PHARMACEUTICAL EXECUTIVE®

# Fast Is Not the Same as Trustworthy: The Data Standing Behind an Effective AI Answer

Exploring the emergence of agentic ecosystems in transforming data accessibility and ensuring reliable answers reach the people making decisions while there is still a choice to make.



*about the authors*

**Kapil Chaddha,** senior director, new product development

**Sarah Bickley,** senior director, global market insights offering

Both with IQVIA

PHARMACEUTICAL COMPANIES ARE not short on market data. Most hold more than their teams can practically use, including sales trends, promotional spend and effort, health care professional (HCP) sentiment, competitive intelligence, pipeline intelligence, loss-of-exclusivity timelines and forecasts.[1] Over the past two years, the industry conversation has focused on what artificial intelligence (AI) could do with all that data. While that question has largely been settled, a more difficult one has emerged: Can the answers be trusted?

Trust is the keystone that holds everything together. When people believe an answer is reliable, more of them use it and decisions move faster. When they don't, the data stay with a few specialists. A large language model (LLM) cannot close that gap alone. The LLM is the brain trained on a vast body of general text, which makes it fluent without making it expert. An AI agent is what gets built around it. The agent uses the LLM to interpret a question, then adds what the model alone lacks, including goals, instructions, a plan for answering and access to the right data. Point that agent at proprietary, well-governed life sciences data rather than generalist sources, and it can support stronger, faster commercial decisions. Data accessibility has far less to do with the interface than with the foundation sitting behind each answer.

## A GOOD ANSWER SUPPORTS MORE ROBUST DECISION-MAKING

Ask a commercial leader what a genuinely usable answer looks like, and the criteria come quickly: relevance, speed, accuracy and data that are connected

rather than scattered. Most of those criteria serve one underlying test: Does the answer help someone make a better decision?

Today, most commercial teams route questions through an analyst who can move between disparate data sets to assemble a view. When a natural-language interface sits on top of connected, validated data, that assembly work no longer requires a specialist or, at minimum, takes a fraction of the time.[1] Someone with no experience in the underlying data can ask a question the way they would ask a market expert, then keep asking follow-ups instead of filing a request and waiting days. Analytical cycles that once took weeks now run inside a single session, moving effort out of planning and analysis and into in-market execution.

The returns compound as adoption extends across the organization. When commercial, medical and other functions query the same connected foundation, they stop reconciling competing extracts and work from a shared source of truth.[1] Only verifiable insights drawn from multiple data sets give users the confidence to defend a recommendation in front of an executive team.

## GOOD ANSWERS ARE HARDER TO REACH THAN THEY LOOK

The data behind a commercial decision rarely live in one place, and the pieces that carry the most signal are often the least structured. Verbatim feedback is a good example: free-text comments on what HCPs recall from interactions with field teams. Any single comment is easy to read. The value sits in comparing them at volume across brands, countries and time. With tens of millions of verbatim records collected across markets for more than a decade, the volume ran well past what manual review could cover, and even a small project to organize HCP feedback by product or country could take weeks.

AI agents have changed the economics of that work. The computing power behind an agentic ecosystem can read unstructured feedback at full scale and set it alongside structured data, which was never practical before. A question about HCP message recall can now be interrogated at scale rather than from a few thousand comments an analyst tagged by hand. The stakes are direct, since the HCP voice translates into actionable commercial intelligence.

The same shift applies to the broader problem of combining sources. Any single analysis, whether a sales trend or promotional spend, has long been possible in isolation. Stitching them together may have consumed weeks. When connected data let those sources be read together, a question like "Why are my patient shares falling?" can surface access barriers, competitor activity and channel performance next to the underlying trend. Doing that reliably is the hard part, and it is where the complexity of a multisource data ecosystem really sits.[1]

## WHY "CLEAN DATA" ARE NOT ALWAYS AGENT-READY

Many organizations assume data maintained to existing industry standards are automatically agent-ready. But foundations built for the "business intelligence" era do not automatically serve an agentic one. Sales data can be standardized across nearly 100 countries and still not work well with an agent. That standardization is valuable and hard won, but it is built for a different job, because preparing data for reporting is not the same as preparing it for reasoning. An agent-ready foundation is one in which data sets can speak to each other, so a question that spans promotion, sentiment and sales doesn't stop at a system boundary.

The knowledge base is what an agent consults to understand your business. It sits around the model rather than inside its training, and it is where most of the real work happens. This includes life sciences context and data dictionaries deep enough to recognize how the same product or acronym appears across every source, since one abbreviation can mean different things in different therapeutic areas.[1] Taxonomies contribute as well, some drawn from public clinical standards and orchestrated for commercial data sets they were never designed to touch. Then there are the business rules,

which sit on harmonized data built against trusted master data and let several specialist agents work under one orchestrator without contradicting each other. There can be thousands of them, keeping each agent inside its purpose instead of guessing. Guardrails work alongside them, defining the questions an agent should not attempt and the data sets it should not touch.

## OVERSIGHT IS BUILT IN, NOT SWITCHED OFF

A trustworthy agentic ecosystem is designed on the assumption that its answers will be checked. Automated checks can evaluate whether a response stays within the defined business rules and guardrails and whether it drew on the right data sets. Those checks catch a great deal but still miss things, which is why human review stays in the operating model. Analysts validate the outputs of each data refresh cycle and field the cases the automated checks let through. Each correction strengthens the guardrails. Accountability for decisions remains in the hands of commercial leaders.

This is what the technology is for. It gives people better insight faster, so their time goes to judgment and to execution in the market instead of to assembling the analysis.

## CONFIDENCE IS STARTING TO SPREAD

The clearest evidence that trust is taking hold shows up in behavior. A year ago, many organizations were still experimenting with agentic ecosystems. That has shifted toward execution. Commercial teams have begun opening these agents live in meetings, running the analysis in the room rather than coming back with an answer a day or two later. Elsewhere, brand leaders are asking for daily market scans, a capability still in development, so field teams can adjust plans as events unfold rather than working to a schedule set weeks earlier.

*It gives people better insight faster, so their time goes to judgment and to execution in the market instead of to assembling the analysis.*

Neither behavior takes hold without confidence in the answers, and each one compresses the distance between a question and a decision. That should be the measure of data accessibility in pharma: whether reliable answers reach the people making decisions while there is still a choice to make.[1] Organizations that invest in connected data, a strong knowledge base and ongoing human oversight reach better decisions faster and from more of the people who need to make them. PE

**Reference**

1. Chaddha K, Jaiswal A. Agentic AI and data accessibility: a strategic shift for life sciences across the product lifecycle. IQVIA Blog. October 2, 2025. https://www.iqvia.com/blogs/2025/10/agentic-ai-and-data-accessibility