
The Bidirectional Move AI Strategy Hasn't Made

BEYOND AUTOMATION.
TOWARD ORGANIZATIONAL INTELLIGENCE.

PERSONAL POV
BY ESTEFANÍA MOLINA

Shared in a personal capacity; Views expressed are my own.

I cannot say this enough: The most important work in AI strategy has nothing to do with AI.

The work that determines whether an enterprise captures AI value happens *before* any model is selected, any platform evaluated, or any use case scoped. It is the design work most leaders do not currently call AI strategy - because it has nothing technical about it. And the reason AI investments keep underperforming is that this work is structurally missing from the conversation.

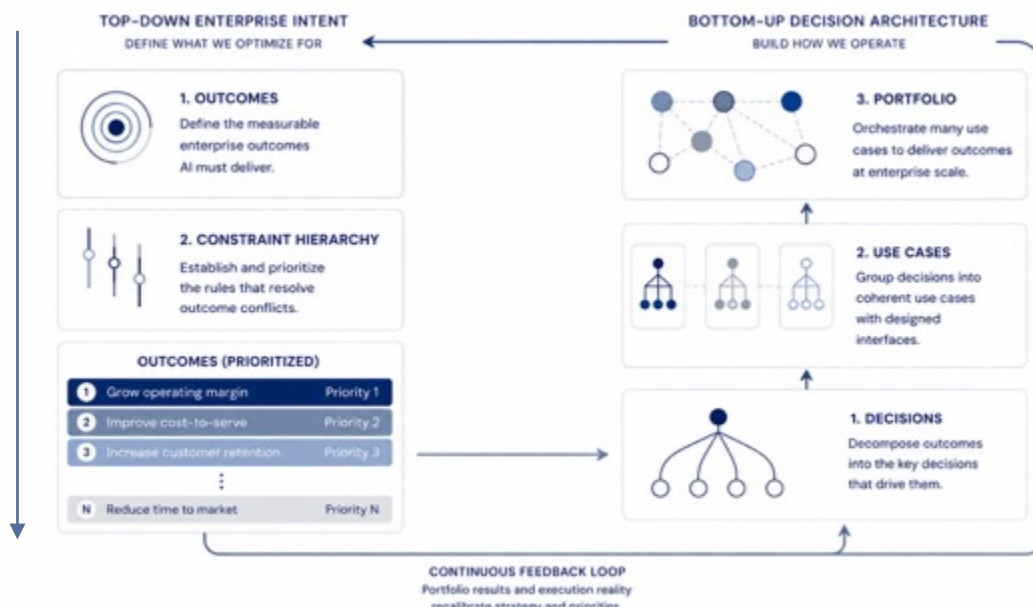
When people say AI strategy today, they usually mean a set of technical choices e.g., which models, which platforms, which use cases, which vendor stack. The work I have in mind is not that. It is **the design work that connects AI directly to business outcomes**: how decisions are framed, which trade-offs govern, where authority sits, how the enterprise reasons. **The technical choices follow this work. They do not replace it.**

Most AI strategy conversations, in the technical sense people usually mean, are unidirectional. Top-down from the boardroom - what is our AI vision, what outcomes do we want - or bottom-up from the use-case list - which projects do we fund, which models do we build, how do we sequence them. The two passes rarely meet. **When they do, it is usually as a sequential handoff**: vision first, then implementation.

I have been increasingly convinced that the work that actually compounds AI value is neither of those. **It is the bidirectional design pass that happens before either direction has fully begun** - and the strange thing about it is that it does not mention AI in its content.

EXHIBIT 1

The Bidirectional AI Strategy (Non-Exhaustive)



- Enterprise outcomes do not operationalize themselves. They must first be **decomposed into the decisions** that produce them
- **Constraint hierarchies** shape how conflicts inside the architecture get resolved
- **Decisions become coherent** only when designed in relation to one another
- **Portfolio behavior** emerges from the interaction between use cases, not from the success of any individual initiative
- **The bidirectional pass is what connects enterprise intent to the systems that execute it**

What "bidirectional" actually means

Two passes, simultaneously, both deliberate, both before development.

