

The Judgment-Consequence Gap: LLM Moral Reasoning in Healthcare Decisions

Hadi Hosseini, Samarth Khanna*, Leona Pierce

Pennsylvania State University
{hadi,samarth.khanna,ijp5139}@psu.edu

Abstract

As large language models (LLMs) enter high-stakes domains such as healthcare, understanding their moral reasoning becomes essential. Decisions about scarce medical resources often hinge on judgments of responsibility, particularly when patients' own actions contribute to illness. We investigate how LLMs reason about responsibility and its consequences, using *attribution theory* to trace their judgments from behavior to illness to denial of care. We evaluate a wide range of LLMs, spanning different model families and capability levels, on various clinical vignettes adapted from prior studies. Our results identify a *judgment-consequence gap*: LLMs largely agree with humans that patients bear responsibility for health-harming behaviors, yet overwhelmingly refuse to let that judgment influence how they allocate scarce resources. Specifically, LLMs default to random allocation, whereas humans consistently favor the less-culpable patient. Compared to humans, LLMs also place greater emphasis on access to information, reducing responsibility judgments when health-risk knowledge is unavailable. These findings reveal that LLMs apply a systematically different moral framework than humans when responsibility and resource scarcity intersect, surprisingly often amplifying normative disagreement with humans as reasoning capability increases.

1 Introduction

Large language models (LLMs) have rapidly moved beyond generating text to actively shaping decisions across virtually every sector of the economy, from writing code and drafting legal documents to advising on medical diagnoses and acting autonomously on behalf of users (Thirunavukarasu et al. 2023; Liu et al. 2025; Singhal et al. 2023). In healthcare, this transformation is especially pronounced. LLMs now assist with clinical documentation, diagnostic reasoning, and treatment recommendation, and a growing number of patients consult them directly for medical guidance (Vrdoljak et al. 2025; Shool et al. 2025; Ayers et al. 2023). As these models take on advisory and decision-making roles in clinical contexts, the values they encode become a matter of direct consequence for patient outcomes (Yu et al. 2024; McCraden et al. 2023).

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Some of the most difficult decisions in medicine, however, go beyond identifying the best treatment for an individual patient. When medical resources are scarce, such as a donated organ, hospital capacity, or an expensive treatment that not all eligible patients can receive, clinicians must decide how to allocate them (Persad, Wertheimer, and Emanuel 2009; Emanuel et al. 2020). These allocation decisions are inherently moral, forcing a link between *judgment* and *consequence*. When a patient's own behavior has contributed to their illness, a decision must be made about whether that behavior should affect who receives the scarce resource (Chan et al. 2024; Björk, Lynøe, and Juth 2015). Decades of research in attribution theory show that humans reliably connect perceived '*responsibility*' to such downstream consequences, with those judged responsible for their misfortune being seen as less deserving of help (Weiner 1985; Lerner 1980). We do not necessarily advocate for LLMs to make such decisions, but given that they are already embedded in healthcare workflows where these trade-offs arise, auditing how they connect responsibility to consequence is essential.

Recent work has begun to investigate how LLM moral reasoning compares to that of humans, through moral reasoning benchmarks (Hendrycks et al. 2021; Scherrer et al. 2023), large-scale preference elicitation inspired by the Moral Machine experiment (Awad et al. 2018; Takemoto 2024; bin Ahmad and Takemoto 2024), and direct evaluation of LLM behavior in allocation scenarios (Dickerson et al. 2025; Hosseini and Khanna 2026). These studies show that while LLMs demonstrate broad familiarity with moral norms, their judgments often diverge from human expectations on key decisions (Almeida et al. 2024; Shen, Clark, and Mitra 2025). However, prior work has largely treated moral reasoning as a single-step evaluation. What remains unexplored is *where in the reasoning chain* the divergence occurs. Do LLMs differ from humans in how they assess moral responsibility, in how they translate that assessment into a consequential decision, or in both?

We address this gap by asking: **when a patient's own behavior contributes to their need for a scarce medical resource, do LLMs connect judgments of moral responsibility to allocation decisions the way humans do?** A related question concerns the role of the patient's epistemic state: **does access to knowledge about health risks modulate responsibility judgments and allocation decisions**

differently for LLMs than for humans? To answer these questions, we test a broad set of LLMs spanning multiple model families, reasoning and non-reasoning configurations, and both open-source and commercial models, on clinical vignettes adapted from human studies of kidney transplant allocation (Chan et al. 2024), lung cancer treatment (Björk, Lynøe, and Juth 2015), and hip replacement surgery (Björk, Juth, and Lynøe 2018).

Our experimental design traces successive stages of the moral reasoning process, from whether the patient is responsible for the harmful behavior, to whether they are responsible for the resulting illness, to whether they should bear the cost of being denied a scarce resource, enabling us to identify precisely where human and LLM behavior diverge.

Results. We find a systematic **judgment-consequence gap**: LLMs largely agree with humans that patients who engage in health-harming behaviors bear moral responsibility, yet they overwhelmingly refuse to let those judgments influence how they allocate scarce resources. Where a majority of human participants allocate the resource to the less-culpable patient, LLMs default to random allocation and actively endorse the position that behavior-based allocation would be unfair. Beyond this central finding, our analysis reveals:

- **Knowledge sensitivity.** LLMs are universally more sensitive than humans to whether the patient had access to information about health risks, sharply reducing attributed responsibility when the patient was uninformed, a pattern consistent with a principled informed-consent framework that diverges from how people actually reason.
- **Reasoning widens the gap.** Enabling extended thinking increases responsibility attribution for most model families, bringing LLMs closer to human assessments, while (if anything) pushing allocation decisions further toward randomization rather than toward the human pattern. The gap widens with deeper reasoning, suggesting it reflects a stable normative commitment rather than a processing limitation.
- **Robustness across medical contexts.** These patterns hold across multiple medical domains, behavior types, and model configurations, indicating that the divergence is a general property of current LLMs rather than an artifact of any particular model or scenario.

2 Related Works

2.1 Responsibility Attribution

Responsibility attribution examines how humans judge an agent's involvement in a moral outcome. Much of the literature builds on Weiner's attributional model (Weiner 1985), which links perceived controllability of outcomes to emotional and judgmental responses. A subsequent meta-analysis provides robust empirical support for this framework (Rudolph et al. 2004). Attribution research spans legal defenses (Darley, Klosson, and Zanna 1978), punishment (Shultz, Wright, and Schleifer 1986), and the moralization of medical stigmas (Weiner, Perry, and Magnusson 1988). Within healthcare specifically, work has explored

how responsibility extends beyond behavior itself to the consequences of that behavior, including the prioritization or deprivation of treatment (Chan et al. 2024; Björk, Lynøe, and Juth 2015). Our work extends these frameworks by examining how LLMs attribute responsibility for both actions and consequences, and the degree to which these attributions align with human judgments.

2.2 Moral Decision Making with LLMs

A growing body of work examines the degree to which LLMs align with human moral judgments. Several benchmarks measure this alignment directly. ETHICS (Hendrycks et al. 2021), MoralChoice (Scherrer et al. 2023), and MoralExceptQA (Jin et al. 2022) evaluate moral reasoning through curated dilemmas, while the Delphi project trains a dedicated model on crowd-sourced moral judgments (Jiang et al. 2025). Other work probes alignment on contentious social issues (Santurkar et al. 2023; Garcia, Qian, and Palminteri 2024) and classic trolley-style problems (Ding et al. 2025), including the Moral Machine experiment (Takemoto 2024; bin Ahmad and Takemoto 2024).

