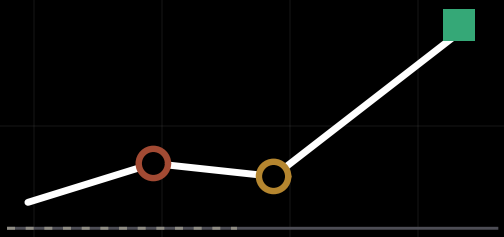
Proofline
ADVISORY



The AI Value Creation Playbook

How to decide what to automate, augment, build, buy, or leave alone.

A practical executive guide to where AI deserves investment, and where it doesn't.

Updated for 2026

Seven pages. Four steps. One number that decides everything.

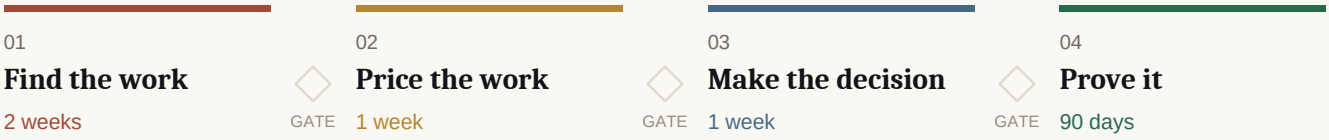
The premise

Why most AI business cases are backwards

Most AI business cases start in the wrong place. Someone sees a capability, a demo, or a competitor announcement, and the organisation goes looking for somewhere to put it. The business case gets written afterwards, to justify a decision that has already been made.

That sequence produces a predictable outcome. The pilot works. The technology does what it promised. And nobody can say what it was worth, because nobody measured what the work cost before it started.

The value question runs the other way. Start with the work your business already does, find out what it costs, and only then ask whether AI changes that number. Most of the time the honest answer for any given workflow is no. That is not a failure of the technology. It is the filter working.



The sequence we use in client engagements, compressed into four steps.

01 **Find the work.** Inventory the recurring workflows where information gets assembled, interpreted, communicated, or turned into a decision.

02 **Price the work.** Establish the unit economics of a single instance, then annualise it. This is the step almost nobody does, and the one that makes everything after it credible.

03 **Make the decision.** Plot impact against feasibility and land on one of five verdicts: automate, augment, build, buy, or leave alone.

04 **Prove it before you scale.** Assess readiness across six dimensions, run to a decision gate, and measure against the baseline from step two.

Four questions this document answers

Which work in my business is actually a candidate?

What is that work costing me today?

Should I automate it, augment it, or leave it alone?

How do I know it worked?

STEP 01

Find the work

Value does not hide in functions or departments. It hides in workflows, and specifically in the ones where a human assembles information, interprets it, communicates it, or turns it into a decision.

Those four verbs are the filter. If a workflow does not involve at least one of them, AI is unlikely to change its economics, however modern the technology sounds.

What makes a candidate

It recurs. Something that happens 2,000 times a month is a candidate. Something that happens twice a year is a project, not a workflow.

It has a definable output. You can describe what finished looks like, and two people would agree on whether an instance is finished.

The inputs already exist. The information a system would need is in your systems, documents, or inboxes. If it only exists in someone's head, you have a knowledge problem to solve first.

A person is the bottleneck. The work waits on human attention rather than on an external party, a regulator, or a clock.

What looks promising but isn't

Work that is genuinely rare. High visibility, low volume. The business case never clears the build cost.

Work where the judgement is the job. AI can prepare the decision. If the judgement itself is the whole job, you are looking at augmentation at best.

Work that is fast because it is broken. Some processes are quick because they skip steps. Automating them scales the shortcut.

Work nobody owns. If you cannot name the person accountable for the outcome, the pilot will not survive contact with the operating business.

A worked inventory

Workflow	Volume/month	Verb	Owner	Candidate
Transaction exception handling	2,400	Interpret	Controller	Yes
Vendor onboarding review	180	Assemble	Procurement	Yes
Board pack preparation	4	Assemble	CFO	No, volume
Pricing approvals	900	Decide	Sales ops	Yes
Annual audit response	1	Assemble	Finance	No, volume
Customer escalation triage	1,600	Communicate	Support lead	Yes

Six workflows, four candidates. That ratio is normal. A first inventory that flags everything as a candidate has not been filtered hard enough.

Fill this in for your business. Print this page. List ten recurring workflows, tag each with its verb, its monthly volume, and its owner, then strike out everything that fails the filter.

STEP 02

Price the work

This is the step that gets skipped, and skipping it is why so many AI initiatives cannot prove they worked. Without a baseline, an improvement can be delivered but not demonstrated.

Pricing a workflow takes four numbers.

Volume	How many instances per month. Count them, do not estimate them. The gap between the two is usually large and usually in the same direction.
Touch time	How long one instance occupies a person, end to end. Measure by watching the work. People report the median and live with the tail.
Rework rate	The share of instances that come back. Often a third of the true cost, and it never appears in the process documentation.
Loaded cost	Salary, benefits, employer taxes, tooling, and management overhead. Not base salary.

Worked example: transaction exception handling

Volume	**2,400 per month**	Clean instance	**0.183 hrs x $52 = $9.53**
Touch time	**11 minutes per instance**	Rework, 18%	**+ $1.72**
Rework rate	**18%**	True cost per instance	**$11.25**
Fully loaded cost	**$52 per hour**	Monthly / annual	**$27,000 / $324,000**

That is one workflow, in one team, that nobody had put a number against.

Then ask what AI actually changes

Assume 70% of exceptions resolve straight through, and 30% still escalate at the same cost. Inference cost is included in the per-instance figure.

