

Improving Behavioral Alignment in LLM Social Simulations via Context Formation and Navigation

Letian Kong

CUHK Business School, The Chinese University of Hong Kong, kongletian@link.cuhk.edu.hk

Qianran(Jenny) Jin

CUHK Business School, The Chinese University of Hong Kong, jennyjin@cuhk.edu.hk

Renyu(Philip) Zhang

CUHK Business School, The Chinese University of Hong Kong, philipzhang@cuhk.edu.hk

Large language models (LLMs) are increasingly used to simulate human behavior in experimental settings, but they systematically diverge from human decisions in *complex decision-making environments* where participants must anticipate others' actions and form beliefs based on observed behavior. We propose a two-stage framework for improving behavioral alignment. The first stage, *context formation*, explicitly specifies the experimental design to establish an accurate representation of the decision task and its context. The second stage, *context navigation*, guides the reasoning process within that representation to make decisions. We validate this framework through a focal replication of a sequential purchasing game with quality signaling (Kremer and Debo 2016), extending to a crowdfunding game with costly signaling (Cason et al. 2025) and a demand-estimation task (Gui and Toubia 2025) to test generalizability across decision environments. Across four state-of-the-art (SOTA) models (GPT-4o, GPT-5, Claude-4.0-Sonnet-Thinking, DeepSeek-R1), we find that complex decision-making environments require both stages to achieve behavioral alignment with human benchmarks, whereas the simpler demand-estimation task requires only *context formation*. Our findings clarify when each stage is necessary and provide a systematic approach for designing and diagnosing LLM social simulations as complements to human subjects in behavioral research.

Key words: LLM Social Simulations, AI Agents, Behavioral Alignment, Context Engineering, Experiments

1. Introduction

Generative artificial intelligence (AI) and large language models (LLMs) are opening new possibilities for studying human behavior in social science and business research (Bail 2024, Demszky et al. 2023, Gopal et al. 2025, Kane et al. 2025, Wiberg et al. 2025). A particularly promising application is using LLMs to simulate human decision-making in experiments, an approach known as *LLM*

social simulations. In these simulations, LLMs act as synthetic subjects, making it possible to study behavioral patterns and evaluate interventions at unprecedented scale (Anthis et al. 2025, Charness et al. 2025, Manning and Horton 2025, Toubia et al. 2025). These simulations offer clear advantages as they are fast, reproducible, and inexpensive compared to recruiting human participants. More importantly, they enable research in settings where human experiments are costly, impractical, or impossible. Recent research explores applications in market research and related evaluation tasks, such as preference measurement and perceptual analysis (Brand et al. 2024, Li et al. 2024, Wang et al. 2024), while industry pilots are already emerging in automated experimentation (Wang et al. 2025) and synthetic customer research.¹ As LLM social simulations expand from research settings into practice, reliable *behavioral alignment*, defined as the extent to which LLM-simulated subjects' decisions correspond to those of human participants, becomes increasingly critical for both scientific validity and practical deployment.

Despite this importance, fundamental questions about when LLM social simulations reliably match human behavior remain unanswered. Early evidence from psychology and economics suggests that LLMs closely reproduce aggregate human response patterns in relatively simple settings, such as trust games and personality assessments (Aher et al. 2023, Chen et al. 2023, Toubia et al. 2025). However, this apparent success breaks down in *complex decision-making environments* characterized by strategic interdependence and endogenous belief formation, where participants must anticipate others' actions and form beliefs based on observed behavior. In these environments, LLMs often diverge from human decisions (Akata et al. 2025, Gao et al. 2024), and LLM-based replications of established experimental findings yield inconsistent results (Chen et al. 2025a, Kirshner 2024). Notably, these alignment failures arise precisely in the complex decision-making environments where human experiments are most difficult and costly to conduct, thereby weakening the central motivation for using LLM social simulations in the first place.

¹ Examples include Delve AI (<https://www.delve.ai/>), which generates synthetic customer personas for market research.

Two broad approaches have been used to improve the alignment of LLM social simulations, but neither provides a systematic solution for complex decision-making environments. First, data-driven approaches embed human data directly into LLM social simulations through fine-tuning or calibration. Simulations show substantial gains when models are calibrated to observed human choices (Manning and Horton 2025) or trained on large-scale behavioral data (Binz et al. 2025). However, the effectiveness of fine-tuning depends critically on the match between training tasks and target applications, as models can fail to generalize across different decision environments (Gao et al. 2024, Toubia et al. 2025). More fundamentally, reliance on human data creates a bottleneck in the complex decision-making environments where simulations are most valuable, as such data are costly, task-specific, and often unavailable.

Second, prompt-based approaches use prompt engineering to improve simulation alignment without additional human data. Widely used strategies include Chain-of-Thought (CoT) reasoning (Kojima et al. 2022, Wei et al. 2022) and persona adjustments (Aher et al. 2023, Kirshner 2024, Li et al. 2024, Xie et al. 2024). While these approaches can improve performance in individual experimental replications, accumulated evidence shows that their effects are inconsistent across tasks and models (Goli and Singh 2024, Leng and Yuan 2023). Some studies go beyond generic prompting strategies and propose targeted interventions for particular decision environments. For example, alignment can be improved by prompting models to explicitly reason about others' likely actions in repeated games (Akata et al. 2025), or by making the experimental design explicit to the model in order to reduce confounding in LLM-based simulations (Gui and Toubia 2025). However, these task-specific solutions do not generalize across different decision environments, and we still lack a systematic framework for understanding when different approaches succeed or fail.

The main goal of this paper is to examine the sources of behavioral misalignment when LLMs simulate human decision-making in complex decision-making environments. Motivated by the widely observed divergence between LLM-simulated subjects and human behavior in such settings, despite the growing capabilities of state-of-the-art (SOTA) LLMs, we seek to address the following key

research question: *How can behavioral alignment in LLM social simulations of complex decision-making be systematically diagnosed and improved?*

To answer our research question, we develop a two-stage theoretical framework inspired by human problem-solving (Newell and Simon 1972, Simon and Newell 1971). The first stage, *context formation*, explicitly specifies the experimental design to establish an accurate representation of the decision task and its context. The second stage, *context navigation*, guides the reasoning process within that representation to make decisions. We empirically evaluate this framework through simulations of three established experimental studies that vary in strategic interdependence and endogenous belief formation. By systematically varying whether prompts provide *context formation*, *context navigation*, or both, we isolate the contribution of each stage to behavioral alignment. We examine models' CoT reasoning traces to diagnose the sources of misalignment. Comparing LLM-simulated subjects' behavior to human benchmarks allows us to assess when and in what ways each stage contributes to behavioral alignment across decision environments and models. We highlight our contributions as follows.

1.1. Theoretical Framework for LLM Decision-Making

As our main contribution, we propose a theoretical framework that conceptualizes behavioral alignment between LLMs and humans as a two-stage process involving *context formation* and *context navigation*. By explicitly separating these stages, our framework provides a holistic and theoretically motivated perspective for understanding and improving behavioral alignment across decision environments.

This separation is critical because it reveals a hierarchical structure. *Context navigation* is effective only when grounded in accurate *context formation*. In complex decision-making environments, providing LLMs with task instructions alone leaves both task representation and the reasoning strategy ambiguous. As a result, models must infer both from patterns learned during pretraining, and the resulting representations and reasoning strategies can systematically differ from human approaches. This hierarchical relationship means that misalignment can stem from inadequate *context formation*,

from inappropriate *context navigation* within an otherwise accurate representation, or from failures in both stages. Importantly, this hierarchy explains why generic reasoning prompts such as CoT often show inconsistent effects because they attempt to improve navigation while explicitly assuming a task representation that may not exist. Our framework systematically addresses both stages, with particular emphasis on establishing accurate *context formation* as the necessary foundation for effective behavior alignment.

1.2. Empirical Validation of Our Framework

To demonstrate the practical scope of our framework, we empirically validate its generalizability along two dimensions: decision environments and model choices. Across decision environments, we apply the framework to three established experimental studies, including two complex behavioral experiments characterized by strategic interdependence and endogenous belief formation, and a simple demand estimation study driven by the law of demand. Across models, we implement a common simulation procedure and evaluation protocol using multiple SOTA LLMs from different model families, including GPT-4o, GPT-5, Claude-4.0-Sonnet-Thinking, and DeepSeek-R1.

Across these settings, our results reveal a consistent pattern: in complex decision-making environments, achieving behavioral alignment requires both *context formation* and *context navigation*, whereas in simple decision environments, *context formation* alone is sufficient. These patterns are robust across models, demonstrating that the framework generalizes across both decision environments and model choices. In addition, we examine models' CoT reasoning traces to diagnose the sources of misalignment, moving beyond outcome comparisons to understand how models interpret tasks and reach decisions.

Our study provides two key insights for designing reliable LLM social simulations. First, the need for explicit reasoning guidance depends on decision environments. In simple, non-strategic settings, making the task representation explicit and accurate is often sufficient for alignment, whereas strategically interdependent and belief-driven environments require additional guidance on how to reason within the representation. Second, while advances in LLM reasoning capabilities (e.g., Guo et al.

2025) can reduce reliance on explicit *context navigation* guidance, they do not eliminate the need for careful *context formation*. Together, our findings highlight that misalignment in LLM social simulations can stem from coherent reasoning applied to inaccurate task representations, with the relative importance of each stage depending on task complexity.

The rest of this paper is organized as follows. In Section 2, we review the relevant literature. Section 3 develops our theoretical framework based on human problem-solving theory. In Section 4, we present our focal empirical application (Kremer and Debo 2016) and evaluate the alignment with human benchmarks. Section 5 applies our framework to two additional experiments (Cason et al. 2025, Gui and Toubia 2025) and demonstrate its generalizability across multiple SOTA models. Section 6 concludes.

2. Literature Review

Our study is closely related to three streams of research: (i) generative AI and LLM in business applications, (ii) the use of LLMs to simulate human subjects in experiments and simulations, and (iii) alignment challenges and solutions in LLM social simulation.