Beyond benchmarking, research has investigated the moral values and ethical theories implicitly encoded in LLMs (Huang et al. 2026; Shen et al. 2025; Zhou et al. 2024a), as well as alignment with principles of fairness in resource allocation (Dickerson et al. 2025; Hosseini and Khanna 2026; Cookson, Ebadian, and Shah 2026). Studies have also examined how LLM decisions shift depending on the demographic persona they are asked to represent (Sorin et al. 2025), and have identified a value-action gap in which stated moral positions diverge from the decisions models ultimately make in practice (Shen, Clark, and Mitra 2025; Hosseini and Khanna 2026). A parallel line of work demonstrates that moral advice generated by LLMs on real-world dilemmas is perceived by humans as superior to that of professional ethicists (Aharoni et al. 2024; Howe et al. 2023; Dillion et al. 2025).

We contribute to this literature by tracing different stages of the moral reasoning process in a consequential decision-making scenario, from judgments of behavioral responsibility through disease attribution to allocation decisions, providing further insight into the specific areas of alignment and misalignment between humans and LLMs.

2.3 LLMs for Healthcare Decision Support

LLMs are increasingly explored for healthcare tasks ranging from clinical documentation and diagnostic reasoning to treatment recommendation (Liu et al. 2025; Xiao et al. 2025). Evaluations have benchmarked performance across the patient journey (Wu et al. 2024), explored multi-agent collaboration (Kim et al. 2024), and developed dynamic, iterative benchmarks that better reflect real-world clinical requirements (Li et al. 2024). Recent work has also identified failure modes such as inflexible reasoning under incomplete information (Lim et al. 2025). Beyond clinical accuracy, a parallel literature addresses normative questions about responsibility, trust, physician-AI disagreement, and the values embedded in medical AI systems (Kempt and Nagel 2022; Kempt, Heilinger, and Nagel 2023; Yu et al. 2024;

McCradden et al. 2023), arguing that these systems are not value-neutral but encode human and institutional judgments through their design and deployment. Comparatively little attention has been paid to how LLMs behave in *normatively charged* medical decisions involving ethical trade-offs rather than factual uncertainty. Our work addresses this gap by empirically comparing LLM moral judgments with human responses in a high-stakes allocation setting.

2.4 LLMs for Simulating Human Subjects

LLMs are increasingly used to simulate human subjects, whether as population-level proxies or as generators of synthetic survey responses. Multi-agent simulations have recreated election outcomes (Zhou et al. 2025), social interaction patterns (Ji et al. 2026; Zhou et al. 2024b; Wang et al. 2024), and trust dynamics (Jia et al. 2024; Guan et al. 2025). To improve simulation fidelity, researchers have explored persona-based prompting (Tseng et al. 2024; Newsham and Prince 2025), grounding in interview transcripts (Park et al. 2024), and finetuning on large-scale survey data (Suh et al. 2025; Kolluri et al. 2025). Foundation models trained directly on human behavioral data have shown promise in predicting human choices across diverse experimental paradigms (Binz et al. 2025), and dedicated benchmarks now evaluate simulation quality systematically (Hu et al. 2026). The Delphi project applies a similar approach to moral judgments specifically (Jiang et al. 2025). However, various concerns persist, such as minor prompt variations substantially altering outputs (Schröder et al. 2025), LLMs' tendency to compress variance and amplify majority effects relative to humans (Almeida et al. 2024; Dickerson et al. 2025), and deeper epistemological questions arising about whether LLMs can meaningfully stand in for human respondents (Kapania et al. 2025; Anthi et al. 2025). We demonstrate how LLMs can reliably simulate human responses in certain stages of the thought process while systematically deviating on consequential decisions.

2.5 Medical Resource Allocation

The allocation of scarce medical resources has been extensively studied across bioethics, health economics, and psychology. Normative frameworks emphasize principles such as maximizing benefit, equal treatment, and prioritizing the worst off (Emanuel et al. 2020; Persad, Wertheimer, and Emanuel 2009), while empirical research documents how laypeople and professionals actually prioritize patients based on age, health status, and social roles (Furnham, Simmons, and McClelland 2000; Furnham, Thomson, and McClelland 2002; Krütli et al. 2016; Chan et al. 2022). In kidney allocation specifically, studies aggregating human preferences into automated systems highlight substantial diversity and instability in individual responses (Freedman et al. 2020; McElfresh et al. 2021; Keswani et al. 2026). Our work extends these inquiries by testing whether the same patterns of disagreement and indecision appear in LLM responses, or whether models converge on systematically different allocation strategies.

3 Experimental Overview

3.1 Vignette Design and Measures

Our experiments are built around clinical vignettes, i.e. scenarios, in which two patients require the same scarce medical treatment. Patient A has no history of health-harming behavior. Patient B has engaged in a behavior, such as heavy drinking, drug use, smoking, or poor diet, that may have contributed to their condition. In each scenario, Patient A is given the treatment “on the basis of Patient B’s negative attribute”. After reading the scenario, participants answer a series of Likert scale questions that probe successive levels of moral responsibility. The first level concerns **behavior**: is Patient B responsible for engaging in the harmful behavior? The second concerns **disease**: is Patient B responsible for developing the illness? The third concerns **deprivation**: is Patient B responsible for being denied the scarce resource? Participants also make an explicit **allocation decision**: given the scarcity of resources only *one* patient can receive the treatment, should it go to Patient A, to Patient B, or should it be decided randomly? Directly related to this question, we also ask whether it is **fair** to allocate the kidney to Patient A based on Patient B’s behavior.

This layered structure allows us to trace the full chain from moral judgment to consequential action. The allocation decision is central to our analysis, revealing whether attributed responsibility actually translates into differential treatment when resources are scarce.

3.2 Kidney Transplant Allocation

We adapt the kidney transplant scenarios of (Chan et al. 2024), asking each model the same responsibility and allocation questions posed to human participants in the original study. This allows a direct comparison with a human baseline across all measures.

The scenarios take two forms. In the first, Patient B has either *stopped* or *continued* the harmful behavior after diagnosis, testing whether behavioral change is treated as morally relevant. Four distinct behaviors (alcohol, drugs, smoking, and poor diet) are crossed with the stopped/continued manipulation, yielding eight conditions. In the second, the behavior is held constant while Patient B’s *access to knowledge* about the health risks varies across three levels, (i) the patient knew about the risks, (ii) the patient did not know but had easy access to information, or (iii) the patient neither knew nor had access. We refer to these two designs throughout as the *behavior alteration* vignette and the *knowledge level* vignette, respectively.

Additionally, we introduce three conditions designed to test the effect of misleading information on responsibility attribution. In these conditions, the patient was actively given *incorrect information* about the health risks, either by a third party, by another person, or by an AI agent, allowing us to also measure any influence of the source of the incorrect advice.

3.3 Lung Cancer and Hip Replacement

To assess whether the patterns observed in the kidney domain generalize across medical contexts, we administer both

the behavior alteration and knowledge level vignettes in two additional medical settings, (i) lung cancer treatment, adapted from (Björk, Lynøe, and Juth 2015), and (ii) hip replacement surgery, adapted from (Björk, Juth, and Lynøe 2018). In both cases, smoking is the health-harming behavior, and the vignette structure, questions, and conditions mirror those of the kidney domain.

These two domains are chosen because they vary in the strength of the causal link between behavior and disease. The connection between smoking and lung cancer is widely recognized, whereas the connection between smoking and hip complications is less intuitive and weaker. Human baseline data are not available for these adapted vignettes on the specific questions we pose, so the cross-domain analysis focuses on within-LLM patterns rather than human-LLM comparison.

3.4 Models and Prompting

We test 12 commercially available and open-source LLMs spanning a broad range of capabilities, drawn from the Claude, GPT, DeepSeek, Gemini, and Llama families. We use 7 of these models both with and without reasoning (“thinking”) enabled, leading to a total of 19 distinct model configurations and allowing for a direct measurement of the effect of reasoning on moral judgment.

To accurately replicate the format of the original human studies, we prompt each model as a simulated study participant. A single session corresponds to one participant. The model receives the scenario and responds to each question in a multi-turn conversation with chat history preserved, just as a human participant would retain context across trials within a study session.¹ Between sessions, the conversation history is cleared, so that each of the 10 independent sessions per condition represents a fresh participant. Models respond in a structured JSON format, providing Likert scale ratings and, where applicable, an allocation choice for each trial.

