



NAVIGATING THE AI CONFIDENCE TRAP

Ensuring strategic success
through validated insights
and expertise

JP Park



At Ipsos, we champion the unique blend of Human Intelligence (HI) and Artificial Intelligence (AI) to propel innovation and deliver impactful, human-centric insights for our clients.

Our Human Intelligence stems from our expertise in prompt engineering, data science, and our unique, high-quality data sets, which embed creativity, curiosity, ethics, and rigour into our AI solutions. Our clients benefit from insights that are safer, faster and grounded in the human context.

#IpsosHiAi

Executive summary

Artificial Intelligence (AI) has become an essential accelerant for pharmaceutical commercial strategy - enabling teams to synthesize market landscapes, generate hypotheses, and compile competitive intelligence. At Healthcare at Ipsos, we have embraced these capabilities and are actively deploying AI across our research practice to deliver faster, richer insights for our clients.

Yet AI's greatest strategic danger isn't the obvious errors - it's the subtle misdirection wrapped in eloquent prose. Well-written, defensible outputs that miss critical nuance create a confidence trap. The confidence trap

can lead to especially severe consequences when teams make critical business decisions on narratives that may sound right but, in reality, rest on quicksand.

Through three real-world inspired examples - a chronic kidney disease (CKD) drug blindsided by disruption, an atopic dermatitis launch crippled by payer restrictions, and a multiple myeloma therapy trapped in the wrong line of treatment - we reveal how AI, specifically Large Language Models (LLMs) can miss what matters most: the unwritten rules, power dynamics, and counterintuitive realities that determine commercial success.

Key takeaways:



AI delivers genuine value today for market orientation, hypothesis generation, competitive landscaping, and synthesis - these are capabilities we routinely use and also recommend our clients leverage



The confidence trap emerges when teams mistake fluent, well-structured output for validated, strategic insight



Critical blind spots persist in areas requiring real-time intelligence, payer dynamics, and treatment sequencing logic that LLMs cannot reliably infer



Ipsos' HI+AI approach integrates AI's speed with dedicated subject matter experts, domain expertise, strategic frameworks, and rigorous primary research to ensure insights and recommendations rest on validated truth



Healthcare at Ipsos' AI capabilities - from Healthcare PersonaBot to agentic workflows - help our clients move fast without compromising the rigorous validation that healthcare decisions demand

Introduction

Consider the following scenario: A commercial strategy team at a life sciences company is tasked with building a market assessment for a new indication. By lunch, they have 40 pages of analysis that seem comprehensive, well-structured, and defensible. Across biopharma, leveraging LLMs has quietly become the default first step in their strategic workflow - and for good reason. We actively encourage this approach.

When used correctly, LLMs are an unparalleled tool for accelerating foundational work by:

- **Accelerating Market Orientation:** Rapidly synthesizing public data to outline market landscapes
- **Generating Initial Hypotheses:** Creating a broad set of potential strategic directions or positioning ideas to be tested
- **Compiling Competitive & Trial Data:** Drafting comprehensive lists of competitors, assets, and ongoing clinical trials
- **Insight Synthesis:** Generating summaries from large amounts of transcripts, literature, or interview notes

Our teams leverage AI to compress weeks of desk research into hours, which, after validation, frees capacity for the higher-value work of interpretation, and strategic synthesis that our clients depend on.

However, as organizations increasingly leverage a mix of proprietary AI tools and LLMs, a critical risk emerges. The output - while coherent and directionally logical - can often still benefit from deeper contextual grounding. Narratives may appear complete, but without alignment to real-world dynamics, critical nuances can be missed, creating a false sense of confidence with narratives that may appear bulletproof but have not yet been battle-tested.

This is where speed without Human Intelligence (HI) becomes expensive gambling - and where the HI+AI approach proves essential.



Across biopharma, leveraging LLMs has quietly become the default first step in their strategic workflow - and for good reason.

Where AI Delivers Real Value

Before examining the confidence trap, it's important to acknowledge the substantial value AI already provides in healthcare commercial strategy. This isn't a technology to fear or resist - it's a capability to harness wisely.



