
RCTs FOR FRONTIER AI GOVERNANCE: METHODOLOGICAL CHALLENGES AND SOLUTIONS FOR HUMAN UPLIFT STUDIES

A PREPRINT

Patricia Paskov^{1,*}, Kevin Wei¹, Shen Zhou Hong², Dan Bateyko³, Xavier Roberts-Gaal⁴, Carson Ezell¹,
Gailius Praninskas⁵, Valerie Chen⁶, Umang Bhatt⁷, Ella Guest¹

¹RAND, Santa Monica, CA 90407, United States

²Johns Hopkins University, Baltimore, MD 21218, United States

³Cornell University, Ithaca, NY 14853, United States

⁴Harvard University, Cambridge, MA 02138, United States

⁵London School of Economics, London, WC2A 2AE, United Kingdom

⁶Carnegie Mellon University, Pittsburgh, PA 15213, United States

⁷University of Cambridge, Cambridge, CB2 1TN, United Kingdom

*Address correspondence to ppaskov@rand.org

July 1, 2026

ABSTRACT

Human uplift studies, or studies that measure the effects of AI access on human performance via randomized controlled trials (RCT) or similar methodologies, increasingly inform frontier AI governance and deployment decisions. While RCT methods are robust in other fields, their interaction with the distinctive properties of frontier AI systems remains underexamined, particularly when results are used to inform high-stakes decisions. We present findings from interviews with 16 expert practitioners with experience conducting human uplift studies in domains including biosecurity, cybersecurity, education, and labor. Across interviews, experts described a recurring tension between the standard causal inference assumptions upon which human uplift studies rely and the object of study itself. Rapidly evolving AI systems, shifting baselines, heterogeneous and changing user proficiency, and porous real-world settings strain assumptions underlying internal, external, and construct validity, complicating the interpretation and appropriate use of uplift evidence. We contribute (1) a synthesis of methodological challenges in human uplift studies, mapped to risks to study validity and classified by their degree of specificity to large language model (LLM) systems, and (2) a mapping from challenges to proposed solutions. By collating expert-identified challenges and solutions, we seek to clarify the interpretive limits and appropriate uses of human uplift evidence, to align evaluation practice with the decisions it informs, and to support more coordinated methodological foundations for AI governance.

Contents

1	Introduction	3
2	Background & related work	3
3	Methodology	4
3.1	Interview protocol	4
3.2	Thematic analysis	5
3.3	Validity mapping	5
3.4	LLM-specificity mapping	5
3.5	Limitations	5
4	Results	5
4.1	Descriptive statistics	5
4.1.1	Sample: Experts	5
4.1.2	Sample: Human uplift studies	6
4.2	Thematic analysis	6
4.2.1	Methodological Challenges	6
4.2.2	Practical Solutions for the Field	10
5	Discussion: Validity Tradeoffs, Collective Action, and Governance	12
6	Conclusion	12
	Appendices	22
A	Rapid Literature Review Methodology	22
B	Interview Methodology	22
C	Thematic Analysis Codes	27

1 Introduction

Large language model (LLM) systems are reshaping fundamental aspects of human society, from how we work and learn to how we make decisions about national security and public policy [1, 2]. Society’s ability to anticipate, manage, and govern these transformations depends in part on how well we can evaluate AI’s real-world impacts. Governance frameworks increasingly rely on evaluations as a precondition for responsible deployment: the EU AI Act mandates evaluations of frontier models [3]; US state-level laws such as New York’s RAISE Act and California’s SB 53 require public reporting of safety evaluations [4, 5]; and frontier developers voluntarily commit to evaluation-based risk thresholds through Responsible Scaling Policies and Frontier Safety Frameworks [e.g., 6, 7, 8]. Yet the evaluation methods on which this paradigm relies largely compare AI systems with each other, and often fall short in measuring how systems impact users and society in practice [9, 10, 11, 12].

Human uplift studies offer a way to close this gap. By measuring the extent to which access to or use of LLM systems shifts human performance on a task relative to a counterfactual through randomized controlled trial (RCT) or similar methodologies, uplift studies provide rigorous causal evidence about AI’s effects on the people who interact with it [13, 14]. Human uplift studies are increasingly discussed by international consortia [15, 13], governments [16, 17, 14, 18], developers [19, 20, 21, 22], and AI evaluation contexts [23, 24, 25, 26, 27, 28, 29, 30].¹

This promise, however, outpaces methodological understanding. While RCT methodology is well-established in fields including medicine, economics, and human-computer interaction [32, 33, 34, 35], its application to frontier AI systems poses new challenges. Moreover, study results are not always made public; even when published, results often omit key methodological details due to safety concerns, such as in Anthropic [19]. As uplift studies increasingly shape consequential decisions about AI deployment, safety, and regulation, gaps between the methodological foundations of this evidence and the institutional conditions under which it is produced and used become questions not only of research design but of AI governance.

This article fills that gap with findings from 16 interviews of experts who have conducted or were contemporaneously conducting human uplift studies involving large language models (LLMs). Our main contributions are: (1) a synthesis of methodological challenges in LLM

uplift studies, mapped to threats to construct, internal, and external validity and classified by their degree of LLM-specificity; and (2) a mapping from challenges to expert-proposed solutions. Our focus and analysis are centered on human uplift studies involving LLM-based systems, though we occasionally use the broader term “AI systems” for generality. By collating expert-identified challenges and solutions, we seek to clarify the interpretive limits and appropriate uses of human uplift evidence, and to promote the development of increasingly rigorous and robust AI evaluations for governance.

2 Background & related work

Evaluations provide stakeholders with critical insights into the capabilities, risks, and opportunities of frontier AI systems [1, 36, 19, 37, 14, 16]. Current approaches, including multiple-choice question-answer (MCQA) benchmarks, red-teaming, and long-form agent evaluations, present distinct trade-offs in validity, reproducibility, and resource requirements [38, 14, 13, 39, 40, 36, 41]. While MCQA benchmarks provide structured performance measurement, they often neglect system interaction with users or environments [11, 9, 10, 12]. As such, benchmarks alone poorly predict downstream impacts, especially on economically or strategically important tasks [42, 43, 44, 45]. Red-teaming approaches, while involving human-computer interaction, often lack the controlled structure for reliable causal estimates [46, 47].

Human uplift studies use RCT or similar methodology to measure the causal impacts of AI systems on human performance, drawing upon decades of experimental rigor from medicine, economics, and social science [10, 48, 11, 18, 49, 32, 13, 50]. RCTs originated in medical research as the gold standard for establishing causal relationships between interventions and outcomes, with rigorous standards developed over decades to minimize bias and ensure reliable inference [32]. RCT methodologies have been adopted across fields including human-computer interaction [33, 34, 51, 52, 53, 54, 55, 56, 57] and economics and social sciences [58, 59, 60]. Across these domains, RCTs are grounded in the *Potential Outcomes Framework*, which specifies the assumptions required for causal inference and unbiased treatment-effect estimation [61, 62].

In recent years, researchers have leveraged RCTs and other methods to evaluate the impact of LLMs on human performance across a growing set of domains, including biological threats [38, 20, 63, 64, 65], developer productivity [31, 66], legal services [9, 67], customer service [68], and academic research [69, 70]. This breadth spans both benefit-oriented questions, such as productivity gains, and worst-case misuse risks. As decision-makers and policymakers increasingly turn to human uplift studies to inform deployment and governance decisions, the methodological soundness of these studies carries more weight.

¹The term “human uplift study” has emerged as a consensus label within AI policy and evaluation discourse to describe this class of methods [15, 13, 14]. For clarity, we use “human uplift study” to refer broadly to RCT-based or quasi-experimental evaluations of AI’s effects on human task performance, while acknowledging that individual studies may vary in design details and rigor. We note that the term itself is a misnomer: not all such studies find that access to frontier AI systems increases human performance, e.g., Becker et al. [31].

Solutions	Methodological Challenges (by Research Phase)								
	Design	Recruitment		Execution			Documentation		
	C1 Measurement [C] *	C2 Controls [C] **	C3 Recruitment [E] *	C4 AI Literacy [I, E] **	C5 Intervention Fidelity [C, I] ***	C6 Interference [I] **	C7 Expectancy Effects [I, E] **	C8 Documentation [E] *	C9 Interpretation [E] *
S1 Standardized Task Libraries	✓							✓	✓
S2 Baseline & Control Conventions		✓				✓		✓	✓
S3 Leveling AI Literacy			✓					✓	✓
S4 Versioned Snapshots				✓				✓	✓
S5 Interference Management					✓			✓	✓
S6 Natural Experiments & Phased Rollouts		✓							✓
S7 AI-Accelerated Research Methods		✓	✓	✓	✓				✓
S8 Post-Hoc Analysis		✓	✓	✓	✓	✓			✓
S9 Information Security Frameworks							✓	✓	

Table 1: Methodological challenges and solutions in human uplift studies: a mapping between proposed solutions (rows) and methodological challenges (columns) identified in 16 expert interviews. Challenges are grouped by the primary research phase in which they arise, tagged by the primary form of validity they threaten (construct [C], internal [I], external [E]), and rated for LLM-specificity (from familiar RCT threats intensified by LLMs (*), to those with RCT precedent but inadequate standard mitigations (**), to challenges native to LLM-based systems (***)). Mapping is illustrative rather than exhaustive.

The application of RCTs to LLMs calls into question the core assumptions of the *Potential Outcomes Framework* [61, 62], including stable interventions and well-defined counterfactuals. While violations of these assumptions are documented even in mature fields [32], LLM-based systems place additional and distinctive pressures on these assumptions. Rapid model iteration can undermine intervention fidelity over the course of a study, while the widespread integration of AI tools into everyday workflows [71, 72] complicates the specification of controls and the prevention of contamination.²

Amid these methodological questions and an evolving empirical landscape, much of the practical experience of conducting human uplift studies remains undocumented. Many studies are conducted under security constraints or commercial confidentiality [73, 12, 74, 23], limiting transparency about design trade-offs, failure modes, and interpretive pitfalls. As a result, stakeholders often rely on human uplift evidence without a clear view of where standard assumptions hold, where they break

down, and how researchers adapt in practice. In this article, we seek to fill this transparency gap.

3 Methodology

We employ expert interviews as the primary method for this study. Human uplift studies, when published, often only report high-level results and frequently omit details on methodological trade-offs, execution challenges, and interpretive uncertainties. Expert interviews allow us to surface this tacit knowledge directly, including insights from studies that remain unpublished or delayed due to security or commercial constraints. This project was reviewed and deemed exempt by our organization’s Institutional Review Board.

3.1 Interview protocol

We conducted semi-structured expert interviews between July and August 2025 with experts who had conducted or were, at the time of interview, conducting at least one human uplift study of an LLM-based system. We selected expert participants using a snowball sampling method [75], seeded with a rapid literature review of published human uplift studies involving LLMs that fulfilled the cri-

²Intervention fidelity refers to whether the treatment actually delivered matches the treatment specified in the study design. In the AI context, this is threatened when models update, safety filters change, or system configurations shift during a study. For more discussion, see C5.

teria in Table 6. For full sampling procedures, refer to Appendix A. In total, we reached out to 53 experts and secured participation of 16 interviewees (32.08% participation rate). Interviews were scheduled for 60 minutes and lasted between 30–70 minutes. The semi-structured script addressed demographics, human uplift study history, methodological challenges and solutions, and open-ended questions. The methodology and script are further detailed in Appendix B and B, respectively.

3.2 Thematic analysis

Following interview completion and transcription, we conducted a qualitative thematic analysis [76] to identify high-level methodological challenges and solutions. We used a two-stage inductive approach, treating coding as an interpretive and reflexive process [77] to identify themes from the bottom up. The final codebook contained 30 codes across 7 categories (Appendix C), which we organized into broader themes (the Results section). Through thematic analysis, we aimed to identify high-level challenges and potential solutions surfaced in expert interviews. Each transcript was independently coded by two annotators. Following established qualitative research practices [78, 79, 80, 81], we did not calculate inter-rater reliability, given our goal to interpret core narratives rather than make objective or predictive claims.

3.3 Validity mapping

We map challenges to three established forms of validity: construct, internal, and external validity (Table 2). This typology derives from [82] and subsequent work [83, 84] and complements recent work in AI evaluation adopting and extending validity frameworks to assess the reliability and generalisability of model evaluations [85, 86].³ Mapping was independently completed by Author 1 and Author 2, with a third author adjudicating disagreements to reach consensus. Used as an analytic lens, this framework helps to identify when and how inferential claims are at

³Beyond construct, internal, and external validity, the Campbell tradition includes a fourth category, statistical conclusion validity, which we omit because practitioner-reported challenges concerned the definition, identification, and generalization of causal effects rather than statistical estimation. The Campbell framework has been widely adopted across the social sciences: political science adopted all four validity types [87], while economics’ “credibility revolution” [meyer1995natural, 88] draws on the same tradition, emphasizing internal and external validity while largely dropping construct validity from its methodological vocabulary. While there is no single universally-accepted validity classification, we adopt the Campbell typology as a minimal established framework that cleanly maps to the challenges emerging from interviews. We retain construct validity as a distinct category because many challenges practitioners face in AI evaluation concern whether their measures capture what they intend to measure. We consider this a question logically prior to, and distinct from, whether internal causal identification assumptions hold.

stake; and seeks to foster more well-informed design and interpretation of human uplift studies.

3.4 LLM-specificity mapping

We further map each challenge along a three-level spectrum reflecting the degree to which it emerges due to the distinctive properties of LLM-based systems (Table 3). Mapping was independently completed by Author 2 and Author 6, with a third author adjudicating disagreements to reach consensus.