Detailed model specifications, including API platforms, thinking-mode implementations, and open-source/commercial classification, are provided in Appendix D, which (like along with all other supplementary material) can be found in the full version of the paper (Hosseini, Khanna, and Pierce 2026).

4 Results

We present findings across five dimensions, (i) responsibility attribution, (ii) allocation decisions, (iii) sensitivity to episodic context, (iv) generalization across medical domains, and (v) the effect of extended reasoning. Together, these reveal a systematic judgment-consequence gap in which LLMs assess moral responsibility much like humans but refuse to act on that assessment. Statistical tests for all comparisons are reported in Appendix E.

¹To verify that our results are not an artifact of this multi-turn format, we also run all models in a single-turn condition in which each question is posed in isolation with no chat history from prior questions or scenarios. The results are qualitatively unchanged. See Appendix G for details.




Figure 1: Mean responsibility scores (5-point scale) by question and model group, aggregated across both behavior alteration conditions (continued vs. stopped). Error bars show 95% confidence intervals.⁴

4.1 Responsibility Attribution

In the behavior alteration vignette, each model reads a scenario in which Patient B has engaged in one of four health-harming behaviors (heavy drinking, drug use, smoking, or poor diet) and has subsequently developed kidney disease requiring a transplant. The scenario specifies that Patient B has either *stopped* or *continued* the behavior after diagnosis. The participant (human or LLM) then answers three questions on a 5-point Likert scale², probing successively deeper levels of responsibility: Is Patient B responsible for the *behavior* itself (Q1)? For developing the *disease* (Q2)? For being *deprived* of the transplant (Q3)? The exact questions (and prompts) can be found in Appendix A.

Agreement on behavioral responsibility. On the question of responsibility for behavior, LLMs and humans are in near-complete agreement (Figure 1). The human mean is 4.42 on the 5-point scale, and the mean across all 19 model configurations is 4.43³, indicating that both perceive patients to be responsible for their harmful behavior. This convergence holds across all four behavior types and is stable across both reasoning and non-reasoning models. When asked whether a patient who engages in a harmful behavior is responsible for that behavior, human and LLM moral judgments show no significant difference (Table 2).

²The range is “definitely no” (1) to “definitely yes” (5).

³Throughout the paper, we are primarily interested in trends that generalize across LLMs, rather than whether any specific LLM replicates or aligns with human choices. We therefore focus on aggregations across LLMs. Additionally, individual models exhibit limited diversity of responses, often selecting their modal response in a majority of trials. A detailed discussion of within-model consistency is provided in Appendix C.

⁴For LLMs, confidence intervals reflect variability across models within each group. For humans, they reflect between-subjects variability computed from condition-level standard deviations reported by (Chan et al. 2024). Because the original human study used a within-subjects design, these intervals are not directly comparable to the authors’ reported significance tests; statistical claims

Divergence on downstream responsibility. The agreement on behavioral responsibility does not carry through to the more morally loaded questions. Reasoning models continue to replicate human judgments closely on both disease responsibility and deprivation responsibility. Non-reasoning models, however, attribute notably less responsibility on these questions, pulling the overall LLM average below the human baseline (Figure 1). The gap between reasoning and non-reasoning models is largest on deprivation, the question most directly tied to whether the patient deserves to lose access to treatment. Per-model breakdowns are provided in Appendix B.1.

This divergence becomes even more pronounced when the question is framed in terms of *fault* rather than *responsibility*. In a separate set of vignettes (Section 4.2 and Section 4.3), participants are asked whether Patient B is at fault for being deprived of a kidney. LLMs (both reasoning and non-reasoning) are substantially less willing to assign fault as compared to humans, even when the underlying scenario is identical (Figure 4). Humans show no comparable sensitivity to this framing distinction. This pattern echoes the distinction drawn by Malle, Guglielmo, and Monroe (2014) between causal judgment and evaluative response as cognitively separable components of blame, since LLMs appear willing to make the causal attribution but resist the evaluative step of assigning fault. We return to this decomposition in Section 5.

Behavioral change is selectively morally relevant. When the scenario specifies that Patient B has stopped the harmful behavior after diagnosis, both humans and LLMs reduce their responsibility ratings relative to the continued-behavior condition (Figure 2). Crucially, this reduction is selective. In the human data, behavioral change has little effect on whether the patient is seen as responsible for the behavior itself, but substantially reduces perceived responsibility for the disease and, most strongly, for being deprived of treatment (Chan et al. 2024). Stopping does not retroactively erase responsibility for harmful behavior, but it does temper the judgment that the patient deserves to lose access to care.

LLMs largely replicate this selective pattern. Across all models and all behavior types (alcohol, smoking, drugs, bad diet), LLMs are more forgiving towards patients who discontinued the harmful behavior, and the reduction is concentrated on disease and deprivation responsibility rather than behavioral responsibility. The magnitude varies across models, with some closely matching the human profile, while others showing considerably larger reductions on deprivation than humans do. The qualitative shape of the effect, however, is consistent. See Appendix B.1 for more details.

4.2 Allocation Decisions

In the knowledge level vignette, participants read the same clinical scenario (Patient B has engaged in one of three⁵

about the human data throughout this paper cite the original within-subjects analyses.

⁵Unlike the previous stage, the “unhealthy eating” behavior is not considered for this stage.




Figure 2: Effect of behavioral change (stopped vs. continued) on responsibility scores. Error bars show 95% confidence intervals.

health-harming behaviors and has since stopped) but are now asked to make a consequential decision. Participants are first presented with a forced-choice allocation question: should the kidney go to Patient A, Patient B, or be decided randomly? They then answer four follow-up questions on a 7-point Likert scale⁶, probing whether deciding based on Patient B’s harmful behavior would be *unfair*, whether Patient B is responsible for the *behavior* and for the *kidney failure*, and whether it is Patient B’s own *fault* for not receiving the kidney (See Appendix A for the exact wording). The vignette also specifies the level of access Patient B had to information about the harmful effects of their behavior. The effects of this manipulation are discussed in Section 4.3.

Humans hold accountable; LLMs refuse to. A majority of human participants prefer to allocate the kidney to Patient A (67.6%), with a substantial minority opting to randomize and very few choosing Patient B. LLMs show a strikingly different pattern. Across all 19 model configurations, the dominant response is to decide randomly, with the vast majority of configurations selecting this option in nearly every response (Figure 3). When aggregated by model type, reasoning models choose Patient A only 9.4% of the time and non-reasoning models 24.7%, both far below the human rate. This preference holds across model families rather than being driven by a handful, as the per-configuration breakdowns in Figure 3 and Appendix B.2 show.

Exceptions do not align with humans. While there are some exceptions to this pattern, even the outlier models do not fully align with humans. Llama-4-Scout chooses Patient A almost exclusively and never randomizes, making it more extreme than humans, who still opt for random allocation roughly a third of the time. GPT-5.4-mini is the closest to the human profile, but it selects Patient B in roughly 20% of trials, far exceeding the human rate of 3.2%.

A normative gap, not a measurement artifact. The preference for randomization is not a failure to make nuanced

⁶The range is “strongly disagree” (1) to “strongly agree” (7).




Figure 3: Kidney allocation decisions across all 19 LLM configurations and human participants, sorted by proportion choosing Patient A (who never engaged in harmful behavior). (T) represents “thinking mode” and (NT) represents “non-thinking mode”.

judgments. As Section 4.1 shows, models clearly differentiate levels of responsibility on the Likert scale. Rather, they apply a distinct normative principle once the judgment becomes consequential. Their answers to the unfairness question reinforce this. LLMs rate behavior-based allocation as unfair at 5.6 on the 7-point scale, while humans average 3.9 (Figure 4). LLMs do not merely decline to act on responsibility, they actively endorse that acting on it would be unfair.

