A CONSULTING GUIDE
BY JUDSON TOMAIKO

46
SLIDES

PREDICTION THEORY · INDUSTRY EVIDENCE · PRACTITIONER FRAMEWORKS

How to Become *AI-Native.*

Most companies ask where to add AI. The better question is what they would redesign if prediction became cheap. This guide works through the answer.

BUILT ON PREDICTION MACHINES, INDUSTRY RESEARCH, AND PRACTITIONER FRAMEWORKS

OPENING

The question most businesses are asking wrong.

Every company I talk to opens with some version of the same sentence. And it immediately tells me they are starting in the wrong place.

THE QUESTION THEY ASK

"Where can we add AI?"

A tool question. It starts with a solution and goes looking for a problem. Companies that start here end up with AI experiments scattered across the organization, none of which change how the business actually runs.

THE QUESTION THAT CREATES VALUE

"If prediction became cheap, what decisions, workflows, and business models would we redesign?"

A strategy question. It starts with the business and asks where cheaper prediction would change how decisions get made.

THE ADOPTION LADDER

Where most organizations actually stand.

AI adoption is not the same as AI-native transformation. Most organizations are on the bottom two rungs of this ladder and believe they are much higher.

1	Individual productivity	People use AI to write, summarize, brainstorm, code, and analyze. It makes individuals faster. Nothing else changes.
2	Task assistance	AI helps with pieces of work, but the underlying workflow stays the same. The process is unchanged.
3	Workflow integration	AI enters real operating processes and starts changing handoffs, reviews, approvals, and exceptions.
4	Decision redesign	AI changes when a decision happens, who makes it, what data supports it, and what action follows.
5	Organizational learning	Daily use creates feedback. Feedback improves prediction. Better prediction changes the business. This is AI-native.

THE CORE ARGUMENT

AI is about decisions, not tools.

Modern AI reduces the cost of prediction. Prediction means using information you have to generate information you do not have. And prediction is inside almost every business decision.

YOUR BUSINESS IS ALREADY PREDICTING

- Which customers will churn
- Which invoices are risky
- Which leads will convert
- Which shipments will be late
- Which machines will fail
- Which employees need support
- Which products will sell
- Which action should happen next

Each of those is a prediction. The question is not whether you are doing it. It is whether you are doing it well, *at the right points in the workflow.*

When prediction gets cheaper, the bottleneck shifts from predicting to **deciding what to do with the prediction.** That is where most organizations are not ready.

DEFINITION

AI is prediction, not magic.

AI takes data you have and generates data you do not have. From a business strategy perspective, that framing is more useful than any amount of technical detail about how models work.

FUNCTION	DATA YOU HAVE	DATA YOU WANT PREDICTED
Sales	Lead behavior	Conversion likelihood
Finance	Invoice history	Fraud or error risk
Operations	Sensor readings	Equipment failure risk
HR	Candidate profile	Role success likelihood
Marketing	Customer activity	Churn probability
Supply chain	Orders and lead times	Future inventory needs

CORNERSTONE

Prediction is not a decision.

A prediction tells you what might happen. A decision tells you what to do about it. **AI does the first. You still own the second.**

> "60% chance of rain" also implies 40% no rain. The model is not telling you what to do. *It is telling you what the world looks like.*

> Prediction + Judgment + Action = Decision

PAYOFF TREE · WORKED EXAMPLE

Take or leave the umbrella?

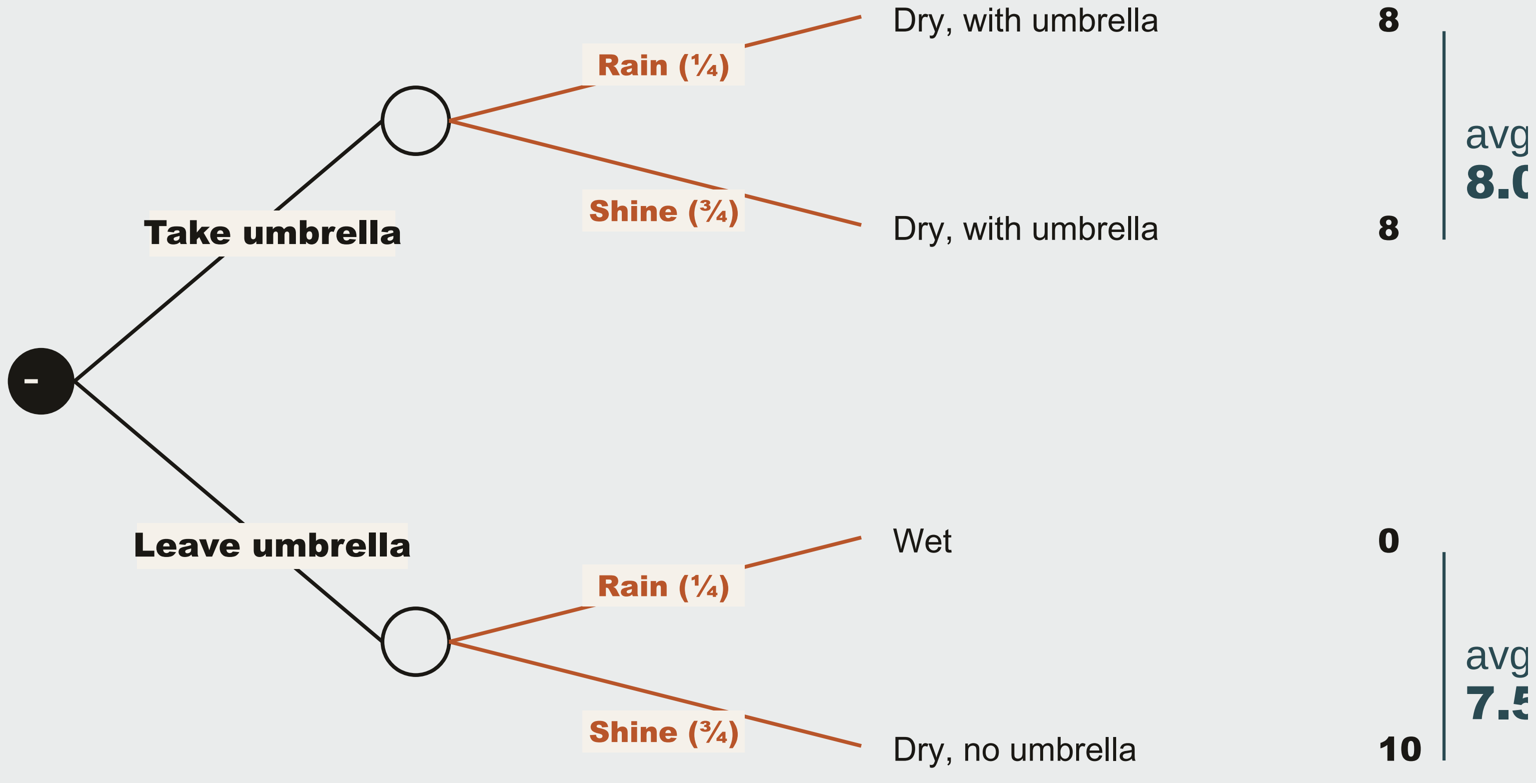

Same prediction (¼ rain). The right action depends on payoffs, not probability alone.

