connectHEOR

# AI IN HEOR

## Who Is The Human in Human-AI?

*Why the most important governance question in HEOR AI is not whether a human is in the loop.*

*The question is whether the right human is, doing the right thing, and whether they know they might already be a ghost.*

**CONNECT AI: HEOR THINKING FOR AN AI-ENABLED ERA**

www.connectheor.com

# Publisher's Note
# Connect AI | Edition 02

Most HEOR teams are not yet using AI in any meaningful sense. There are pilots. There are demonstrations. There are working groups and vendor conversations and internal mandates to explore AI that have been on agendas for eighteen months without producing anything with a governance structure around it.

This is not a criticism. Caution in HEOR is professional judgment. The consequences of getting AI wrong are measured in failed submissions and eroded credibility, not in a product recall or a service outage. The bar is higher here than almost anywhere else AI is being deployed.

But something is already happening, quietly, at the team level. Your analysts are using ChatGPT to draft report sections. Someone used Claude to synthesise a clinical study report last week. Copilot is helping write model technical documents. These are not formal deployments. There is no governance framework around them. Nobody has decided which person is responsible for the judgment calls embedded in those outputs. And because the tools feel familiar, like a very good search engine or a very fast junior analyst, nobody has asked whether the human reviewing the output is genuinely capable of challenging what the AI decided.

There is a more uncomfortable truth underneath this. Most directors and VPs leading HEOR functions were already drifting away from the operational detail of the work before AI arrived. That drift is not new. Teams grew. Deliverable volumes increased. The natural consequence was that senior leaders moved progressively further from the substance of individual outputs, staying close to strategy and programmes while the technical and scientific detail was handled by the team. That is a normal and arguably appropriate evolution of a senior role.

AI does not change that. But it removes the last structural reason for a senior leader to stay close to individual outputs, which was the effort and uncertainty involved in producing them. When an analyst spent three days drafting an evidence synthesis, the director was naturally drawn into the process. When Claude produces the same synthesis in an afternoon, the director receives a polished document and the natural pull towards engagement is gone. The distance that took years to develop in the manual world is now the default from the first use of the tool.

# Publisher's Note
# Connect AI | Edition 02

This is where the governance question becomes a leadership question. AI will not reduce the amount of work in HEOR. It will change the distribution of it. The execution layer, the drafting, the synthesis, the structuring, will be handled increasingly by AI. The judgment layer, the framing choices, the interpretive decisions, the strategic positioning of evidence, will not. Those will concentrate in the hands of the people who are senior enough to make them. Directors and VPs who have drifted away from the detail of HEOR work do not need to return to execution. They need to return to judgment, and to understand that AI is about to make that return both more important and more demanding than it has ever been.

This edition of Connect AI is about one question inside all of that: who is the human? Specifically, which person, with which capability, is genuinely in the loop. Because the most dangerous governance failure we are already seeing is not a system failure. It is a human one. The human who was supposed to own the reasoning is still named in the sign-off chain. They stopped genuinely engaging weeks ago. Nobody noticed.

We published a hybrid intelligence framework in 2025 in peer-reviewed jounral that introduced two modes of human-AI collaboration and the governance conditions each requires. This paper builds on that work to name a third mode and to describe the specific failure pattern that connects all four governance problems we are already seeing in early HEOR AI adoption. We call it the Governance Ghost.

—

Tushar Srivastava
Chief AI Strategist, ConnectHEOR

SECTION 1

# You Already Have a Governance Ghost.

Think about the last time someone on your team used ChatGPT, Copilot, or Claude for a piece of HEOR work. A literature summary. A model technical report section. A synthesis of clinical evidence for a dossier. A gap analysis for an evidence planning meeting. Now think about the person whose name is associated with that output: the reviewer, the director who approved it, whoever signed off. Ask yourself honestly: did that person genuinely interrogate the specific judgment calls embedded in the AI output? Not whether it read well. Not whether the references checked out. But: does this framing of the clinical evidence reflect the right interpretation? Does this summary of the comparator landscape match what we would actually defend to a payer? Did the AI omit something that changes the conclusion?

