

Responsible AI in Commercial Strategy: Protecting IP While Moving Faster

A practical point of view on using generative AI in biotech commercial work while protecting confidential information, preserving human judgment, and reducing IP risk.

Generative AI is already part of how many teams research, draft, summarize, and accelerate work. But adoption is outpacing trust. For biotech commercial leaders, that gap matters. AI can accelerate synthesis, pressure-testing, drafting, and content development. But it can also create avoidable risk if teams paste confidential strategy into public tools, over-trust unverified output, or lose track of what content is original, adapted, or machine-generated.

The right goal is not to use AI everywhere or avoid it entirely. It is to use AI where it adds value, while keeping human accountability, IP protection, and decision quality intact.

A practical position for commercial teams

At Route Consulting, generative AI should be treated as a tool for acceleration, not a substitute for judgment. It can support research, synthesis, outlining, drafting, and structured ideation. It should not replace strategic thinking, expert review, or final decision-making.

A useful operating principle is simple: **AI may assist the work, but people remain accountable for the work.**

Where AI can add value safely

Used responsibly, AI can help commercial teams move faster in a few specific ways:

- summarizing large sets of public information
- drafting first-pass outlines, briefing notes, and discussion prompts
- identifying decision questions, assumptions, and scenario branches

- accelerating synthesis of meeting notes or research inputs
- improving content workflow efficiency before human refinement

These are high-value uses because they reduce manual effort without handing over final commercial judgment.

What should never go into public AI tools

The most important policy decision is usually not what AI can generate. It is what your team is allowed to input. Do not enter the following into public or non-approved AI systems:

- confidential commercial strategy
- unpublished research or internal analyses
- client-sensitive materials
- proprietary market models
- pricing assumptions
- draft board or leadership materials
- nonpublic competitive assessments
- regulated, personal, or otherwise sensitive data

Careless use can expose confidential information, compromise decision quality, or create downstream legal and security issues.

A simple responsible-use policy

- 1 **Approved use cases.** Generative AI may be used for public-information synthesis, brainstorming and structured ideation, first-draft outlines, writing support for nonfinal internal drafts, summarization of nonconfidential materials, and workflow acceleration where human review is built in.
- 2 **Restricted or prohibited use cases.** Generative AI should not be used to make final strategic recommendations without expert review, draft regulated claims without verification, process confidential information in non-approved tools, produce final external-facing materials without human editing and approval, or make legal, regulatory, scientific, or access decisions autonomously.

- 3 **Human review is mandatory.** Every AI-assisted output should be reviewed for factual accuracy, relevance to the commercial context, bias or omissions, unsupported claims, language that overstates certainty, and alignment with brand, access, and stakeholder realities.
- 4 **Protect IP and ownership.** Keep a clear record of what was created by a person, what was AI-assisted, what source materials were used, and what edits and approvals were made before publication.
- 5 **Use approved platforms only.** If your organization uses AI, it should do so through tools with known policies on data retention, training use, user access controls, enterprise security, and auditability.

What this means for biotech commercial work

For biotech commercial teams, responsible AI use is not just a technology issue. It is a credibility issue. Commercial strategy often touches sensitive territory: market assumptions, launch timing, competitive vulnerabilities, access logic, stakeholder narratives, and internal decision-making. These are exactly the kinds of materials that should be handled carefully.

The most effective approach is selective use with firm guardrails: use AI to accelerate thinking, not replace it; keep confidential inputs out of non-approved systems; require human review on all substantive outputs; document provenance and ownership; and apply stricter review to external-facing or IP-relevant work.

Final point of view

The question is no longer whether generative AI can help commercial teams work faster. It can. The more important question is whether teams are using it in a way that protects confidential information, preserves human authorship, and maintains decision quality. Responsible AI is not about slowing work down. It is about creating a system where speed does not come at the cost of trust, IP, or judgment.