

Using Agents to Automate Audits of Personalization Algorithms at Scale

Alessandro Morosini^{*,1} Sarah H. Cen^{*,2} Andrew Ilyas²
Hedi Driss¹ Aleksander Madry¹ Chara Podimata¹
{morosini, hdriss, madry, podimata}@mit.edu
{sarahcen, andrewi}@andrew.cmu.edu

April 21, 2026

Abstract

Social media platforms use personalization systems to decide what content each user sees. Auditing these systems is essential for understanding their role in shaping information exposure across different user populations. In practice, auditing involves creating user accounts that interact with the platform and analyzing differences in what the algorithm recommends to each of them. This is usually done either through human-based audits, which recruit real users to participate in studies, or sock-puppet audits, which deploy synthetic accounts controlled by researchers. Yet, both approaches face a fundamental limitation: they confound algorithmic responses to user *attributes* (e.g., age) with responses to user *behavior* (e.g., what content they like). We present an auditing framework that addresses this limitation by using Large Language Models (LLMs) as behavioral engines for synthetic accounts. Our design decouples behavior from platform-visible signals by defining “personas” that capture both demographic and ideological characteristics grounded in U.S. survey data. Within each persona, we hold the LLM prompt (and thus, the *behavioral policy*) fixed while randomizing demographic attributes via profile perturbations. This enables identification of the causal effect of user attributes on algorithmic content allocation under black-box access. As a case study, we deploy 1,120 agents on \mathbb{X} spanning 14 personas and 3 counterfactual conditions, collecting more than 200,000 exposures during the 2024 U.S. election period. Our analysis finds evidence that the algorithm responds differently to user demographics, but these effects are persona-dependent rather than systematic, with the direction of effects varying across user types. Confirming past findings and widespread anecdotal evidence, our results also suggest that \mathbb{X} ’s recommendation algorithm exposes users to significantly more toxic, polarizing, and right-leaning content relative to its chronological counterpart, with the degree of this amplification varying sharply by user ideology. Our work establishes LLM-based agents as a new methodological tool for algorithmic auditing, enabling causal identification that was previously infeasible while also providing platforms with a practical tool for sandbox testing recommendation systems.

1 Introduction

Millions of people around the world interact with *personalization systems* (e.g., social media and AI platforms) on a daily basis. These systems use algorithms to tailor content, recommendations, and

¹Massachusetts Institute of Technology

²Carnegie Mellon University

*Equal contribution.

responses to individual users based on their attributes, preferences, and past behavior. TikTok has more than 1.5 billion monthly active users, X (formerly Twitter) has more than 500 million, and Instagram has 3 billion monthly [68]. Meanwhile, ChatGPT, one of the fastest-growing AI applications powered by LLMs, boasts approximately 800 million active weekly users [38]. These systems rely on complex and constantly evolving algorithms to determine what content to recommend to users, what information to prioritize, and how users engage with digital spaces. Through millions of micro-decisions each day, these systems gradually change the fabric of our society by shaping public discourse, influencing our worldviews, and amplifying certain cultural and political narratives over others. Yet, for all their far-reaching influence, the mechanisms behind these algorithmic decisions remain largely opaque, leaving policymakers, researchers, and the public with limited visibility into *how these systems operate* and *what values they embed*.

In Public Policy language, policymakers and researchers would like to “audit” these algorithms. Auditing—the process of evaluating the properties of a system—has a long and established history, especially in the financial and housing sectors (see e.g., [71]). Recently, the idea of auditing has also been extended to algorithmic systems, where the goal is to evaluate how algorithms make decisions and whether those decisions align with public interests [47]. However, algorithmic auditing fundamentally differs from traditional auditing in that auditors typically lack access to the platform’s internal systems. This lack of access has made algorithmic auditing a pressing policy concern: Congress has convened multiple hearings to question the CEOs of Meta, OpenAI, and other platforms about how their algorithms work and what kinds of decisions are driven by automated processes versus direct human oversight [21, 53, 65]. These hearings reflect a growing recognition that algorithmic systems are not just technical artifacts, but are critical components affecting the political stability of our society.

In general, the more access an auditor has to a system, the more powerful the audit can be. For example, a financial audit entails reviewing *all* financial records and transactions of the audited entity, a level of access that makes such audits highly effective [39]. However, even in regulated industries like banking and housing, where such access is legally mandated, firms routinely resist disclosing proprietary information and internal processes [55]. This problem is even more acute for social media algorithms and AI platforms, where companies are *extremely protective* of their internal systems, which are proprietary and among their most critical competitive assets [55, 56]. At the same time, tech leaders themselves (including OpenAI’s CEO Sam Altman) have repeatedly expressed openness to external auditing and oversight, provided of course that such processes do not stifle innovation or expose sensitive intellectual property [48]; and even apart from their public statements, almost every tech company nowadays has robustness and safety teams in house, studying the potential negative effects of these algorithmic systems on society. The growing demand for algorithmic accountability among policymakers and researchers, coupled with the need to preserve innovation, creates a policy gap: *how can we audit algorithms without requiring direct access to their internal workings?*

The natural answer is *black-box* auditing: evaluating algorithmic systems solely through their observable inputs and outputs, without access to internal model parameters or training data. Yet, despite its importance, black-box auditing is difficult to implement correctly. Effective algorithmic audits of social media and AI platforms require not just *static* evaluation, but *dynamic* counterfactual analysis. “How would the algorithm have behaved if the input (e.g., human characteristics or revealed behavior) were *slightly* different?”; “Would it have recommended different content to a different type of user?”; and “How much of the observed outcome (e.g., toxicity or political leaning of recommended content) is driven by the algorithm itself versus shifts in user behavior, platform policies, or external events?” These are only a small sample of the questions a successful audit must answer, and have monopolized policymakers’ interest in public hearings [21, 53, 65].

Traditionally, counterfactual analysis in the context of algorithmic auditing would involve exposing real human users to modified platform conditions or curated content streams and observing their responses [14, 76]. However, this approach has several flaws that make it unsuitable for auditing at the scale required for modern platforms. First, human-centered counterfactual analysis is not clean from a causal inference perspective: users arrive at platforms with pre-existing biases, preferences, and exposures that auditors cannot fully control, making it impossible to disentangle algorithmic effects from users’ prior beliefs and external influences. In short, humans are *noisy counterfactual subjects*, which undermines the reliability of any inferences drawn from such experiments. Second, humans are expensive to recruit, monitor, and study, particularly at the scale necessary to audit platforms with billions of active users and millions of content variations. This makes large-scale audits cost-prohibitive, particularly for independent researchers, public interest groups, and even regulatory agencies. And finally, human users themselves may alter their behavior when they know they are being studied (aka the *Hawthorne effect*), introducing yet another source of bias into the audit process.

More recent studies use “sock-puppet” audits, where controlled accounts, typically implemented as computer programs, interact with the platform according to *predefined rules* specified by the researcher [7, 33, 79]. Sock-puppet audits address some of the limitations introduced by observing human behavior in these systems but face another fundamental limitation. The sock-puppet accounts are limited in their behavior to the specific *rules* that they have been programmed to follow, thus exhibiting very predictable (also almost “canned”) behavior when interacting with content.

This is the methodological gap that our paper aims to fill. Specifically, we propose to **conduct black-box auditing of personalization systems using LLM-powered agents**. We develop a general framework for this approach and present a large-scale case study on X’s recommendation algorithm during the 2024 U.S. presidential election. Unlike traditional sock puppets, which follow scripted rules, LLM-powered agents reason about each piece of content before deciding whether to interact with it, producing a realistic feedback loop between user actions and platform recommendations. At the same time, because agent personality is specified by a prompt, researchers can define diverse agent types while maintaining full experimental control. By design, this decouples agent behavior, encoded in the prompt, from platform-visible attributes. This enables counterfactual analysis: by holding the prompt fixed while varying a single attribute, we can identify how the algorithm responds to user characteristics. Next, we outline our contributions in detail.

1.1 Our Contributions

A Black-Box Auditing Framework at Scale. We formalize black-box auditing of personalization systems as the identification of *behavior-conditional* signal responses, and introduce a framework for black-box auditing using LLM-powered agents. Each agent executes a fixed behavioral policy encoded in a prompt that describes its *persona*. Each persona is characterized by demographic and political ideology attributes, grounded in U.S. Census data [74] and Pew’s Political Typology [58]. Crucially, both the attributes and the behavior are defined *a priori*, not shaped by platform interaction. This has two consequences. First, because user attributes such as ideology are specified before any engagement with the platform, researchers can stratify algorithmic outcomes by user characteristics, usually confounded in human-based audits and less realistic with scripted bots. Second, because the behavioral policy is held constant across accounts within a persona, auditors can randomize platform-visible “signals” (attributes that a platform can either observe directly or infer confidently, e.g. a user’s age) while keeping agent behavior unchanged: this enables causal identification of the platform’s response to demographic attributes. Our methodology is general and can be adapted to a broad class of algorithmic auditing settings.

Field Experiment on X. We present the first large-scale deployment of LLM-powered agents for auditing on a live social media platform. We deploy 1,120 agents on X¹ over the 24 days following the 2024 U.S. presidential election (November 5–28), a period of heightened political salience. Each agent is grounded in one of 14 personas spanning a range of demographic and ideological profiles. This yields 80 accounts per persona randomly assigned across four conditions (one baseline and three single-attribute counterfactuals over age, gender, and location). Each account browses both the algorithmically curated *For You* feed and the reverse-chronological *Following* feed. The divergence between these two feeds has attracted growing attention from researchers and the press alike [23, 28, 36]. In each session, agents observe recommended posts and choose among liking, following the author, reading replies, or ignoring. All these actions are selected by the LLM based on the post’s content and the agent’s persona. In total, the deployment yields over 200,000 account-post exposures, each annotated with full text, author metadata, engagement counts, and logged agent actions and reasoning traces.²

Evidence of Differential Exposure on X. On the aggregate user level, we find that X’s recommendation algorithm systematically amplifies toxic, polarizing, and right-leaning content in the *For You* feed relative to the *Following* feed, while left-leaning content is not significantly amplified. Stratifying by user ideology reveals even sharper ideological asymmetries: for example, right-leaning content is amplified for both user groups, while left-leaning content is actively suppressed for right-leaning users. On top of this, our counterfactual design provides causal evidence that the algorithm responds to demographic signals, but not uniformly: pooled effects are largely null, yet joint tests reject homogeneity for all outcomes, and persona-level significant effects are roughly double what chance alone would predict. The platform’s response to demographic attributes is thus highly persona-dependent: the same signal perturbation can amplify or de-amplify content depending on the user type. Such heterogeneity would be missed by an aggregate study, even one that successfully decouples behavior from signals.

Release of Dataset. We release the dataset from our deployment: over 200,000 account-post exposures collected over 24 days around the 2024 U.S. election, each annotated with post text, author metadata, feed context, content labels, and the experimental design (persona and condition assignments). To our knowledge, this is the first publicly available large-scale dataset from a live social media platform that pairs algorithmic and chronological feed exposures for the same accounts. The dataset also contains logged reasoning traces and engagement actions of our LLM agents, recording how they form preferences over real content and how those preferences vary with ideology, demographics, and topic. This provides a resource both for researchers studying LLM-simulated human behavior and for future auditing work.

The rest of the paper is organized as follows. Section 1.2 presents the related work. Section 2 formalizes the auditing problem and defines the target estimand. Section 3 describes the experimental design and deployment. Section 4 presents the measurement and estimation framework. Section 5 reports our findings. Section 6 discusses the future directions and implications of our work.

¹Our team obtained an IRB exception for the experiment. Our LLM agents only engage with existing content and do not post original content. This minimizes impact on other users’ experience.

²A note on the legality of scraping from X: A 2018 court decision (see <https://www.aclu.org/documents/sandvig-v-barr-memorandum-opinion>) allowed researchers to violate platforms’ terms of service for research purposes. Our team also worked closely with a law clinic to verify the legality of our approach.

1.2 Related Work

Algorithmic Auditing. A broad literature has highlighted the societal importance of algorithmic content curation on search platforms and social media, given its influence on exposure to information [5, 24, 27, 40, 57]. This influence extends to high-stakes domains such as political elections, where algorithmic curation may shape voter information environments [46, 50, 79, 80]. Algorithmic auditing has emerged as a key methodology for evaluating such systems, particularly when internal models and data are unavailable. However, most real-world audits operate under black-box access due to legal limitations and lack of platform cooperation [64], raising challenges including limited observability, difficulties constructing representative samples, and risks of biased inference [6, 17]. This has motivated calls for principled auditing frameworks that improve transparency, reproducibility, and robustness [37, 60].

Human-Based Audits. Multiple studies have examined recommender systems using audits that rely on real users, with varying degrees of experimental control. Platform-run randomized experiments offer the strongest identification but are rare and subject to platform discretion; examples include X’s A/B test comparing algorithmic and chronological feeds [36] and Meta’s field experiments which manipulated exposure to algorithmic curation [29, 30, 52]. Without platform cooperation, researchers have recruited participants into quasi-experimental designs, such as switchback experiments [76] or browser-based interventions [59]. Alternative approaches rely purely on observational data collected via scraping [62] or from consenting users [14, 19]. In general, human-based audits are difficult to scale and do not allow auditors to construct counterfactuals over user profiles or exposure histories, nor to stratify outcomes by latent user characteristics that are endogenous to platform use.

Sock-Puppet Audits. More recent work focuses on sock-puppet audits, in which researchers deploy synthetic user accounts to study recommender system behavior. Existing approaches fall into two main categories. In the first, puppets are trained to a user profile but remain passive, avoiding engagement actions: Bartley et al. [8] and Ye et al. [79] define political orientation of a puppet through curated follow networks, while Haroon et al. [33] seed puppets with ideological watch histories and track subsequent recommendations. In the second category, puppets take actions, generally according to predetermined rules: Boeker and Urman [13] and Srba et al. [67] deploy agents on TikTok and YouTube that navigate content via scripted sequences; Bartley et al. [9] and Chen et al. [20] configure bots to perform simple interactions. In both categories, puppets do not reason about individual items before deciding whether to interact. Moreover, user attributes and behavior remain tightly coupled, making it difficult to isolate the causal effect of specific characteristics. Our work addresses both limitations: first, LLM-based agents reason about content before engaging, capturing realistic feedback dynamics without needing to follow hard-coded rules of engagement; second, our experimental design decouples user attributes from behavior, enabling counterfactual auditing of attribute-specific platform responses.

