

THE 2ND PILLAR OF
The Grounded Enterprise

The AI ROI Playbook



*A Leader's Framework for Measuring,
Defending, and **Killing** Enterprise AI
Investments*

Alan Hancock

CHIEF TECHNOLOGY OFFICER EDITION

EXECUTIVE PREVIEW

The AI ROI Playbook

*A Leader's Framework for Measuring, Defending, and **Killing** Enterprise
AI Investments*

ALAN HANCOCK

CHIEF TECHNOLOGY OFFICER EDITION

Preface: The Three Lenses of Enterprise AI Value

Enterprise AI proposals rarely die because the technology failed. They die in the budget committee, and they die for a reason that has almost nothing to do with model quality. The person who kills them is usually the Chief Financial Officer. The CFO is not being a Luddite. The CFO is applying a discipline the pitch never accounted for.

Here is the mismatch that ends most of these projects. When an engineering or product team pitches an AI deployment, they lead with transformation—new capability, competitive advantage, revenue that has not arrived yet. When the CFO evaluates that same pitch, they are looking for something else entirely: hard cash impact, priced-in downside, and a baseline they can audit. Neither side is wrong about the facts. They are looking through different lenses, and the project falls into the gap between them.

We wrote this book to close that gap, and the work starts here. Every dollar of enterprise AI value falls into one of three tiers. The tiers are not interchangeable. They do not carry equal weight in a CFO's mind, and mixing them up is the single most common reason a sound project never gets funded.

The fatal move is easy to make and hard to survive. A team justifies a large efficiency outlay with a story about strategic transformation, never prices the risk the deployment introduces, and hands the committee a single ROI number that answers none of the three questions a CFO is actually asking. The number looks precise. It is answering the wrong question in the wrong tier. The fix is not a better number. It is naming which tier you are in and measuring it on that tier's own terms.

Tier 1: Efficiency and Direct Cost Reduction

Hard cash and reclaimed capacity. This is the tier a CFO trusts most, because it lands on the income statement in language finance already speaks. It counts the labor hours reclaimed when a manual workflow is automated, the spend consolidated when one grounded internal system replaces a shelf of fragmented point tools, and the compute cost stripped out when wasted token overhead is cut. Tier 1 is where you build credibility. A claim here is defensible, measurable, and—if you locked a real baseline first—hard to argue with. It is also the tier most pitches underplay, because it is less thrilling than the alternative and far harder to fake.

Tier 2: Quality, Compliance, and Risk Avoidance

The value of what did not go wrong. This is the tier engineers forget and CFOs never do. In a regulated or high-stakes pipeline, preventing one bad output—a fabricated citation in a legal filing, a wrong figure in a financial disclosure—can be worth more than automating tens of thousands of routine tasks. A single six-figure penalty avoided outweighs a year of steady automation. Tier 2 is difficult precisely because you are pricing an absence. The breach that never happened. The drift you caught before a customer did. The downtime you prevented.

Absence resists a spreadsheet, and that is exactly why undisciplined pitches leave it blank—which is how a headline productivity win quietly becomes a net loss when the ignored risk finally lands.

Tier 3: Strategic Growth and New Capability

Net-new value that did not exist before. This is the tier every pitch leads with, and the value here is genuine—but it is the least defensible in a budget committee, because it trades in possibility rather than measured fact. Tier 3 covers capabilities that were simply unavailable at any price without these systems: running a live query across decades of unstructured legacy contracts, or compressing a proposal cycle from three weeks to a few hours to bid on deals you previously had to walk away from. Tier 3 is where the strategic story belongs, and it is a story worth telling. It should never carry the funding case by itself.

Read the three together and the discipline of this book comes into focus. Tier 1 earns the CFO's trust, Tier 2 protects the firm from the downside that turns wins into losses, and Tier 3 justifies the ambition—but only once the first two have done their work. A pitch that leads with Tier 3 and never touches the other two is not a business case. It is a wish with a slide deck.

So we will hold you to one rule from this page forward. If you cannot map an initiative to at least one of these three tiers, and you cannot audit its progress inside a defined ninety-day window, you should not fund it. Not yet. That rule—tier it, baseline it, or kill it—is the spine that runs through all fourteen chapters, and it is the reason the person holding the checkbook will keep reading past this page.

Introduction: The Denominator Problem

Most explanations for disappointing AI returns blame the technology. The model underdelivered. The use case was wrong. The vendor oversold. Sometimes that is true. Far more often, the technology worked fine and the measurement failed—and it failed in a specific, findable place. It failed in the denominator.

Return on investment is a fraction. Gains on top, cost on the bottom. The entire enterprise-AI conversation has fixated on the top. Vendors project gains, teams promise gains, boards interrogate gains. Meanwhile the bottom of the fraction—what you are truly spending, and what you measure that spending against—goes almost unexamined. That is the error. Most AI business cases are already broken at the denominator before anyone argues about the gains at all.

We can state the thesis of this book in one sentence, and we will spend the rest of it earning the sentence. Enterprise AI ROI fails not because the technology underperforms, but because organizations miscalculate the denominator, misattribute the gains, and keep funding what they should have killed. Three failures, one root. Every one of them is a failure to measure honestly, and every one is fixable with discipline rather than genius. You do not need better technology to fix this. You need better arithmetic and the nerve to act on it.

