ORTENT WHITEPAPER

Fund Learning, Not Faith

A CFO-approvable way to back AI-enabled revenue.

By Andrew Wyatt, Founder, Ortent Advisory Ltd

September 2026 • A diagnostic for CFOs, PE operating partners, and chairs

Contents

Executive summary

Every board of a growth-stage SaaS business is having a version of the same conversation this quarter. The CEO has an AI-enabled revenue idea. The CFO cannot fund it because the business case does not compute. The chair asks whether the alternative to funding is to do nothing, and everyone in the room knows the answer to that is worse.

The pattern is not scepticism. It is category confusion. The business case template every CFO has spent twenty years building is designed to value efficiency investments. It computes ROI on an activity you are already doing and are proposing to do more cheaply. It cannot value a revenue play whose evidence base does not yet exist.

> *The board approves what the business case shows. The business case shows efficiency AI. The revenue AI dies in committee before it is ever built.*

This paper gives CFOs and boards a way to fund AI-enabled revenue without abandoning the discipline that made them CFOs and board directors in the first place. The move is not to run bigger numbers through the same template. It is to fund learning, not faith. Stage the investment into four decision points. Attach kill criteria and evidence gates to each. Approve the next stage on evidence produced by the previous one.

The mechanic is the Reinventors Playbook. Four revenue patterns that AI new-revenue plays actually fall into. A six-axis lens that qualifies which pattern applies. A four-stage funding ladder (Frame, Probe, Pilot, Scale) with defined evidence gates and kill criteria between each stage. The CFO funds a Probe on twelve weeks and £75K. The Probe either clears its evidence gate or it does not. If it clears, Pilot is approved on the strength of the Probe data, not on the strength of the CEO's belief. If it does not clear, the kill criteria the CFO wrote in month one make the wind-down a clean decision rather than a political one.

The paper is not about approving more AI investments. It is about approving them the right way, with the right evidence, at the right cost. Growth-stage businesses that fund AI-enabled revenue this way do not spend less. They spend the same amount and learn faster. They kill the losers earlier and scale the winners with confidence. The eventual exit values reflect that discipline.

The killer instrument sits in Section 3. The four patterns, the six-axis lens, and the four-stage ladder together. Read it. Score the AI new-revenue idea currently on your board deck. Come back to Section 4 for how to run each stage.

Why the problem exists

Two things are true at the same time in every growth-stage SaaS business this year, and the tension between them is where the AI revenue conversation dies.

The first thing that is true is that everyone can now cut cost with AI. Customer support ticket handling, marketing content, sales-engineering documentation, code review, meeting summaries. The efficiency wins are real, the ROI is defensible in an eight-week payback, and the CFO can model the numbers on a spreadsheet the CEO recognises. That is the AI investment that keeps getting approved.

The second thing that is true is that almost nobody in the growth-stage SaaS market is making new money from AI. Not because the opportunity does not exist. The opportunity is larger than the efficiency wins by an order of magnitude. But new-money AI initiatives require investment on the strength of a hypothesis, not a proof. The business case template every CFO has spent twenty years refining does not know how to value a hypothesis. It knows how to value an activity you are already doing and are proposing to do cheaper.

> *So new-money AI dies in the business case. The CFO cannot approve it in good faith because the evidence base does not exist. The CEO cannot produce the evidence base without the funding.*

The chair, if the chair is honest, watches the same conversation happen quarterly and does not have the language to break the loop. The gap is not risk appetite. Growth-stage SaaS

chairs have plenty of risk appetite. It is disciplined risk appetite. The gap is that no board pack currently in circulation contains an instrument that lets a chair approve staged risk against evidence gates, rather than binary risk against a single business case.

This paper is that instrument.

The other pattern worth naming. The CEOs who eventually make new money from AI are not the ones with the biggest AI budgets or the flashiest partnerships. They are the ones who structured their AI new-revenue investments as a portfolio of small, staged bets, each with a kill criterion. They spent the same aggregate money as their competitors. They just spent it in the right sequence. Their competitors funded one large bet on the strength of a single conviction and either won it or lost it. The Reinventors funded four small bets, killed two of them fast on evidence, and pushed the two that cleared into serious scale investment.