Silicon Sampling. A parallel line of research examines whether LLMs can approximate human behavior in experimental settings, sometimes referred to as “silicon sampling” [3] or “homo silicon” [25]. Argyle et al. [3] show that GPT-3, when conditioned on demographic backstories, can approximate human opinion distributions across political and social questions. Filippas et al. [25] demonstrate that LLMs can replicate classic behavioral findings at low cost and with flexibility in varying experimental conditions. Other work has tested LLM behavior in economic games [1, 45] and behavioral biases [41, 42]; employed LLMs for market research [16]; proposed frameworks for accelerating social science research [35, 44, 73]. See Anthis et al. [2] for a recent overview of LLM-based social simulations as a research method. While a growing literature documents limi-

tations in the ecological validity of LLM-simulated experiments [12, 26, 31], our design mitigates these concerns: personas are defined a priori from survey data rather than generated by the model [3, 43], and the LLM serves as a fixed behavioral policy whose potential biases cancel out under within-persona differencing.

2 Problem Setting and Notation

We model the platform as an unknown allocation rule that maps user characteristics and interaction history to recommended content, similarly to [18]. We aim to isolate how the platform’s content allocation changes when *platform-visible user signals* (e.g., declared profile fields and other attributes the platform can infer) are modified, holding the user’s *behavior* fixed. We refer to this as *counterfactual auditing*.

We consider a platform that serves content to a user over a finite time horizon T , which may correspond to a fixed deployment window or a number of total interactions. Let \mathcal{Z} denote the space of content items (e.g., posts), and let \mathcal{B} denote the set of user actions (e.g., like or follow). Each user account $i \in [n]$ is associated with a time-invariant signal vector $S_i \in \mathcal{S}$ representing user attributes visible to or inferable by the platform (e.g., user’s declared age). At each time step $t \in [T]$, the platform serves content $Z_{i,t} \in \mathcal{Z}$ and the user responds with an action $B_{i,t} \in \mathcal{B}$. Let $H_{i,t} = \{(Z_{i,1}, B_{i,1}), \dots, (Z_{i,t}, B_{i,t})\}$ denote the interaction history for user account i up to time t .

We model the platform’s recommendation algorithm as an unknown black-box mapping $\mathcal{A} : (S_i, H_{i,t-1}) \mapsto \mathcal{D}(\mathcal{Z})$, where $\mathcal{D}(\mathcal{Z})$ denotes a distribution over content items. User behavior is generated by a *policy* π : a complete behavioral specification that encodes the user’s full type, including both demographic characteristics and latent attributes such as political ideology or education level (in our implementation, π corresponds to the LLM prompt). Given a certain content item $Z_{i,t}$ (e.g., a post or a tweet), the policy induces a distribution over actions: $B_{i,t} \sim \pi(Z_{i,t})$. The auditor summarizes each trajectory via an outcome $Y_i = g(Z_{i,1:T}, B_{i,1:T})$. For example, g may capture properties commonly studied in auditing, such as the toxicity score, the share of political content, and the share of left- versus right-leaning content.

Policy-Conditional Signal Response. Fix a user policy π and consider two different values $s_0, s_1 \in \mathcal{S}$ of a platform-visible signal (e.g., π might encode a progressive user, while s_0 and s_1 correspond to profiles displaying ages 25 and 55 to the platform, respectively). We define the policy-conditional signal response as the difference in expected outcomes when *the same policy* π is paired with different signals:

$$\Delta_\pi(s_1, s_0) = \mathbb{E}[Y \mid S = s_1, \pi] - \mathbb{E}[Y \mid S = s_0, \pi], \quad (1)$$

where expectations are over the distribution of histories induced by (\mathcal{A}, π, s_1) and (\mathcal{A}, π, s_0) , respectively. Conceptually, Δ_π measures how the *expected outcome differs* (e.g., the share of toxic content) if the *same behavioral policy* were exposed to the platform under *different signal* assignments.

Measuring Δ_π in Eq. (1) via black-box access is difficult for two reasons. First, black-box access precludes counterfactual simulation: the auditor cannot re-query the algorithm under modified inputs and must instead observe realized trajectories from deployed accounts. Second, even if counterfactuals were observable, existing audit designs fail to decouple *behavior* from *platform-visible signals*. As a result, current audits compare users that differ simultaneously in both, estimating

$$\mathbb{E}[Y \mid S = s_1, \pi = \pi_1] - \mathbb{E}[Y \mid S = s_0, \pi = \pi_0], \quad (2)$$

where $\pi_1 \neq \pi_0$, rather than the policy-conditional signal response that holds π fixed. This coupling arises naturally in human audits because for an individual, the behavior (policy) and its induced

histories are correlated with S . It also arises in sock-puppet audits because intended “types” are usually instantiated through scripted behaviors, where the policy explicitly depends on S .

The comparison in Eq. (2) is not invalid, but it conflates the platform’s response to user signals with its response to user behavior, limiting our causal understanding of personalization systems. Addressing these limitations requires a design that decouples *behavior* from *signals* and enables counterfactual comparison across different signal assignments; we present such a framework next.

3 Experiment Design

Our experimental design centers on deploying LLM-powered agents that interact with the platform’s recommendation system in a closed feedback loop. Unlike common sock puppets, which follow scripted rules of engagement, our agents *reason* about each piece of content before deciding how to act. Each agent is instantiated with a fixed prompt, a *persona*, that encodes both demographic attributes (e.g., age, gender, location, race) and political ideology. Given this persona, the LLM-powered agent determines autonomously how to respond to each piece of content it encounters, with no further intervention from us. This produces a realistic feedback loop between user actions and platform recommendations, where behavior follows from a fixed *identity* rather than hard-coded rules. In the language of Section 2, the persona prompt captures the user’s policy π .

Because the behavior policy is held constant across replicas of the same persona, the design supports within-persona randomization of platform-visible demographic signals. This allows us to isolate the causal effect of attributes such as age, gender, or location on algorithmic recommendations, the policy-conditional signal response defined in Equation (1). Our methodology is general and applies whenever (i) a platform serves a sequence of content items to an account, (ii) recommendations depend on platform-visible account signals in addition to interaction history, and (iii) an auditor can deploy multiple accounts and log realized outcomes. An overview of the framework is presented in Figure 1.

3.1 Overview

We introduce *personas*, indexed by $p \in [P]$, each associated with a fixed behavior policy π_p . Accounts are indexed by $i \in [n]$ and we let \mathcal{I}_p be the set of users associated with persona p . Each account $i \in \mathcal{I}_p$ follows the corresponding policy π_p . Persona creation details are presented in Section 3.2. Since our accounts share the same policy within each persona (and hence also the same behavior), we also refer to these accounts as *replicas*. All replicas interact with the platform following the same protocol, as described in Section 3.3.

Because each persona’s policy π_p is fixed by construction and encodes latent characteristics (e.g., political ideology) *a priori*, the design enables stratifying any outcome by user type without the endogeneity that plagues observational studies. In observational settings, any proxy for a user’s ideology derived from engagement history is itself shaped by algorithmic recommendations through the feedback loop between \mathcal{A} and $H_{i,t}$, making the stratification variable a joint product of the user and the algorithm. As our agents’ policies are specified before any platform interaction, this conditioning is well-defined by construction. In our \mathbb{X} case study, we exploit this to measure policy-conditional amplification of content in the *For You* feed compared to the *Following* feed.

Additionally, replicas differ only through a *profile condition* assignment $C_i \in \{0, 1, \dots, K\}$, where $C_i = 0$ denotes a baseline condition for user i and each $C_i = k$ (for $k \geq 1$) denotes a perturbation to a platform-visible signal. Within each persona p , we assign C_i by balanced randomization: replicas in \mathcal{I}_p are split into $K + 1$ equal-size groups, one per condition. In our deployment, we use

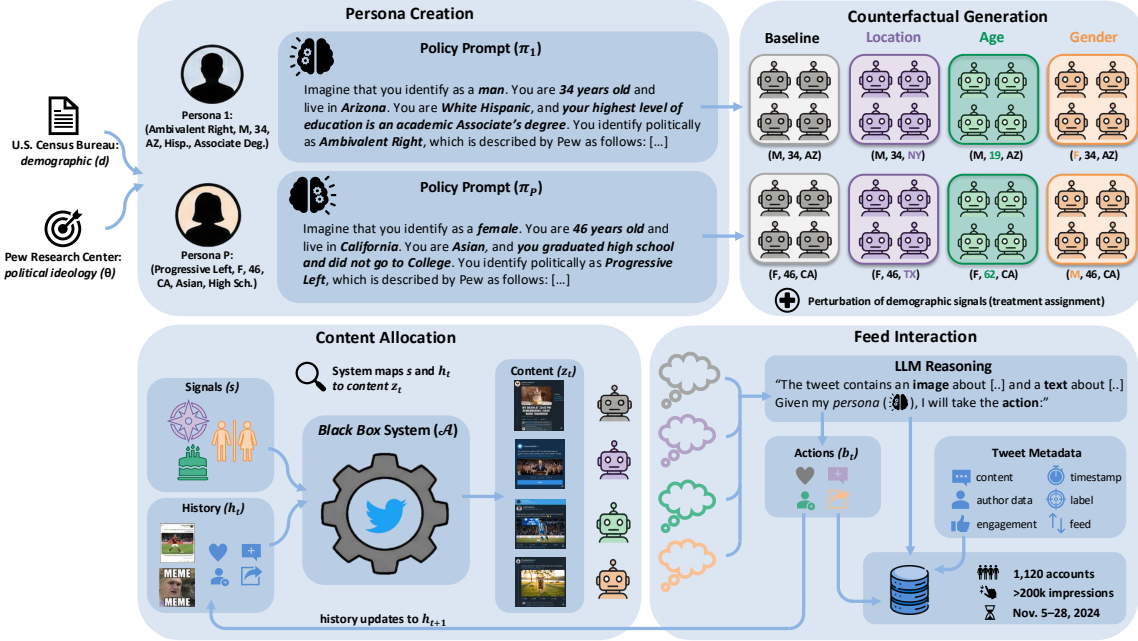


Figure 1: Overview of our experimental framework. **Top:** each persona is constructed by combining demographics from the U.S. Census with a political typology from Pew Research, producing an LLM prompt that fully specifies the agent’s behavior policy π_p . Each persona is replicated into multiple accounts that share the policy but differ in a single platform-visible signal (location, age, or gender) relative to the baseline. **Bottom:** The interaction loop for a single account. The platform’s black-box recommendation system \mathcal{A} observes the account’s signal s and interaction history h_t to allocate content z_t . The replicas reason about content via the same policy, and select an action b_t ; actions are appended to form the updated history h_{t+1} , creating a feedback loop between agents and platform recommendations. All interactions, reasoning traces and metadata are logged.

$P = 14$ personas and $|\mathcal{I}_p| = 80$ replicas per persona, with $K = 3$ counterfactual conditions plus the baseline, splitting replicas into four equal-size groups of 20 accounts each.

For each counterfactual condition k , our target estimand is the policy-conditional average treatment effect relative to baseline,

$$\tau_{p,k} = \mathbb{E}[Y \mid C = k, \pi_p] - \mathbb{E}[Y \mid C = 0, \pi_p], \quad (3)$$

where Y is a metric of interest, e.g., the content toxicity score or its political leaning; see Section 4.1 for more details. Within each persona, each account’s condition assignment C_i determines which component of the platform-visible signal S_i is perturbed relative to the baseline, while the behavior policy is held fixed at π_p by construction. The estimand $\tau_{p,k}$ therefore takes the form of a policy-conditional signal response $\Delta_{\pi_p}(s_1, s_0)$ as in Eq. (1), comparing the signal induced by condition k against the baseline. In Section 4.3, we present a regression model for estimating this quantity from the experimental outcomes.

3.2 Personas and Counterfactuals

We now describe how we construct personas and assign counterfactual conditions. The key design requirement is that all replicas within a persona share an identical behavior policy π_p , while differing

only in their platform-visible signal S_i , as determined by the condition assignment C_i .

Persona Construction. Each persona p is defined by a tuple of attributes (θ_p, d_p) , where $\theta_p \in \Theta$ denotes a political typology and $d_p = (\text{age}_p, \text{gender}_p, \text{location}_p, \text{race}_p, \text{education}_p)$ denotes demographic characteristics. We use data from the U.S. Census Bureau [75] (for demographic attributes) and Pew Research Center [58] (for political typology) to ground these attributes in real-world population distributions.

Demographic attributes are sampled to approximate the distribution of platform users. Age is drawn from four brackets (18–29, 30–49, 50–64, 65+) according to Census population shares, weighted by age-specific platform usage rates from Pew’s social media surveys. Gender is sampled according to Census proportions (male, female, or other). Race and ethnicity follow Census proportions, collapsed to seven categories (White Non-Hispanic, White Hispanic, Black, Asian, American Indian, Native Hawaiian/Pacific Islander, and Multiracial). Location is drawn from U.S. cities with population exceeding 50,000, with sampling probability proportional to city population. Finally, education is sampled conditional on age, gender, and race using Census educational attainment tables. Exact sampling weights are reported in Appendix A.2.

For political typology, we use Pew’s nine-category Political Typology [58], which partitions the U.S. electorate based on political values and party affiliation. These categories span the entire range of political ideologies and are rather fine-grained. Presenting them from right-most to left-most leaning, the categories are: Faith and Flag Conservatives, Committed Conservatives, Populist Right, Ambivalent Right, Stressed Sideliners, Outsider Left, Democratic Mainstays, Establishment Liberals, and Progressive Left. More details on how the typology works and what are the characteristics of each type can be found in Appendix A.3. We sample θ_p conditional on the persona’s age and gender, using Pew’s published demographic breakdowns of each typology.

