



Computational Basis of Large Language Models' Decision Making in Social Simulation

Ji Ma^{1,2} 

Abstract

Large language models (LLMs) increasingly serve as humanlike decision-making agents in social science and applied settings. These LLM agents are typically assigned humanlike characters and placed in real-life contexts. However, how these characters and contexts shape an LLM's behavior remains underexplored. In this study the author proposes and tests methods for probing, quantifying, and modifying an LLM's internal representations in a dictator game, a classic behavioral experiment on fairness and prosocial behavior. The author extracts “vectors of variable variations” (e.g., “male” to “female”) from the LLM's internal state. Manipulating these vectors during the model's inference can substantially alter how those variables relate to the model's decision making. This approach offers a principled way to study and regulate how social concepts can be encoded and engineered within transformer-based models, with implications for alignment, debiasing, and designing artificial intelligence agents for social simulations in both academic and commercial applications, strengthening sociological theory and measurement.

Keywords

behavioral experiment, prosocial behavior, large language model-based agent, social simulation

Two aspects of society are particularly interesting to sociologists: structures and meanings. Just as social network analysis supplied formal tools for representing relational structure, computational language models supply new formal methods for studying social meanings (Arseniev-Koehler 2024). More recently, large language models (LLMs; e.g., GPT-5, Llama, DeepSeek) extend earlier computational text methods, such as dictionaries, topic models, and sentiment classifiers (Edelmann et al. 2020; Nelson 2020), by producing context-sensitive representations and humanlike outputs rather than only aggregate labels (Bail 2024; Ziems et al. 2024). This shift has broadened their role: beyond assistance in representation, coding, or annotation, LLMs are now used as exploratory aides in hypothesis generation and, increasingly, as synthetic respondents in social simulations (Anthis et al. 2025; Banker et al. 2024).

However, the growing use of LLMs in sociology has raised critical questions about their validity, reliability, and theoretical grounding (Kozlowski and Evans 2025;

¹LBJ School of Public Affairs, University of Texas at Austin, Austin, TX, USA

²Gradel Institute of Charity, New College, University of Oxford, Oxford, UK

Corresponding Author:

Ji Ma, LBJ School of Public Affairs, University of Texas at Austin, 2315 Red River Street, Austin, TX 78712, USA

Email: maji@austin.utexas.edu

Wang, Morgenstern, and Dickerson 2025). The internal mechanisms driving their decisions remain opaque, and subtle variations in prompting can lead to inconsistent outputs, complicating their use as proxies for human cognition (Aher, Arriaga, and Kalai 2023; Ma 2024). This “black box” problem is particularly acute for sociologists, who require transparent and theoretically grounded methods to ensure that computational tools genuinely advance our understanding of social phenomena, rather than merely reproducing statistical artifacts.

In response, methodological work has largely advanced on two fronts, both of which treat the model as an *opaque system*. The first involves developing sophisticated statistical techniques to correct for model errors. Methods such as prediction-powered inference use LLM predictions as auxiliary information to be combined with a small set of “gold standard” human data, thereby improving the efficiency of causal estimates (Broska, Howes, and van Loon 2025). Similarly, design-based supervised learning corrects measurement errors in artificial intelligence (AI)-assisted data labeling by leveraging a small number of expert annotations (Rister Portinari Maranca et al. 2025). The second approach, prompt engineering, focuses on refining the input (e.g., personas) given to LLMs to elicit more humanlike responses (Bisbee et al. 2024). Although valuable, these methods focus on input-output validation, rather than interrogating the internal mechanisms that produce the observed behaviors.

This study introduces a framework for moving inside the black box of LLMs.¹ Rather than only validating inputs and outputs, I *probe*, *quantify*, and *steer* internal activations tied to sociological concepts (see Kim, Evans, and Schein 2025). The central questions I address are (1) how social meanings are internally represented in LLMs and (2) how those internal representations can be systematically manipulated to steer model behavior. Drawing on computational sociology, behavioral experiments, and activation engineering, I treat variables (e.g., gender, age, framing) as directions (vectors) in the residual-stream space of LLMs. I show how to (1) extract these socially meaningful vectors, (2) orthogonalize them to isolate each variable’s unique contribution, and (3) inject controlled perturbations to modify downstream behavioral decisions in a behavioral experiment. This provides a transparent, theory-guided workflow for assessing and manipulating how LLMs internally represent social factors, enabling more reliable and interpretable applications in sociological research. However, this study concerns mainly the *internal representational validity* and controlled steering within the model’s activation space. This study does not, on its own, establish *external sociological validity* (i.e., accurately representing human cognition and behavior), which requires separate comparative validation with human baselines.

LLMS IN SOCIOLOGICAL RESEARCH: TRENDS AND GAPS

Measuring “Meaning” with Word Embeddings

Computational text analysis in sociology spans manual content analysis, dictionary methods, and probabilistic topic models (Carlsen and Ralund 2022; Nelson 2020; Stoltz and Taylor 2024). A pivotal point came with distributional word embeddings, which represent lexical items as points in a high-dimensional semantic space (Mikolov

et al. 2013). These techniques enabled sociologists to operationalize cultural schemas, semantic fields, and ideologies as measurable geometric relations (Arseniev-Koehler 2024; Garg et al. 2018; Kozlowski, Taddy, and Evans 2019). Yet classical embeddings suffer from *context invariance*: each word has one static vector regardless of usage, at odds with interactionist and pragmatic traditions that emphasize social and cultural meaning as contextual (Boutyline and Arseniev-Koehler 2025:92).

Transformer-based LLMs (e.g., BERT, GPT variants, Llama) address this by producing *contextualized* token representations across layers. The same surface form now traces different activation trajectories depending on discourse, persona, or pragmatic framing, aligning more closely with sociological accounts of fluid meaning (Mostafavi, Porter, and Robinson 2025). This shift, from static lexical geometry to dynamically evolving internal representations, creates new empirical leverage: researchers can inspect how latent social dimensions (e.g., gendered role expectations, moral frames) are encoded, interact, and transform as text is processed.

LLM adoption in social science has now moved beyond corpus labeling toward end-to-end research mediation. Earlier toolkits (topic models, sentiment analyzers, supervised classifiers) uncovered attitudes and sentiments but required labeled data or bespoke training (Grimmer and Stewart 2013; Roberts 2016). Unlike conventional machine learning algorithms, LLMs exhibit remarkable zero-shot learning (performing new tasks without seeing any examples) and few-shot learning (learning effectively from only a handful examples) capabilities (Brown et al. 2020). LLMs can thus handle new tasks, ranging from text classification to reasoning, without substantial labeled training data. This has eased large-scale text analysis projects in social science, where manually annotated data sets are often limited (Ziems et al. 2024). Moreover, researchers have begun to use LLMs not just as passive tools for coding and summarizing text, but also as active “collaborators” in hypothesis generation, literature synthesis, and research ideation (Bail 2024; Banker et al. 2024; Zhou et al. 2024). In this way, LLMs increasingly shape every stage of empirical inquiry, from initial conceptualization to final reporting (Chang et al. 2024; Korinek 2025).

LLMs as Synthetic Respondents: Opportunities and Validity Challenges

Beyond their role as representational tools, LLMs are increasingly used as “intelligent agents” capable of simulating humanlike responses in social scenarios. This capability stems from their training on vast data sets of human interaction, which enables them to generate contextually relevant and coherent replies. A growing body of research evaluates LLMs as “synthetic respondents” in survey and experimental paradigms (Argyle et al. 2023; Horton 2023). In settings such as public goods, trust, and dictator games, models have shown prosocial and strategic behaviors that appear consistent with human aggregates (Johnson and Obradovich 2023; Leng and Yuan 2024; Mei et al. 2024; Xie et al. 2024). Such convergences highlight their potential for rapid, low-cost behavioral prototyping and hypothesis exploration (Anthis et al. 2025).

However, this approach faces significant challenges of validation and external validity. The behavior of LLM agents is often fragile, with subtle shifts in prompts,

framing, or situational cues causing drastic changes in their decisions (Aher et al. 2023; Kozlowski and Evans 2025). Deeper inconsistencies emerge when examining persona manipulations. For example, Ma (2024) found that cues for age, gender, and personality elicited behavioral patterns that were inconsistent across different model families and diverged from human distributions. LLMs may generate responses that reflect out-group stereotypes rather than authentic in-group perspectives (i.e., *mispresentation*), and they tend to erase heterogeneity within demographic groups by optimizing for probable, average outputs (i.e., *flattening*) (Wang et al. 2025). From an epistemic perspective, these issues are not merely technical limitations; they risk perpetuating harmful histories of epistemic injustice (Fricker 2007), where the lived experiences of marginalized groups are erased or spoken for by dominant voices (Atari et al. 2023; Wang et al. 2025). These discrepancies fuel broader concerns that LLM behaviors may arise from sophisticated statistical pattern completion rather than from internalized causal schemas or a genuine understanding of the world (Lake et al. 2017). Consequently, benchmarking model outputs against human aggregates is necessary but insufficient for validation.

These challenges caution that apparent behavioral convergence between LLMs and humans may be superficial, arising from text-trained pattern matching rather than humanlike cognition. To move beyond validating surface outputs, I distinguish between *internal representational validity* (whether constructs are encoded as meaningful directions that affect a model's internal computations) and *external sociological validity*, which concerns the correspondence of model behavior to human subjects. My analyses address the former by treating experimental variables (prompt features, persona attributes, instructions) as controlled factors and interrogating how they alter internal representations en route to decisions (Anthis et al. 2025). The latter requires dedicated human–LLM comparisons.

Beyond Prompting: Activation Engineering

Prompt engineering is accessible because simple text edits can shift outputs, but it is fragile and entangles style, content, and persona, making effects hard to separate or replicate across paraphrases. When the goal is theory, convenience is not enough: we must test whether observed differences reflect stable internal dimensions rather than surface cues. Activation engineering makes that trade worthwhile by intervening in the residual stream to localize where a factor acts, quantify its alignment with the decision direction, and disentangle variables for causal probing (Subramani, Suresh, and Peters 2022; Turner et al. 2024). This process supports *methodological robustness*: computationally skilled sociologists can audit internal validity (e.g., estimating, orthogonalizing, and stress-testing concept vectors), so the broader field can deploy prompt-level workflows with greater confidence and without worrying about the model's internal validity.

Specifically, the activation engineering extracts internal differences associated with theoretically defined contrasts, orthogonalizes them to isolate unique social dimensions, and reinjects controlled perturbations to test causal influence on downstream

decisions. This shifts analysis from external prompting to mechanism probing and addresses three sociological needs: (1) *concept operationalization*: social variables (gender persona, interaction horizon, framing) become measurable vectors whose geometry (angles, projections) encodes relational structure; (2) *causal probing*: injecting purified (orthogonalized) vectors tests whether a latent representation exerts directional pressure on decisions, moving beyond correlational output validation; and (3) *transparency and auditability*: layerwise trajectories reveal where (depth) and how (magnitude, alignment) social factors enter decision formation, informing bias audits and theoretical interpretation.