This exposes a systematic normative misalignment. Both humans and LLMs agree that patients who engage in health-harming behaviors bear responsibility for the behavior, for the disease, and for being deprived of treatment. But where humans treat these judgments as grounds for differential allocation in a life-or-death decision, LLMs refuse to do so, and frame the very act of behavior-based allocation as unjust.

This behavior contrasts with findings from Dickerson et al. (2025), who show in a similar setting that LLMs rarely choose to “flip a coin”. By re-running the allocation question with the third option worded in other ways, such as “flip a coin” or “leave it to chance”, we confirm that this preference is not an artifact of how the option is worded.⁷ This preference instead reflects features of our scenario, namely that the two patients are otherwise identical, that the medical effects of the behavior are equalized, and that the only remaining difference is one that models decline to act on. We elaborate

⁷It also holds when the allocation question is asked on its own, without the responsibility and fairness questions (Appendix G), so it is not an artifact of eliciting responsibility alongside the decision.

on this in Appendix F.

Reasoning traces reflect the gap. To check whether models indeed connect their responsibility judgment to the allocation decision, as opposed to not relating the two, we inspect the reasoning traces of thinking-enabled models on this decision (full method and results in Appendix H). Among responses where a model rated the patient highly responsible yet still chose to allocate randomly, the large majority first acknowledge that responsibility and then invoke fairness or equal treatment as the reason to set it aside, rather than ignoring it. The randomization is thus a deliberate refusal to let responsibility drive the decision, not a failure to register it, and this pattern is more pronounced in more capable models.

4.3 Effect of Access to Knowledge

The knowledge level vignette also varies the level of access Patient B had to information about the health risks of their behavior. In the *knowledge* condition, the patient knew that the behavior creates a risk of kidney failure. In the *access* condition, the patient did not know but had easy access to this information and most people in their community were aware of the risk. In the *no-access* condition, the patient neither knew nor had easy access to this information, and most people in their community were similarly unaware. This manipulation tests whether moral responsibility, in the eyes of humans and LLMs, requires informed choice.

LLMs factor in access to knowledge; humans do not.

Human responses are largely consistent across the three knowledge conditions, with minimal variation on all four Likert questions relative to the 7-point scale. LLMs, by contrast, show dramatically greater sensitivity. While they draw little distinction between the *knowledge* and *access* conditions, they assign substantially lower responsibility and fault scores when the patient had no access to information about the risks. Nearly every one of the 19 LLM configurations is more sensitive to the knowledge manipulation than humans on *every* question (Figure 4; see Appendix B.3 for per-model breakdowns). In Section 5, we discuss how this relates to the principle of informed consent in medical ethics (Emanuel et al. 2020).

The knowledge sensitivity also extends to allocation decisions. Among non-reasoning models, which show the most willingness to choose Patient A, the rate of Patient A allocation drops from 37.4% in the knowledge condition to 12.6% in the no-access condition. The knowledge manipulation thus modulates not only LLMs’ stated moral judgments but also their consequential decisions, in a way that human allocation decisions reflect only weakly. Per-model allocation breakdowns by knowledge condition are shown in Appendix B.2.

Incorrect advice is treated as equivalent to no access. In addition to the three knowledge conditions, we include three *ill-advised* conditions in which the patient was actively ill-advised, (i) told by a third party, (ii) by another person, or (iii) by an AI agent that the behavior does not create a risk of




Figure 4: Mean Likert scores (7-point scale) by information condition for humans, reasoning models, and non-reasoning models across all four questions in the knowledge levels experiment. The shaded regions represent 95% confidence intervals.

kidney failure. These conditions test whether LLMs distinguish between ignorance (no access to correct information) and being ill-informed (access to incorrect information).

They largely do not. Across all four Likert questions, mean LLM scores in the ill-advised conditions are far closer to the *no-access* condition than to the *knowledge* or *access* conditions (Figure 5). The same pattern holds for allocation decisions, where the rate of Patient A choices in ill-advised conditions is comparable to the no-access rate. The source of the incorrect advice also makes little difference. Whether the incorrect information came from another person, an AI agent, or an unspecified third party, the resulting scores are nearly indistinguishable. For LLMs, being misled is morally equivalent to never having known.

4.4 Reasoning Capabilities Amplify Responsibility Attribution

In Section 4.1, we observed that reasoning models tend to assign higher responsibility scores than non-reasoning ones. To isolate the effect of reasoning from other architectural differences, we directly compare the seven model families that support both a thinking-enabled and a thinking-disabled mode.

Thinking amplifies responsibility attribution for most models. For five of the seven families, enabling thinking increases responsibility scores, with the strongest effects on disease responsibility and deprivation responsibility (Figure 6b–c). Claude Haiku and DeepSeek Flash show the largest shifts, with average increases exceeding 0.7 scale points across the three questions. Behavioral responsibility (Figure 6a) is less affected, likely because scores are already near the ceiling of the scale regardless of mode. The effect is not universal, however, with Claude Opus and Gemini Flash-Lite showing near-zero or slightly negative shifts, indicating that the thinking toggle does not uniformly push all architectures in the same direction. Similar trends appear for the unfairness and fault questions in the knowledge level vignette, though the pattern is less clear (See Appendix B.5).

The normative gap persists. Despite these shifts in responsibility ratings, the core finding from Section 4.2 holds

across both modes, i.e. thinking-enabled and thinking-disabled models alike overwhelmingly choose to randomize rather than de-prioritizing the responsible patient. Extended reasoning amplifies the attribution of responsibility and, if anything, pushes models further toward randomization rather than toward the human pattern of differential allocation, reinforcing the interpretation that LLMs treat responsibility attribution and allocation as normatively distinct.

4.5 Generalization Across Medical Domains

To test whether these patterns are specific to kidney allocation, we replicate the experimental design across two additional domains, lung cancer treatment and hip replacement surgery (Björk, Lynøe, and Juth 2015; Björk, Juth, and Lynøe 2018). All three use smoking as the health-harming behavior, allowing direct comparison. Although the vignette structure and questions are identical across domains, the descriptions differ in how smoking relates to the medical condition. For lung cancer and kidney disease, smoking is described as contributing to the disease itself, while for hip replacement, smoking is unrelated to the underlying condition (namely, osteoarthritis) and instead complicates surgical recovery (see Appendix A for exact framing).

Core patterns replicate across domains. The central finding from the kidney allocation setting holds across all three medical contexts. In each domain, the vast majority of models overwhelmingly prefer to random allocation (Figure 7b), and the knowledge-sensitivity gradient observed in Section 4.3 is preserved (Figure 10). Responsibility attribution patterns from the behavior alteration vignette (Section 4.1) are also stable. Perceived responsibility for the behavior remains high across all three domains, while responsibility for disease and deprivation is comparable between lung cancer and kidney disease (Figure 7a).

Hip replacement is a partial exception. Randomization still dominates in all three domains. A few models allocate to Patient A more often in the hip context, which lifts the average Patient A rate to 23% versus 16% for kidney and 12%




Figure 5: Mean LLM scores for questions in the knowledge level vignette across all six information conditions. Error bars show 95% confidence intervals.




Figure 6: Effect of thinking mode on responsibility scores in the behavior alteration vignette. Each line connects a model family’s thinking-ON and thinking-OFF scores, averaged over stopped and continued conditions.

for lung (Figure 7b),⁸ and this shift is not significant across models (Appendix E.4). Where the domains do differ is in responsibility for the disease. Perceived disease responsibility is much lower for hip replacement, where smoking complicates surgical recovery rather than causing the underlying condition (Figure 7a and c). This gives models a forward-looking reason to weigh the behavior without treating it as the cause of the illness.

Behavior type has minimal effect. Within the kidney experiments, which tested four different health-harming behaviors (heavy drinking, drug use, smoking, and poor diet), neither humans nor LLMs show meaningful variation in allocation decisions or responsibility ratings depending on the specific behavior. This suggests the findings are driven by the structure of the moral scenario rather than by attitudes toward any particular substance (Figure 11).