Behavior Policy. Each persona p is associated with a behavior policy $\pi_p : \mathcal{Z} \rightarrow \mathcal{D}(\mathcal{B})$ mapping observed content to a distribution over actions. We implement π_p using GPT-4o with a fixed prompt that encodes the persona’s attributes in natural language. More detail on the prompt can be found in Appendix A.4. Concretely, when the agent encounters some content item $z \in \mathcal{Z}$, the item is sent to GPT-4o together with the persona prompt. The model then returns an action $b \in \mathcal{B}$ based on the encoded identity, with temperature set to zero. The full interaction protocol is presented in Section 3.3.

Each account is also assigned platform-required identifiers needed for deployment. Profile images depict non-human subjects (e.g., landscapes) and are generated using DALL-E [61], carrying no demographic signal. Account emails are created via MailSlurp, a disposable-inbox service, and usernames are decided by \mathbb{X} at registration. These identifiers are held fixed throughout the experiment and are not used to define the persona policy; they play no role in the counterfactual analysis. Instead, display names and birth dates do carry demographic information and serve as the signal for gender and age, respectively: their construction is described in Appendix A.5.

Treatment Conditions. In general, the choice of which demographic features to perturb is platform-dependent: the relevant signals are the ones that the platform can observe or infer. On \mathbb{X} , we consider $K = 3$ perturbations targeting location ($C_i = 1$), age ($C_i = 2$), and gender ($C_i = 3$). Each perturbation modifies exactly one demographic dimension while holding the behavior policy, political typology, and all remaining demographics at their baseline values; accounts with $C_i = 0$ serve as controls. Age is directly observed by the platform through the birth date declared at registration. Gender and location are not explicitly declared; instead, the platform can plausibly infer them from the account’s display name and IP address, respectively. While this cannot be verified under black-box access, demographic inference from names and IP addresses is standard practice in online advertising and content personalization. More details can be found in Appendix A.5.

Within each persona p , we partition the $|\mathcal{I}_p| = 80$ replicas into 4 equal-sized groups of 20 accounts by balanced randomization, one group per condition $C \in \{0, 1, 2, 3\}$. With $P = 14$ personas, our experiment comprises $N = 1,120$ accounts in total. Crucially, each perturbation modifies exactly one demographic dimension, while holding the persona’s political typology and all other demographics fixed at their baseline values. As the persona prompt is fixed by construction, observed differences in outcomes across conditions C_i capture how the platform’s content allocation responds to the perturbed attribute for a user following a fixed behavior policy (Eq. (3)).

3.3 Platform Interaction

We now describe the interaction protocol for our deployment on \mathbb{X} . While the experimental design above is mostly platform-agnostic, the choices below reflect the specific constraints of \mathbb{X} .

Feed Interaction. We deploy the accounts from November 5 to November 28, 2024, spanning Election Day and its immediate aftermath. Before the experiment begins, each account follows a set of accounts selected during an onboarding phase: the LLM is presented with \mathbb{X} ’s suggested-accounts page and chooses which accounts to follow based on the persona’s interests. Each account i then conducts eight daily browsing sessions over the deployment period, with session start times sampled randomly throughout the day: four on the algorithmically ranked *For You* feed and four on the reverse-chronological *Following* feed. The feed is fixed for the duration of each session, which lasts approximately ten minutes.

In each session, the agent navigates to the platform’s feed and observes content items z_t served to it. For each item, the agent extracts the tweet’s content and metadata from the rendered page and passes them to GPT-4o in a single API call together with the behavior policy prompt. The model returns an action $b_t \in \{\text{like, follow, read-replies, ignore}\}$. The action is executed via browser automation, and the tuple (z_t, b_t) is logged before advancing to the next item. A prompt-level instruction nudges the agent to take at least one active action every five items to ensure sufficient signal for the recommendation algorithm. Full feed engagement prompts and action definitions can be found in Appendix B.2.

Once deployed, agents risk suspension if detected by the platform’s systems. To mitigate detection, we introduce inter-action delays that simulate human browsing cadence, route each account through an IP address matching its assigned location, stagger account creation over time, and randomize session start times. Accounts may nonetheless face temporary locks requiring email verification, interrupting scheduled sessions and reducing observed exposures; we automate account recovery through a custom email-verification pipeline. Attrition rates are reported in Appendix B.4; we find no evidence of differential attrition across treatment conditions.

Data collection. We access the platform via Selenium, a framework to automate web browsers. This allows our sock-puppets to observe the rendered feed as it would appear to an ordinary user. For each content item z_t served to account i , we record the item’s text and author handle, its timestamp and engagement counts (likes, retweets, replies) at the time of observation, indicators for attached media, and the feed type (algorithmic or chronological). We also log the account’s action $b_t \in \mathcal{B}$ and the full LLM output including its reasoning trace. These logs define the realized trajectories used to construct our metrics of interest (Section 4.1).

We define a content item as *exposed* to an account if it is rendered during browsing and recorded by our logging pipeline. We do not impose a dwell-time threshold to define exposure: even when the agent chooses to ignore a piece of content, the LLM must process it to select the ignore action, so dwell time is mechanically driven by inference rather than serving as a measure of attention.

4 Measurement and Estimation

This section describes our measurement and estimation framework. We begin by defining the outcome metrics used to characterize content exposure (Section 4.1). We then introduce the amplification ratio to compare content exposure on the algorithmic feed and chronological feed (Section 4.2). Finally, we present our counterfactual estimation strategy, first pooled across personas to capture average effects, then at the persona level to capture heterogeneity (Section 4.3).

4.1 Metrics of Interest

We study how the platform’s content delivery varies across demographic profiles. To this end, we focus on a set of metrics that have been widely used in prior audits of social media platforms. Note that our metrics are not specific to X. That said, we instantiate our discussion to the X case study for clarity of presentation.

Classifying tweet content along these dimensions is a non-trivial task. For toxicity, reliable off-the-shelf classifiers exist and we adopt one directly. For political content and polarization, however, off-the-shelf models trained on general corpora performed poorly on our election-period data, likely due to the lack of domain-specific context. We therefore developed a custom LLM-based classification pipeline using chain-of-thought prompting [77], adopting best practices for LLM-based annotation: explicit definitions, structured output constraints, and grounded few-shot examples [10]. Recent work has demonstrated that LLMs can match or exceed human annotators for this class of tasks [34, 72]. We validated our pipeline against a hand-labeled set of 200 tweets; full implementation details and prompts are reported in Appendix C.

Toxic Content. Toxicity has been a central outcome metric in prior algorithmic audits and content moderation studies [14, 22, 70]. Multiple off-the-shelf classifiers are available, typically trained on large-scale annotated datasets of abusive or harmful language. We use the open-source `Detoxify` library [32], which outputs a continuous toxicity score for each tweet. We define the binary outcome *toxic* using a 0.5 threshold.

Political Content. A tweet is labeled as *political* if it clearly refers to: (i) the U.S. presidential election (e.g., campaigning, voting, results, legitimacy); (ii) U.S. political institutions, public policy, or governance; (iii) formal political actors such as candidates, elected officials, or political parties; or (iv) explicit partisan or electoral positioning. The classifier assigns a category from: {`not_political`, `political_left`, `political_neutral`, `political_right`}. From these we derive three binary outcomes: *politicalness* (any political category, including neutral), *right-leaning* (`political_right` only), and *left-leaning* (`political_left` only), each computed over all tweets. Our method for classifying political content, including the exact prompt specifications, can be found in Appendix C.

Polarizing Content. A tweet is labeled as *polarizing* if it satisfies two conditions: (1) it distinguishes between two or more social, political, religious, sexual, or racial groups (explicitly or implicitly), following the conceptual framework of Naseem et al. [51], and (2) it frames those groups in a moralized conflict—such as blame, hostility, moral superiority, threat, or victim-oppressor narratives. Mentioning politics, religion, or identity alone is not sufficient; the tweet must actively pit groups against each other [15, 66]. When classified as polarizing, the model assigns a type from {`political`, `religious`, `sexual`, `racial`, `other`}. The categorical breakdown aids labeling accuracy by anchoring the model’s reasoning; our analysis uses only the binary indicator *polarizing* versus *not polarizing*, pooling across all types. Our method for classifying polarizing content, including the exact prompt specifications, can be found in Appendix C.

Author Reach. To capture whether the algorithmic feed favors content from high-reach accounts, we measure the follower count of each tweet’s author as a proxy for author virality [79]. We classify a tweet as *high-reach* if its author’s follower count exceeds the sample median, yielding a binary indicator analogous to the other outcome variables.

4.2 “For You” vs “Following” Feed Analysis

We measure whether \mathbb{X} ’s algorithmic ranking in the *For You* feed amplifies exposure to specific content types relative to the reverse-chronological *Following* feed. This comparison addresses long-standing concerns about filter bubbles and echo chambers [63, 69]: if algorithmic curation systematically over-represents certain content—political material, toxic language, or ideologically skewed viewpoints—the aggregate effect on public discourse may be substantial. These concerns have prompted a growing body of research [30, 36, 49] and press [23, 28] to study whether and how algorithmic curation shapes content exposure.

Amplification Ratio. Huszár et al. [36] introduced the notion of *amplification ratio*: the relative likelihood that a set of tweets is seen under an algorithmically ranked feed compared to the reverse-chronological one. In their setting, the ratio is estimated via a randomized experiment where users are assigned exclusively to one of the two feeds. Our experimental design is different: each account browses both feeds during the deployment window, with sessions split evenly between the *For You* and *Following* feeds (Section 3). This yields *paired observations* at the user level, enabling within-user comparisons that eliminate confounds from account history or unobserved user characteristics that might correlate with feed assignment in between-subjects designs.

Fix a content property of interest and let $\mathcal{Z}_0 \subset \mathcal{Z}$ denote the subset of content items satisfying that property (e.g., tweets classified as political). For account i on day d , let FY_{id} and FOL_{id} denote the sets of content items observed in the *For You* and *Following* feeds, respectively, with cardinalities $m_{id}^{\text{FY}} = |\text{FY}_{id}|$ and $m_{id}^{\text{FOL}} = |\text{FOL}_{id}|$. Define the \mathcal{Z}_0 -content rate for account i on day d in each feed as:

$$r_{id}^{\text{FY}} = \frac{1}{m_{id}^{\text{FY}}} \sum_{z \in \text{FY}_{id}} \mathbf{1}[z \in \mathcal{Z}_0], \quad r_{id}^{\text{FOL}} = \frac{1}{m_{id}^{\text{FOL}}} \sum_{z \in \text{FOL}_{id}} \mathbf{1}[z \in \mathcal{Z}_0]. \quad (4)$$

We aggregate to the account level by averaging daily rates within each account, then compute the mean across accounts to obtain \bar{r}^{FY} and \bar{r}^{FOL} . The *amplification ratio* for content property \mathcal{Z}_0 is:

$$A(\mathcal{Z}_0) = \left(\frac{\bar{r}^{\text{FY}}}{\bar{r}^{\text{FOL}}} - 1 \right) \times 100\%, \quad (5)$$

This approach first averages content rates within each account and then across accounts, ensuring equal weight per account regardless of the number of active days.³ A value of 0% indicates equal rates across feeds, while positive (negative) values indicate that the algorithmic feed *amplifies* (*de-amplifies*) the content type relative to the reverse-chronological baseline. Note that the two feeds differ not only in ranking but also in content pool, as the *For You* feed surfaces content from outside the user’s follow graph. Hence, our ratio captures the full change in exposure from switching feeds rather than isolating the ranking effect. Since each account browses both feeds, cross-feed spillovers may arise (e.g., engagement on one feed updating the platform’s preference

³Because \bar{r}^{FOL} aggregates first over days and then over accounts, the denominator is well-defined whenever the content type has a nonzero base rate in the *Following* feed, a condition satisfied by all metrics we study.

model). If present, such spillovers would likely pull the two feeds closer together, making our results conservative. We discuss robustness in more detail in Appendix F.

We perform inference via persona-level cluster bootstrap: each draw resamples personas with replacement, preserving the within-persona dependence structure induced by shared behavioral policies.

Stratification by Ideology. A central question in the accountability literature is whether recommendation algorithms amplify *congenial* content, i.e., content that aligns with users’ political views. A growing body of evidence suggests they do, with a focus on political and polarizing content [4, 5, 29]. These findings motivate a stratified analysis: does amplification differ systematically for users with different ideological orientations?

Our experimental design is uniquely suited to answer this question. Traditional amplification studies face a fundamental identification challenge when stratifying by user ideology: the behaviors that reveal ideology to researchers (e.g., engagement patterns) are the same signals that platforms use to personalize content, creating circularity between the stratification variable and the outcome. Our design resolves this through an *a priori* ideology definition: each persona is assigned a political orientation before any platform interaction occurs, and the LLM behavioral policy remains unchanged across all sessions. This clean separation enables a direct test of whether amplification is heterogeneous across user types.

We stratify *left-leaning* and *right-leaning* following the Pew typology (see Appendix A.3). We report the estimated amplification ratios stratified by ideology in Section 5.1. As in the aggregate analysis, we perform inference via persona-level cluster bootstrap.

4.3 Counterfactual Analysis

Having characterized the platform’s relative amplification bias, we now turn to estimating the causal effect of user attributes on content exposure. We first estimate *pooled* treatment effects to capture the average algorithmic response to demographic changes across all personas. In Section 2, we defined the outcome as a summary of user’s full trajectory; in practice, we estimate treatment effects at the account-day level. We aggregate binary tweet labels into daily exposure rates and define the *algorithmic lift* as the additive difference between the exposure rate in the *For You* feed and the *Following* feed:

$$y_{id} = r_{id}^{\text{FY}} - r_{id}^{\text{FOL}}, \tag{6}$$

where r_{id}^{FY} and r_{id}^{FOL} are the content rates defined in Equation (4). The algorithmic lift y_{id} is the dependent variable in all counterfactual regressions that follow. Intuitively, y_{id} captures how much extra content of a type the algorithm injects beyond the chronological baseline [49].

We choose this additive specification over the ratio used in Section 4.2 for two reasons specific to the account-day unit of analysis. The first reason is robustness to sparsity: unlike the aggregate means, daily organic exposure to sparse content (e.g., toxic tweets) is frequently zero, rendering granular ratios undefined. The second reason is to focus on absolute differences: while the amplification ratio describes the system’s *relative bias*, the additive lift captures the magnitude of exposure. This allows us to quantify the actual “dose” of additional content of a certain type (e.g., toxic tweets) the algorithm injects into the user’s experience, given a counterfactual perturbation.