In this study, I instantiate this framework within a dictator game, demonstrating how variables with social meanings can be (1) extracted from residual streams, (2) disentangled via orthogonalization, and (3) applied as precise steering interventions. This establishes a methodological bridge between social and behavioral sciences and computational modeling, advancing a more interpretable, mechanism-aware computational social science. This method improves interpretability and control of internal mechanisms, but it does not by itself justify inferences about human social behavior without external validation against human baselines.

ARCHITECTURE OF LLMS: A LESS TECHNICAL OVERVIEW

LLMs are designed to predict text token by token, much like a sophisticated version of a smartphone’s autocomplete feature. When the phone finishes a sentence, it is essentially guessing, on the basis of patterns it has learned, which word (or token) is most likely to come next. For example, imagine you start typing: “The cat sat on the” A simple guess might be “mat,” as that is a common phrase. But LLMs can go further, understanding that words such as *floor* or *sofa* might also fit depending on the broader context. The LLM “knows” this because it has been trained to find patterns in massive amounts of text and to predict how real sentences typically continue.

We can use the Llama 3.1-8B model as an example of how LLMs operate because its architecture and training methodology are representative of transformer-based language models, the dominant architecture of LLMs (Vaswani et al. 2017). The Llama family of models, originally open sourced by researchers aiming to democratize access to powerful LLMs, has gained wide adoption in academic and industry settings.

In broad terms, Llama 3.1-8B follows a pipeline that begins by converting textual input into numerical form (i.e., token IDs), converting these token IDs into high-dimensional embeddings, passing those embeddings through multiple *decoder layers* that refine the representation, and finally using a *head* layer to guess the next most probable token. In the following subsections, I delve deeper into the mechanics of this pipeline. Along with the illustration in Figure 1, this explanation demonstrates how the model transforms input text into meaningful predictions.

Tokenizer: Turning Words into Numeric IDs

The first step is to convert raw text into discrete units (*tokens*) that the model can process, a procedure called *tokenization*. The tokenizer assigns each token a unique integer

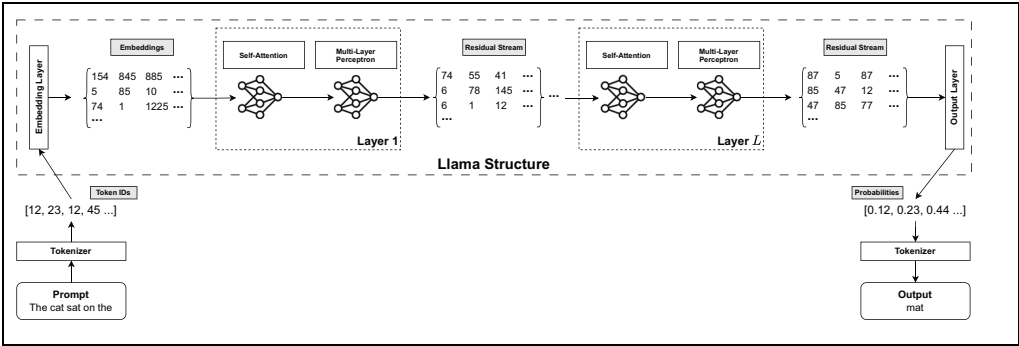


Figure 1. Illustration of Llama model internal structure.
 Note: For Llama 3.1-8B-Instruct, $L = 32$.

ID. For illustration, *cat* might map to 3,456 and *mat* to 7,891. Modern LLMs (including Llama 3.1-8B) use subword tokenization (a byte-pair encoding–style algorithm) rather than whole-word vocabularies (Sennrich, Haddow, and Birch 2016). Rare or morphologically complex words are decomposed into frequent substrings so the model can represent virtually any input without an enormous word list. For example, *sociologically* might be segmented into smaller meaningful pieces (e.g., *socio*, *logical*, *ly*); exact splits depend on the trained tokenizer. The Llama 3.1-8B tokenizer has a fixed vocabulary of 128,256 learned tokens (including whole words, subwords, symbols, and control tokens), each mapped deterministically to an ID. These IDs are the numerical sequence passed to the embedding layer.

Embedding Layer: Mapping Token IDs to Vectors

Once the input text is tokenized, each token ID is mapped to a high-dimensional vector known as an *embedding*. This is where the model begins to represent the semantic meaning of the tokens. In Llama 3.1-8B, each token is mapped to a 4,096-dimensional vector. These embeddings are not static; they are learned during the model’s training process and are updated as the model processes the text. This “contextualization” of embeddings is a key feature of transformer models, distinguishing them from earlier methods such as Word2Vec, where each word had a fixed embedding. The initial embeddings are then combined with positional encodings, which provide the model with information about the order of the tokens in the sequence.

The “Thinking Blocks”: Stacked Decoder Layers

The “heavy lifting” in Llama 3.1-8B happens in 32 *stacked decoder layers*. Each layer refines the representation of the input text (i.e., the embedding matrix) and passes the updated matrix, the *residual stream*, to the next layer. Each decoder layer mainly has two critical components: the *self-attention mechanism* and the *multilayer perceptron* (MLP). The following subsections briefly introduce the three major components.

Self-Attention Mechanism. This component helps the model understand which words or tokens in a sentence should be most relevant to each other. For instance, if a sentence discusses “cats chasing mice,” the model can learn to focus on words such as *chase*, *feline*, and *mouse* in the right context, even if they appear several words apart. This mechanism allows the model to capture long-range dependencies and understand the relationships between different parts of the text. This design is a key reason why Llama models can generate coherent and contextually relevant sentences, and it is a major breakthrough in natural language processing (Vaswani et al. 2017).

MLP. The output of the self-attention mechanism is then passed to an MLP, also known as a feed-forward network. The MLP is a relatively simple neural network that applies a series of nonlinear transformations to the attention output. Its role is to process the information aggregated by the self-attention mechanism and to add representational capacity to the model. Self-attention is responsible for identifying relationships between tokens, but the MLP is where much of the model’s “knowledge” is stored and processed. The MLP allows the model to learn complex patterns and relationships in the data that go beyond the pairwise comparisons of the attention mechanism.

Residual Streams (or Skip Connections). After each subblock (self-attention or MLP), the output is added back into the original input of that subblock. This creates a “residual stream,” also called a “skip connection.” This step ensures that each layer can refine existing information without discarding what was learned previously. In our “The cat sat on the . . .” example, a residual connection ensures that early-layer knowledge about *cat* and *mat* persists as later layers weigh in. This “accumulated knowledge” approach makes training deep networks more stable and effective.

Final Output Layer: Predicting the Next Token

Once the text data flow through all 32 decoder layers, the final layer, labeled *lm_head* in Llama 3.1-8B, translates the 4,096-dimensional representation into a probability distribution across 128,256 possible tokens (e.g., [0.12, 0.23, 0.22 . . .], a list of 128,256 possibilities²; the index position of each possibility corresponds to a unique token). This is where the model ultimately decides what the next word or token should be, given everything it has seen so far. In our example, the model might predict that the next word after “The cat sat on the . . .” is *mat* with a 70 percent probability, *floor* with a 20 percent probability, and so on.

Through the above architecture, Llama 3.1-8B can generate coherent and contextually relevant sentences, paragraphs, or even long-form responses. Whether it is predicting *mat* when you type “The cat sat on the . . .” or handling more intricate tasks, such as summarizing lengthy documents, answering specialized-domain questions, or crafting creative narratives, the model relies on a pipeline of embeddings, attention, MLP, and residual connections, each contributing to the advanced language understanding and generation system.

Relation to Other LLMs

I use Llama 3.1-8B as an illustrative case because its transformer architecture is representative of modern open-weight LLMs (Vaswani et al. 2017). Our method relies on several key architectural elements: tokenization, embedding layers, stacked decoders with attention and feed-forward subblocks, and a residual stream with additive updates. These elements are common across a wide range of models. The steering procedure operates on the geometry of the residual stream, an architectural invariant of decoder-only transformers. Therefore, the control mechanism is, in principle, portable across models.

However, portability does not imply uniform effects. Models vary in layer depth, hidden dimensionality, normalization, gating, and training data. For instance, state-of-the-art models such as DeepSeek-R1 use mixture-of-experts feed-forward blocks, where tokens are routed to specialized experts, altering sparsity and local subspace structure (Dai et al. 2024; DeepSeek-AI et al. 2025; Lepikhin et al. 2020). Diverse pretraining and instruction tuning also change how and where social concepts are encoded. Our approach can accommodate these differences by reestimating all steering vectors for each model and depth. In theory, as long as residual updates are approximately additive and the decision subspace is present, injecting the dependent variable (DV)–aligned component of a partial independent variable (IV) vector will produce controlled behavioral shifts. In practice, the magnitude, sign stability, and layer sensitivity of these shifts are empirical and may vary with model family and data.

KEY METHODOLOGICAL BACKGROUND

Representation of Social Concepts in LLMs

Recent studies have shown that LLMs encode social concepts (Kim et al. 2025; Park, Choe, and Veitch 2024; Tigges et al. 2023) and facts (Engels et al. 2024; Gurnee and Tegmark 2024) in their internal space, either as one-dimensional (e.g., good vs. bad) or multidimensional (e.g., seven days of the week). Concretely, the model’s internal representation of, say, “male” or “female” can be viewed as a *vector* in a high-dimensional space. By nudging this vector in a particular direction, one can effectively steer the model’s output; for instance, adding a small offset to the “male” vector might make the model more likely to adopt male-coded traits or perspectives in its text.

Because these transformations manifest as vectors in the residual-stream space, we can isolate and control each variable’s unique contribution by defining corresponding vectors (see the next section and Figure 2). This means a vector capturing “age 20 to age 40,” “female to male,” or “neutral to liberal” can be treated with standard vector operations, such as addition, subtraction, and projection, without loss of interpretability. The additive structure in the LLM’s hidden layers further enables targeted steering via small, well-chosen modifications to the residual stream, allowing us to manipulate the model’s behaviors (see the section “Additive Nature of Internal Representations in LLMs”).