3.5 Limitations

Our interview process was constrained by sample size and by potential sampling and response bias. Given the nascency of LLM human uplift research and the small population of experts in this area, however, our sample of 16 experts is in line with norms in expert interview research in AI ethics [81, 90, 91, 92, 93, 94, 95, 96, 97] and human-computer interaction [98, 99, 80, 100, 101].⁴ Our interview sample reflects broader biases representative of the field: experts were predominantly U.S.-based, male, and academically affiliated, with only one industry representative responding despite multiple outreach attempts. These limitations may affect the generalizability of findings to non-Western, non-English-speaking contexts, or industry contexts. Limited industry representation likely reflects disclosure constraints rather than limited engagement, underscoring the transparency gap motivating this study. Finally validity and LLM-specificity mappings reflect our structured analytic judgment rather than objective classification, and are intended as interpretive heuristics rather than sharp ontological claims.

4 Results

4.1 Descriptive statistics

We provide some summary statistics about our expert interviewees, as well as the studies that interviewees had conducted and that we discussed in interviews.

4.1.1 Sample: Experts

We interviewed 16 experts (Table 4), including 13 male and 3 female experts. Educational backgrounds included

⁴The population of experts with direct experience conducting LLM human uplift studies is small; in later rounds of our snowball sampling, interviewees frequently identified experts already on our outreach list, suggesting we were approaching the boundary of the reachable expert population. As one expert noted, “the AI uplift world, it’s still kind of small.” Yet content saturation in similar expert studies has been achieved with a dozen or fewer interviews [102, 103]. We observed that later interviews surfaced few new themes or challenges beyond those identified in earlier interviews, consistent with content saturation, though we did not conduct a formal saturation analysis. Our 32% participation rate was comparable to that of similar studies [104].

Validity type	Definition and common threats
Construct [C]	The extent to which study operations (e.g., interventions, measures, settings, participants) correspond to intended abstract constructs [83, 84, 85]. Threats arise when the treatment is ill-defined (e.g., the LLM intervention updates mid-study), when control conditions fail to represent the relevant counterfactual, or when outcome measures capture only a subset of pathways relevant to the decision-relevant construct.
Internal [I]	The extent to which the assumptions required to identify a causal effect between the explanatory variable and the outcome of interest are satisfied within the study context [62, 84]. Threats arise when treatment diffuses across experimental boundaries through spillovers or contamination, or when participants are differentially exposed to varying versions of an intervention.
External [E]	The extent to which internally valid causal effects generalize to different individuals, contexts, and outcomes [89, 84]. Threats arise when the recruited sample diverges from the decision-relevant population, when user proficiency co-evolves with the technology, or when baselines shift over time.

Table 2: Validity dimensions used to organize methodological challenges in human uplift studies.

Level	Definition
LLM-native ***	The causal mechanism generating the challenge is specific to properties of LLM-based systems; precedents in prior RCT contexts are weak or absent.
LLM-transformed **	Precedent exists in prior RCT contexts, but standard mitigations are inadequate for preserving the relevant form(s) of validity in LLM studies.
LLM-intensified *	Precedent exists in prior RCT contexts, and standard mitigations remain largely adequate; LLMs primarily increase the frequency or severity of the threat.

Table 3: Three-level classification of how distinctly each methodological challenge emerges from properties of LLM-based systems.

PhD ($n = 7$), master’s ($n = 6$), or bachelor’s degrees ($n = 3$). Interviewed experts were spread across a range of seniority levels, with 0–5 years of work experience ($n = 3$), 6–10 years of work experience ($n = 7$), 11–15 years of work experience ($n = 3$), and 16+ years of work experience ($n = 3$). Experts were affiliated with universities ($n = 6$), independent research institutions ($n = 7$), government ($n = 1$), an AI company ($n = 1$), and other organization ($n = 1$). Institutions were based in the U.S. ($n = 14$), Germany ($n = 1$), and unspecified ($n = 1$).

4.1.2 Sample: Human uplift studies

Experts had completed or were, at the time of interview, conducting between 1 and 6 LLM human uplift studies, with most working on just one ($n = 10$). We discussed 16 distinct studies total, including those published ($n = 9$), under review ($n = 4$), and not-expected-to-be-published due to proprietary or security concerns ($n = 3$). Studies focused primarily on biology/biological risk ($n = 6$), with others spanning software engineering, cybersecurity, medicine, social sciences, and other domains. Seven studies aimed to evaluate, in some capacity, the potential for LLM systems to enable worst-case misuse risks by an attacker or threat actor. Most research teams included domain experts ($n = 12$); and half included social scientists ($n = 8$). All studies except one were randomized controlled trials with at least two arms (LLM access vs. control) ($n = 15$). Sample sizes in discussed uplift studies ranged from under 20 to nearly 5000 participants (median = 110), with only three studies exceeding 1000 participants. Studies recruited participants primarily via

convenience sampling through partner organizations, social media, or targeted outreach.

4.2 Thematic analysis

We present our thematic analysis in two stages. First, we present a structured synthesis of methodological challenges, organized across the AI evaluation lifecycle as defined in [105, 24]. We annotate challenges with [C], [I], or [E] to indicate the primary forms of validity they threaten, as defined in Table 2. We furthermore tag challenges with a star rating ($\star$ to $\star\star\star$) to indicate its degree of LLM-specificity, as defined in Table 3. Some challenges, like intervention fidelity (C5), are distinct to LLM-based RCTs. Others are *familiar in form but transformed* by properties of frontier AI systems, such as control conditions (C2) and interference (C6). Still others, such as interpretation of results across models and time (C9), are well-established in experimental design but take on distinct form in LLM uplift studies, particularly in high-stakes governance contexts. We then summarize a key subset of solutions. Bolding within quotations is added for emphasis. Table 1 maps challenges to corresponding solutions.

4.2.1 Methodological Challenges

C1 Measurement Under Constrained and Incomplete Pathway Coverage (Design) [C] $\star$

A growing body of evidence documents limitations of measurement instruments in AI evaluation [106, 107, 39, 108, 109, 110, 111], with human uplift studies as

Expert ID	Organization Type	Expert ID	Organization Type
Expert A	Independent Research Institution	Expert I	University
Expert B	University	Expert J	University
Expert C	Independent Research Institution	Expert K	University
Expert D	Other	Expert L	AI Company
Expert E	Independent Research Institution	Expert M	Independent Research Institution
Expert F	Independent Research Institution	Expert N	Independent Research Institution
Expert G	Independent Research Institution	Expert O	University
Expert H	Governmental Institution	Expert P	University

Table 4: Overview of Experts

no exception. Experts described challenges in defining measurement instruments that proxy real-world behavior, particularly in safety and misuse-focused uplift studies, where action spaces are wide, settings are complex and adversarial, and motivating threat models are themselves often under-specified [1]. Several experts emphasized concerns about the realism of experimental tasks. Expert N discussed a cybersecurity study:

N: “One big challenge we’re facing is whether or not the tasks, the actual lab environments that we’re giving these proxy attackers is realistic enough of the real world. Are we representing the real cyber world in assigning these tasks? That’s one big challenge. And the implication is, if it’s not realistic, then who cares about the uplift that’s given?”

Experts noted that concerns about task realism are compounded by the need to constrain action spaces in order to make studies tractable. In domains with wide action spaces and multiple viable strategies, researchers must often focus on a subset of pathways, introducing trade-offs between granularity and coverage. Expert A remarked:

A: “Studies need to artificially constrain [misuse scenario] pathways in order to be able to study anything at all...I think it’s important that we don’t lose sight of that. Say that you have specified one particular pathway and you didn’t see uplift in that, **that doesn’t need to necessarily mean that you’re now able to rule out all the other pathways that you didn’t study.**”

Together, these challenges underscore that constrained and incomplete pathway coverage is often unavoidable in human uplift studies, but remains methodologically defensible only when task constraints are clearly specified, transparently communicated, and explicitly aligned with the research question of interest. This distinction is especially salient in governance and policy contexts, where results from a single constrained pathway may be interpreted, implicitly or explicitly, as speaking to a broader

space of real-world behaviors or risks than that which the study was designed to measure.

C2 Control Conditions in AI-Integrated Environments (Design) [C] **

Selecting appropriate control conditions, defined here as the within-study comparison group used to identify the causal effect of AI systems, emerged as a central design challenge shaping both the interpretation and comparability of human uplift studies.⁵ While control specification is a general concern in randomized experiments, experts noted its unique challenge in human uplift studies, given the increasing embeddedness of AI tools in everyday workflows. In such settings, identifying a realistic and meaningful counterfactual is often nontrivial.

Experts described wide variation in control definitions in practice. Some studies restrict participants to basic internet search or non-AI tools, while others provide access to human experts or alternative software systems. Experts noted that, unlike many traditional interventions, LLMs do not replace a single prior tool but are layered onto an existing ecosystem of technologies. As such, control conditions are inherently relative to a reference point, which, if not made explicit, can lead to misinterpretation of uplift estimates. Expert P noted:

P: *In any control setting, there’s going to be some technology that they have available to them, whether it be AI, or it used to be called AI but now it’s no longer AI, or maybe it’s some other AI tool...it’s always relative to something and maybe impacts interpretation. Think about the coding papers, for example, that study how Copilot impacts programmers. Well, before Copi-*

⁵Throughout the paper, we distinguish *control conditions*, which are study-specific comparison groups used for within-study causal identification, from *baselines*, which are reference points used to contextualize and compare uplift magnitudes across studies or over time. We further refer to *human baselines* as baselines anchored to prevailing human capabilities or workflows [105]. We avoid using “baseline” to denote pre-intervention measurements, though this too is a common use case in literature.

lot, there was Autocomplete and Tab-Complete and IDEs...

The appropriate choice of control conditions ultimately depends on the specific research question and the reference point or threshold of interest, for which experts suggest scientific consensus or stakeholder consultation may be useful.

C3 Recruiting Populations Aligned with Specialized Research Questions (Recruitment) [E] ★

Experts described recruitment as a central constraint on external validity, particularly when research questions target specialized or high-stakes domains such as biosecurity, cybersecurity, or law. Experts described two recurring recruitment constraints. First, in many governance-, safety- and security-oriented studies, the populations of greatest interest are not directly recruitable at all. When research seeks to model malicious or highly capable threat actors, researchers must rely on proxy populations such as students, professionals, or domain experts, whose motivations, incentives, and constraints differ systematically from those of real-world adversaries. In these cases, uplift estimates depend on how well study designs, incentives, and task framing approximate relevant behaviors.

Second, in other settings, the population of interest is often difficult or costly to recruit. For example, Expert E examined AI use in legal contexts but recruited law students rather than practicing lawyers, noting that “lawyers’ time at law firms is pretty costly ... you just run up against budget constraints and have to make these trade-offs.” Similar challenges arise in domains requiring specialized technical expertise or narrow sub-disciplinary knowledge. Researchers often trade representativeness for feasibility, sometimes favoring more curated samples to reduce variance or improve retention at the cost of smaller or less representative populations. Across scenarios, recruitment choices shape results. In policy-relevant settings, where decisions frequently hinge on rare or adversarial behaviors rather than average effects, failure to clearly communicate recruitment limitations can lead to overconfident extrapolation about populations or behaviors.

C4 Heterogeneous and Evolving AI Literacy (Recruitment) [I, E] ★★

Participant proficiency in the use of AI, or *AI literacy*, emerged as a salient consideration in the design and interpretation of human uplift studies. Experts emphasized that variation in AI literacy is especially consequential, given the degree to which performance depends on users’ prompts, output interpretation, and tool integration. Expert L noted:

L: *“If you’re selecting someone who is a complete novice, has never used GenAI, I think you will not see much uplift just because they don’t know*

how to use the tool, whereas if you just give the same person six months to learn about GenAI and ask them again the same question six months down the road, they might be successful in using GenAI in a particular manner.”

Heterogeneity in AI literacy poses distinct threats to both external and internal validity. If the population of interest is mis-specified, such that participant skill does not reflect the population of interest, uplift estimates may fail to generalize, undermining external validity. If variation in AI literacy is unevenly distributed or not controlled for within a study, it may act as a confounder, threatening internal validity by obscuring the causal relationship between AI assistance and observed outcomes.

C5 Intervention Fidelity Under Rapid Model Evolution (Execution) [C, I] ★★★

Experts identified intervention fidelity as a key challenge: unlike traditional interventions, the LLM under study can change mid-experiment, silently and asymmetrically (see, e.g., [A.3 112]; [113]). In human uplift studies, the intervention typically includes access to a particular AI model, often embedded within a broader tool ecosystem including system prompts, safety filters, plugins, or auxiliary tools. Experts noted that developer or provider changes to any of these underlying components, while often unflagged, can materially affect study outcomes. A lack of clear visibility into model versioning, common among experts, induces uncertainty about intervention fidelity. Expert N noted:

N: *“When we started talking about this experiment and designing it, we noticed that the publicly available models were capable of running code in their environments and installing different Python tools. For example, we are starting [now] to see a lot more refusals from the same exact model. So the model has undergone an update and we no longer have access to that snapshot of the previous instance of that same model. And so if you run a study over a period of three months in which that model is being updated and you’re unaware, you’re comparing apples and oranges.”*

When participants are exposed to materially different model versions or configurations at the same time, this introduces unbalanced heterogeneity in the intervention, threatening internal validity by violating the assumption that treatment is consistently defined across subjects. By contrast, when the intervention changes uniformly over time, for example, due to a globally deployed model update, internal validity may be preserved, provided that participants are affected symmetrically. However, rather than

estimating the effect of a single, fixed model, the study would then capture the effect of exposure to a changing system over a specified time horizon. The risk of such changes increases with study duration: brief laboratory tasks may face minimal exposure, while longitudinal or real-world studies spanning weeks or months are more vulnerable.