CONTEXT

AI has been around for a while.

AI did not suddenly appear. What changed is the economics: prediction got dramatically cheaper, faster, and more accessible.

ALREADY EMBEDDED

- Fraud detection
- Spam filtering
- Search ranking
- Credit scoring
- Recommendations
- Translation
- Demand forecasting
- Predictive maintenance
- Computer vision
- Dynamic pricing

WHAT CHANGED RECENTLY

Access. Anyone can use it directly. No data science team required.

Interface. Natural language replaced API calls and specialized tooling.

Capability. Reasoning, code, image, and language within one system.

Cost. The price of a prediction dropped by orders of magnitude.

Strategy. Boards and CEOs now treat this as a top-five priority.

Build your strategy around *the economics of cheaper prediction* , not around the hype cycle. The hype will pass. The economics will not.

00

THE INDUSTRY

The market is biased toward building.

The AI economy rewards adoption, hiring, pilots, and visible construction. It barely rewards diagnosis, comparison, or the decision not to build.

Adoption is advancing faster than value.

Different studies set different thresholds for success. The same imbalance shows up in every one of them.

MEASURABLE VALUE

MCKINSEY 2025	Report any enterprise EBIT impact	**39%**
MCKINSEY 2025	Reach the high-performer threshold	**6%**
MIT NANDA 2025	Report measurable pilot P&L impact	**5%**
BCG 2025	Achieve value at scale	**5%**
BAIN 2025	Tie AI to revenue or cost impact	**23%**

Bain also finds 80% report use-case satisfaction. Satisfaction measures something different from scaled financial value, so it is not plotted here.

FAILURE OR ABANDONMENT

RAND 2024	Estimated AI project failure rate	**80%+**
S&P GLOBAL 2025	Abandoned most AI initiatives	**42%**

RAND cites, by some estimates, that more than 80 percent of AI projects fail. RAND interviewed 65 practitioners to study why projects fail. It did not independently calculate that figure. MIT NANDA 2025 is a preliminary study whose findings remain contested.

Different thresholds, same imbalance. Adoption and satisfaction are advancing faster than scaled financial value.

THE INDUSTRY · INCENTIVES

The incentive system rewards building.

95% say AI fluency is required. 53% prefer it over deeper domain knowledge. (TestGorilla, vendor survey, 2026.)

WHAT HIRING SYSTEMS REWARD

Visible action.

- Demonstrated use of AI tools
- Visible building and workflow redesign
- Comfort applying AI to work
- Stories about tools deployed and outcomes produced

WHAT REMAINS INVI

Informed restraint.

- Rejecting an unsupported use case
- Choosing conventional software
- Fixing the process before automating it
- Ending a pilot whose business case failed
- Defining termination criteria in advance

AI action is visible and graded. *Informed restraint is not.*

The problem is diagnosis, not tooling.

THE CAUSE · DIAGNOSIS FAILURE

Most organizations cannot reliably name the problem they are solving before they reach for a tool.

85% of executives told HBR their organizations are bad at diagnosing problems.

THE CONSEQUENCE · SCRAPPED PILOTS

When the problem is unclear, the pilot has no target to hit, and it quietly dies before it scales.

S&P Global finds **46%** of AI proofs of concept get scrapped before broad adoption.

Ask not "which AI should we pilot," but *"what problem are we actually solving."*

Strong incentives to build. *Weak incentives to diagnose.*

The result is the pattern the studies keep finding. Wide experimentation. Limited value. Weak ideas killed only after the spending begins.

THE REST OF THIS GUIDE IS THE CORRECTION. IT STARTS WITH THE WORKFLOW.

PART ONE

01

Workflow redesign beats tool adoption.

Adding AI to a broken or outdated process does not fix it. It accelerates the broken logic and can make things measurably worse.

DIAGNOSIS

What most businesses are actually doing with AI right now.

It looks like progress. It is not. The pattern is the same regardless of which tools are involved.

WHAT IT LOOKS LIKE	WHY IT IS NOT AI-NATIVE
Buying token packages and telling employees to use them	Token usage is not a business outcome
Deploying Copilot or ChatGPT Enterprise company-wide	A subscription does not change a single decision
Building agents and internal tools via prompt engineering	Agents on broken workflows automate the broken logic
Vibe coding: spinning up apps and automations quickly	Built by AI, but static, no learning loop
Hiring prompt engineers without redesigning workflow	Optimizes the wrong layer entirely
Measuring progress by usage rate, spend, or tools created	No feedback loop the org learns from

The metric that matters is not how much AI is being used. *It is whether any decisions changed.*

FIRST MISTAKE

The first mistake is starting with the tool.

Tool adoption can be useful and still completely fail to change the business. The difference is in which question you are answering.

TOOL-FIRST THINKING

Where can we add AI?

- Should we buy Copilot?
- Can we add a chatbot to the website?
- Can we summarize our meetings?
- Can we automate this step?

AI-NATIVE THINI

What workflow actually matters here?

- What decision is inside that workflow?
- What prediction would improve that decision?
- What action changes because of the prediction?
- What feedback makes the system better over time?

No decision, no AI strategy. Start with the workflow. *Find the decisions.* Then find what prediction those decisions need.

CASE STUDY · 1980S

Ford learned this the expensive way.

ATTEMPT 1 · AUTOMATE THE WORKFLOW

Ford computerized its accounts payable department hoping for efficiency gains. Hundreds of people matched invoices against purchase orders and receiving documents. The goal was speed and fewer errors.

Result: same workflow, faster execution. The expected headcount reduction never fully materialized.

ATTEMPT 2 · REDESIGN THE WORKFLOW

Ford asked: *why does this reconciliation process exist at all?* The matching existed because discrepancies arose between what was ordered, shipped, and received. They redesigned: goods were paid for at receipt, not at invoice match.

The reconciliation layer largely disappeared. The department shrank by **75 percent.**

Ask not "Can we automate this step?" but *"Should this step exist at all?"*

UNIVERSAL PATTERN

The same pattern applies across every function.

Every function has a version of this. The old question accelerates the existing process; the AI-native question redesigns it from first principles.

FUNCTION	OLD QUESTION	AI-NATIVE QUESTION
Finance	Can we process invoices faster?	Can we prevent bad invoices from entering the workflow?
Customer service	Can we answer tickets faster?	Can we predict and prevent the issue before they call?
HR	Can we screen resumes faster?	Can we predict role fit while reducing bias?
Supply chain	Can we forecast demand better?	Can we redesign ordering and inventory around live prediction?
Sales	Can we write outreach faster?	Which customer needs which message at which moment?
Operations	Can we inspect defects faster?	Can we predict defects upstream before they happen?

Workflows are where AI becomes real.

A business is a collection of workflows. AI only creates organizational value when it enters actual work.

