



The Next Generation of Forecasting

Companies have made their forecasts more accurate. Far fewer have made them more useful. Closing that gap is now the central planning challenge.

AT A GLANCE

- Planning technology has advanced considerably faster than the planning operating model, and the distance between the two is where most of the expected value is lost.
- The decisive metric is no longer forecast error. It is the lag between a material change in assumptions and a corresponding change in a decision — a number most companies have never measured.
- AI in planning today is augmentation, not autonomy: it accelerates preparation, analysis, and insight, while human judgment still owns the decision.
- Accuracy gains of 20% to 40% are achievable, but they improve business performance only if they change decisions — and that depends on four things that have little to do with the algorithm: context, trust, integration, and adoption.

It is a question that is increasingly putting planning leaders on the defensive: if the forecasting models are so much better, why is the plan still wrong by March? And why, when everyone can see that it is wrong, does so little change until the next scheduled review?

The answer is that most companies have solved the wrong half of the problem. Over the past decade, demand and financial forecasting have improved substantially — driver-based models, machine learning, external signal integration. What has not improved at anything like the same rate is the speed at which a changed forecast becomes a changed decision. Planning has become more accurate without becoming more responsive, and in volatile categories, responsiveness is worth more than precision.

This distinction matters because it reassigns the problem. A forecast error is a data science issue. A six-week lag between knowing and deciding is an operating model issue, and it is owned by the business, not by the analytics team. The next generation of forecasting will be defined by companies that treat it as the latter.

The market moves continuously, but planning is still periodic.

Consider the shape of a typical planning year. Consumer behavior, category dynamics, competitive moves, customer negotiations, macroeconomic conditions, and technology shifts all move continuously. Against that, the planning calendar offers a handful of moments: the annual plan, a first-quarter review, a mid-year reforecast, perhaps a third-quarter course correction.

Between those moments, the company is navigating on a map drawn at a point in time. The plan is not wrong when it is written; it ages, and the rate at which it ages has accelerated. A quarterly cadence was adequate when the underlying volatility was quarterly. For most consumer, health, and pharmaceutical businesses, it no longer is.

The consequence is not only slow reaction. It is an erosion of the plan's authority. Teams know the numbers are stale, so they build private spreadsheets, local assumptions, and workarounds. Decisions migrate out of the planning system and into email threads and side meetings. But a decision made outside the system is a decision the system can never learn from.

The market moves continuously. Planning still responds periodically.

Signals shift every week; the planning calendar opens a handful of windows a year.



Exhibit 1 | The operating environment moves continuously; the planning calendar opens only a few windows a year.

Technology has outpaced the operating model

The instinct is to close this gap with more technology, and most large companies have. Advanced planning systems, machine learning forecasts, and increasingly digital twins of the supply chain or the commercial P&L are in production somewhere in the enterprise.

The operating model around them has barely moved. Planning still runs on periodic cycles. Data still crosses functions through manual handoffs. Decisions are still made, and frequently unrecorded, outside the tools that produced the analysis.

Planning technology has advanced faster than the planning operating model

The distance between the two is where value quietly leaks away.




Exhibit 2 | The value realization gap: planning technology has advanced faster than the operating model around it.

This is the value realization gap, and it explains a pattern many executives will recognize: a forecasting capability that demonstrably works in a pilot, delivers measurable accuracy gains, and still does not move the company's financial trajectory. The model improved, but the decision process did not.

Technology creates value when it becomes part of how the company plans and decides — not when it sits alongside that process producing better inputs nobody has time to act on.

There is a structural reason this recurs. Forecasting improvements are typically delivered as a project with a technical owner and a technical success metric, while the planning process they are meant to improve belongs to finance, commercial, and supply functions whose calendars, incentives, and handoffs the project never touches. A new capability is dropped into an unchanged process and, predictably, the process wins. Six months later the model is still running, the planners are still rebuilding their own spreadsheets, and the business case has been quietly reframed as efficiency.

Three stages of planning maturity: periodic, dynamic, continuous

It is worth being precise about what "next generation" means, because the term is currently used to describe at least three different states.

Traditional periodic planning extrapolates history, assembles functional plans, tests a small number of fixed scenarios, and makes decisions outside the system. Most companies still have part of their planning estate here.

Dynamic planning is what good looks like today: driver-based forecasts rather than extrapolation, internal and external signals integrated into a single view, scenario modeling that can be run on demand, and genuine cross-functional alignment on one set of assumptions. This is achievable with current technology, and it is where the leading consumer and pharmaceutical companies are concentrating their efforts today.

A continuous decision system is the step change – and a future promise. Planning becomes a closed loop – sense, forecast, simulate, decide, learn – running continuously rather than on a calendar. Replanning is triggered by material change rather than by date, and decisions and their outcomes are captured inside the system and feed the next cycle.

From producing a plan to continuously navigating the business

Each step adds capability; the last one changes the operating model itself.




Exhibit 3 | Three stages of planning maturity, from periodic planning to a continuous decision system.

