

FAQs:

What are the Risks and Mitigation Strategies in Using Gen-AI in HEOR?



Shilpi Swami, Judit Banhazi, Hanan Irfan, Tushar Srivastava | Affiliations: ConnectHEOR, London, UK | Email: shilpi.swami@connectheor.com

Introduction:

Despite the potential of generative artificial intelligence (Gen-AI) in health economics outcomes research (HEOR), concerns remain about data privacy, transparency, and methodological integrity. This frequently asked questions (FAQ)-style abstract aims to consolidate common questions related to Gen-AI risks and propose mitigation strategies for its safe and effective adoption in HEOR.

Methods:

A targeted literature review (2019-2024) was conducted, focusing on academic publications, industry white papers, and real-world implementation that addressed Gen-AI applications in HEOR. Key themes were extracted and categorized as “FAQs,” each detailing specific risks and outlining potential mitigation measures.

Discussion and Conclusion:

AI is not a human replacement, but collaboration which must be guided by strong governance, transparent systems, and consistent human oversight. The list presented here is not exhaustive, and further clinical review and broader testing are recommended. Future work will expand the risk library and evaluate the framework across disease areas and real-world data streams.

“ This study demonstrates that Gen-AI can strengthen HEOR workflows, but only when every “what-if” risk is systematically mitigated through a Human-in-the-Loop, governance-driven framework that embeds clinical context, transparency, and regulatory compliance. ”

“ AI ≠ Human replacement. It is about how well we understand the risks and apply mitigation strategies to streamline workflows and transform the way of decision-making. ”

Abbreviations:

AI, Artificial intelligence; API, Application Programming Interface; FAQ, Frequently Asked Questions; Gen-AI, Generative AI; HEOR, Health economics Outcomes Research

References:

1. Swami S, Srivastava T, Babiy V. Generative AI: The Next Frontier in Health Economic Model Conceptualization. ISPOR Europe 2024;  
2. Swami S, Srivastava T. Human Vs. Machine: AI Frameworks and Recommendations in HEOR. Poster MSR41, ISPOR Europe 2024;  
3. Srivastava T, Swami S. Thinking Enough? Evaluating LLM Reasoning Algorithms in HEOR, Poster MSR233, ISPOR Europe 2024;  
4. Fox MS. AI and expert system myths, legends, and facts. IEEE Expert: Intelligent Systems and Their Applications. 1990;5(1):8–20.

What If?

- Gen-AI delivers irrelevant insights?
- Our knowledge base is incomplete?
- Compute or LLM capacity is insufficient?
- Prompts yield shallow or off-target outputs?
- AI hallucinates or embeds bias?
- AI reasoning is a black box?
- Privacy or regulatory rules are breached?
- Governance or oversight is missing?
- AI produces critical errors?
- AI benefits can't be demonstrated?
- Gen-AI remains siloed from existing tools?
- Over-reliance weakens human expertise?
- AI advances outpace our practices?
- Gen-AI misinterprets clinical context?
- AI contradicts existing evidence or assumptions?
- Teams resist adopting Gen-AI tools?
- Regulatory bodies question AI-derived evidence?
- Outputs vary unpredictably between runs?
- Gen-AI tools become platform-dependent or inaccessible?

Mitigation Strategy

- Run small, focused pilots and iterate before scaling.
- Curate PICO, HTA guidelines, and public trial data in a structured Graph-RAG.
- Secure scalable cloud GPUs and choose a large-context model (e.g., GPT-4o).
- Translate each workflow step into Chain-of-Thought prompts and refine.
- Use diverse data sources and mandate expert review of every output.
- Log Chain-of-Thought and prompt-output pairs for full traceability.
- Anonymize data, restrict access, and audit against GDPR/HIPAA/HTA.
- Establish a Human-in-the-Loop workflow and a Gen-AI steering committee.
- Set error flags, re-run with adjusted prompts, and require expert sign-off.
- Track KPIs- time saved, edit rates, and concordance with real data.
- Integrate via APIs and provide no-code interfaces for seamless use.
- Rotate AI/human lead roles and schedule regular skills refresh.
- Maintain a governance group to pilot new models and update protocols.
- Provide disease-specific context (e.g., PICO), relevant clinical guidelines and expert validation.
- Cross-check AI outputs against literature and expert consensus during review.
- Offer training, demonstrate pilot results, and integrate into familiar workflows.
- Document AI reasoning steps and validation process; publish methodology where possible.
- Use fixed prompts, set temperature control in LLM, and log versions for reproducibility.
- Prioritize open platforms, exportable outputs, and fallback workflows for resilience.

Financial Disclosure: Authors are employees of ConnectHEOR Limited. No external funding received to conduct this research and no conflict of interest to declare.

