EARNED AUTOMATION FRAMEWORK

Evaluating where AI belongs at your bottlenecks

Most AI strategies fail because they apply AI too broadly, too early. The better approach is to first identify operational bottlenecks as candidates for AI transformation. This framework evaluates where AI belongs once they're identified.

Tenets

Evaluate the situation, not the technology. AI capabilities change quarterly. Your bottleneck's characteristics don't.

Automation is earned, not rationalized. Invest in the situation before investing in the tool.

Organizational pressure to adopt is not a business case for doing so. The pressure is real, but the case still has to be.

Less surface area, more leverage. Narrowing where you apply AI is a feature, not a limitation.

Using the Framework: Three Axes

Evaluate each bottleneck through three sequential dimensions. Each builds on the last and narrows the field.

1 STRATEGIC DIFFERENTIATION

Does this bottleneck encode how the business wins?

If the work is a source of competitive advantage, replacing it with generic AI risks regressing your output toward the industry mean. This axis forces leaders to protect what wins the game.

LOW — INDUSTRY STANDARD	HIGH — DEFINES YOUR EDGE

2 VARIANCE SENSITIVITY

What is the cost of being wrong here?

AI is fundamentally probabilistic. The question is not whether errors occur, but whether the situation can tolerate them. This is a measure of how forgiving the environment is, not how accurate the AI is.

LOW — ERRORS ABSORBED	HIGH — ERRORS CASCADE

3 COST OF REVERSAL

If this goes wrong, what does it cost to undo?

Evaluated through two lenses: **operational** (how hard to undo) and **perception** (what it signals about your company). If either lens is high cost, the axis is high cost. One is enough.

LOW — EASILY REVERSED	HIGH — HARD TO UNDO OR UNSAY

Reading the Landscape

Each filter produces a binary result: high or low. Three binary filters create eight possible combinations, and each describes a distinct situation with its own risk profile and recommended approach. This is not a scoring system, and you are not averaging across dimensions. You are classifying a bottleneck based on what is true about it.

The situations range from full automation to full stop. Most land somewhere in between, which is where the real decisions live. The value of this framework is not in confirming the easy calls. It is in forcing clarity on the ambiguous ones.

SEE REVERSE →

The Terrain

Where your bottleneck lands determines how to proceed.

D: Differentiation **V:** Variance **C:** Cost of Reversal

LOW HIGH

1. THE GREEN LIGHT

AUTOMATE AGGRESSIVELY

D: LOW V: LOW C: LOW

Classic AI-friendly work. Boring, bounded, and forgiving. Nobody outside the building sees it, errors are cheap, and none of it is why anyone buys from you.

RECOMMENDATION
Automate aggressively. No special guardrails needed.

EXAMPLES
Report formatting, expense categorization, scheduling, data cleanup, document routing

2. THE QUIET RISK

PILOT QUIETLY

D: LOW V: LOW C: HIGH

Mechanically safe but culturally visible. The work isn't core and errors are tolerable, but the decision to use AI here sends a signal that's hard to walk back.

RECOMMENDATION
Pilot quietly. Measure perception as carefully as efficiency.

EXAMPLES
Performance reviews, support triage, onboarding content, sensitive internal comms

3. THE WINGMAN

AUGMENT, DON'T REPLACE

D: LOW V: HIGH C: LOW

AI can help, but the cost of a bad output is real. The saving grace is that mistakes are easy to catch and cheap to fix before they matter.

RECOMMENDATION
Augment, don't replace. Humans in the loop on every decision that matters.

EXAMPLES
Sales call prep, design exploration, competitive synthesis, RFP drafts, code generation

4. THE THIRD RAIL

EXTREME CAUTION

D: LOW V: HIGH C: HIGH

Errors are expensive, visible, and become screenshots. The work isn't what differentiates the business, but getting it wrong can define you.

RECOMMENDATION
Extreme caution. If AI is involved, it stays backstage. A human signs off on everything.

EXAMPLES
Brand voice at scale, content moderation, sensitive customer comms, crisis comms

5. THE BLIND SPOT

CONSTRAIN HEAVILY

D: HIGH V: LOW C: LOW

Each individual output looks fine. But over months, AI pulls everything toward generic. Your messaging starts sounding like everyone else's and you can't pinpoint when it happened.

RECOMMENDATION
Constrain heavily. Tight guardrails and active monitoring for quality drift.

EXAMPLES
Product messaging, sales narratives, editorial tone, UX writing, thought leadership

6. THE LIABILITY ZONE

DO NOT AUTOMATE

D: HIGH V: LOW C: HIGH

Occasional errors won't break anything. But the decision to automate signals what you value, and once stakeholders know a human was removed, that perception is nearly impossible to reverse.

RECOMMENDATION
Do not automate. This is not an AI problem. Fix the process instead.

EXAMPLES
Account management, culture comms, editorial curation, career pathing, partner scoring

7. THE TEMPTATION TRAP

SLOW DOWN

D: HIGH V: HIGH C: LOW

The most dangerous because it looks the most attractive. Low reversal cost creates false confidence. By the time you see the damage, it's embedded in downstream decisions.

RECOMMENDATION
Slow down. Test in parallel. Don't scale until parallel results converge.

EXAMPLES
Pricing strategy, investment thesis, roadmap prioritization, M&A screening

8. THE RED ZONE

DON'T

D: HIGH V: HIGH C: HIGH

Everything is high. The situation cannot tolerate errors, reversal is expensive or impossible, and the signal damage is severe. This is where pressure to automate runs hardest into reality.

RECOMMENDATION
Don't. Name the risk explicitly. Make the cost visible to decision-makers.

EXAMPLES
Loan decisioning, autonomous hiring/firing, insurance adjudication, regulatory submissions

Changing the Classification

These are not permanent labels. If a bottleneck lands in a restrictive archetype, the question is not "how do we justify automating anyway." It's "what would we need to change about the process, the oversight, or the recovery path to move this into a different archetype?" Adding human review can move Variance from high to low. Building a fast correction process can move Cost of Reversal from high to low. You don't change the AI. You change the situation around it.