Miscalculating the Denominator

The denominator holds two numbers, and most business cases get both wrong in the same direction. The first is cost—the real, fully loaded cost of the AI system, which is almost always understated. Model API charges are the visible tip. Underneath sit data pipeline maintenance, evaluation, human review, the observability needed to catch silent failures, and the ongoing labor of keeping the thing working as the world shifts around it. A business case that counts only the visible cost is dividing by a number far too small, which makes the return look far too good.

The second number is the baseline—the pre-AI reality the system is supposed to beat. This is the one enterprises skip almost entirely. If you never measured what a process cost, how long it took, and how good it was before the AI arrived, you have nothing honest to compare against afterward. You end up reconstructing the baseline from memory, and memory flatters the choice you already made. A wrong cost and a phantom baseline compound. Together they can turn a genuinely marginal project into a slide that promises triple-digit returns, and neither error is visible unless someone insists on locking the numbers first. Chapters 3 and 4 are where we lock them.

Misattributing the Gain

The numerator has its own failure mode, and it is subtler because it hides inside numbers that are technically real. A team measures that an AI tool saves time. The time saving is genuine. Then they convert those saved

hours to dollars at a loaded labor rate and book the result as return. The conversion is where the fiction enters. Time saved is not money saved unless something structural changes—a role not backfilled, a contract ended, additional revenue-generating work absorbed on the same cost base. Absent that change, the hours dissipate into slightly easier days and the income statement never moves.

There is a second misattribution that runs alongside the first. A metric improves after a deployment, and the improvement gets credited entirely to the AI, when other things changed at the same time—a process redesign, a seasonal shift, a team that simply got more practiced. Correlation gets booked as causation, and the AI takes credit for gains it only partly caused. Honest attribution means isolating what the tool actually moved and naming the structural change that turned motion into money. We build that discipline in Chapter 3. Until you have it, your numerator is a hopeful estimate wearing the costume of a measurement.

Refusing to Kill

The third failure is not arithmetic. It is nerve. Even organizations that measure cost and gain honestly tend to keep funding projects that are not working, because killing a project feels like admitting failure and extending it feels like patience. So pilots linger. Deployments that stopped returning value months ago stay in the budget because no one wants to be the person who ends them. The portfolio silts up with the undead, each one a small, permanent drain that no single meeting ever decides to stop.

A disciplined portfolio treats killing as a routine, healthy act—capital protected, attention freed, the system working exactly as designed. That posture is impossible to hold after the fact. It has to be built in before a project starts, as explicit exit criteria agreed upon while everyone is still calm and no one is defending their reputation. The willingness to kill early is what separates a governed portfolio from an expensive museum of good intentions. Chapter 8 gives you the mechanics. The mindset you will have to supply yourself, and this book will keep pushing you toward it.

Who This Book Is For

We wrote this for the person who signs, or who has to defend the signature. If you hold budget authority over enterprise AI spending, or you stand in front of a board and account for it, this book is aimed directly at you. Much of it will also serve the technical leaders who build the business cases, because a case built to survive the scrutiny described here is simply a better case. But the primary reader is the one whose name is on the decision.

That focus shapes the tone. We are not going to sell you on AI, and we are not going to talk you out of it. Both postures are lazy. The useful stance is the one a good CFO already takes toward every other use of the firm's capital: genuine openness paired with an insistence on evidence. You can be excited about this technology and still refuse to fund a project that cannot name its baseline. In fact, that combination is the only one that survives a serious budget committee, and it is the combination this book is built to give you.

How the Book Is Built

The five parts trace a single path, from why this matters to what you do Monday morning. Each part assumes the one before it.

Part I, which begins on the next page, makes the macro case and delivers the diagnosis—why standing still is its own expensive bet, and the three ways measurement fails once you start moving. Part II builds the financial engine: the full-stack cost ledger, the taxonomy that converts vague benefit into hard accounting entries, and the risk adjustment that prices what could go wrong. This is where the three lenses from the preface acquire real machinery. Tier 1 efficiency and Tier 2 risk avoidance both live here, in numbers you can defend.

Part III installs the operating system—how you intake and rank competing proposals, gate them through a lifecycle with real kill points, and measure whether anyone is actually using what you deployed. Part IV turns the framework loose on specific domains: engineering, operations and support, and knowledge work, each with its own metrics and its own traps. Part V brings it home to the executive: how to defend an AI budget to a board in language finance respects, and a ninety-day plan that converts everything preceding it into action you can start immediately.

Read straight through, the parts compound. Read out of order, most will still stand on their own, because we have tried to make each artifact usable the day you reach it. What none of them will do is work without the discipline underneath. The tools are only as honest as the baseline you feed them, which is why the argument starts where it does—with the cost of doing nothing, and the measurement failures that make even decisive action invisible on the income statement.

Keep reading.

You've just read the thesis. The full playbook turns it into machinery — fourteen chapters that build the cost ledger, price the risk, install the kill criteria, and hand you a ninety-day plan you can start Monday.

ORDER THE FULL PLAYBOOK →

Part of a two-book method. Its companion, **The Grounded Enterprise** by Duane Cash, shows you how to *build* enterprise AI that deserves to be trusted; this book shows you how to prove the spend was worth it.

TheGroundedEnterprise.com

Alan Hancock is a technology leader who has delivered more than \$180M in solutions to Fortune 500 companies and government agencies.