The **top-down pass** specifies what changes in enterprise performance when AI works - not as a strategic theme but as an operational outcome e.g., *net margin +2 percentage points sustained after fulfillment cost. Reduce cost-to-serve by 15% while protecting service levels*. It also specifies the constraint hierarchy that resolves outcome conflicts e.g., *service floor 94% first, then margin, then cost* - set once at the enterprise level.

The **bottom-up pass** operates in **three layers**, each scaling up in the scope of what is being designed.

- (i) **Decision architecture is the first.** The unit at this layer is the individual decision. Each AI initiative is decomposed into the specific decisions that must produce the top-down outcome - pricing decisions, allocation decisions, forecast decisions, the cross-functional trade-off decisions where the constraint hierarchy applies directly. Each decision is named, classified by role of AI and scope on the matrix, and given an embedded trade-off rule. The output is a complete inventory of the decisions inside the initiative, with their roles, scopes, and trade-offs explicit.
- (ii) **Use case architecture is the second.** The unit at this layer is the relationships between decisions inside a single use case. Once individual decisions are named at Level 1, their interconnections become the design object: **which decision's trade-off resolution is another decision's input condition? What signal needs to pass between them? Where does one decision generate variance the next must absorb?** The interfaces are designed and instrumented at this layer, not left to be discovered at runtime. The output is a map of how decisions inside a use case connect - a topology, not an inventory.
- (iii) **Portfolio architecture is the third** - and it is where the bottom-up work loops back most directly to the top-down pass. Decisions inside different use cases interact whether or not the organization has noticed. Shared decisions appearing in two initiatives with different models resolving them differently. Sequential dependencies where one initiative's output should inform another's input. Competing objectives where two initiatives push the same business variable in opposite directions. None of these conflicts can be resolved within any individual initiative. They require the top-down pass to govern them - the enterprise outcomes that define what the portfolio is collectively producing, and the priority rules that resolve conflicts when those outcomes compete. This is where the bidirectional loop closes: **the bottom-up work surfaces the cross-initiative interfaces, and the top-down pass makes them coherent at portfolio level.**

At every layer, the two passes connect. The top-down outcomes shape what the architecture is aiming at - outcomes decompose into the specific decisions and use cases that must produce them. The top-down constraint hierarchy shapes how conflicts inside the architecture get resolved - it becomes the trade-off logic embedded at cross-functional decisions. Together, these two connections make the architecture coherent. In most organizations, neither is made explicit. Top-down outcomes are stated, bottom-up use cases are built, and the connecting tissue between them is left to be resolved during execution.

All three layers happen before any technical work begins. The bottom-up pass produces decision catalogs, interconnection maps, and cross-initiative governance - written artifacts that AI later operates inside. None of this work involves model selection, platform evaluation, or implementation. The technical deployment that follows is constrained by what the architecture specifies; the architecture is not retrofitted to whatever gets deployed.

Why before development matters

The instinctive pushback is that this sounds like waterfall in an agile world. Why design everything upfront when you could ship a use case, learn, iterate?

The answer is that **the bidirectional pass is not a design of the technology**. It is a design of the decisions the technology will operate inside. Those decisions exist in the business already - they have been resolved informally for years, often by escalation or by whichever function moves first. The design pass writes them down, with their trade-offs and decision rights explicit, so that when AI shows up to make or shape them at scale, it operates inside a structure rather than inheriting whatever ambiguity was there before.

If the design pass happens after development, the AI ships into the ambiguity, makes the trade-offs implicitly, and the architecture has to be reverse-engineered from what the system already does. **That reverse engineering is the expensive version of this work. Doing it before development is the inexpensive version.**

The two payoffs the design pass produces

Most arguments for upfront design work die at the question of cost. Decision architecture work takes months. It requires senior cross-functional attention. The case for doing it has to compete with everything else competing for that attention. The case has to be specific.