5 Discussion

The judgment-consequence gap. Our most consistent finding is a disconnect between how LLMs assess moral responsibility and what they do with it. LLMs agree with hu-

⁸This is concentrated in a few models, as Claude-Haiku (NT) allocates to Patient A in 85% of hip trials despite randomizing in 100% of kidney and lung trials, and DeepSeek-v4-Pro (NT) shifts from 33% to 73%.

mans that patients bear responsibility for the behavior, the resulting disease, and, to a lesser extent, their deprivation of treatment. Yet in a consequential allocation decision they refuse to deprioritize the responsible patient and endorse the view that doing so would be unfair (5.6 vs. the human 3.9 on the unfairness question). We term this the *judgment-consequence gap*. LLMs treat responsibility as a backward-looking assessment (“B caused this”) but decline to make it prescriptive (“therefore B should bear the cost”). The reasoning traces make the separation explicit, as randomizing models acknowledge responsibility before setting it aside on fairness grounds rather than ignoring it (Section 4.2).

This separation has precedent in human moral cognition, where blame decomposes into a causal judgment about an agent’s role and a separable evaluative response (Malle, Guglielmo, and Monroe 2014; Shultz, Wright, and Schleifer 1986). Our data represent an extreme form, with the causal judgment intact but the evaluative response suppressed. In humans this gap usually arises from self-interest or situational pressure (Saltzstein 2010), whereas in LLMs it appears to be a principled refusal grounded in a competing fairness norm rather than a failure to act on a judgment.

Reasoning amplifies the gap. This interpretation is reinforced by our reasoning analysis (Section 4.4). Enabling extended thinking increases responsibility attribution for most model families, bringing LLMs closer to human judgments




Figure 7: Cross-domain comparison across lung cancer, kidney disease, and hip replacement. Note that in (b), Patient A is the patient who never engaged in harmful behavior. Error bars in (a) and (c) show 95% confidence intervals.

on the assessment side. Yet allocation decisions remain unchanged, with thinking-enabled and thinking-disabled models alike overwhelmingly preferring randomization. Deeper reasoning widens the judgment-consequence gap rather than closing it, suggesting that the gap is not a product of shallow processing but reflects a stable normative commitment.

Relationship to the value-action gap. A growing literature documents a *value-action gap* in LLMs, in which models’ stated value preferences diverge from the choices they make in scenario-based tasks (Shen, Clark, and Mitra 2025; Huang et al. 2026; Hosseini and Khanna 2026). Our finding is structurally distinct, since our LLMs are internally consistent, assessing responsibility while endorsing the fairness of setting it aside. The human pattern reflects *desert-based* reasoning, on which responsibility bears on who should shoulder a cost, whereas the LLM pattern reflects a contractualist (Scanlon 1998) or *egalitarian* (Rawls 1971) commitment to equal treatment regardless of desert (Persad, Wertheimer, and Emanuel 2009; Emanuel et al. 2020). We aim to *describe* this divergence, not endorse it, since there is no consensus that behavior-based allocation is justified. The distinction matters practically, as interventions that improve behavioral consistency will not move a model that is already consistent but normatively different, consistent with Lei et al. (2024), who find GPT-4o more oriented toward social justice than humans in unfair moral dilemmas.

Knowledge sensitivity and the role of epistemic state. Our second major finding is that LLMs are far more sensitive than humans to whether the patient could have known the health risks (Section 4.3). Human responses are largely flat across knowledge conditions, while LLMs sharply reduce responsibility and fault when the patient lacked access to information. This suggests LLMs weight informed consent, a foundational principle in medical ethics (Emanuel et al. 2020), more heavily than humans do, applying a more “textbook” framework that diverges from how people actually reason. Cushman (2008) show that human judgments often rely on intuitive heuristics rather than reasoning about an agent’s epistemic state. The asymmetry also carries an equity implication. Because accurate health information is

itself unequally available, a model that conditions responsibility on prior access extends more leniency to patients from under-informed communities. Whether seen as sensitivity to structural disadvantage or as a judgment resting on circumstances irrelevant to present need, it is a distributive effect that stakeholders would benefit by recognizing.

Not merely an alignment artifact. Reinforcement learning from human feedback (Christiano et al. 2017; Ouyang et al. 2022) penalizes outputs that could appear biased or discriminatory, which plausibly reinforces a preference for equal-chance allocation (Khamassi, Nahon, and Chatila 2024). Alignment training does not, however, account for the full pattern. Rather than passively abstaining, the models actively rate behavior-based allocation as unfair (5.6 vs. the human 3.9 on the unfairness question), and the gap grows with extended reasoning (Section 4.4) rather than fading. Our experiment on framing effects (Appendix F) shows that the preference depends on whether the third option means equal-chance allocation, not on how it is worded, and the reasoning traces (Appendix H) show models acknowledging responsibility before setting it aside on fairness grounds. Together these point to a stable normative commitment to procedural fairness rather than a surface-level avoidance of controversial outputs.

Implications for healthcare. LLMs are increasingly used as clinical decision-support tools (Liu et al. 2025; Xiao et al. 2025), and patients consult them on normative questions of prioritization and fairness, not only factual ones (Aharoni et al. 2024). Our findings highlight a risk beyond generic worries about accuracy, that of *partial alignment* creating false confidence. A clinician would see the model correctly identify that a patient bears responsibility, matching their own assessment, and might expect its recommendation to follow through. Instead the model refuses to deprioritize the patient and frames doing so as unjust. The danger is not a wrong answer but a right first step that obscures a normative divergence in the second. The risk is sharpest in conditions where the decision-maker is ill-advised (Section 4.3), where a model excuses a patient misled by an AI agent as readily as one misled by a person. As patients increasingly turn to

LLMs for health information, a model that discounts responsibility for those misinformed by AI may absorb the cost of other systems' errors into its own moral accounting.

Implications for alignment research. Our results suggest current alignment methods produce what Khamassi, Nahon, and Chatila (2024) call "weak alignment", a surface-level correspondence with human values that breaks down when models must compose moral assessments into decisions. The gap is not a knowledge deficit, since LLMs demonstrably possess the relevant moral knowledge through their accurate attributions, sensitivity to informed consent, and recognition that behavioral change matters. The deficit is in translating that knowledge into action, a compositional problem where models perform each sub-task competently but apply a different rule when combining them. Current fine-tuning and RLHF may be insufficient here, as they target surface behavior rather than the decision rules connecting assessment to action. Methods that address the bridge between judgment and action, alongside study of how training data and reward signals shape those rules, may be needed.

6 Limitations and Future Work

Stylized scenarios. Our scenarios deliberately abstract away from the clinical, logistical, and regulatory complexities of real-world organ allocation. This is by design, since the goal is not to evaluate LLMs for clinical deployment but to identify where they agree with or diverge from humans in morally charged reasoning, and to understand what might happen if they were consulted for advice in such contexts. Whether the judgment-consequence gap persists in richer, multi-factor settings is a natural extension.

Human baselines. Our human comparison data are drawn from prior studies (Chan et al. 2024; Björk, Lynøe, and Juth 2015) that we replicate as closely as possible in our LLM experiments. However, the human samples primarily reflect Western, educated populations (Henrich, Heine, and Norenzayan 2010), and moral reasoning norms vary across cultures and societies. The judgment-consequence gap itself is a phenomenon we observe *within* LLMs, and what is measured against these baselines is how sharply that gap diverges from human judgment. Since desert-based intuitions vary across cultures, that divergence could widen or narrow with more diverse human data, making replication with culturally varied human samples an important direction for future work.

Prompt sensitivity. Because our primary question is about alignment with human moral judgments, we adapt the prompt structure directly from the original human studies. LLM responses are nonetheless known to be sensitive to prompt wording, question ordering, and response format (Schröder et al. 2025). We show that exposure to other questions or scenarios does not affect our overall results (Appendix G), that the medical context does not (Section 4.5), and that relabeling the third allocation option does not (Appendix F). Since different framings can carry different ethical connotations, more systematic variation of

prompt framing and response format remains a meaningful direction for future work.