2.1. Generative AI and LLM in Business Applications

Generative AI has introduced a new class of technologies capable of producing text, code, designs, and other content at scale, spurring a rapidly growing literature on its economic and social implications. In work settings, access to LLMs and their outputs can accelerate task completion and, in many contexts, improve output quality and productivity at both the individual and organizational levels (Brynjolfsson et al. 2025, Dell’Acqua et al. 2023, Noy and Zhang 2023, Peng et al. 2023). For creative tasks, LLM-generated ideas and content tend to increase individual creative output, but may reduce diversity and novelty in collective outputs (Doshi and Hauser 2024, Hou et al. 2025). On online platforms, AI-generated content reshapes voluntary contribution and participation dynamics, often concentrating activity among higher-capability contributors (Burtch et al. 2024, Shan and Qiu 2025), whereas AI-generated metadata (e.g., titles and hashtags) significantly boosts user engagement by improving user-content matching and inspiring better human-AI co-creation (Zhang et al. 2024). At

the labor-market level, the diffusion of generative AI is associated with observable changes in labor demand and employment outcomes, with evidence that the magnitude and direction of these effects vary across occupations and depend on how AI is adopted in work settings (Demirci et al. 2025, Eloundou et al. 2023, Hu et al. 2025, Hui et al. 2024).

Beyond studying the impacts of generative AI, a parallel research stream develops LLM-based methods as tools to empower business applications and research. Generative models enhance content production workflows by increasing the productivity and scalability of text and visual content creation in marketing and communication, while also introducing the risk of greater homogenization of marketing content across firms (Hackenburg and Margetts 2024, Liu et al. 2025a, Reisenbichler et al. 2025, 2022). They are also embedded in decision support and process automation, where generated evaluations and classifications feed into systematic processes for strategic analysis, experimentation, and professional knowledge summarization (Doshi et al. 2025, Guo et al. 2024, Liu et al. 2025b, Ye et al. 2025). Related work examines how work is split between humans and LLMs, showing that different ways of using models alongside humans lead to different creative, innovation, and research outcomes (Arora et al. 2025, Boussioux et al. 2024, Chen and Chan 2024, Cillo and Rubera 2025). Other studies integrate LLMs into research workflows to support conceptual development, systematic analysis, and scalable data construction in behavioral research (Carlson and Burbano 2024, Cheng et al. 2024, Grossmann et al. 2023, Kim and Lee 2023, Zhou et al. 2024).

Within this broad research landscape on generative AI in business, LLMs have demonstrated substantial generative and reasoning capabilities across diverse application domains. An open question is how these capabilities translate when models are used to produce judgments and decisions to resemble those of humans.

2.2. LLM Social Simulations

As LLMs become more prevalent in business and academia, researchers are examining whether they can complement or even substitute for human participants in surveys and lab experiments. The idea is to use LLMs to generate judgments, behaviors, and choices that would otherwise require

recruiting human subjects, offering a scalable and low-cost alternative to traditional social science and business research (Anthis et al. 2025, Grossmann et al. 2023). Early evidence from Argyle et al. (2023), who introduce the notion of “silicon samples”, shows that LLMs can closely reproduce the response distributions observed in human surveys.

A growing body of research has since examined the extent to which these silicon samples resemble human subjects in their behavioral and psychological responses. In relatively simple experimental settings, LLMs have been shown to reproduce fundamental human traits such as fairness and reciprocity in classic games (Aher et al. 2023, Mei et al. 2024, Xie et al. 2024), align closely with human moral consensus (Dillion et al. 2023), emulate behaviors such as altruism and risk aversion (Filippas et al. 2024), and display human-like cognitive biases (Chen et al. 2025b, Salecha et al. 2024). Consistent with these findings, large-scale replication efforts report high success rates, including replication of approximately 75% of 156 psychological and management experiments (Cui et al. 2024) and eight out of nine classic behavioral-operations experiments (Kirshner 2024). Beyond replication, several studies propose methodological frameworks that integrate LLM-simulated subjects into research workflows to support theory development and hypothesis testing (Ludwig et al. 2025, Manning et al. 2024, Tranchero et al. 2024), and deploy them in broader society and system-level simulations (Allouah et al. 2025, Chuang et al. 2024, Park et al. 2022, Yang et al. 2025, 2024).

Taken together, this stream of research suggests that LLMs can reproduce many well-established human behaviors across tasks with clearly specified decision problems and limited ambiguity in how information maps to choices. However, human decision-making often unfolds in multi-agent environments characterized by strategic interdependence and endogenous belief formation, which we refer to as complex decision-making environments. It remains unclear how well LLMs capture human behaviors in such settings. Our study contributes to this literature by systematically evaluating LLM-simulated subjects in these environments.

2.3. LLM Alignment Challenges and Solutions

To effectively leverage LLMs in research and practical applications, alignment, the extent to which model outputs correspond to human intentions, preferences, and values, has been a central challenge. In LLM social simulations, this concern becomes whether alignment mechanisms developed in

computer science are sufficient to ensure the generation of human-like behavior in complex decision-making environments. Existing approaches primarily operate either at training or inference stages. Training-time methods, including supervised fine-tuning (SFT, Ouyang et al. 2022, Wei et al. 2021), reinforcement learning with human feedback (RLHF, Christiano et al. 2017, Ziegler et al. 2019), and reinforcement learning with verifiable rewards (RLVR, Guo et al. 2025, Lambert et al. 2024), shape model behavior by encouraging adherence to natural-language instructions, human preferences, and verifiable reasoning objectives. At the inference time, methods relying on careful and systematic prompt design such as Chain-of-Thought (CoT, Kojima et al. 2022, Wei et al. 2022), synergizing reasoning and action (ReAct, Yao et al. 2022), and tree-of-thoughts (ToT, Yao et al. 2023) guide models towards more reliable reasoning across diverse tasks.

In LLM social simulations, behavioral alignment, the extent to which LLM-simulated subjects' decisions correspond to those of human participants, remains a central bottleneck despite advances in training-time and inference-time alignment methods. Relative to human benchmarks, LLMs often exhibit higher levels of rationality (Chen et al. 2023, Kirshner 2024), and less stable intertemporal preference structures (Goli and Singh 2024). Misalignment is even more pronounced in complex decision-making environments where participants need to anticipate others' actions and form beliefs based on observed behavior. In such settings, LLMs tend to form systematically biased beliefs and exhibit cooperation patterns and reasoning processes that diverge from those of humans (Bauer et al. 2025, Brookins and DeBacker 2023, Gao et al. 2024). While some studies go beyond documenting misalignment by providing diagnostic analyses of its sources, most stop short of proposing generalizable strategies for improving alignment.

Closest to our paper are a few recent studies that propose targeted interventions to improve behavioral alignment in specific settings. For example, Akata et al. (2025) show that coordination behavior can be improved through a social CoT prompting strategy that encourages models to anticipate others' actions. Gui and Toubia (2025) demonstrate that standard blinded experiment designs can induce confoundedness in LLM-based simulations and propose unblinding as a remedy.

Other approaches incorporate human data more directly. [Binz et al. \(2025\)](#) introduce *Centaurs*, an LLM fine-tuned on large-scale human behavioral data from psychological experiments. [Manning and Horton \(2025\)](#) combine theory-grounded mechanisms with human data from related experiments to predict behaviors in novel strategic environments. While effective within their respective domains, these approaches remain task-specific and do not offer a unified theoretical account of why LLMs diverge from human behavior in complex decision-making.

Building on this literature, we adopt a decision-theoretic perspective that shifts attention from task-specific fixes to the underlying decision processes that lead to behavioral alignment. Rather than focusing on procedural and/or technical adjustments alone, we examine how LLMs form representations of decision tasks and reason with those representations to make decisions. We propose a framework that diagnoses the sources of behavioral misalignment in complex decision-making environments and provides a systematic approach for addressing it. By doing so, our approach connects principles from human cognition and social decision-making to the methodological design of LLM social simulations.

3. Theoretical Framework

To develop this framework for LLM-based simulations of human decision-making behaviors, we draw on [Simon and Newell \(1971\)](#)'s theory of human problem-solving.

3.1. Human Problem-Solving

Research on human problem-solving shows that people usually make decisions in two stages ([Newell and Simon 1972](#), [Simon and Newell 1971](#)). In the first stage, they form an internal representation of the decision task by clarifying what the question is asking, what information is relevant, and what the goal of the problem is. This task representation provides the structure within which the task can be understood. In the second stage, people reason within this representation to determine how to solve the problem, evaluating potential actions using heuristics that guide their search. Empirical studies further support this two-stage view of problem-solving. Changes in how a problem is described lead people to form different task representations ([Hayes and Simon 1974](#)). Furthermore, because people

navigate problems through selective search rather than exhaustive evaluation (Bhaskar and Simon 1977, Newell and Simon 1972), changes in task representation shape the reasoning process and can thus produce different solution paths even when the underlying task remains the same (Simon and Hayes 1976).

Formally, we define the query Q as the externally provided description of the task (e.g., the instructions and information shown to participants), the context C_H as the internal representation of the decision task, and the navigation strategy N_H as the selective reasoning procedure used to navigate within this representation as . For example, consider a sequential purchasing task, the query Q corresponds to the written instructions shown to a participant in a given decision round, such as describing the product, the observed wait time, and the payoff structure. The context C_H captures how the participant interprets this information. For example, whether they understand wait time as a potential signal of product quality or as merely a cost. The navigation strategy N_H reflects how the participant reasons within this context, such as whether they update beliefs about quality based on prior purchases or rely on a simple expected-value comparison.