In most cases, the honest answer is: probably not in any depth. The output looked plausible. The tone was right. The reviewer made a few edits, approved it, and moved on. That is not unreasonable. These tools are designed to produce outputs that feel complete and authoritative. **But felt complete is not the same as was interrogated.** And the difference matters the moment an HTA body or a senior leader asks a question that requires a human to stand behind the specific reasoning in the document.

This is the Governance Ghost forming. Not in a formal AI deployment. In the everyday use of tools your team already has access to. The ghost is the reviewer who is named as responsible for an AI-assisted output but who stopped genuinely interrogating what the AI contributed. They are still approving. They have not genuinely engaged for weeks. Nobody has noticed.

> ***The Governance Ghost is not a bad actor. It is a good person placed in a governance role that the workflow no longer supports them exercising. The audit trail looks fine. The human oversight is nominal.***

This matters because the outputs produced with AI assistance today, the SLR summaries drafted in ChatGPT, the evidence syntheses shaped by Claude, the model rationale sections polished by Copilot, are the inputs to submissions that HTA bodies will scrutinise. When a NICE EAG question arrives asking why a specific clinical endpoint was framed the way it was, or why a particular study was given the weight it received in the synthesis, the Governance Ghost cannot answer it. Not because they lack the capability. Because they read the output rather than interrogating it.

The question worth sitting with is not whether a human reviews every AI-assisted output. Every team will say yes. **The question is whether the person reviewing each piece of AI-assisted work genuinely know what is expected of them.** If the answer is they check it before it goes out, you may already have a ghost.

There is a wider shift underneath this worth naming directly. AI will not reduce the amount of work in HEOR. It will change the distribution of it. The execution layer, the drafting, the structuring, the synthesis, will be handled increasingly by AI tools. The judgment layer will not. The choices about how to frame evidence, which studies to weight and why, how to position a value story for a specific HTA context: these will concentrate. They will require more from the people responsible for them, not less, because AI produces outputs that look complete and the temptation to ratify rather than interrogate is structurally built into every AI-assisted workflow.

This means the reset required is not primarily about AI governance processes. It is about leadership mindset. Directors and VPs who have naturally moved away from the operational detail of HEOR work over the past years need to understand that AI does not continue that trajectory. It reverses it, selectively. Not a return to execution. A return to judgment. The question is whether the people who are supposed to own the judgment layer are genuinely present in it, or whether they have drifted into governance roles that the workflow no longer supports them exercising.

? *The ghost test: think of the last three pieces of work your team produced with AI assistance. An SLR section, an evidence summary, a model rationale document. For each one, identify the two or three most important judgment calls embedded in the output. Who owns each one? Could that person explain right now why those choices were made, what the output would have looked like if the AI had framed it differently, and defend those choices to a payer or HTA body?*

SECTION 2

# How Ghosts Are Made: Four Patterns.

The Governance Ghost does not emerge from negligence. It emerges from a predictable set of organisational patterns that are each individually reasonable and collectively produce a governance structure that looks right and functions poorly. Four patterns create ghosts. Most organisations in early AI adoption are running at least two of them simultaneously.

### The Pilot Lead

The senior manager or associate director who championed AI adoption on the team and has been running the experiments, testing ChatGPT on SLR abstracts, using Claude to draft model rationale sections, exploring Copilot for clinical evidence summaries. They know what the tools can do. They have become the informal AI authority on the team, and when AI-assisted outputs need reviewing, they are the natural choice. The problem is that their relationship with these tools was built on capability demonstration, watching them produce impressive outputs under favourable conditions. When Claude drafts an evidence summary that frames the comparator landscape a particular way, the Pilot Lead checks whether it is factually accurate. They are unlikely to ask whether that framing would hold up under payer scrutiny, or whether a different characterisation of the evidence would have been more strategically defensible. That question lives outside the territory where their AI enthusiasm operates. Over time, their reviews become accuracy checks rather than strategic interrogations. The ghost forms gradually.