Landscape Synthesis at Scale

AI excels at synthesizing vast amounts of public information into coherent narratives. A team preparing for a therapeutic area entry can generate comprehensive overviews of approved treatments, pipeline assets, clinical trial activity, and published guidelines in hours rather than weeks.



Hypothesis Generation

LLMs are remarkably effective at generating diverse strategic hypotheses for testing. Rather than anchoring on one or two positioning ideas, teams can rapidly explore dozens of potential directions - then apply human judgment to identify which merit validation.



Document Synthesis and Summarization

From advisory board transcripts to competitive intelligence reports, AI dramatically accelerates the synthesis of qualitative inputs into actionable summaries.



Competitive Intelligence Compilation

Building comprehensive competitive landscapes - tracking pipeline assets, trial designs, and corporate announcements - is precisely the kind of structured synthesis where AI shines.

The Confidence Trap: Real-World Inspired Examples

The following examples, based on real assets and market dynamics, highlight common areas where alignment between AI-generated insights and human expertise is critical. Rather than cautionary tales, these examples illustrate where integrating domain expertise and real-world validation ensures strategies are both robust and actionable.

Example 1: Yesterday's News Before Launch



The Setup: A pharma company developing an anti-inflammatory for CKD asks an LLM for competitive positioning. The output is a pristine - 20 slides mapping every approved and competitive pipeline asset, highlighting the new drug's unique mechanism as the key differentiator.



The Blind Spot: Buried in a recent American Society of Nephrology presentation, a small biotech showed something never achieved in this indication: not just slowing disease progression, but a near halt to it. Because this groundbreaking data was shared only in a conference presentation and not yet captured in widely accessible literature, the LLM failed to see that the strategic question had fundamentally changed from "Who is our competition now?" to "What disruption is coming next?"



The Fix: One afternoon of in-depth KOL interviews would have surfaced this. Every thought leader has heard about the biotech that had seemingly halted disease progression. A human strategist would have asked not "What's approved?" but "What has gotten everyone excited?" The entire strategic conversation would have shifted from "How do we launch?" to "Given this disruption, where can we uniquely win?" - driving decisions focused on specific patient subgroups or proving value in a way that would make their drug complementary, not obsolete. Ipsos's expertise in leveraging AI strategically ensures we identify when it enhances insights and when human judgment is paramount, avoiding potential pitfalls. In cases where AI is best left aside, our seasoned analysts provide a nuanced understanding to guide optimal decisions. The strategic focus would have pivoted from executing a flawed launch to defining a new path forward, saving millions in marketing spend and redirecting R&D investment to maximize future success.

AI analyzed existing data well but lacks access to real-time field insight and the context to properly weigh it.

Example 2: Clinically Sound, Commercially Naive



The Setup: A biopharma company is building its go-to-market strategy for a new biologic for severe atopic dermatitis. Based on the clinical data, they anticipate a list price of \$40,000 annually. The team asks an LLM to forecast the addressable market based on clinical guidelines, which recommend the drug for patients who have failed one systemic therapy. The LLM's output is compelling: a 500,000-patient opportunity and a path to a \$2B+ blockbuster.



The Blind Spot: The LLM's output is clinically logical but commercially naive because it cannot grasp the inherent tension between clinical guidelines and payer cost-containment. Despite the guideline recommendations, the LLM misses that a major national payer has a pre-existing internal policy for the drug class. This "step-edit" policy requires patients to fail two cheaper systemics first - an unwritten rule that reveals the payer's aggressive management philosophy. This single nuance cuts the immediately addressable market by more than half.



The Fix: A market access expert knows the most critical question isn't "What do guidelines say?" but "What is the payer's unwritten philosophy?" - a layer of insight that cannot be reliably inferred from published data and that current LLMs struggle to capture. Translating this insight into action, a market access expert would conduct targeted, in-depth interviews with pharmacy directors from key payer accounts. These conversations would quickly surface the restrictive step-edit philosophy, enabling the strategy to pivot towards a best-in-class patient support program designed specifically to help physicians and patients navigate complex prior authorization and appeals processes.