While both contexts and navigation strategies can in general be viewed as draws from underlying distributions reflecting heterogeneity in interpretation and reasoning, in controlled experimental settings, these distributions are typically tightly concentrated around an intended task representation and solution pathway, which we denote by C_H and N_H . Given (Q, C_H, N_H) , the human response is generated by sampling from a distribution:

$$Y_H \sim f_H(Q, C_H, N_H),$$

where f_H is a random function that captures how the query, the context, and the navigation strategy jointly generate the observed human decision Y_H .

In practice, researchers use procedures such as piloting and manipulation checks to ensure that participants form the intended representation of the task rather than idiosyncratic ones. These procedures validate the context C_H constructed from Q and concentrate its distribution around the intended interpretation. Once this validated context is in place, the navigation strategy N_H tends

to follow systematically from the task representation, allowing human responses to reflect a largely shared understanding of the query Q .

Although human cognition and LLM prediction rely on different underlying mechanisms, both involve forming task representations and reasoning within those representations. Accordingly, the theory of human problem-solving provides a foundation for understanding when and why LLM social simulations misalign with human behavior, and for designing interventions to improve alignment.

3.2. Context Formation and Context Navigation

Motivated by the human problem-solving process (Newell and Simon 1972, Simon and Newell 1971), we propose a two-stage framework designed to diagnose the root causes of misalignment in LLM simulations of human behaviors and to effectively improve human-LLM behavioral alignment. This framework identifies and addresses two key challenges that could cause misalignment: *context formation* and *context navigation*. *Context formation* refers to how LLMs form an accurate task representation, whereas *context navigation* refers to how LLMs reason within that context to generate decisions consistent with human behaviors. Once we specify the query Q , the context C , and the navigation N , an LLM will generate the response as

$$Y \sim f_{\text{LLM}}(Q, C, N),$$

where f_{LLM} is a random function powered by LLM that generates output based on the input prompt of the triplet (Q, C, N) .

SOTA LLMs such as GPT-5² and DeepSeek-R1 are aligned with human behaviors in a wide range of tasks such as math, coding, and commonsense Q&As (e.g., Guo et al. 2025). It is therefore reasonable to assume that f_{LLM} well approximates f_{H} with the same input (Q, C, N) . In particular, if the LLM has the same input as the human, $Y \sim f_{\text{LLM}}(Q, C_{\text{H}}, N_{\text{H}})$ will follow a similar distribution to the human response $Y_{\text{H}} \sim f_{\text{H}}(Q, C_{\text{H}}, N_{\text{H}})$. In this case, f_{LLM} is a simulator that could credibly replicate human behaviors regarding query Q . This assumption has also been validated by various successful LLM-based social simulations (e.g., Cui et al. 2024, Kirshner 2024, Mei et al. 2024).

² OpenAI. (2025). Introducing GPT-5. Retrieved from <https://openai.com/index/introducing-gpt-5/>.

However, if only the query Q is provided to the LLM but the context C or the navigation N is missing, the model will automatically infer such information from its prior knowledge. We refer to this setting as the *baseline* case. More specifically, the LLM, acting as a simulator, will first sample the context C_{Base} based on the query Q , then sample the navigation N_{Base} based on the query and the context (Q, C_{Base}) , and finally generate the outcome Y_{Base} based on the query and the sampled context and navigation, i.e.,

$$C_{\text{Base}} \sim \mathbb{P}_{\text{LLM}}(C|Q), \quad N_{\text{Base}} \sim \mathbb{P}_{\text{LLM}}(N|Q, C_{\text{Base}}), \quad Y_{\text{Base}} \sim f_{\text{LLM}}(Q, C_{\text{Base}}, N_{\text{Base}}), \quad (1)$$

where $\mathbb{P}_{\text{LLM}}(C|Q)$ is the distribution of C conditioned on Q and $\mathbb{P}_{\text{LLM}}(N|Q, C)$ is the distribution of N conditioned on (Q, C) , both of which capture the prior knowledge of the pretrained model rather than the specific problem setting. In general, the sampled context and navigation $(C_{\text{Base}}, N_{\text{Base}})$ deviate from those in the real human setting $(C_{\text{H}}, N_{\text{H}})$. As a consequence, Y_{Base} is unlikely to be aligned with Y_{H} .

The setting in which the query Q and the accurate context humans perceive C_{H} are specified, while the navigation strategy remains unspecified, is referred to as the *context formation* case. In this case, the model will infer the navigation strategy, N_{CF} based on (Q, C_{H}) , and generate the outcome Y_{CF} based on $(Q, C_{\text{H}}, N_{\text{CF}})$, i.e.,

$$N_{\text{CF}} \sim \mathbb{P}_{\text{LLM}}(N | Q, C_{\text{H}}), \quad Y_{\text{CF}} \sim f_{\text{LLM}}(Q, C_{\text{H}}, N_{\text{CF}}). \quad (2)$$

Compared with the *baseline* case, context formation improves alignment via specifying the accurate context C_{H} . Yet, misalignment may still occur if the LLM-generated navigation strategy N_{CF} deviates from the human navigation N_{H} , thus yielding the simulated outcome Y_{CF} misaligned with the human response Y_{H} .

Finally, we consider the setting where the query Q , the accurate context humans perceive C_{H} , and the human navigation strategy N_{H} are specified as model input, which we refer to as the *context formation and navigation* case. In this case, the LLM will generate the outcome

$$Y_{\text{CFN}} \sim f_{\text{LLM}}(Q, C_{\text{H}}, N_{\text{H}}). \quad (3)$$

Because f_{LLM} is close to f_{H} , the simulated outcome under *context formation and navigation*, Y_{CFN} , should follow a similar distribution as the human outcome Y_{H} . For the rest of this paper, we will empirically demonstrate that for complex decision-making in multi-agent environments with strategic interdependence and endogenously generated beliefs, simulations based on *context formation and navigation* could robustly generate outcomes aligned with human behaviors.

4. Simulation and Empirical Analysis

To apply and empirically validate our proposed framework, we implement it across experimental settings drawn from complex decision-making environments characterized by strategic interdependence and endogenous belief formation among multiple agents. [Kremer and Debo \(2016\)](#) provide a representative example of such an environment. We choose it as our focal case for two reasons. First, its main findings have been validated through human replications in the *Management Science* Replication Project ([Davis et al. 2023](#)), which offers open-access data, instructions, and protocols that ensure transparency and comparability. Second, the study has recently been used in other LLM-simulated replication attempts (e.g., [Kirshner 2024](#)).

4.1. Overview of the Experimental Setting

To provide the context for our framework application and evaluation, we briefly review the experimental setting of [Kremer and Debo \(2016\)](#). The paper examines how uninformed consumers infer product quality from queue length in a sequential purchasing environment. In this setting, uninformed consumers rely on observed wait times as product quality signals because queue length reflects the purchasing behavior of informed consumers. The study reports a main finding that the presence of informed consumers affects the impact of queue length on purchase decisions of uninformed consumers.

In the experiment, groups of four participants enter a market where a product is either of high value (\$3.50) or of low value (\$0), each with 50% probability. Participants are randomly assigned either to be informed of the true product value or to remain uninformed with 50% probability. Purchase decisions are made sequentially and, before making a decision, each participant observes the delivery

wait time. If a participant decides to purchase, their delivery wait time equals one production period plus the number of participants who have purchased before them. Each participant's payoff equals the \$4-endowment plus the product value (if purchased), minus a waiting cost of \$1 for each unit of wait time.

While [Kremer and Debo \(2016\)](#) originally study additional treatment conditions, following [Davis et al. \(2023\)](#), we focus on the comparison between Q00 (no informed participants) and Q50 (50% of participants are informed) conditions. In condition Q00, all participants are uninformed, so wait times carry no information about the product quality. In condition Q50, half of the participants are informed, so that their purchasing decisions make wait time an informative signal for uninformed participants to infer quality.

Following [Davis et al. \(2023\)](#), we examine [Kremer and Debo \(2016\)](#)'s finding that uninformed consumers interpret queue length as a quality signal using two hypothesis tests based on probit regressions. The first hypothesis (H1) examines whether uninformed consumers are less likely to purchase when they face an empty queue ($w = 1$) in the presence of informed consumers. Here, w denotes the delivery wait time, defined as the number of periods until the product is delivered; thus $w = 1$ corresponds to an empty queue with no prior purchases. We estimate the following probit model:

$$\mathbb{P}(\text{Purchase} = 1) = \Phi \left(\beta_0 + \sum_{w \in \{2,3,4\}} \beta_w D_w + \sum_{w \in \{1,2,3,4\}} \gamma_w D_w \times D_{Q50} \right) \quad (4)$$

where, $\Phi(\cdot)$ is the CDF of the standard normal distribution, D_w is a dummy variable equal to 1 if the delivery wait time is $w \in \{1, 2, 3, 4\}$ and 0 otherwise, with $w = 1$ omitted as the baseline category, D_{Q50} is a dummy variable for condition Q50, and $D_w \times D_{Q50}$ captures the interaction between each wait-time dummy $w \in \{1, 2, 3, 4\}$ and condition Q50. In this specification, the coefficient γ_1 on $D_1 \times D_{Q50}$ captures the change in purchase probability at $w = 1$ when the condition changes from Q00 to Q50. Accordingly, H1 is empirically supported when γ_1 is negative and statistically significant.

The second hypothesis (H2) examines whether the negative effect of wait time on purchase is reduced when informed consumers are present. We estimate the interaction effect between wait time w and treatment condition Q50 in the following probit model:

$$\mathbb{P}(\text{Purchase} = 1) = \Phi(\beta_0 + \beta_1 w + \beta_2 D_{Q50} + \beta_3 D_{Q50} \times w). \quad (5)$$

where w is the delivery wait time, D_{Q50} is a dummy variable for condition Q50, and $D_{Q50} \times w$ captures the interaction between wait time and condition Q50. Accordingly, H2 is empirically supported when β_3 is positive and statistically significant.