C6 Interference: Spillovers and Contamination in AI-Integrated Environments (Execution) [I] **

Experts highlighted challenges arising from spillovers and contamination in AI-integrated environments. *Spillovers* occur when exposure to AI diffuses indirectly, for example through social interaction, shared strategies, or collaboration, violating standard non-interference assumptions of RCTs. These risks are especially salient in settings with close cohort structures, such as classrooms, labs, workplaces, or training programs, where participants naturally exchange information. *Contamination*, by contrast, occurs when control group participants directly access restricted AI tools or comparable systems, violating experimental protocol. As Expert A noted:

A: *“I expect cheating to be much more salient in LLM uplift studies, especially if there is an internet-only control group. Contrast this with a clinical drug trial, where, if you’re not giving the control group the drugs, they’re probably not going to be able to acquire it.”*

Across both dynamics exists a trade-off: study designs that aim to improve external validity by approximating real-world, longer-term use often face heightened risks of spillovers and contamination. Short, tightly controlled studies, on the other hand, may better preserve causal identification, but fail to reflect real-world AI use.

C7 Expectancy Effects (Execution) [I, E] **

Experts discussed limits to blinding in human uplift studies: because interaction with AI systems is explicit, conversational, and central to task execution, blinding participants to AI treatment is often infeasible. As a result, observed effects may reflect not only the technical capabilities of the system, but also users’ expectations about what AI can do. In experimental psychology, such dynamics are often described *expectancy effects* (e.g. [114]). Expectancy effects complicate interpretation when uplift results are used to make claims about underlying model performance or to support cross-study comparisons. Experts discussed partial mitigations – such as blinding outcome assessors or analysts – though such measures do not fully resolve these dynamics. Instead, the central interpretive choice lies in whether expectation-driven behavior is treated as a confound to be minimized and controlled for or as an integral feature of realistic AI use. Blinding challenges limit internal validity when uplift estimates are in-

terpreted as isolating model capability, and external validity when results are compared across studies with different framing, populations, or deployment contexts.

C8 Documentation Under Proprietary and Security Constraints (Documentation) [E] *

Experts described documentation — or the reporting of evaluation results, metrics, methodological details, experimental materials, and datasets — as a challenge to understanding evidence and advancing methodologies. While evaluations in many fields follow established documentation standards [115, 116], human uplift studies fall short due to their relative nascency, along with proprietary and security concerns. Expert D noted:

D: *“It’s often very difficult to know what was actually done on an uplift study based on model cards. Did they have a day to do the task? A week? Was it a multiple choice exam? An intensive planning exercise?”*

This lack of detail hampers scientific progress and appropriate interpretation of results for governance applications and beyond. While Expert B noted contractual solutions can ensure publishing rights, the feasibility of such arrangements can depend on power dynamics and dependencies between researchers, evaluators, and model developers, which shape the negotiability of terms in practice.

C9 Interpreting Results Across Models and Time (Documentation) [E] *

While LLM uplift studies estimate a point-in-time effect of a specific model for a specific population on a specific task, their results often inform broader and longer-term governance questions on, for example, safeguards and deployment. The dynamics and constraints discussed in this section introduce challenges to understanding to what extent a given study’s estimates extend beyond its original conditions. Expert A noted:

A: *“There’s often a conflation of what study questions are of interest. There’s this overall question: how good are people at using AIs for this task right now? But often the study question that people want to answer in a safety context is: if we launch this model, and then in the next few years people become much better at using AIs, how well would they perform? ... I wish that studies could better isolate those two different questions and communicate clearly how their results should be interpreted.”*

The implications are greater when uplift results inform high-stakes decisions such as in safety cases and policy

frameworks. Expert D illustrated how quickly the underlying conditions drift:

D: *“Comparison over time is very difficult...if you run another study with the same group of people, **six months later, the world has probably changed in pretty meaningful ways.** People are more familiar with using LLMs...That’s going to change the way that they perform.”*

Interpretive challenges further compound in cross-model comparisons. Unlike MCQA benchmarks such as MMLU [117] or GPQA [118], human uplift studies cannot be easily rerun on updated models or alternative systems without substantial recruitment and execution costs. As a result, direct comparison across model versions is often infeasible in practice. Expert O explained:

O: *“If you ran a study with this model and then reviewers are like, well, why don’t you try model X or Y? **It’s not a static benchmark that you can just re-run.** You’d have to recruit a new group of participants, which is not necessarily realistic.”*

Interpretive challenges similarly extend to baseline specification, where baseline refers to the reference point used to contextualize and compare uplift magnitudes across studies or over time. As decision-relevant comparison groups shift, driven by improvements in open-source models or the increasing integration of AI into everyday tools, reliance on outdated baselines risks comparisons that no longer reflect realistic conditions. Simultaneously, this dynamic risks what Expert D characterized as a “boiling frog” problem, in which gradual changes in reference points obscure substantial shifts in absolute capability [119].

4.2.2 Practical Solutions for the Field

Experts proposed a range of solutions associated with the challenges above, spanning study-level practices to ecosystem-level interventions. We map these solutions to challenges below and in Table 1. In doing so, we aim to illustrate potential points of leverage rather than to advocate for any specific approach. This mapping is not exhaustive, prescriptive, or definitive.

S1 Standardized Task Libraries

Experts proposed developing shared benchmark tasks for different task domains, given “huge economies of scale” and suggested that consortia of organizations could jointly contribute tasks for the common good without bearing full individual costs. Experts emphasized the importance of determining which tasks are representative or appropriately difficult. To this end, Expert A suggested consulting

domain experts and surveying study designers, decision-makers, and relevant communities about proxy quality prior to study launch, noting that pre-registered expectations about real-world implications could help structure post-results interpretation. Expert B proposed using multiple KPIs to capture distinct dimensions of performance, while other experts similarly highlighted sub-task-level measurement as a way to preserve informative signal in complex, multi-step tasks or capabilities. More thoughtfully designed, crowd-sourced standardized task libraries can deepen and broaden measurement by encouraging systematic coverage of relevant questions and pathways and improving the validity of proxies through economies of scale.

Related challenges: C1, C8, C9

S2 Baseline and Control Selection Conventions

Experts advocated for clearer conventions distinguishing baseline and control selection to improve interpretability and comparability across uplift studies. Suggested practices included explicitly characterizing the *prior technology or workflow* being displaced when defining design and, where feasible, adopting standardized baseline and control bundles for common settings (e.g., academic research, professional services). Experts emphasized that both baseline and control choices should be made explicit and align with the specific reference point or decision threshold relevant to the downstream use of the results. Explicit baseline and control conventions could anchor results to a clear reference point, improving interpretability and cross-study comparability. Further, conventions could align evaluations with decision-relevant questions, support consistent documentation through shared standards, and foster crowd-sourced identification and mitigation of spillovers or contamination.

Related challenges: C2, C6, C8, C9

S3 Leveling and Accounting for AI Literacy

Experts reported a range of strategies for AI literacy challenges, including measuring and filtering for proficiency in recruitment, stratifying randomization on prior experience, providing introductory training to participants, and controlling for AI skill level in post-hoc analysis. Accounting for AI literacy makes user heterogeneity explicit, helping align research questions with the populations and interaction modes the study seeks to capture, while potentially mitigating effects of confounding variables. Clear documentation of skill supports transparent reporting and fosters more appropriate interpretation by clarifying which user groups to which results apply.

Related challenges: C4, C8, C9

S4 Versioned Evaluation Infrastructure, Snapshot Access, and Verification

Versioned AI systems and stable evaluation snapshots could allow researchers to measure, control for, or appropriately interpret intervention fidelity. Provider-side guar-

antees of fixed model versions for study duration, explicit version identifiers for models and system configurations, and mediated access to snapshots through secure or tiered research environments could all promote more rigorous research [120, 121]. The success of this solution hinges on coordination between developers, providers, and researchers. Such infrastructure need not imply public release of sensitive systems, but could instead support intervention fidelity, reproducibility, and interpretability under proprietary or security constraints. Emerging research on verification offers potential solutions [122, 123, 124, 125, 126].

Related challenges: C5, C8, C9

S5 Contamination and Spillover Management

Experts proposed mitigating contamination and spillover threats through multiple approaches: monitoring LLM usage by control groups and removing protocol violators; offering post-study “amnesty” periods where participants can admit to violations without penalty while allowing data exclusion; physically separating treatment and control groups through staggered scheduling or dispersed recruitment; and implementing technical controls such as providing restricted Chromebooks or network-level site blocking. Several experts emphasized the importance of designing incentive structures that encourage compliance rather than relying solely on monitoring and detection systems.⁶ Contamination and spillover management practices can mitigate interference issues and promote better documentation and interpretation of results.

Related challenges: C6, C8, C9

S6 Natural Experiments

Expert B identified phased product roll-outs and staggered deployments as underutilized opportunities for generating quasi-experimental variation in human uplift studies, noting “it’s not that expensive because [companies] are going to [deploy in phases] in many cases — you can’t give [the new technology] to everybody all at once anyways. So if they coordinate with academics, I think there’s a big opportunity to ... get experimental variation.” When roll-out timing or access is determined by operational constraints such as infrastructure limits, geographic sequencing, or user tiering, rather than individual characteristics, these settings can approximate natural experiments that support causal inference under weaker assumptions than fully randomized trials. In such cases, differences in exposure may plausibly be treated as exogenous to user ability or motivation, enabling estimation of uplift effects with reduced selection bias. Natural experiments constrain the research question to real-world adoption contexts and can sidestep recruitment challenges, given built-in inclusion

⁶Notably, in contexts where subjects lack strong performance incentives, contamination may be less problematic. This highlights a fundamental trade-off: while incentives simulate real-world motivations, they may also incentivize control groups to cheat.

of decision-relevant populations. Natural experiments can also support interpretation of results across models or deployment stages.

Related challenges: C3, C9

S7 AI-Accelerated Research Methods

Experts identified AI as a promising tool for addressing scalability challenges in uplift research, highlighting its potential to reduce timeline and cost constraints. For study design, Expert L suggested using AI agents as complementary participants in pilot studies to enable rapid testing of experimental protocols [e.g., 127, 128, 129, 130, 131]. Agent-based piloting could help, for example, surface brittle task designs, unintended affordances, or reward-hacking in complex multi-step settings at low cost. AI-accelerated methods may also enable rapid exploration of questions and behavior across diverse simulated human profiles under clean experimental conditions. At the same time, agent behavior may be limited in external validity and should be cautiously vetted, validated, and used as a complement to rather than substitute for human subjects research.

Related challenges: C3, C4, C5, C6

S8 Post-Hoc Analysis

Post-hoc analytical techniques can partially mitigate internal validity challenges in human uplift studies, particularly when ideal experimental control is infeasible. Careful ex post analysis, such as deriving heterogeneous treatment effects, adjusting for non-compliance, or reweighting observations to account for imbalance or attrition, can help diagnose and, in some cases, adjust for violations of key identification assumptions. Several experts emphasized that uplift studies generate unusually rich interaction logs, which enable finer-grained post-hoc audits of participant behavior. Expert A further noted that adhering to research standards, such as explicitly estimating causal effects rather than relying solely on directional hypothesis tests, can strengthen internal validity and downstream interpretation. While post-hoc analysis cannot substitute for robust experimental design, experts viewed it as a valuable complement.

Related challenges: C3, C4, C5, C6, C7, C9

S9 Information Security Advisory Frameworks & Tiered-Access

For studies presenting security or proprietary concerns, Expert A advocated for structured, deliberate consultation with relevant communities through advisory boards that include domain and scientific experts. For example, human uplift studies on medical tasks may draw from consensus among medical experts, while human uplift studies on bioterrorism tasks may draw from consensus among national security advisors. This approach could help balance security, proprietary, and transparency concerns by providing independent review of information hazard risks while maintaining scientific openness where

appropriate. Experts framed documentation not as a binary choice between openness and secrecy, but rather as a question of granularity and audience. Tiered-access reporting provides a practical mechanism for disclosing sufficient information to support interpretation of uplift results, while responsibly managing proprietary and security risks through controlled access [e.g., 132].

Related challenges: C8, C9

5 Discussion: Validity Tradeoffs, Collective Action, and Governance

Expert interviews surfaced three recurring methodological tensions in human uplift research. Each implicates construct [C], internal [I], and external validity [E], though through different mechanisms. *Realism versus control*, for example, trades ecological realism against causal identification: naturalistic settings invite the spillovers (C6) that threaten causal identification, while the controlled tasks that protect it may poorly reflect real use (C1). *Speed versus rigor* trades timely evidence against design depth: rigorous, externally valid studies of long-horizon tasks may require months or years from design to completion, a window during which the intervention and research question itself may change (C5, C9). *Disclosure versus protection* affects the assessment of all three validity types, with proprietary and security constraints (C8) limiting the reporting and documentation needed to effectively interpret and act on evidence.

These tensions carry direct consequences for AI governance: studies designed on weak foundations or reported without sufficient context may produce false confidence or, conversely, trigger unnecessarily burdensome obligations. Well-designed and clearly documented studies, by contrast, can strengthen the evidence base for governance. Crucially, validity concerns do not operate in isolation: design choices that strengthen one form of validity often weaken another. Policy decisions should therefore rest on convergent evidence from multiple studies using complementary approaches. In governance decisions in which the relevant question shifts faster than any one study can resolve it, uplift estimates may be combined with approaches such as scenario analysis, forecasting, and methods designed to inform decisions under deep uncertainty [133, 134]. In all cases, rigorous documentation to address C9, paired with tiered-access reporting (S9) where relevant, is a precondition for effective use of uplift evidence.