***Knows what the AI produced. Has stopped asking whether it should have framed things differently.***

### The Signing Director

The VP or director who appears as the final approver on AI-assisted outputs. They have strategic context. They understand what is at stake for the programme. What the current workflow does not give them is the proximity to engage seriously with what the AI actually contributed. An SLR synthesis section that would previously have taken a researcher three days to draft has arrived in their inbox in an afternoon. It reads well. The citations look right. They approve it with a few comments on the framing. The specific judgment calls embedded in that synthesis, which studies the AI chose to foreground, how it characterised the heterogeneity of evidence, what it implied about the comparator, are not visible at the level they are engaging. The volume of AI-assisted outputs arriving for approval means there is no structural space to go deeper. This is not a failure of diligence. It is a failure of workflow design. In the language of the three modes defined in Section 3, the Signing Director is nominally in the loop but functionally on the loop: providing supervisory oversight rather than genuine interrogation of individual decisions.

***Has the authority to own the output. Has not been given the conditions to own what the AI put into it.***

### The Inherited Committee

The review structure that pre-dated AI has been extended, without redesign, to cover AI-assisted outputs. The medical reviewer checks clinical accuracy. The market access lead checks payer framing. The publications lead checks scientific rigour. The global evidence lead approves the narrative. Each person reviews their slice of the AI-assisted document. Nobody reviews the whole. Nobody asks: did the AI synthesis of this evidence reflect a coherent evidence strategy, or did it pattern-match from its training in a way that happens to read well but

does not reflect the actual evidentiary position we want to take? Those questions sit between the slices. The committee covers them collectively. No individual owns them. When a medical affairs leader asks why the evidence summary positioned the data the way it did, the committee cannot answer as a unit. The person who speaks will be the most confident, not the most informed.

***Distributed review creates the appearance of accountability. The gap is between the slices, where no one is looking.***

**The Governance Ghost**

The most dangerous pattern and the hardest to name. This person was genuinely engaged at the start, they helped set the guidelines for AI use, they reviewed the first few AI-assisted outputs carefully, they gave detailed feedback on what was acceptable and what needed more human input. As AI-assisted work became routine on the team, their engagement became progressively more nominal. They are still listed as the approver. They still sign off. But their review has become a read-through rather than an interrogation. The workflow no longer creates the expectation of genuine challenge, and the volume of AI-assisted output means a deep review feels disproportionate. They have not been removed from the loop. The loop has been redesigned around their absence without anyone deciding that is what happened. The audit trail shows consistent human oversight. What it does not show is that oversight long ago became sign-off.

***Still in the loop on paper. No longer in the loop in practice. The document history conceals both.***

Each of these patterns can produce a ghost. The Pilot Lead becomes one when their technical reviews substitute for strategic interrogation. The Signing Director becomes one when timeline pressure makes genuine engagement structurally impossible. The Inherited Committee produces a collective ghost: everyone reviewed it, nobody owned it. And the Governance Ghost is the endpoint all three patterns share if the organisation does not design against them deliberately.

! LEADERSHIP SIGNAL

*Before any AI-assisted HEOR output goes to an HTA body: for each major framing or methodological choice, can you name the specific person who genuinely interrogated it, not just approved it, but interrogated it, and confirm they are still in the governance structure in a way that would allow them to defend it under clarification?*

SECTION 3

# Three Modes. One Governance Mistake.

The hybrid intelligence framework we published in Advances in Therapy introduced two distinct modes of human-AI collaboration for HEOR workflows. The governance implications of that framework point toward a third. All three are legitimate. All three are being misapplied. And the misapplication in every case produces the same outcome: a ghost where a genuine human owner should be.