That difference in structure is the difference between a business that catches the AI revenue wave and a business that watches it pass while efficiency numbers grow.

The Reinventors Playbook is not a strategy document. It is a funding discipline. The CFO owns it. The CEO reports against it. The board reviews it quarterly. Ninety minutes of leadership time when a new AI revenue idea is proposed. Fifteen minutes at each subsequent stage gate. That is the ask.

The Reinventors Playbook

> *Read this section before anything else. Score your current AI new-revenue idea here. Then read Section 4 to understand what the score means.*

The Playbook has three components. Read them in order. They compose.

Component one. The four revenue patterns.

AI new-revenue plays fall into four patterns. Pick the pattern the current idea fits before doing anything else. If the idea does not fit any of the four cleanly, that is your first useful piece of information: the idea is under-specified.

Productised Expertise. The business currently sells time from its most senior experts. AI encodes what those experts know and turns it into a product a much larger population can buy at a fraction of the cost. Examples: legal advisory turning senior-partner review into an AI-assisted subscription service, medical diagnostic firms turning consultant time into decision-support software, tax advisory turning specialist review into workflow tools. The pattern's economics improve as more customers use it because the expert layer is amortised.

New Product Surface. The business currently has a product with a defined surface area. AI adds a new capability, a new module, or a new adjacent product that customers can buy on top of the existing product. Examples: a CRM adding an AI meeting-summary product, a compliance platform adding an AI regulatory-change monitor, an EHR system adding an AI clinical-documentation assistant. The pattern's economics improve as the existing customer base opts in.

New Pricing Model. The business currently prices on seats, tiers, or usage. AI enables a new pricing structure that captures value differently: outcome-based, workflow-based, or transaction-based. Examples: a legal-tech SaaS moving from seat pricing to outcome-based pricing on a matter-count basis, a customer-service platform moving from tier pricing to cost-per-resolution, a marketing platform moving to attribution-based revenue-share. The pattern's economics improve because the price rises to meet the value delivered.

New Reachable Segment. The business currently serves a defined segment because the delivery cost per customer sets a floor. AI lowers delivery cost enough to make a previously unreachable segment economically viable. Examples: an enterprise-focused SaaS becoming SMB-viable through AI-driven implementation, a specialist services firm entering a mid-market segment through AI-enabled delivery, a compliance product entering non-regulated adjacent industries through AI-driven onboarding. The pattern's economics work

because the new segment's aggregate value exceeds the additional cost.

Every AI new-revenue idea worth funding fits one of these four patterns. If an idea seems to fit two, it is two ideas, not one, and each should be funded separately. If an idea does not fit any, it is not yet a revenue idea. It is a research hypothesis. Fund it as research, not as revenue.

Component two. The six-axis qualifying lens.

Once you have the pattern, score the idea on six axes. Each axis gets a 1-to-5 score. Four of the axes measure fit. Two measure risk. The total is what determines which stage of the funding ladder the idea enters at.

Axis 1. Value (fit). How large is the value the idea creates for the customer, per unit of usage or per contract, compared to the customer's current alternative? A score of 5 means the customer would happily pay materially more than the current alternative. A score of 1 means the customer would only switch if the price were lower.

Axis 2. Cost-to-Serve (fit). How favourable is the unit economics of delivery once the AI capability is in place? A score of 5 means margin is meaningfully better than the business's current product margin. A score of 1 means the margin is worse than current product margin and improves only at implausible scale.

Axis 3. Willingness-to-Pay (fit). How ready is the target buyer to pay for the value created, given their current budget lines and buying process? A score of 5 means there is an existing budget line the offering slots into directly. A score of 1 means the buyer would need to create a new budget category and defend it internally.

Axis 4. Reach (fit). How large is the addressable customer population who would actually buy the offering at a plausible price? A score of 5 means the reachable population exceeds ten times the number of customers required for the business case. A score of 1 means the reachable population is smaller than the business case requires.

Axis 5. Regulation (risk). How much regulatory or compliance work stands between the current state and the offering being sold to a paying customer? A score of 5 means no material regulatory friction. A score of 1 means the regulatory work is longer than the funding ladder itself and is unlikely to clear inside two years.