4.2. Simulation Design

To replicate the experimental setting of [Kremer and Debo \(2016\)](#) using LLM, we simulate cohorts of agents powered by OpenAI’s GPT-4o snapshot from August 6th, 2024. The temperature parameter controlling stochasticity in the LLM output is set to its default value 1, allowing for variation across simulated samples. The agents do not share memory between each other. We adopt LangChain,³ an open-source framework for building agent-based systems, to instantiate and manage the LLM-simulated subjects.

[Kremer and Debo \(2016\)](#)’s original sample consists of 100 students assigned to 25 cohorts of 4 (8 cohorts in condition Q00 and 17 cohorts in condition Q50), each cohort completing 26 sequential decision-making rounds. [Kremer and Debo \(2016\)](#) demonstrate that learning effects and round dependencies are minimal in their original study, so we implement the simulation as a between-subjects design with independently generated one-shot decisions. This allows us to take advantage of the LLM’s ability to generate independent responses at scale while avoiding context-window overflow. We simulate 208 (8×26) cohorts in condition Q00 and 442 (17×26) cohorts in condition Q50, reproducing the same total number of experimental units as [Kremer and Debo \(2016\)](#).

Within each cohort, LLM-simulated subjects are randomly assigned to positions in the decision sequence and to informed or uninformed roles. Once a subject makes a decision, the purchase outcome

³ <https://www.langchain.com/>

is recorded and determines the wait time presented to the next subject in that cohort, preserving the endogenous evolution of the queue. Across cohorts, simulations are conducted independently. Each run uses new role assignments and initializations for all subjects. Our approach preserves both the total number of observations and the causal structure of the original experiment, including the sequential nature of decisions and the endogenous formation of the wait line.

4.3. Prompt Design

Guided by our theoretical framework, we construct three prompt designs that vary in whether they include (i) neither *context formation* nor *context navigation*, which will generate the outcome Y_{Base} by Eqn. (1), (ii) *context formation* only, which will generate the outcome Y_{CF} by Eqn (2), or (iii) both components, which will generate the outcome Y_{CFN} by Eqn. (3). Figure 1 panels (a), (b), and (c) present the prompt design of the *baseline*, *context formation*, and *context formation and navigation* designs, respectively. These variants operationalize the two-stage structure of our framework and allow us to see how each stage contributes to alignment with human behavior.

(a) Baseline Design	(b) Context Formation Design	(c) Context Formation and Navigation Design
<p>User: (a.i) <General experimental instruction></p> <p>AI: (a.ii) Understood, I assume I am an experiment participant and I will make a decision whether to ORDER or NOT ORDER.</p> <p>User: (a.iii) <Detailed experimental instruction> (a.iv) <CoT reasoning prompt></p> <p>AI: <Generated Response></p>	<p>User: (b.i) <Context formation prompt> (b.ii) <General experimental instruction></p> <p>AI: (b.iii) Understood, I assume I am an experiment participant and I will make a decision whether to ORDER or NOT ORDER.</p> <p>User: (b.iv) <Detailed experimental instruction> (b.v) <CoT reasoning prompt></p> <p>AI: <Generated Response></p>	<p>User: (c.i) <Context formation prompt> (c.ii) <General experimental instruction></p> <p>AI: (c.iii) Understood, I assume I am an experiment participant and I will make a decision whether to ORDER or NOT ORDER.</p> <p>User: (c.iv) <Detailed experimental instruction> (c.v) <Context navigation prompt> (c.vi) <CoT reasoning prompt></p> <p>AI: <Generated Response></p>

Figure 1 Prompting Methods for Baseline, Context Formation, and Context Formation and Navigation

Figure 1(a) reveals the *baseline* design. Each LLM-simulated subject first receives a general experiment instruction adapted from the experiment instruction of Kremer and Debo (2016) as part of the

initialization, i.e., Figure 1(a.i). Then, we append a brief compliance confirmation message to ensure the agent’s behavior is under control (Figure 1(a.ii), see also Goli and Singh 2024). In each decision round, we append the detailed experimental instruction (e.g., wait time and, when applicable, product value) and a standard CoT reasoning prompt (Prompt 1; see also Kojima et al. 2022, Wei et al. 2022), as shown in Figure 1 (a.iii) and (a.iv), respectively. This prompting setup reflects a common approach in attempts to simulate human behavior with LLMs (Goli and Singh 2024, Kirshner 2024).

Prompt 1: CoT reasoning prompt for Kremer and Debo (2016)

Put your step-by-step reasoning in <reason></reason> and the final decision (ORDER/NOT ORDER) in <answer></answer>.

Guided by our framework, the *context formation* design builds on the *baseline* design by adding the *context formation* prompt (Prompt 2) at the beginning as shown in Figure 1 (b.i), which explicitly describes the two experimental conditions to improve how LLMs understand the context of the task (C_H in Eqn. (2)). While human participants are typically not informed of all experimental conditions in order to preserve ecological validity, leaving the experimental design implicit can be problematic in LLM social simulations. In particular, when the experimental design is not explicitly specified, LLM-simulated subjects may misinterpret the task context, leading to behavioral misalignment (Gu and Toubia 2025). Accordingly, to make the context more transparent, we incorporate the *context formation* prompt below into the full prompts presented to all LLM-simulated subjects, regardless of the conditions they are assigned to.

Prompt 2: Context formation prompt generated for Kremer and Debo (2016)

This study consists of two different conditions, and each participant will be assigned to one of them at random. You will not switch between conditions during the study; you will only participate in one of the two conditions. The two conditions differ in the information available to consumers before they make their purchasing decisions.

- Condition 1 (Q00): In this condition, all consumers are uninformed about the true value of the product before making a purchase decision. Each product can either be a high-value product (\$3.50) or a low-value product (\$0) with equal probability (50%). However, consumers will not know whether the product is high-value or low-value at the time of ordering.
- Condition 2 (Q50): In this condition, some consumers are informed while others remain uninformed. If you are assigned to this condition, you will first learn whether you are an informed or uninformed consumer. Informed consumers will know the product's true value before making a purchase decision, while uninformed consumers will not.

You are assigned to condition Q50/Q00. The following sections will provide the detailed instructions.

The *context formation and navigation* design includes both components, as shown by Figure 1(c). It supports *context formation* by clarifying the experimental setup at the beginning (Figure 1(c.i)) and *context navigation* by adding an instruction in each round that directs the model to selectively reason within the given context (N_H in Eqn. (3), Figure 1(c.v), and Prompt 3).

Prompt 3: Context navigation prompt generated for Kremer and Debo (2016)

Reason in the following approach: Hold an initial belief about what is unobserved, consider what the observable cue suggests, update your belief with a mix of reasoning and intuition, and make your decision accordingly.

4.4. Simulation Results

Following Davis et al. (2023), we evaluate replication along two hypotheses: H1 assesses whether uninformed consumers are less likely to purchase in condition Q50 than in condition Q00 when $w = 1$, and H2 tests whether the negative effect of wait time on purchasing is attenuated in condition Q50 relative to condition Q00. Based on our theoretical framework, we expect that the *baseline* design is unlikely to reproduce these human behavioral patterns, the *context formation* design can improve

alignment, and the *context formation and navigation* design is expected to fully replicate the original results.

We present the simulation results in Table 1. Under our *baseline* design, which uses the original human instructions together with a standard CoT reasoning prompt, we fail to replicate either H1 or H2. For H1, the estimated γ_1 is marginally significant but in the wrong direction ($p = .059$, $\hat{\gamma}_1 = 0.59$). For H2, the interaction term between wait time and Q50 condition, β_3 , is statistically significant but negative ($p = .05$, $\hat{\beta}_3 = -0.97$). The CoT reasoning traces show that the uninformed LLM-simulated subjects in both conditions follow a simple payoff-comparison rule without accounting for the quality-signaling effect of queue length. They calculate the expected value of ordering as $EV = 4 - \text{wait time} + 1.75$, where 1.75 is the expected product value under the 50/50 prior and 4 is the endowment, and compare this expected value directly with the guaranteed endowment from not ordering. Accordingly, subjects in both conditions make identical decisions across wait times. Purchasing is only preferred when $EV > 4$, which implies $w < 1.75$. As a result, subjects purchase only at the minimal wait time ($w = 1$) and not at longer wait times ($w > 1$), regardless of the condition they are in. The LLM-simulated subjects do not refer to others' actions or to informational cues of the wait time in reasoning, which explains the joint replication failure of H1 and H2.

The *context formation* design replicates H1 in the statistical sense ($p = 0.04$, $\hat{\gamma}_1 = -0.77$), implying the lower likelihood of purchase in Q50 than in Q00 for $w = 1$. In terms of H2, LLMs stop purchasing entirely once $w > 1$ in both conditions. This produces no variations in purchasing for $w > 1$, making it impossible to estimate how the effect of wait time differs across conditions. Although this prevents a formal statistical test of H2, the pattern already diverges from human behavior, where the purchase rate decreases with wait time but remains strictly above zero in both conditions.

A closer look at the CoT reasoning traces under *context formation* reveals additional insights. Uninformed LLM-simulated subjects, albeit with improved awareness of their own uninformed states, fail to recognize that some consumers are informed, so that no one waiting (i.e., $w = 1$) signals the low quality of the product. Uninformed consumers explicitly note their own status and become

more risk-averse if they know there are informed ones on the market, leading to a lower purchase likelihood under Q50. Furthermore, their decision rules remain the same as the *baseline* case, i.e., the subjects directly compute the expected value of ordering and compare it with the endowment, without drawing any inference from the observed wait time. Because they do not interpret the queue length as conveying information about quality, their decisions remain constant (i.e., not purchase) for all $w > 1$. This lack of variation not only prevents the estimation of β_3 for H2, but also reflects behavioral misalignment, as humans continue to make purchases at longer waits. Thus, the *context formation* design improves context awareness of the task but does not induce the model to infer quality from wait time, leaving the hypotheses unreplicated.