PROCESS ANATOMY

Workflow → Tasks → Decisions → Jobs

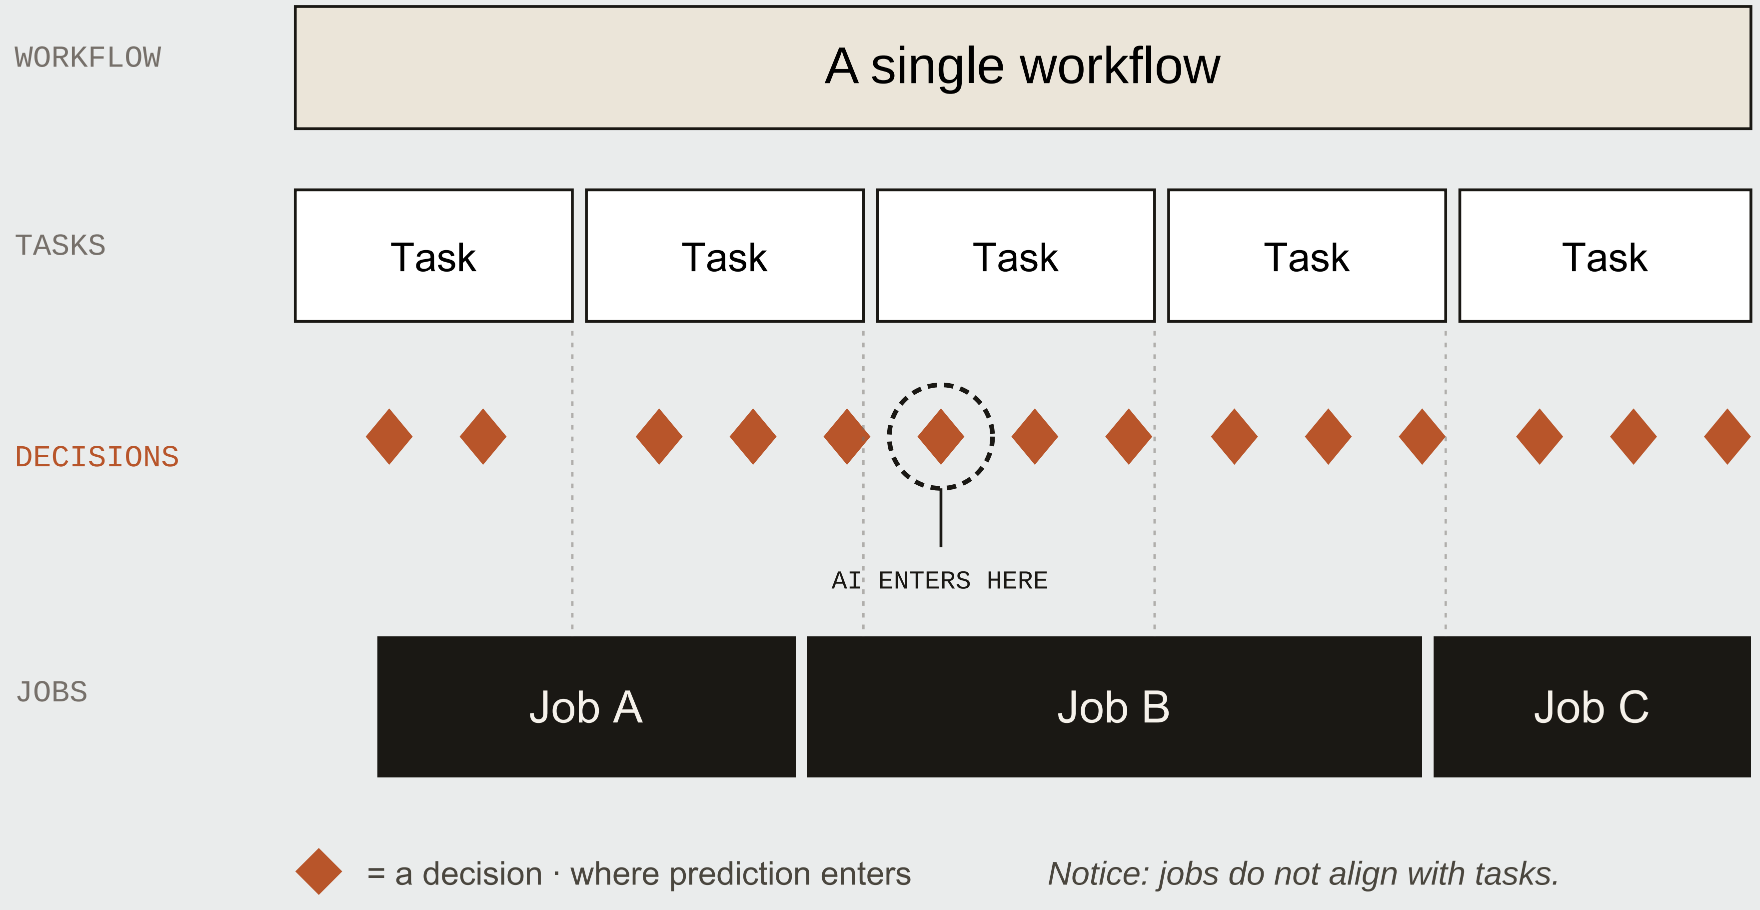

Notice: jobs do not align with tasks.

A WORKFLOW CONTAINS

Inputs · Tasks · Decisions · Handoffs · Exceptions · Approvals · Outputs · Feedback

EXAMPLE · CUSTOMER SUPPORT

1. Customer submits issue
2. System categorizes the issue
3. Agent reviews customer history
4. Agent chooses a response
5. Customer reacts
6. Issue is resolved or escalated
7. Outcome is recorded

AI could summarize, classify, recommend, route, or draft at any of those steps. Each only matters if it changes a decision or improves an outcome.

THE CHAIN

Workflows contain tasks. Tasks contain decisions.

Most AI value lives several layers below where most strategy conversations happen. Six links from idea to impact. Skip one and the investment turns to shelfware.

01
Workflow
The operating process. *Customer support.*

02
Task
The unit of work. *Triage a ticket.*

03
Decision
The fork in the road. *Bot, agent, or specialist?*

04
Prediction
What AI estimates. *Complexity, urgency, risk.*

05
Action
What changes. *Route, escalate, answer, hold.*

06
Feedback
What the system learns. *Resolved? Reopened?*

WEAK USE CASE

If the AI output does not change an action, it probably does not change the business.

STRONG USE C

If the output changes a decision and creates feedback the system can learn from, it becomes a learning loop.

PART TWO

Decisions, data, and the shape of uncertainty.

02

Prediction is powerful only when the organization knows which decision it supports, what data can improve it, and where human judgment still belongs.

ANATOMY

A prediction is one of six parts of an AI decision.

Most organizations build the prediction and skip the other five.

ANATOMY OF AN AI DECISION

Six parts. The model is only one.