Axis 6. Reversibility (risk). How reversible is the investment if the offering fails? A score of 5 means the investment produces reusable assets, capabilities, or IP that transfer to other initiatives. A score of 1 means the investment is bespoke to this offering and produces nothing salvageable on failure.

Add the six scores. Total is between 6 and 30. That total determines the starting stage on the funding ladder.

- **6 to 12: Frame only.** The idea is under-specified. Fund three weeks of framing work by two people. If the score improves through framing, the idea can enter Probe. If it does not, kill it.
- **13 to 18: Probe.** The idea is credible enough to test the riskiest assumption with a small, cheap experiment. Fund twelve weeks and £50K to £100K.
- **19 to 24: Pilot.** The idea has a strong hypothesis but not yet a paying customer. Fund six months and £200K to £500K to build a functional first version, sign a paying pilot customer, and prove the unit economics.
- **25 to 30: Scale.** The idea has evidence from a prior deployment or from a comparable market that removes most of the hypothesis risk. Fund the full build.

The scoring is not designed to be precise. It is designed to force a conversation about which assumptions are hypotheses and which are evidence, and to make the CFO comfortable that the starting funding matches the actual state of evidence, not the CEO's conviction.

Component three. The four-stage funding ladder.

Each stage has a defined activity, a defined evidence gate, and a defined kill criterion. The next stage is only approved if the evidence gate for the current stage clears. The kill criterion is written on the day funding is approved, before any work starts. That timing matters: kill criteria written after evidence exists tend to move to protect the investment. Kill criteria written before evidence exists hold.

> STAGE One
> **Frame**

Three weeks, two people, £15K to £25K. Define the customer, define the value, define the price, name the six-axis assumptions that need testing. Output is a written brief the CFO and CEO agree on. Evidence gate: score has improved from initial Playbook run. Kill criterion: score does not improve, or the framing exposes an assumption that cannot be tested inside a Probe budget.

STAGE Two

Probe

Twelve weeks, three to five people, £50K to £100K. Test the single riskiest assumption from the Frame stage. This is not a working product. It is a targeted experiment. If the riskiest assumption is willingness-to-pay, the Probe is a paid pilot conversation with three prospects. If the riskiest assumption is cost-to-serve, the Probe is a technical prototype that stress-tests the delivery model on real inputs. Evidence gate: the specific assumption tested is either validated or invalidated with defensible data. Kill criterion: the assumption cannot be tested inside the Probe budget, or the test invalidates it.

STAGE Three

Pilot

Six months, seven to ten people, £200K to £500K. Build a functional first version, sign one paying customer, run it in production, and prove the unit economics. This is the first stage where a real product ships to a real customer. Evidence gate: the paying customer renews or expands, unit economics are within 20% of the model, and the second sales conversation is measurably easier than the first. Kill criterion: the paying customer does not renew, or the unit economics are more than 40% worse than the model, or the second sales conversation is materially harder than the first (which usually means the value story does not travel).

STAGE Four

Scale

Twelve to twenty-four months, funding sized to the model. Build the go-to-market, scale the customer base, and grow the revenue line. This is the stage where the offering becomes a real revenue line of the business. Evidence gate is now the normal board KPI review: ARR growth, gross margin, net revenue retention. Kill criterion becomes strategic rather than tactical: does the revenue line still fit the strategy, and is it hitting the trajectory the model required.

The discipline of the ladder is that stage gates are non-negotiable. A Pilot that has not cleared its evidence gate cannot proceed to Scale. The CFO holds the gate. The CEO does the work. The chair reviews the gate decisions quarterly. This is not a bureaucratic constraint. It is the mechanism that turns AI revenue investment from a single act of faith

into a portfolio of learning.

The framework in depth

Section 3 gave you the instrument. This section explains how to actually run each component.

Which of the four patterns should you pick?

The pattern is not a matter of preference. It is a matter of what the AI capability actually enables in your specific business. Two tests.

Test one: what does the AI capability do that your current product does not? If the AI capability encodes expertise your business already sells at high price, you are looking at **Productised Expertise**. If it adds a new function to your existing surface, you are looking at **New Product Surface**. If it lets you charge for value your current pricing structure does

not capture, you are looking at **New Pricing Model**. If it lowers your delivery cost enough to serve a previously unreachable customer, you are looking at **New Reachable Segment**.