Understanding the three modes is also the foundation for understanding the four profiles in Section 2. The Signing Director is a ghost because they have drifted into HOTL for decisions that require HITL. The Governance Ghost is a ghost because their nominal AITL role no longer involves genuine direction. The vocabulary of modes makes the governance failure specific enough to address.

**Srivastava T, Irfan H, Babiy V, Swami S. (2025). Integration of Generative AI with Human Expertise in HEOR: A Hybrid Intelligence Framework. Advances in Therapy, 42, 4103–4130.**

*The framework introduces HITL (AI leads; human validates) and AITL (human leads; AI augments) as two implementable approaches with distinct governance requirements. The choice between modes must be deliberate and matched to the nature of the decision, not inferred from what the tool can produce fastest.*

| | Human-in-the-Loop (HITL) | AI-in-the-Loop (AITL) | Human-on-the-Loop (HOTL) |
|---|---|---|---|
| **Who leads** | AI leads; human validates | Human leads; AI augments | AI operates; human supervises |
| **Best for in HEOR** | AI drafts SLR sections, abstracts, evidence summaries, report write-ups | Human defines evidence strategy, key messages, value narrative, endpoint rationale; AI assists execution | Reformatting, translation, or quality-checking of settled, approved content |
| **Human must be able to** | Identify when the AI framing is strategically wrong, not just factually inaccurate | Set direction clearly enough that AI cannot misinterpret the evidence position | Define exception triggers and intervene when outputs fall outside defined bounds |
| **Primary risk** | Automation bias: ratifying a well-written output rather than interrogating its choices | Vague human direction producing an AI output that reads well but reflects nobody's actual strategic intent | Governance Ghost: the loop exists on paper but oversight has become purely nominal |
| **Failure signal** | Nobody can explain why the AI framed the evidence a particular way | AI analysis was completed before the human evidence strategy was settled | Named reviewer has not substantively engaged since the pilot |

The governance failure is not choosing the wrong mode in principle. It is failing to assign a mode at all, and watching governance drift toward HOTL for decisions that require HITL or AITL. In practice, this drift is happening right now in organisations that are using ChatGPT or Copilot for HEOR tasks without anyone asking: what mode is this? Who is responsible for the judgment calls in this output? When the analyst uses Claude to draft a clinical evidence synthesis for a global value dossier, is the reviewer in HITL mode, with the expectation and capability to genuinely challenge what the AI concluded, or are they in HOTL mode, reading the output and approving it because it sounds right?

Healthcare has already seen where HOTL drift leads in higher-stakes AI deployments. In clinical decision support systems, clinicians who were initially close to the AI implementation became progressively more distant as the system moved into routine use. Genuine oversight degraded, through entirely normal workflow pressure, into periodic ratification. The governance record showed continuous human involvement. The quality of that involvement had fundamentally changed. HEOR does not have AI systems as sophisticated as clinical decision support yet. But the teams that are building habits around ChatGPT and Copilot today are building the governance culture that will govern much more consequential AI tools in two years.

That is why the mode question matters now, when the stakes of getting it wrong are still relatively contained. The patterns being set in pilots and early adoption will be the patterns applied when AI starts touching the work that goes to HTA bodies.

> ***The mode you set in the pilot is the mode you will defend in the submission. Governance drift toward HOTL for decisions that require HITL is not a future risk. It is already happening in organisations that describe themselves as responsible adopters.***

! **LEADERSHIP SIGNAL**

*For each way your team currently uses AI tools, drafting SLR sections in ChatGPT, synthesising evidence in Claude, writing model documentation in Copilot, assign a mode explicitly. For each: which mode is it? Which mode does it require? And is the human currently reviewing those outputs genuinely doing what that mode demands of them?*

SECTION 4

# What the Right Human Must Actually Be Able to Do.

Once the ghost problem is named, the question becomes: what does the right human, the non-ghost, actually look like? The hybrid intelligence framework identifies human expertise as foundational to any effective HEOR AI system. But expertise is doing too much work in that sentence unless it is made specific. Three capabilities need to be present at the moments that matter. They are not personality traits or seniority levels. They are professional practices that can be assessed, assigned, and required.