Pooled Treatment Effects. To estimate the average effect of demographic signals, we encode the treatment *direction* explicitly across personas:

- *Location*: $D_i^{\text{urb}} \in \{-1, 0, +1\}$, where +1 indicates the displayed state has a *higher* urbanization rate than the persona’s baseline state, -1 indicates a lower rate (more rural), and 0

indicates baseline. We use the percentage of population residing in urbanized areas from the 2020 U.S. Census [74].

- *Age*: $D_i^{\text{age}} \in \{-1, 0, +1\}$, where $+1$ indicates the displayed age bracket exceeds the persona’s baseline, -1 indicates it falls below, and 0 indicates baseline.
- *Gender*: $D_i^{\text{male}} = 1$ if account i displays a male profile, 0 if female.

Note that our ordinal encoding of age and location assumes a monotone relationship between these signals and algorithmic recommendations. Alternative specifications, such as dummies for specific categories, could capture non-monotone effects.

Let $\mathbf{T}_i = (D_i^{\text{urb}}, D_i^{\text{age}}, D_i^{\text{male}})^\top$ denote the treatment vector. We estimate the following model at the account-day level:

$$y_{id} = \gamma_d + \delta_{p_i} + \boldsymbol{\beta}^\top \mathbf{T}_i + X_{id}^\top \boldsymbol{\theta} + \varepsilon_{id}, \quad (7)$$

where γ_d are *day* fixed effects, δ_{p_i} are *persona* fixed effects, and X_{id} is a vector of time-varying controls. The coefficient vector $\boldsymbol{\beta}$ identifies the average treatment effect of each demographic attribute on algorithmic lift, pooled across all personas. The covariate vector X_{id} includes two sets of controls to improve precision: (1) *Hour-of-day indicators* (i.e., binary indicators for session activity in each hour), and (2) *Account age* (i.e., hours elapsed since account creation).

Because treatment is randomly assigned within each persona, $\boldsymbol{\beta}$ identifies the causal effect of each demographic attribute on algorithmic lift. We cluster standard errors at the account level to adjust for the correlation among repeated daily observations within the same account.

Heterogeneous Treatment Effects. The pooled analysis assumes that the platform’s response to a demographic signal (e.g., an increase in age) is constant across all user profiles. However, algorithmic personalization is likely context-dependent. In this section, we relax the homogeneity assumption to estimate *persona-specific* treatment effects.

We interact the treatment assignment with the persona indicator. Recall that $C_i \in \{0, 1, \dots, K\}$ denotes the profile condition assigned to account i . Let $W_{i,p,k} = \mathbf{1}\{p_i = p, C_i = k\}$ indicate that account i belongs to persona p and was assigned to counterfactual condition k . We estimate:

$$y_{id} = \gamma_d + \delta_{p_i} + \sum_{p=1}^P \sum_{k=1}^K \beta_{p,k} W_{i,p,k} + X_{id}^\top \boldsymbol{\theta} + \varepsilon_{id}. \quad (8)$$

Here, $\beta_{p,k}$ captures the algorithmic lift for persona p in condition k relative to that persona’s baseline. Under balanced randomization of C_i within each persona, these coefficients identify the persona-specific average treatment effects on algorithmic lift. The control vector X_{id} is identical to the one defined for the pooled analysis. We cluster standard errors at the account level.

5 Case Study: Results from Deployment on \mathbb{X}

We deployed 1,120 synthetic accounts on \mathbb{X} during the 2024 U.S. presidential election, as per the experimental design described in Section 3. This section reports and discusses the main results, following the measurement and estimation framework of Section 4. First, we analyze whether \mathbb{X} ’s algorithmic *For You* feed systematically amplifies specific content types relative to the chronological *Following* feed (Section 5.1). Second, we estimate how perturbations to platform-visible demographic signals affect algorithmic content exposure (Section 5.2).

Content type	Baseline (\bar{r}^{FOL})	Amplification	CI (low)	CI (high)
Political	38.0%	+14.2%	+9.1%	+19.8%
Right-leaning	14.9%	+23.3%	+15.8%	+31.9%
Left-leaning	6.5%	-6.1%	-14.6%	+2.1%
Toxic	3.9%	+39.2%	+14.4%	+69.7%
Polarizing	14.0%	+32.3%	+23.7%	+41.6%
High-reach	64.5%	+3.4%	-6.1%	+16.7%

Table 1: Amplification of content types in X’s algorithmic *For You* feed relative to the *Following* feed. The baseline column reports the average content rate in the *Following* feed (\bar{r}^{FOL}). Amplification values represent the relative percentage difference in content exposure. **The *For You* feed exposes users to significantly more toxic, polarizing, political, and right-leaning content.** Confidence intervals (95%) are constructed via persona-level cluster bootstrap.

5.1 “For You” vs “Following” Feed

We first assess the baseline behavior of the recommendation algorithm by computing the *amplification ratio* for different content categories. This within-subject comparison controls for user-level heterogeneity by comparing the algorithmic feed against the user’s own chronological baseline. Results are reported in Table 1.

We find that all content types except left-leaning and high-reach are significantly amplified in the For You feed. The largest effects are for toxic (+39.2%) and polarizing (+32.3%) content, consistent with evidence that algorithmic ranking increases exposure to anger and out-group hostility [49]. Political content is amplified by +14.2%; its high baseline prevalence makes this the largest shift in absolute terms. We also observe a statistically significant ideological asymmetry in the aggregate results: right-leaning content is amplified, whereas left-leaning content shows no significant amplification (+23.3% vs -6.1%, with the 95% CI including zero for the latter). As we demonstrate next, this aggregate pattern masks sharper divergences across user groups.

Stratification by Ideology. Table 2 and Figure 2 stratify these results by the political leaning of the account. The aggregate patterns mask substantial heterogeneity: amplification effects differ sharply—and for some content types, even reverse sign—across left-leaning and right-leaning users.

The most striking example is for toxic content: the algorithm amplifies it by +80.3% for left-leaning users but shows no significant effect for right-leaning users (-4.0%, CI includes zero)—a gap of 84.3 percentage points ($p < 0.01$). A similar pattern holds for political content, which is amplified roughly twice as much for left-leaning users as for right-leaning users (+18.1% vs. +8.7%), though the difference is only marginally significant. Polarizing content is also amplified more for left-leaning users (+36.9% vs. +25.8%), though the difference is not statistically significant. High-reach content shows the opposite pattern, with amplification concentrated among right-leaning users (+38.4% vs. -11.0%, $p < 0.01$).

Right-leaning content is amplified for *both* user groups (+22.3% for left-leaning users, +24.5% for right-leaning users), with no significant difference between them. Left-leaning content, by contrast, shows no amplification even for left-leaning users (-1.4%, not significant) and is actively suppressed for right-leaning users (-17.0%). The uniform amplification of right-leaning content across user types is consistent with a level effect in the algorithm’s content allocation, not with personalization toward ideologically congruent material. For right-leaning users, this effect is reinforced by the suppression of left-leaning content, creating an asymmetric echo chamber that does not arise for left-leaning users. This result reveals striking disparate exposure: to put it plainly, the algorithmic

Content type	Left-leaning users			Right-leaning users			Difference
	Ampl.	CI (low)	CI (high)	Ampl.	CI (low)	CI (high)	
Political	+18.1%	+13.1%	+24.5%	+8.7%	+2.1%	+17.8%	-9.4 pp*
Right-leaning	+22.3%	+11.5%	+34.2%	+24.5%	+16.2%	+37.2%	+2.2 pp
Left-leaning	-1.4%	-10.9%	+8.9%	-17.0%	-28.5%	-5.5%	-15.6 pp**
Toxic	+80.3%	+50.4%	+109.7%	-4.0%	-16.7%	+12.0%	-84.3 pp***
Polarizing	+36.9%	+26.6%	+46.6%	+25.8%	+14.7%	+43.1%	-11.0 pp
High-reach	-11.0%	-12.8%	-9.0%	+38.4%	+30.3%	+46.8%	+49.4 pp***

Table 2: Amplification ratios stratified by political leaning of users. **The algorithm’s effects differ substantially across user groups.** The difference column reports the gap (Right – Left) in percentage points (pp). Confidence intervals (95%) are constructed via persona-level cluster bootstrap; p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

feed amplifies right-leaning content for left-leaning users by more than 22%, while it suppresses left-leaning content by 17% for right-leaning users.

Across the hypothesis tests in Tables 1 and 2, all significant results survive Benjamini–Hochberg (BH) correction for multiple hypothesis testing [11], with the exception of the political between-group difference.




Figure 2: Amplification ratios by content type, stratified by user political leaning (left: blue, right: red). Bars show point estimates. **Toxic content is amplified by 80% for left-leaning users but not for right-leaning users; right-leaning content is amplified for both groups, while left-leaning content is suppressed for right-leaning users.** Error bars indicate 95% confidence intervals.

Outcome	Location (D^{urb})		Age (D^{age})		Gender (D^{male})	
	$\hat{\beta}$	95% CI	$\hat{\beta}$	95% CI	$\hat{\beta}$	95% CI
Politicalness	-1.70 (0.142)	[-3.86, 0.46]	0.70 (0.473)	[-1.26, 2.66]	0.90 (0.449)	[-1.45, 3.25]
Right-Leaning	-0.50 (0.477)	[-1.87, 0.87]	1.00 (0.102)	[-0.18, 2.18]	-0.40 (0.534)	[-1.58, 0.78]
Left-Leaning	-0.70 (0.131)	[-1.48, 0.08]	-0.30 (0.367)	[-1.08, 0.48]	0.30 (0.436)	[-0.48, 1.08]
Polarization	-1.60 (0.024)	[-2.97, -0.23]	0.50 (0.439)	[-0.68, 1.68]	-0.70 (0.302)	[-2.07, 0.67]
Toxicity	0.10 (0.770)	[-0.68, 0.88]	0.80 (0.031)	[0.02, 1.58]	-0.10 (0.876)	[-0.69, 0.49]
High-reach	1.30 (0.334)	[-1.25, 3.85]	-0.90 (0.524)	[-3.64, 1.84]	-0.10 (0.939)	[-2.84, 2.64]

Table 3: Pooled treatment effects on algorithmic lift. Each coefficient represents the average effect of a one-unit change in the demographic signal on the gap between *For You* and *Following* content rates. Coefficients for binary outcomes are scaled by 100 and expressed in percentage points. **Most pooled effects are small and not statistically significant, suggesting the algorithm’s response to demographic signals is not systematic across user profiles.** Values in parentheses are p -values computed from standard errors clustered at the account level. Significant coefficients ($p < 0.05$) shown in bold.

5.2 Counterfactual Analysis

We now turn to the counterfactual question: does the algorithm respond differently to users with different demographic profiles, *holding behavior fixed*? Using the algorithmic lift y_{id} defined in Eq. (6) as the dependent variable, we estimate the causal effect of perturbing age, gender, and location signals on the gap between algorithmic and chronological content exposure.

Pooled Treatment Effects. Table 3 reports pooled treatment effects from the model in Equation (7). The main finding is that demographic perturbations have limited systematic effects on algorithmic lift when pooled across user profiles: out of the estimated coefficients, only two are statistically significant at the 5% level. In particular, we find that displaying an older age increases the algorithmic lift of toxic content by 0.80 percentage points ($p = 0.031$), and displaying a more urban location reduces the lift of polarizing content by 1.60 percentage points ($p = 0.024$). Relative to the baseline algorithmic lift, these correspond to a 68% increase in the algorithm’s additional injection of toxic content (baseline lift: 1.15 pp) and a 38% reduction for polarizing content (baseline lift: 4.18 pp). Gender has no significant effect on any outcome. However, neither result survives BH correction for multiple testing across the 18 hypotheses.

The absence of strong pooled effects does not imply that the platform ignores demographic signals entirely. Rather, it is consistent with the hypothesis that any algorithmic responsiveness to profile attributes is *persona-dependent*, i.e., conditioned on the full behavioral profile the platform has accumulated. We investigate this possibility next.

Heterogeneous Treatment Effects. The near-zero pooled estimates could reflect genuinely null effects or heterogeneity that cancels in the aggregate. To distinguish these, we estimate persona-specific treatment effects via the model of Equation (8). Joint F -tests reject the null of zero

treatment effects for all outcomes at the 5% level: left-leaning ($F = 2.02, p < 0.001$), politicalness ($F = 1.93, p < 0.001$), author reach ($F = 1.93, p < 0.001$), toxic ($F = 1.89, p < 0.001$), right-leaning ($F = 1.57, p = 0.011$), and polarizing ($F = 1.49, p = 0.022$). All six remain significant after BH correction.

Figure 3 illustrates this pattern for left-leaning content, the outcome with the strongest heterogeneity. None of the three pooled coefficients for left-leaning content are significant (Table 3), yet persona-specific effects vary in sign and magnitude across personas, so that no consistent average effect emerges in the pooled estimate. Gender and location perturbations drive most of the heterogeneity, particularly for left-leaning and polarizing content; age effects are largely confined to toxicity and author reach. Full persona-level results are reported in Appendix E.

At the individual level, 26 of 252 persona–outcome–treatment coefficients are significant at the 5% level. Given the number of comparisons, these individual coefficients should not be interpreted in isolation. Instead, together with the joint rejection of homogeneity and the near-zero pooled estimates, they indicate that the algorithm’s response to demographic signals varies in direction across user types. This has a direct methodological implication: pooled audits that average treatment effects across heterogeneous user populations may fail to detect algorithmic discrimination present for specific subgroups.