Test two: if you were to sell the offering tomorrow, what would the customer buy? A subscription to expert knowledge? A new module on an existing platform? A different way of paying for what they already buy? A version of your product they previously could not afford? The answer names the pattern.

Businesses that struggle to pick a pattern usually have an idea that is trying to be two things at once. That is your first useful diagnostic. Split the idea into two. Fund each separately. Kill the weaker one earlier.

How to score the six axes honestly.

The six-axis lens works if the scoring is honest. Two ways to make it honest.

First, do not score alone. The CEO scores the idea, the CFO scores the idea, one other senior operator scores the idea. Do it independently, without conferring. Compare. Where the CEO scored 5 and the CFO scored 2, that axis is where your view of the state and the state itself have diverged. That is the axis to test in Probe.

Second, if you cannot cite specific evidence for a score of 4 or 5, downgrade to 3. Evidence means something you can point to: a customer conversation with a named person, a comparable market data point from a public source, a technical experiment result from your own team. If the score depends on your belief that the number is right, the score is a 3 at best.

The scoring exists to make the funding decision defensible in a board meeting six months later, when the Probe has either cleared or failed. If the scoring was honest and evidence-based on day one, the outcome is defensible either way: the Probe cleared and the decision was right, or the Probe failed and the kill criteria the CFO wrote on day one make the wind-down clean. If the scoring was based on conviction and belief, the outcome is a political problem regardless of which way the evidence lands.

How the funding ladder actually works in practice.

Three failure modes that break the ladder discipline, and how to avoid them.

Failure mode one: Frame gets skipped. The CEO is enthusiastic, the CFO is under pressure, the board wants to see action. The team goes straight to Probe. What happens next is that the Probe wastes twelve weeks testing the wrong assumption because Frame was the stage where you were meant to identify which assumption was riskiest. Avoid this failure

mode by treating Frame as mandatory. It is three weeks and £25K. The cost of skipping it is a wasted Probe.

Failure mode two: Pilot drift. The Pilot signs a paying customer at a discount that makes the unit economics work on paper but not in reality. The customer renews because the discount is baked in. The evidence gate technically clears. Scale is approved. Twelve months later the offering discovers that no other customer will pay the un-discounted price and the business has a scale problem that only becomes visible at scale. Avoid this failure mode by writing the evidence gate on the day the Pilot starts, and specifying that the paying customer must be at a price the CFO would defend as the unit price at scale.

Failure mode three: Scale without kill criteria. By the time an offering reaches Scale, the executive team is emotionally committed. Nobody wants to write down what would cause them to stop investing. That is exactly why the kill criteria for Scale need to be written before the Scale funding is approved, and reviewed quarterly. The kill criteria at this stage are usually strategic (does this fit the strategy) or financial (does this hit the model trajectory), not tactical. But they need to exist and be reviewed. Businesses that skip this stage produce zombie revenue lines that consume capital for years past their expiry date.

How much does the whole ladder cost?

For a typical growth-stage SaaS business at £20M to £50M ARR considering one AI new-revenue idea, the ladder from Frame to end of Pilot costs between £250K and £625K over roughly nine months. Scale funding is sized to the specific model. That is the total commitment for one idea taken through disciplined stage gates.

Compare that to the alternative pattern. A business that funds an AI new-revenue idea as a single approval on the strength of a business case typically commits £1M to £2M in year one on a plan that has not been evidence-tested. If the plan works, the outcome is similar. If it does not, the difference is a large write-down eighteen months in versus a small kill decision three months in.

> *The staged model is not slower. It is faster to signal, faster to kill, and faster to redirect. Businesses that run the staged model on three ideas simultaneously will typically kill two and scale one inside twelve months, and the one that scales will have twenty months of head start over a competitor that ran a single unstaged bet.*

Applied cross-industry examples

Three examples, three different revenue patterns, three different Playbook outcomes. All names and specifics generalised. The patterns are drawn from actual engagements.