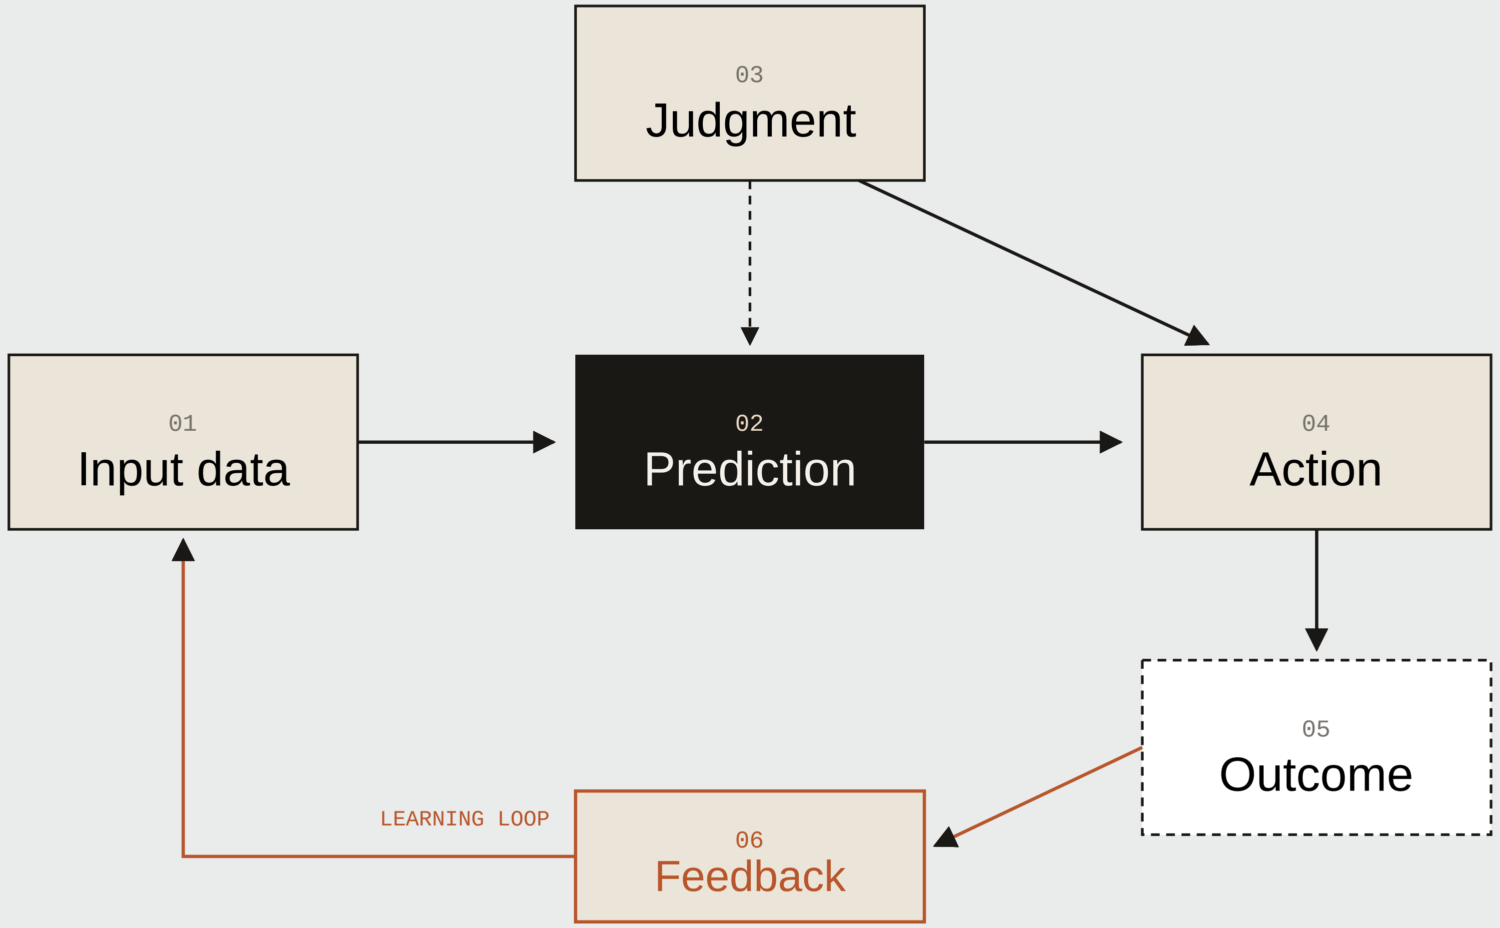

1 · Input data
What the system observes. Quality of prediction is bounded by what goes in.

2 · Prediction
What the system estimates. The probability about missing information.

3 · Judgment
What the organization values. How you weigh different kinds of mistakes.

4 · Action
What changes in the world. The prediction is only as useful as the action.

5 · Outcome
What happens as a result. Did the action produce the intended effect?

6 · Feedback
What the system learns. Outcome data flows back to improve future predictions.

A model predicting 80% churn does not decide whether to discount, call, fix the issue, or let the customer leave. *Each path has different costs. That is judgment's job.*

VALUE

A changed decision is not yet a valuable one.

Changing a decision is the mechanism. Value is the result. Three ways it shows up, and one test for whether it will.

01 · COST OR CAPACITY

Spend less, handle more

The same work costs less, or the same team absorbs more volume without adding headcount.

02 · REVENUE OR GROWTH

Sell more, keep more

Better prediction wins more business, converts more leads, or retains more customers.

03 · RISK OR QUALITY

Fewer errors, safer outcomes

The organization avoids mistakes, catches problems earlier, or produces safer results.

04 · THE READINESS

Five questions

1. What baseline improves
2. Which metric moves
3. How the output moves it
4. Who owns the resulting action
5. When the value becomes visible

If the team can describe the output but not the value mechanism, *the project is not ready.*

VALUE · FUNDING

Spend to save, or spend to make?

Two kinds of AI investment. One funds the other. Framework from IBM's AI Value Creators.

RENOVATION · SPEND TO SAVE

Optimize what already exists.

- Automation and optimization of existing work
- Low-risk internal use cases
- Freed budget becomes the innovation fund

In AI Value Creators, savings from an operational AI application in potato processing funded a more ambitious growth initiative.

INNOVATION · SPEND TO

Change what you sell.

- Prediction changes what you sell or how you decide
- Requires workflow redesign, not tool insertion
- Funded by renovation savings, phase to phase

Renovation funds innovation.

DATA MATURITY

Most organizations have descriptive data. Very few have what AI needs to learn.

AI needs different kinds of data at different points in the workflow. Most have the first type in abundance. Very few have the fourth.

01 · DESCRIPTIVE	02 · DIAGNOSTIC	03 · PREDICTIVE	04 · PRESCRIPTIVE
What happened?	**Why did it happen?**	**What is likely next?**	**What should we do?**
You know your outcomes but not their causes. Most companies have this in abundance and think it is enough.	You understand what drove the outcome. Requires connecting data across systems and asking harder questions than a dashboard can.	You can estimate future outcomes from current patterns. Where most AI investment lives. Requires the diagnostic layer to work.	You know which action produces the best outcome. Requires feedback data showing what interventions actually worked. Most organizations do not have this.
Q3 sales were down 12% in the Northeast.	*The drop followed a pricing change in that region.*	*Without intervention, Q4 will not recover.*	*Adjust regional pricing or reallocate sales effort.*

If you cannot identify where data enters your workflow, where it gets lost, and where feedback returns to the model, you have *scattered data, not an AI strategy.*

METHODS

You cannot redesign what you have not mapped.

Before selecting a model, understand how work actually happens. Not how it is supposed to happen. Those two things are almost never the same.

