

Where AI Stops

Where AI products work, where they hit a wall, and what that means for how you lead.

by Victoria Rudi

Key points

→ **The real boundary is formalizability.** AI stops where the problem can't be specified and the environment won't hold still long enough to model. That boundary doesn't depend on how good the model is.

→ **The matrix sorts AI products by what they do.** Four quadrants, two axes: bounded vs. unbounded problem, predictable vs. unpredictable environment. Each quadrant represents a different kind of value: efficiency, new capability, better decision inputs, or territory no product can own yet.

→ **Problems move.** A crisis starts unbounded and unpredictable, gets investigated and structured, and migrates toward quadrants where AI becomes load-bearing. It also runs in reverse: a regulatory shock can blow a bounded process wide open overnight.

Introduction

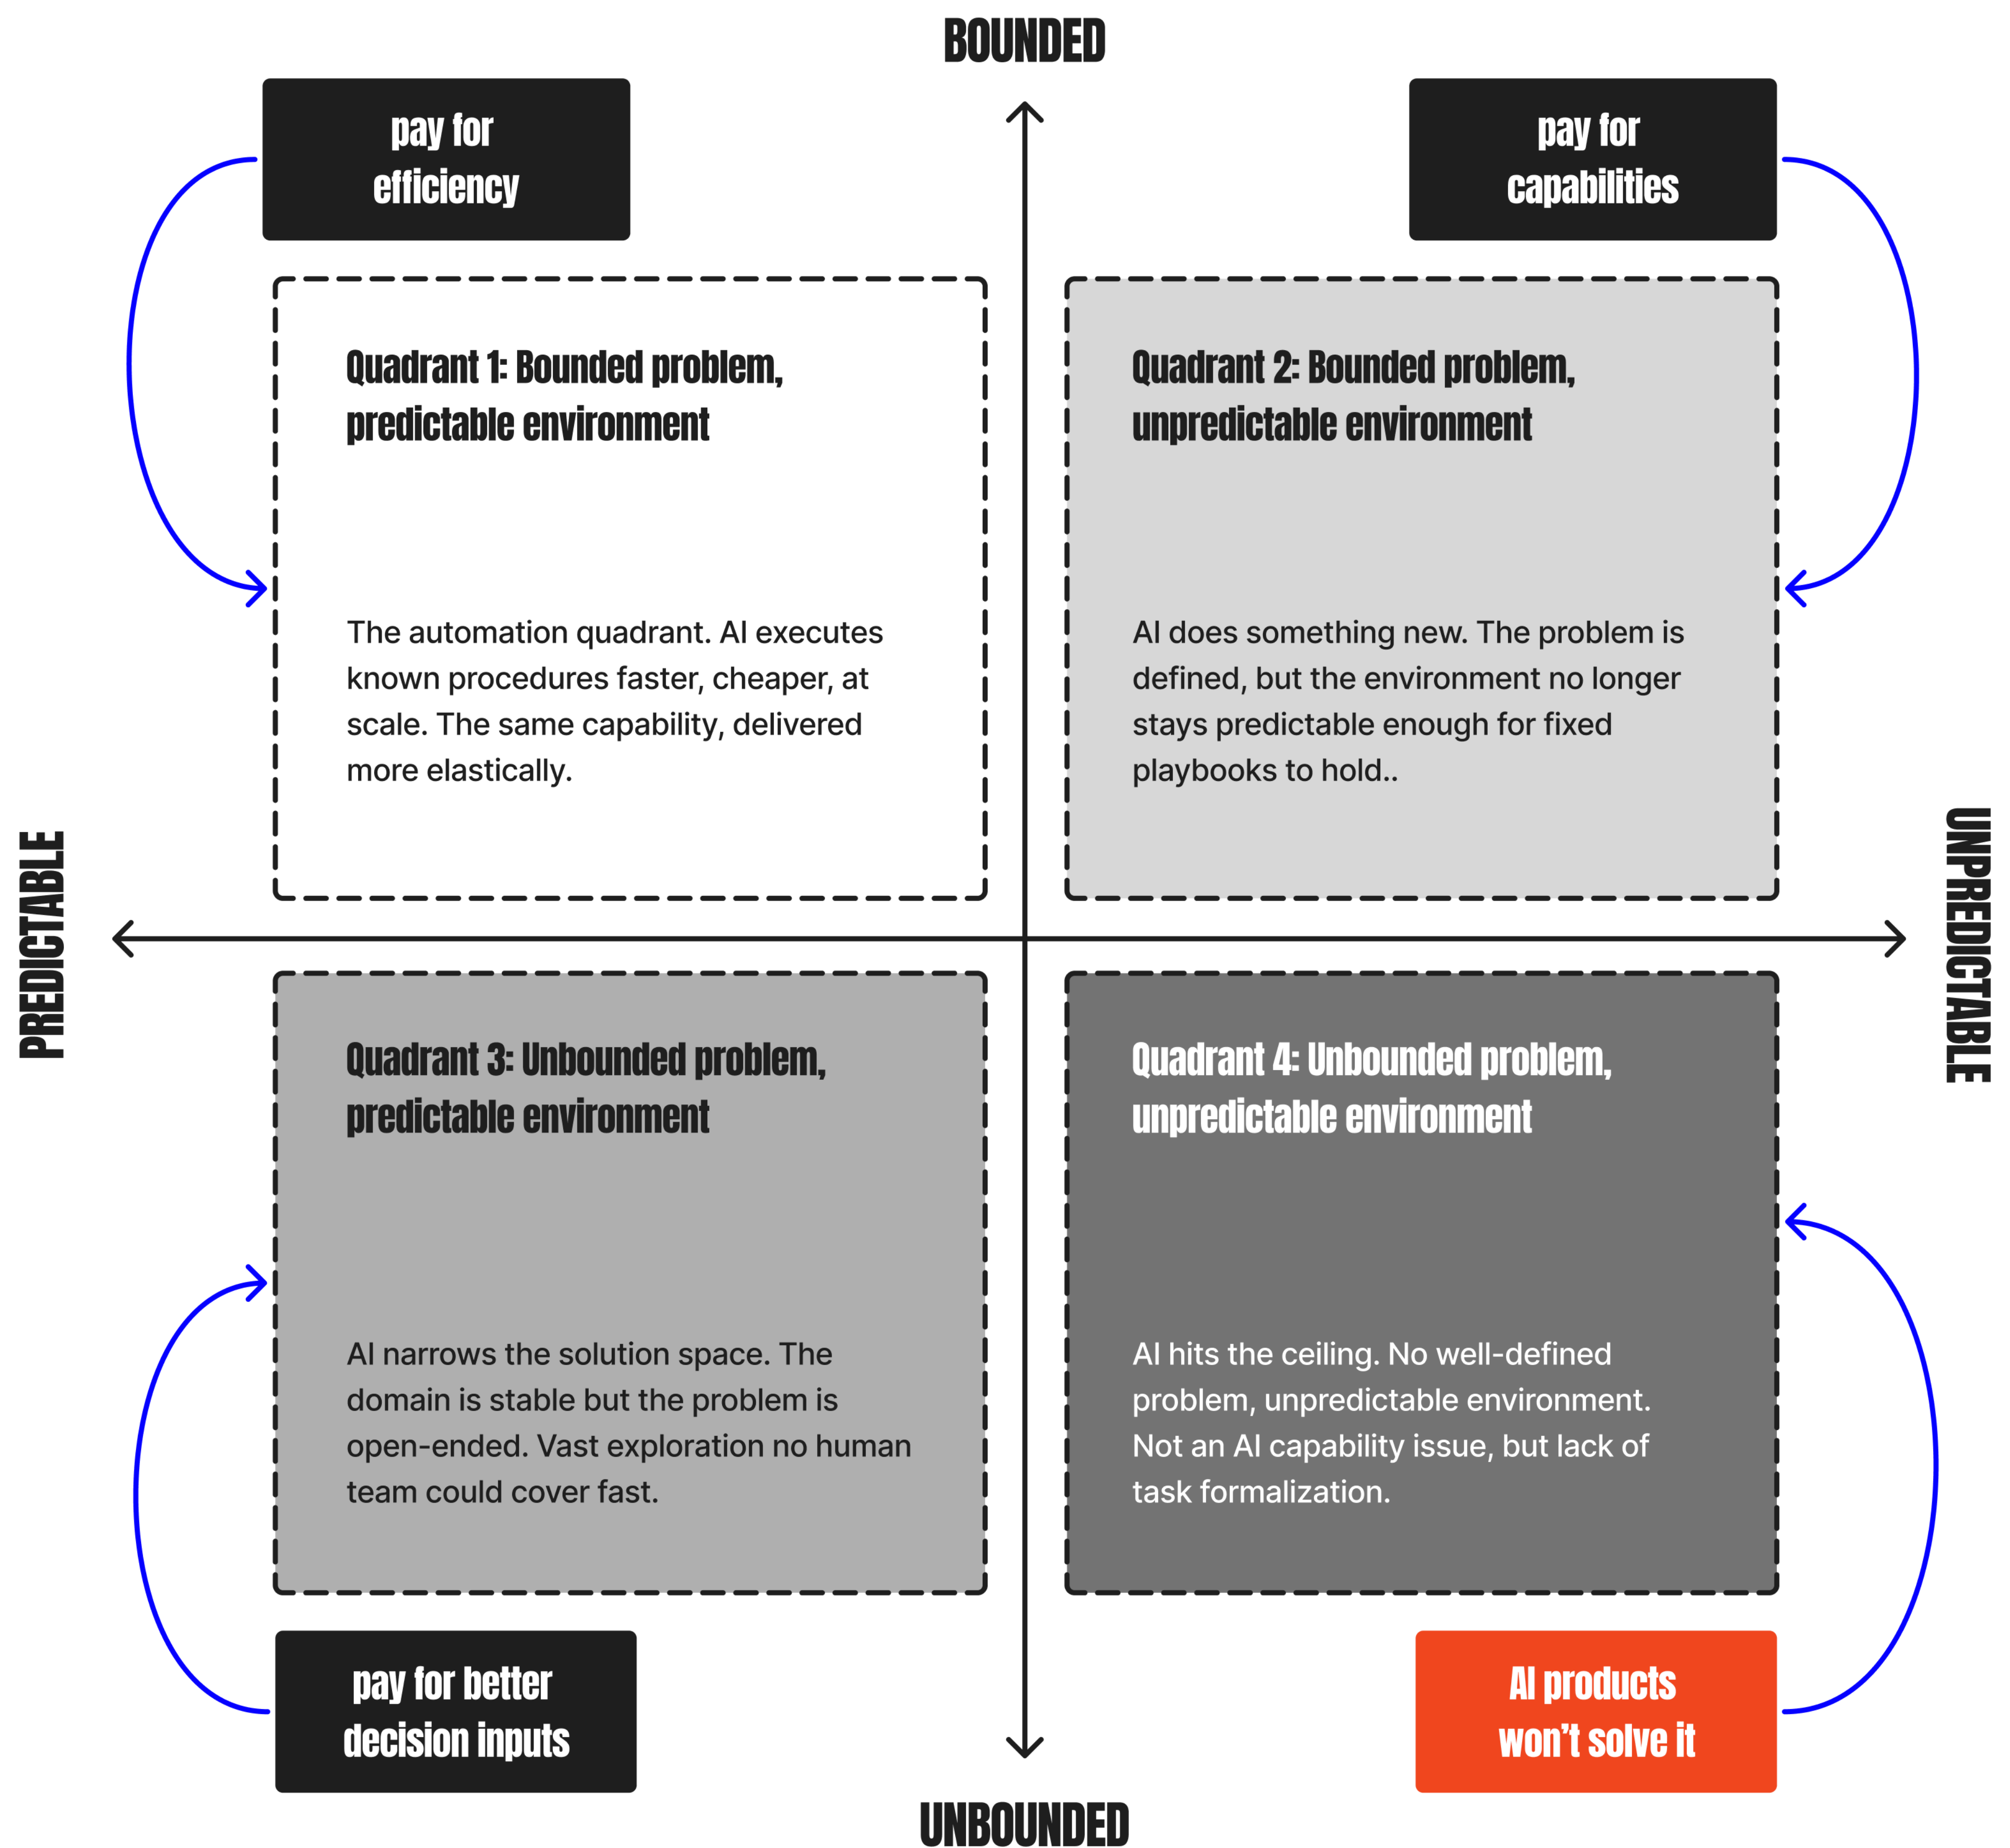
The AI product landscape sorts into four quadrants based on two questions:

- Is the problem bounded or unbounded?
- Is the environment predictable or unpredictable?

Most AI products live in the first three quadrants. The fourth, unbounded problem, unpredictable environment, is where AI hits a ceiling.

This is not a capability gap that closes with better models, but a permanent boundary: AI cannot take ownership of problems that haven't been formed yet.

The most consequential decisions usually begin there: before the problem is defined, before the variables are clear, before any tool can take over.



Where AI products work, sorted by problem type and environment stability.

The AI product matrix

EXPLAINER

- *Bounded:* The problem parameters are clear, the steps are known, playbooks exist.
- *Unbounded:* The problem requires investigation, the path to solving it is unclear, the problem itself may be ill-defined or too vast to formalize.

- *Predictable:* Humans or systems behave within a known enough range. The dynamics may vary, but the range of what might happen is knowable.
- *Unpredictable:* Humans or systems do not stay within a stable enough range for fixed playbooks to hold. They may react, shift, or actively defeat the model.

Quadrant 1: Bounded problem, predictable environment

This is the automation quadrant. AI does the existing thing faster, cheaper, and at scale.

The problem is already codified. Freight billing, compliance monitoring, appointment scheduling, logistics coordination. These workflows had established protocols and rules before AI arrived. The parameters are clear, the steps are known, playbooks exist.

The environment stays within a known range. The inputs can vary. A freight driver says something slightly outside the script, a compliance case has an unusual detail, but the variation is bounded. Nothing in the environment is trying to defeat the model or shift the rules mid-process.

Where SaaS required every path to be pre-coded, AI handles variation within the bounds without needing every branch explicitly programmed. It flexes inside the known range.

The most advanced version of this is agentic infrastructure: the user states an intent and the agent composes the entire execution path without an interface, manual configuration, or translation layer.

But the quadrant does not change. The problem is still bounded. The environment is predictable.

Quadrant 2: Bounded problem, unpredictable environment

This is the new capability quadrant. The problem is defined, but the environment no longer stays predictable enough for fixed playbooks to hold.

SaaS can track the data. It can give you a dashboard. It can send you notifications about your deal stalling. What it can't do is read the specific situation unfolding in real time and generate a response adapted to it. That is the threshold AI crosses in this quadrant.

EXAMPLES

- **Cybersecurity platforms** that build a behavioral baseline for each network and flag deviations attackers specifically design to avoid detection.
- **Sales platforms** that read the specific dynamics of each deal and generate a play adapted to that exact situation, instead of a universal playbook.
- **Visa's fraud detection system** builds a behavioral profile for each cardholder and catches patterns designed to look legitimate.

These AI products operate inside a defined problem space, but in environments that do not stay predictable enough for fixed rules to hold. This was previously the exclusive domain of human judgment. For the first time, AI can do a meaningful version of this work.

Quadrant 3: Unbounded problem, predictable environment

This is the exploration quadrant. AI is powerful here by narrowing the space.

The environment is stable and well-mapped. The rules governing the domain are known: the laws of chemistry don't shift, the genetic code doesn't reorganize, financial models operate within established logic.

But the problem is no longer bounded. There is no defined procedure to execute, no single correct path. The solution space is vast, and the question is open.

Examples:

- **AlphaFold** predicts 3D structures from amino acid sequences.
- **Chemistry42** explores chemical space for novel compounds, narrowing vast possibilities.
- **Anaplan** generates scenarios and runs what-if analyses across massive datasets.

These AI products deconstruct open-ended problems, explore multiple pathways, and extract probabilistic insight from overwhelming complexity. They narrow the possibility space.

Quadrant 4: Unbounded problem, unpredictable environment

This is where AI products hit a ceiling. AI requires a sufficiently defined problem and/or a stable enough environment to operate. In Quadrant 4, neither condition holds. The problem is still being formed, the environment is shifting while you investigate, and the relevant variables are still emerging.