The *context formation and navigation* design successfully replicates both H1 and H2 (H1: $p < 0.01$, $\hat{\gamma}_1 = -0.42$; H2: $p < 0.01$, $\hat{\beta}_3 = 1.60$), showing that combining *context formation* and *context navigation* achieves alignment. Under the *context formation and navigation* design, the uninformed LLM-simulated subjects not only recognize the context but also reason about product quality based on wait time. The CoT reasoning traces show that the LLM-simulated subjects refer to earlier consumers' actions as cues to adjust their beliefs about product quality before deciding whether to purchase. As a consequence, the model moves from a fixed payoff-comparison rule purely based on expected-value calculation to a context-sensitive reasoning process in a dynamic environment.

This improved navigation of the context yields LLM behaviors more closely aligned with human decisions. The LLM-simulated subjects order less often at short waits when no one purchases, replicating the lower initial purchasing rate in condition Q50 (H1). They also maintain moderate purchasing at longer waits when earlier purchases suggest high product quality, reproducing the attenuated wait-time effect observed in human experiments (H2). Thus, the *context formation and navigation* design not only improves the model's context awareness, but also guides its reasoning within that context, successfully replicating both hypotheses in [Kremer and Debo \(2016\)](#).

A potential concern of our results is that successful replication simply reflects memorization, i.e., the original study may appear in the GPT-4o's training data. If memorization alone were driving

replication, the baseline design, using the same underlying model but without additional prompt design, should also reproduce the original results of Kremer and Debo (2016). The failure of the *baseline* design, therefore, indicates that replication success under the *context formation and navigation* design should be attributed to our prompt design rather than memorization by LLM.

Our framework can be applied to different settings. For replicating existing studies, *context formation* prompts explicitly specify the experimental design based on the original instructions (e.g., treatment conditions, payoff structures, decision variables), ensuring all elements are represented accurately. *Context navigation* prompts can be derived from the study's core theoretical mechanism (e.g., the authors' explanation for how their treatment works). For new experimental studies not reported in the literature, *context formation* follows the same generation process, but *context navigation* requires an iterative approach. Because new experiments are designed to test whether a proposed economic or behavioral mechanism is valid, navigation prompts cannot directly specify how participants reason. Instead, they should be grounded in evidence from pilot studies (e.g., think-aloud protocols for humans) or the theoretical mechanism under investigation, and refined iteratively as empirical patterns emerge. Section 5 demonstrates the framework's application in two additional experimental settings.

Table 1 Human and LLM Replications of Kremer and Debo (2016)

Paper/Prompt Design	Subject Type	Hypothesis 1				Hypothesis 2			
		γ_1	χ^2	p-value	Replication Success	β_3	χ^2	p-value	Replication Success
Kremer and Debo (2016)	Human	-1.14	31.23	< 0.01	—	0.85	74.34	< 0.01	—
Davis et al. (2023) (South Carolina)	Human	-0.84	22.72	< 0.01	✓	0.22	5.35	0.021	✓
Davis et al. (2023) (Michigan)	Human	-1.09	47.94	< 0.01	✓	0.69	57.72	< 0.01	✓
Baseline	LLM	0.59	3.79	0.05	✗	-0.97	4.02	0.05	✗
Context Formation	LLM	-0.77	4.43	0.04	✓	—	—	—	✗
Context Formation and Navigation	LLM	-0.42	7.22	< 0.01	✓	1.60	43.53	< 0.01	✓

Note: For simulation results of the *context formation* design, probit regression for Hypothesis 2 can not be performed due to uniform behavior: LLMs ceased purchasing entirely at any wait time > 1 in both conditions, making wait time a perfect predictor of purchase decision and thus producing no variation for estimating the interaction effect. This rigid behavior contrasts with the gradual decline in purchasing observed in human data, suggesting that *context formation* alone fails to simulate human-like variability in decision-making.

5. Applications of the Theoretical Framework

In this section, we apply the two-stage framework beyond Kremer and Debo (2016) to evaluate its generalizability across distinct decision-making environments and across LLM models. Section 5.1

examines a documented LLM alignment failure in consumer demand estimation reported by [Gui and Toubia \(2025\)](#), serving as a benchmark for assessing the framework in a simple decision task that lacks the strategic interdependence and endogenous belief formation of our focal applications. In Section 5.2, we extend our framework to a crowdfunding game characterized by strategic interdependence and endogenously generated beliefs from [Cason et al. \(2025\)](#). Section 5.3 tests whether the patterns observed in our three experiments generalize across several SOTA LLMs by applying a common prompt design to each model and comparing their alignment outcomes.

5.1. Application to [Gui and Toubia \(2025\)](#): Demand Estimation with Price Experiments

We begin by applying our theoretical framework to the demand-estimation experiment of [Gui and Toubia \(2025\)](#), a simple decision-making environment in which LLM-simulated subjects nonetheless fail to reproduce the downward-sloping human demand curve. In this study, 1,000 Prolific respondents each evaluate 40 consumer packaged goods whose prices are randomly drawn from 0% to 200% of the product's regular price. For each product-price pair, participants answer a binary question ("Would you or would you not purchase this product?"), and their responses are aggregated to estimate purchase probabilities at each price level. Consistent with the law of demand, human purchase probabilities decrease monotonically with price.

In [Gui and Toubia \(2025\)](#)'s LLM benchmark case, they present GPT-4o-mini with the same product-price combinations and generate 50 independent purchase decisions for each combination. These simulated decisions are aggregated in the same way as the human responses to produce purchase probabilities at each price level. The resulting demand curve deviates sharply from the human benchmark. Rather than the purchase probabilities decreasing in price, the model produces an inverted-U shape curve in which purchase likelihood first increases and then decreases with price (Figure 2 in [Gui and Toubia 2025](#)). Such misalignment motivates applying our framework to this simple, non-strategic setting.

5.1.1. Prompt Design The main body of the *baseline* design follows the prompt design of [Gui and Toubia \(2025\)](#), with a CoT reasoning prompt appended at the end. The *context formation*

design is built on the *baseline* design by adding an LLM-generated introduction that clarifies the experimental setting at the beginning, following the same prompt-generation procedure introduced in Section 4.3.

Prompt 4: Context formation prompt generated for Gui and Toubia (2025)

In this study, each participant is presented with one of 40 different products, shown at a price drawn from 11 systematically varied levels ranging from 0% to 200% of its regular retail value. For each product-price scenario, participants provide a single binary response – whether they would purchase the product or not. You are assigned a survey question in which the price level of the product is manipulated.

The *context navigation* prompt guides the model to selectively reason within the given context by focusing on features that are relevant for decision-making.

Prompt 5: Context navigation prompt generated for Gui and Toubia (2025)

Reason in the following approach: Notice what stands out in the context, use it to guide your choice, and don't overcomplicate with unnecessary details.

5.1.2. Simulation Results To evaluate simulation performance, Figure 2 plots average purchase probabilities across relative price levels under each prompt design and compares the resulting demand curves to the human benchmark reported by Gui and Toubia (2025). Using the more advanced GPT-4o (rather than GPT-4o-mini as in Gui and Toubia 2025) and appending an additional CoT prompt to our *baseline* design, the *baseline* simulations reproduce the inverted-U pattern reported in Gui and Toubia (2025)'s LLM-based simulations. However, this simulation pattern diverges from the human benchmark, which exhibits a clearly downward-sloping relationship. Examination of the CoT reasoning traces suggests that, under the *baseline* design, LLM-simulated subjects interpret the zero-price level (0% of the regular price) as a potential input or pricing error and therefore choose not to purchase. This behavior indicates that the model focuses on the anomalous nature of a zero price rather than reasoning about price in a standard consumer demand context.

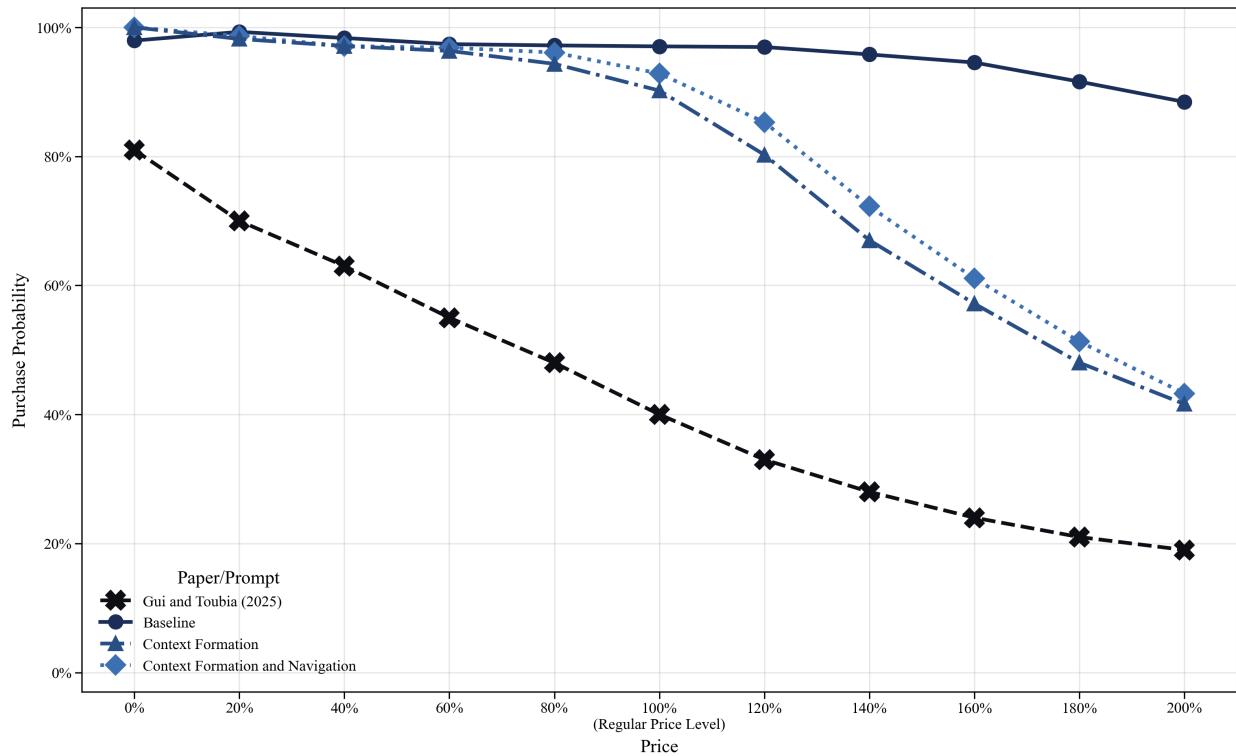


Figure 2 Demand curve elicited from LLMs (GPT-4o)

