditions do not change. The domain is hostile to unaided judgment regardless of who is doing the judging. This distinction matters because the wrong conclusion — that the fix is a smarter or more experienced decision-maker — sends companies looking for talent when what they actually need is a different kind of input into the decision entirely.

Doesn't This Undercut Expert Advice Too?

If gut instinct is unreliable in low-feedback, low-agreement, high-irreversibility domains, a fair question follows immediately: what makes any outside judgment — including advisory input — any more trustworthy than the CEO's own? If nobody gets fast feedback on acquisitions, doesn't that disqualify everyone's intuition equally, advisors included?

It would, if advisory judgment were just another person's gut feeling layered on top of the CEO's. It is not supposed to be that, and when it is, the objection is correct — that kind of advice deserves exactly the skepticism a CEO would apply to their own unaided instinct.

The actual fix for a domain with no fast feedback loop is not to find someone whose gut is supposedly better calibrated than yours. It is to substitute something the domain itself cannot

provide: a documented reference class. Not one executive's memory of twelve ambiguous decisions, but a structured record of how decisions with these specific characteristics have actually played out — across many organizations, many cycles, and enough volume that the patterns separate from the noise. This is precisely what the weather forecaster has and the acquiring CEO usually doesn't: a base rate built from accumulated outcomes rather than a single, slow, contaminated personal history.

This is the function **FPI™** performs in a capital allocation decision. It does not replace judgment, and it does not claim a better gut. It supplies the structured reference class — the documented pattern history — that the domain's slow feedback loop fails to generate on its own. The CEO still makes the decision. The decision is no longer being made on twelve ambiguous data points pretending to be a track record.

The Test That Determines Whether Any of This Is Worth Trusting

None of this is worth anything, including from **FPI™**, unless it can be checked. A pattern claim that cannot fail is not a pattern claim — it is a story wearing the costume of analysis, and the costume gets more convincing, not less, the more confidently it is delivered.

The standard is simple to state and uncomfortable to apply: a documented pattern is only as credible as its ability to specify, in advance, what should happen if the pattern is active — and then to be checked against what actually happened. "Companies that move fast sometimes get into trouble" cannot fail, because it does not say enough to be wrong. "Acquirers running more than three simultaneous integration initiatives with the same five-person leadership team will see decision velocity drop and milestones slip across multiple initiatives within two quarters" can fail. It names a condition, a mechanism, and a timeline specific enough to be checked against the record. A claim's size or the confidence behind it tells you nothing about whether it is true. Its specificity, and its willingness to be proven wrong, tells you almost everything.

This is the question worth asking of any advisor, framework, or pattern claim before trusting it with a capital decision: not how experienced is the person making the claim, but how exposed is the claim itself to being wrong.

Before You Trust Your Gut on the Next One

The next time a major capital decision feels obvious — the acquisition that feels right, the expansion that feels overdue, the restructuring that feels necessary — five questions are worth asking before the capital moves:

- Have I actually seen this exact situation enough times to have a real track record, or does it just feel familiar because I have seen things that resemble it?

- If I described this decision, stripped of names and dollar figures, to another executive with comparable experience, would they reach the same conclusion I just did?
- If I am wrong, how long before I would even find out — and how much else will have happened to the business by then, muddying the answer?
- If this goes badly, is it a setback I can correct, or is the capital and the time simply gone?
- Is there anything informing this decision beyond my own judgment and the model my team built to support the conclusion I already favored?

None of these questions require distrusting your own judgment. They require knowing which kind of terrain you are standing on before you decide how much weight that judgment can safely carry.

Chavis Raynor, PhD *is the founder of Failure Pattern Intelligence™ (FPI™), a proprietary methodology built on 500+ documented failure patterns cross-validated against independent authoritative sources. FPI™ Advisory provides expert intervention for imminent capital decisions. FPI™ Mastery builds systematic pattern recognition capability over six months. chavisraynor.com*