### Submission-Ready Authority

The ability to walk into a NICE EAG clarification, a G-BA scientific advice meeting, or a payer discussion and defend a specific piece of evidence framing as a deliberate expert choice, because the person speaking genuinely believes it was the right choice and can reconstruct the reasoning behind it. Right now, in the context of where HEOR AI actually is, this capability is needed for outputs like: why the clinical evidence synthesis framed the comparative efficacy data the way it did, why a particular study was weighted as it was in the evidence summary, why the value narrative led with this endpoint rather than that one. These are judgment calls. They may have been shaped or drafted by ChatGPT or Claude. The human reviewer must be able to own them as if they made them deliberately, because to anyone outside the team, they did. This capability develops through direct exposure to HTA scrutiny and payer engagement. It is the capability most at risk of being absent from AI review structures because it does not look technical.

***If no one reviewing an AI-assisted piece of work could defend every major framing choice in it to a payer today, the work is not yet defensible.***

### Named Assumption Ownership

A specific name against a specific judgment call in an AI-assisted output, not the document as a whole, but the individual choices within it, with a clear understanding of why that choice was made and what the output would have looked like if the AI had framed it differently. In AITL mode, this is the human's primary contribution: they directed the framing, they set the strategic intent, the AI executed it. In HITL mode, it is the critical safeguard: the reviewer must actively adopt ownership of the AI's choices as if they made them deliberately, rather than treating them as the tool's responsibility. When a manager asks why did we characterise the evidence this way, the answer cannot be that is what Claude produced. Someone on the team needs to own it and be able to explain the reasoning.

***Name the person responsible for every major framing choice before the output leaves the team. If you cannot name them, the governance is missing its centre.***

### The Adversarial Practice

A structured habit assigned to a named person whose role is to ask three questions at every AI-assisted review before sign-off. For an evidence synthesis or SLR section: what did the AI not include that a sceptical payer would notice? What framing choice, if challenged, would expose a weakness in the evidence narrative? And what would this document look like if the AI had been prompted differently, and would we be comfortable with that version too? These questions take fifteen minutes. They almost never get asked, because nobody has been assigned to ask them and the organisational reward is for outputs that move quickly, not for outputs that

are interrogated thoroughly. Automation bias, the measurable tendency to accept plausible-looking AI outputs without adequate scrutiny, is mitigated by structured interrogation with assigned ownership, not by good intentions.

***Assign someone to ask the three questions at every AI-assisted review. If that role does not exist, automation bias is already operating on your team.***

These three capabilities are almost never concentrated in one person, and the people who hold them are rarely the same people currently occupying the governance positions. The person with submission-ready authority is often the most senior and the least operationally proximate to the work. The person most willing to take named assumption ownership is often too junior to carry it under external scrutiny. The person best placed to run the adversarial practice is often not in the governance structure at all.

It is also worth being clear that named assumption ownership and submission-ready authority, while related, are distinct roles. The person who owns a specific framing assumption does not need to be the same person who could defend the full submission in front of an HTA body. Named assumption ownership is about granular accountability within the team. Submission-ready authority is about the organisational capability to stand behind the whole. Both need to exist. Neither substitutes for the other.

The practical action is specific and immediate. Before your next AI-assisted workflow produces an output that goes beyond the team, answer this: who in your team could walk into a NICE EAG clarification today and own each major AI-assisted framing choice? If that person is not currently named in the review structure, that is the first thing to fix. Not the AI system. Not the process diagram. The human gap.

