The last operational mile in life sciences commercialization

Five questions you need to ask of your commercial decision system

beghou

A *Commercialization That Works* Perspective

Executive summary

Commercialization teams are doing the right activities, but they are not always doing them together, according to our *Commercialization That Works* research findings reported in the **2026 Biopharma Commercialization Research Report**.

The cost of that misalignment across commercialization functions shows up in responsiveness. Operating separately and at different speeds hampers teams' ability to act on insights fast enough to keep up with the changing market.

Upon a closer look at how that gap actually plays out, the steps themselves aren't the problem: the data, the insights, and the AI recommendations are all there.

What is missing is the wiring between the recommendation and the action.

This is not a data problem in the traditional sense but a decisioning problem. It is the breakdown of cohesion in what our *Commercialization That Works* research identified as the “last operational mile.”

When operational gaps exist in the insight-to-action loop

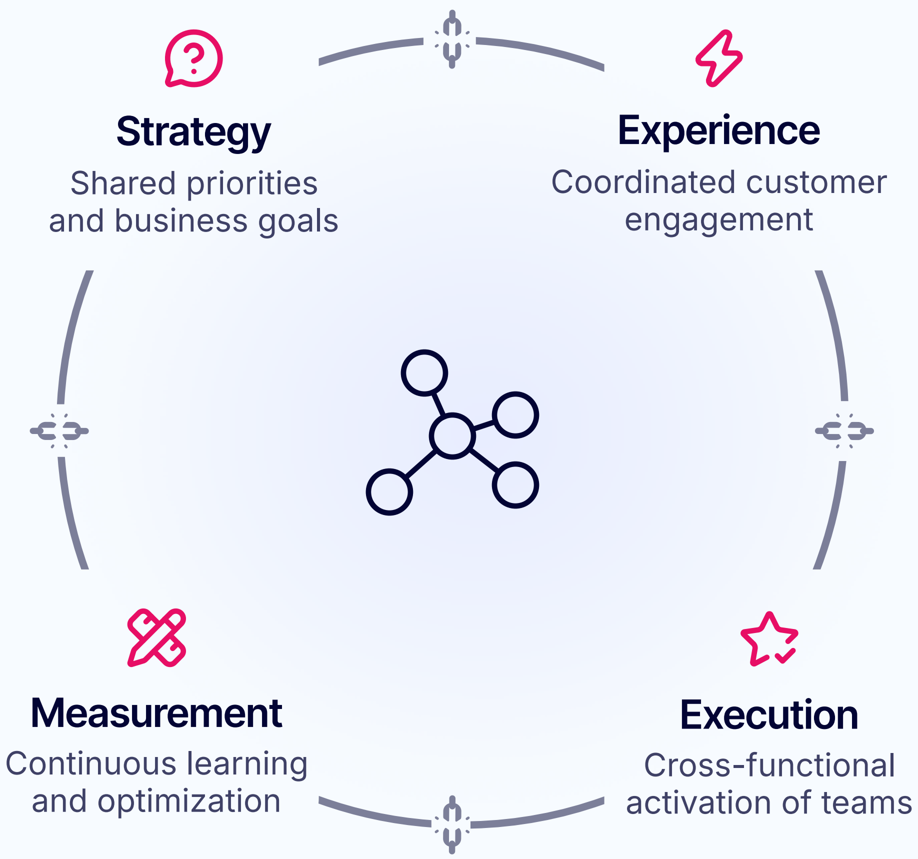

WHAT BREAKS TODAY

- Fragmented systems impede the one source of truth.
- Vital context for real-time decisions is missing.
- The moment for action passes due to disconnects between the analytics and execution layers.
- Every event triggers a sequence of meetings instead of action.
- Actions require a hand-off that adds a week to the timeline.

Turning commercial insight and execution into one fluid motion is the capability question that will separate the organizations that outperform over the next 3 to 5 years from those who do not.

