From Surveillance to Manipulation: How Data Collection Enables Targeted Disinformation

Antony Seabra, Claudio Cavalcante, Sergio Lifschitz

[1]Departamento de Informatica - PUC-Rio

{amedeiros, cfraga, sergio}@inf.puc-rio.br

Abstract. *Surveillance capitalism has fundamentally transformed the digital economy by commodifying personal data to predict and influence behavior. While its economic and ethical implications have been widely studied, less attention has been given to its role in enabling the dissemination of disinformation. This paper explores how the same mechanisms that underpin surveillance capitalism, such as ubiquitous data collection, behavioral profiling, and algorithmic targeting, increasingly mediated by Artificial Intelligence, are leveraged to propagate fake news with unprecedented precision. Drawing from case studies, and empirical research, we demonstrate how actors exploit personalized data to craft and disseminate manipulative content aimed at specific psychological and ideological profiles. By analyzing collected results, we demonstrate that surveillance-based advertising also enables large-scale manipulation, posing serious risks to democratic integrity, public trust, and digital governance.*

1. Introduction

The concept of surveillance capitalism, as introduced by [Zuboff 2019], has emerged as one of the most influential frameworks for understanding structural transformations in the digital economy. Grounded in the large-scale and continuous extraction of behavioral data, this economic model aims to predict and shape human behavior for commercial gain, generating new forms of power and informational asymmetry. [Zuboff 2019] describes a system in which human experience is converted into raw material for data-driven business practices, often invisible to the user. Subsequent studies have reinforced this diagnosis, showing how digital platforms systematically monitor clicks, searches, social interactions, consumption patterns, and even physical displacements through the aggregation of geolocation data, mobile sensors, and browser tracking technologies [Gonzalez 2021, Seabra et al. 2024].

Empirical experiments demonstrate that, as individuals browse the web or interact with connected devices, their data is continuously collected and shared with dozens of third-party services. The tracking of navigation patterns, search queries, and product browsing history forms the foundation for commercial profiling, which fuels highly effective targeted advertising systems [Libert 2015]. These profiles reveal consumption habits, product preferences, behavioral rhythms, and recurring interests, enabling platforms to optimize engagement and advertising revenue through personalized content delivery.

This article posits that the very same mechanisms used to infer commercial interests are also capable of extracting sensitive dimensions such as political preferences, emotional vulnerabilities, and ideological inclinations [Brunton and Nissenbaum 2015]. These profiles are not only exploited for commercial advertising, but more frequently for opaque purposes such as disinformation campaigns [Tufekci 2014, Isaak and Hanna 2018]. We argue that this convergence between surveillance infrastructure and informational manipulation directly threatens public trust in journalism, democratic institutions, and civic deliberation.

The objective of this study is to empirically demonstrate how the current architecture of the web enables and amplifies the dissemination of manipulative content. To this end, we conducted controlled browsing experiments, packet-level traffic analysis, and behavioral profiling based on user interaction data. Through this evidence, we seek to uncover the interplay between surveillance capitalism and targeted disinformation, contributing to the broader debate on the ethical governance of digital ecosystems.

2. Background

This section provides the conceptual foundation for our study by reviewing the key phenomena that underpin our investigation: surveillance capitalism as a data-driven economic model, the technical ecosystem that enables the capture and interception of digital communications, and the mechanisms through which misinformation and disinformation are disseminated in online environments.

2.1. Surveillance Capitalism

The term *surveillance capitalism*, popularized by Zuboff [Zuboff 2019], describes a novel economic logic in which human behavior is systematically monitored, extracted, and commodified. In this model, digital platforms capture large amounts of behavioral data, such as clicks, search queries, location histories, social interactions, and even biometric signals, which are then analyzed and repurposed to predict and influence future actions. Initially developed to optimize advertising, this data-driven infrastructure has evolved into a complex network of tracking, profiling, and behavioral targeting that extends across virtually all domains of online activity.

The core mechanism of surveillance capitalism involves the transformation of user experience into behavioral surplus: data not required for service provision but valuable for predictive analytics. This surplus is sold to third parties or used internally to refine algorithmic personalization and advertising systems. The result is a feedback loop in which more data enables more precise predictions, driving both economic value and influence over user behavior.

This asymmetry of information and control has raised profound concerns regarding autonomy, privacy, and democratic integrity. Although most users are unaware of the extent and implications of this surveillance, the infrastructure built to serve commercial purposes has proven to be adaptable to other domains, including political influence and information manipulation.