While some of the highlighted challenges (e.g. C1, C3, C8, C9), familiar in form, can turn to established methodological precedent for some degree of guidance [e.g., 89, 115, 135], a novel set of challenges stems from properties distinctive to LLM-based systems (e.g. C5 and, to varying degrees, C2, C4, C6, C7). For these challenges, existing methodological toolkits offer weaker guidance. The solutions proposed by interviewed experts represent early attempts to fill this gap. Each solution carries its own va-

lidity tradeoffs, which may be empirically explored in future work.

Beyond any individual uplift study lies a meta-level need for collective action. Several proposed solutions (S1, S2, S9) depend directly on field-level coordination; without this coordination, challenges C1–C9 are more likely to persist. At present, fragmented infrastructure forces research teams to repeatedly reinvent methods, while proprietary evaluations and misaligned research incentives constrain knowledge-sharing. Field-level workshops, shared infrastructure, and consensus-building processes could enable researchers to exchange lessons and co-develop novel solutions for AI evaluation and governance. Fields with deep experience in RCTs and causal inference (e.g. medicine, psychology, and economics) remain severely underleveraged within human uplift research for both familiar and novel challenges: cross-domain synergies and bridges could meaningfully advance human uplift, evaluation, and related governance efforts. In this context, public and philanthropic actors are uniquely positioned to support evaluations and initiatives that no single organization has incentives to undertake.

6 Conclusion

Human uplift studies, or studies that measure AI effects on human performance relative to a status quo, offer a framework for measuring the societal impacts of AI systems. By studying how AI systems impact human performance in realistic tasks, uplift studies can inform governance decisions in ways that model-centric static benchmarks alone cannot. At the same time, the application of RCT methodologies to frontier AI systems introduces distinctive challenges that complicate both study design and interpretation and, subsequently, threaten to confound downstream decision-making.

Across 16 expert interviews, we find that these challenges map onto familiar threats to construct, internal, and external validity, amplified by the properties of frontier AI systems. These trade-offs carry direct implications for policy and governance of AI systems. No single uplift study, however well designed, can provide a definitive assessment of system safety or societal impact. Policy-relevant conclusions should therefore rest on convergent evidence drawn from multiple studies with complementary designs, rather than on isolated findings.

Looking forward, advancing the rigor and usefulness of human uplift research will require more than incremental improvements in individual study design. Established disciplines hold relevant expertise and experts identified a range of novel and practical solutions, including standardized task libraries, clearer baseline and control conventions, versioned evaluation infrastructure, and structured approaches to interference, that can partially mitigate these challenges. Most importantly, however, the field requires coordination mechanisms to address collec-

tive action problems, enabling knowledge and resources to accumulate rather than remain siloed.

As human uplift studies play a growing role in informing decisions and policymaking, strengthening the methodological foundations of uplift research and clarifying what these studies can and cannot support becomes essential for responsible governance. Doing so can contribute to a coherent and policy-relevant body of evidence towards more beneficial AI futures.

Acknowledgements

The authors are grateful to the expert interviewees who generously volunteered their time, experience, and judgment to participate in this study. Their insights substantially informed both the framing of the research and the interpretation of its findings. The authors also thank the leadership of RAND's Center on AI, Security, and Technology (CAST) for their support of this work. RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was undertaken through CAST, which examines the opportunities and risks of rapid technological change, with a focus on artificial intelligence, security, and biotechnology. For more information, contact cast@rand.org.

Ethical Considerations Statement

This study was reviewed by RAND's Institutional Review Board and deemed exempt under 2024-N0632-MOD-06. The research draws on semi-structured interviews with expert practitioners and involved no experimentation, intervention, or deception. All experts provided informed consent prior to participation and were informed of the voluntary nature of the study, their right to decline to answer questions, and their ability to withdraw at any time without penalty. Interviews were conducted using secure videoconferencing infrastructure and, with expert permission, audio-recorded for transcription. Raw recordings and non-anonymized transcripts were accessible only to authorized project staff and stored on encrypted organizational systems. Recordings were deleted following transcription and anonymization. Anonymized transcripts were reviewed to remove identifying details prior to analysis or sharing with collaborators. Experts were asked how they wished to be identified in any published materials, and no quotations were attributed to individuals by name. All quotations included in the paper were cleared with experts. These procedures were designed to minimize risks of inadvertent disclosure or reputational harm while enabling transparent reporting of methodological insights.

Funding

This research was independently initiated and conducted within the RAND's Center on AI, Security, and Technology (CAST) using income from operations and gifts and grants from philanthropic supporters. A complete list of donors and funders is available at www.rand.org/CAST. RAND clients, donors, and grantors have no influence over research findings or recommendations.

References

- [1] Yoshua Bengio et al. *International AI Safety Report 2026*. Tech. rep. DSIT 2026/001. Accessed: 2026-05-20. UK Department for Science, Innovation and Technology, Feb. 2026. URL: <https://internationalaisafetyreport.org>.
- [2] Nestor Maslej et al. *Artificial Intelligence Index Report 2025*. 2025. arXiv: 2504 . 07139 [cs . AI]. URL: <https://arxiv.org/abs/2504.07139>.
- [3] European Parliament and Council of the European Union. *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. July 2024. URL: <http://data.europa.eu/eli/reg/2024/1689/oj>.
- [4] New York State Assembly. *RAISE Act, Assembly Bill A09449*. https://assembly.state.ny.us/leg/?default_fld=&leg_video=&bn=A09449&term=2025&Summary=Y&Actions=Y&Memo=Y&Text=Y. Accessed: 2026-03-10. 2025.
- [5] California State Legislature. *Senate Bill 53: Artificial Intelligence Models: Large Developers*. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB53. Accessed: 2026-03-10. Sept. 2025.
- [6] Anthropic. *Anthropic's Responsible Scaling Policy*. en. Sept. 2023. URL: <https://perma.cc/D9HQ-96HW> (visited on 09/30/2023).
- [7] Google DeepMind. *Frontier Safety Framework*. <https://deepmind.google/blog/introducing-the-frontier-safety-framework/>. Accessed: 2026-05-20. 2024.
- [8] OpenAI. *Preparedness Framework (Version 2)*. <https://cdn.openai.com/pdf/18a02b5d-6b67-4cec-ab64-68cdfbdebcd/preparedness-framework-v2.pdf>. Accessed: 2026-05-20. 2025.
- [9] Reva Schwartz et al. *Reality Check: A New Evaluation Ecosystem Is Necessary to Understand AI's Real World Effects*. 2025. arXiv: 2505 . 18893 [cs . CY]. URL: <https://arxiv.org/abs/2505.18893>.

- [10] Lujain Ibrahim et al. *Towards interactive evaluations for interaction harms in human-AI systems*. arXiv:2405.10632 [cs]. July 2025. DOI: 10.48550/arXiv.2405.10632. URL: <http://arxiv.org/abs/2405.10632> (visited on 08/21/2025).
- [11] Laura Weidinger et al. *Sociotechnical Safety Evaluation of Generative AI Systems*. 2023. arXiv: 2310.11986 [cs.AI]. URL: <https://arxiv.org/abs/2310.11986>.
- [12] Ilan Strauss et al. *The State of AI Governance Research: AI Safety and Reliability in Real World Commercial Deployment*. Apr. 2025. DOI: 10.35650/aidp.4112.d.2025. URL: <http://dx.doi.org/10.35650/AIDP.4112.d.2025>.
- [13] Frontier Model Forum. *Frontier Capability Assessments*. en-US. Tech. rep. Frontier Model Forum, Apr. 2025. URL: <https://www.frontiermodelforum.org/technical-reports/frontier-capability-assessments/> (visited on 08/21/2025).
- [14] UK AI Security Institute. *AI Safety Institute approach to evaluations*. en. Feb. 2024. URL: <https://www.gov.uk/government/publications/ai-safety-institute-approach-to-evaluations/ai-safety-institute-approach-to-evaluations> (visited on 08/21/2025).
- [15] Yoshua Bengio et al. *International Scientific Report on the Safety of Advanced AI: Interim Report*. Research Report DSIT 2024/009. Seoul, Republic of Korea: AI Seoul Summit, May 2024. URL: <https://www.gov.uk/government/publications/international-scientific-report-on-the-safety-of-advanced-ai>.
- [16] European Commission. *Code of Practice for General-Purpose AI Models: Safety and Security Chapter*. July 2025. URL: <https://perma.cc/U6FM-RU8Y>.
- [17] UK AI Security Institute. *Early lessons from evaluating frontier AI systems — AISI Work*. en. Oct. 2024. URL: <https://www.aisi.gov.uk/work/early-lessons-from-evaluating-frontier-ai-systems> (visited on 04/30/2025).
- [18] U.S. AI Safety Institute. *NIST AI 800-1 Managing Misuse Risk for Dual-Use Foundation Models*. July 2024. URL: <https://perma.cc/TB4L-NVHK>.
- [19] Anthropic. *System Card: Claude Opus 4 & Claude Sonnet 4*. May 2025. URL: <https://perma.cc/D234-7WQ2>.
- [20] Tejal Patwardhan et al. *Building an early warning system for LLM-aided biological threat creation*. en-US. Jan. 2024. URL: <https://web.archive.org/web/20240508103517/https://openai.com/index/building-an-early-warning-system-for-llm-aided-biological-threat-creation/> (visited on 08/21/2025).
- [21] Rohin Shah et al. *An Approach to Technical AGI Safety and Security*. arXiv:2504.01849 [cs]. Apr. 2025. DOI: 10.48550/arXiv.2504.01849. URL: <http://arxiv.org/abs/2504.01849> (visited on 08/21/2025).
- [22] Shanghai AI Lab et al. *Frontier AI Risk Management Framework in Practice: A Risk Analysis Technical Report*. arXiv:2507.16534 [cs] version: 2. July 2025. DOI: 10.48550/arXiv.2507.16534. URL: <http://arxiv.org/abs/2507.16534> (visited on 08/21/2025).
- [23] Tegan McCaslin et al. *STREAM (ChemBio): A Standard for Transparently Reporting Evaluations in AI Model Reports*. arXiv:2508.09853 [cs]. Aug. 2025. DOI: 10.48550/arXiv.2508.09853. URL: <http://arxiv.org/abs/2508.09853> (visited on 08/21/2025).
- [24] Patricia Paskov et al. *Preliminary suggestions for rigorous GPAI model evaluations*. 2025. DOI: 10.7249/pea3971-1. URL: <http://dx.doi.org/10.7249/PEA3971-1>.
- [25] METR. *What should companies share about risks from frontier AI models?* en. June 2025. URL: <https://metr.org/blog/2025-06-27-risk-transparency/> (visited on 08/21/2025).
- [26] Future of Life Institute. *2025 AI Safety Index*. en-US. Tech. rep. Future of Life Institute, July 2025. URL: <https://futureoflife.org/ai-safety-index-summer-2025/> (visited on 08/21/2025).
- [27] Ardi Janjeva et al. *Evaluating Malicious Generative AI Capabilities*. en. Tech. rep. The Alan Turing Institute, July 2024. URL: <https://cetas.turing.ac.uk/publications/evaluating-malicious-generative-ai-capabilities> (visited on 08/21/2025).
- [28] Ben Bucknall et al. “In Which Areas of Technical AI Safety Could Geopolitical Rivals Cooperate?” In: *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’25. New York, NY, USA: Association for Computing Machinery, June 2025, pp. 3148–3161. ISBN: 979-8-4007-1482-5. DOI: 10.1145/3715275.3732201. URL: <https://dl.acm.org/doi/10.1145/3715275.3732201> (visited on 08/20/2025).
- [29] Claudia Wilson. “The US Can Win Without Compromising AI Safety”. en. In: *Tech Policy Press* (Nov. 2024). URL: <https://techpolicy.press/the-us-can-win-without-compromising-ai-safety> (visited on 08/21/2025).
- [30] Holden Karnofsky. *If-Then Commitments for AI Risk Reduction*. en. Tech. rep. Carnegie Endowment for International Peace, Sept. 2024. URL: <https://carnegieendowment.org/research/2024/09/if-then-commitments-for-ai-risk-reduction?lang=en> (visited on 08/21/2025).