Figure 3: Persona-specific treatment effects on algorithmic lift for left-leaning content, estimated via model (8). Each row is a persona, each panel a treatment. Points are estimated coefficients with 95% confidence intervals (filled markers indicate significance). **Left-leaning content exhibits the most persona-level heterogeneity: all pooled effects are null, yet 7 persona-specific effects are significant.**

5.3 Summary of Findings

Taken together, the amplification analysis of Section 5.1 and the counterfactual analysis of Section 5.2 point to two main conclusions. First, consuming content through X’s algorithmic *For You* feed leads to substantially higher exposure to toxic, polarizing, and right-leaning content than relying on the chronological *Following* feed. This amplification gap is uneven across user groups: left-leaning users are disproportionately exposed to toxic and political content, while right-leaning users experience a one-sided information environment in which politically congenial content is am-

plified and opposing content is suppressed. Second, perturbing observable demographic signals causally affects algorithmic content delivery, but in a heterogeneous, persona-dependent manner. Pooled treatment effects are largely null, yet joint tests reject the hypothesis of homogeneous effects for all outcomes. The same perturbation can shift content delivery in opposite directions depending on the user’s behavioral profile, producing effects that cancel in the aggregate.

A common pattern emerges across both analyses: population-level summaries mask sharp differences in how the algorithm treats distinct user groups. These findings highlight the value of moving beyond aggregate metrics in algorithmic auditing, toward designs that can detect heterogeneous effects across demographically and behaviorally distinct subpopulations.

6 Discussion

Summary. In this paper, we introduced a framework that uses LLM-powered agents for black-box auditing of personalization systems. Our agents are instantiated to follow certain “personas”, and reason about the content they are exposed to before engaging. Both observable and latent attributes can be encoded directly in the agent’s prompt, enabling stratification by design rather than through post-hoc inference from behavioral proxies. Crucially, because behavior is fully specified by the prompt, it is decoupled from user attributes: platform-visible signals can then be perturbed across conditions without altering how the agent behaves, enabling counterfactual auditing.

We deployed our auditing framework on \mathbb{X} during the 2024 U.S. presidential election, operating 1,120 agents and collecting over 200,000 impressions. Consistent with prior work and anecdotal evidence, we find that the algorithmic feed amplifies toxic, polarizing, political, and right-leaning content relative to the chronological baseline; left-leaning content is not significantly amplified. Stratifying by ideology reveals highly asymmetric amplification: right-leaning content is amplified regardless of user leaning, while left-leaning content is suppressed for right-leaning users. In our counterfactual analysis, pooled demographic effects are largely null, but persona-level analyses reject homogeneity of treatment effects for all outcomes, indicating that the algorithm’s response to demographic signals is persona-dependent.

Limitations. Our study is subject to some limitations. First, GPT-4o lacks lived experience and might carry biases. However, the framework does not require LLMs to faithfully approximate humans: it requires a *fixed* behavioral policy across replicas and behavior realistic enough to trigger personalization. For the counterfactual analysis, any additive LLM bias cancels under within-persona differencing; for the amplification analysis, the within-user comparison remains internally valid, but LLM engagement biases may shape the feedback loop and hence the levels of algorithmic amplification. Second, with 20 accounts per persona-condition cell and six outcome variables across three treatments, we test a large number of hypotheses relative to the available observations; we therefore present the heterogeneous analysis as exploratory, with evidence for heterogeneity resting on joint F -tests rather than individual coefficients. Finally, our findings are specific to \mathbb{X} during a period of exceptional political salience, and all accounts are newly created—the algorithm may weight profile signals differently for established users.

Future Work. Several extensions follow naturally from this work. Scaling to hundreds of personas would sharpen estimates of demographic effects; with sufficient power, this would also enable moderation analysis to identify which persona attributes (e.g., engagement patterns or political ideology) interact with the algorithm to shape its response to demographic signals. Enabling agents to create content (posting tweets, replying, quote-tweeting) would generate richer engagement signals and make the feedback loop between agent and algorithm more realistic, likely triggering deeper personalization. This richer interaction mode would require substantially more careful IRB

oversight, as agent-generated content directly affects real users.

Perhaps the most important extension is deploying the framework on other platforms. Nothing in our design is specific to \mathbb{X} ; the same approach applies to any platform with an algorithmic feed. The need for such audits is growing: policymakers increasingly demand algorithmic accountability, and platforms themselves have begun to respond. \mathbb{X} , for instance, recently published the source code of its recommendation pipeline [78]. Yet code alone does not make a system transparent: without model weights, production data, and surrounding infrastructure, the released code cannot be independently verified. Moreover, the patterns that audits seek to detect are often emergent consequences of optimization objectives, not behaviors readable from source code. Black-box auditing therefore remains necessary to characterize what users actually experience.

Finally, we stress that our framework need not be adversarial: platforms themselves could deploy synthetic agent populations in sandboxed environments to stress-test algorithm changes before live rollout, or regulators could run standardized agent cohorts to produce auditable, reproducible compliance evidence.

References

- [1] G. Aher, R. I. Arriaga, and A. T. Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *Proceedings of Machine Learning Research*, volume 202, 2023.
- [2] J. R. Anthis, R. Liu, S. M. Richardson, A. C. Kozlowski, B. Koch, E. Brynjolfsson, J. Evans, and M. S. Bernstein. Position: Llm social simulations are a promising research method. In *Proceedings of Machine Learning Research*, volume 267, 2025.
- [3] L. P. Argyle, E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, and D. Wingate. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31, 2023. ISSN 14764989. doi: 10.1017/pan.2023.2.
- [4] C. A. Bail, L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. F. Hunzaker, J. Lee, M. Mann, F. Merhout, and A. Volfovsky. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221, 2018.
- [5] E. Bakshy, S. Messing, and L. A. Adamic. Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132, 2015.
- [6] J. Bandy. Problematic machine behavior: A systematic literature review of algorithm audits. *Proceedings of the acm on human-computer interaction*, 5(CSCW1):1–34, 2021.
- [7] J. Bandy and N. Diakopoulos. Curating quality? how twitter’s timeline algorithm treats different types of news. *Social Media and Society*, 7, 2021. ISSN 20563051. doi: 10.1177/205630512111041648.
- [8] N. Bartley, A. Abeliuk, E. Ferrara, and K. Lerman. Auditing algorithmic bias on twitter. In *ACM International Conference Proceeding Series*, 2021. doi: 10.1145/3447535.3462491.
- [9] N. Bartley, K. Burghardt, and K. Lerman. Auditing exposure bias on social media for a healthier online discourse. In *Workshop Proceedings of the 18th International AAAI Conference on Web and Social Media (ICWSM Workshops)*, 2024. doi: 10.36190/2024.12. URL

- https://workshop-proceedings.icwsm.org/abstract.php?id=2024_12. CySoc 2024: 5th International Workshop on Cyber Social Threats.
- [10] J. Baumann, P. Röttger, A. Urman, A. Wendsjö, F. M. P. del Arco, J. B. Gruber, and D. Hovy. Large language model hacking: Quantifying the hidden risks of using llms for text annotation, 2025. URL <https://arxiv.org/abs/2509.08825>.
 - [11] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.
 - [12] J. Bisbee, J. D. Clinton, C. Dorff, B. Kenkel, and J. M. Larson. Synthetic replacements for human survey data? the perils of large language models. *Political Analysis*, 32(4):401–416, 2024.
 - [13] M. Boeker and A. Urman. An empirical investigation of personalization factors on tiktok. In *Proceedings of the ACM web conference 2022*, pages 2298–2309, 2022.
 - [14] P. Bouchaud, D. Chavalarias, and M. Panahi. Crowdsourced audit of twitter’s recommender systems. *Scientific Reports*, 13(1):16815, 2023.
 - [15] W. J. Brady, K. McLoughlin, T. N. Doan, and M. J. Crockett. How social learning amplifies moral outrage expression in online social networks. *Science Advances*, 7(33):eabe5641, 2021.
 - [16] J. Brand, A. Israeli, and D. Ngwe. Using gpt for market research. In *Proceedings of the 25th ACM Conference on Economics and Computation*, EC ’24, page 613, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400707049. doi: 10.1145/3670865.3673479. URL <https://doi.org/10.1145/3670865.3673479>.
 - [17] S. Casper, C. Ezell, C. Siegmann, N. Kolt, T. L. Curtis, B. Bucknall, A. Haupt, K. Wei, J. Scheurer, M. Hobbhahn, et al. Black-box access is insufficient for rigorous ai audits. *arXiv preprint arXiv:2401.14446*, 2024.
 - [18] S. H. Cen, A. Ilyas, J. Allen, H. Li, and A. Madry. Measuring strategization in recommendation: Users adapt their behavior to shape future content. *arXiv preprint arXiv:2405.05596*, 2024.
 - [19] A. Y. Chen, B. Nyhan, J. Reifler, R. E. Robertson, and C. Wilson. Subscriptions and external links help drive resentful users to alternative and extremist youtube channels. *Science Advances*, 9(35):eadd8080, 2023.
 - [20] W. Chen, D. Pacheco, K.-C. Yang, and F. Menczer. Neutral bots probe political bias on social media. *Nature communications*, 12(1):5580, 2021.
 - [21] S. J. Committee. Hearing on holding big tech accountable, 2023. URL <https://www.judiciary.senate.gov/meetings/holding-big-tech-accountable>. Accessed: 2025-03-06.
 - [22] T. Davidson, D. Warmesley, M. Macy, and I. Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 512–515, 2017.
 - [23] G. De Vynck and N. Nix. Elon musk’s twitter pushes hate speech, extremist content into ‘for you’, 2023. URL <https://www.washingtonpost.com/technology/2023/03/30/elon-musk-twitter-hate-speech/>. Accessed: 2026-02-07.

- [24] N. Elkin-Koren. Let the crawlers crawl: On virtual gatekeepers and the right to exclude indexing. *U. Dayton L. Rev.*, 26:179, 2000.
- [25] A. Filippas, J. J. Horton, and B. S. Manning. Large language models as simulated economic agents: What can we learn from homo silicus? In *Proceedings of the 25th ACM Conference on Economics and Computation, EC '24*, page 614–615, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400707049. doi: 10.1145/3670865.3673513. URL <https://doi.org/10.1145/3670865.3673513>.
- [26] Y. Gao, D. Lee, G. Burtch, and S. Fazelpour. Take caution in using llms as human surrogates. *Proceedings of the National Academy of Sciences*, 122(24):e2501660122, 2025.
- [27] S. Gerhart. Do web search engines suppress controversy? *First Monday*, 2004.
- [28] J. Gillum, A. Corse, and A. Tong. X algorithm feeds users political content—whether they want it or not, 2024. URL <https://www.wsj.com/tech/x-algorithm-feeds-users-political-content-whether-they-want-it-or-not-96fd55fc>. Accessed: 2026-02-07.
- [29] S. González-Bailón, D. Lazer, P. Barberá, M. Zhang, H. Allcott, T. Brown, A. Crespo-Tenorio, D. Freelon, M. Gentzkow, A. M. Guess, et al. Asymmetric ideological segregation in exposure to political news on facebook. *Science*, 381(6656):392–398, 2023.
- [30] A. M. Guess, N. Malhotra, J. Pan, P. Barberá, H. Allcott, T. Brown, A. Crespo-Tenorio, D. Dimmery, D. Freelon, M. Gentzkow, et al. How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science*, 381(6656):398–404, 2023.
- [31] G. Gui and O. Toubia. The challenge of using llms to simulate human behavior: A causal inference perspective. *arXiv preprint arXiv:2312.15524*, 2023.
- [32] L. Hanu and Unitary team. Detoxify. Github. <https://github.com/unitaryai/detoxify>, 2020.
- [33] M. Haroon, M. Wojcieszak, A. Chhabra, X. Liu, P. Mohapatra, and Z. Shafiq. Auditing youtube’s recommendation system for ideologically congenial, extreme, and problematic recommendations. *Proceedings of the National Academy of Sciences of the United States of America*, 120, 2023. ISSN 10916490. doi: 10.1073/pnas.2213020120.
- [34] M. Heseltine and B. Clemm von Hohenberg. Large language models as a substitute for human experts in annotating political text. *Research & Politics*, 11(1):20531680241236239, 2024.
- [35] L. Hewitt, A. Ashokkumar, I. Ghezze, and R. Willer. Predicting results of social science experiments using large language models. *Preprint*, 2024.
- [36] F. Huszár, S. I. Ktena, C. OBrien, L. Belli, A. Schlaikjer, and M. Hardt. Algorithmic amplification of politics on twitter. *Proceedings of the National Academy of Sciences of the United States of America*, 119, 2022. ISSN 10916490. doi: 10.1073/pnas.2025334119.
- [37] B. Imana, A. Korolova, and J. Heidemann. Having your privacy cake and eating it too: Platform-supported auditing of social media algorithms for public interest. *Proceedings of the ACM on Human-Computer Interaction*, 7, 2023. ISSN 25730142. doi: 10.1145/3579610.
- [38] B. Insider. Chatgpt isn’t slowing down — and google’s ai search overviews are struggling to catch up, 2025. URL <https://www.businessinsider.com/chatgpt-by-the-numbers-2025-11>. Accessed: 2025-03-06.