2.2. Communications and Interceptions in the Web Ecosystem

The technical feasibility of surveillance capitalism depends on a vast, distributed infrastructure of communication protocols and data exchange mechanisms that silently capture and propagate user information across the web. When users interact with digital services, their devices continuously transmit data to a multitude of servers, ad networks, and third-party trackers. These communications are facilitated by protocols such as HTTP/HTTPS and are often enriched with metadata including cookies, user-agent strings, referral URLs, geolocation data, and device identifiers.

This flow of data is not confined to direct interactions with content providers. A single web request may involve dozens of embedded third-party scripts, trackers, and ad exchanges, each extracting behavioral signals that contribute to detailed user profiles. The opacity and decentralization of this ecosystem make it difficult for users to discern what data is collected, by whom, and for what purposes [Libert 2015, Christl 2017].

To study this hidden layer of the web, researchers have adopted interception techniques such as the use of `mitmproxy`, which allows for the decryption and inspection of HTTPS traffic. These tools offer critical visibility into the mechanisms of data flow and personalization, providing empirical evidence of how user information is captured and repurposed in real time. This technical layer is essential for understanding how behavioral profiles are constructed and how content-including advertisements and news-is algorithmically tailored based on inferred preferences.

2.3. Misinformation and Disinformation

The digital infrastructure initially designed for commercial personalization has increasingly been co-opted for the dissemination of misleading and manipulative content. *Misinformation* refers to false or inaccurate information that is spread without intent to deceive, while *disinformation* denotes deliberately deceptive content created and propagated with the purpose of manipulating beliefs or behaviors [Tufekci 2014, Isaak and Hanna 2018].

These forms of informational distortion thrive in algorithmic environments that prioritize engagement over accuracy. Social media platforms, search engines, and recommendation systems often amplify emotionally charged or polarizing content, which tends to perform better in attention-driven metrics. The same microtargeting techniques used in commercial advertising can be repurposed to direct specific narratives to ideologically segmented audiences, reinforcing cognitive biases and ideological echo chambers.

High-profile cases such as the Facebook-Cambridge Analytica scandal have demonstrated how behavioral data can be used to craft and deliver political messages tailored to psychological and ideological vulnerabilities. Research shows that such tactics were employed in the 2016 U.S. presidential election and the Brexit referendum, with disinformation campaigns targeting users with precision-engineered messages designed to influence electoral outcomes.

Understanding how surveillance-based infrastructures are weaponized for information manipulation is central to addressing the broader implications for democratic resilience, media trust, and public discourse. In this context, our study seeks to empirically expose the pipeline through which data collection, behavioral profiling, and disinformation dissemination converge.

3. Related Work

A growing body of research has examined the technological, economic, and sociopolitical foundations of surveillance capitalism and its impact on privacy and autonomy. Zuboff's seminal work [Zuboff 2019] provides the conceptual basis for understanding how digital platforms transform human behavior into a proprietary source of revenue through data extraction and predictive analytics. Building on this foundation, [Gonzalez 2021] and [Christl 2017] offer in-depth analyses of the institutional and infrastructural design of corporate surveillance systems, highlighting the opacity and lack of user control inherent to these practices.

In the advertising context, [Libert 2015, Seabra et al. 2024] found that most modern websites incorporate third-party tracking mechanisms that monitor user behavior for commercial profiling. These systems support microtargeting strategies powered by detailed personal data, including web navigation, geolocation, and social media activity, allowing advertisers to deliver precisely crafted messages to individual users.

More recent studies have shown how these same profiling and targeting mechanisms have been repurposed for political influence. [Tufekci 2014] discusses how algorithmic curation engineers the public sphere in ways that fragment collective discourse and erode democratic consensus. [Brunton and Nissenbaum 2015] argue for the use of obfuscation techniques to resist digital tracking, though they acknowledge the imbalance of power between users and platforms.

The Facebook-Cambridge Analytica scandal exposed how seemingly trivial data, such as Facebook likes or browsing patterns, can be used to infer psychological traits and politically segment the population with high precision [Isaak and Hanna 2018]. These tactics have been linked to polarizing electoral campaigns, including the Brexit referendum and the 2016 U.S. presidential election, as documented in parliamentary reports and journalistic investigations.