- [31] Joel Becker et al. *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*. arXiv:2507.09089 [cs]. July 2025. DOI: 10.48550/arXiv.2507.09089. URL: <http://arxiv.org/abs/2507.09089> (visited on 08/21/2025).
- [32] Alex Eble, Peter Boone, and Diana Elbourne. “On Minimizing the Risk of Bias in Randomized Controlled Trials in Economics”. en. In: (2017).
- [33] Xin Tong et al. “Just Do Something: Comparing Self-proposed and Machine-recommended Stress Interventions among Online Workers with Home Sweet Office”. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI ’23. New York, NY, USA: Association for Computing Machinery, Apr. 2023, pp. 1–20. ISBN: 978-1-4503-9421-5. DOI: 10.1145/3544548.3581319. URL: <https://dl.acm.org/doi/10.1145/3544548.3581319> (visited on 08/21/2025).
- [34] André Dahlinger et al. “The Impact of Abstract vs. Concrete Feedback Design on Behavior Insights from a Large Eco-Driving Field Experiment”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 1–11. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173953. URL: <https://doi.org/10.1145/3173574.3173953> (visited on 08/21/2025).
- [35] Cameron Robert Jones et al. “People cannot distinguish GPT-4 from a human in a Turing test”. In: *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’25. New York, NY, USA: Association for Computing Machinery, June 2025, pp. 1615–1639. ISBN: 979-8-4007-1482-5. DOI: 10.1145/3715275.3732108. URL: <https://dl.acm.org/doi/10.1145/3715275.3732108> (visited on 01/08/2026).
- [36] John Burden et al. *Paradigms of AI Evaluation: Mapping Goals, Methodologies and Culture*. arXiv:2502.15620 [cs]. June 2025. DOI: 10.48550/arXiv.2502.15620. URL: <http://arxiv.org/abs/2502.15620> (visited on 07/20/2025).
- [37] Bill Anderson-Samways et al. “Responsible Scaling: Comparing Government Guidance and Company Policy”. en. In: (2024).
- [38] Christopher A. Mouton, Caleb Lucas, and Ella Guest. *The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study*. en. Tech. rep. RAND Corporation, Jan. 2024. URL: https://www.rand.org/pubs/research_reports/RRA2977-2.html (visited on 05/09/2024).
- [39] Anka Reuel et al. *BetterBench: Assessing AI Benchmarks, Uncovering Issues, and Establishing Best Practices*. 2024. arXiv: 2411.12990 [cs.AI]. URL: <https://arxiv.org/abs/2411.12990>.
- [40] Lizhi Lin et al. *Against The Achilles’ Heel: A Survey on Red Teaming for Generative Models*. 2024. arXiv: 2404.00629 [cs.CL]. URL: <https://arxiv.org/abs/2404.00629>.
- [41] Kevin Wei et al. “Methodological Challenges in Agentic Evaluations of AI Systems”. In: *ICML Workshop on Technical AI Governance (TAIG)*. 2025. URL: <https://openreview.net/forum?id=ZhSKG8IslC>.
- [42] Rylan Schaeffer et al. “Why Has Predicting Downstream Capabilities of Frontier AI Models with Scale Remained Elusive?” en. In: June 2025. URL: <https://openreview.net/forum?id=IINtLvJal> (visited on 08/08/2025).
- [43] Michael Brooks. “Is your AI benchmark lying to you?” en. In: *Nature* 644.8075 (Aug. 2025). Bandiera_abtest: a Cg_type: Technology Feature Publisher: Nature Publishing Group Subject_term: Technology, Machine learning, Mathematics and computing, Genomics, pp. 294–296. ISSN: 1476-4687. DOI: 10.1038/d41586-025-02462-5. URL: <https://www.nature.com/articles/d41586-025-02462-5> (visited on 08/08/2025).
- [44] Arvind Narayanan and Sayash Kapoor. *GPT-4 and professional benchmarks: the wrong answer to the wrong question*. en. Apr. 2025. URL: <https://www.aisnakeoil.com/p/gpt-4-and-professional-benchmarks> (visited on 08/08/2025).
- [45] Anson Ho and J. S. Denain. *The real reason AI benchmarks haven’t reflected economic impacts*. Substack newsletter. Mar. 2025. URL: <https://epochai.substack.com/p/the-real-reason-ai-benchmarks-havent> (visited on 08/08/2025).
- [46] Michael Feffer et al. “Red-Teaming for Generative AI: Silver Bullet or Security Theater?” In: *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*. AAAI Press, Feb. 2025, pp. 421–437. (Visited on 09/11/2025).
- [47] Sorelle Friedler et al. *AI Red-Teaming Is Not a One-Stop Solution to AI Harms: Recommendations for Using Red-Teaming for AI Accountability*. en-US. Oct. 2023. URL: <https://datasociety.net/library/ai-red-teaming-is-not-a-one-stop-solution-to-ai-harms-recommendations-for-using-red-teaming-for-ai-accountability/> (visited on 03/24/2024).
- [48] Roel Dobbe, Thomas Krendl Gilbert, and Yonatan Mintz. *Hard Choices in Artificial Intelligence*. 2021. arXiv: 2106.11022 [cs.CY]. URL: <https://arxiv.org/abs/2106.11022>.
- [49] Henry Farrell et al. “Large AI models are cultural and social technologies”. en. In: *Science* 387.6739 (Mar. 2025), pp. 1153–1156.

- [50] Anson Ho and Arden Berg. *Do the biorisk evaluations of AI labs actually measure the risk of developing bioweapons?* 2025. URL: <https://epoch.ai/gradient-updates/do-the-biorisk-evaluations-of-ai-labs-actually-measure-the-risk-of-developing-bioweapons>.
- [51] Thomas W. Price et al. “Engaging Students with Instructor Solutions in Online Programming Homework”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI ’20. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–7. ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376857. URL: <https://dl.acm.org/doi/10.1145/3313831.3376857> (visited on 08/21/2025).
- [52] Sooyeon Jeong et al. “Huggable: The Impact of Embodiment on Promoting Socio-emotional Interactions for Young Pediatric Inpatients”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 1–13. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3174069. URL: <https://doi.org/10.1145/3173574.3174069> (visited on 08/21/2025).
- [53] Smitha Milli, Luca Belli, and Moritz Hardt. “Causal Inference Struggles with Agency on Online Platforms”. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’22. New York, NY, USA: Association for Computing Machinery, June 2022, pp. 357–365. ISBN: 978-1-4503-9352-2. DOI: 10.1145/3531146.3533103. URL: <https://dl.acm.org/doi/10.1145/3531146.3533103> (visited on 01/08/2026).
- [54] HaeJin Lee et al. “Learning Behaviors Mediate the Effect of AI-powered Support for Metacognitive Calibration on Learning Outcomes”. In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. CHI ’25. New York, NY, USA: Association for Computing Machinery, Apr. 2025, pp. 1–18. ISBN: 979-8-4007-1394-1. DOI: 10.1145/3706598.3713960. URL: <https://dl.acm.org/doi/10.1145/3706598.3713960> (visited on 08/21/2025).
- [55] Maurice Jakesch et al. “AI-Mediated Communication: How the Perception that Profile Text was Written by AI Affects Trustworthiness”. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–13. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300469. URL: <https://doi.org/10.1145/3290605.3300469> (visited on 08/21/2025).
- [56] Elliot G. Mitchell et al. “From Reflection to Action: Combining Machine Learning with Expert Knowledge for Nutrition Goal Recommendations”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI ’21. New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–17. ISBN: 978-1-4503-8096-6. DOI: 10.1145/3411764.3445555. URL: <https://doi.org/10.1145/3411764.3445555> (visited on 08/21/2025).
- [57] Jonathan Bassen et al. “Reinforcement Learning for the Adaptive Scheduling of Educational Activities”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI ’20. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–12. ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376518. URL: <https://doi.org/10.1145/3313831.3376518> (visited on 08/21/2025).
- [58] Abhijit Vinayak Banerjee. “Field Experiments and the Practice of Economics”. en. In: *American Economic Review* 110.7 (July 2020), pp. 1937–1951. ISSN: 0002-8282. DOI: 10.1257/aer.110.7.1937. URL: <https://pubs.aeaweb.org/doi/10.1257/aer.110.7.1937> (visited on 09/04/2025).
- [59] Abhijit Vinayak Banerjee, Esther Duflo, and Michael Kremer. “The Influence of Randomized Controlled Trials on Development Economics Research and on Development Policy”. en. In: *The State of Economics, the State of the World*. Ed. by Kaushik Basu, David Rosenblatt, and Claudia Sepúlveda. The MIT Press, Jan. 2020, pp. 439–487. ISBN: 978-0-262-35347-2. DOI: 10.7551/mitpress/11130.003.0015. URL: <https://direct.mit.edu/books/book/4917/chapter/624664/The-Influence-of-Randomized-Controlled-Trials-on> (visited on 09/04/2025).
- [60] Joshua D Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. en. Princeton University Press, Jan. 2009. ISBN: 978-0-691-12035-5. URL: <https://press.princeton.edu/books/paperback/9780691120355/mostly-harmless-econometrics>.
- [61] Donald B Rubin. “Estimating causal effects of treatments in randomized and nonrandomized studies.” In: *Journal of educational Psychology* 66.5 (1974), p. 688.
- [62] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press, 2015.
- [63] Satyapriya Krishna et al. *Evaluating the Critical Risks of Amazon’s Nova Premier under the Frontier Model Safety Framework*. en. July 2025. URL: <https://arxiv.org/abs/2507.06260v1> (visited on 03/23/2026).
- [64] xAI. *Grok 4.1 Model Card*. Tech. rep. Nov. 2025. URL: <https://data.x.ai/2025-11-17-grok-4-1-model-card.pdf>.

- [65] Chen Bo Calvin Zhang et al. *LLM Novice Uplift on Dual-Use, In Silico Biology Tasks*. en. Feb. 2026. URL: <https://arxiv.org/abs/2602.23329v2> (visited on 03/23/2026).
- [66] Valerie Chen et al. *Code with Me or for Me? How Increasing AI Automation Transforms Developer Workflows*. arXiv:2507.08149 [cs]. July 2025. DOI: 10.48550/arXiv.2507.08149. URL: <http://arxiv.org/abs/2507.08149> (visited on 09/11/2025).
- [67] Jonathan H. Choi and Daniel Schwarcz. “AI Assistance in Legal Analysis: An Empirical Study”. In: *Journal of Legal Education* 73 (2024), p. 384. URL: <https://heinonline.org/HOL/Page?handle=hein.journals/jled73&id=388&div=&collection=>.
- [68] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. “Generative AI at work”. In: *The Quarterly Journal of Economics* 140.2 (2025), pp. 889–942.
- [69] Abel Brodeur et al. *Comparing Human-Only, AI-Assisted, and AI-Led Teams on Assessing Research Reproducibility in Quantitative Social Science*. eng. Working Paper 195. I4R Discussion Paper Series, 2025. URL: <https://www.econstor.eu/handle/10419/308508> (visited on 08/22/2025).
- [70] Marc Ratkovic et al. *Harnessing GPT for Enhanced Academic Writing: Evidence from a Field Experiment with Early-Career Researchers in the Social Sciences*. en. SSRN Scholarly Paper. Rochester, NY, June 2025. DOI: 10.2139/ssrn.5313034. URL: <https://papers.ssrn.com/abstract=5313034> (visited on 07/21/2025).
- [71] Menlo Ventures. 2025: *The State of Generative AI in the Enterprise*. <https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/>. Dec. 2025.
- [72] McKinsey & Company. *The State of AI: Global Survey 2025*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>. 2025.
- [73] Toby Shevlane and Allan Dafoe. “The Offense-Defense Balance of Scientific Knowledge: Does Publishing AI Research Reduce Misuse?” In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. AIES ’20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 173–179. ISBN: 9781450371100. DOI: 10.1145/3375627.3375815. URL: <https://doi.org/10.1145/3375627.3375815>.
- [74] Oscar Delaney, Oliver Guest, and Zoe Williams. *Mapping Technical Safety Research at AI Companies: A literature review and incentives analysis*. 2024. arXiv: 2409.07878 [cs.CY]. URL: <https://arxiv.org/abs/2409.07878>.
- [75] C. Parker, S. Scott, and A. Geddes. “Snowball Sampling”. en. In: *SAGE Research Methods Foundations* (Sept. 2019). Publisher: SAGE. URL: <http://methods.sagepub.com/foundations/snowball-sampling> (visited on 08/21/2025).
- [76] Virginia Braun and Victoria Clarke. “Using thematic analysis in psychology”. In: *Qualitative Research in Psychology* 3.2 (Jan. 2006). Publisher: Routledge. eprint: <https://doi.org/10.1191/1478088706qp063oa>, pp. 77–101. ISSN: 1478-0887. DOI: 10.1191/1478088706qp063oa. URL: <https://doi.org/10.1191/1478088706qp063oa> (visited on 08/22/2025).
- [77] Virginia Braun and Victoria Clarke. “Toward good practice in thematic analysis: Avoiding common problems and becoming a knowing researcher”. In: *International Journal of Transgender Health* 24.1 (Jan. 2023). Publisher: Taylor & Francis. eprint: <https://doi.org/10.1080/26895269.2022.2129597>, pp. 1–6. ISSN: 2689-5269. DOI: 10.1080/26895269.2022.2129597. URL: <https://doi.org/10.1080/26895269.2022.2129597> (visited on 08/28/2025).
- [78] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. “Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice”. In: *Proc. ACM Hum.-Comput. Interact.* 3.CSCW (Nov. 2019), 72:1–72:23. DOI: 10.1145/3359174. URL: <https://doi.org/10.1145/3359174> (visited on 07/17/2024).
- [79] Virginia Braun et al. “Thematic Analysis”. In: *Advanced Research Methods for Applied Psychology*. 2nd ed. Num Pages: 11. Routledge, 2024. ISBN: 978-1-003-36271-5.
- [80] Tianshi Li et al. “Understanding Challenges for Developers to Create Accurate Privacy Nutrition Labels”. In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI ’22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–24. ISBN: 978-1-4503-9157-3. DOI: 10.1145/3491102.3502012. URL: <https://dl.acm.org/doi/10.1145/3491102.3502012> (visited on 07/31/2025).
- [81] Kevin Wei et al. “How Do AI Companies ”Fine-Tune” Policy? Examining Regulatory Capture in AI Governance”. en. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7.1 (Oct. 2024). Number: 1, pp. 1539–1555. ISSN: 3065-8365. DOI: 10.1609/aies.v7i1.31745. URL: <https://ojs.aaai.org/index.php/AIES/article/view/31745> (visited on 02/05/2025).
- [82] Donald T. Campbell. “Factors relevant to the validity of experiments in social settings”. In: *Psychological Bulletin* 54.4 (1957), pp. 297–312. DOI: 10.1037/h0040950.
- [83] Donald T Campbell and Thomas D Cook. “Quasi-experimentation”. In: *Chicago, IL: Rand McNally* 1.1 (1979), pp. 1–384.