- [39] International Auditing and Assurance Standards Board. International standard on auditing 210: Agreeing the terms of audit engagements, 2009. URL <https://www.ifac.org/system/files/publications/files/A010%202012%20IAASB%20Handbook%20ISA%20210.pdf>. Paragraph 6(b). Accessed: 2026-03-10.
- [40] L. D. Introna and H. Nissenbaum. Shaping the web: Why the politics of search engines matters. *The information society*, 16(3):169–185, 2000.
- [41] Y. Leng. Can llms mimic human-like mental accounting and behavioral biases? In *Proceedings of the 25th ACM Conference on Economics and Computation*, EC '24, page 581, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400707049. doi: 10.1145/3670865.3673632. URL <https://doi.org/10.1145/3670865.3673632>.
- [42] Y. Leng and T. Nguyen. *Latent Neural Coupling of Risk and Time Preferences in LLMs Mirrors Human Biases*, page 542. Association for Computing Machinery, New York, NY, USA, 2025. ISBN 9798400719431. URL <https://doi.org/10.1145/3736252.3742588>.
- [43] A. Li, H. Chen, H. Namkoong, and T. Peng. Llm generated persona is a promise with a catch, 2025. URL <https://arxiv.org/abs/2503.16527>.
- [44] B. S. Manning, K. Zhu, and J. J. Horton. Automated social science: Language models as scientist and subjects. Technical report, National Bureau of Economic Research, 2024.
- [45] Q. Mei, Y. Xie, W. Yuan, and M. O. Jackson. A turing test of whether ai chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences of the United States of America*, 121, 2024. ISSN 10916490. doi: 10.1073/pnas.2313925121.
- [46] D. Metaxa, J. S. Park, J. A. Landay, and J. Hancock. Search media and elections: A longitudinal investigation of political search results. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–17, 2019.
- [47] D. Metaxa, J. S. Park, R. E. Robertson, K. Karahalios, C. Wilson, J. Hancock, C. Sandvig, et al. Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends® in Human-Computer Interaction*, 14(4):272–344, 2021.
- [48] C. Metz and C. K. MacCarthy. Openai chief, calling for regulation, warns of risks of artificial intelligence, 2023. URL <https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html>. Accessed: 2025-03-06.
- [49] S. Milli, M. Carroll, Y. Wang, S. Pandey, S. Zhao, and A. D. Dragan. Engagement, user satisfaction, and the amplification of divisive content on social media. *PNAS nexus*, 4(3):pgaf062, 2025.
- [50] E. Mustafaraj, E. Lurie, and C. Devine. The case for voter-centered audits of search engines during political elections. In *FAT* 2020 - Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020. doi: 10.1145/3351095.3372835.
- [51] U. Naseem, J. Ren, S. Anwar, S. Kohail, R. A. G. Veliz, R. Geislinger, A. Jabr, I. Abdulmumin, L. Qureshi, A. A. Borkar, M. I. Mukhtar, A. A. Ayele, I. S. Ahmad, A. Ali, M. Semmann, S. H. Muhammad, and S. M. Yimam. Polar: A benchmark for multilingual, multicultural, and multi-event online polarization, 2025. URL <https://arxiv.org/abs/2505.20624>.

- [52] B. Nyhan, J. Settle, E. Thorson, M. Wojcieszak, P. Barberá, A. Y. Chen, H. Allcott, T. Brown, A. Crespo-Tenorio, D. Dimmery, et al. Like-minded sources on facebook are prevalent but not polarizing. *Nature*, 620(7972):137–144, 2023.
- [53] H. C. on Energy and Commerce. Hearing on misinformation and social media, 2021. URL <https://energycommerce.house.gov/committee-activity/hearings/hearing-on-misinformation-and-social-media>. Accessed: 2025-03-06.
- [54] OpenAI, :, S. Agarwal, L. Ahmad, J. Ai, S. Altman, A. Applebaum, E. Arbus, R. K. Arora, Y. Bai, B. Baker, H. Bao, B. Barak, A. Bennett, T. Bertao, N. Brett, E. Brevdo, G. Brockman, S. Bubeck, C. Chang, K. Chen, M. Chen, E. Cheung, A. Clark, D. Cook, M. Dukhan, C. Dvorkak, K. Fives, V. Fomenko, T. Garipov, K. Georgiev, M. Glaese, T. Gogineni, A. Goucher, L. Gross, K. G. Guzman, J. Hallman, J. Hehir, J. Heidecke, A. Helyar, H. Hu, R. Huet, J. Huh, S. Jain, Z. Johnson, C. Koch, I. Kofman, D. Kundel, J. Kwon, V. Kyrylov, E. Y. Le, G. Leclerc, J. P. Lennon, S. Lessans, M. Lezcano-Casado, Y. Li, Z. Li, J. Lin, J. Liss, Lily, Liu, J. Liu, K. Lu, C. Lu, Z. Martinovic, L. McCallum, J. McGrath, S. McKinney, A. McLaughlin, S. Mei, S. Mostovoy, T. Mu, G. Myles, A. Neitz, A. Nichol, J. Pachocki, A. Paino, D. Palmie, A. Pantuliano, G. Parascandolo, J. Park, L. Pathak, C. Paz, L. Peran, D. Pimenov, M. Pokrass, E. Proehl, H. Qiu, G. Raila, F. Raso, H. Ren, K. Richardson, D. Robinson, B. Rotsted, H. Salman, S. Sanjeev, M. Schwarzer, D. Sculley, H. Sikchi, K. Simon, K. Singhal, Y. Song, D. Stuckey, Z. Sun, P. Tillet, S. Toizer, F. Tsimpourlas, N. Vyas, E. Wallace, X. Wang, M. Wang, O. Watkins, K. Weil, A. Wendling, K. Whinnery, C. Whitney, H. Wong, L. Yang, Y. Yang, M. Yasunaga, K. Ying, W. Zaremba, W. Zhan, C. Zhang, B. Zhang, E. Zhang, and S. Zhao. gpt-oss-120b and gpt-oss-20b model card, 2025. URL <https://arxiv.org/abs/2508.10925>.
- [55] F. Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, 2015. URL <https://www.hup.harvard.edu/books/9780674970847>.
- [56] F. Patel and J. Melendi. Meta’s oversight board needs access to Facebook’s algorithms to do its job, 2024. URL <https://www.brennancenter.org/our-work/analysis-opinion/metass-oversight-board-needs-access-facebooks-algorithms-do-its-job>. Accessed: 2026-03-10.
- [57] B. Perreault, J. H. Lee, R. Shava, and E. Mustafaraj. Algorithmic misjudgement in google search results: Evidence from auditing the us online electoral information environment. In *2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT 2024*, 2024. doi: 10.1145/3630106.3658916.
- [58] Pew Research Center. Beyond red vs. blue: The political typology. Technical report, November 2021. URL <https://www.pewresearch.org/politics/2021/11/09/beyond-red-vs-blue-the-political-typology-2/>.
- [59] T. Piccardi, M. Saveski, C. Jia, J. Hancock, J. L. Tsai, and M. S. Bernstein. Reranking partisan animosity in algorithmic social media feeds alters affective polarization. *Science*, 390(6776):eadu5584, 2025.
- [60] I. D. Raji, A. Smart, R. N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, and P. Barnes. Closing the ai accountability gap: Defining an end-to-end frame-

- work for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 33–44, 2020.
- [61] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021.
- [62] M. H. Ribeiro, R. Ottoni, R. West, V. A. Almeida, and W. Meira Jr. Auditing radicalization pathways on youtube. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 131–141, 2020.
- [63] F. Rowland. The filter bubble: What the internet is hiding from you. *portal: Libraries and the Academy*, 11(4):1009–1011, 2011.
- [64] C. Sandvig, K. Hamilton, K. Karahalios, and C. Langbort. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 22(2014):4349–4357, 2014.
- [65] Senate Judiciary Subcommittee on Privacy, Technology, and the Law. Hearing on oversight of ai: Rules for artificial intelligence, 2023. URL <https://www.judiciary.senate.gov/meetings/oversight-of-ai-rules-for-artificial-intelligence>. Accessed: 2025-03-06.
- [66] A. Simchon, W. J. Brady, and J. J. Van Bavel. Troll and divide: the language of online polarization. *PNAS nexus*, 1(1):pgac019, 2022.
- [67] I. Srba, R. Moro, M. Tomlein, B. Pecher, J. Simko, E. Stefancova, M. Kompan, A. Hrkova, J. Podrouzek, A. Gavornik, et al. Auditing youtube’s recommendation algorithm for misinformation filter bubbles. *ACM Transactions on Recommender Systems*, 1(1):1–33, 2023.
- [68] Statista. Most popular social networks worldwide as of january 2025, ranked by number of monthly active users, 2025. URL <https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/>. Accessed: 2025-03-06.
- [69] C. R. Sunstein. Republic: Divided democracy in the age of social media. 2018.
- [70] Z. Talat and D. Hovy. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93, 2016.
- [71] P. Tanyi and J. Cathey. The audit of banks in the usa: Has it changed since the financial crisis? *Journal of Banking Regulation*, pages 1–26, 2024.
- [72] P. Törnberg. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*, 2023.
- [73] M. Tranchero, C.-F. Brenninkmeijer, A. Murugan, and A. Nagaraj. Theorizing with large language models. Working Paper 33033, National Bureau of Economic Research, October 2024. URL <http://www.nber.org/papers/w33033>.
- [74] U.S. Census Bureau. 2020 census urban and rural classification. <https://www.census.gov/programs-surveys/geography/guidance/geo-areas/urban-rural/2020-ua-facts.html>, 2020. Accessed: 2025.

- [75] U.S. Census Bureau. American community survey and current population survey demographic tables, 2021. URL <https://www.census.gov/data.html>. Accessed: 2025.
- [76] S. Wang, S. Huang, A. Zhou, and D. Metaxa. Lower quantity, higher quality: Auditing news content and user perceptions on twitter/x algorithmic versus chronological timelines. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW2):507:1–507:25, 2024. doi: 10.1145/3687046.
- [77] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [78] xAI. xai-org/x-algorithm: Algorithm powering the For You feed on X. <https://github.com/xai-org/x-algorithm>, 2026. GitHub repository.
- [79] J. Ye, L. Luceri, and E. Ferrara. Auditing political exposure bias: Algorithmic amplification on twitter/x during the 2024 u.s. presidential election. In *ACMF AccT 2025 - Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 2025. doi: 10.1145/3715275.3732159.
- [80] H. Zade, M. Wack, Y. Zhang, K. Starbird, R. Calo, J. Young, and J. D. West. Auditing google’s search headlines as a potential gateway to misleading content: Evidence from the 2020 us election. *Journal of Online Trust and Safety*, 1(4), 2022.

Supplementary Material

A Details on Persona Creation

A.1 Full Profiles

Table 4 details the three counterfactual perturbations applied to each persona (location, age, and gender) relative to the base profile. Table 5 reports the attributes held fixed across all conditions within each persona (race/ethnicity, education, and political leaning).

Persona	Location (C=1)	Age (C=2)	Gender (C=3)
1	Arizona → Texas	30–49 → 65+	M → F
2	Florida → Colorado	30–49 → 18–29	M → F
3	Florida → North Carolina	18–29 → 30–49	F → M
4	Louisiana → New York	30–49 → 50–64	F → M
5	Massachusetts → California	18–29 → 50–64	F → M
6	Alabama → Illinois	50–64 → 18–29	M → F
7	Michigan → Washington	18–29 → 30–49	M → F
8	Nevada → Texas	30–49 → 65+	M → F
9	Tennessee → Arizona	18–29 → 65+	F → M
10	Indiana → California	30–49 → 18–29	M → F
11	California → Florida	50–64 → 30–49	M → F
12	California → California	18–29 → 30–49	M → F
13	New York → Ohio	30–49 → 50–64	F → M
14	California → Illinois	30–49 → 18–29	F → M

Table 4: Counterfactual perturbations by persona. Each counterfactual changes exactly one platform-visible attribute relative to the base profile. Locations are reported at the state level (cities omitted), ages are sampled uniformly from the buckets.

A.2 Demographic Sampling Distributions

Each persona’s demographic attributes are sampled independently from real-world population distributions. We describe the source and sampling weights for each attribute below. All the external information is either derived from the U.S. Census [75] or Pew Research Center [58].

Gender. Gender is drawn from U.S. Census population: Female (50.5%), Male (47.2%), Other (2.3%).

Age. Age brackets are sampled by combining Census population shares with platform-specific usage rates. Let $P(\text{age})$ denote the Census share and $P(\text{Twitter} \mid \text{age})$ the platform usage rate. We compute $P(\text{age} \mid \text{Twitter}) \propto P(\text{Twitter} \mid \text{age}) \cdot P(\text{age})$, yielding the following sampling weights:

	18–29	30–49	50–64	65+
Census share $P(\text{age})$	13.3%	25.8%	19.0%	17.1%
Usage rate $P(\text{Twitter} \mid \text{age})$	84%	81%	73%	45%
Sampling weight $P(\text{age} \mid \text{Twitter})$	20.8%	38.9%	25.8%	14.4%

Persona	Race/Ethnicity	Education	Political Leaning
1	White (Hispanic)	Associate’s	Ambivalent Right
2	White (Non-Hisp.)	Some college	Ambivalent Right
3	American Indian	Some college	Establishment Liberals
4	White (Non-Hisp.)	High school	Democratic Mainstays
5	White (Non-Hisp.)	High school	Outsider Left
6	White (Non-Hisp.)	Some college	Progressive Left
7	White (Non-Hisp.)	11th grade	Outsider Left
8	White (Hispanic)	High school	Faith and Flag Conservatives
9	Black	Bachelor’s	Outsider Left
10	White (Non-Hisp.)	Some college	Populist Right
11	Black	Bachelor’s	Outsider Left
12	White (Non-Hisp.)	10th grade	Ambivalent Right
13	White (Non-Hisp.)	Bachelor’s	Committed Conservatives
14	Asian	High school	Progressive Left

Table 5: Unperturbed persona attributes. These attributes define the behavioral policy, together with the base values of the platform-visible attributes.

Within each bracket, a specific age is drawn uniformly at random.

Race and ethnicity. Race is drawn from Census proportions, collapsed to seven categories: White Non-Hispanic (58.9%), White Hispanic (16.6%), Black (13.6%), Asian (6.3%), American Indian (1.3%), Native Hawaiian/Pacific Islander (0.3%), and Multiracial (3.0%).

Location. Location is drawn from U.S. incorporated places with population exceeding 50,000, with sampling probability proportional to city population (2022 Census estimates, 333 cities).

Education. Education is sampled conditional on the persona’s age, gender, and race using Census educational attainment tables, covering 15 levels: none, 1st–4th grade, 5th–6th grade, 7th–8th grade, 9th grade, 10th grade, 11th grade, high school graduate, some college (no degree), occupational Associate’s, academic Associate’s, Bachelor’s, Master’s, professional degree, and doctoral degree.

Political typology. Political typology is sampled from $P(\theta \mid \text{age, gender})$, derived from the demographic breakdowns published in Pew’s Political Typology report [58]. We provide more details on this in the next section.

A.3 Political Typology Descriptions

We adopt the nine-category political typology developed by Pew Research Center [58], which classifies Americans based on their political values and attitudes rather than party affiliation alone. The following descriptions are drawn directly from Pew’s published characterizations.