Example one. Legal-tech SaaS at £18M ARR. Productised Expertise pattern.

The business sold time from senior partner reviewers to law firm clients on a fixed-fee-per-review basis. The CEO wanted to turn the senior-partner review into an AI-assisted service that a much larger population of law firms could buy at a subscription price, dropping the cost per review to something the mid-market could afford.

The Playbook run scored the idea at 22 out of 30. Value axis 5 (customers would pay materially more than the current alternative at the mid-market price point). Cost-to-Serve 4 (unit economics improve as more firms use it). Willingness-to-Pay 3 (existing budget lines for compliance and review would need to expand). Reach 4 (the mid-market law firm population is roughly twelve times the current customer count). Regulation 3 (some professional-services regulatory work required). Reversibility 3 (the AI capability transfers to related legal offerings if this one fails).

Starting stage: Pilot. The scoring was strong enough that Probe was not necessary, but not strong enough to go straight to Scale.

The Pilot ran for six months and £320K. Signed one mid-market law firm as the paying pilot customer at the price the CFO had modelled for scale. Unit economics landed within 12%

of the model. Second sales conversation was measurably easier than the first (the pilot customer served as the reference for the next three deals). Evidence gate cleared.

Scale was approved. Two years later, the offering was a £8M ARR revenue line growing at 60% year-on-year. Gross margin was materially better than the business's core product. The offering became the primary growth driver of the eventual exit valuation eighteen months after that.

The counter-factual is worth noting. Had the business funded this as an un-staged £2M investment, the outcome would have been similar in the long run but the executive team's confidence at the six-month mark would have been much lower. The stage gate discipline gave the CEO the language to defend the investment in front of the board every quarter with evidence rather than belief.

Example two. Life sciences SaaS at £32M ARR. New Product Surface pattern.

The business had a compliance and quality platform used by mid-sized pharmaceutical manufacturers. The CEO wanted to add an AI-enabled QC layer as a new product surface: continuous quality monitoring built on top of the compliance platform, sold as a separate module.

The Playbook run scored the idea at 15 out of 30. Value axis 4 (customers would pay meaningfully more for continuous QC than for periodic reviews). Cost-to-Serve 2 (the compute cost of continuous monitoring at real customer scale was closer to the pricing than expected). Willingness-to-Pay 3 (existing quality budget could expand but requires internal justification). Reach 3 (the addressable population is smaller than the CEO initially claimed). Regulation 2 (the QC layer would need FDA clearance in the US and equivalent in other jurisdictions). Reversibility 1 (the offering was bespoke to this customer segment and produced few reusable assets on failure).

Starting stage: Probe. The riskiest assumption was Cost-to-Serve. The Regulation risk was significant but was known upfront and could be de-risked in a Probe.

The Probe ran for ten weeks and £85K. A technical team built a prototype that ran on actual customer data at three-quarters of scale. The cost of compute and inference materially exceeded the assumed unit economics. Even with plausible improvements to model efficiency over eighteen months, the offering could not be profitable at any price the market would bear.

The kill criteria the CFO had written on day one triggered. The offering was killed cleanly at the end of the Probe. Total investment: £100K including Framing.

The paper covers this as a positive kill. The offering would have consumed £2M to £3M over eighteen months in the un-staged funding model before the same cost-to-serve problem revealed itself, at which point the business would have been organisationally and reputationally committed and the kill decision would have been politically much harder. The staged model produced the same answer for one-twentieth of the cost and one-sixth of the time.

The counter-factual matters here too. The team that ran the Probe went on to identify a different AI capability that leveraged the same core platform without the compute cost problem. That capability entered Frame six months later. Twenty months after that, it was in Pilot with a signed paying customer.

Example three. B2B SaaS at £42M ARR. New Pricing Model pattern.

The business sold marketing automation software to mid-market companies on a per-seat licence. The CEO wanted to shift to outcome-based pricing on a subset of the customer base: charging based on qualified leads delivered rather than seats provisioned.