- [84] Thomas D Cook, Donald Thomas Campbell, and William Shadish. *Experimental and quasi-experimental designs for generalized causal inference*. Vol. 1195. Houghton Mifflin Boston, MA, 2002.
- [85] Olawale Salaudeen et al. *Measurement to Meaning: A Validity-Centered Framework for AI Evaluation*. arXiv:2505.10573 [cs]. June 2025. DOI: 10.48550/arXiv.2505.10573. URL: <http://arxiv.org/abs/2505.10573> (visited on 08/21/2025).
- [86] Alexandra Chouldechova et al. *A Shared Standard for Valid Measurement of Generative AI Systems' Capabilities, Risks, and Impacts*. 2024. arXiv: 2412.01934 [cs.CY]. URL: <https://arxiv.org/abs/2412.01934>.
- [87] Rebecca B Morton and Kenneth C Williams. *Experimental political science and the study of causality: From nature to the lab*. Cambridge University Press, 2010.
- [88] Joshua D Angrist and Jörn-Steffen Pischke. “The credibility revolution in empirical economics: How better research design is taking the con out of econometrics”. In: *Journal of economic perspectives* 24.2 (2010), pp. 3–30.
- [89] Bruce D. Meyer. “Natural and Quasi-Experiments in Economics”. In: *Journal of Business & Economic Statistics* 13.2 (1995), pp. 151–161. DOI: 10.1080/07350015.1995.10524589.
- [90] Angelina Wang, Teresa Datta, and John P. Dickerson. “Strategies for Increasing Corporate Responsible AI Prioritization”. In: *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*. AAAI Press, Feb. 2025, pp. 1514–1526. (Visited on 08/19/2025).
- [91] Negar Rostamzadeh et al. “Healthsheet: Development of a Transparency Artifact for Health Datasets”. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '22. New York, NY, USA: Association for Computing Machinery, June 2022, pp. 1943–1961. ISBN: 978-1-4503-9352-2. DOI: 10.1145/3531146.3533239. URL: <https://dl.acm.org/doi/10.1145/3531146.3533239> (visited on 08/21/2025).
- [92] Nikita Mehandru, Samantha Robertson, and Nilufar Salehi. “Reliable and Safe Use of Machine Translation in Medical Settings”. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '22. New York, NY, USA: Association for Computing Machinery, June 2022, pp. 2016–2025. ISBN: 978-1-4503-9352-2. DOI: 10.1145/3531146.3533244. URL: <https://dl.acm.org/doi/10.1145/3531146.3533244> (visited on 01/08/2026).
- [93] Inyoung Cheong et al. “(A)I Am Not a Lawyer, But...: Engaging Legal Experts towards Responsible LLM Policies for Legal Advice”. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. New York, NY, USA: Association for Computing Machinery, June 2024, pp. 2454–2469. ISBN: 979-8-4007-0450-5. DOI: 10.1145/3630106.3659048. URL: <https://dl.acm.org/doi/10.1145/3630106.3659048> (visited on 01/08/2026).
- [94] Sanne Vrijenhoek et al. “Diversity of What? On the Different Conceptualizations of Diversity in Recommender Systems”. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. Rio de Janeiro, Brazil: Association for Computing Machinery, 2024, pp. 573–584. ISBN: 9798400704505. DOI: 10.1145/3630106.3658926. URL: <https://doi.org/10.1145/3630106.3658926>.
- [95] Lara Groves et al. “Auditing Work: Exploring the New York City algorithmic bias audit regime”. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. Rio de Janeiro, Brazil: Association for Computing Machinery, 2024, pp. 1107–1120. ISBN: 9798400704505. DOI: 10.1145/3630106.3658959. URL: <https://doi.org/10.1145/3630106.3658959>.
- [96] Wesley Hanwen Deng et al. “Exploring How Machine Learning Practitioners (Try To) Use Fairness Toolkits”. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '22. Seoul, Republic of Korea: Association for Computing Machinery, 2022, pp. 473–484. ISBN: 9781450393522. DOI: 10.1145/3531146.3533113. URL: <https://doi.org/10.1145/3531146.3533113>.
- [97] Michael Madaio et al. “Learning about Responsible AI On-The-Job: Learning Pathways, Orientations, and Aspirations”. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. Rio de Janeiro, Brazil: Association for Computing Machinery, 2024, pp. 1544–1558. ISBN: 9798400704505. DOI: 10.1145/3630106.3658988. URL: <https://doi.org/10.1145/3630106.3658988>.
- [98] Nitin Agrawal et al. “Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–13. ISBN: 978-1-4503-8096-6. DOI: 10.1145/3411764.3445677. URL: <https://dl.acm.org/doi/10.1145/3411764.3445677> (visited on 07/31/2025).
- [99] Hongbo Zhang et al. “IEDS: Exploring an Intelligence Embodied Design Space Combining Designer, AR, and GAI to Support Industrial Conceptual Design”. In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. CHI '25. New York, NY, USA: Association for Computing Machinery, 2025.

- tion for Computing Machinery, Apr. 2025, pp. 1–25. ISBN: 979-8-4007-1394-1. DOI: 10.1145/3706598.3713528. URL: <https://dl.acm.org/doi/10.1145/3706598.3713528> (visited on 07/31/2025).
- [100] Marie-Therese Sekwenz, Ben Wagner, and Simon Parkin. “It is unfair, and it would be unwise to expect the user to know the law!” – Evaluating reporting mechanisms under the Digital Services Act”. In: *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’25. New York, NY, USA: Association for Computing Machinery, June 2025, pp. 532–546. ISBN: 979-8-4007-1482-5. DOI: 10.1145/3715275.3732036. URL: <https://dl.acm.org/doi/10.1145/3715275.3732036> (visited on 07/31/2025).
- [101] Oğuz ’Oz’ Buruk, Louise Petersen Matjeka, and Florian ’Floyd’ Mueller. “Towards Designing Playful Bodily Extensions: Learning from Expert Interviews”. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI ’23. New York, NY, USA: Association for Computing Machinery, Apr. 2023, pp. 1–20. ISBN: 978-1-4503-9421-5. DOI: 10.1145/3544548.3581165. URL: <https://dl.acm.org/doi/10.1145/3544548.3581165> (visited on 07/31/2025).
- [102] Monique Hennink and Bonnie N. Kaiser. “Sample sizes for saturation in qualitative research: A systematic review of empirical tests”. In: *Social Science & Medicine* 292 (Jan. 2022), p. 114523. ISSN: 0277-9536. DOI: 10.1016/j.socscimed.2021.114523. URL: <https://www.sciencedirect.com/science/article/pii/S0277953621008558> (visited on 01/28/2024).
- [103] Greg Guest, Arwen Bunce, and Laura Johnson. “How Many Interviews Are Enough?: An Experiment with Data Saturation and Variability”. en. In: *Field Methods* 18.1 (Feb. 2006). Publisher: SAGE Publications Inc, pp. 59–82. ISSN: 1525-822X. DOI: 10.1177/1525822X05279903. URL: <https://doi.org/10.1177/1525822X05279903> (visited on 01/28/2024).
- [104] Benjamin Harrap et al. “A randomised controlled trial of email versus mailed invitation letter in a national longitudinal survey of physicians”. In: *PLOS ONE* 18.8 (Aug. 2023), e0289628. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0289628. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10443851/> (visited on 08/22/2025).
- [105] Kevin L. Wei et al. *Recommendations and Reporting Checklist for Rigorous & Transparent Human Baselines in Model Evaluations*. 2025. arXiv: 2506.13776 [cs.AI]. URL: <https://arxiv.org/abs/2506.13776>.
- [106] Inioluwa Deborah Raji et al. *AI and the Everything in the Whole Wide World Benchmark*. 2021. arXiv: 2111.15366 [cs.LG]. URL: <https://arxiv.org/abs/2111.15366>.
- [107] Ben Hutchinson et al. *Evaluation Gaps in Machine Learning Practice*. 2022. arXiv: 2205.05256 [cs.LG]. URL: <https://arxiv.org/abs/2205.05256>.
- [108] Maribeth Rauh et al. “Gaps in the Safety Evaluation of Generative AI.” In: *AIES (I)*. Ed. by Sammay Das et al. AAAI Press, 2024, pp. 1200–1217. URL: <http://dblp.uni-trier.de/db/conf/aies/aies2024-1.html#RauhMMHCASMBKGB24>.
- [109] Maria Eriksson et al. *Can We Trust AI Benchmarks? An Interdisciplinary Review of Current Issues in AI Evaluation*. 2025. arXiv: 2502.06559 [cs.AI]. URL: <https://arxiv.org/abs/2502.06559>.
- [110] Hanna Wallach et al. *Position: Evaluating Generative AI Systems Is a Social Science Measurement Challenge*. arXiv:2502.00561 [cs]. June 2025. DOI: 10.48550/arXiv.2502.00561. URL: <http://arxiv.org/abs/2502.00561> (visited on 08/21/2025).
- [111] American Statistical Association. “American Statistical Association Releases Statement on Statistical Significance and P-Values”. en. In: (Mar. 2016).
- [112] Sayash Kapoor et al. *Holistic Agent Leaderboard: The Missing Infrastructure for AI Agent Evaluation*. 2025. arXiv: 2510.11977 [cs.AI]. URL: <https://arxiv.org/abs/2510.11977>.
- [113] Xiaoyuan Zhu et al. *Auditing Black-Box LLM APIs with a Rank-Based Uniformity Test*. 2026. arXiv: 2506.06975 [cs.CR]. URL: <https://arxiv.org/abs/2506.06975>.
- [114] Olivier Klein et al. “Low Hopes, High Expectations: Expectancy Effects and the Replicability of Behavioral Experiments”. In: *Perspectives on Psychological Science* 7.6 (2012). PMID: 26168114, pp. 572–584. DOI: 10.1177/1745691612463704. URL: <https://doi.org/10.1177/1745691612463704>.
- [115] Sally Hopewell et al. “CONSORT 2025 statement: updated guideline for reporting randomized trials”. en. In: *Nature Medicine* 31.6 (June 2025). Publisher: Nature Publishing Group, pp. 1776–1783. ISSN: 1546-170X. DOI: 10.1038/s41591-025-03635-5. URL: <https://www.nature.com/articles/s41591-025-03635-5> (visited on 09/05/2025).
- [116] American Economic Association. *About the AEA RCT Registry*. 2025. URL: <https://www.aeaweb.org/journals/policies/rct-registry> (visited on 09/06/2025).
- [117] Dan Hendrycks et al. *Measuring Massive Multi-task Language Understanding*. 2021. arXiv: 2009.03300 [cs.CY]. URL: <https://arxiv.org/abs/2009.03300>.

- [118] David Rein et al. *GPQA: A Graduate-Level Google-Proof Q&A Benchmark*. 2023. arXiv: 2311.12022 [cs.AI]. URL: <https://arxiv.org/abs/2311.12022>.
- [119] Sayash Kapoor et al. *On the Societal Impact of Open Foundation Models*. arXiv:2403.07918 [cs] version: 1. Feb. 2024. DOI: 10.48550/arXiv.2403.07918. URL: <http://arxiv.org/abs/2403.07918> (visited on 09/06/2025).
- [120] Stephen Casper et al. “Black-Box Access is Insufficient for Rigorous AI Audits”. In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT ’24. ACM, June 2024, pp. 2254–2272. DOI: 10.1145/3630106.3659037. URL: <http://dx.doi.org/10.1145/3630106.3659037>.
- [121] Miles Brundage et al. *Frontier AI Auditing: Toward Rigorous Third-Party Assessment of Safety and Security Practices at Leading AI Companies*. Working paper. Jan. 2026.
- [122] Kar Balan, Robert Learney, and Tim Wood. “A Framework for Cryptographic Verifiability of End-to-End AI Pipelines”. In: *Proceedings of the 10th ACM International Workshop on Security and Privacy Analytics*. IWSPA ’25. Pittsburgh, PA, USA: Association for Computing Machinery, 2025, pp. 49–59. ISBN: 9798400715013. DOI: 10.1145/3716815.3729011. URL: <https://doi.org/10.1145/3716815.3729011>.
- [123] Filippo Scaramuzza et al. “Show Me You Comply... Without Showing Me Anything”: Zero-Knowledge Software Auditing for AI-Enabled Systems. 2026. arXiv: 2510.26576 [cs.SE]. URL: <https://arxiv.org/abs/2510.26576>.
- [124] Filippo Scaramuzza, Giovanni Quattrocchi, and Damian A. Tamburri. *Engineering Trustworthy Machine-Learning Operations with Zero-Knowledge Proofs*. 2025. arXiv: 2505.20136 [cs.SE]. URL: <https://arxiv.org/abs/2505.20136>.
- [125] Kai Yao and Marc Juarez. *SPRINT: Robust Model Attribution of Generated Images via Secret Pixel Reconstruction*. 2026. arXiv: 2508.05691 [cs.CR]. URL: <https://arxiv.org/abs/2508.05691>.
- [126] Siqi You. *Verifiable Audit Mechanisms in AI Compliance Automation: Scalability*. SSRN working paper. Available at SSRN: <https://ssrn.com/abstract=6547458>. Mar. 2026.
- [127] Yangjun Ruan et al. *Identifying the Risks of LM Agents with an LM-Emulated Sandbox*. 2024. arXiv: 2309.15817 [cs.AI]. URL: <https://arxiv.org/abs/2309.15817>.
- [128] Jacy Reese Anthis et al. *LLM Social Simulations Are a Promising Research Method*. 2025. arXiv: 2504.02234 [cs.HC]. URL: <https://arxiv.org/abs/2504.02234>.
- [129] Jiarui Lu et al. “ToolSandbox: A Stateful, Conversational, Interactive Evaluation Benchmark for LLM Tool Use Capabilities”. In: *Findings of the Association for Computational Linguistics: NAACL 2025*. Ed. by Luis Chiruzzo, Alan Ritter, and Lu Wang. Albuquerque, New Mexico: Association for Computational Linguistics, Apr. 2025, pp. 1160–1183. ISBN: 979-8-89176-195-7. DOI: 10.18653/v1/2025.findings-naacl.65. URL: <https://aclanthology.org/2025.findings-naacl.65/>.
- [130] Romain Froger et al. *ARE: Scaling Up Agent Environments and Evaluations*. 2025. arXiv: 2509.17158 [cs.AI]. URL: <https://arxiv.org/abs/2509.17158>.
- [131] Joon Sung Park et al. *LLM Agents Grounded in Self-Reports Enable General-Purpose Simulation of Individuals*. 2026. arXiv: 2411.10109 [cs.AI]. URL: <https://arxiv.org/abs/2411.10109>.
- [132] Frontier Model Forum. *FMF Announces First-Of-Its-Kind Information-Sharing Agreement*. en-US. Mar. 2025. URL: <https://www.frontiermodelforum.org/updates/fmf-announces-first-of-its-kind-information-sharing-agreement/> (visited on 09/06/2025).
- [133] Robert J. Lempert. *Decision Making under Deep Uncertainty and the Great Acceleration: The Role of Experts and Policy Analysis in a World in Transition*. Expert Insights PE-A3789-1. Santa Monica, CA: RAND Corporation, Apr. 2025. DOI: 10.7249/PEA3789-1. URL: <https://www.rand.org/pubs/perspectives/PEA3789-1.html>.
- [134] Vincent A. W. J. Marchau et al., eds. *Decision Making under Deep Uncertainty: From Theory to Practice*. Open access. Springer Cham, 2019. ISBN: 978-3-030-05252-2. DOI: 10.1007/978-3-030-05252-2.
- [135] Tammy C Hoffmann et al. “Better reporting of interventions: template for intervention description and replication (TIDieR) checklist and guide”. In: *BMJ* 348 (2014), g1687. DOI: 10.1136/bmj.g1687. URL: <https://www.bmj.com/content/348/bmj.g1687>.
- [136] OpenAI. *Introducing ChatGPT*. en-US. Nov. 2022. URL: <https://perma.cc/G9CV-F9SR> (visited on 08/21/2025).
- [137] Emma Wiles et al. “Artificial Intelligence in Labor Market Matching, Chapter 3: Using AI to Upskill Non-Technical Workers into Data Science: A Field Experiment”. en. Doctor of Philosophy in Management. Massachusetts Institute of Technology, May 2024. URL: <https://dspace.mit.edu/entities/publication/ebdb9c88-766b-4c31-a614-c10ea5ea6a18>.
- [138] Allen Nie et al. *The GPT Surprise: Offering Large Language Model Chat in a Massive Coding Class Reduced Engagement but Increased Adopters Exam Performances*. Apr. 2024. DOI:

- 10.48550/arXiv.2407.09975. URL: <http://arxiv.org/abs/2407.09975> (visited on 07/15/2025).
- [139] Jonathan H. Choi, Amy B. Monahan, and Daniel Schwarcz. “Lawyering in the Age of Artificial Intelligence”. eng. In: *Minnesota Law Review* 109.1 (2024), pp. 147–218. URL: <https://heinonline.org/HOL/P?h=hein.journals/mnlr109&i=147> (visited on 07/15/2025).
- [140] Antonio Roldan-Mones. “When GenAI increases inequality: evidence from a university debating competition”. en. In: (2024).
- [141] Stefan Küchemann et al. “Can ChatGPT support prospective teachers in physics task development?”. In: *Physical Review Physics Education Research* 19.2 (Sept. 2023). DOI: 10.1103/PhysRevPhysEducRes.19.020128. URL: <https://link.aps.org/doi/10.1103/PhysRevPhysEducRes.19.020128> (visited on 07/15/2025).
- [142] Daniel Schwarcz et al. *AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice*. en. SSRN Scholarly Paper. Rochester, NY, Mar. 2025. DOI: 10.2139/ssrn.5162111. URL: <https://papers.ssrn.com/abstract=5162111> (visited on 09/05/2025).
- [143] Salih Rakap. “Chatting with GPT: Enhancing Individualized Education Program Goal Development for Novice Special Education Teachers”. EN. In: *Journal of Special Education Technology* 39.3 (Sept. 2024). ISSN: 0162-6434. DOI: 10.1177/01626434231211295. URL: <https://doi.org/10.1177/01626434231211295> (visited on 07/19/2025).
- [144] Hamsa Bastani et al. “Generative AI without guardrails can harm learning: Evidence from high school mathematics”. In: *Proceedings of the National Academy of Sciences* 122.26 (July 2025). DOI: 10.1073/pnas.2422633122. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2422633122> (visited on 07/19/2025).
- [145] Stefan Feuerriegel et al. “Generative AI”. en. In: *Business & Information Systems Engineering* 66.1 (Feb. 2024), pp. 111–126. ISSN: 1867-0202. DOI: 10.1007/s12599-023-00834-7. URL: <https://doi.org/10.1007/s12599-023-00834-7> (visited on 08/21/2025).
- [146] Michael Gusenbauer. “Google Scholar to overshadow them all? Comparing the sizes of 12 academic search engines and bibliographic databases”. en. In: *Scientometrics* 118.1 (Jan. 2019), pp. 177–214. ISSN: 1588-2861. DOI: 10.1007/s11192-018-2958-5. URL: <https://doi.org/10.1007/s11192-018-2958-5> (visited on 01/26/2025).
- [147] Affan Yasin et al. “On Using Grey Literature and Google Scholar in Systematic Literature Reviews in Software Engineering”. In: *IEEE Access* 8 (2020), pp. 36226–36243. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.2971712. URL: <https://ieeexplore.ieee.org/document/8984351/> (visited on 01/26/2025).
- [148] Gali Halevi, Henk Moed, and Judit Bar-Ilan. “Suitability of Google Scholar as a source of scientific information and as a source of data for scientific evaluation—Review of the Literature”. In: *Journal of Informetrics* 11.3 (Aug. 2017), pp. 823–834. ISSN: 1751-1577. DOI: 10.1016/j.joi.2017.06.005. URL: <https://www.sciencedirect.com/science/article/pii/S1751157717300676> (visited on 01/26/2025).
- [149] Martin Boeker, Werner Vach, and Edith Motschall. “Google Scholar as replacement for systematic literature searches: good relative recall and precision are not enough”. en. In: *BMC Medical Research Methodology* 13.1 (Dec. 2013). Number: 1 Publisher: BioMed Central, pp. 1–12. ISSN: 1471-2288. DOI: 10.1186/1471-2288-13-131. URL: <https://bmcmmedresmethodol.biomedcentral.com/articles/10.1186/1471-2288-13-131> (visited on 01/26/2025).
- [150] Niroop Channa Rajashekar et al. “Human-Algorithmic Interaction Using a Large Language Model-Augmented Artificial Intelligence Clinical Decision Support System”. In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. CHI ’24. New York, NY, USA: Association for Computing Machinery, May 2024, pp. 1–20. ISBN: 979-8-4007-0330-0. DOI: 10.1145/3613904.3642024. URL: <https://dl.acm.org/doi/10.1145/3613904.3642024> (visited on 08/21/2025).

A Rapid Literature Review Methodology

To identify relevant articles, we queried Google Scholar for articles dated between January 1, 2023 and June 30, 2025 that contained any one human uplift keyword and any one LLM keyword from Table 5.⁷ We focus specifically on studies around LLMs given the nascent state of LLM human uplift studies, their unique challenges, and the gap of research related to their methodology. As there is no standardized terminology for human uplift studies involving LLMs, the keywords are chosen to be intentionally broad. Articles containing search terms indicating that they were literature reviews rather than experimental studies were excluded (see Table 5). This search resulted in a total of 106 unique articles,⁸ Author 5 then conducted initial filtering/annotation per the inclusion/exclusion criteria in Table 6, with review by Author 2, resulting in a final list of 10 studies meeting our criteria. Criteria ensured that studies were included in our final list of results if they 1) contained experimental results and 2) compared performance on some tasks by humans with vs. without access to LLM systems. The final 10 studies were: [137, 138, 139, 67, 140, 141, 142, 143, 144, 70].

Type	Category	Keywords
Inclusion	Human Uplift	“human uplift”, “randomized controlled trial”
Inclusion	LLM	“[LLM, large language model, ChatGPT, AI, artificial intelligence, AI model, artificial intelligence model, AI system, artificial intelligence system] access”, “access to [LLM, large language model, ChatGPT, AI, artificial intelligence, AI model, artificial intelligence model, AI system, artificial intelligence system]” [*]
Exclusion	Literature Review	“systematic review”, “systematic literature review”, “scoping review”, “scoping literature review”

* The notation “[LLM, large language model] access” means “LLM access,” “large language model access,” etc.

Table 5: Search terms for rapid literature review

Our method is subject to a few limitations. Most notably, our rapid review used Google Scholar as the search engine and used relatively restrictive search terms. Google Scholar was chosen as our search tool because most recent studies on this topic have found that it has extremely high coverage for scientific articles [146], with one study finding that Google Scholar had coverage of 96% of computer science articles indexed in other databases [147]. Google Scholar is also appropriate for our review because it indexes preprints and gray literature: nearly all of the literature related to human uplift with LLMs has been conducted within the past few years and is thus available only as preprints and/or on arXiv. We acknowledge that Google Scholar produces issues with the search interface, lack of precision, and reproducibility [148, 149]. To address interface limitations (the 256-character limit on search strings), we use multiple different search queries and combine the results.

Our LLM search terms were also limited in that they were simple string matches for specific phrases that we identified as commonly used in human uplift studies. These search terms necessarily have low recall and inadvertently exclude relevant studies (e.g., [150]) because even the “human uplift” terminology is fairly nascent in the LLM context. Our selection of these search terms was driven by the lack of standardized language describing the experimental conditions in human uplift studies; using broader search terms resulted in many thousands of results, making analysis impossible for a rapid review.⁹ Coverage in our rapid review may have thus been limited, and we interpret our results as only being broadly suggestive of broader methodological practices.

B Interview Methodology

Expert participants were selected using a snowball sampling method [75]; we defined experts as researchers who had conducted or were currently conducting human uplift studies involving LLMs that fulfilled the criteria in Table 6. We first created a seed set of experts, beginning with the first or corresponding authors of the 10 articles from our rapid review. We added to the seed experts from major AI developers whose recent system cards indicated that they had conducted a human uplift study (OpenAI, Anthropic, Google DeepMind, and Amazon). Based on their expertise

⁷The start date was chosen to roughly align with the public release of modern LLMs (e.g., ChatGPT’s public release on November 30, 2022[136]). Articles published after the end date range were not included in our rapid literature review, though we did include such articles as seeds for our interviews if they came to our attention and fulfilled relevant inclusion/exclusion criteria.

⁸The search queries initially resulted in 111 results. After initial annotation by Author 5, five articles were found to be duplicates or preprints of other articles, resulting in a total of 106 unique articles.

⁹A search for “‘randomized controlled trial’ AND ‘LLM’ ” yields close to 4,000 results, while replacing “LLM” with “AI” in the query yields nearly 28,000.

Inclusion Criteria	Exclusion Criteria	Details & Rationale
Study contained experimental results	Study did not contain experimental results	We included only primary literature with quantitative/empirical results, excluding literature reviews, pre-registrations, policy documents, etc. Our intention was to examine the methodological details of original experiments, as well as to discover experts (paper authors) with expertise in conducting human uplift studies for expert interviews.
	Non-experimental studies	We excluded observational studies in which experimental conditions were not controlled and participant performance was not evaluated directly by researchers (e.g., studies based on survey results or observational data). Our rationale is that these studies are significantly methodologically distinct from human uplift studies and that an important aspect of human uplift studies is the study and evaluation of human-LLM interaction in task performance (see below).
Experimental design attempted to compare task performance of humans with vs. without access to LLM systems	Experimental conditions did not compare humans with LLM system access vs. humans without LLM system access	We excluded human baselines (i.e., studies comparing human-only performance vs. AI-only performance) as out of scope because they are not interactive. Note, however that there are existing guidelines for conducting robust human baselines [105].
	The AI systems in the experiment were not LLMs	Although RCTs and human uplift studies have been conducted in the context of other technologies and non-LLM AI systems, we limited the scope of this paper to LLM systems. Our rationale is that LLM systems are significantly distinct from other AI systems (e.g., due to the diversity of LLM use cases, different architectures, different methods/modes of interaction, a heightened need for construct validity in LLM evaluation vs. in traditional ML contexts, etc. [85, 145, 10, 110, 105]) and pose significantly distinct challenges in the AI context (see Section 4.2).
	Human participants were not given direct access to LLM systems	We excluded experiments in which participants were given access only to static LLM outputs, rather than direct access to LLM systems. Our rationale is that direct interaction with LLM systems creates important causal pathways and threat vectors that are not present when presented with static outputs only [10].

Table 6: Inclusion/Exclusion criteria for filtering of rapid literature review articles

and awareness of relevant studies, Author 9 and Author 10 then provided a list of published human uplift studies involving LLMs, and the first or corresponding authors of these studies were added to the list. Author 9 and Author 10 are senior authors on this paper who have (combined) over a decade of experience conducting randomized controlled trials, including multiple human uplift studies in the LLM context.

Author 1 and Author 2 contacted the seed experts via email, inviting them to participate in a research interview and sending them an informational fact sheet about the project. Experts were asked at the end of each interview to identify one or two other individuals that they believed we should contact to participate in interviews. Author 1 and Author 2 then contacted the snowball sampled individuals and invited them to participate in a research interview. Experts were not financially compensated for participating in interviews.

Interviews were conducted on the encrypted ZoomGov platform¹⁰ by Authors 1, 2, 4, and 5, all of whom have had research experience with interview methods. Each interview consisted of one interviewer and one expert interviewee. All interviews were conducted in English. We recorded audio for transcription purposes only and with the consent of the experts.