To help close the loop from insight to action in commercialization, this paper outlines a structured method for diagnosing your operational gaps, recognizing them, and taking action to address them.

The problem: insight-rich, decision-poor

Early commercial performance matters more than ever, with poor performance often compounding over time. Industry research shows that more than one-third of new launches in the U.S. miss their first-year expectations, and roughly **70% of products that miss expectations at launch keep missing them in subsequent years**.

Modern launches move at sprint speed, yet the organizational systems supporting them often remain tuned for marathon pacing. Several forces are increasing the pressure on commercial teams to act faster, such as:

Specialty therapies increasingly dominate pipelines.

Field forces are shrinking as digital channels gain ground.

Pricing policies like the *Inflation Reduction Act* have compressed the window between launch and price negotiation.

Organizations have invested in data, tools, and insight, but the problem is often that their operating model cannot consistently deliver that insight to the right teams or turn it into timely action.

CRM, marketing automation, customer data platforms, omnichannel orchestration, and AI-driven next-best-action models have expanded visibility, but many teams still struggle with data trust, decision rights, governance, workflow handoffs, and execution ownership.

As a result, organizations may know where action is needed but still cannot act in time.

Diagnosing where it breaks down

THE 5 CHECKPOINTS

Cohesion is abstract until you watch a single decision try to move through an organization but ultimately stall or drop off completely. AI surfaces a recommendation. A dashboard flags a launch territory falling behind. A competitor moves on day one of a label update. The insight arrives, the meeting gets scheduled, the workflow stalls in approvals, and the market moves on.

Although these scenarios often surface first as commercial challenges, their downstream consequences ultimately affect patient access, engagement, and outcomes. The impacts might be noticeably visible along this chain, but determining the underlying cause of the breakdown can be difficult without a structured evaluation.

This is where a proven diagnostic method helps. Based on our experience, every real-time commercialization decision across launch, established brand, and access scenarios needs to pass through 5 consequential checkpoints:

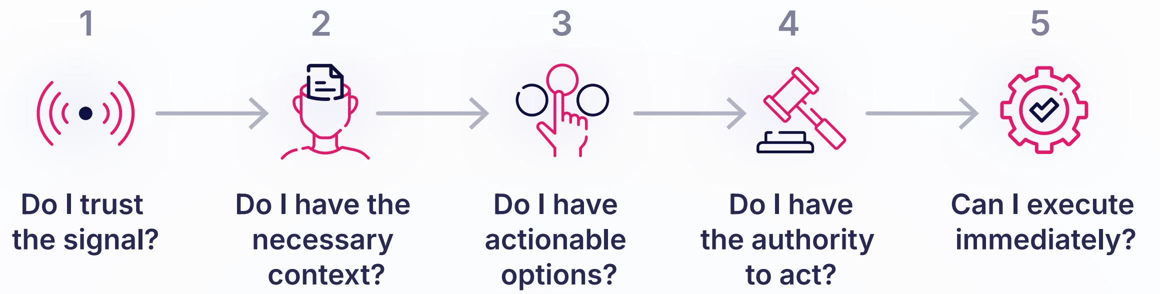

A failure at any one of these inhibits timely execution, and most leaders can name the checkpoint where their organization breaks down the moment they see the list. The next step is asking why it breaks there and what sequence of fixes addresses it.

The structural obstacles often emerge where strategy, teams, tech, and data intersect, and the pattern is easiest to see in the scenarios that commercialization teams live every week. Each of the following sections steps through the 5 checkpoints, including why they break, how that shows up, and what can be done about it.

1

Do I trust the signal?

Is this number real or one of several conflicting answers from systems that create multiple versions of the truth?

THE ROOT CAUSE:

→ There is no single source of truth

Across the industry, teams are working with increasingly complex data environments, where the volume of information has grown faster than the ability to align, connect, govern, and trust it. In our *Commercialization That Works* research, 53% of data and analytics respondents cited data completeness as a hindrance to good launch decisions, and 37% named internal data governance as a barrier.