The critical difference sits in that last word. In periodic planning, every cycle starts over: assumptions are rebuilt, the rationale behind last quarter's call is lost, and institutional knowledge lives in whoever still holds the role. In a continuous decision system, learning is built into the system rather than restarted every cycle. Over one to three years, that compounding knowledge is worth more than any single-point improvement in model accuracy.

What AI actually does in planning today

There is considerable noise about autonomous planning agents. The current practical picture is less sexy, but more useful. AI is augmenting planning work; it is not running it.

In sensing, AI monitors internal and external information and flags anomalies and material changes. In the driver layer, it summarizes and interrogates data and explains what moved a forecast and why. In scenarios, it generates inputs and accelerates what-if comparison. In decisions, it surfaces risks and implications and prepares recommendations for human review. In learning, it captures decisions, rationale, and outcomes.

What these have in common is that AI is accelerating preparation, analysis, and insight — the work that consumes most of a planner's week — while human judgment still owns the decision. In our experience, that division is not a transitional compromise. The companies getting the most from AI in planning are not those trying to remove people from the loop; they are those using AI to give people materially better options, far earlier, with the reasoning visible horizon.

This also explains why the current generation of planning assistants matters more than it first appears. A tool that answers "what changed in the July forecast for our top three segments, and why?" in under three minutes does not replace a planner, but it removes the two days of data assembly to get the answer, freeing planner's time to do something more valuable.

The capability sorts into five stages, and the sequence is not optional.

AI-enabled planning evolves in stages — each layer builds on the one below it

Stages one and two are real today. Three and four are emerging. Five is the destination.



Exhibit 4 | AI-enabled planning evolves in cumulative stages; each layer depends on the capabilities below it.

Task assistance and workflow augmentation are real and scaling now. Specialized and coordinated planning agents are emerging: agents that monitor a defined set of signals, challenge assumptions, run scenarios, recommend actions for review, and exchange context across functions to resolve cross-functional trade-offs. The continuous decision system sits above them.

The path is cumulative. Advanced AI planning depends on the data, context, workflows, and trust established at the stages below it.

Companies that have attempted to skip to coordinated agents without clean data, shared context, integrated workflows, and established trust have ended up rebuilding the earlier stages afterward, under worse conditions and with less patience from the business.

The hardest barriers are not the AI

In the programs we see, AI capability is rarely the biggest constraint. Instead, it is the human factors that present hard challenges

The hardest barriers to scaling AI-enabled planning are usually not the AI

Where AI-enabled planning stalls, across the sense → analyze → decide → execute flow.




Exhibit 5 | The four barriers that most often stall AI-enabled planning: context, trust, integration, and adoption.

Context. Data arrives late or inconsistently, external signals stay disconnected from internal ones, and the AI has no reliable picture of the enterprise it is reasoning about. An AI without context produces fluent, confident, unusable output.

Trust. Forecast changes are hard to explain, assumptions are hard to trace, and recommendations arrive without a rationale a planner can interrogate. Explainability is not a compliance nicety here; it is the mechanism by which a recommendation becomes a decision. If a finance director cannot see which driver moved and by how much, the number does not get used.

Integration. Planning systems are disconnected, AI insights sit outside the workflow in a separate tool or deck, and manual handoffs reintroduce precisely the latency the AI was bought to remove.

Adoption. Decision rights are unclear, teams work from different assumptions, and decisions are not captured — so the loop never closes and the system never learns.

These are organizational and architectural problems rather than model problems. AI capability alone does not create planning value; context, trust, workflow, and learning do.

What planning leaders should do

Three priorities follow for anyone accountable for a planning transformation.

Measure response time, not only error. Forecast accuracy is necessary but not sufficient. The more revealing metric is the lag between a material change in assumptions and a corresponding change in a decision — pricing, promotion, supply, allocation. Most companies have never measured it. The number is usually sobering, and it is the number that moves the P&L.

Treat explainability as an adoption requirement. The practical test is whether a planner can move from a total forecast to the specific driver impacts behind it — execution, external factors, consumer trends — and then see where an action would change the outcome. A forecast that says value sales will grow 51% is a claim. A forecast that is transparent in attributing that growth to impact of price, distribution, media spend, and consumer sentiment, is something a commercial director can argue with or align on. This “why” behind the numbers is what gives forecast credibility and drives adoption.

Keep the receipts: capture decisions where they are made. The highest-return change available to most planning organizations is also the least technical: record the decision, the assumptions behind it, the alternatives considered, and the outcome, inside the planning environment. This is the raw material for institutional learning and the precondition for any agentic capability that follows.

The pattern is consistent across the markets and categories we work in. Accuracy gains of 20% to 40% are real and material, but they only translate into better business performance when they change decisions. The companies that make that conversion have shortened the distance between knowing and deciding. And bridging this distance from knowledge to decision is ultimately the point of next-generation forecasting.

Lana Klein is Practice Head, Decision Analytics at Cloverpop, where she advises global consumer, health, and pharmaceutical companies on AI-enabled forecasting, planning, and decision intelligence. This article is adapted from a Cloverpop executive roundtable on the next generation of planning.



Contact Us

Lana Klein

Practice Head, Decision Analytics

lane@cloverpop.com



Connect on LinkedIn



www.cloverpop.com