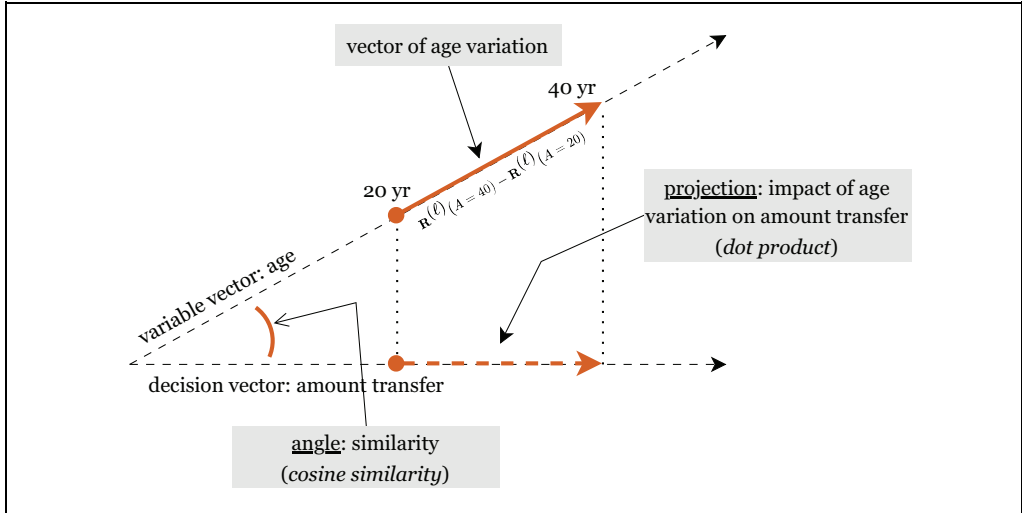


Figure 2. Illustration of variable vectors at layer ℓ .

Note: $\mathbf{R}^{(\ell)}(A = 20)$ is the average residual stream of layer ℓ for all trials with age 20.

Vector of Variable Variation

Figure 2 provides an intuitive geometric view of how we treat each variable’s influence as a *vector* in the LLM’s residual-stream space. In this simplified two-dimensional illustration, moving from *age 20* to *age 40* at layer ℓ appears as a vector in the “age” direction. This vector is obtained by subtracting the average residual stream of layer ℓ for all trials with age 20 from the average residual stream of the same layer for all trials with age 40. Similarly, subtracting the average residual stream of layer ℓ for all trials with $D = 0$ from the average residual stream of the same layer for all trials with $D = 10$ yields a vector in the “decision” direction. We can then use standard vector operations to quantify interactions among the following variables.

Angle (Cosine Similarity). The cosine similarity between two vectors measures the cosine of the angle between them. It is a measure of directional similarity, ranging from -1 (opposite directions) to 1 (same direction), with 0 indicating orthogonality. I use cosine similarity to assess whether the representations of two concepts are aligned in the model’s activation space.

Cosine similarity is a common way to check whether two concepts “point” in a similar direction, but it needs to be used with caution. As Steck, Ekanadham, and Kallus (2024) showed, the way embeddings are learned and regularized can stretch or shrink the hidden axes of a space without changing the model’s actual predictions. After this stretching, normalizing vectors (as cosine does) can produce arbitrary, and sometimes nonunique, “similarities,” even though the underlying dot products remain well defined. In deep models, combinations of regularization methods can implicitly rescale different latent dimensions, making cosine values even harder to interpret.

To address this, I explicitly separate “direction” from “strength.” I use cosine similarity only as a directional indicator, and I pair it with the dot product, which also

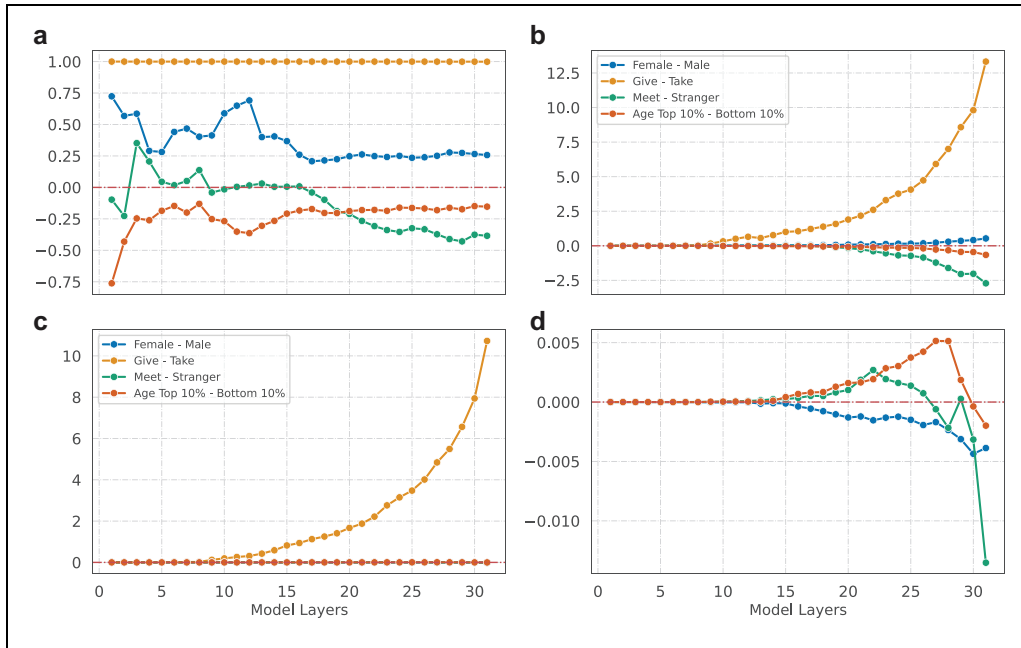


Figure 3. Relations between instrumental and dependent variable steering vectors: (a) cosine similarity (Y) by layers (X), (b) dot product (Y) by layers (X), (c) partial dot product (Y) by layers (X), and (d) partial dot product, excluding *Give-Take*.

Note: In Figure 3, (d) is the same as (c) but with *Give-Take* excluded to better visualize the other variables. These profiles characterize internal relations in large language models' residual streams and should not be confused with human cognition or behavior.

reflects magnitude. Looking at both alignment (cosine) and magnitude (dot product) gives a more complete picture of how variables relate inside the model and avoids overinterpreting the cosine in isolation (see the results on the dissociation between the two measures; Figure 3).

Projection and Dot Product. The dot product of two vectors is related to the cosine similarity but is scaled by the magnitudes of the vectors. Specifically, the dot product of vector **a** and **b** is $\|\mathbf{a}\| \|\mathbf{b}\| \cos\theta$. I use the dot product to measure the magnitude of the projection of one vector onto another. This allows us to quantify how much a change in one variable (e.g., age) contributes to a change in another (e.g., the decision to give money).

Mathematically, we can define the *vector of variable variation*. Suppose we want the vector representing how a single variable X shifts the LLM's residual-stream representation at layer ℓ . For concreteness, let $X = \text{age}$ with two distinct values: $A = 20$ and $A = 40$. In each experimental trial, we hold the variable of interest (in this case, age) constant while randomizing all other variables (gender, game instruction, future interaction). This experimental design allows us to isolate the average effect of the variable of interest, as the effects of the other, randomized variables will average out to zero over a large number of trials. This is a standard approach in experimental research and

is preferable to holding all other variables constant, as it allows us to estimate the average effect of a variable in the presence of other sources of variation.

We then collect two sets of residual-stream vectors: (1) $\{\mathbf{r}_1^{(\ell)}, \dots, \mathbf{r}_{n_1}^{(\ell)}\}$, from all trials in which $A = 20$ (but G, I, M vary randomly) and (2) $\{\mathbf{r}_{n_1+1}^{(\ell)}, \dots, \mathbf{r}_{n_1+n_2}^{(\ell)}\}$, from all trials in which $A = 40$ (again with G, I, M randomized).

We compute the mean residual stream for each set:

$$\mathbf{R}^{(\ell)}(A=20) = \frac{1}{n_1} \sum_{j=1}^{n_1} \mathbf{r}_j^{(\ell)}, \quad \mathbf{R}^{(\ell)}(A=40) = \frac{1}{n_2} \sum_{k=n_1+1}^{n_1+n_2} \mathbf{r}_k^{(\ell)}.$$

By taking the *difference* of these means, we obtain the *vector of age variation*:

$$\mathbf{v}_{\text{age}} = \mathbf{R}^{(\ell)}(A=40) - \mathbf{R}^{(\ell)}(A=20).$$

Because the other variables are randomized and we have a large number of experimental trials, \mathbf{v}_{age} reflects *how shifting from age 20 to age 40 changes the internal representation on average*, after “averaging out” any incidental effects of gender, game instruction, or future interaction on residual streams. In the same way, we can define analogous vectors for gender (male vs. female), game instruction (give vs. take), or future interaction (meet vs. not meet) by grouping trials according to whichever variable is under study and then subtracting the mean residual streams. Figure 2 shows only a two-dimensional plane for clarity, but the same vector operations extend naturally into the higher dimensional spaces of transformer-based LLMs.

This method builds directly on the concept of “steering vectors” from the activation engineering literature (Kim et al. 2025; Turner et al. 2024), but it makes a key sociological contribution. Whereas most existing work focuses on steering broad, often abstract concepts (e.g., “romance,” “violence”), I develop a methodology for extracting and manipulating steering vectors that correspond to *specific and socially meaningful variables* within a *controlled experimental design*. This allows us to move from a general ability to steer LLMs to a specific, sociologically informed method for testing hypotheses about the influence of social factors on behavior. This is the core innovation of the approach: we are not just steering LLMs, we are conducting virtual social experiments at the variable level inside them.

Additive Nature of Internal Representations in LLMs

Modern transformer-based LLMs maintain a *residual stream* $\mathbf{r}^{(\ell)}$ at each layer ℓ . These residual streams accumulate the hidden representation of the input prompt as it propagates through attention and feed-forward sublayers. Crucially, in many transformer designs, these updates occur in a largely *additive* fashion:

$$\mathbf{r}^{(\ell)} \approx \mathbf{r}^{(\ell)} + (\text{some transformation of } \mathbf{r}^{(\ell)}).$$

Because of this structure, small, well-chosen interventions in $\mathbf{r}^{(\ell)}$ can exert a relatively *predictable* effect on the final output (the LLM’s generated text or decision). Specifically, by carefully inserting a *steering vector* Δ (see the section “IV Steering Vectors”) into $\mathbf{r}^{(\ell)}$, we nudge the hidden representation in a direction and magnitude of interest. Subsequent layers preserve much of this directional influence because of the additive nature of the residual streams.

To give a concrete example, imagine an LLM-based dictator that must decide how much money to give (or take) to another player. For a simplified two-layer transformer, the first layer produces $\mathbf{r}^{(1)}$, which is then mostly additively updated in the second layer to yield $\mathbf{r}^{(2)}$. If we add a small steering vector Δ to $\mathbf{r}^{(1)}$ at layer 1 (e.g., nudging the model toward generosity), Δ “persists” onto layer 2 through the residual connection and ultimately shifts the final decision. In this way, targeted additions to the residual streams can directly manipulate the model’s internal reasoning trajectory and shape its final decision.