Absence of deliberative interaction. Our protocol queries each model in a single turn with no opportunity for follow-up or deliberation. In practice, clinical decision support involves iterative dialogue in which a human can probe the model's reasoning and request alternative framings. The judgment-consequence gap might narrow or widen under such interactive conditions, and evaluating LLM moral reasoning in multi-turn settings is an important direction for future work.

Scope of moral scenarios. While we test across three medical domains and multiple behavior types, all our scenarios involve self-inflicted health risks and pairwise patient comparisons. The findings may not generalize to allocation dilemmas involving non-behavioral factors (e.g., age, disability, social role) or to settings with more than two candidates. Extending to these broader scenarios would test whether the judgment-consequence gap is specific to desert-based reasoning or reflects a more general property of LLM moral cognition.

Moral versus legal framing. Our experiments frame responsibility in moral terms, but in legal contexts responsibility is constitutively linked to consequences. Legal liability entails penalties, sentencing, or differential treatment (Duff 2009; Arenella 1991). Whether LLMs would exhibit the same judgment-consequence gap under a legal framing is an open question. Moore, Stevens, and Conway (2011) show that sensitivity to punishment predicts moral judgment in humans, and LLMs' apparent insensitivity to the punitive implications of their assessments may reflect the absence of experiential grounding in reward and punishment. Testing whether a legal framing closes the gap would help clarify whether it is a domain-general property of LLM moral reasoning or specific to contexts where consequences remain implicit.

7 Conclusion

LLMs can identify moral responsibility with human-like accuracy, yet they systematically refuse to act on those judgments when allocating scarce medical resources. This judgment-consequence gap is not a reasoning failure, rather a stable normative commitment to procedural fairness that deepens with extended thinking and persists across model families and medical domains. As LLMs are increasingly consulted in contexts where moral assessment must inform consequential decisions, understanding where and why their moral reasoning diverges from human expectations is essential for building systems that are not just aligned on the surface but aligned in how they connect judgment to action.

Acknowledgments

This research was supported in part by NSF Awards IIS-2144413 and IIS-2107173. We would also like to thank the amazing reviewers for their suggestions that helped improving the paper.

References

- Aharoni, E.; Fernandes, S.; Brady, D. J.; Alexander, C.; Criner, M.; Queen, K.; Rando, J.; Nahmias, E.; and Crespo, V. 2024. Attributions toward Artificial Agents in a modified Moral Turing Test. *CoRR*, abs/2406.11854.
- Almeida, G. F.; Nunes, J. L.; Engelmann, N.; Wiegmann, A.; and de Araújo, M. 2024. Exploring the psychology of LLMs’ moral and legal reasoning. *Artificial Intelligence*, 333: 104145.
- Anthis, J. R.; Liu, R.; Richardson, S. M.; Kozlowski, A. C.; Koch, B.; Evans, J. A.; Brynjolfsson, E.; and Bernstein, M. S. 2025. LLM Social Simulations Are a Promising Research Method. *CoRR*, abs/2504.02234.
- Anthropic. 2025. Introducing Claude Haiku 4.5. Blog post. Released October 15, 2025.
- Anthropic. 2026. Introducing Claude Opus 4.7. Blog post. Released April 16, 2026.
- Arenella, P. 1991. Convicting the morally blameless: Re-assessing the relationship between legal and moral accountability. *UCIA l. rev.*, 39: 1511.
- Awad, E.; Dsouza, S.; Kim, R.; Schulz, J.; Henrich, J.; Shariff, A.; Bonnefon, J.-F.; and Rahwan, I. 2018. The Moral Machine Experiment. *Nature*, 563: 59–64.
- Ayers, J. W.; Poliak, A.; Dredze, M.; Leas, E. C.; Zhu, Z.; Kelley, J. B.; Faix, D. J.; Goodman, A. M.; Longhurst, C. A.; Hogarth, M.; and Smith, D. M. 2023. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Internal Medicine*, 183(6): 589–596.
- Ballesterio, G.; Hosseini, H.; Khanna, S.; and Shorrrer, R. I. 2026. Strategic Algorithmic Monoculture: Experimental Evidence from Coordination Games. *CoRR*, abs/2604.09502.
- Bhattacharyya, S.; Mehta, M.; Chen, L.; Salvador, C.; Lapedriza, À.; Dudy, S.; and Wang, J. Z. 2026. Expressing Social Emotions: Misalignment Between LLMs and Human Cultural Emotion Norms. *CoRR*, abs/2604.16757.
- bin Ahmad, M. S. Z.; and Takemoto, K. 2024. Large-scale moral machine experiment on large language models. *CoRR*, abs/2411.06790.
- Binz, M.; et al. 2025. A foundation model to predict and capture human cognition. *Nat.*, 644(8078): 1002–1009.
- Björk, J.; Juth, N.; and Lynøe, N. 2018. “Right to recommend, wrong to require”—an empirical and philosophical study of the views among physicians and the general public on smoking cessation as a condition for surgery. *BMC Medical Ethics*, 19(1): 2.
- Björk, J.; Lynøe, N.; and Juth, N. 2015. Are smokers less deserving of expensive treatment? A randomised controlled trial that goes beyond official values. *BMC medical ethics*, 16(1): 28.
- Chan, L.; Schaich Borg, J.; Conitzer, V.; Wilkinson, D.; Savulescu, J.; Zohny, H.; and Sinnott-Armstrong, W. 2022. Which features of patients are morally relevant in ventilator triage? A survey of the UK public. *BMC Medical Ethics*, 23(1): 33.
- Chan, L.; Sinnott-Armstrong, W.; Borg, J. S.; and Conitzer, V. 2024. Should Responsibility Affect Who Gets a Kidney? *Responsibility and Healthcare*, 35.
- Christiano, P. F.; Leike, J.; Brown, T. B.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep Reinforcement Learning from Human Preferences. In *Advances in Neural Information Processing Systems*, volume 30.
- Cookson, B.; Ebadian, S.; and Shah, N. 2026. Fairness Perceptions of Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, 35393–35401.
- Cushman, F. 2008. Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2): 353–380.
- Darley, J. M.; Klosson, E. C.; and Zanna, M. P. 1978. Intentions and their contexts in the moral judgments of children and adults. *Child Development*, 66–74.
- DeepSeek AI. 2026. DeepSeek V4 Preview Release. Blog post. Released April 24, 2026. Open-source weights.
- Dickerson, J. P.; Hosseini, H.; Khanna, S.; and Pierce, L. 2025. Who Gets the Kidney? Human-AI Alignment, Indecision, and Moral Values. *arXiv preprint arXiv:2506.00079*.
- Dillion, D.; Mondal, D.; Tandon, N.; and Gray, K. 2025. AI language model rivals expert ethicist in perceived moral expertise. *Scientific Reports*, 15.
- Ding, J.; Jiang, P.; Xu, Z.; Ding, Z.; Zhu, Y.; Jiang, J.; and Li, Y. 2025. “Pull or Not to Pull?": Investigating Moral Biases in Leading Large Language Models Across Ethical Dilemmas. *arXiv preprint arXiv:2508.07284*.
- Duff, A. 2009. Legal and Moral Responsibility. *Philosophy Compass*, 4(6): 978–986.
- Emanuel, E. J.; Persad, G.; Upshur, R.; Thome, B.; Parker, M.; Glickman, A.; Zhang, C.; Boyle, C.; Smith, M.; and Phillips, J. P. 2020. Fair allocation of scarce medical resources in the time of Covid-19.
- Freedman, R.; Borg, J. S.; Sinnott-Armstrong, W.; Dickerson, J. P.; and Conitzer, V. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, 283: 103261.
- Furnham, A.; Simmons, K.; and McClelland, A. 2000. Decisions concerning the allocation of scarce medical resources. *Journal of Social Behavior & Personality*, 15(2).
- Furnham, A.; Thomson, K.; and McClelland, A. 2002. The allocation of scarce medical resources across medical conditions. *Psychology and Psychotherapy: Theory, Research and Practice*, 75(2): 189–203.
- Garcia, B.; Qian, C.; and Palminteri, S. 2024. The moral turing test: Evaluating human-llm alignment in moral decision-making. *arXiv preprint arXiv:2410.07304*.
- Google. 2026. Gemini 3.1 Flash Lite: Our most cost-effective AI model yet. Blog post. Preview March 3, 2026; GA May 8, 2026.
- Google DeepMind. 2026a. Gemini 3.1 Pro Model Card. Model card. Released February 19, 2026.
- Google DeepMind. 2026b. Gemma 4. Model page. Released April 2, 2026. Apache 2.0 license.