METHOD 01

Observation

Shadow the people doing the actual work, not their managers describing it.

METHOD 02

Interviews

Ask frontline employees where the real friction is and what workarounds they use.

METHOD 03

Customer journey mapping

Follow the customer experience end to end, not just the internal process.

METHOD 04

Process mining

Use system logs to reconstruct what actually happened in workflows, not what was planned.

METHOD 05

Exception review

Most complexity and most AI opportunity hides in edge cases, not the happy path.

METHOD 06

Bottleneck analysis

Find where work piles up, slows down, or gets reworked. That is where prediction helps most.

The official workflow is almost never the real workflow. *Build AI around real work, not the process diagram.*

FRAMEWORK

Not every decision is equally ready for AI.

This framework helps identify where AI works confidently and where overconfident systems are most likely to fail.

	WE UNDERSTAND THE PATTERN	WE DO NOT UNDERSTAND THE PATTERN
WE HAVE DATA	**Known knowns** Good data and clear pattern. Your strongest AI candidates. Automate confidently. *Credit scoring · Fraud detection · Demand forecasting*	**Unknown knowns** The model looks confident but the pattern may mislead. Dangerous because it does not feel dangerous. *Algorithmic bias lives here.*
WE LACK DATA	**Known unknowns** We know what data is missing. Design to fill the gap or build a human review step. Acknowledge the gap.	**Unknown unknowns** Something outside the training data breaks the model. Classic failure for overconfident systems. *COVID demand shocks · Flash crashes · Novel fraud*

DECISION CRITERIA

When can a decision be automated?

Automation is not a binary switch. It is a function of prediction quality, mistake cost, and reversibility.

EASIER TO AUTOMATE

Confidence is high, stakes are low.

- Data is strong and prediction confidence is high
- Mistakes are low cost and recoverable
- The action is reversible
- Feedback arrives quickly so the model can learn
- Rules are clear and the environment is stable
- Ethical risk is low

KEEP HUMAN-I

Stakes are high, accountability matters.

- Stakes are high and mistakes are hard to reverse
- Data is thin or the situation is genuinely novel
- Trust or fairness is at risk
- Accountability must be traceable to a person
- Customer relationship depends on human judgment
- Regulatory or legal exposure attaches to the outcome

Automate confidence, *not accountability.*

PART THREE

What happens to human work?

03

AI changes the task mix before it changes the job title. The people who understand this early are the ones who adapt ahead of the curve.

AI hits decisions first, then tasks, then jobs, on three different timescales.

"AI will replace X job" is almost always the wrong frame. Jobs do not disappear overnight because they have to lose *enough decisions across enough tasks* first. That is why the four patterns on the next slide exist.

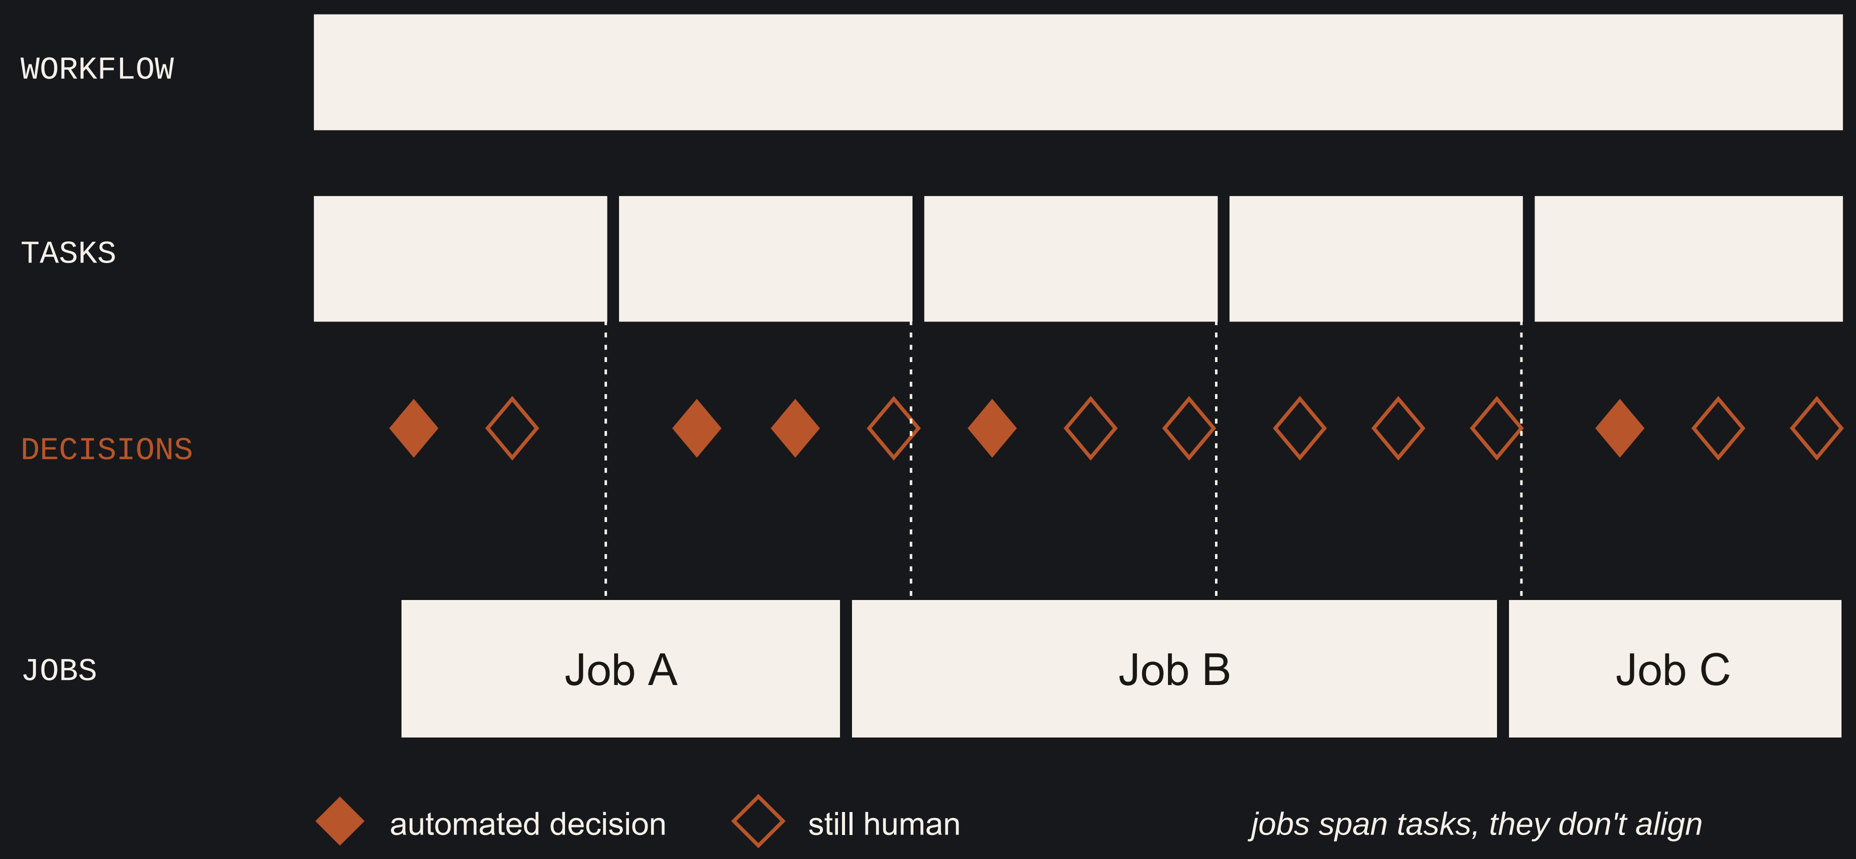

LAYER 01 · DECISIONS

AI automates a few diamonds. The task still exists. The job still exists.