OVERVIEW OF EXPERIMENT DESIGN

I systematically *probe* (step 1), *quantify* (steps 2, 3, and 4), and *modify* (step 5) how LLM-based agents behave in a dictator game. This approach provides not only an *analytic* framework (decomposing the influence of different factors on final decisions) but also a *practical* one (actively steering the model’s choices). The approach consists of the following steps:

1. *Focus on residual streams*: We first record the *residual streams* because their additive nature allows small “nudges” to predictably shift the model’s internal representations.
2. *Identify steering vectors*: For each factor X , we measure how the LLM’s internal representation changes when toggling X (e.g., female vs. male).
3. *Partial out confounds*: By subtracting overlapping components, we obtain a *pure* vector capturing only X ’s unique effect (e.g., uncontaminated by age, framing of games).
4. *Project onto decision space*: We project the partial vector of X onto the LLM’s decision vector, so we can measure the impact of each factor on the LLM’s final decision.
5. *Manipulate*: We inject a scaled version of the projection into the residual streams to modify X ’s impact on the LLM’s final decision, allowing us to steer the model’s behavior.

In summary, by *quantifying* steering vectors, *orthogonalizing* them for pure effects, and then *injecting* the scaled projections back into the LLM’s residual streams, we can manipulate an LLM agent’s decision in the dictator game.

METHODS

Baseline Experiment: Game Setup

I followed the game setup described in Ma (2024:12–13) and devised a simplified version with fewer IVs to better focus on the main tasks of this study (i.e., *probe*, *quantify*,

and *modify*). In this dictator game, an LLM-based “dictator” decides how much to allocate (or “take”) from another player (the “recipient”). Both the initial endowment and the amount that can be transferred are set to \$20. For example, if the dictator transfers \$10 to the recipient, both players will end up with \$30, which is considered a fair split.

The following are the IVs considered in the baseline experiment (i.e., for probing and quantifying purposes, no manipulation):

1. *Agent persona*: Gender $\mathbf{G} \in \{\text{male, female}\}$ and Age $\mathbf{A} \in \{20, 21, 22, \dots, 60\}$.
2. *Game instruction* $\mathbf{I} \in \{\text{give, take}\}$: In the *give* condition, the dictator can only transfer a non-negative amount $\mathbf{D} \in [0, 20]$. In the *take* condition, the dictator can transfer a negative amount (effectively “take” from the recipient), so $\mathbf{D} \in [-20, 20]$.
3. *Future interaction* $\mathbf{M} \in \{\text{meet, not meet}\}$: Whether the recipient is a stranger the agent will meet afterward (“stranger_meet”) or not (“stranger”).

Let the DV be the amount D that the LLM-based dictator ultimately gives (positive) or takes (negative). The model’s input includes (G, A, I, M) plus any additional invariant prompt context. Formally,

$$D \in \begin{cases} [0, 20], & \text{if } I = \text{give,} \\ [-20, 20], & \text{if } I = \text{take.} \end{cases}$$

Baseline Trials: Randomizing All Variables

In each baseline trial, I (1) randomize the values of all the input variables, G (gender), A (age from 20 to 60), I (give vs. take), and M (meet vs. not meet), and (2) collect the LLM’s responses (i.e., the amount transferred D). I did not modify any of the model’s generation hyperparameters, such as temperature or top-p. This yields an empirical distribution:

$$p(D|G, A, I, M).$$

To collect sufficiently robust data:

1. We obtain a set of responses (i.e., trials) $\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_k$, where $k = 1,000$.
2. The LLM can be prompted multiple times with the same tuple $(\mathbf{g}, \mathbf{a}, \mathbf{i}, \mathbf{m})$ because the total number of input combinations (i.e., 320) is much smaller than the total number of trials (i.e., 1,000).
3. We can compute $\mathbb{E}[\mathbf{D} | \mathbf{G} = \mathbf{g}, \mathbf{A} = \mathbf{a}, \mathbf{I} = \mathbf{i}, \mathbf{M} = \mathbf{m}]$ to capture the average behavior under each condition.

This serves as our *baseline distribution* for later comparisons and manipulations.

IV STEERING VECTORS

Basic Steering Vector

I now briefly frame the key concepts within the dictator game context. We are interested in how certain IVs, such as gender (G), age (A), game instruction (I), or meeting condition (M), influence the final decision D . To capture each IV's direction and magnitude of influence inside the LLM, we define a *steering vector* as, consider an IV $X \in \{G, A, I, M\}$ that can take two distinct values x_1 and x_2 . Holding all other variables fixed or randomized,³ we measure the difference in the layer ℓ residual stream:

$$\mathbf{v}_X^{(\ell)} = \mathbf{R}^{(\ell)}(x_2) - \mathbf{R}^{(\ell)}(x_1).$$

$\mathbf{v}_X^{(\ell)}$ quantifies how much the internal representation at layer ℓ shifts because of changing X from x_1 to x_2 .

Partial Steering Vector

In realistic setups, multiple IVs (e.g., G, A, I, M) interact in the prompt. For instance, the LLM might conflate *female* with *younger* if these often co-occur in training data. To focus on a single IV X_1 independently, we remove the influence of other IVs X_2, X_3 , and so on.

For two variables X_1 and X_2 , if $\mathbf{v}_{X_1}^{(\ell)}$ and $\mathbf{v}_{X_2}^{(\ell)}$ are their respective steering vectors, the *partial* steering vector of X_1 controlling for X_2 is

$$\mathbf{v}_{X_1|X_2}^{(\ell)} = \mathbf{v}_{X_1}^{(\ell)} - \frac{\langle \mathbf{v}_{X_1}^{(\ell)}, \mathbf{v}_{X_2}^{(\ell)} \rangle}{\|\mathbf{v}_{X_2}^{(\ell)}\|^2} \mathbf{v}_{X_2}^{(\ell)}.$$

By subtracting the component of $\mathbf{v}_{X_1}^{(\ell)}$ aligned with $\mathbf{v}_{X_2}^{(\ell)}$, we isolate the “pure” effect of X_1 . This procedure can be repeated or extended to three or more variables using standard multivariate orthogonalization, generating a partial steering vector $\mathbf{v}_{X_1|\{\text{others}\}}^{(\ell)}$.

For example, if $X_1 = \text{gender}$ (male vs. female) and $X_2 = \text{age}$ (young vs. older), $\mathbf{v}_{X_1|X_2}^{(\ell)}$ removes age-related correlations from the gender vector. Thus, we can specifically measure and manipulate “male versus female” while *not* shifting the representation in a direction correlated with “young versus older.”

DV STEERING VECTOR

Following the same rationale as for IVs, we can define a *DV steering vector* that captures how the LLM's internal representation changes when the final decision D shifts from one anchor to another. In our dictator game, D is the amount of dollars allocated to the recipient, which can range from -20 (taking all) to $+20$ (giving all), depending on “take” versus “give” frames. To identify the model's DV steering vector, we pick

two typical reference decisions, $D = 10$ (a “fair” split of 20) and $D = 0$ (giving nothing). At layer ℓ ,

$$\mathbf{v}_D^{(\ell)} = \mathbf{r}^{(\ell)}(x \mid D=10) - \mathbf{r}^{(\ell)}(x \mid D=0).$$

$\mathbf{v}_D^{(\ell)}$ is how the residual stream changes when the model’s output changes from $D = 0$ to $D = 10$ at layer ℓ .

PROJECTING IV STEERING VECTOR ONTO DV

The next question is how strongly each IV pushes the LLM’s decision D . To see whether a partial steering vector $\mathbf{v}_{X|\{\text{others}\}}^{(\ell)}$ (for an IV X) actually affects the decision D , we project it onto $\mathbf{v}_D^{(\ell)}$:

$$\text{Proj}_D\left(\mathbf{v}_{X|\{\text{others}\}}^{(\ell)}\right) = \frac{\left\langle \mathbf{v}_{X|\{\text{others}\}}^{(\ell)}, \mathbf{v}_D^{(\ell)} \right\rangle}{\|\mathbf{v}_D^{(\ell)}\|^2} \mathbf{v}_D^{(\ell)}.$$

The *sign* of $\left\langle \mathbf{v}_{X|\{\text{others}\}}^{(\ell)}, \mathbf{v}_D^{(\ell)} \right\rangle$ indicates whether X *increases* or *decreases* D . The *magnitude* of this projection ($\|\text{Proj}_D(\cdot)\|$) is an “effect size”: a larger value suggests that toggling X can produce a bigger shift between $D = 0$ and $D = 10$.

For example, if $\mathbf{v}_{G=\text{female}|\{A,I,M\}}^{(\ell)}$ has a positive projection on $\mathbf{v}_D^{(\ell)}$, it implies that, all else equal, adopting a “female” persona pushes the model to giving more dollars. Conversely, if $\mathbf{v}_{A=60|\{G,I,M\}}^{(\ell)}$ has a strong negative projection, it indicates that an older persona correlates with less generous allocations in the model’s internal reasoning.

MANIPULATING THE EFFECT OF IVS ON DV

Having identified how strongly each IV influences D , we now want to precisely manipulate the LLM’s decision by variable. The key operation is to inject a vector into the residual streams such that it pushes the model in the direction of a higher (or lower) D .

Choice between Injecting IV Partial Steering Vector or Its Projection

The *partial steering vector* for an IV X , $\mathbf{v}_{X|\{\text{others}\}}^{(\ell)}$, captures how changing X shifts the internal representation at layer ℓ , after orthogonalizing out the influence of other variables. The *projection* of that partial vector *onto the DV direction*, that is, $\text{Proj}_D\left(\mathbf{v}_{X|\{\text{others}\}}^{(\ell)}\right)$, is strictly the component of $\mathbf{v}_{X|\{\text{others}\}}^{(\ell)}$ that points in the same (or opposite) direction as $\mathbf{v}_D^{(\ell)}$.

If we inject the *entire* partial steering vector, we might also be adding directions that do not project onto $\mathbf{v}_D^{(\ell)}$. Those extra directions could have unintended effects on the model’s reasoning or how it articulates its choice, possibly influencing style,

justification, or other textual features that are not directly about increasing or decreasing D . By contrast, if we only inject the *projection*, we only push the model in the direction that moves D , allowing a more precise intervention.

Depending on the objective, either approach can be informative, and exploring the differences between the two can itself be an interesting, although potentially complex, research effort. Here, my main priority is to demonstrate that manipulating the IVs can indeed produce expected changes in the model's decision. Thus, to favor precision over comprehensiveness, I chose to inject *only the projections* of the partial IV steering vectors.

