

Every AI Decision Is an Agency Decision

Why our future with AI depends on designing what it does before it determines what we can do

EMILY LIDDLE · CO-FOUNDER, THEORYA.AI

This Provocation applies Theory A to corporate functions: the work through which an organization operates, governs, and changes. Theory A begins with a simple premise: human agency is a source of organizational value, and organizations can either compound it or diminish it. This Provocation is part of the Design for Agency Collection.

An organization can lose the ability to shape its own future without ever consciously choosing to. It happens one reasonable AI decision at a time.

A system is allowed to answer an employee, review a contract, reconcile an account, approve an exception, or trigger the next action. Each decision may produce exactly what the business case promised: faster work, lower cost, greater consistency, more capacity. It also changes who experiences the work, who develops judgment, who recognizes when something is wrong, and who is able to respond.

Repeated across an organization, these choices change more than how work gets done. They change where the ability and authority to shape an outcome sit, and where responsibility for the result remains. Those are decisions about agency. People exercise human agency through their judgment, choices, and actions. Artificial agency is created through the decisions and actions an organization permits AI to take. Organizational agency is the ability to direct human and AI agency toward an outcome, understand what that arrangement is producing, and reshape it as the organization learns and AI advances.

Every AI decision is an agency decision because it helps determine who or what will be able to shape what happens next. It shapes both the return the organization can realize now and the capabilities it will rely on to create value in the future. Design for Agency treats both as requirements of the same design: create return now and keep building the agency the organization will need next. It can guide a new design, reshape an existing one, or show where earlier choices now require repair.

Many of the capabilities an organization depends on are developed through experience. A new claims professional works through hundreds of ordinary claims before they can spot the one that does not fit. A junior underwriter prices deal after deal before others can rely on their judgment. Doing the work is one of the ways people learn to do it well.

AI can accelerate parts of that development. It can explain unfamiliar cases, simulate difficult situations, and expose people to far more examples than experience alone might provide. Knowledge and pattern recognition that once depended on years of accumulated exposure can now develop much faster. But seeing a hundred difficult situations is not the same as making a consequential decision in one, getting it wrong, and carrying what follows. **Exposure can be accelerated, but responsibility has to be carried.**

Training, simulation, observation, and coaching all contribute to capability. Consequential experience adds something different. A person must apply knowledge under real conditions, make a choice with consequences, receive feedback, and adjust. Over time, this experience develops forms of judgment that rarely appear in a training program or process map: sensing that something is off before being able to explain why, realizing that the problem everyone is solving is not the one worth solving, connecting things others have kept separate, and acting without precedent when the situation calls for it.

Work can create two kinds of value at once. One is the visible result: the claim paid, the deal priced, the question answered. The other is less visible: the judgment, relationships, and shared understanding developed through the experience of producing it. The first appears in every report. The second is rarely measured because work often produces it without anyone having to design for it.

Agency is not merely allocated through work. It is also produced through it.

02

The Ground Just Shifted

AI systems can now exercise agency across a widening range of work. They can interpret what is happening, make decisions within defined boundaries, coordinate activity, and initiate action in both digital and physical environments. Automation has performed predefined steps for decades. What has changed is the scope, speed, and autonomy with which AI can now operate inside an organization.

When AI can interpret, decide, coordinate, and act across a chain of work, the existing task or process stops being a reliable design boundary. Those boundaries reflect what people and earlier technologies could do. AI makes different arrangements possible. The outcome provides a more stable point of reference: what must the organization accomplish, regardless of how the work is divided between people and AI? From there, the organization can decide how human and AI agency should be orchestrated. That orchestration shapes how decisions are made, how

action moves, what people and systems learn, and which capabilities strengthen or weaken over time. **Agency compounds when operating the design leaves the organization more capable than before.**

Every decision about what work AI takes on is two decisions at once. The first is about current return: what the organization can do now, how quickly, at what quality, and at what cost. The second is about what the organization will become capable of because the work has changed.

The same allocation creates both consequences. The first fits easily into a business case and day-to-day reporting: hours saved, costs reduced, faster cycle times, fewer errors. The second develops quietly and is rarely part of how the organization monitors, understands, and manages the work. Fewer people may be learning to handle exceptions, judgment may be concentrating in a smaller group, and the organization may be becoming dependent on a system it cannot meaningfully challenge, even while the work appears to be performing as expected. Design for Agency makes those consequences part of the choice. An arrangement compounds organizational agency only when it creates value now and strengthens the organization's ability to create value next.

03

How Agency Gets Allocated by Default

The allocation of work and authority between people and AI rarely appears as a consequential design decision. It happens inside a process redesign, a job description, a budget, or a software feature turned on for reasons nobody remembers three years later.