The Playbook run scored the idea at 25 out of 30. Value axis 5 (customers would pay materially more when the price scales with outcome). Cost-to-Serve 5 (AI-driven attribution actually improves the cost of measuring outcomes). Willingness-to-Pay 4 (marketing budgets are structured to pay for outcomes, not seats). Reach 4 (the AI capability made the outcome measurement defensible enough for a large portion of the customer base to accept). Regulation 4 (no material regulatory friction beyond standard contract structure). Reversibility 3 (the offering produces reusable measurement infrastructure).

Starting stage: Scale. The scoring was strong enough that the business had effectively de-risked the offering already through eighteen months of measurement infrastructure investment that pre-dated the pricing decision.

Even so, the funding was structured with a six-month evaluation period on the first cohort of accounts that opted into outcome pricing. The evaluation cleared. The Pilot on the outcome-priced accounts delivered 3x the previous seat revenue on the same accounts. Unit economics were significantly better than the seat model. Twelve months into Scale, one-third of the customer base had moved to outcome pricing and the revenue line had grown by 40% on the same customer count.

The lesson here is that a strong Playbook score does not skip the discipline. The Scale investment was still structured with a six-month evaluation period and defined kill criteria. The kill criteria did not trigger. But their existence made the initial Scale approval a defensible board decision rather than an act of faith.

Before / after: the May board meeting

The CEO of a PE-backed B2B SaaS business at £28M ARR opened the May 2025 board meeting with a proposal to fund an AI-enabled revenue product. The product was ambitious. The CEO's ask was £1.6M in year-one investment on the strength of a business case that assumed a first paying customer at month nine and £4M ARR from the new line by end of year two.

The CFO could not fund it. Not because he was against the idea, but because the business case was not a business case. It was a business plan. Every number in the model was a hypothesis. The willingness-to-pay assumption was based on three exploratory customer conversations. The unit economics assumption was based on a technical estimate the engineering team was not confident in. The reach assumption used a market size number from a public analyst report the CFO knew to be aspirational. The CFO named these in the meeting. The CEO defended each assumption. The chair called time and said the board would revisit at the next meeting.

At the next board meeting six weeks later, the CEO came back with the same business case. The CFO's response was the same. The chair called time again. The pattern was on the way to becoming a stalemate, and the CEO was on the way to becoming frustrated enough to start looking at whether the business would still be his eighteen months later.

The chair suggested a different approach. He introduced the Reinventors Playbook and asked the CEO and CFO to run it together over the following two weeks.

The two of them ran the four-pattern selection first. The idea fit New Product Surface:

a new module on the existing platform, sold to the existing customer base. They ran the six-axis lens together and independently. Their scores diverged by seven points on the aggregate, mostly on the Willingness-to-Pay and Cost-to-Serve axes. The average score was 17.

Starting stage per the ladder: Probe. Twelve weeks. £85K.

The Probe was structured to test the riskiest assumption from the axis divergence: willingness-to-pay. The Probe was three targeted paid conversations with prospects, each of whom was asked to pay a small research fee for the opportunity to shape the offering. Two of the three prospects paid the research fee. The third did not.

The evidence gate: two out of three prospects paid. That cleared the willingness-to-pay assumption.

Pilot was approved at £340K over six months. The Pilot built a functional first version of the module and signed one paying customer at the price point the CFO had modelled for scale. The Pilot ran into two of the three predictable Pilot problems (implementation was harder than estimated, and one specific integration required more engineering than planned) but delivered a renewing paying customer at the end of six months at the modelled price.

Scale was approved twelve months after the original May board meeting proposal. Total investment through Frame, Probe, Pilot: £445K. The un-staged proposal would have committed £1.6M in the same period.

Two years after Scale approval, the offering was a £5.2M ARR revenue line, materially exceeding the original business case model. Gross margin was better than the core product. The offering was the differentiated growth story that anchored the exit conversation twelve months later.

> *The lesson is not that the CFO was right and the CEO was wrong. Both were partly right. The lesson is that the Playbook gave them a shared instrument.*

The CFO could approve staged funding against evidence gates rather than approve a single business case on faith. The CEO could invest in the product he believed in without asking the CFO to abandon financial discipline. The chair could sponsor the sequence rather than adjudicate a stalemate. Everybody kept their integrity intact and the business made new money.