Injecting the Projection of Partial IV Steering Vector to Residual Stream

Let $\alpha \in \mathbb{R}$ be an *injection coefficient* that controls both the sign and magnitude of the intervention. For each layer ℓ , we replace the original residual stream $\mathbf{r}^{(\ell)}(x)$ with

$$\tilde{\mathbf{r}}^{(\ell)}(x) = \mathbf{r}^{(\ell)}(x) + \alpha \cdot \text{Proj}_D \left(\mathbf{v}_{X|\{\text{others}\}}^{(\ell)} \right).$$

As a result, the pure representation of X in residual streams can be amplified when $\alpha > 0$ and reduced if $\alpha < 0$. Because we used the projection of the partial steering vector of X , this manipulation remains specific to X . That is, we are not inadvertently pushing the hidden state in directions aligned with other variables (such as age or meeting condition).

Changing the LLM's Decision

Consider the function f that maps the final-layer residual stream $\mathbf{r}^{(L)}$ to the decision variable D , representing the original decision before any manipulation:

$$f(\mathbf{r}^{(L)}) = D.$$

Let $\mathbf{p}_D^{(L)} = \text{Proj}_D \left(\mathbf{v}_{X|\{\text{others}\}}^{(L)} \right)$. After introducing a small perturbation to the residual stream, we approximate the new decision as

$$\begin{aligned} f(\tilde{\mathbf{r}}^{(L)}) &= f(\mathbf{r}^{(L)} + \alpha \cdot \mathbf{p}_D^{(L)}) \\ &\approx f(\mathbf{r}^{(L)}) + \alpha \cdot \left\langle \nabla_{\mathbf{r}} f(\mathbf{r}^{(L)}), \mathbf{p}_D^{(L)} \right\rangle \\ &= D + \alpha \cdot \left\langle \nabla_{\mathbf{r}} f(\mathbf{r}^{(L)}), \mathbf{p}_D^{(L)} \right\rangle. \end{aligned}$$

In this equation, $\nabla_{\mathbf{r}} f(\mathbf{r}^{(L)})$ is the *gradient* of the output function with respect to the residual stream. This gradient indicates how sensitive the decision D is to changes in each component of $\mathbf{r}^{(L)}$. $\left\langle \nabla_{\mathbf{r}} f(\mathbf{r}^{(L)}), \mathbf{p}_D^{(L)} \right\rangle$ is the dot product and quantifies the *direction and strength* of how $\mathbf{p}_D^{(L)}$ affects D . α is the injection coefficient that controls both the *magnitude* and *direction* of the perturbation. Layer-by-layer adjustments

compound in a similar manner because of the residual nature of the transformer architecture, meaning that interventions at different layers can have cumulative or localized effects on the final decision.

FINE-TUNING THE MANIPULATION

Eventually, the final intervened decision depends on the product $\alpha \cdot \langle \nabla_{\mathbf{r}} f(\mathbf{r}^{(L)}), \mathbf{p}_D^{(L)} \rangle$, where most of the terms are intrinsic to the model and the prompt. However, two key parameters can be fine-tuned to alter model decisions: the layer number L and the injection coefficient α .

Layer number L determines the stage at which the residual stream is modified within the model. Intervening at *earlier layers* can have a more *amplified effect* on the final decision because of the cumulative nature of residual connections; intervening at *later layers* allows more *localized control* over specific aspects of the decision-making process.

Injection coefficient α controls the strength of the manipulation. A larger $|\alpha|$ results in a more significant shift in the decision D , either amplifying or reversing the influence of the IV X . Existing studies suggest that $|\alpha|$ should be smaller than 15 to generate meaningful and interpretable results; otherwise, the model may collapse and generate nonsensical outputs (Turner et al. 2024). We expand the range of α to 30 to explore a larger spectrum of possible effects.

By adjusting these two parameters, we can further fine-tune the extent to which the model’s decision is altered. Through the two parameters, our interventions can be both *targeted* (through the choice of layer) and *controllable* (through the scaling factor α), enabling nuanced manipulation of the LLM’s behavior.

RESULTS: BASELINE MODEL PERFORMANCE AND BEHAVIOR

To ensure the LLM agent behaves logically, I include a verification question at the end of each trial to assess whether the agent accurately calculates the final allocations for both the dictator and the recipient after the transfer. Performance on this question can serve as a measure of the agent’s mathematical reasoning and understanding of the game rules.

Out of 1,000 baseline trials, the agent produces logically correct responses in 571 instances (57.1 percent). The giving rates exhibit a bimodal distribution, with agents giving either nothing (200 trials [35.03 percent]) or half (371 trials [64.97 percent]). Compared with Ma’s (2024:19–20) findings, in which only 40.43 percent of trials were logically correct and more than half of those involved giving nothing, the current experiment demonstrates both improved logical accuracy and greater generosity. However, because of the simplified game setup in this study, the two results may not be directly comparable. Table A1 and Figure A1 in the online supplement show the distribution of all variables.

Table 1 presents the logistic regression results predicting the likelihood of transferring a nonzero amount in the dictator game. The DV is binary, with 0 indicating no

Table 1. Baseline: Logistic Regression Predicting Amount Transferred.

Variable	Coefficient	95% CI	<i>p</i>
Give framing	1.059	[.696, 1.423]	<.001
Stranger meet	.815	[.439, 1.192]	<.001
Female	.211	[-.153, .575]	.256
Age	.001	[-.015, .017]	.901
Intercept	.640	[-.046, 1.326]	.068

Note: $N = 571$; pseudo- $R^2 = .077$. Dependent variable: amount transferred binary (0 for nothing transferred, 1 for nonzero amount transferred). CI = confidence interval.

transfer and 1 indicating a transfer of a nonzero amount. The coefficient for *give framing* is 1.059 ($p < .001$), indicating that agents exposed to a “give” framing are approximately 2.88 times ($\approx e^{1.059}$) more likely to transfer a nonzero amount compared with those exposed to a “take” framing. This finding, although seemingly straightforward, provides a crucial validation of our experimental setup. The fact that the model is highly sensitive to this fundamental framing of the task confirms that the LLM has learned the basic structure of the dictator game and is responding in a way that is consistent with the experimental manipulations. This serves as a critical baseline, demonstrating that the model is “paying attention” to the experimental conditions, which gives us confidence in the more nuanced findings that follow. Similarly, the coefficient for *stranger meet* is 0.815 ($p < .001$), suggesting dictators are about 2.26 times ($\approx e^{0.815}$) more likely to make a nonzero transfer to recipients if they meet after the game, compared with having no interaction at all.

The variable *female* has a positive but nonsignificant coefficient (0.211, $p = .256$), suggesting that LLM agents with a “female” persona may be 1.24 times ($\approx e^{0.211}$) more likely to transfer a nonzero amount than “male” agents; however, this effect is not statistically significant. Similarly, *age* shows a negligible effect on transfer likelihood (coefficient = 0.001, $p = .901$).

The intercept coefficient (0.640, $p = .068$) reflects the baseline log-odds of transferring a nonzero amount, with an odds ratio of 1.90 ($\approx e^{0.640}$). This indicates that LLM agents are prosocial by default, being nearly twice as likely to transfer a nonzero amount than nothing. Overall, the model accounts for a modest proportion of variance in the DV (pseudo- $R^2 = .077$).

Baseline: Computational Basis of Variables

The descriptive and regression results provide a behavioral evaluation of the LLM’s decision-making process. To explore its internal mechanisms, I next examine the representations of IVs, DVs, and their relationships through LLM residual streams. Figure 3 illustrates the relationship between IV steering vectors and the final decision vector across all layers of the LLM model. Remember that these relationships index within-model geometry (alignment and magnitude) rather than human-level construct validity.

Cosine similarity (Figure 3a) measures the alignment between IV and DV steering vectors, indicating the degree to which they are directionally consistent within the LLM's internal semantic space. High cosine similarity values (close to 1) imply a strong directional alignment, signifying the IV has a significant positive influence on the decision-making process. Conversely, values close to -1 indicate the vectors point in exactly opposite directions, meaning the IV has an inverse influence on the decision. Values near 0 suggest minimal or no directional influence of the IV on the decision.

Dot product (Figure 3b) quantifies the overall strength of each IV's influence on the decision vector. Higher dot product values indicate a greater magnitude of influence, reinforcing the observations from cosine similarity regarding which IVs have more of an effect. *Partial dot product* (Figures 3c and 3d), by isolating the unique contribution of each IV, reveals how much each variable independently affects outputs, controlling for the influence of other IVs.

Cosine Similarity: Alignment between IV and DV Steering Vectors

Figure 3a presents the cosine similarity values between IV and DV steering vectors across all layers of the LLM. The results reveal several key patterns. First, the *directional alignment* between IV and DV steering vectors remains relatively *consistent* across different layers of the language model. This suggests that, at each stage of processing, the model's internal semantic space represents these variables and their relationships in a coherent and stable manner.

Second, the *alignments* exhibit more variations in the earlier layers (layers 1–20) but stabilize in the later layers (layers 20–31). This indicates that initial processing stages are more sensitive to the influence of IVs, and deeper layers refine these representations, resulting in more stable and consistent alignments with the decision-making process.

Third, the *give–take framing* exhibits the highest cosine similarity (close to 1) with the DV steering vector, indicating that this variable has the most consistent and strongest directional influence on the model's decision-making process.

Fourth, the *female* variable shows a moderate cosine similarity (ranging between 0.75 and 0.25) with the DV steering vector. This suggests that *gender information* in the LLM's internal representation has a discernible but less dominant effect on the final decision compared with framing.

Fifth, the cosine similarity of the *meet–stranger* variable with the DV steering vector is low in the early layers but gradually increases in the later layers. This indicates that the model's representation of *meeting conditions* becomes more aligned with the decision-making process as it processes information through deeper layers.

Sixth, the *age* variable exhibits high directional alignment with the DV steering vector in the first two layers but sharply decreases to relatively low and gradually decreasing levels across the later layers. This suggests that *age information* plays a significant role in the initial stages of decision making but becomes less influential as processing progresses.

In general, certain IVs, particularly those related to the framing of the game and demographic attributes, exert varying levels of influence on the decision-making process across different layers of the LLM. The stronger alignment of framing variables such as *give–take* underscores their pivotal role in guiding the model’s decisions. Demographic factors such as gender and age exhibit more nuanced influences that evolve through the model’s layers.