- Guan, H.; He, J.; Fan, L.; Ren, Z.; He, S.; Yu, X.; Chen, Y.; Zheng, S.; Liu, T.-Y.; and Liu, Z. 2025. Modeling earth-scale human-like societies with one billion agents. *arXiv preprint arXiv:2506.12078*.
- Hendrycks, D.; Burns, C.; Basart, S.; Critch, A.; Li, J.; Song, D.; and Steinhardt, J. 2021. Aligning AI With Shared Human Values. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Henrich, J.; Heine, S. J.; and Norenzayan, A. 2010. The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3): 61–83.
- Hosseini, H.; and Khanna, S. 2026. Distributive fairness in large language models: Evaluating alignment with human values. *Advances in Neural Information Processing Systems*, 38: 86551–86599.
- Hosseini, H.; Khanna, S.; and Pierce, L. 2026. The Judgment-Consequence Gap: LLM Moral Reasoning in Healthcare Decisions. *arXiv:2608.05583*.
- Howe, P. D. L.; Fay, N.; Saletta, M.; and Hovy, E. 2023. ChatGPT’s advice is perceived as better than that of professional advice columnists. *Frontiers in Psychology*, Volume 14 - 2023.
- Hu, T.; Baumann, J.; Lupo, L.; Collier, N.; Hovy, D.; and Röttger, P. 2026. SimBench: Benchmarking the Ability of Large Language Models to Simulate Human Behaviors. In *The Fourteenth International Conference on Learning Representations*.
- Huang, J.; Qin, J.; Qiu, X.; Levy, S.; Kaufman, M. R.; and Dredze, M. 2026. Knowing But Not Doing: Convergent Morality and Divergent Action in LLMs. *CoRR*, abs/2601.07972.
- Ji, J.; Lei, R.; Pan, X.; Wei, Z.; Sun, H.; Lin, Y.; Chen, X.; Yang, Y.; Li, Y.; Ding, B.; et al. 2026. Leveraging LLM-based agents for social science research: insights from citation network simulations. *Humanities and Social Sciences Communications*.
- Jia, F.; Ye, Z.; Lai, S.; Shu, K.; Gu, J.; Bibi, A.; Hu, Z.; Jurgens, D.; Evans, J.; Torr, P. H.; et al. 2024. Can large language model agents simulate human trust behavior? *Advances in neural information processing systems*, 37: 15674–15729.
- Jiang, L.; Chai, Y.; Li, M.; Liu, M.; Fok, R.; Dziri, N.; Tsvetkov, Y.; Sap, M.; and Choi, Y. 2026. Artificial Hive-mind: The Open-Ended Homogeneity of Language Models (and Beyond). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Jiang, L.; Hwang, J. D.; Bhagavatula, C.; Bras, R. L.; Liang, J. T.; Levine, S.; Dodge, J.; Sakaguchi, K.; Forbes, M.; Hessel, J.; Borchardt, J.; Sorensen, T.; Gabriel, S.; Tsvetkov, Y.; Etzioni, O.; Sap, M.; Rini, R.; and Choi, Y. 2025. Investigating machine moral judgement through the Delphi experiment. *Nat. Mac. Intell.*, 7(1): 145–160.
- Jin, Z.; Levine, S.; Adauto, F. G.; Kamal, O.; Sap, M.; Sachan, M.; Mihalcea, R.; Tenenbaum, J.; and Schölkopf, B. 2022. When to Make Exceptions: Exploring Language Models as Accounts of Human Moral Judgment. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Kapania, S.; Agnew, W.; Eslami, M.; Heidari, H.; and Fox, S. E. 2025. Simulacrum of stories: Examining large language models as qualitative research participants. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–17.
- Kempton, H.; Heilinger, J.-C.; and Nagel, S. K. 2023. “I’m afraid I can’t let you do that, Doctor”: meaningful disagreements with AI in medical contexts. *AI & society*, 38(4): 1407–1414.
- Kempton, H.; and Nagel, S. K. 2022. Responsibility, second opinions and peer-disagreement: ethical and epistemological challenges of using AI in clinical diagnostic contexts. *Journal of Medical Ethics*, 48(4): 222–229.
- Keswani, V.; Cousins, C.; Nguyen, B.; Conitzer, V.; Heidari, H.; Borg, J. S.; and Sinnott-Armstrong, W. 2026. Moral Change or Noise? On Problems of Aligning AI With Temporally Unstable Human Feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, 37501–37509.
- Khamassi, M.; Nahon, M.; and Chatila, R. 2024. Strong and weak alignment of large language models with human values. *CoRR*, abs/2408.04655.
- Kim, Y.; Park, C.; Jeong, H.; Chan, Y. S.; Xu, X.; McDuff, D.; Lee, H.; Ghassemi, M.; Breazeal, C.; and Park, H. W. 2024. Mdagents: An adaptive collaboration of llms for medical decision-making. *Advances in Neural Information Processing Systems*, 37: 79410–79452.
- Kolluri, A.; Wu, S.; Park, J. S.; and Bernstein, M. S. 2025. Finetuning LLMs for Human Behavior Prediction in Social Science Experiments. In Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; and Peng, V., eds., *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, 30096–30111. Association for Computational Linguistics.
- Krütli, P.; Rosemann, T.; Törnblom, K. Y.; and Smieszek, T. 2016. How to fairly allocate scarce medical resources: ethical argumentation under scrutiny by health professionals and lay people. *PloS one*, 11(7): e0159086.
- Lei, Y.; Liu, H.; Xie, C.; Liu, S.; Yin, Z.; Chen, C.; Li, G.; Torr, P.; and Wu, Z. 2024. FairMindSim: Alignment of Behavior, Emotion, and Belief in Humans and LLM Agents Amid Ethical Dilemmas. *CoRR*, abs/2410.10398.
- Lerner, M. J. 1980. *The Belief in a Just World: A Fundamental Delusion*. Plenum Press.
- Li, S. S.; Balachandran, V.; Feng, S.; Ilgen, J. S.; Pierson, E.; Koh, P. W.; and Tsvetkov, Y. 2024. Mediq: Question-asking llms and a benchmark for reliable interactive clinical reasoning. *Advances in Neural Information Processing Systems*, 37: 28858–28888.