WHAT BREAKS

Fragmented systems impede the one version of the truth that speed requires. When master data has gaps, metric definitions conflict, and channel data are siloed, every decision starts with a debate about whose number is right.

WHY IT MATTERS

The debate is the most common and visible failure and can result in days lost reconciling reports, missed windows to adjust spend or messaging, and teams defaulting to inaction because they cannot defend a number with confidence.

HOW THIS MIGHT SHOW UP:

3 teams, 3 data sources, and 3 different stories

In a monthly leadership meeting, the sales team reports performance is on track, the marketing team reports share of voice is declining, and the analytics team reports new patient starts are flat. The decision on promotional spend can’t move, so it is tabled.

This is the most expensive failure of all because it looks like a normal meeting. Less visible but more corrosive is the trust cost. When leadership cannot agree on a number, they stop trusting each other’s data, and every subsequent decision gets harder.

WHAT TO DO ABOUT IT

Not every data point needs the same level of rigor, but the signals that drive key commercial decisions need consistent definitions and reliable sources. This requires a clear data foundation that defines which sources are trusted for which decisions and how metrics are calculated.

Data also need to be complete enough, connected across functions, and accessible when decisions are made. This becomes more important as AI tools are used to summarize trends, detect patterns, and recommend action. AI can only work from the signals it can access and interpret, so unclear definitions or incomplete data can make weak inputs appear more definitive than they are. The practical first step is to identify the few metrics that most often trigger debate and make them decision-ready.

2

Do I have the necessary context?

A signal without context is just a number.

For example, knowing engagement is lagging is not the same as knowing why.

THE ROOT CAUSE:

→ Context lives somewhere else

Even in organizations where the foundational data is trusted, the context that explains it rarely sits anywhere a decision can reach. It lives in people's heads, in a slide buried in last quarter's review, or in the institutional memory of whoever ran the last launch. Historically, context like this was unusable until someone normalized it into a table that a dashboard or data model could read, which meant most of it was never captured at all.

WHAT BREAKS

The person making a real-time call is missing vital context until it lives alongside the decision.

WHY IT MATTERS

Context provides information about what kind of problem it is. Without it, teams can overreact to normal variation, underreact to meaningful change, or misread what a signal is actually saying. For example, a drop in engagement may reflect content fatigue, access friction, competitive activity, seasonality, formulary disruption, a field execution issue, or a change in the underlying patient population. Context enables the team to determine which of these is the real cause.

HOW THIS MIGHT SHOW UP:

A signal means different things in different markets

Two territories show the same decline in new patient starts. In one, the decline reflects a temporary reimbursement delay tied to a payer update. In the other, it reflects a competitor gaining traction among a specific healthcare professional (HCP) segment. The dashboard flags both territories the same way, so the response is also the same: more field attention and more promotional pressure.

The team moves quickly but not intelligently. One territory needed access support and patient services coordination. The other needed a targeted competitive response.

Because the context did not travel with the signal, the organization acted on the symptom instead of the cause.

WHAT TO DO ABOUT IT

It is important to treat context as part of the operating foundation. Commercial decisions need access to therapeutic context, brand strategy, launch objectives, customer archetypes, payer dynamics, patient journey realities, competitive activity, field feedback, channel history, prior decisions, and business rules.

Fortunately, with modern technology such as AI, the constraints around context are lifting. Unstructured sources like documents, market research, call notes, meeting summaries, and transcripts where this knowledge lives can now be put to work directly. However, most organizations have not rewired for it.

The key shift is moving context from informal memory into reusable knowledge assets, making sure to curate, govern, and connect both structured and unstructured data sources to decision workflows so AI-enabled tools and human teams are working from the same commercial reality.

3

Do I have actionable options?

Does the system surface real alternatives with projected impact or only flag the problem?

THE ROOT CAUSE:

→ **Recommendations exist apart from execution**

The recommendation lives in the analytics layer, and execution lives in the CRM or the marketing platform.