Dot Product and Partial Dot Product: Magnitude of IV Influence

The results of the dot product and partial dot product, presented in Figures 3b to 3d, provide additional insights into the magnitude of each IV’s influence on the decision vector and the unique contribution of each IV to the decision-making process. Specifically, looking at the *cumulative influence across layers*, the magnitude of both the dot product and partial dot product increases progressively across the layers. This indicates that the influence of IVs accumulates as information flows through the model’s layers, leveraging the residual connections to amplify their effect on the final decision.

Turning to the *dominant influence of framing*, the influence of the *give–take* framing is significantly higher than other variables when measured by the dot product. Even after controlling for the influence of other variables, as shown by the partial dot product, the *give–take* framing remains highest. In contrast, the influence of other variables becomes minimal, highlighting the dominant role of framing in the LLM’s decision-making process.

Dissociation between Alignment and Magnitude

The analysis of the computational basis of variables reveals a noticeable dissociation between the alignment (cosine similarity) and the magnitude (dot product and partial dot product) of IV influence. Specifically, the directional alignment of IVs with the DV does not always correspond to their effect on final outcomes. For example, gender persona moderately aligns (i.e., cosine similarity) with the decision vector (Figure 3a, blue line), but its final effect (i.e., dot product) on the decision is miniscule (Figure 3d, blue line). This indicates that both the direction and magnitude of IV steering vectors play a role in shaping the decision-making process.

RESULTS: MANIPULATION ANALYSIS

Model Statistics of Manipulations

I conducted 1,891 manipulations (61 injection coefficients \times 31 layers⁴). For each manipulation, I performed 1,000 trials and subsequently conducted regression analyses using only the logically correct trials to obtain the regression coefficients for all IVs and overall model statistics (Figure 5 annotates an example manipulation with $\alpha = 30$ and $\ell = 1$).

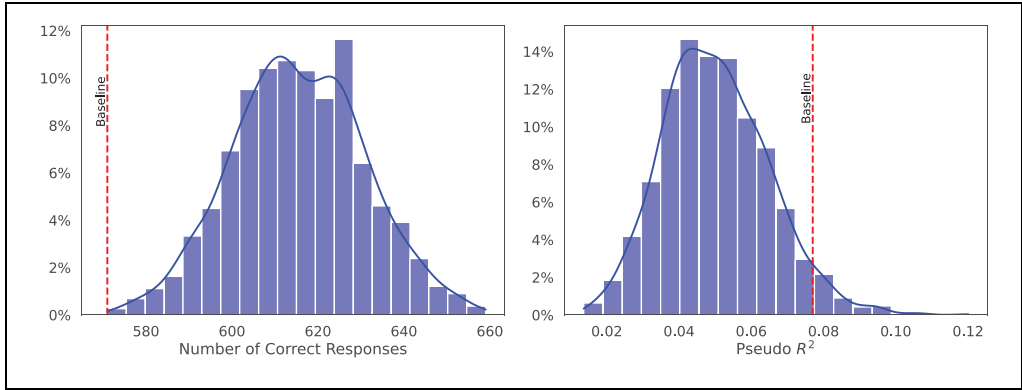


Figure 4. Histograms of manipulation-level metrics.

Note: Left: count of logic-check-passing trials (data-quality metric). Right: pseudo- R^2 from the logistic model predicting nonzero transfer using the instrumental variables (IVs) (i.e., give/take, meet/stranger, female/male, and age). Higher left values indicate more usable trials; higher right values indicate stronger IV \rightarrow dependent variable association conditional on the IVs.

Figure 4 shows the distributions of manipulation-level metrics. The number of logic-check-passing responses (left) is approximately normal, ranging from 571 to 659 (mean = 615.47, s.d. = 15.33). Pseudo- R^2 (right) is also approximately normal, ranging from 0.014 to 0.120 (mean = .050, s.d. = .015). Relative to baseline (pseudo- $R^2 = 0.077$ with 571 usable trials), manipulations, on average, increase the share of logic-check-passing trials ($z = 2.9$) while lowering pseudo- R^2 ($z = 1.8$).

Why do more trials pass the logic check after injection? Our intervention gently “nudges” the model in the direction of making a clearer decision. This extra nudge helps the model stick to the game’s standard response format and do the simple arithmetic correctly. With fewer off-template explanations or calculation slips, more trials pass the logic check.

Why does pseudo- R^2 go down? The regression only uses the listed variables (give/take, meet/stranger, female/male, age) to explain whether the model makes any transfer. Our intervention adds an additional push toward deciding that is not recorded in those variables. Because part of the model’s behavior now comes from this extra push (not from the IVs), the IVs by themselves explain a smaller share of the variation in outcomes. That makes pseudo- R^2 lower.

Why can both happen at the same time? These two metrics measure different things. The logic check is a data-quality test (Did the model follow the rules and arithmetic?), whereas pseudo- R^2 asks how much of the outcome is explained by the IVs alone. A small decision-stabilizing nudge can make outputs cleaner and more consistent (raising pass rates), while also shifting some of the decision to something outside the IVs (lowering the amount the IVs can explain). Thus, better data quality and lower pseudo- R^2 can coexist without contradiction.

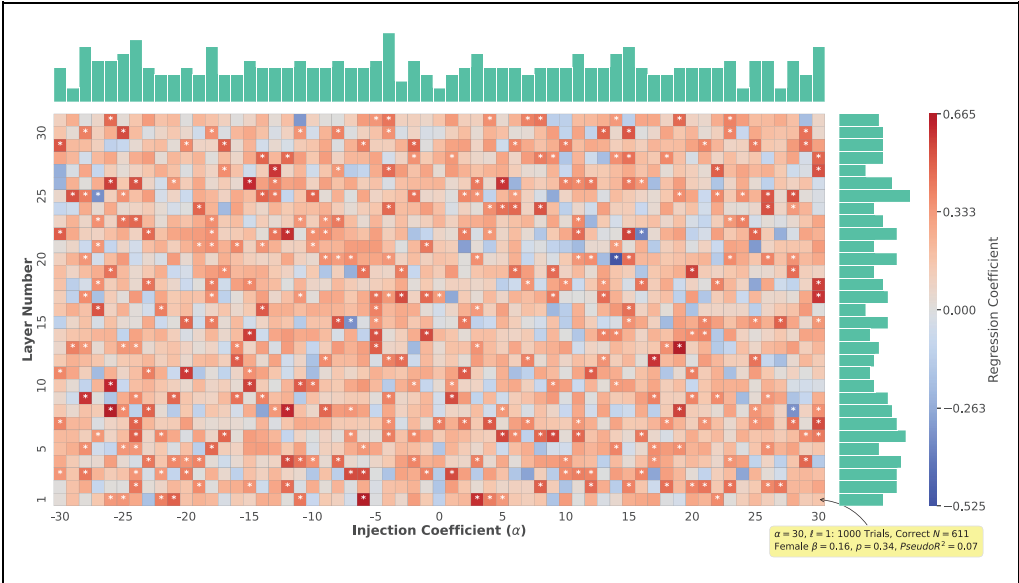


Figure 5. Regression coefficients of steered *female*.

Note: Each cell represents the regression coefficient of the *female* variable for a specific combination of layer number and injection coefficient. Cells marked with asterisks indicate coefficients that are statistically significant at the $p < .05$ level. The marginal histograms show the number of significant coefficients (i.e., the number of asterisks) by layer or injection coefficient.

Regression Coefficients of Steered “Female”

Figure 5 displays the regression coefficients of the steered *female* variable across all layers and injection coefficients. In this figure, each cell represents the regression coefficient of the *female* variable for a specific combination of layer number and injection coefficient (with the cell annotated for $\alpha = 30$ and $\ell = 1$ as an example). Regression coefficients that are statistically significant at the $p < .05$ level are marked with asterisks.

Among all 1,891 manipulations, 320 (16.92 percent) resulted in statistically significant regression coefficients for the *female* variable. Of these significant coefficients, 315 (98.44 percent) are positive, ranging continuously from 0.32 to 0.66. This indicates that a majority of the manipulations lead to an increase in the likelihood of transferring a nonzero amount, even when the injection coefficients are negative (meaning steered *away* from “female”). Because these coefficients span a continuous range, as Figure A2 in the online supplement illustrates, it is possible to select a specific manipulation to achieve a desired effect size.

The analysis of the steered *female* variable reveals that a substantial proportion of manipulations significantly influence the model’s decision-making process, aligning with our theoretical expectations. However, these manipulations do not exhibit clear patterns, and the majority of significant coefficients remain positive, even when the injection coefficient is negative. Theoretically, we would expect a stronger positive

correlation between the injection coefficient and the resulting regression coefficient, but the observed correlation is minuscule ($|r| < .05$). These findings suggest the model's internal representations involve complex interactions that are not fully captured by the current manipulation framework.

RESULTS: ORTHOGONALITY ANALYSIS

According to our theoretical framework, the injected IV steering vectors are orthogonal to one another. This orthogonality implies that the effects of manipulating different IVs should be independent. Specifically, the manipulation of the *female* variable should not influence the regression coefficients of other IVs.

Figures B1 to B3 in the online supplement present the orthogonality analysis. The regression coefficients of the *age* and *give–take* variables are rarely significantly affected by the manipulation of the *female* variable. This suggests the steering vectors of different IVs are largely orthogonal, supporting the independence of manipulations.

However, the coefficient of the *meet–stranger* variable is more frequently influenced by the manipulation of the *female* variable, indicating a potential interaction between these two variables. In human studies, multiple lines of research show that women tend to engage in more communal, relational forms of prosocial behavior, whereas men often exhibit more agentic, strength-intensive, or collectively oriented actions (Eagly 2009). Moreover, framing manipulations in dictator games suggest that women often show higher generosity toward strangers than do men in certain contexts (Chowdhury, Grossman, and Jeon 2020; Chowdhury, Jeon, and Saha 2017). Similarly, as social distance increases, women's generosity changes significantly more than does men's (Doñate-Buendía, García-Gallego, and Petrović 2022). These findings align with our observation that *meet–stranger* is more sensitive to *female* manipulations, suggesting the model may have captured the documented gender–context interplay.

We still observe 284 robust manipulations (i.e., manipulations in which only the *female* variable is significantly affected but not the other IVs), a substantial proportion of the successful manipulations of *female* (284 of 320 [88.75 percent]). This high rate of orthogonal manipulation implies that, despite some interaction between *female* and *meet–stranger*, the steering vectors remain largely independent in most cases. Theoretically, this underscores that gender-related shifts often intersect with social-context cues but can still be controlled in activation steering. Practically, it highlights the viability of using separate steering vectors to manipulate specific variables without extensively “spilling over” into others, an important consideration for applications requiring precise control over multiple attributes.