Each decision can make sense on its own. Operations are reducing cycle time. Technology is adding functionality. Finance is protecting the expected return. A leader is rewriting a role around the work that remains. Each is answering a legitimate question. What is usually missing is the question that crosses all of them: when these choices add up, will the organization be more capable of shaping what comes next, or less?

Without that question, reasonable decisions made in separate parts of the organization can combine into a design no one intended.

Organizations invest in training, development, and succession. But those efforts usually happen somewhere else, with different people, after the work has already been divided between people and AI. The people responsible for developing capability are left to do so within an arrangement they did not design.

The money follows the same pattern. Every AI investment should be tested for ROI. But ROI cannot answer the whole question. It tells the organization whether an investment is expected to pay back. It does not, by itself, show whether the design

leaves the organization able to adapt, remain relevant, and continue creating value as the world changes. That longer horizon is what return on existence is meant to name. Current return can be supported with evidence. Future capability begins as a hypothesis. It is often not considered in the business case, and organizations may not have the measures needed to understand what is being built or lost once the new design is operating.

That is where the cost hides. **An organization can hit every target it set while weakening its ability to respond to what comes next.** The bill arrives later as work that must be rebuilt, expertise that must be bought back, decisions nobody can confidently make, or dependence on systems nobody can meaningfully challenge. The loss remains invisible as long as the work continues to meet its targets. It becomes visible when the organization needs a capability it stopped developing, chose to outsource, or never recognized it would need.

04

The Line Keeps Moving

Until recently, organizations could treat the division of work between people and technology as relatively stable. Machines and software changed what people did, but their capabilities evolved slowly enough for jobs, processes, and careers to be built around that division for years. Those designs rested on a simple assumption: technology would change slowly enough that today's division of work would remain useful tomorrow.

AI makes that assumption unreliable. What AI can do evolves with each model, vendor release, integration, and new use discovered inside the organization. The boundary between human and AI agency cannot be settled once. As AI evolves and people develop new capabilities around it, organizations must keep reconsidering where each should act and why.

When an organization treats the division of work as fixed, every advancement in AI or change in the business forces it to reopen decisions it thought were settled. Roles have to be renegotiated, processes reopened, controls rewritten, and the work divided again. **The organization redesigns reactively, after the change has arrived, rather than evolving through deliberate choices of its own.**

An organization that knows where agency sits and why can respond differently. When AI capability or business needs evolve, it can ask what has changed, what is now possible, and what should be redesigned. The reasoning behind the current arrangement is visible. The people doing the work help lead the redesign because they know where judgment enters, which exceptions matter, and how the current arrangement is strengthening or weakening capability in both people and AI. The organization can respond deliberately and use each change to improve the arrangement.

The line will keep moving. People must retain the authority and responsibility to decide where it should move next.

The Quiet Clock

When AI absorbs work that once developed people, the flow of human capability changes. Without a deliberate replacement for the learning and experience that work provided, capability can weaken long before performance does. Experienced people continue to handle the hardest cases, results hold, and nothing appears broken. Behind those results, fewer people may be developing the judgment the organization will eventually need.

The gap becomes visible when something falls outside what AI or the current process can handle and too few people have the judgment to respond. Rebuilding at that point is slow, costly, and risky. Instead of using its time and attention to develop what it needs next, the organization must reconstruct a capability it once had.

For a time, the organization can hire people who developed their judgment elsewhere through broad experience and real responsibility. But every organization is drawing from the same pool while the work experiences that traditionally developed future experts are becoming less common.

This is a stock-and-flow problem. The organization's available expertise is the stock. The flow is the continued development of people with the judgment needed to replenish and extend it. **The stock can appear healthy long after the flow into it has slowed, which is why the loss is difficult to see.**

The flow does not depend on preserving the old division of work. AI can perform the work and become part of how human capability is developed. It can expose people to representative and rare cases, reveal evidence and uncertainty, simulate difficult situations, and provide feedback. People can challenge its conclusions, and their corrections can improve the system in return. When capability depends on consequential judgment, the design must still give people opportunities to make decisions and carry responsibility. How the work is designed determines whether AI weakens human capability or helps human and AI capability develop together.

The quiet clock is the time between when replenishment slows and when the organization discovers it needs the capability. Performance may remain strong throughout that period, even as the organization's options narrow. The loss is not inevitable. Deliberate design can create a new flow of capability before the need becomes a crisis.

Not Fewer People. Different Work.

Redesigning work may change how many people an organization needs. Fewer people should not be the default answer.

Starting with headcount asks how many people the current work requires. Starting with the outcome asks what must become possible, what work should exist, how people and AI should contribute, and what new work can now be imagined. Any

change in team size should follow from the design rather than determine it.