**Srivastava et al. (2025). Hybrid Intelligence Framework. Advances in Therapy.**
*The framework is explicit that human input must validate and refine AI outputs, not merely ratify them. Organisational readiness and the active cultivation of human expertise are identified as primary enablers of responsible AI integration, not secondary safeguards.*

SECTION 5

# Five Things to Do Before Monday.

The following are not a governance framework. They are five specific actions that address the ghost problem directly. None of them requires new technology. All of them require a leadership decision.

### 1. Run the ghost test on your current review structure.

For every person currently named as the reviewer or approver for AI-assisted HEOR outputs, SLR sections, evidence summaries, dossier content, model rationale documents, ask the ghost test: when did they last genuinely interrogate a specific AI framing choice rather than just approving the output? If the honest answer is at the start or I am not sure, that person may already be a ghost. This is a workflow design question, not a performance question. The ghost usually exists because nobody redesigned the review expectation as AI-assisted work became routine.

### 2. Assign a mode to every type of AI-assisted task.

For each way your team uses AI tools, drafting SLR sections, synthesising clinical evidence, drafting model documentation, preparing gap analyses, assign a mode explicitly. HITL is appropriate when the AI drafts and the human must genuinely interrogate the framing, not just the accuracy. AITL is required when the human must set the strategic intent and evidence position first, and AI executes under that direction. At this stage of HEOR AI evolution, HOTL is appropriate only for low-stakes, settled tasks. The error is not using HOTL for routine tasks. The error is allowing HITL and AITL tasks to drift into HOTL review because the output looks polished.

### 3. Name the owner of every major framing choice.

Before any AI-assisted output goes beyond the team, every significant framing or interpretive choice embedded in it should have a named human owner: a person who owns why that characterisation of the evidence was made, what the output would have looked like if framed differently, and who can explain the choice to a payer or medical affairs leader. Not the team reviewed it. A person. This act of naming is itself a governance act, it forces the question of whether the right person is actually responsible, and frequently reveals that the person best placed to own a specific judgment is not the person currently signing off.

### 4. Assign someone to ask the adversarial questions.

Create an explicit role in every AI-assisted review: the person whose job is to ask what the AI did not include, what framing choice a sceptical payer would challenge, and what the document would look like if the AI had been prompted differently. This role should be assigned by name before the review, not filled by whoever happens to raise a concern. It should be treated as the most valuable contribution in

the review process. The evidence on automation bias is clear: awareness alone does not prevent it. Structured interrogation with individual accountability does.

### 5. Name the person with submission-ready authority for each output category.

For every category of AI-assisted HEOR output, clinical evidence summaries, value narratives, SLR syntheses, model rationale documents, identify the person who has the submission-ready authority to defend every significant framing choice in front of a payer or HTA body. This is a different role from the named assumption owner in step 3. The assumption owner is accountable for a specific choice within the document. The submission-ready authority is the person who can stand behind the document as a whole under external questioning. Both roles need to be named. Both need to be in the review chain before the output goes anywhere near a submission.

> **Srivastava et al. (2025). Hybrid Intelligence Framework. Advances in Therapy.**
> *Social influence and organisational readiness are identified as key enablers of responsible AI integration in HEOR. Building a culture where human judgment is actively valued, where the person who surfaces a governance gap is recognised rather than experienced as creating delay, is a precondition for the hybrid intelligence framework to work as designed.*

SECTION 6

# The Six Questions That Tell You Where You Are.

These questions are not a maturity framework. They are a direct test of whether your current governance structure is real or nominal. They are worth working through with a specific recent piece of AI-assisted work in mind.