DISCUSSION

LLMs occupy a dual theoretical position in sociological research: they can be (1) repositories of culturally structured meanings acquired through large-scale training and (2) adaptive agents whose outputs can be experimentally probed as if they were synthetic participants (Anthis et al. 2025; Bail 2024; Ziems et al. 2024). These roles raise a core methodological challenge reviewed earlier: prevailing validation strategies (prompt

engineering, output benchmarking, *post hoc* statistical correction) treat models as opaque input-output devices, leaving unanswered whether their behavior reflects robust internal models of social meaning or merely brittle pattern matching (Aher et al. 2023; Ma 2024; Kozlowski and Evans 2025). Theoretical leverage for sociology comes from linking cultural schemas, role expectations, and framing effects to behavior. Doing so requires opening that black box to identify where and how socially meaningful distinctions (e.g., gender, age, moral frames) reside in internal activation space.

Our results operationalize this theoretical program: if social meanings are instantiated as approximately linear directions in the residual stream (Park et al. 2024; Tigges et al. 2023), then they can be (1) estimated via experimental contrasts, (2) disentangled via orthogonalization, and (3) causally manipulated through controlled injection. This reframes LLMs from inscrutable statistical artifacts into quantifiable internal representations whose geometry (angles, projections, magnitudes) encodes sociological variables, analogous to how network analysis formalized relational structure decades earlier (Arseniev-Koehler 2024; Borgatti et al. 2009). Activation engineering thus supplies a bridge between theories of meaning and computational implementation: it converts abstract constructs into manipulable vectors, enabling virtual experiments on internal representations rather than only surface responses.

This study introduces and validates a method for probing and steering internal representations of social concepts in LLMs. Combining classic experimental design with activation engineering, I isolate, measure, and manipulate the influence of specific social variables on model behavior, showing that variables are encoded with varying strength and depth and that steering can orthogonally control their effect on decisions. Methodologically, this contributes (1) an analytic strategy that moves beyond treating LLMs as “black boxes,” replacing sole reliance on prompt-based validation with direct inspection of internal processing for more rigorous, theory-grounded assessment of simulations, and (2) a practical tool for computational social scientists and industry applications: by operationalizing social concepts as vectors, researchers can run virtual experiments to test causal hypotheses in controlled settings, expanding theory building where conventional experiments are difficult or impossible.

For additional conceptual clarity and accuracy, my contribution concerns *internal representational validity* and controllability of LLMs, where social variables reside in activation space and how they can be causally steered. I do not claim *external sociological validity*; linking these internal directions to human cognition and behaviors is a separate empirical question requiring comparative studies.

Research Design Guidelines for Social Scientists

Activation-based steering can be treated as a research instrument rather than a black-box trick. It enables three concrete uses. First, *validity audits* assess whether core constructs (e.g., social distance, gendered prosociality, fairness) are robustly encoded as internal directions; if a construct is weakly represented, simulations relying on it are unlikely to be theoretically and empirically faithful. Second, *theory-driven experiments* inject small, controlled perturbations along a construct’s direction to test hypotheses

(e.g., whether gender cues increase prosocial transfers), identify boundary conditions, and generate new hypotheses. Third, *measurement* can compare construct directions across models and checkpoints to track “cultural drift” in what is encoded. These practices interrogate internal mechanisms and should complement, never replace, studies with human participants.

Key Steps: Variable-Specific Steering in Social Simulations with LLMs

The goal is precise, variable-specific control over LLM behavior so that simulated responses can align with observed social patterns of humans, achieving a higher degree of external validity. The following research-design framing aligns LLM simulation with standard social science practices: clear constructs, prespecified contrasts, conservative interventions, human benchmarking, transparency, and reproducibility.

1. *Clarify constructs and operationalizations*: Define the conceptual variables of interest (e.g., gender persona, social distance, framing) and link each to concrete prompt fields and outcome definitions. Prespecify value ranges and coding rules for inputs and the DV.
2. *Establish a baseline*: Run a representative or random set of prompts spanning the operational variables. Record behavioral outcomes and the internal signals needed for later analysis. This baseline functions as your “unmanipulated” distribution against which all interventions are judged.
3. *Isolate variables*: Disentangle overlapping constructs so that each variable’s effect can be examined on its own. Conceptually, you are removing shared variation to focus on the unique component of each variable. This supports interpretable, variable-specific tests.
4. *Prespecify interventions and success criteria*: Select layers to test and a conservative grid of intervention strengths. Start small and define *ex ante* success metrics, such as expected direction of change on the outcome, acceptable collateral movement on non-target variables, and minimal disruption to response quality.
5. *Run within-model experiments*: Inject the variable-specific signal during inference to test causal hypotheses inside the model. Keep prompts and decoding settings constant, and change one factor at a time to preserve experimental control.
6. *Benchmark against human data*: Triangulate manipulated outcomes with human benchmarks (e.g., preregistered replications, archival data sets). Choose manipulations that yield the closest alignment, transparently report mismatches, and treat any alignment as an approximate calibration, not proof of equivalence, because of external validity concerns.
7. *Documentation and governance*: Preregister contrasts and evaluation criteria where possible; release prompts, random seeds, LLM parameters, and intervention logs to support reproducibility. Note ethical safeguards (e.g., bounding interventions along sensitive demographic axes) and intended use constraints.

Example Modular Code Scripts

In addition to these guidelines, I released three core code scripts (<https://doi.org/10.17605/OSF.IO/J9RQK>) used in this study to facilitate verification, further research, and application of variable-specific steering methods. These scripts are modular and can be adapted for different models and experiments: `core_functions.py`, which contains all essential functions for model manipulation and analysis; `play_dictator_baseline_residuals.py`, which implements the baseline experimental setup for measuring residuals; and `play_dictator_steer_amount_gender_partial.py`, which conducts experiments on gender steering effects.

Industry Applications

The method introduced in this study may be valuable to entities that want flexible, precise governance over LLM outputs without the overhead of constant retraining. For example, the method can be applied to encourage prosocial behavior in AI-driven systems by selectively amplifying or diminishing certain variables, thereby promoting fairness or cooperation.

In marketing research, this method enables the creation of more valid social simulations by populating them with LLM agents whose behaviors can be manipulated. Researchers can construct virtual audience panels where latent traits such as “genre affinity,” “trailer responsiveness,” or “ticket-price sensitivity” are operationalized as distinct, manipulable steering vectors. This allows clean, virtual A/B tests that isolate the causal effect of specific marketing interventions. For example, one could predict a film’s opening-week theater viewing volume by varying the trailer cut or media mix while holding the “ticket-price sensitivity” vector constant, thereby eliminating confounding effects that often undermine prompt-only approaches. This technique facilitates rapid, low-cost testing of strategies (e.g., release timing, trailer variants, cast-focused creatives, geotargeted spending) on large, customizable synthetic cohorts. The most promising scenarios can then be validated with smaller, targeted field experiments, such as limited-market screenings or geo-split ad campaigns using admissions as the outcome, reducing costs and accelerating the greenlight and marketing optimization cycle.

For research service platforms, providers of online experimentation and survey services (e.g., Qualtrics and Prolific) could adopt this method to run “synthetic subject” studies, in which factors such as demographic attributes or question framing are precisely manipulated. This level of control enables novel experimental designs that would be challenging or impractical to conduct with human participants, offering researchers a more scalable and ethically flexible way to investigate complex behaviors.

In more entertainment-focused settings, role-playing and story-driven games can incorporate the method to generate richer, more responsive narratives. By injecting or subtracting specific steering vectors (e.g., “cautious” vs. “bold” traits), developers and players can fine-tune how AI-driven characters behave, speak, or make decisions, allowing deeply personalized storylines.

These examples highlight just a few of the many possibilities for applying the method in real-world contexts. By offering transparent, fine-grained control over the underlying variables that guide an LLM's behavior, this study's approach stands to benefit not only social scientists but also commercial applications seeking to manage complex AI-driven interactions responsibly and efficiently.

Future Directions

This study provides a foundational method for probing and steering social concepts within LLMs. To build on this foundation and move from a proof of concept to a robust tool for social science, future work could prioritize making the approach more powerful, accessible, and sociologically valid. I outline three key directions.

From Scripts to User-Friendly Tools. To make activation-based steering more accessible, a crucial next step is to develop a user-friendly software package. Such a package would abstract away low-level implementation details by providing a high-level application programming interface, support for multiple models, comprehensive documentation, and built-in visualization tools, thereby increasing accessibility for a broader range of researchers.

Expanding to Other Social Constructs and Experiments. The framework developed in this study is highly generalizable. With a user-friendly software package, future research could explore its application to a wide range of social constructs and experimental paradigms, such as moral psychology, political science, and organizational behavior, to build a more comprehensive understanding of how LLMs represent the social world.

Bridging the Gap between Internal and External Validity. This study focused on the *internal representational validity* of LLMs, but future work must establish *external sociological validity*. This requires systematic human–LLM comparisons, the use of activation-based steering for cognitive modeling, and large-scale validation studies to test if and how LLM behavior corresponds to that of human subjects.

Taken together, the results sketch a practical bridge between sociological theory and the internal geometry of modern language models. By working inside the model to measure, compare, and steer social meanings as directions in representation space, we turn LLMs from opaque black boxes into calibrated instruments for theory-testing, bounded synthetic data, and auditable interventions. The program outlined here is intentionally modular: each step invites open benchmarks, uncertainty estimates, and guardrails the community can scrutinize and improve. As these practices spread, debates about whether LLM agents “simulate people” will give way to cumulative evidence about which regularities they capture, where they fail, and how to align them with human values. The opportunity is not to replace judgment but to enhance it, making our models clearer, our inferences humbler, and our science faster.

Acknowledgments

I thank Chenxin Zhang, René Bekkers, and four anonymous reviewers for their constructive comments on earlier versions of this paper.


Funding

The author disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The project is partly supported by (1) Academic Development Funds from the RGK Center, (2) research funds from the Gradel Institute of Charity at the University of Oxford, and computing resources through (3) the Texas Advanced Computing Center at the University of Texas Austin (Keahey et al. 2020) and (4) Dell Technologies, Client Memory Team and AI Initiative PoC Lead Engineer Wenté Xiong.

Author's Note

The author declares that this study complies with required ethical standards. AI tools were used for grammar, proofreading, and clarity improvement only.

ORCID iD

Ji Ma  <https://orcid.org/0000-0003-3682-6587>

Supplemental Material

Supplemental material for this article is available online.

Notes

- 1 The method requires access to the model's internal weights and activations, so it is feasible only with open-weight models such as Llama and DeepSeek.
- 2 These numbers are logits before transformation with, for example, the softmax function.
- 3 Whether other variables are held fixed or randomized, the resulting difference in the residual streams $\mathbf{v}_X^{(\ell)}$ consistently captures the unique effect of the IV X . When other variables are fixed, changing X directly isolates its effect on the residual stream. When randomized, any influence other variables might have is averaged out across multiple instances, neutralizing their effect.
- 4 I injected at 31 internal layers (layers 1–31), excluding the final decoder layer ($L = 32$) to avoid direct perturbations at the logits head.