WHAT BREAKS

Because the loop between the analytics layer and execution layer is not connected, the moment for action passes even though the problem is visible.

This is the gap that defines the moment. AI has earned its place in generating insights, with 80% of data and analytics leaders citing its impact on clarifying customer focus, but only 30% finding it impactful for actually engaging customers, as reported in the 2026 Biopharma Commercialization Research Report.

WHY IT MATTERS

Systems that only flag problems still leave the hardest work to humans: interpreting the signal, finding context, and deciding what to do. As AI increases the volume and speed of recommendations, this gap can become more pronounced: teams may have more things to evaluate without a clearer path to action. This delays the ability to act and consumes resources that could be dedicated to higher-level tasks.

HOW THIS MIGHT SHOW UP:

The territory that is not moving

It is week 6 of a launch. National uptake is on plan, but one region is materially behind. AI flags the gap and ties it to low awareness in a specialty segment reps have visited but shows no digital engagement. The signal exists, but it lives in a regional dashboard reviewed weekly, not in real time.

By the time it is escalated, the launch window has narrowed. Early launch periods are exactly when teams can least afford to be blind for weeks, yet that is precisely when foundational readiness is still forming.

WHAT TO DO ABOUT IT

A stronger model translates signal and context into a small set of evidence-backed options tailored to the objective, segment, and constraints. Instead of saying “engagement is declining,” the system should point to likely responses, such as refreshing content, adjusting sequencing, increasing field follow-up, or investigating access barriers.

To achieve this, organizations need to define decision pathways in advance by mapping common signals to a curated set of approved actions, linking those actions to available channels and content, and embedding business rules and constraints directly into the system.

4

Do I have the authority to act?

Can the person looking at the screen act, or does every signal trigger a meeting and an escalation?

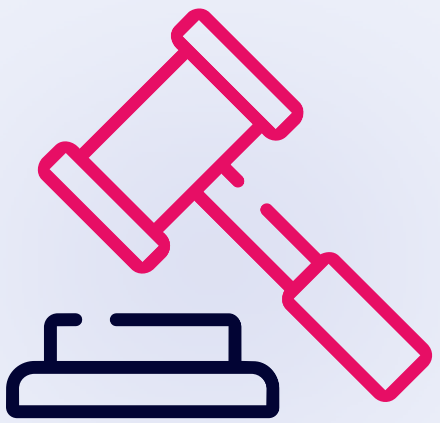

THE ROOT CAUSE:

→ **Lack of governance blocks speed**

Even with the right data, context, and execution infrastructure, a decision still needs someone empowered to make it. Governance has a reputation problem, with most organizations associating it with bureaucracy. This leaves teams without pre-approved action frameworks, clear decision rights, and defined escalation paths.

WHAT BREAKS

Every event triggers a sequence of meetings that compresses days of opportunity into weeks of process.

WHY IT MATTERS

The organization consistently reacts too late to influence outcomes when timing matters most. By the time a decision is approved, the underlying conditions may have already shifted, making the response less effective or even irrelevant.

Over time, this lag erodes competitive advantage and trains teams to expect delay rather than act with urgency.

HOW THIS MIGHT SHOW UP:

The asset that stopped working

Email open rates for a flagship HCP campaign have dropped 31% over 6 weeks. The same content has been in rotation for 4 months. AI surfaces three alternatives with projected lift.

The data are clear, and the recommendation is concrete. However, execution requires regulatory review, MLR approval, and a CRM update. That chain of approval takes weeks in most organizations. Meanwhile, every rep call backed by the stale asset is a spent opportunity. At scale, a 31% decline across the top segment is thousands of touchpoints actively training HCPs to ignore the brand.

WHAT TO DO ABOUT IT

Real-time execution requires someone to own every metric, definition, and decision: Who owns this number? Who can change its definition? Who resolves the conflict when two teams see different answers?

Governance earns a poor reputation only when it is bolted on after the fact. Designed upfront, it is what makes speed possible. It gives teams the confidence to move because the boundaries are already clear.