| | |
|---|---|
| **01** | Think of the last three pieces of work your team produced with ChatGPT, Copilot, Claude, or any other AI tool. For each one: who owns the most important framing or interpretive choices in that output, and could that person explain those choices to a payer or medical affairs leader today? |
| **02** | For each person currently named as reviewer or approver for AI-assisted work on your team: are they genuinely in the loop, interrogating framing choices, or on the loop, reading and approving? And for the tasks they oversee: does their level of involvement match what those tasks actually require? |
| **03** | In your most recent AI-assisted evidence synthesis or SLR section: who asked what the AI left out that a sceptical reviewer would notice? Who checked whether the framing would hold up under HTA or payer challenge? If nobody did, the adversarial practice does not exist on your team yet. |
| **04** | Is your current review process for AI-assisted outputs designed for what AI actually produces, outputs that look complete, read authoritatively, and embed judgment calls that are not always visible, or is it the same review process you applied to manually produced work, extended without redesign? |
| **05** | If a payer or HTA body asked tomorrow why a recent AI-assisted evidence synthesis framed the clinical data the way it did, who would answer? And are you confident they can explain the reasoning as if they made every framing choice deliberately? |
| **06** | Since your team started using AI tools regularly, has the depth of human engagement in reviewing those outputs increased, stayed the same, or decreased? If it has decreased, you are watching a ghost form. |

The question that matters most is the last one. AI adoption naturally accelerates. Governance naturally habituates. The gap between them is where ghosts live. The organisations that close that gap proactively, that redesign their governance structures as AI workflows mature rather than assuming the pilot-phase engagement will persist, are the ones whose AI-assisted submissions will hold up when the scrutiny arrives.

CLOSING

# Not Less Work. Different Work.

Every organisation using AI tools in HEOR will eventually face the same moment. Someone will ask about a specific piece of AI-assisted work. An evidence summary, a clinical synthesis, a dossier section. It will not be a question about accuracy. It will be a question about judgment: why was the evidence framed this way? Why was this study given the weight it received? Why did the synthesis reach this conclusion and not that one? It is a question that requires a human, a specific named capable human, to stand behind a choice and explain it as if they made it deliberately.

That moment is the test. Not of the AI tool. Of the governance structure built around it. And the organisations that pass that test are not the ones whose teams use the most sophisticated AI. They are the ones that understood what AI actually changes: not the amount of work in HEOR, but its distribution. Execution moves toward AI. Judgment moves toward the people senior enough to own it. That is not a burden. It is a clarification of what senior HEOR leadership is for.

The directors and VPs who will lead HEOR functions well in an AI-enabled environment are not those who manage AI adoption from a distance. They are those who recognise that the natural drift away from operational detail, a drift that has been building for years as teams grew and volumes increased, is no longer appropriate for the judgment layer of the work. AI has not made senior HEOR roles less important. It has made the judgment component of those roles more concentrated, more visible, and less possible to defer. The reset required is not a governance process. It is a leadership mindset shift: a return to the layer of the work that AI cannot do, and that only someone with submission-ready authority, named assumption ownership, and the habit of adversarial interrogation is equipped to own.

HEOR AI governance will not be defined by which organisations got to ChatGPT first. It will be defined by which organisations understood the redistribution, reset their leadership expectations accordingly, and built governance structures where the humans responsible for judgment are genuinely present in it. And which ones discover the gap when the question arrives that their Governance Ghost cannot answer.

> ***The question is not whether your AI tool is good enough. It is whether the human responsible for what it produced is still genuinely there, or whether you have a ghost.***

## Connect AI

Connect AI explores how life science industry must evolve in an AI-enabled world.
The series focuses on leadership decisions, governance, and operating models - not tools or technology in isolation.

Each edition examines a specific dimension of how evidence is governed, defended, and owned as AI becomes embedded in HEOR workflows.

Companion paper: Who Is The Human in Human-AI? An HTA Perspective. Available separately from ConnectHEOR.

Key reference: Srivastava T, Irfan H, Babiy V, Swami S. Integration of Generative AI with Human Expertise in HEOR: A Hybrid Intelligence Framework. Adv Ther. 2025;42:4103–4130. doi:10.1007/s12325-025-03273-w

A ConnectHEOR publication

**WWW.CONNECTHEOR.COM**