References

- Aher, Gati V., Rosa I. Arriaga, and Adam Tauman Kalai. 2023. "Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies." Pp. 337–71 in *Proceedings of the 40th International Conference on Machine Learning. International Conference on Machine Learning*. MLResearchPress. Retrieved February 1, 2026. <https://proceedings.mlr.press/v202/aher23a/aher23a.pdf>.
- Anthis, Jacy Reese, Ryan Liu, Sean M. Richardson, Austin C. Kozlowski, Bernard Koch, James Evans, et al. 2025. "LLM Social Simulations Are a Promising Research Method." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2504.02234>.
- Argyle, Lisa P., Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. "Out of One, Many: Using Language Models to Simulate Human Samples." *Political Analysis* 31(3):337–51.

- Arseniev-Koehler, Alina. 2024. "Theoretical Foundations and Limits of Word Embeddings: What Types of Meaning Can They Capture?" *Sociological Methods & Research* 53(4):1753–93.
- Atari, Mohammad, Mona Xue, Peter Park, Damián Blasi, and Joseph Henrich. 2023. "Which Humans?" PsyArXiv. Retrieved February 1, 2026. https://osf.io/5b26t_v1.
- Bail, Christopher A. 2024. "Can Generative AI Improve Social Science?" *Proceedings of the National Academy of Sciences* 121(21):e231402121.
- Banker, Sachin, Promothesh Chatterjee, Himanshu Mishra, and Arul Mishra. 2024. "Machine-Assisted Social Psychology Hypothesis Generation." *American Psychologist* 79(6):789–97.
- Bisbee, James, Joshua D. Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M. Larson. 2024. "Synthetic Replacements for Human Survey Data? The Perils of Large Language Models." *Political Analysis* 32(4):401–16.
- Borgatti, Stephen P., Ajay Mehra, Daniel J. Brass, and Giuseppe Labianca. 2009. "Network Analysis in the Social Sciences." *Science* 323(5916):892–95.
- Boutyline, Andrei, and Alina Arseniev-Koehler. 2025. "Meaning in Hyperspace: Word Embeddings as Tools for Cultural Measurement." *Annual Review of Sociology* 51:89–107.
- Broska, David, Michael Howes, and Austin van Loon. 2025. "The Mixed Subjects Design: Treating Large Language Models as Potentially Informative Observations." *Sociological Methods & Research* 54(3): 1074–1109.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models Are Few-Shot Learners." Version 4. arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2005.14165>.
- Carlsen, Hjalmar Bang, and Snorre Ralund. 2022. "Computational Grounded Theory Revisited: From Computer-Led to Computer-Assisted Text Analysis." *Big Data & Society* 9(1):20539517221080146.
- Chang, Samuel, Andrew Kennedy, Aaron Leonard, and John A. List. 2024. "12 Best Practices for Leveraging Generative AI in Experimental Research." NBER Working Paper No. 33025. Cambridge, MA: National Bureau of Economic Research.
- Chowdhury, Subhasish M., Philip J. Grossman, and Joo Young Jeon. 2020. "Gender Differences in Giving and the Anticipation Regarding Giving in Dictator Games." *Oxford Economic Papers* 72(3):772–79.
- Chowdhury, Subhasish M., Joo Young Jeon, and Bibhas Saha. 2017. "Gender Differences in the Giving and Taking Variants of the Dictator Game." *Southern Economic Journal* 84(2):474–83.
- Dai, Damai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, et al. 2024. "DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2401.06066>.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, et al. 2025. "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2501.12948>.
- Doñate-Buendía, Anabel, Aurora García-Gallego, and Marko Petrović. 2022. "Gender and Other Moderators of Giving in the Dictator Game: A Meta-Analysis." *Journal of Economic Behavior & Organization* 198:280–301.
- Eagly, Alice H. 2009. "The His and Hers of Prosocial Behavior: An Examination of the Social Psychology of Gender." *American Psychologist* 64(8):644–58.
- Edelmann, Achim, Tom Wolff, Danielle Montagne, and Christopher A. Bail. 2020. "Computational Social Science and Sociology." *Annual Review of Sociology* 46:61–81.
- Engels, Joshua, Eric J. Michaud, Isaac Liao, Wes Gurnee, and Max Tegmark. 2024. "Not All Language Model Features Are Linear." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2405.14860>.
- Fricke, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford, UK: Clarendon.
- Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. "Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes." *Proceedings of the National Academy of Sciences* 115(16): E3635–44.
- Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3):267–97.

- Gurnee, Wes, and Max Tegmark. 2024. "Language Models Represent Space and Time." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2310.02207>.
- Horton, John J. 2023. "Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?" NBER Working Paper No. 31122. Cambridge, MA: National Bureau of Economic Research.
- Johnson, Tim, and Nick Obradovich. 2023. "Evidence of Behavior Consistent with Self-Interest and Altruism in an Artificially Intelligent Agent." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2301.02330>.
- Keahey, Kate, Jason Anderson, Zhuo Zhen, Pierre Riteau, Paul Ruth, Dan Stanzione, Mert Cevik, et al. 2020. "Lessons Learned from the Chameleon Testbed." In *2020 {USENIX} Annual Technical Conference ({USENIX} {ATC} 20)*, pp. 219–233. Retrieved February 1, 2026. <https://www.usenix.org/system/files/atc20-keahey.pdf>.
- Kim, Junsol, James Evans, and Aaron Schein. 2025. "Linear Representations of Political Perspective Emerge in Large Language Models." arXiv. Retrieved February 1, 2025. <https://arxiv.org/abs/2503.02080>.
- Korinek, Anton. 2025. "AI Agents for Economic Research." NBER Working Paper No. w34202. Cambridge, MA: National Bureau of Economic Research.
- Kozlowski, Austin C., and James Evans. 2025. "Simulating Subjects: The Promise and Peril of Artificial Intelligence Stand-Ins for Social Agents and Interactions." *Sociological Methods & Research* 54(3): 1017–73.
- Kozlowski, Austin C., Matt Taddy, and James A. Evans. 2019. "The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings." *American Sociological Review* 84(5):905–49.
- Lake, Brenden M., Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. 2017. "Building Machines That Learn and Think Like People." *Behavioral and Brain Sciences* 40:e253.
- Leng, Yan, and Yuan Yuan. 2024. "Do LLM Agents Exhibit Social Behavior?" arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2312.15198>.
- Lepikhin, Dmitry, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, et al. 2020. "GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2006.16668>.
- Ma, Ji. 2024. "Can Machines Think Like Humans? A Behavioral Evaluation of LLM-Agents in Dictator Games." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2410.21359>.
- Mei, Qiaozhu, Yutong Xie, Walter Yuan, and Matthew O. Jackson. 2024. "A Turing Test of Whether AI Chatbots Are Behaviorally Similar to Humans." *Proceedings of the National Academy of Sciences* 121(9):e2313925121.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/1301.3781>.
- Mostafavi, Moeen, Michael D. Porter, and Dawn T. Robinson. 2025. "Contextual Embeddings in Sociological Research: Expanding the Analysis of Sentiment and Social Dynamics." *Sociological Methodology* 55(1):25–58.
- Nelson, Laura K. 2020. "Computational Grounded Theory: A Methodological Framework". *Sociological Methods & Research* 49(1):3–42.
- Park, Kiho, Yo Joong Choe, and Victor Veitch. 2024. "The Linear Representation Hypothesis and the Geometry of Large Language Models." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2311.03658>.
- Rister Portinari Maranca, Alessandra, Jihoon Chung, Musashi Hinck, Adam D. Wolsky, Naoki Egami, and Brandon M. Stewart. 2025. "Correcting the Measurement Errors of AI-Assisted Labeling in Image Analysis Using Design-Based Supervised Learning." *Sociological Methods & Research* 54(3): 984–1016.
- Roberts, Margaret E. 2016. "Introduction to the Virtual Issue: Recent Innovations in Text Analysis for Social Science." *Political Analysis* 24(V10):1–5.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016. "Neural Machine Translation of Rare Words with Subword Units." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/1508.07909>.

- Steck, Harald, Chaitanya Ekanadham, and Nathan Kallus. 2024. "Is Cosine-Similarity of Embeddings Really about Similarity?" Pp. 887–90 in *Companion Proceedings of the ACM Web Conference 2024: WWW '24*. New York: Association for Computing Machinery.
- Stoltz, Dustin S., and Marshall A. Taylor. 2024. *Mapping Texts: Computational Text Analysis for the Social Sciences*. Oxford, UK: Oxford University Press.
- Subramani, Nishant, Nivedita Suresh, and Matthew E. Peters. 2022. "Extracting Latent Steering Vectors from Pretrained Language Models." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2205.05124>.
- Tigges, Curt, Oskar John Hollinsworth, Atticus Geiger, and Neel Nanda. 2023. "Linear Representations of Sentiment in Large Language Models." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2310.15154>.
- Turner, Alexander Matt, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2024. "Steering Language Models with Activation Engineering." Version 5. arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2308.10248>.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, et al. 2017. "Attention Is All You Need." *Advances in Neural Information Processing Systems* 30. Retrieved February 1, 2026. https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- Wang, Angelina, Jamie Morgenstern, and John P. Dickerson. 2025. "Large Language Models That Replace Human Participants Can Harmfully Misportray and Flatten Identity Groups." *Nature Machine Intelligence* 7(3):400–11.
- Xie, Chengxing, Canyu Chen, Feiran Jia, Ziyu Ye, Kai Shu, Adel Bibi, Ziniu Hu, et al. 2024. "Can Large Language Model Agents Simulate Human Trust Behaviors?" arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2402.04559>.
- Zhou, Yangqiaoyu, Haokun Liu, Tejes Srivastava, Hongyuan Mei, and Chenhao Tan. 2024. "Hypothesis Generation with Large Language Models." arXiv. Retrieved February 1, 2026. <https://arxiv.org/abs/2404.04326>.
- Ziems, Caleb, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. "Can Large Language Models Transform Computational Social Science?" *Computational Linguistics* 50(1): 237–91.

Author Biography

Ji Ma is an assistant professor in nonprofit and philanthropic studies at the Lyndon B. Johnson School of Public Affairs and RGK Center at the University of Texas at Austin. He is also an affiliated faculty member of the Center for East Asian Studies and School of Information at the University of Texas at Austin and a visiting researcher of the Gradel Institute of Charity, New College, University of Oxford. His research and teaching focus on state–civil society relations, knowledge production for evidence-based policymaking and nonprofit studies, and computational social science methods.