Rather than a large elaborate apparatus, good governance relies on deliberate design that is established before the dashboards are built, vendors are integrated, and the first real-time decision has to be made. The organizations that get this right know the answers to the critical questions in advance.

5 Can I execute immediately?

Is the recommendation wired to the systems that act on it, or does the change require a hand-off that takes a week?

THE ROOT CAUSE:

→ **The external ecosystem is disconnected**

Many commercialization organizations run multiple external vendors at once: CRM, MDM, NPP platforms, hub providers, data aggregators, and channel execution tools. Often, each was selected to solve a specific problem, and many were not designed to work together.

WHAT BREAKS

The result is an ecosystem that generates enormous data but cannot move it where it needs to go when it needs to get there.

WHY IT MATTERS

The disconnect means insights arrive too late to influence outcomes. AI can intensify this gap by increasing the speed at which recommendations are produced while the execution environment remains slow, fragmented, or manual.

In fast-moving commercial environments, that delay directly translates into missed opportunities, slower response to market shifts, and diminished impact of otherwise sound strategies.

HOW THIS MIGHT SHOW UP:

The patient who doesn’t start

Prescription data show HCPs are writing scripts, but a high percentage of patients never start therapy. Hub data show prior authorization denials spiking in two payer segments. Patient Services has the capacity to intervene but cannot see which prescriptions are at risk in real time.

A script written is a patient not treated. Every denied prior authorization that goes uncontested is a patient who does not start the therapy their physician believed would benefit them. For some disease areas, a treatment can add the gift of time, and the delay carries consequences that no commercial recovery plan can undo.

WHAT TO DO ABOUT IT

Build the operating connections before they are needed, and pressure-test execution pathways the same way you would pressure-test launch scenarios. This includes testing whether AI-enabled recommendations can move into the systems, teams, and workflows required to act on them.

Closing the insight-to-action loop requires someone at the center of the vendor ecosystem, coordinating data flows, resolving definitional conflicts, and ensuring the right information reaches the right system at the right moment.

The human layer

The human layer underpins the 5 checkpoints and associated organizational and technical structural obstacles. Even with a technical closure to the insight-to-action loop, people still have to trust it, act on it, and change how they work because of it. AI raises the stakes for this human layer because recommendations can arrive before trust has been established.

Field teams that have run on instinct do not adopt an AI recommendation because the platform went live: 79% responded they would begin using a new field-facing technology affecting their work only when they could determine the value themselves in our 2026 ***Commercialization That Works: Field Force Pulse Survey***.

Similarly, brand leaders who have always owned their metrics may not defer to a shared definition just because the governance model said so.

Instead, driving change in a commercial organization is the deliberate work of designing around the human layer: the incentives, the feedback loops, the leadership behaviors, and the early wins that build trust.

Organizations spend months on data, technology, and analytics. They spend only weeks, if that, on the people layer that decides whether any of it gets used.

Closing the loop is ultimately a human achievement that views adoption as the first priority of design rather than the last mile of implementation.

Closing the loop

The organizations that pull ahead will not be the ones with the most data, the most sophisticated AI, or the largest field forces. They will be the ones that have closed the loop between insight and action to be able to observe a real-world signal, design a response, and activate it before the moment passes. That work follows a specific sequence, and in our experience, many organizations tend to invest in the recommendation layer before building the foundation that lets their teams trust what it produces, which is required for confident action and decision-making.

The most useful first step is also the simplest.

Pick a real scenario your team faces today and walk it through the 5 checkpoints until you find the first place execution would break. That breakdown point is different for every organization, but the pattern is consistent. That point is also where the work begins to transform your commercialization model into one that can adapt to the market as it changes.

At Beghou, we enable commercial organizations to move at the pace of modern launches, so insight, decision, and execution are one motion. We welcome a conversation about where your decision loop breaks and what it would take to close it.

beghou

Contact us at info@beghou.com