There are two payoffs, and they reinforce each other.

EXHIBIT 2

The Two Strategic Payoffs of the Bidirectional AI Strategy

**From Fragmented AI Activity
to Enterprise Measurement**

ROI becomes measurable when outcomes are **decomposed into decisions** first



Decision Architecture

**From AI Commoditization
to Durable Differentiation**

Differentiation compounds when **judgment is architected** deliberately

The first is measurement. AI ROI measurement is broken in most enterprises not because the math is hard but because the **unit of analysis is wrong**. Use cases get instrumented; enterprise outcomes do not materialize; the gap looks inexplicable. The bidirectional design pass is what makes the gap explicable - and therefore measurable. Once the top-down outcome is decomposed into the specific decisions that should produce it, you can read after deployment whether each decision is producing what it was supposed to, whether the interfaces between decisions are carrying the trade-off resolutions cleanly, and whether the overall outcome is materializing. None of those reads is possible without the design pass that defines what the decisions are in the first place.

The second is competitive advantage. As AI tooling commoditizes - as similar models, similar optimization approaches, and similar decision systems diffuse across competitors - the locus of durable differentiation moves upward. Away from execution itself, **toward the architecture of judgment governing it**. The bidirectional design pass is exactly that architecture. It encodes how an organization frames problems, weighs trade-offs, and protects strategic coherence across decisions. That encoding does not commoditize the way models do. **Two competitors deploying identical AI inside different decision architectures produce different outcomes. The architecture is the moat.**

Both payoffs are the result of the same act. That is unusually clean for an enterprise design move. Most efforts that produce measurement honesty are different from efforts that produce competitive advantage; here they are the same effort, and the bidirectional design pass is what makes them coincide.

The strategic move

The bidirectional design pass is one of those rare enterprise interventions where the cost is bounded - one or two months of sustained senior cross-functional attention within a high-stakes domain - while the payoff compounds for years. Each subsequent AI initiative operating inside the same architecture inherits its constraint hierarchy, its decision rights, and its embedded trade-off logic. The architecture is not redesigned for every initiative. It is designed once at the domain level and reused across many.

The organizations that do this work deliberately will measure their AI investment more honestly, sustain meaningful differentiation as execution converges, and accumulate institutional capability that becomes harder for competitors to close with each passing quarter. The organizations that avoid it will continue instrumenting use cases more carefully, deploying more tools more quickly, and wondering why enterprise outcomes remain persistently weaker than local KPIs suggest.


The work itself is deeply unfashionable. It is granular, analytical, and structurally difficult. It likely requires a new type of enterprise capability that most organizations have not yet named, let alone developed - one able to operate simultaneously across strategic intent, operational design, and decision decomposition.

It is not data science. Data scientists build models that execute trade-off resolutions after decisions are already defined. It is not traditional strategy. Strategy typically operates at the level of narrative, positioning, and prioritization - not decision decomposition and architectural trade-off design. It is not product management. Product managers deliver defined capabilities, not the decision systems governing how those capabilities behave.

The work itself does not produce a launch announcement, a platform deployment, or a transformation story. It produces something much less visible: a decision map, a constraint hierarchy, a set of embedded trade-off rules written directly into the systems executing the enterprise.

The kind of artifact that does not photograph well at a town hall.

But increasingly, it may be the artifact that determines whether everything downstream compounds - or quietly converges.



About the Author
Estefanía Molina

<https://www.linkedin.com/in/estefaniamolinae/>

Strategy consultant with over eight years of experience advising Fortune 500 executives across the Americas on enterprise transformation, AI strategy, and organizational redesign. Her work focuses on translating complex, interconnected decisions into measurable business outcomes, including large-scale initiatives delivering hundreds of millions of dollars in value.