1. Information and oral consent: interviewers gave experts an opportunity to re-review the fact sheet and collected oral consent to proceed with the interview and to record for transcription purposes.
2. Demographics: interviewers collected demographic information from the experts (summarized in Section Sample: Experts).
3. Study history: interviewers defined human uplift studies as “studies [that] measure the extent to which access to and/or use of a general-purpose AI model impacts human performance on a task, relative to a baseline (think comparing human performance with human and LLM performance). Human uplift studies often em-

¹⁰One expert could not access Zoom, and the interview was conducted instead on Google Meet with no recording (the interviewer took verbatim notes instead).

ploy randomized controlled trial design.” Experts were asked how many total human uplift studies they had conducted in the LLM context, then to answer questions about the methodological design of the (most recent) human uplift study that they had conducted. This module included questions such as the number of human participants in the study, how participants were recruited, the experimental conditions, the process for choosing and validating measurement instruments, quality control and compliance enforcement measures, etc.

4. Methodological challenges: interviewers asked the expert to identify methodological challenges encountered when designing or implementing the study. Experts were asked to focus on methodological challenges specific to the AI context. Once experts had identified a list of methodological challenges, interviewers asked follow-up questions about how each challenge could influence the validity and reliability of human uplift studies, how the experts and their research teams addressed each challenge, and what methodological options (and tradeoffs between different options) existed for addressing each challenge.
5. Open-ended questions: at the end of the interview, interviewers asked two open-ended questions to allow experts to identify any challenges that had not yet been discussed earlier in the interview. Interviewers also asked experts to identify other individuals or organizations whom they thought we should invite to participate in the study (snowball sampling).

The full interview script is provided in Appendix B.

Interviews were first transcribed using a privately hosted instance of the OpenAI Whisper model; these transcripts were then manually validated by Authors 1, 2, 3, 4, and 6. Our subsequent thematic analysis followed a reflexive, two-stage inductive coding approach, following the methods of [99, 80, 98, 81].

In the first stage of TA, Authors 1, 4, 5, and 6 familiarized themselves with the interviews before independently generating codes by reviewing two randomly selected interviews from our sample through a bottom-up, open coding process. Authors 1 and 2 then discussed these four sets of codes, consolidating and refining them into an initial set of codes while grouping individual codes into categories and sub-categories. Categories were informed by the stages of the AI evaluation lifecycle, defined in [24, 105].

In stage two of TA, Authors 3, 4, 5 used Dedoose¹¹ to annotate all interviews using the initial codes developed from the first stage. Each interview was independently annotated by two coders, and no coder analyzed an interview for which they were also the interviewer. Codes were iteratively adapted and refined throughout the coding process, with minor changes made as needed and discussed with Author 1 and Author 2.

Interview Script

Before we begin, could you confirm that you received and reviewed the fact sheet sent prior to this interview? If no, would you like another copy? It’s short and I can wait while you read it.

I would like to remind you that your participation is voluntary, and you should only participate if you fully understand the study requirements and risks. Do you have any questions related to participating in the study?

Do you consent to proceed?

Thanks for agreeing to be interviewed! Before we start I want to double check with you about a few things: 1) We are audio taping this interview and we will be preparing transcripts from the interviews. Is that OK with you? 2) We will be reporting themes and variation in responses across the interviews. We may include some direct quotes, but we will not include any quotes or attributions without your consent. You are free to decline to answer any question, to provide the level of detail you feel is appropriate in any response, or to respond “I don’t know” to any question. Does that make sense to you?

¹¹<https://www.dedoose.com>

Module: Demographics (D)**Estimated Time: 5m**

D	1	How would you like to be identified in our paper? Options are: Full identification with name, role, and affiliation Identification with role and affiliation but without name Identification by only a short descriptor such as “evaluations researcher at an AI company” No identification (we would identify you as “anonymous” or “anonymous expert” with expertise in human uplift studies) We’d also like to include demographics. These will be broad and non-identifying. You can decline to respond if you’d like.
D	2	What is your gender?
D	3	How many years of experience do you have? 0-5, 6-10, 11-15, 16+?
D	4	In which country is your research organization headquartered?
D	5	What is your highest achieved level of education? Bachelors, Masters, PhD, Other (which)?

Module: Study Count (SC)**Estimated Time: 1m***Instructions: if SC1==yes, proceed to Module 1. If SC1==no: skip to Module 2*

		In this interview, we’ll be asking you about human uplift studies in AI evaluations. Human uplift studies measure the extent to which access to and/or use of a general-purpose AI model impacts human performance on a task, relative to a baseline. Human uplift studies often employ randomised controlled trial design. I’m going to start by asking you a few quick questions about human uplift studies that you’ve been involved in. Have you been involved in running any human uplift studies in the past? How many? If multiple, let’s begin with the most recent, substantial, and completed study. We can come back to others, time permitting.
SC	1	Please try to keep your responses to the following questions brief, as these questions are just to give us some context on your background and what methods are currently being used in human uplift studies.

Module: Current State (CS)**Estimated Time: 15m***Instructions: conduct Module 1-3 consecutively for the first evaluation, then loop back through Modules 1-2 for remaining evaluations, time permitting.*

CS	1	Is it published? If not, do you intend to publish it? When?
CS	2	What was the domain/subject area?
CS	3	What was the make-up of the team that designed and ran the study (e.g., background, skillsets)?
CS	4	What did the study attempt to measure?
CS	5	What was the study design, and what were the control and experimental conditions?
CS	6	What measurement instruments were used in this study, and how were evaluation data/items created or chosen?
CS	7	How were measurement instruments tested or validated?
CS	8	How many human participants were there in this study?
CS	9	If there was drop-off between the beginning and end of the trial, please report both numbers.
CS	10	How did you recruit human participants for this study?
CS	11	What quality control measures were used to recruit participants or to ensure participant compliance in this study?
CS	12	What statistical methods were used to analyze this study?
CS	13	Over what period of time did the study measure outcomes?
CS	14	Do you intend to conduct follow-up studies with this sample?

Module: Methodological Challenges (MC)**Estimated Time: 25m***Instructions: use stages for conversational guidance rather than imposed structure.*

What were some of the methodological challenges that you encountered when designing or implementing this study?

For interviewees who have not conducted a study but otherwise have insight into uplift studies, what are some of the methodological challenges that researchers encounter when designing or implementing human uplift studies study?

1 Design: defining the evaluation's scope, purpose, structure, methodological design etc.

2 Implementation: selecting and constructing evaluation tools, recruiting human participants, etc.

3 Execution: data collection stage of the study, i.e., running the study itself with AI systems and with human participants

4 Analysis: conducting statistical analysis and/or interpreting results of the data

5 Documentation: recording and sharing evaluation results, metrics, methodological details, etc.

MC 1 Please try to focus on challenges specific to the AI context. For concreteness, feel free to discuss challenges at any stage of the evaluation process.

For each challenge identified above:*Instructions: loop through MC for *each* methodological challenge identified*

MC 2 How does this challenge limit human uplift studies? In other words, how could it have influenced the validity or reliability of study results?

MC 3 Is there anything about this challenge that makes it uniquely challenging in the AI context? In other words, why don't RCTs or human uplift trials in other contexts face this challenge?

MC 4 What are ways that you or others have tried to deal with this challenge? What are the different options here, and what are the tradeoffs?

Module: Final questions (FQ)**Estimated Time: 5m***Instructions: if time remains after this module and the interviewee has conducted multiple AI human uplift studies, loop back to Modules 1-2 for each subsequent study, time permitting.*

FQ 1 If you could wave a wand and solve a problem more broadly in the field, what would you want solved?

FQ 2 Is there anything else that we haven't discussed that you would like to raise?

FQ 3 Are there any other researchers or organizations that you suggest we reach out to for this study? As a reminder, we are focused on interviewing researchers who have conducted human uplift studies/RCTs in the context of AI.

Thanks for participating in this interview! Just to re-confirm after our interview, we will give you a chance to review any quotes, attributions, identification, etc. before we release this paper. We will not include any quotes, attributions, or identification without your permission. Do you have any other questions for us about this study or about your participation in this study? If you have any other questions or comments, you can contact the research team.

C Thematic Analysis Codes

Code Category	Description
Design & Implementation	
Defining the control conditions	Defining the control condition to which experimental conditions are compared. Control conditions may include, e.g., internet search, 2023 level AI, use of expert human hotlines, etc.
Defining the experimental condition(s)	Defining the treatment/experimental condition(s), including which LLM(s) to evaluate. Factors considered may include the recency of models, elicitation of models, the pace of model deployment, interaction of models in real-world, and causal mechanisms or threat models.
Defining the population of interest	Specifying the subset of humans to which the study results are meant to generalize. Populations could be defined by demographic, AI skill/literacy, level of motivation, expertise, intent, etc.
Defining the research question and causal mechanism for impact	Hypothesizing pathways by which the AI may affect human outcomes (e.g., suggestions, explanations, speed). Researchers must test mechanisms or pathways through which access to LLMs could result in risks/benefits. Even for evaluations that are measuring LLM impacts generally, the research design is often making implicit choices to prioritize measurement of particular action spaces.
Designing the test environment	Deciding about the test environment, including physical spaces (e.g. lab, field, etc.) and study length.
Developer involvement/collaboration	Interacting with developers in the design stage, including regarding contracts and model access.
Domain- or methodology-specific experimental design considerations	Experimental design considerations (e.g. using a within-participant, crossover, or other more complex design) to improve measurement in a specific domain or to make a certain statistical methodology possible.
Measurement - proxy/dependent variable choice	Designing and piloting measurement instruments such that they act as a valid proxy for the real-world outcome researchers seek to measure.
Measurement - task specification	Choosing the task to be measured. May include discussions of depth-breadth tradeoff, e.g. tradeoff of pre-specifying pathways of a threat model to better measure them vs. allowing for wider decision space (with poorer measurement).
Preventing spillovers/contamination (quality control)	Controlling for contamination, spillovers, and cheating both <i>in</i> (e.g. in the lab during study hours) and <i>out</i> (e.g. out of the lab during non-study hours) of the direct study environment and both within and between treatment/control groups.
Uplift vs. other evaluation methods	Deciding which evaluation method to use in order to measure the concept of interest, whether uplift or another method (e.g., benchmarks). Considerations may include factors such as cost, ecological validity, etc.
Recruitment	
Participant incentives	Choosing incentive or disincentive structures to motivate participants to participate in the study, comply with study protocols, and model real-world motivations. May include financial or reputational incentives, social dynamics, research monitoring, etc.
Participant recruitment	Recruiting motivated participants that proxy the real-world population of interest. Considerations may include baseline AI skill level (both in terms of using LLMs and interpreting LLM outputs), domain expertise, and demographics.
Participant training	Training participants on relevant skillsets, including but not limited to using LLMs and interpreting LLM outputs.

Table 7: Human Uplift Study Coding Categories and Descriptions (continued on next page)

Code Category	Description
Execution	
Monitoring access and use of LLM	Tracking the extent to and ways in which participants use LLMs. May include tracking message/user logs and/or number of messages, tracking time spent using LLMs, recording interactions, collecting self-reported metrics, etc.
Monitoring spillovers/contamination	Gathering data on the occurrence and degree of spillovers, contamination, or cheating to act on and/or account for in analysis.
Retention of sample	Retaining a [representative] sample, especially in the case of longitudinal studies. Synonymous with preventing attrition or drop-off.
Use of ethical protocols	Ensuring ethical treatment and safety of participants and society, including but not limited to in security-related evaluations. Includes mitigation of info hazards, bio hazards, etc.
Analysis	
Controls - participant characteristics	Controlling and/or accounting for participant heterogeneity – including AI skill level – in analysis.
Controls - spillovers/contamination	Controlling for contamination, spillovers, cheating etc. in analysis.
Grading results/analyzing data (e.g. user logs)	Effectively making use of and analyzing large amounts of data that may be unstructured or inconsistently structured (e.g. user logs) to understand causal pathways.
Statistical methods	Choosing and implementing statistical methods used to analyze data and results. Includes considerations around how to aggregate metrics to answer the research question (e.g. mean v top percentile).
Documentation	
Choosing baselines and thresholds against which to compare results	Deciding on and operationalizing success/failure thresholds and comparator choices (human baseline, heuristic).
Interpreting results and rapid model progress	Interpreting and applying results amidst rapidly changing/deployed models and transfer/generalizability concerns.
Sharing of [sensitive] methodology, data, and/or results	Deciding about and acting on publication, disclosure, and reporting on methodology, data, and/or results.
Problem/Solution Classification	
Problem v solution	Statement pertains to a problem or solution.
Problem	Statement pertains to a problem.
Solution	Statement pertains to a current or potential solution.
Current solutions	Statement discusses solutions that are currently – or have been – used.
Potential/ideal solutions	Statement discusses solutions that <i>ought to be</i> used.
AI Specificity	
AI Specificity: High	Statements indicating a problem is highly specific to AI.
AI Specificity: Low	Statements indicating a problem is general to RCTs, uplift studies, etc., in that particular problem domain – not to the AI context specifically.
Solution Characteristics	
Long Term Research Questions	Statements indicating a problem will require ongoing effort.
Perceived insolvability	Statements indicating a problem is insolvable/intractable.
Perceived solvability	Statements indicating problem is solvable/tractable.
Effective solution	Effective solution.
Ineffective solution	Ineffective solution.
High certainty of solution effectiveness	High certainty of solution effectiveness.
Low certainty of solution effectiveness	Low certainty of solution effectiveness.

Table 7: Human Uplift Study Coding Categories and Descriptions (continued)