Escalated instances	720 per month at $11.25
Straight-through cost	included, under $0.50 per instance
New monthly cost	$9,072
Annual saving	$215,000
Run cost, platform and monitoring	$24,000 per year
Net annual benefit	**$191,000**
Build cost	$180,000
Payback	**11 months**

Eleven months is a real business case, and the kind of number that survives a CFO reading it. If you cannot fill in the four inputs, you are not ready to fund the workflow. And the saving is only real if you decide in advance what happens to the freed capacity.

> **Fill this in for your business.** Take your top candidate from the previous page and price it. Volume, touch time, rework, loaded cost. Then annualise.

STEP 03

Make the decision

Plot each priced candidate on two axes. Business impact, from step two. Feasibility, meaning how hard it is to change this workflow given your data, systems, and organisation. Five verdicts sit on that grid, and four of them are not “build an AI system”.

Automate	High impact, high feasibility, and an error is recoverable. The system handles the instance end to end and a person reviews exceptions. This is the smallest quadrant, and organisations consistently try to put too much in it.
Augment	High impact, but the judgement is consequential or contested. AI prepares the decision and a person makes it. Faster preparation is usually where most of the value was anyway. When in doubt, this is the answer.
Buy	The workflow is not distinctive to your business. Someone sells this. The cost of buying is loss of control over the roadmap, which matters only if this workflow is a source of advantage.
Build	The workflow is distinctive, the data is yours, and the advantage compounds. Right less often than it feels like it is, and it carries an ongoing obligation rather than a one time cost. Budget for the second year.
Leave alone	The economics do not clear, the process is not stable enough to improve, or a cheaper non-AI fix exists. Legitimate and frequently correct. A playbook that never returns this verdict is a sales document.

The order of operations that saves money

Before you decide to automate, ask whether the process should be simplified instead. A workflow with eleven steps, four handoffs, and a rework rate of 18% does not primarily have an AI problem. Automating it preserves the eleven steps in software, where they become harder to change. Fix the process first where you can. Then automate what is left. The sequence is cheaper and the result is more durable.

Client outcomes

Engagements led by Proofline's principal, following this sequence.

+300%	**Audience growth**	Digital media organisation, through AI marketing automation
+50%	**Loan origination volume**	Specialty consumer finance, following AI decisioning and underwriting implementation
−70%	**Monthly accounting effort**	Same digital media organisation, after workflow redesign
+40%	**Net margin**	AI-enabled operations platform
−30%	**Readmission rates**	Healthcare platform, targeted patient populations

Figures reported by the client organisations. Names anonymised.

STEP 04

Prove it before you scale

A pilot succeeds under conditions production does not offer: clean data someone prepared, an engaged team who chose to take part, a stable workflow, and no governance review in the way. Production removes all four. The pilot did not fail. It was never a test of the thing that needed testing.

Six dimensions decide whether it survives

The weakest dimension is your constraint. It governs the whole effort regardless of how strong the others are.

Data foundation	Is the information a system would rely on available and trustworthy enough?
Process readiness	Is the work defined and measured well enough to prove it got better?
Technology fit	How easily does data move between systems, and how fast can a workflow change?
Talent and adoption	Do new ways of working become standard practice once introduced?
Governance and risk	Can consequential decisions be reviewed, and who approves the tools?
Leadership alignment	Do leaders agree on the outcome, and on what happens to the freed capacity?

The 90-day roadmap

Days 1–14	**Baseline** Measure the workflow. Write the number down where the sponsor can see it. Gate: Is there enough economic value here to continue?
Days 15–45	**Narrow build** One workflow. One team. Real data, production conditions, no special handling. Instrument it from day one. Gate: Is the improvement real under normal conditions?
Days 46–75	**Operate it** Hand it to the people who will run it. The workarounds they invent tell you what the design got wrong. Gate: Has adoption held without the project team in the room?
Days 76–90	**Decide** Expand, hold, or stop. All three are acceptable, and naming “stop” in advance makes the gate real.

Executive readiness scorecard

Score each dimension 0 to 100 on how ready you are today, not where you intend to be. Your lowest score is your constraint. Address it before funding anything else.

Dimension	Score	Evidence for that score
Data foundation		
Process readiness		
Technology fit		
Talent and adoption		
Governance and risk		
Leadership alignment		

My constraint is: _____ What has to change first: _____

Print this page and fill it out with your leadership team, separately, before you compare answers. Where your scores disagree is more informative than where they agree. For a scored version, the assessment at prooflineadvisory.com takes six minutes.

About the principal



Mitesh Vashee

Founder and Principal, Proofline Advisory

Mitesh is an AI, product and business transformation executive with more than two decades of experience building companies, launching products, leading transformation programs, and operating at executive level.

He also led technology transformation at a leading nonprofit during a period of 150% revenue growth to approximately $400M. His experience spans early-stage ventures and large enterprises across financial services, healthcare, media, and B2B software.

Proofline grew out of that experience: start with the business problem, establish the economics, and use AI only where it materially improves the outcome.

TRACK RECORD

20+
Years of experience

3x
Startup founder

130+
Products launched

3 F100s
AI and technology innovation hubs built

SCALE

$3B+
Transformation engagements led

$8B
Banking-sector merger

2B+
Global product and platform reach

220+
Person teams led

Trusted by

Rynoh · pwc · accenture · experian · usbank · Houston Landing · Walmart · NIKE

Where to go next

Take the AI Opportunity and Readiness Assessment. Fourteen questions, six dimensions, about six minutes, and you get your constraint identified immediately.

prooflineadvisory.com/assessment

Or start a conversation: info@prooflineadvisory.com. We reply within two working days. The first conversation is 30 minutes and it isn't a pitch.