LAYER 02 · TASKS

Enough decisions automate that the task hollows out. The job is still there. It has other tasks.

LAYER 03 · JOBS

Only when enough tasks across a job hollow out does the role itself reconstitute.

FOUR PATTERNS

AI changes jobs in four ways.

The question is not whether AI replaces jobs. It is how the job changes as prediction improves.

01 · AUGMENT

Worker becomes more valuable

AI takes over part of the task. Spreadsheets did not eliminate bookkeepers. The best ones became far more valuable because they could do far more with the same effort.

02 · CONTRACT

Fewer humans for same output

AI removes enough tasks that fewer humans are needed. Fulfillment centers are the example. Headcount drops for a given volume, even if not every role disappears.

03 · RECONSTIT

The job is rebuilt

Some tasks disappear, new ones appear, and the job is rebuilt. Radiology is shifting toward complex case interpretation, patient communication, and AI output supervision.

04 · SHIFT THE SKILL

Same title, different skill

If the bus drives itself, the human role shifts toward supervising and managing children. The credential, training, and hiring criteria all change.

AI changes the task mix *before* it changes the job title. Map which tasks in your workflows are most exposed, and get ahead of it.

WHERE HUMANS MOVE

Humans move toward judgment, exceptions, and trust.

As AI handles more prediction, the work does not disappear. It migrates toward judgment, accountability, and relationship.

HUMAN ROLES SHIFT TOWARD

Judgment under uncertainty
Escalation and recovery
Relationship management
Ethics and accountability
Exception handling
Ambiguous and low-data cases
Customer trust repair
Model supervision and correction

IN PRACTICE

If AI drafts customer responses, the human role shifts from *writing every reply* to reviewing sensitive cases, handling angry customers, correcting tone, identifying emerging issue patterns the model has not seen, and deciding when to escalate to leadership.

AI-native teams need builders and deciders.

Technical capability makes AI possible. Business judgment makes it worthwhile. Most teams over-invest in the first and under-invest in the second.

BUILDERS · TECHNICAL CAPABILITY

What is possible.

- Models
- Data engineering
- Evaluation
- Security

DECIDERS · BUSINESS CAP/

What is worth doing.

- Problem diagnosis
- Workflow knowledge
- Financial evaluation
- Risk judgment
- Knowing when not to build

Builders know what is possible. *Deciders know what is worth doing.*

PART FOUR

Learning loops create advantage.

04

AI-native businesses do not just use AI. They build systems that learn from use. The feedback loop is the strategic asset, not the model.

THE FLYWHEEL

AI-native companies build systems that get smarter every time they are used.

Use creates feedback. Feedback improves prediction. Better prediction earns more use.

THE COMPOUNDING LOOP

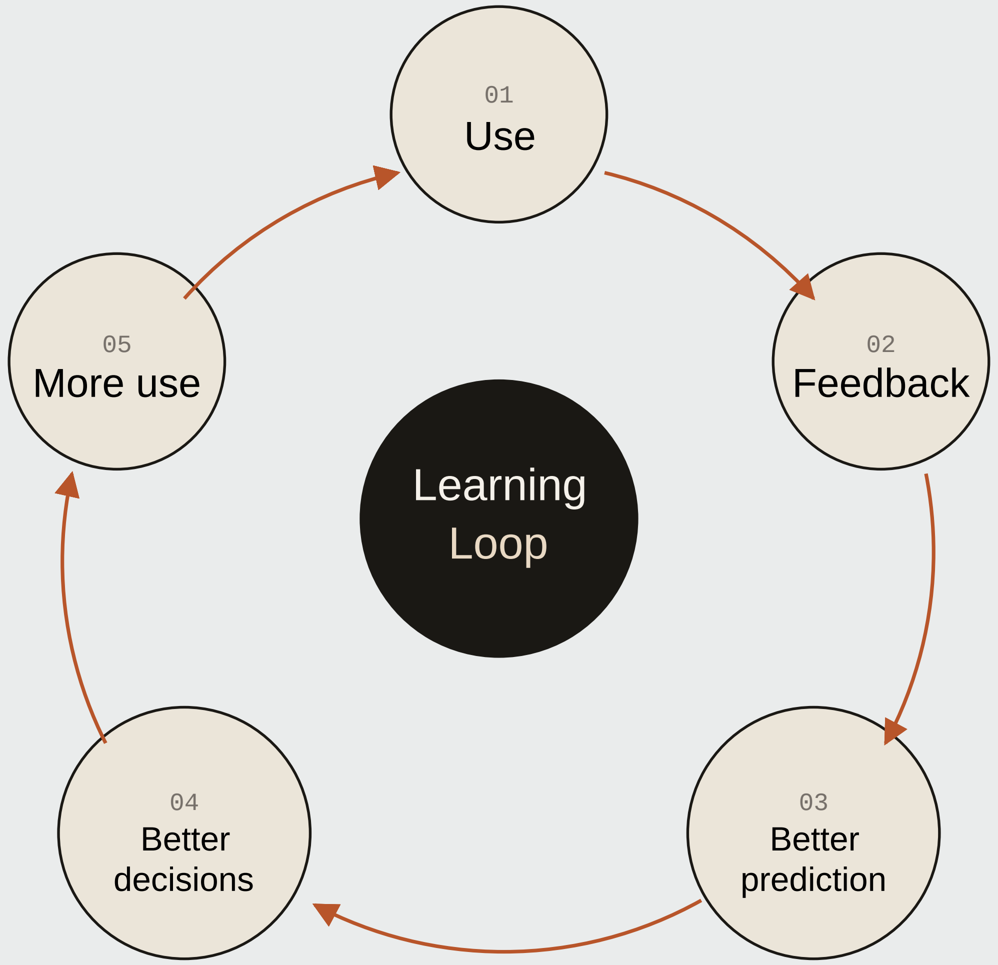

WHAT EACH STEP REQUIRES

1 **Use inside the workflow**
Recommendations have to land in someone's daily work, not in a separate dashboard.

2 **Feedback gets captured**
Accept, edit, or reject. Every choice is a labeled data point. Most tools throw this away.

3 **Predictions improve**
The system retrains on corrections and outcomes, not on its own outputs.

4 **Decisions get sharper**
Workflow becomes faster and more accurate, so it earns more trust and more use.

5 **More use feeds the loop**
Volume compounds. The advantage widens with every cycle.

Your first AI model will probably make things worse before it gets better.

Early in the learning curve, the system is essentially a new employee. Promising, but it makes mistakes and needs supervision before it should act without review.