CONSIDER THIS EXAMPLE

In the telecom A2P market, regulators set pricing rules designed for a world that no longer exists, hyperscalers shift traffic to cheaper channels the moment an operator raises rates, aggregators profit from the opacity that operators suffer from, and the chaos itself inflates KPIs in ways that make the real problem harder to see.

No single variable can be isolated. Models can't hold the situation stable long enough to optimize against it. The problem is still being formed by the collision of competing incentives, and the environment changes shape every time a player moves.

AI can assist with sub-tasks inside Quadrant 4 work: research, synthesis, pattern detection, and scenario generation. But there is a line it cannot cross. AI cannot take ownership of the problem-forming function itself.

That means deciding what the investigation actually is, determining which variables matter, and judging whether the current frame is revealing the situation or distorting it. Those acts depend on criteria that emerge in real time, from inside the situation, as the investigation unfolds.

AI operates on problems that can be sufficiently specified and/or in environments that stay stable enough to model. Quadrant 4 is defined by the absence of both conditions.

The problem cannot yet be specified, and the act of specifying it is itself the work. The environment is shifting while that work unfolds, rewriting the constraints before any solution can stabilize.

The bottleneck is not computation or AI capability. The problem does not yet exist in a form any system can optimize against, and the ground it sits on will not hold still long enough for one to try.

NOTE

Reality modeling: world models, simulation engines, systems that build and update internal representations of how environments behave, does push into the territory. They can handle environments that shift, because the model updates as the environment moves. That addresses half of the Q4 problem: the unstable environment part.

But the other half is harder. The problem-forming function: deciding what the investigation actually is, which variables matter, whether your current frame is revealing the situation or distorting it, that requires something prior to modeling. You have to know what to model. You have to judge whether the model you're building is capturing the right thing. And those judgments emerge from inside the situation, in real time, through contact with the mess.

Problems move

A real situation does not sit in one quadrant permanently. Problems migrate.

A crisis begins in Quadrant 4: unbounded, unpredictable. You investigate, diagnose, structure, locate the problem. Once the problem is defined and the variables identified, parts of it move to Quadrant 3, where AI can model scenarios and optimise across known parameters.

Once processes are established and codified, the work moves further into Quadrant 2 or Quadrant 1, where AI can operate autonomously within bounded conditions. And within Q1, the migration continues: what begins as a procedure a human executes becomes a procedure an agent executes, and eventually a procedure an agent composes from stated intent alone.

But migration also runs in reverse. A regulatory shock can take a bounded, automated process and blow it open. A competitor's move can invalidate the model your entire strategy was built on.

A market collapse can turn a well-understood domain into unfamiliar terrain overnight. What was Q1 yesterday can become Q4 today. The framework is not a ladder of progress, but a map of where a situation lives right now, and situations move in both directions.

The implication

The AI landscape is being sold as a capability story: what AI can do, what it will do next, and what it will replace. That framing is useful for investors and engineers, but it can be misleading for the executive who needs to make decisions now.

The more useful question is: where does AI's expansion stop?

It stops at the boundary of formalizability. When the problem cannot be specified in advance, the relevant variables are still emerging, the criteria for judging a good response only become visible through the investigation itself, and the environment is shifting while that investigation unfolds, no product can take ownership of the work.

Quadrant 1 buys efficiency. Quadrant 2 buys a genuinely new capability. Quadrant 3 buys better decision inputs. But the decisions that matter most often begin in unbounded, unpredictable settings.

How to use this matrix

Three applications for the executive:

- **When evaluating an AI product:** Locate it on the matrix. Ask which quadrant it serves. A Q1 product that claims to solve Q4 problems is overselling. A Q3 product being purchased for Q1 efficiency is misallocated. Match the tool to the quadrant.
- **When diagnosing a situation:** Identify which quadrant your problem currently lives in. If you're in Q4, stop looking for a product to own it, staff it with human judgment and use AI for the sub-tasks. If the problem has migrated to Q3 or Q2, that's when tools become load-bearing.
- **When allocating resources:** Map your spending across quadrants. Most organizations over-invest in Q1 automation and under-invest in the human capacity to handle Q4. The highest-risk situations live where no product can help, and that's usually where the least structured support exists.