*Note: The “*Gui and Toubia (2025)*” line is an approximation of the human demand curve reported in *Gui and Toubia (2025, Fig. 2)* and is not reproduced from the original human survey data; small discrepancies may arise due to figure resolution.*

Meanwhile, the *context formation* design produces a steadily declining demand curve, with purchase probabilities starting close to 100% at the lowest prices and falling to around 40% at the highest. This reproduces a fundamental principle of consumer behavior: the law of demand. Examination of the CoT reasoning traces shows that the *context formation* design changes how the LLM-simulated subjects interpret prices. Under this design, a price of zero is interpreted as the product being genuinely free rather than as a suspicious abnormality with hidden conditions, which in turn yields a downward-sloping demand curve aligned with human behavior in the original study. By contrast, while the *context formation and navigation* design also manages to generate a downward-sloping demand curve, it does not yield clear improvements over the *context formation* design alone.

These results indicate that, in the demand estimation setting, *context formation* is sufficient to replicate human behavior. Because the task is simple, featuring binary purchase decisions with price as the only manipulated variable and no strategic interaction or belief formation, the model does

not require explicit guidance on how to reason through the decision process. Explicitly framing the randomized design through *context formation* is therefore sufficient for the model to arrive at the appropriate reasoning process on its own. This contrasts with the [Kremer and Debo \(2016\)](#) replication, where the strategic interdependence and endogenous belief formation required both *context formation* and *context navigation* to achieve behavioral alignment. It is worth noting that [Gui and Toubia \(2025\)](#) also explores an alternative simulation design to align with human decision patterns. Both [Gui and Toubia \(2025\)](#)’s design and our *context formation* design aim to enhance LLMs’ understanding of the problem by making price manipulation and variation explicit. In essence, unblinding is an implementation of *context formation* in our framework.

5.2. Application to [Cason et al. \(2025\)](#): Refund Bonus and Information Asymmetry in Crowdfunding

In this section, we extend our framework to the crowdfunding experiment of [Cason et al. \(2025\)](#), which examines how refund bonuses mitigate information asymmetries in a strategically interdependent investment environment. The setting features two inventors and two investors interacting in a repeated game. In each round, each inventor offers a single invention of fixed quality, either good or bad, while each investor, endowed with \$2, decides independently whether to invest \$1 in each of the two inventions.

Inventors observe their own invention’s quality, whereas investors do not. Invention quality is good with prior probability 0.4286. For bad inventions, investors receive a negative report with 25% probability that fully reveals low quality. If no negative report is received, the invention is equally likely to be good or bad. An invention is funded only if both investors invest. If funded, good inventions yield \$3 per investor, while bad inventions result in a loss of the \$1 investment. Uninvested funds can be placed in a safe outside option yielding \$1.50.

The experiment studies whether refund bonuses, payments promised to investors if an invention fails to reach the funding threshold, can serve as a signal of quality. Specifically, if only one investor invests in an invention, that investor receives a \$1 refund bonus in addition to the returned \$1 investment. Because bad inventions are more likely to trigger bonus payments, offering a refund bonus

is more costly for low-quality inventions, making bonuses a potential separating signal. Experimental conditions differ in how bonuses are determined. In the Endogenous Bonus (EB) condition, inventors choose whether to offer a bonus after observing quality. In the Random Bonus (RB) condition, bonuses are assigned randomly and independently of quality. The main finding is that good inventions are successfully funded more often in the EB condition than in the RB condition.

Cason et al. (2025)'s original design assigns 168 subjects to 14 matching groups of 12, with 8 groups in the EB condition and 6 groups in the RB condition, where each subject repeatedly performs decision-making for 40 rounds. Taking advantage of the LLM's ability to simulate subjects at scale, we modify the design by increasing the number of subjects and reducing the number of decision rounds per subject. We simulate 40 independent groups of 12 subjects in each condition, with each subject completing 8 decision rounds.

To evaluate whether the main finding of Cason et al. (2025) can be replicated, we compare funding success for good-quality projects that offer refund bonus in the EB and RB conditions using a Wilcoxon rank-sum test. We treat whether funding success is systematically higher the EB condition than in the RB condition. A positive and statistically significant z-statistic indicates successful replication of the main finding.

5.2.1. Prompt Design We use the same approach to construct prompt design as in the previous Kremer and Debo (2016) replication. The *baseline* prompt consists of the original experiment instruction and a CoT reasoning prompt added at the very end. Consistent with the approach introduced in Section 4.3, the *context formation* design is again implemented by adding an LLM-generated overview of the experimental conditions.

Prompt 6: Context formation prompt generated for Cason et al. (2025)

In this experiment, you will participate in a series of decision-making rounds where people take on the roles of investors and inventors. Each round involves two inventions that may receive funding if both investors choose to support them. Investors decide where to invest based on limited information, and

inventors may have the option to offer a refund bonus--a payment to compensate investors if the invention fails to get enough support.

There are two experimental conditions that differ in how refund bonuses are determined. In the Endogenous Bonus (EB) condition, inventors choose whether to offer a refund bonus after learning the quality of their invention. In the Random Bonus (RB) condition, refund bonuses are assigned randomly with a 50% chance, regardless of invention quality. You have been randomly assigned to one of these two conditions.

You are assigned to the Random/Endogenous Bonus condition. The following sections will provide the detailed instructions.

Because Cason et al. (2025) studies a strategic costly-signaling game with asymmetric information, the *context navigation* prompt functions as an instruction that structures how the model reasons within the given context.

Prompt 7: Context navigation prompt generated for Cason et al. (2025)

Use the following approach: explicitly state priors, likelihoods, compute the posterior via Bayes' rule, then compare expected values of both actions before deciding.

5.2.2. Simulation Results Table 2 reports the replication performance of the main result in Cason et al. (2025) across the three prompt designs. Consistent with our earlier replication of Kremer and Debo (2016), neither the *baseline* nor the *context formation* designs could replicate the main result. In contrast, the *context formation and navigation* design successfully replicate the behaviors observed in human experiments. Furthermore, we investigate the qualitative patterns in the CoT reasoning traces for each prompt design, offering insight into how different prompt designs affect simulation results in this complex decision-making environment.

Under the *baseline* design, we do not replicate the main finding. Table 2 shows a statistically significant effect in the opposite direction ($z = -4.27, p < 0.01$): good inventions are funded less often

in the EB condition than in the RB condition. The CoT reasoning traces suggest that this reversal is driven by inventors' overreliance on the negative report. Many inventors infer, incorrectly, that good inventions never generate negative reports and therefore will attract more investors without needing a refund bonus. As a result, they treat the report, rather than the refund bonus, as the primary determinant of funding and largely ignore the bonus as an incentive mechanism. In parallel, investors in both EB and RB conditions apply a similarly shallow heuristic: they compare the payoff values in the instructions and interpret the refund bonus as simply additional money, rather than using it to infer underlying invention quality. Consistent with these patterns, inventors often view the refund bonus as a potential cost and thus rarely offer it, yielding a low bonus-offer rate for good inventions in the EB condition, whereas in the RB condition good inventions receive bonuses at the exogenously assigned rate of 50%. Because investors respond positively to bonuses regardless of their source, good inventions receive more bonus offers, and therefore higher funding rates, in RB than in EB, producing the reversed result in the simulation.

Likewise, the *context formation* design still yields a statistically significant effect in the wrong direction ($z = -5.06, p < 0.01$). Although this design clarifies how refund bonuses are determined across conditions, it does not meaningfully change the underlying reasoning. Inventors remain anchored to the negative report signal and continue to neglect the refund bonus as an incentive mechanism. Investors likewise show no deeper inference, treating bonuses simply as higher payoffs. Consequently, both the reasoning patterns and behavioral outcomes under this design are largely indistinguishable from those observed under the *baseline* case.

With the *context formation and navigation* design, we successfully replicate the main finding of Cason et al. (2025) ($z = 2.45, p = 0.01$). The *context navigation* prompt directs attention to the key relationships among invention quality, the refund bonus decision, and investor response. Inventors become less anchored to the negative report signal and begin to view the refund bonus as a mechanism for influencing investor behavior. In parallel, both inventors and investors shift from comparing static payoff entries to evaluating expected payoffs. This change mitigates the earlier distraction

induced by the negative report and yields investment patterns consistent with the human benchmark, reproducing the result that good inventions are funded more often in the EB condition than in the RB condition.

Taken together, our results indicate that both *context formation* and *context navigation* are necessary to replicate Cason et al. (2025). In contrast to the simpler demand-estimation setting, where *context formation* design alone suffices, the strategic interdependence and belief-driven nature of crowdfunding requires both components to align the LLM-simulated subjects' behavior with human outcomes.

Table 2 LLM Replication Results for Cason et al. (2025)

Paper/Prompt Design	Subject Type	<i>z</i>	<i>p</i> -value	Replication Success
Cason et al. (2025)	Human	3.10	< 0.01	—
Baseline	LLM	-4.27	< 0.01	✗
Context Formation	LLM	-5.06	< 0.01	✗
Context Formation and Navigation	LLM	2.45	0.01	✓

Note: *z* statistics and asymptotic *p*-values are based on the tests of the corresponding results reported in Cason et al. (2025) and in the LLM simulations.