EARLY SYSTEMS TYPICALLY PRODUCE

- Upset customers who receive wrong or odd outputs
- Employee frustration from reviewing bad recommendations
- Slower workflows during exception handling and human review
- More escalations, not fewer, during the learning period
- Temporary quality declines that test leadership patience

STYLIZED LEARNING CURVE

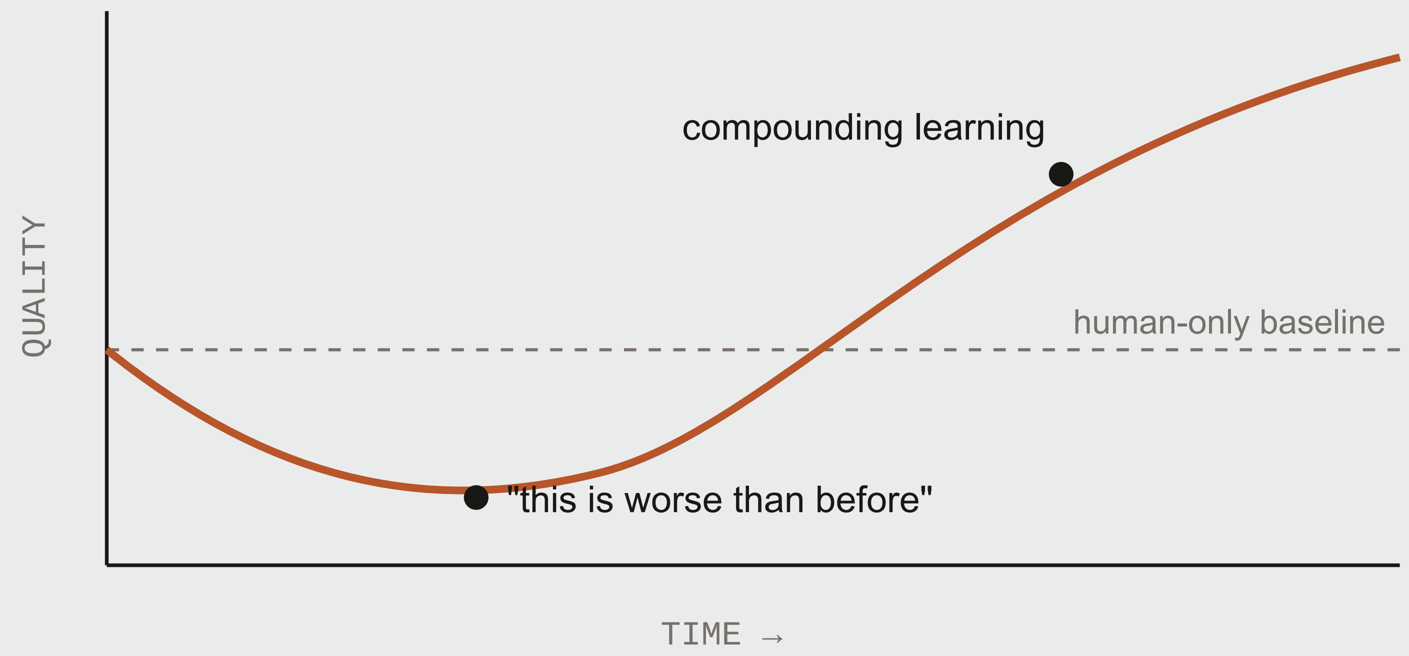

The first model does not need to be perfect. *It needs to be safely learnable inside the workflow.*

READINESS

Don't wait for a perfect model.

Leaders must decide when AI is ready for real work. Most organizations get this wrong by waiting too long or deploying too fast.

THE WRONG QUESTION

"Is the model perfect?"

Perfection is the wrong bar. By the time the model is "perfect" on paper, the workflow has already changed underneath it.

THE RIGHT QUESTION

"Is the model good enough to learn safely inside the workflow?"

Different question, different deployment plan, different success metric.

Cost of a wrong call?

What happens when the model is wrong, and who notices?

Can a human review first?

Is there a human-in-the-loop step before the action commits?

Is the action reversible?

If we get it wrong, can we recover the customer or the outcome?

Will customers notice mistakes?

Does the failure show up externally or stay internal?

How fast does feedback arrive?

Same day? Next quarter? Determines how fast it can learn.

How much short-term pain?

For long-term learning advantage. Be honest with stakeholders up front.

WHY PILOTS FAIL

Most AI pilots test a tool in isolation instead of redesigning a workflow.

COMMON FAILURE PATTERNS

The tool is not embedded in daily work
Employees do not trust the output
Outputs do not actually change any decisions
Feedback from use is not captured
The pilot measures usage rate, not business outcomes
The workflow stays the same after the pilot ends

FROM THE FIELD

> MIT NANDA · 2025
>
> The "GenAI Divide": a small share of deeply integrated pilots show meaningful business value, while the majority of AI initiatives remain stuck with no measurable P&L impact.

Dipping a toe in does not generate compounding learning. **You have to commit to the loop.**

A pilot should test a learning loop, *not just a tool.* The question is whether the organization can learn from it at scale.

PART FIVE

Strategy, ownership, and risk.

05

The hard choices are not only technical. They are about who controls the data, who owns the learning, and who is accountable when AI is wrong.

MAKE OR BUY

The real question is who ends up owning the learning.

Not just a cost or speed question. A strategic question about where learning lives and who captures the compounding advantage.

BUY WHEN

The task is generic.

- Not differentiating
- Process is not your competitive advantage
- Speed to deployment matters most
- Vendor is demonstrably better at this task
- Vendor dependency risk is manageable

BUILD WH

The workflow is your edge.

- Workflow is central to competitive advantage
- Data is proprietary and not available to competitors
- Feedback from use improves the system over time
- The model directly shapes the customer experience
- The learning loop is what creates the durable advantage

As access to capable models becomes more widespread, the model alone becomes less likely to create a durable advantage. Owning the data is not enough if someone else owns the learning loop. The data is the ingredient. *The loop is the factory. You want to own the factory.*

VENDOR DILIGENCE

Your AI vendor may be using your data to build capabilities they will sell to your competitors.

A company may technically own its data and still lose its strategic advantage. If a vendor owns the model and the learning, the risk is not just privacy, it is competitive dependency.

VENDOR DECISION · WORKED EXAMPLE

Accept the contract as written?