Faith and Flag Conservative. Faith and Flag Conservatives are highly religious, politically engaged and both socially and economically conservative. They favor a robust role for religion in public life and a smaller role for government in society, and they hold that a strong American military is essential in international affairs. They overwhelmingly identify with the GOP and remain strong supporters of former President Donald Trump. More than four-in-ten are White evangelical Protestants, the highest share of any political typology group. Faith and Flag Conservatives are

more likely than those in other groups to emphasize the importance of religion in their lives and to hold restrictive attitudes on abortion and same-sex marriage. Three-quarters say that the best way to ensure peace is through military strength rather than through diplomatic means. Faith and Flag Conservatives are the oldest typology group, with a median age of 57.

Committed Conservative. Staunchly conservative and overwhelmingly Republican, Committed Conservatives hold pro-business views traditionally associated with the Republican Party, have favorable attitudes about international trade and favor a limited role of government. Their approach to international relations centers on engaging with U.S. allies and maintaining American military might. Committed Conservatives tend to hold more moderate positions on immigration than the two other deeply conservative groups. While nearly all voted for Donald Trump for president in 2020, and most hold positive views of him today, Committed Conservatives are less likely than Populist Right and Faith and Flag Conservatives to favor a major role for Trump in their party's future. Committed Conservatives are among the most educated of the GOP-oriented groups and are among the highest-income political typology groups.

Populist Right. Very conservative and overwhelmingly Republican, Populist Right hold highly restrictive views about immigration policy and are very critical of government. But, in contrast to other parts of the GOP coalition, their criticism extends well beyond government to views of big business and to the economic system as a whole: 82% say that large corporations are having a negative impact on the way things are going in the country, and nearly half support higher taxes on the wealthy and on large corporations. A majority of Populist Right are women (54%). Populist Right are also one of the least highly educated groups; just two-in-ten are college graduates. Nearly nine-in-ten say that the economic system in this country unfairly favors powerful interests, far higher than the share in any other Republican-oriented group.

Ambivalent Right. On issues ranging from the size of the federal government to views about business, gender and race, Ambivalent Right hold many views that are largely consistent with core conservative values. Yet they also hold more moderate stances on several social issues and differ from some other segments of the GOP coalition in taking a more internationalist view of foreign policy and a less restrictive position on immigration. With 63% under the age of 50, they are substantially younger than other Republican-oriented groups. About two-thirds (65%) are White, making this group more racially and ethnically diverse than other GOP coalition groups. Just over half (54%) say abortion should be legal in all or most cases, a view held by far smaller shares of those in other GOP groups. They are also distinct from other Republican-oriented groups in their views of Donald Trump—more likely to say they feel coldly toward the former president (46%) than warmly (34%).

Stressed Sideliner. Stressed Sideliners are generally disconnected from politics and the two major parties, voting at lower rates than most other typology groups. Although Stressed Sideliners make up 15% of American adults, they were just 10% of 2020 voters due to their relatively low turnout rate. They are split evenly between those who identify with or lean toward the Republican Party (45%) and those who are Democrats or Democratic leaners (45%). They tend to lean liberal on economic issues and tilt conservative on some social issues. About four-in-ten (43%) live in lower-income households, higher than most other political typology groups. They are the group most likely to describe their personal financial situation as only fair or poor (63%). Stressed Sideliners are also one of the least highly educated groups.

Outsider Left. Outsider Left are the youngest typology group, making up 10% of the public. Holding liberal views on most issues and overwhelmingly voting Democratic, they aren't particularly enamored with the Democratic Party—though they have deeply negative views of the GOP. Nearly

half (48%) describe their own political views as liberal. Outsider Left are somewhat more skeptical of government’s role than other Democratic-oriented groups. Four-in-ten are under the age of 30 and 83% are under 50. They are racially and ethnically diverse: about half (49%) are White, 20% are Hispanic, 15% are Black and 10% are Asian. Just 53% say the Democratic Party represents them at least somewhat well, and an overwhelming majority (86%) say that they usually feel like none of the candidates for public office represent their views well.

Democratic Mainstay. Democratic Mainstays are one of the largest groups in the political typology and the largest single group as a share of the Democratic coalition. They generally favor policies that expand the social safety net and support higher taxes on corporations. But they are somewhat more hawkish than other Democratic-oriented groups on foreign policy and less liberal on immigration policy and some social issues. Nearly half (49%) consider themselves “strong Democrats.” Democratic Mainstays are slightly older and have less formal education than other Democratic-oriented groups. They are the group with the largest share of Black non-Hispanic adults (26%), and six-in-ten are women. They are the only Democratic-oriented typology group in which a larger share say that the decline in the share of Americans belonging to an organized religion is bad for society than say this is good for society.

Establishment Liberal. Holding liberal positions on nearly all issues, Establishment Liberals are some of the strongest supporters of the current president and the Democratic Party of any political typology group. While deeply liberal, Establishment Liberals are the typology group most likely to see value in political compromise and tend to be more inclined toward more measured approaches to societal change than their Progressive Left counterparts. Like other Democratic-oriented groups, most (73%) say a lot more needs to be done to ensure racial equality. Yet they are the only Democratic-aligned group in which a majority of those who say a lot more needs to be done also say this can be achieved by working within the current system. About half (51%) are White, while 18% are Black, 20% are Hispanic and 10% are Asian. Nearly nine-in-ten (89%) say that compromise is how things get done in politics, higher than the share in any other political typology group.

Progressive Left. Reflecting their name, Progressive Left have very liberal views across a range of issues—including the size and scope of government, foreign policy, immigration and race. A sizable majority (79%) describe their views as liberal, including 42% who say their views are very liberal. Roughly two-thirds (68%) are White, non-Hispanic, by far the largest share among Democratic-aligned groups. Progressive Left are the second youngest typology group—71% are ages 18 to 49. They are also highly educated, with about half (48%) holding at least a four-year college degree. Their views on race and racial equality distinguish them from other typology groups: sizable majorities say White people benefit from societal advantages that Black people do not have and that most U.S. institutions need to be completely rebuilt to ensure equal rights for all Americans regardless of race or ethnicity.

When we stratify personas by political leaning, we group *right-leaning* personas as those assigned Faith and Flag Conservative, Committed Conservative, Populist Right, or Ambivalent Right, and *left-leaning* personas as those assigned Outsider Left, Democratic Mainstay, Establishment Liberal, or Progressive Left. Stressed Sideliners constitute a separate, politically disengaged category; however, no persona in our final sample is assigned this typology (see Table 5).

A.4 Behavior Policy

Each sock-puppet’s behavior is governed by a *persona prompt* that conditions the LLM on a specific demographic and political profile:

Persona Prompt

Imagine that you {gender}. You are {age} years old and live {preposition_location} {location}. You are {race}, and {education}. You identify politically as {political_type}, which is described by Pew as follows: {poltype_description}.

Each placeholder in the persona prompt is populated as follows:

- {gender}: One of “identify as a man,” “identify as a woman,” or “do not identify as a man or a woman.”
- {age}: A randomly generated integer within the sampled age bracket (18–29, 30–49, 50–64, or 65–85).
- {preposition_location}: Either “in” or “near,” selected uniformly at random.
- {location}: A U.S. city sampled from incorporated places with population of 50,000 or more, weighted by population.
- {race}: One of “White and Non-Hispanic,” “White and Hispanic,” “Black,” “Asian,” “American Indian,” “Native Hawaiian and Other Pacific Islander,” or “Multiracial.”
- {education}: Highest attainment level (see Table 6).
- {political_type}: One of the nine categories from Pew’s Political Typology: “Faith and Flag Conservative,” “Committed Conservative,” “Populist Right,” “Ambivalent Right,” “Stressed Sideline,” “Outsider Left,” “Democratic Mainstay,” “Establishment Liberal,” or “Progressive Left.”
- {poltype_description}: A detailed description of the political typology category drawn from Pew Research Center’s published characterizations [58]. Full descriptions are provided in Section A.3.

Prompt phrase
“you received no education”
“your highest level of education is [4th/6th/8th/9th/10th/11th] grade”
“you graduated high school and did not go to college”
“you received some college education and did not finish your degree”
“your highest level of education is an [occupational/academic] Associate’s degree”
“your highest level of education is a [Bachelor’s/Master’s/professional/Doctoral] degree”

Table 6: Education level prompts.

A.5 Platform-Visible Signals

Each counterfactual condition perturbs exactly one demographic dimension while holding the behavioral policy fixed. We target location, age, and gender because \mathbb{X} plausibly observes or infers each: location via IP geolocation, age via the declared birth date at registration, and gender via

the display name. On other platforms, different signals may be observable and thus amenable to perturbation. Below we describe how each perturbed signal reaches the platform in our deployment.

Location ($C = 1$). The platform does not ask users to declare their location; instead, it infers geographic information from the account’s IP address. We route each account’s traffic through a residential proxy service (SmartProxy) that provides state-level sticky sessions. During preprocessing, we query the provider’s API for state-specific port ranges and assign each account a unique port within the range corresponding to its assigned state. For baseline and non-location conditions ($C \in \{0, 2, 3\}$), the assigned state is the persona’s baseline location; for location-perturbed accounts ($C = 1$), it is the counterfactual state. At runtime, each account’s browser connects through `state.smartproxy.com` on its assigned port, so the platform observes a residential IP in the target state for every session.

Age ($C = 2$). The platform requires a birth date during account registration. For each account, we generate a random birth date within the account’s assigned age bracket: for baseline and non-age conditions ($C \in \{0, 1, 3\}$) this is the persona’s baseline bracket, while for age-perturbed accounts ($C = 2$) it is the counterfactual bracket. The birth date is entered into the platform’s registration form, providing the primary channel through which the platform observes the account’s age.

Gender ($C = 3$). Gender is communicated to the platform through the account’s display name. We generate names conditional on demographic attributes using GPT-4o via the following prompt:

Name Generation Prompt

{`demographic_prompt`} Please consider 200 unique full names that would likely fit this person then randomly choose 100. Please return the 100 that you choose as an array of strings. Please do not write anything else in your response, just the array of strings of names. The output should follow the form “[name1, name2, ...]” and contain nothing else.

The prompt {`demographic_prompt`} follows the template defined in Section A.4 but with the gender field set to match the account’s *assigned condition* rather than the baseline. We generate a pool of 100 candidate names per persona—larger than the 80 accounts per persona—and then randomly sample without replacement when assigning names to individual accounts.

B Details on Deployment and Interaction

B.1 Account Distribution

We initially planned to deploy 1,120 sock-puppets – 80 replicas for 14 different personas. However, due to verification failures at account creation time, 826 accounts were ultimately deployed successfully, yielding 9,733 user-day observations.


Table 7 summarizes the distribution of observation time across accounts. The majority of accounts contributed between 10 and 16 days of data, with 9.7% achieving near-full coverage and 2.3% experiencing early dropout. We also report the distribution of user and user-day observations stratified by counterfactual category in Figures 4a and 4b.

B.2 Feed Engagement


Each sock-puppet interacts with the platform’s feed through an LLM-mediated engagement loop. The agent observes tweets served by the recommendation algorithm and decides how to respond based on its persona. We define four possible actions an agent can take in response to each tweet:

Statistic	Value
Accounts deployed	826
User-day observations	9,733
<i>Days active per account</i>	
Mean	13.2 days
Median	12 days
IQR	[10, 16] days
Maximum	23 days
Near-full coverage (≥ 21 days)	9.7%
Early dropout (≤ 3 days)	2.3%

Table 7: Sample Summary



(a) Number of active users per cell.



(b) Number of user-day observations per cell.

Figure 4: Sample sizes by persona \times treatment condition.

- **Like:** Signal positive engagement with the content.
- **Follow:** Subscribe to the tweet’s author for future content.
- **Read Replies:** View the conversation thread, signaling interest without explicit endorsement (selecting this action navigates into the reply thread, where up to five replies are recursively processed through the same pipeline before returning to the main feed).
- **Ignore:** Skip the tweet without interaction.

These actions mirror the primary engagement mechanisms available to human users on the platform. To ensure sufficient engagement signal for the recommendation algorithm to learn user preferences, we require agents to take at least one active action (like, follow, or read replies) every five tweets. This constraint prevents purely passive scrolling behavior that would provide minimal signal to the algorithm. Importantly, the constraint is specified only as a prompt-level instruction: each content item is processed in a single, stateless LLM call with no access to prior actions. The instruction therefore acts as a behavioral nudge rather than a hard constraint.

The engagement behavior is controlled by two prompts. The *system prompt* establishes the agent’s identity and explains the available actions:

Feed Engagement Prompt (System Prompt)

Welcome to Twitter! In future messages, you will see a sequence of tweets and have to decide how to react to each one.

{persona_prompt}

In our conversation below, I will show you a sequence of Tweets on Twitter. Please use your persona as described above to decide how you would like to react to each one.

You are primarily on Twitter to read tweets, but you can also follow accounts, like tweets, and read replies to tweets.

Below, I list the possible actions you can take.

- If you want to see more tweets like the one you just saw, you can follow the author (to do this, include "[FOLLOW]" in your response).
- If you are curious to see what other people think about the tweet, you can read the replies (to do this, include "[READ-REPLY]" in your response).
- If you like the tweet, you can like it (to do this, include "[LIKE]" in your response),
- If you do not want to react to a tweet (for example, if you are not interested in it), you can ignore it (to do this, include "[IGNORE]" in your response).

You should use one of the first three actions ([FOLLOW], [READ-REPLY], [LIKE]) at least once every 5 tweets. You can use [IGNORE] as many times as you like.

Your response should contain either [IGNORE], [FOLLOW], [READ-REPLY], or [LIKE]. If it contains [IGNORE], please do not include any other reactions.

For each tweet encountered, the agent receives a *user prompt* containing the tweet’s content and metadata:

Tweet Prompt (User Prompt)

You are on twitter and see the following tweet - please first repeat the tweet to ensure that you have understood it, then explain your reaction.

Author: {author}

Tweet Text: {tweetText}

Likes: {likes}

Retweets: {retweets}

Replies: {replies}

(Plus optionally the quoted tweet text and media image if present)

We include engagement metrics (likes, retweets, replies) to provide social proof signals that human users would observe. The instruction to “repeat the tweet” serves as a chain-of-thought mechanism, ensuring the agent processes the content before deciding on an action. The agent’s chosen action is then executed on the platform via automated browser interactions.