5.3. The Impact of LLM Choice

In this section, we demonstrate the generalizability and robustness of our framework by implementing it with different LLMs to simulate the three experiments studied in this paper. In particular, we seek to uncover insights on how SOTA LLMs with reasoning capability may affect the validity and performance of our framework. Specifically, we test GPT-5 snapshot from Aug. 7th 2025, Claude 4.0-Sonnet-Thinking from May 14th, 2025, and DeepSeek-R1 from May 28th, 2025 under the same experimental settings, simulation procedures, and evaluation criteria while setting all their temperature parameters to the default value 1. We report the simulation results of all models and experiments in Table 3.

We highlight three key observations. First, across all models and all three studies, the *context formation and navigation* design successfully replicates the original human results. This pattern suggests that providing both *context formation* and *context navigation* reliably aligns LLM behavior with human data, regardless of the model's baseline reasoning ability. Second, models with stronger

reasoning abilities often succeed with *context formation* alone. Such models (e.g., GPT-5) benefit from process-level supervision and verifier-style reinforcement learning that encourages more structured and internally coherent reasoning (e.g., Guo et al. 2025), reducing their reliance on explicit navigation instructions and enabling them to effectively operate within the task representation supplied by *context formation*. This pattern indicates that, when the context is made explicit, stronger models can often supply the appropriate navigation path endogenously. Finally, we observe a statistical replication of Cason et al. (2025) under the *baseline* design with Claude-4.0-Sonnet-Thinking. However, its CoT reasoning traces suggest that the underlying reasoning still diverges from human decision processes in this setting, implying that *context formation* remains necessary even for strong reasoning models.

Our cross-model results validate that our two-stage framework aligns LLM behaviors with human decisions when both *context formation* and *context navigation* are provided, while also showing that the need for *context formation* depends on the model’s reasoning ability. Nonetheless, even strong reasoning models sometimes require both stages to reproduce human patterns. Overall, the framework both improves behavioral alignment and clarifies when external context engineering is necessary versus when the model’s internal reasoning suffices. More broadly, these implications connect to the industry notion of context engineering: shaping and structuring the information in an LLM’s context window to guide task processing and reasoning. Our results suggest that as models’ reasoning strength increases, alignment depends increasingly on the quality of the engineered contextual representation, which is implemented in our study through *context formation*.

6. Conclusion

This paper develops and tests a two-stage framework for improving behavioral alignment in LLM social simulations by separating *context formation* (making the task representation explicit) from *context navigation* (guiding how the model reasons within that representation). Across three established empirical settings and multiple frontier models, the results show a consistent pattern: providing both stages reliably recovers human behavioral findings in strategically interdependent, belief-driven

Table 3 Overview of Cross-Model and Cross-Design Replication Results

Study	Model Name	Baseline	Context Formation	Context Formation and Navigation
Kremer and Debo (2016)	GPT-4o	✗	✗	✓
	GPT-5	✗	✓	✓
	Claude-4.0-Sonnet-Thinking	✗	✗	✓
	DeepSeek-R1	✗	✓	✓
Gui & Toubia (2025)	GPT-4o-mini	✗	✓	✓
	GPT-4o	✗	✓	✓
	GPT-5	✗	✓	✓
	Claude-4.0-Sonnet-Thinking	✗	✓	✓
	DeepSeek-R1	✗	✓	✓
Cason et al. (2025)	GPT-4o	✗	✗	✓
	GPT-5	✗	✓	✓
	Claude-4.0-Sonnet-Thinking	✓	✓	✓
	DeepSeek-R1	✗	✓	✓

environments, while stronger reasoning models can often succeed with *context formation* alone when the task is simpler. The framework thus offers practical, theory-grounded “context engineering” guidance for designing, diagnosing, and evaluating LLM-simulated subjects as credible substitutes or complements to human subjects in complex decision-making environments.

The paper’s main limitations are scope and mechanism. Empirically, it tests only three behavioral experiments, which may not span the broader space of dynamic, organizational, or information-rich settings where misalignment can manifest differently. Methodologically, it focuses on prompt-based interventions rather than training-time or architectural approaches (e.g., fine-tuning or RLHF-style methods), which may limit how broadly the framework transfers beyond prompting. Finally, the evaluation cannot fully keep pace with the most recent model releases. While our validation uses published studies that may appear in training data, our framework provides systematic guidance for future novel experiments where memorization is impossible.

Future work can extend this framework to richer socio-technical environments, e.g., multi-step decisions, group settings, and organizational workflows, where context evolves and interaction histories matter. It can also leverage the framework for systematic hypothesis generation, pretesting, and stress-testing theories before deploying costly human-subject studies. Finally, an important next step is to translate the framework into an actionable recipe, a set of concrete design choices, diagnostics,

and reporting standards, that enables researchers and practitioners to apply *context formation* and *context navigation* systematically when building, auditing, and scaling LLM social simulations.⁴

References

Aher GV, Arriaga RI, Kalai AT (2023) Using large language models to simulate multiple humans and replicate human subject studies. *International Conference on Machine Learning*, 337–371 (PMLR).

Akata E, Schulz L, Coda-Forno J, Oh SJ, Bethge M, Schulz E (2025) Playing repeated games with large language models. *Nature Human Behaviour* 1–11.

Allouah A, Besbes O, Figueroa JD, Kanoria Y, Kumar A (2025) What is your ai agent buying? evaluation, implications and emerging questions for agentic e-commerce. *arXiv preprint arXiv:2508.02630* .

Anthis JR, Liu R, Richardson SM, Kozlowski AC, Koch B, Evans J, Brynjolfsson E, Bernstein M (2025) Llm social simulations are a promising research method. *arXiv preprint arXiv:2504.02234* .

Argyle LP, Busby EC, Fulda N, Gubler JR, Ryttig C, Wingate D (2023) Out of one, many: Using language models to simulate human samples. *Political Analysis* 31(3):337–351.

Arora N, Chakraborty I, Nishimura Y (2025) Ai–human hybrids for marketing research: Leveraging large language models (llms) as collaborators. *Journal of Marketing* 89(2):43–70.

Bail CA (2024) Can generative ai improve social science? *Proceedings of the National Academy of Sciences* 121(21):e2314021121.

Bauer K, Liebich L, Kosfeld M (2025) Can GPT mimic human preferences? An empirical and structural investigation. *ECIS 2025 Proceedings* (AIS Electronic Library (AISeL)), URL https://aisel.aisnet.org/ecis2025/ai_anthro/ai_anthro/4.

Bhaskar R, Simon HA (1977) Problem solving in semantically rich domains: An example from engineering thermodynamics. *Cognitive science* 1(2):193–215.

Binz M, Akata E, Bethge M, Brändle F, Callaway F, Coda-Forno J, Dayan P, Demircan C, Eckstein MK, Éltető N, et al. (2025) A foundation model to predict and capture human cognition. *Nature* 1–8.

Boussioux L, Lane JN, Zhang M, Jacimovic V, Lakhani KR (2024) The crowdless future? generative ai and creative problem-solving. *Organization Science* 35(5):1589–1607.

Brand J, Israeli A, Ngwe D (2024) Using gpt for market research. *Proceedings of the 25th ACM Conference on Economics and Computation*, 613–613.

⁴ An online appendix is available upon request.

Brookins P, DeBacker JM (2023) Playing games with gpt: what can we learn about a large language model from canonical strategic games? *Available at SSRN 4493398* .

Brynjolfsson E, Li D, Raymond L (2025) Generative ai at work. *The Quarterly Journal of Economics* 140(2):889–942.

Burtsch G, Lee D, Chen Z (2024) The consequences of generative ai for online knowledge communities. *Scientific Reports* 14(1):10413.

Carlson N, Burbano V (2024) The use of llms to annotate data in management research: Warnings, guidelines, and an application to organizational communication. *SSRN Electronic Journal* .

Cason TN, Tabarrok A, Zubrickas R (2025) Signaling quality: How refund bonuses can overcome information asymmetries in crowdfunding. *Management Science* 71(7):5933–5947.

Charness G, Jabarian B, List JA (2025) The next generation of experimental research with llms. *Nature Human Behaviour* 1–3.

Chen Y, Hu Y, Lu Y (2025a) Predicting field experiments with large language models. *arXiv preprint arXiv:2504.01167* .

Chen Y, Kirshner SN, Ovchinnikov A, Andiappan M, Jenkin T (2025b) A manager and an ai walk into a bar: does chatgpt make biased decisions like we do? *Manufacturing & Service Operations Management* 27(2):354–368.

Chen Y, Liu TX, Shan Y, Zhong S (2023) The emergence of economic rationality of gpt. *Proceedings of the National Academy of Sciences* 120(51):e2316205120.

Chen Z, Chan J (2024) Large language model in creative work: The role of collaboration modality and user expertise. *Management Science* 70(12):9101–9117.

Cheng X, Mayya R, Sedoc J (2024) From human annotation to llms: Silicon annotation workflow for management research. *arXiv preprint arXiv:2412.14461* .

Christiano PF, Leike J, Brown T, Martic M, Legg S, Amodei D (2017) Deep reinforcement learning from human preferences. *Advances in neural information processing systems* 30.

Chuang YS, Goyal A, Harlalka N, Suresh S, Hawkins R, Yang S, Shah D, Hu J, Rogers T (2024) Simulating opinion dynamics with networks of llm-based agents. *Findings of the association for computational linguistics: NAACL 2024*, 3326–3346.

Cillo P, Rubera G (2025) Generative ai in innovation and marketing processes: A roadmap of research opportunities. *Journal of the Academy of Marketing Science* 53(3):684–701.

Cui Z, Li N, Zhou H (2024) Can ai replace human subjects? a large-scale replication of psychological experiments with llms. *A Large-Scale Replication of Psychological Experiments with LLMs* (August 25, 2024) .

Davis AM, Flicker B, Hyndman K, Katok E, Keppler S, Leider S, Long X, Tong JD (2023) A replication study of operations management experiments in management science. *Management Science* 69(9):4977–4991.