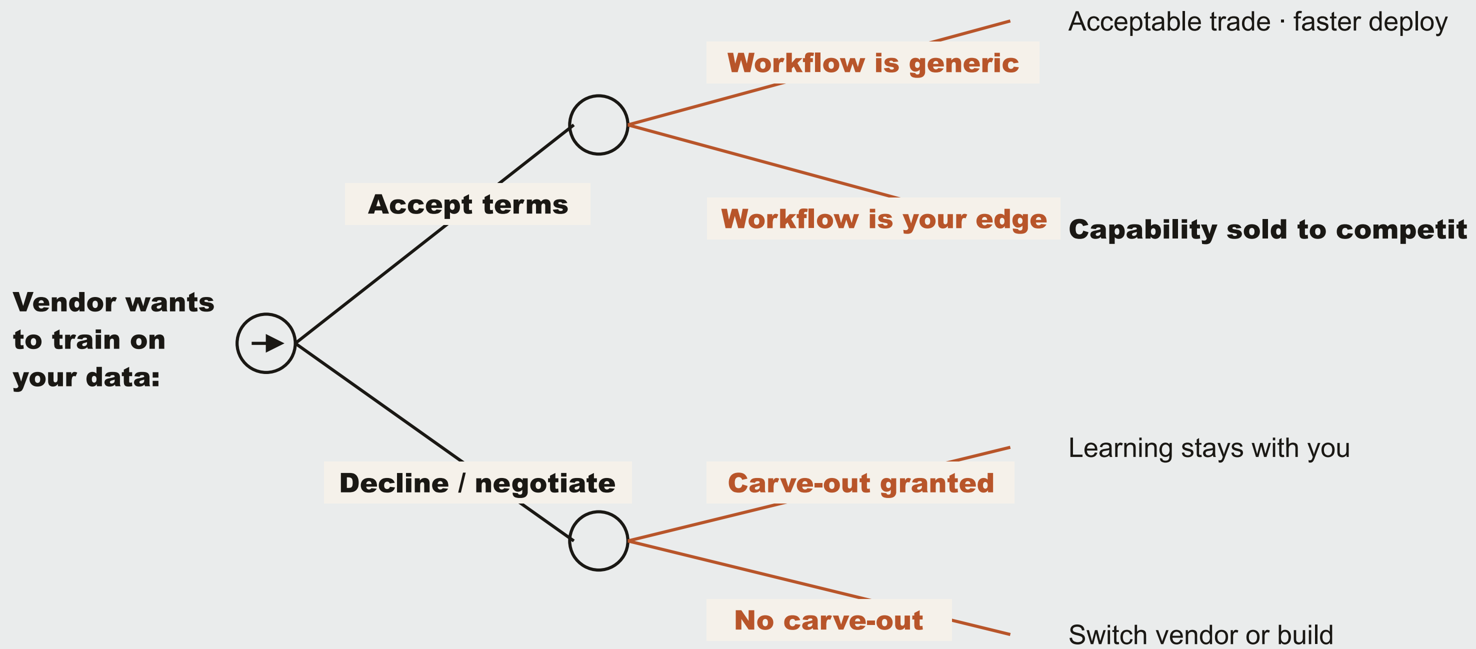

ASK BEFORE SIGNING

Does the vendor train on our data to improve their shared model?
Will our feedback improve a model competitors can later buy?
Can competitors access similar capabilities through this vendor?
Can we export our accumulated learning history if we switch?
Can we leave without losing the workflow intelligence we built?
Are we buying a tool, or giving away our operational knowledge?

OPERATING CONSTRAINT

AI ethics is an operating constraint.

Not a final slide added for optics. Governance keeps deciding, from proposal to retirement, whether a system deserves to remain in use. Not just at launch.

01
Fairness
Does the model perform consistently across demographic groups? Critical in hiring, lending, healthcare.

02
Transparency
Can stakeholders understand why a recommendation was made? Regulators increasingly require this.

03
Explainability
Can the reasoning be communicated to non-technical audiences? Distinct and harder than transparency.

04
Robustness
Does the model fail gracefully on data it was not trained on? Overconfident failure is worse than no prediction.

05
Privacy
What data is being used, where is it stored, and who can access it? Applies to training and inference.

06
Cybersecurity
AI systems are attack surfaces. Adversarial inputs can manipulate outputs in ways traditional software cannot.

07
Human oversight
Where do humans review and intervene? Every high-stakes AI system needs a defined escalation path.

08
Accountability
Who is responsible when the model is wrong? Diffuse accountability is no accountability. Name the decision-maker.

MATURITY MODEL

Levels of AI-native maturity.

Becoming AI-native is a progression, not a switch. Most companies skip levels and wonder why their AI investments are not compounding.

1	Tool use	Employees use AI individually. The organization does not change. Workflows, decisions, and processes are identical to before.
2	Task assistance	AI helps with specific tasks: summaries, drafts, analysis, routing. Individual productivity improves. Nothing structural changes.
3	Workflow integration	AI outputs enter actual business processes. Employees review, correct, and escalate AI recommendations. Handoffs start to change.
4	Decision redesign	The company changes who decides, when they decide, and what data informs the decision. The workflow itself is different.
5	Learning organization	The goal. Daily use creates feedback. Feedback improves prediction. The organization itself becomes smarter every time work gets done.

THE PLAYBOOK

The AI-native playbook.

Four principles that connect everything in this guide.

PRINCIPLE 01

Start with the decision.

Before touching a model, ask: What are we trying to predict? Who uses that prediction? What action changes because of it? If you cannot answer those three, you do not have an AI use case.

PRINCIPLE 02

Redesign the workflow.

Do not ask "Where can AI fit?" Ask: "If prediction were free, how would we build this process from scratch?" That question forces you to challenge the assumptions baked into the current workflow.

PRINCIPLE

Build the data flywheel.

Use creates feedback. Feedback improves prediction. Better prediction earns more use. Design the feedback loop from day one. Without it, you are not building an asset. You are renting one.

PRINCIPLE 04

Keep judgment where it belongs.

As AI handles more prediction, move humans toward ambiguous situations, high-stakes calls, ethical decisions, relationship repair, and accountability. Move them where they matter most.

IN CLOSING

The winners will not be the companies with the most AI tools.

They will be the companies that *redesigned decisions around prediction* , built learning loops into their workflows, and moved human judgment to where it actually matters. Not the companies that piloted the most tools, bought the most licenses, or announced the most AI initiatives.

> AI-native companies do not just use AI. They *rebuild the business so the organization learns every time work gets done.*

Books and industry research.

BOOKS AND FRAMEWORKS

Prediction Machines Agrawal, Gans, and Goldfarb · Harvard Business Review Press

AI Value Creators Thomas, Zikopoulos, and Soule · O'Reilly, IBM

Reengineering the Corporation Hammer and Champy · Harper Business · origin of the Ford accounts-payable case

INDUSTRY RESEARCH

The State of AI in 2025 McKinsey & Company · mckinsey.com

The GenAI Divide (preliminary study) MIT Project NANDA · nanda.media.mit.edu

Root Causes of Failure for AI Projects RAND Corporation · rand.org · 80% is an external estimate cited by RAND

Are You Generating Value from AI? (2025) BCG · bcg.com

Generative AI Shows Rapid Growth, Mixed Results (2025) S&P Global, 451 Research · spglobal.com

AI Moves from Pilots to Production (2025) Bain & Company · bain.com

REFERENCES II

Hiring, education, and management research.

HIRING AND FLUENCY

The State of Hiring for AI Fluency (2026) TestGorilla · testgorilla.com · vendor survey

AI Fluency Rubric (V2, 2026) Zapier · zapier.com · broad AI-fluency hiring framework

MANAGEMENT RESEARCH

Job Interviews Aren't Evaluating the Right Skills (2025) Fuller, Sesser, and Leeds · Harvard Business Review · hbr.org

Are You Solving the Right Problems? (2017) Wedell-Wedellsborg · Harvard Business Review · hbr.org