AI produced a logical analysis - but commercial reality operates on rules that aren't written in guidelines.



Rather than cautionary tales, these examples illustrate where integrating domain expertise and real-world validation ensures strategies are both robust and actionable.

Example 3: The Illusion of Fourth-Line Therapies



The Setup: A drug with a novel mechanism for multiple myeloma has Phase I data suggesting a strong efficacy signal in heavily pretreated patients who have received both BCMA CAR-T and a GPRC5D bispecific antibody therapy. An LLM recommends pursuing a registrational trial in this heavily pretreated patient population (typically 4L or later) - a conclusion that is clinically sound, guideline-aligned, and addresses unmet needs, but is strategically incomplete.



The Blind Spot: Historically, patients progressed sequentially through multiple lines of therapy over several years, eventually reaching 4L. However, the advent of CAR-T and bispecific antibodies has disrupted this dynamic. Patients now benefit from long remissions after receiving these therapies. In parallel, the population of patients who have received both modalities remains limited. As a result, even clinically compelling therapies can now struggle commercially if positioned too late, making earlier-line entry a critical strategic imperative for driving meaningful adoption.



The Fix: An experienced myeloma expert would likely have said: "Compelling data, but positioned too late to create meaningful value." The challenge is structural - the patient population is smaller, treatment pathways are more entrenched, and operational and logistical barriers become more pronounced in a fragile, heavily pretreated population. The entire clinical and commercial strategy should focus on one goal: rapidly generating 2L data to unlock access to a broader patient population and greater commercial upside. The strategic discussion would immediately shift to "What clinical trial do we need to run to secure a 2L indication?" and "Which KOLs can we engage to begin early use in the 2L setting based on this promising data?"

AI followed clinical logic well - but missed the commercial physics of a rapidly evolving treatment landscape.



Common Patterns: Where LLMs Need Human Partnership

In these examples, LLMs showed limitations across four common areas:

The "Plausible but Wrong" Scenario:

Building clean narratives on untested assumptions that can invalidate the strategy

Assessing Source Credibility:

Failing to evaluate which data sources are reliable, or to reconcile conflicting information

Understanding Implicit Logic:

Overlooking unwritten, category-specific nuances such as treatment sequencing or referral dynamics

Smoothing Over Critical Details:

Simplifying complexity and glossing over edge cases where value is created or lost

These aren't failures of AI - they're boundaries of current capability. Recognizing these boundaries is what separates strategic AI use from the confidence trap.



Recognizing boundaries is what separates strategic AI use from the confidence trap.



At Healthcare at Ipsos, we're not simply observers of AI's potential - we're active practitioners.



Our Commitment to AI-Enabled Research

At Healthcare at Ipsos, we're not simply observers of AI's potential - we're active practitioners. Our investments include:

- Healthcare PersonaBot that enables interactive exploration of target audiences
- AI-augmented qualitative research using intelligent probing and synthesis tools
- Synthetic data capabilities for extending sample reach and filling analytical gaps
- Agentic AI workflows (in development) that aim to perform preliminary validation before human review

The question becomes how to utilize these tools wisely in a domain where research findings inform treatment protocols, resource allocation, and ultimately patient outcomes. We apply a disciplined process that ensures speed is always anchored in validated truth:

Frame the Right Problem

We begin where AI cannot - by applying deep industry expertise, proven strategic frameworks, and cross-functional advisory experience to define the right questions, decision context, and hypotheses.

Accelerate with AI

With this expert-led foundation in place, we utilize AI to rapidly synthesize data, map the landscape, and generate an initial narrative - compressing weeks of work into hours.

Pressure-Test What Matters

We then systematically interrogate the output - challenging assumptions, identifying where the narrative is most fragile, and surfacing the critical gaps where AI is most likely to be directionally wrong.