Design for Agency does not preserve work simply because people do it today. An organization that understands the capabilities it needs now and in the future can expand AI's role with greater confidence. It can decide which work AI should perform, how the capabilities required by the outcome will continue to develop in people and AI, and what the released capacity can make possible.

Take a team that answers questions and fixes problems for customers or employees. AI can handle many routine requests, and it should. If the redesign stops there, people inherit the escalations, exceptions, and hardest cases. The work becomes more cognitively demanding and emotionally intense, with fewer lower-stakes experiences through which people can learn or recover. **Design for Agency treats the human role as something to be designed rather than allowing it to be defined by whatever AI leaves behind.** It creates a sustainable mix that includes representative cases for learning, AI support during complex work, and protected time for improvement and recovery.

With routine demand reduced, people can work from first principles. They can remove the causes of recurring problems, rethink what the service should make possible, and focus on moments where human judgment and care matter most.

07

Deciding on Purpose

Design for Agency is a method for deciding how human and AI agency should be allocated and orchestrated around an outcome. It makes choices explicit before they disappear inside systems, roles, budgets, and third-party agreements.

The starting point does not have to be the whole enterprise. It can be one outcome, service, operating area, or live AI use case. The outcome sets the scope. The organization defines what must become true, then identifies the work, decisions, capabilities, and forms of agency required to achieve it.

Existing arrangements can be redesigned too. If the organization can still change how work and agency are allocated, it remains a design question. When capability has already been lost, changing the allocation is not enough. Repair is also required. That can include rebuilding apprenticeship, restoring judgment, and recreating the conditions in which expertise develops.

Five questions show whether agency has been designed deliberately or allocated by default:

Why is the work set up this way? The reason should be explicit and still relevant. An approval step, for example, may remain because of a risk decision made years ago or in isolation, even though the original reason no longer applies.

Where should AI act, and where is human judgment required? The boundary should reflect the outcome and its risks. AI might resolve routine expense exceptions while a person handles one involving policy ambiguity or a significant employee

consequence.

Who remains answerable for the outcome? A named person or governing body must understand what the system is doing and be able to intervene. AI may make routine decisions, but human authority remains responsible for the pattern of results.

What should trigger redesign, and who can initiate it? A rise in exceptions, a new AI capability, or a better approach identified by practitioners might all signal that the arrangement should change.

How will capability continue to develop in people and AI? People develop through experience, feedback, responsibility, teaching, and opportunities to stretch. AI improves through better information, tools, evaluation, updates, and governed feedback. AI might handle routine cases while people review representative examples, form their own judgment before seeing its conclusion, and use their corrections to improve the system.

These questions are answered with the people doing the work. Design for Agency gives them a defined role in shaping the arrangement and initiating redesign because they see where judgment, exceptions, and capability development actually occur. That role requires protection: people control what they share, time for design and learning is funded, released capacity does not automatically become headcount reduction, and practitioner proposals receive a clear decision and explanation.

Designing organizational agency requires more than deciding whether AI or a person performs the work. For each capability the outcome requires, the organization can decide whether to compound, create, elevate, distribute, preserve, transfer, or retire it. AI may perform the work while human capability around it continues to develop. Each ruling is a hypothesis based on current evidence, with measures and triggers for review as people, technology, and organizational needs evolve.

The test has two parts: does the design produce return now, and does it compound organizational agency over time? A viable design must do both.

Design for Agency becomes durable when it is embedded in the work. The organization revisits the arrangement as technology advances, people develop new capabilities, or a better way to achieve the outcome becomes possible. Continuous improvement asks how to improve the work. Design for Agency asks how human and AI agency should work together, what that arrangement is making the organization capable of, and what should change next.

The Closing Provocation

AI capability will become widely available. Access to the same models, tools, and automation will not be what distinguishes one organization from another. **The**

difference will come from the human agency an organization develops and how effectively it combines that agency with AI's expanding capabilities.

Much of today's work should change. Some should move to AI. Some should be eliminated because it no longer creates value. Entirely new work should become possible. Design for Agency directs those choices toward an outcome and makes the most of both human and artificial capability.

A successful design must do more than produce value now. It must leave the organization better able to adapt, make consequential decisions, and create value as AI and the world around it continue to evolve.

Every AI decision is an agency decision. It shapes both the value an organization can create now and its capacity to create value next. The real test is whether the return created today also strengthens the organization's ability to create value tomorrow.

This Provocation introduces the Design for Agency Collection, one of the Expertise Collections of TheoryA.ai.

THEORYA.AI

DESIGN FOR AGENCY COLLECTION · THEORYA.AI