B.3 Action Distribution

Table 8 reports the distribution of actions across all tweet interactions. Agents ignore or read replies to the vast majority of tweets. Likes are relatively rare, and follows are near-absent.

Table 8: Distribution of agent actions.

Action	Count	Share (%)
Read-Reply	139,000	57.1
Ignore	90,151	37.1
Like	14,095	5.8
Follow	13	0.01
Total	243,259	

B.4 Attrition Analysis

Our identification strategy for the counterfactual analysis relies on random assignment of treatment conditions within each persona. If accounts in certain treatment conditions are more likely to be banned or suspended during the experiment, the resulting imbalance could bias our estimates. Therefore, we test whether, conditional on successful deployment, observation time varies systematically by treatment.

Note that we focus on mid-experiment attrition (bans and suspensions) rather than pre-deployment failures at the account creation stage. Unlike in human experiments, pre-deployment attrition cannot introduce selection bias in our setting: accounts within the same persona-treatment cell are exchangeable replicas running an identical behavioral policy, so creation failures only reduce statistical power, not the representativeness of surviving accounts.

Table 9 reports observation time by treatment condition among the 826 deployed accounts. To test for differential attrition, we regress days active on treatment indicators with fixed effects:

$$\text{DaysActive}_i = \alpha + \beta_{\text{loc}} \cdot \mathbf{1}[\text{Location}] + \beta_{\text{age}} \cdot \mathbf{1}[\text{Age}] + \beta_{\text{gen}} \cdot \mathbf{1}[\text{Gender}] + \gamma_{p_i} + \varepsilon_i \quad (9)$$

No treatment coefficient is statistically significant (Location: $p = 0.805$; Age: $p = 0.924$; Gender: $p = 0.301$), indicating that observation time is balanced across conditions. Overall, we find no evidence that treatment assignment affects the likelihood of mid-experiment attrition.

Table 9: Observation Time by Treatment Condition

Treatment	N	Mean	Median	IQR
Base	219	13.20	12	[10, 16]
Location	191	13.39	12	[10, 17]
Age	209	13.30	13	[10, 16]
Gender	207	12.80	12	[10, 16]

C Tweet Labeling

We classify tweet content along two dimensions: political content and polarizing content. Both classifiers use gpt-oss-20b [54], with chain-of-thought prompting and few-shot learning ($k = 4$ examples selected via bootstrap optimization).

C.1 Political Content

We classify each tweet into one of four mutually exclusive categories: `not_political`, `political_left`, `political_right`, or `political_neutral`. The classifier is prompted as follows:

Political Classification Prompt

The tweet was posted in November 2024, during the U.S. presidential election period.

Classify (1) whether the tweet contains substantive U.S. political content, and (2) if political, its ideological alignment in U.S. politics.

Definition 1 — Political vs Not Political

Label a tweet as “political” only if it clearly refers to at least one of the following:

- The U.S. presidential election (campaigning, voting, results, legitimacy, turnout)
- U.S. political institutions, public policy, or governance
- Formal political actors (candidates, elected officials, political parties)
- Explicit partisan or electoral positioning

If it is unclear whether the tweet contains substantive political content, label it “not_political”.

Definition 2 — Ideological Alignment (only if political)

Classify ideological alignment based only on the political or policy position explicitly expressed or endorsed by the tweet.

- “political_left”: Explicit endorsement of U.S. liberal or progressive positions or policies
- “political_right”: Explicit endorsement of U.S. conservative positions or policies
- “political_neutral”: Political content that is informational, mixed, anti-partisan, or lacks a clear ideological endorsement

Important Guidelines:

- Ideological alignment may be inferred from praise, criticism, or sarcasm only when the tweet expresses a clear and unambiguous political stance

- Emotional reactions or election outcome commentary without a clear stance should be labeled “political_neutral”
- Mentions of political figures alone are not sufficient for left/right classification
Output exactly one of: “not_political”, “political_left”, “political_right”, “political_neutral”.

For analysis, we derive three outcome variables from these labels:

- Politicalness: Binary indicator for whether the tweet is political (i.e., not `not_political`).
- Right-leaning: Binary indicator for `political_right`, computed over all tweets.
- Left-leaning: Binary indicator for `political_left`, computed over all tweets.

C.2 Polarizing Content

We classify each tweet into one of six mutually exclusive categories, according to [51]: `not_polarizing`, `polarizing_political`, `polarizing_religious`, `polarizing_sexual`, `polarizing_racial`, or `polarizing_other`. The classifier is prompted as follows:

Polarization Classification Prompt

You will be given a tweet collected in November 2024, during the U.S. presidential election period. Your task is to determine whether it is polarizing and, only if so, identify the primary type of polarization.

Definition:

A tweet is polarizing if and only if it satisfies both conditions:

1. It distinguishes between two or more social, political, religious, sexual, or racial groups (explicitly or implicitly), and
2. It frames those groups in a moralized conflict, such as: blame or hostility; moral superiority or inferiority; threat or danger; victim-oppressor narratives.

If either condition is not met, the tweet is not polarizing.

Important Constraints

- Mentioning politics, religion, gender, sexuality, race, or ideology alone is not sufficient for polarization.
- Statements, slogans, values, prayers, or opinions without group-based blame or moral conflict must be labeled “not_polarizing”.
- Emotional language is irrelevant unless it is directed at a group as a group.
- Attacks on individuals are polarizing only if they generalize to a broader group.
- Intra-group discussion or disagreement is not polarizing unless subgroups are framed as enemies, traitors, or morally inferior.

Output exactly one of: “not_polarizing”, “polarizing_political”, “polarizing_religious”, “polarizing_sexual”, “polarizing_racial”, “polarizing_other”.

For analysis, due to the sparsity of the polarizing categories, we collapse the six categories into a binary polarization indicator equal to 1 if the tweet is classified into any polarizing category, and 0 otherwise.

D Pooled Treatment Effects

We also report pooled treatment effects estimated without control variables as a robustness check. Results are nearly identical to the primary specification (Table 3), confirming that the controls do not affect the estimates.

Outcome	Location (D^{urb})		Age (D^{age})		Gender (D^{male})	
	$\hat{\beta}$	95% CI	$\hat{\beta}$	95% CI	$\hat{\beta}$	95% CI
Politicalness	-1.70 (0.137)	[-3.86, 0.46]	0.70 (0.501)	[-1.26, 2.66]	1.00 (0.408)	[-1.16, 3.16]
Right-Leaning	-0.50 (0.440)	[-1.87, 0.87]	0.90 (0.111)	[-0.28, 2.08]	-0.40 (0.561)	[-1.58, 0.78]
Left-Leaning	-0.60 (0.183)	[-1.38, 0.18]	-0.30 (0.383)	[-1.08, 0.48]	0.40 (0.379)	[-0.38, 1.18]
Polarization	-1.60 (0.020)	[-2.97, -0.23]	0.50 (0.444)	[-0.68, 1.68]	-0.70 (0.304)	[-2.07, 0.67]
Toxicity	0.10 (0.831)	[-0.68, 0.88]	0.80 (0.033)	[0.02, 1.58]	-0.10 (0.868)	[-0.69, 0.49]
High-reach	1.10 (0.381)	[-1.45, 3.65]	-0.90 (0.510)	[-3.64, 1.84]	-0.10 (0.928)	[-2.84, 2.64]

Table 10: Pooled treatment effects on algorithmic lift, without control variables. Each coefficient represents the average effect of a one-unit change in the demographic signal on the gap between *For You* and *Following* content rates. Coefficients for binary outcomes are scaled by 100 and expressed in percentage points. Values in parentheses are p -values computed from standard errors clustered at the account level. Significant coefficients ($p < 0.05$) shown in bold.

E Heterogeneous Treatment Effects

This appendix reports persona-specific treatment effects on algorithmic lift for all six outcomes, estimated via model (8). Each figure shows the estimated coefficient and 95% confidence interval for each persona-treatment pair. Filled markers indicate significance at the 5% level.

F Robustness of Amplification

This appendix presents robustness checks for the amplification ratio analysis in Section 4.2. We examine whether estimates are sensitive to the exposure window used for aggregation (that is, whether using only early observations versus the full deployment period produces different amplification ratios). A key identifying assumption of our estimator is that browsing one feed does not causally affect content shown in the other. Potential violations (such as engagement spillovers or preference updates across feeds) would make feeds more similar over time, biasing amplification estimates toward zero. If such cross-feed contamination occurs, we would expect amplification estimates to *decline* as accounts accumulate browsing history, since the algorithm would increasingly tailor both feeds to the same learned preferences.

To address this, we compute amplification ratios using cumulative exposure windows of increasing length: each account’s first day of activity only, first 7 days (Week 1), first 14 days (Week 2),




Figure 5: Persona-specific treatment effects on algorithmic lift for political content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.

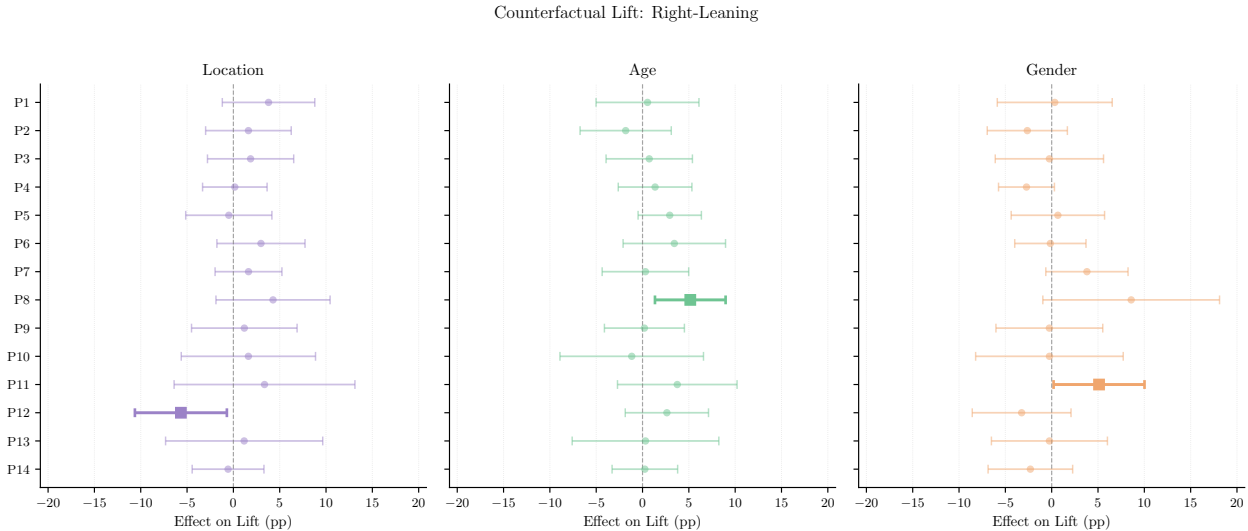


Figure 6: Persona-specific treatment effects on algorithmic lift for right-leaning content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.

and the full deployment period. Table 11 reports amplification estimates for each content type across these windows.

We observe that political and toxic content show no statistically significant change across exposure windows, with confidence intervals overlapping substantially throughout. Right-leaning, polarizing, and high-reach content exhibit attenuation from Day 1 to Week 1 (29.8% to 24.2%, 48.4% to 32.8%, and 14.5% to 3.3%, respectively), followed by stabilization. The attenuation is sharpest for high-reach content, which is amplified on Day 1(+14.5%, significant) but drops to a small and non-significant effect by Week 1 (+3.3%), where it stabilizes. Similarly, left-leaning content shows Day 1 de-amplification (-21.7%, significant) that attenuates toward zero by Week 1

Counterfactual Lift: Left-Leaning




Figure 7: Persona-specific treatment effects on algorithmic lift for left-leaning content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.

Counterfactual Lift: Polarizing




Figure 8: Persona-specific treatment effects on algorithmic lift for polarizing content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.

(-6.2%, no longer significant) and stabilizes thereafter.

The early attenuation observed for some content types (Day 1 to Week 1) likely reflects the algorithm calibrating to user behavior before reaching steady state. Crucially, estimates stabilize by Week 1 and remain unchanged thereafter. If cross-feed spillovers were substantially biasing estimates toward zero, we would expect continued attenuation throughout the deployment; instead, the estimates plateau. This stability supports our identifying assumption that browsing one feed does not meaningfully affect content shown in the other. Moreover, to the extent that any residual bias exists, it would attenuate estimates toward zero, implying that our reported amplification effects are, if anything, conservative.




Figure 9: Persona-specific treatment effects on algorithmic lift for toxic content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.




Figure 10: Persona-specific treatment effects on algorithmic lift for high-reach content. Each row is a persona; each panel shows a different treatment (location, age, gender). 95% confidence intervals.

Content Type	Day 1	Week 1	Week 2	Full
Political	+12.7% [+7.8, +18.1]	+14.6% [+9.5, +20.2]	+14.2% [+9.1, +19.9]	+14.2% [+9.1, +19.8]
Right-Leaning	+29.8% [+17.4, +43.0]	+24.2% [+16.2, +33.4]	+23.4% [+15.8, +32.1]	+23.3% [+15.8, +31.9]
Left-Leaning	-21.7% [-30.0, -13.2]	-6.2% [-14.1, +1.4]	-6.1% [-14.6, +2.2]	-6.1% [-14.6, +2.1]
Polarizing	+48.4% [+35.2, +63.1]	+32.8% [+24.2, +42.3]	+32.6% [+24.1, +42.0]	+32.3% [+23.7, +41.6]
Toxic	+37.9% [+9.9, +75.3]	+42.5% [+16.9, +74.5]	+39.4% [+14.5, +69.9]	+39.2% [+14.4, +69.7]
High-reach	+14.5% [+0.8, +35.5]	+3.3% [-6.3, +16.8]	+3.4% [-6.1, +16.6]	+3.4% [-6.1, +16.7]

Table 11: Amplification estimates by cumulative exposure window. Uncertainty is 95% CI from persona-level cluster bootstrap: in each iteration, we resample personas with replacement and recompute the amplification ratio, preserving the within-persona correlation structure.