Where the Playbook does not apply

The Reinventors Playbook is designed for growth-stage SaaS businesses between roughly £10M and £100M ARR funding AI new-revenue investments. Four situations where the Playbook does not translate cleanly.

Deep-tech businesses with year-long R&D cycles

The four-stage ladder assumes stages measured in weeks and months. Deep-tech businesses often have research phases measured in quarters and years. The pattern still applies (stage funding against evidence gates), but the stage durations and budgets need to be adapted to the cycle length. Frame becomes a quarter, Probe becomes six to nine months, Pilot becomes twelve to eighteen months. The mechanic holds. The specific numbers do not.

Pure efficiency AI plays

Efficiency AI is different. The business case template every CFO already uses is designed for efficiency plays. Do not force efficiency AI through the Reinventors Playbook. Fund it conventionally with an ROI calculation and a payback period. The Playbook exists specifically for the class of investment that the conventional business case cannot value.

Portfolio companies with no CFO discipline

The Playbook requires a CFO who can hold stage gates rigorously and write kill criteria that will actually kill the initiative if triggered. Portfolio companies that do not have that discipline in the first place will run the Playbook as theatre: the stages happen, the evidence gates get soft-approved, the kill criteria get quietly ignored. The Playbook exposes weak CFO discipline rather than compensating for it. If the CFO cannot hold the gates, the Playbook does not work.

Businesses in distress

In a distress situation, funding staged learning is a luxury the balance sheet cannot afford. Distress businesses should not be funding AI new-revenue plays at all. They should be stabilising the existing revenue base. Come back to the Playbook once the business is stable and the balance sheet permits multi-quarter investment horizons.

Two additional situations that do not disqualify the Playbook but are worth flagging.

Post-acquisition integration. Businesses inside the first twelve months of a significant acquisition often want to fund AI new-revenue investments to justify the deal. The Playbook still applies, but be honest about which axes are being scored on evidence and which are being scored on strategic conviction. Post-acquisition strategic conviction has a well-observed tendency to inflate scores, particularly on Reach and Willingness-to-Pay. Consider adding a third external scorer to the Playbook run to reduce the confirmation bias.

Businesses with a specific defensible AI moat. If the business already has meaningful AI capability, data, or model advantage that gives it a genuinely defensible moat, the Playbook scoring will produce higher scores across Value and Cost-to-Serve than would be typical. That is legitimate. The stage-gate discipline still matters even at Scale scores, because the offering still has to prove product-market fit at a specific price with a specific customer. Do not let strong scoring skip the evidence gates. It just means the stages get funded more generously per stage.

Next steps

Read the paper. This paper explains why the Playbook scores what it scores, walks through each of the four patterns and six axes in depth, and shows three worked cross-industry examples.

Run the Playbook on your current AI new-revenue idea. Pick the pattern. Score the six axes. Determine the starting stage on the ladder. Write the kill criteria for the current stage before any funding is approved.

Run it with your CFO. The Playbook does not work if the CFO does not own the funding gates. Score the idea independently. Compare. The axes where you scored 5 and the CFO scored 2 are the interesting ones. That is where evidence needs to be gathered.

Run the interactive version. The Ortent Reinventors Diagnostic lives at ortent.co/tools/reinventors/prompt. Pattern selection, six-axis scoring, starting-stage recommendation, draft Probe brief.

Book a working session. If the Playbook surfaced something you want a second pair of eyes on, book a session at ortent.co/contact. Forty-five minutes. No pitch, no deck, no slides. Just the paper, your scoring, and the specific stage-gate design your CFO can hold.

Subscribe to The Growth Chair. Weekly notes for founders, boards, and investors trying to make growth-stage SaaS actually scale. One issue a week.

> *That is the whole services page.*

ortent.

Andrew Wyatt is founder of Ortent Advisory Ltd, an operator-led advisory for growth-stage SaaS businesses in Life Sciences, Digital Health, and AI markets. Four exits across three decades, including Lotus to IBM, Paragon to Phone.com, Apertio to Nokia at $240M, and Clearswift to Lyceum. Board Advisor, Non-Executive Director, and Fractional Chief Growth Officer. London.

ortent.co • linkedin.com/in/andrewcwyatt • © 2026 Ortent Advisory Ltd