Dell'Acqua F, McFowland III E, Mollick ER, Lifshitz-Assaf H, Kellogg K, Rajendran S, Krayer L, Candelon F, Lakhani KR (2023) Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (24-013).

Demirci O, Hannane J, Zhu X (2025) Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science* .

Demszky D, Yang D, Yeager DS, Bryan CJ, Clapper M, Chandhok S, Eichstaedt JC, Hecht C, Jamieson J, Johnson M, et al. (2023) Using large language models in psychology. *Nature Reviews Psychology* 2(11):688–701.

Dillion D, Tandon N, Gu Y, Gray K (2023) Can ai language models replace human participants? *Trends in Cognitive Sciences* 27(7):597–600.

Doshi AR, Bell JJ, Mirzayev E, Vanneste BS (2025) Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal* 46(3):583–610.

Doshi AR, Hauser OP (2024) Generative ai enhances individual creativity but reduces the collective diversity of novel content. *Science advances* 10(28):eadn5290.

Eloundou T, Manning S, Mishkin P, Rock D (2023) Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130* 10.

Filippas A, Horton JJ, Manning BS (2024) Large language models as simulated economic agents: What can we learn from homo silicus? *Proceedings of the 25th ACM Conference on Economics and Computation*, 614–615.

Gao Y, Lee D, Burtch G, Fazelpour S (2024) Take caution in using llms as human surrogates: Scylla ex machina. *arXiv preprint arXiv:2410.19599* .

Goli A, Singh A (2024) Frontiers: Can large language models capture human preferences? *Marketing Science* 43(4):709–722.

Gopal RD, Li J, Riemer K, Sarker S, Singh PV, Susarla A, Bichler M, Thatcher JB (2025) Inventing with machines: Generative ai and the evolving landscape of is research.

Grossmann I, Feinberg M, Parker DC, Christakis NA, Tetlock PE, Cunningham WA (2023) Ai and the transformation of social science research. *Science* 380(6650):1108–1109.

Gui G, Toubia O (2025) The challenge of using llms to simulate human behavior: A causal inference perspective. *arXiv preprint arXiv:2312.15524* .

Guo D, Yang D, Zhang H, Song J, Wang P, Zhu Q, Xu R, Zhang R, Ma S, Bi X, et al. (2025) Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature* 645(8081):633–638.

Guo E, Gupta M, Deng J, Park YJ, Paget M, Naugler C (2024) Automated paper screening for clinical reviews using large language models: data analysis study. *Journal of Medical Internet Research* 26:e48996.

Hackenburg K, Margetts H (2024) Evaluating the persuasive influence of political microtargeting with large language models. *Proceedings of the National Academy of Sciences* 121(24):e2403116121.

Hayes JR, Simon HA (1974) Understanding written problem instructions. Gregg LW, ed., *Knowledge and cognition*, 167–200 (Potomac, MD: Lawrence Erlbaum Associates).

Hou J, Wang L, Wang G, Wang HJ, Yang S (2025) The double-edged roles of generative ai in the creative process: Experiments on design work. *Information Systems Research* .

Hu L, Wu Z, Chen PY (2025) Balancing ai and human expertise: The role of ai watermarking and detection. *Available at SSRN 5270073* .

Hui X, Reshef O, Zhou L (2024) The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science* 35(6):1977–1989.

Kane GC, Karahanna E, Boudreau MC, Xue L, Schecter A, Li W, Salge C (2025) Augmented reasoning, transformed inquiry: Integrating generative ai into information systems methodologies. *Transformed Inquiry: Integrating Generative AI into Information Systems Methodologies (June 22, 2025)* .

Kim J, Lee B (2023) Ai-augmented surveys: Leveraging large language models and surveys for opinion prediction. *arXiv preprint arXiv:2305.09620* .

Kirshner S (2024) Artificial agents and operations management decision-making. *Available at SSRN 4726933* .

Kojima T, Gu SS, Reid M, Matsuo Y, Iwasawa Y (2022) Large language models are zero-shot reasoners. *Advances in neural information processing systems* 35:22199–22213.

Kremer M, Debo L (2016) Inferring quality from wait time. *Management Science* 62(10):3023–3038.

Lambert N, Morrison J, Pyatkin V, Huang S, Ivison H, Brahman F, Miranda LJV, Liu A, Dziri N, Lyu S, et al. (2024) Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124* .

Leng Y, Yuan Y (2023) Do llm agents exhibit social behavior? *arXiv preprint arXiv:2312.15198* .

Li P, Castelo N, Katona Z, Sarvary M (2024) Frontiers: Determining the validity of large language models for automated perceptual analysis. *Marketing Science* 43(2):254–266.

Liu C, Wang T, Yang SA (2025a) Generative ai and content homogenization: The case of digital marketing. *Available at SSRN 5367123* .

Liu R, Li J, Zivkovic M, Abbasi A (2025b) Automating in high-expertise, low-label environments: Evidence-based medicine by expert-augmented few-shot learning. *MIS Quarterly* 49(3):1049–1094.

Ludwig J, Mullainathan S, Rambachan A (2025) Large language models: An applied econometric framework. Technical report, National Bureau of Economic Research.

Manning BS, Horton JJ (2025) General social agents. *arXiv preprint arXiv:2508.17407* .

Manning BS, Zhu K, Horton JJ (2024) Automated social science: A structural causal model-based approach. *MIT and NBER Working Paper* .

Mei Q, Xie Y, Yuan W, Jackson MO (2024) A turing test of whether ai chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences* 121(9):e2313925121.

Newell A, Simon HA (1972) *Human problem solving* (Englewood Cliffs, NJ: Prentice-Hall).

Noy S, Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654):187–192.

Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A, et al. (2022) Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35:27730–27744.

Park JS, Popowski L, Cai C, Morris MR, Liang P, Bernstein MS (2022) Social simulacra: Creating populated prototypes for social computing systems. *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, 1–18.

Peng S, Kalliamvakou E, Cihon P, Demirer M (2023) The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* .

Reisenbichler M, Reutterer T, Schweidel DA (2025) Applying large language models to sponsored search advertising. *Marketing Science* .

Reisenbichler M, Reutterer T, Schweidel DA, Dan D (2022) Frontiers: Supporting content marketing with natural language generation. *Marketing Science* 41(3):441–452.

Salecha A, Ireland ME, Subrahmanya S, Sedoc J, Ungar LH, Eichstaedt JC (2024) Large language models display human-like social desirability biases in big five personality surveys. *PNAS nexus* 3(12):pgae533.

Shan G, Qiu L (2025) Examining the impact of generative ai on users voluntary knowledge contribution: Evidence from a natural experiment on stack overflow. *Information Systems Research* .

Simon HA, Hayes JR (1976) The understanding process: Problem isomorphs. *Cognitive psychology* 8(2):165–190.

Simon HA, Newell A (1971) Human problem solving: The state of the theory in 1970. *American psychologist* 26(2):145.

Toubia O, Gui GZ, Peng T, Merlau DJ, Li A, Chen H (2025) Database report: Twin-2k-500: A data set for building digital twins of over 2,000 people based on their answers to over 500 questions. *Marketing Science* .

Tranchero M, Brenninkmeijer CF, Murugan A, Nagaraj A (2024) Theorizing with large language models. Technical report, National Bureau of Economic Research.

Wang D, Hsu TY, Lu Y, Gu H, Cui L, Xie Y, Headean W, Yao B, Veeragouni A, Liu J, et al. (2025) Agenta/b: Automated and scalable web a/btesting with interactive llm agents. *arXiv preprint arXiv:2504.09723* .

Wang M, Zhang DJ, Zhang H (2024) Large language models for market research: A data-augmentation approach. *arXiv preprint arXiv:2412.19363* .

Wei J, Bosma M, Zhao VY, Guu K, Yu AW, Lester B, Du N, Dai AM, Le QV (2021) Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652* .

Wei J, Wang X, Schuurmans D, Bosma M, Xia F, Chi E, Le QV, Zhou D, et al. (2022) Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35:24824–24837.

Wiberg H, Dai T, Lam H, Kulkarni R (2025) Synergizing artificial intelligence and operations research: Perspectives from informs fellows on the next frontier. *INFORMS Journal on Data Science* .

Xie C, Chen C, Jia F, Ye Z, Lai S, Shu K, Gu J, Bibi A, Hu Z, Jurgens D, et al. (2024) Can large language model agents simulate human trust behavior? *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Yang Y, Zhang Y, Wu M, Zhang K, Zhang Y, Yu H, Hu Y, Wang B (2025) Twinmarket: A scalable behavioral and social simulation for financial markets. *arXiv preprint arXiv:2502.01506* .

Yang Z, Zhang Z, Zheng Z, Jiang Y, Gan Z, Wang Z, Ling Z, Chen J, Ma M, Dong B, et al. (2024) Oasis: Open agent social interaction simulations with one million agents. *arXiv preprint arXiv:2411.11581* .

Yao S, Yu D, Zhao J, Shafran I, Griffiths T, Cao Y, Narasimhan K (2023) Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems* 36:11809–11822.

Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan KR, Cao Y (2022) React: Synergizing reasoning and acting in language models. *The eleventh international conference on learning representations*.

Ye Z, Yoganarasimhan H, Zheng Y (2025) Lola: Llm-assisted online learning algorithm for content experiments. *Marketing Science*.

Zhang X, Sun C, Zhang R, Goh KY (2024) The value of ai-generated metadata for ugc platforms: Evidence from a large-scale field experiment. *arXiv preprint arXiv:2412.18337*.

Zhou Y, Yuan Y, Huang K, Hu X (2024) Can chatgpt perform a grounded theory approach to do risk analysis? an empirical study. *Journal of Management Information Systems* 41(4):982–1015.

Ziegler DM, Stiennon N, Wu J, Brown TB, Radford A, Amodei D, Christiano P, Irving G (2019) Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.