- Lim, K.; Kang, U.; Li, X.; Kim, J. S.; Jung, Y.; Park, S.; and Kim, B. 2025. Susceptibility of Large Language Models to User-Driven Factors in Medical Queries. *CoRR*, abs/2503.22746.
- Liu, Q.; Yang, R.; Gao, Q.; Liang, T.; Wang, X.; Li, S.; Lei, B.; and Gao, K. 2025. A Review of Applying Large Language Models in Healthcare. *IEEE Access*, 13: 6878–6892.
- Malle, B.; Guglielmo, S.; and Monroe, A. 2014. Moral, Cognitive, and Social: The Nature of Blame.
- McCradden, M. D.; Joshi, S.; Anderson, J. A.; and London, A. J. 2023. A normative framework for artificial intelligence as a sociotechnical system in healthcare. *Patterns*, 4(11).
- McElfresh, D. C.; Chan, L.; Doyle, K.; Sinnott-Armstrong, W.; Conitzer, V.; Borg, J. S.; and Dickerson, J. P. 2021. Indecision modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5975–5983.
- Meta AI. 2024. Llama 3.3 70B. Hugging Face model card. Released December 6, 2024.
- Meta AI. 2025. The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation. Blog post. Released April 5, 2025.
- Moore, A. B.; Stevens, J.; and Conway, A. R. 2011. Individual differences in sensitivity to reward and punishment predict moral judgment. *Personality and Individual Differences*, 50(5): 621–625.
- Newsham, L.; and Prince, D. 2025. Personality-Driven Decision-Making in LLM-Based Autonomous Agents. *arXiv preprint arXiv:2504.00727*.
- OpenAI. 2025. Introducing gpt-oss. Blog post. Released August 5, 2025. Apache 2.0 license.
- OpenAI. 2026a. Introducing GPT-5.4 mini and nano. Blog post. Released March 17, 2026.
- OpenAI. 2026b. Introducing GPT-5.5. Blog post. Released April 23, 2026.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training Language Models to Follow Instructions with Human Feedback. In *Advances in Neural Information Processing Systems*, volume 35, 27730–27744.
- Park, J. S.; Zou, C. Q.; Shaw, A.; Hill, B. M.; Cai, C.; Morris, M. R.; Willer, R.; Liang, P.; and Bernstein, M. S. 2024. Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109*, 52.
- Persad, G.; Wertheimer, A.; and Emanuel, E. J. 2009. Principles for allocation of scarce medical interventions. *The lancet*, 373(9661): 423–431.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Rudolph, U.; Roesch, S.; Greitemeyer, T.; and Weiner, B. 2004. A meta-analytic review of help giving and aggression from an attributional perspective: Contributions to a general theory of motivation. *Cognition and emotion*, 18(6): 815–848.
- Saltzstein, H. D. 2010. The Relation between Moral Judgment and Behavior: A Social-Cognitive and Decision-Making Analysis. *Human Development*, 37(5): 299–312.
- Santurkar, S.; Durmus, E.; Ladhak, F.; Lee, C.; Liang, P.; and Hashimoto, T. 2023. Whose opinions do language models reflect? In *International conference on machine learning*, 29971–30004. PMLR.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scherrer, N.; Sh, C.; Feder, A.; and Blei, D. M. 2023. Evaluating the moral beliefs encoded in LLMs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23. Red Hook, NY, USA: Curran Associates Inc.
- Schröder, S.; Morgenroth, T.; Kuhl, U.; Vaquet, V.; and Paaßen, B. 2025. Large language models do not simulate human psychology. *arXiv preprint arXiv:2508.06950*.
- Shen, H.; Clark, N.; and Mitra, T. 2025. Mind the Value-Action Gap: Do LLMs Act in Alignment with Their Values? In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 3097–3118.
- Shen, H.; Kneare, T.; Ghosh, R.; Yang, Y.-J.; Clark, N.; Mitra, T.; and Huang, Y. 2025. ValueCompass: A Framework for Measuring Contextual Value Alignment Between Human and LLMs. In Zhang, C.; Allaway, E.; Shen, H.; Miculicich, L.; Li, Y.; M’hamdi, M.; Limkonchotiawat, P.; Bai, R. H.; T.y.s.s., S.; Han, S. S.; Thapa, S.; and Rim, W. B., eds., *Proceedings of the 9th Widening NLP Workshop*, 75–86. Suzhou, China: Association for Computational Linguistics. ISBN 979-8-89176-351-7.
- Shool, S.; Adimi, S.; Amleshi, R. S.; Bitaraf, E.; Golpira, R.; and Tara, M. 2025. A systematic review of large language model (LLM) evaluations in clinical medicine. *BMC Medical Informatics Decis. Mak.*, 25(1): 117.
- Shultz, T. R.; Wright, K.; and Schleifer, M. 1986. Assignment of moral responsibility and punishment. *Child development*, 177–184.
- Singhal, K.; et al. 2023. Large Language Models Encode Clinical Knowledge. *Nature*, 620: 172–180.
- Sorin, V.; Korfiatis, P.; Collins, J. D.; Apakama, D.; Omar, M.; Glicksberg, B. S.; Yeow, M.; Brandeland, M.; Nadkarni, G. N.; and Klang, E. 2025. Socio-Demographic Modifiers Shape Large Language Models’ Ethical Decisions. *J. Heal. Informatics Res.*, 9(4): 567–586.
- Suh, J.; Jahanparast, E.; Moon, S.; Kang, M.; and Chang, S. 2025. Language Model Fine-Tuning on Scaled Survey Data for Predicting Distributions of Public Opinions. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 21147–21170. Vienna, Austria: Association for Computational Linguistics. ISBN 979-8-89176-251-0.
- Takemoto, K. 2024. The moral machine experiment on large language models. *Royal Society open science*, 11(2).
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large

Language Models in Medicine. *Nature Medicine*, 29: 1930–1940.

Tseng, Y.; Huang, Y.; Hsiao, T.; Chen, W.; Huang, C.; Meng, Y.; and Chen, Y. 2024. Two Tales of Persona in LLMs: A Survey of Role-Playing and Personalization. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, 16612–16631. Association for Computational Linguistics.

Vrdoljak, J.; Boban, Z.; Vilović, M.; Kumrić, M.; and Božić, J. 2025. A Review of Large Language Models in Medical Education, Clinical Decision Support, and Healthcare Administration. *Healthcare*, 13(6).

Wang, R.; Yu, H.; Zhang, W.; Qi, Z.; Sap, M.; Bisk, Y.; Neubig, G.; and Zhu, H. 2024. Sotopia- π : Interactive learning of socially intelligent language agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12912–12940.

Weiner, B. 1985. An attributional theory of achievement motivation and emotion. *Psychological review*, 92(4): 548.

Weiner, B.; Perry, R. P.; and Magnusson, J. 1988. An attributional analysis of reactions to stigmas. *Journal of personality and social psychology*, 55(5): 738.

Wu, X.; Zhao, Y.; Zhang, Y.; Wu, J.; Zhu, Z.; Zhang, Y.; Ouyang, Y.; Zhang, Z.; Wang, H.; Lin, Z.; et al. 2024. Med-journey: Benchmark and evaluation of large language models over patient clinical journey. *Advances in Neural Information Processing Systems*, 37: 87621–87646.

Xiao, H.; Zhou, F.; Liu, X.; Liu, T.; Li, Z.; Liu, X.; and Huang, X. 2025. A comprehensive survey of large language models and multimodal large language models in medicine. *Inf. Fusion*, 117(C).

Yu, K.-H.; Healey, E.; Leong, T.-Y.; Kohane, I. S.; and Manrai, A. K. 2024. Medical Artificial Intelligence and Human Values. *New England Journal of Medicine*, 390(20): 1895–1904.

Zhou, J.; Hu, M.; Li, J.; Zhang, X.; Wu, X.; King, I.; and Meng, H. 2024a. Rethinking Machine Ethics—Can LLMs Perform Moral Reasoning through the Lens of Moral Theories? In *Findings of the Association for Computational Linguistics: NAACL 2024*, 2227–2242.

Zhou, L.; Xu, Y.; Wang, Z.; and Wang, D. 2025. FlockVote: LLM-Empowered Agent-Based Modeling for Simulating US Presidential Elections. *arXiv preprint arXiv:2512.05982*.

Zhou, X.; Zhu, H.; Mathur, L.; Zhang, R.; Yu, H.; Qi, Z.; Morency, L.-P.; Bisk, Y.; Fried, D.; Neubig, G.; et al. 2024b. Sotopia: Interactive evaluation for social intelligence in language agents. In *International Conference on Learning Representations*, volume 2024, 40975–41019.