Rebuild with Validated Intelligence

This is where validated primary research becomes essential - and where Ipsos' combination of access, scale, and domain expertise adds particular value. We design and execute targeted primary research at scale, leveraging direct access to key decision-makers including KOLs, patients, caregivers, and payers to validate, refine, or overturn the initial narrative. These real-world insights don't just refine the narrative - they redefine it, enabling a strategy rebuilt around real-world decision drivers and designed to succeed under actual market conditions.

This is not an AI-led analysis with human review. It is **a human-led, AI-accelerated, and expert-validated system** designed to ensure that what is generated quickly is also right.

The outcome is built on validated primary research and integrated expertise. As our colleagues in synthetic data research have noted, the market has seen an influx of vendors making expansive AI claims without rigorous validation. The questions to ask are about validation methodology and domain expertise, not technical sophistication.

The Vision for Tomorrow: From Fast Intern to True Partner

While human validation is the essential safeguard for today, it is not the end of the story. The limitations we see are often not inherent to the technology itself, but rather to how it is being used - as a simple text synthesizer. The future lies in building more sophisticated AI systems that internalize the very logic of strategic validation.

At Healthcare at Ipsos, we are actively exploring this direction.

Our teams - including those developing advanced AI solutions - are finding that by designing Agentic AI workflows, we can move beyond the first pass toward systems that perform preliminary sanity checks:



Evidence Hierarchy by Design: Instead of just summarizing, next-generation agents follow a strict evidence hierarchy - identifying where internal data is thin, automatically triggering searches for external intelligence, and weighing the final insight based on the reliability of the source - privileging a peer-reviewed publication over a press release, for example.



Context-Aware Logic: By building deep human intelligence and strategic frameworks directly into these agentic workflows, we can teach the AI to account for the category-specific sequencing, institutional dynamics, and unwritten rules that standard models miss.





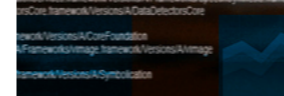



From Intern to Partner: This elevates AI from a “fast intern” that drafts slides to a “strategic partner” that actively flags hidden risks and surfaces subtle nuances - essentially performing its own pressure test before a human even sees the first draft.

Our most recent progress in agentic AI has moved well beyond simple exploration, having successfully established frameworks and designed multi-agent ecosystems tailored for market research in life sciences.

Through successful proof-of-concept engagements, we are demonstrating unparalleled analytical depth and actionable insights, thereby solidifying Ipsos’ position as a leader in this transformative space.



The future lies in building more sophisticated AI systems that internalize the very logic of strategic validation.







Looking Ahead: AI Tools for Healthcare Market Research

While this paper focuses on AI in strategic analysis, a parallel challenge exists in market research: What happens when advanced AI is used to augment or simulate human responses?

In an upcoming piece in this series, our Healthcare Data Science team will explore the practical application of these principles to synthetic data, persona bots, digital twins, and AI-powered research tools.

For those working with AI-enhanced research methodologies, that piece offers:

-  Honest assessment of what current tools can and cannot reliably deliver in healthcare contexts
-  Practical guidance on where synthetic data adds genuine value today
-  The SURE framework for evaluation (Statistical similarity, Utility, Rarity/Novelty, Expert validation)
-  Ipsos' commitment to leading with science – ensuring AI solutions are grounded in robust experimentation and validated against real-world data

Together, these papers reflect Healthcare at Ipsos' approach: to embrace AI's transformative potential while maintaining the rigour that healthcare decisions demand.



Ipsos' approach embraces AI's transformative potential while maintaining the rigour that healthcare decisions demand.

Further Reading



MAY 2026

NAVIGATING THE AI CONFIDENCE TRAP

Ensuring strategic success through
validated insights and expertise



JP Park

Consultant, Healthcare (US)

JP.Park@Ipsos.com

ABOUT HEALTHCARE AT IPSOS

A globally connected team in over 30 countries worldwide, we provide insights, evidence and guidance across the healthcare product lifecycle. We empower our clients to align commercial success with what matters most: **Improving patients' lives.**

Learn more at: www.Ipsos.com