These studies underscore the shift from commercial data usage to sociopolitical manipulation, revealing a gray zone between legitimate marketing and engineered disinformation. Our work contributes to this literature by offering up-to-date empirical evidence on how the same surveillance-based infrastructure continues to be used, silently and automatically, to facilitate the spread of misinformation and disinformation at scale.

4. Methodology

To investigate how data collection techniques are exploited for disinformation, we designed a multi-phase experimental framework that simulates realistic web browsing behavior under controlled and reproducible conditions. Our study focuses specifically on the Brazilian political information ecosystem, capturing how different ideological profiles interact with the web and how these interactions influence the content and advertisements they receive. The methodology, illustrated in figure 1, consists of four sequential stages: multidimensional user profile creation, web traffic interception and data logging, content and source analysis, and comparative evaluation across profiles with distinct ideological configurations.

4.1. Multidimensional User Profile Creation

In the first phase, we developed three simulated user personas based on a multidimensional profiling strategy rather than a binary political spectrum. Instead of labeling profiles as simply conservative, progressive, or centrist, we assigned issue-specific stances across a range of ideological dimensions. These included economic policy (e.g., free-market vs. redistributionist), social values (e.g., traditionalist vs. progressive), environmental

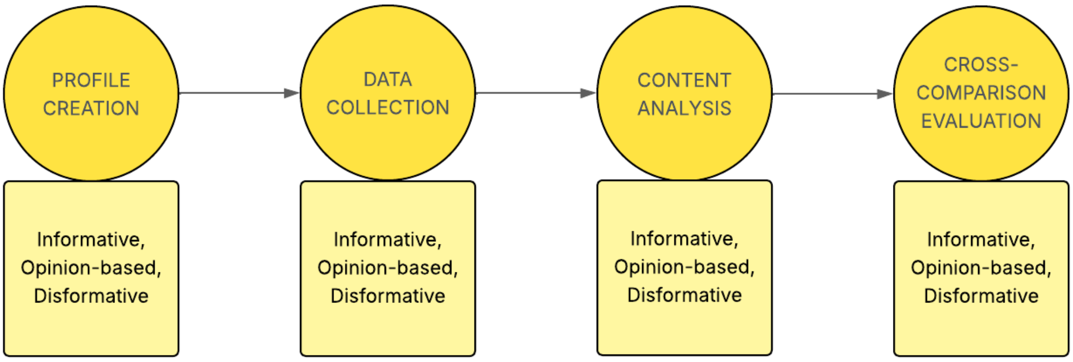
PROFILE CREATION
DATA COLLECTION
CONTENT ANALYSIS
CROSS-COMPARISON EVALUATION
Informative, Opinion-based, Disformative
Informative, Opinion-based, Disformative
Informative, Opinion-based, Disformative
Informative, Opinion-based, Disformative

Figure 1. Research Methodology

attitudes (e.g., climate activist vs. climate skeptic), public health perspectives (e.g., pro-vaccine vs. vaccine-hesitant), and immigration (e.g., open borders vs. restrictionist). This modeling approach reflects the complexity of real-world political identities, where individuals often hold mixed views depending on the issue. For instance, a persona might exhibit progressive environmental views while maintaining conservative social beliefs.

4.2. Multidimensional User Profile Design

Three simulated user personas based on multidimensional characteristics were constructed. Rather than relying on a binary ideological spectrum, each profile reflects a distinct combination of stances on economic, social, institutional, and media dimensions relevant to the Brazilian political landscape. Profile A is defined by pro-market economic preferences and socially conservative values. Profile B represents a progressive activist with left-leaning positions across both economic and social domains. Profile C embodies a politically ambivalent and distrustful individual, prone to populist and anti-establishment content. Together, these personas capture a diverse set of ideological configurations that enable the analysis of how personalized content delivery adapts to complex political behaviors.

Profile A represents a user with liberal economic views and conservative social values. This persona supports privatization, deregulation, and market-oriented reforms, while also aligning with traditional family norms, strong support for law enforcement, and skepticism toward progressive social movements. In terms of media consumption, Profile A frequently engages with content from outlets such as Jovem Pan, Gazeta do Povo, and conservative influencers on YouTube and Telegram.

Profile B reflects a progressive, left-leaning activist. This user supports social justice movements, public healthcare and education, wealth redistribution, and environmental protection. Profile B is critical of police violence, defends civil liberties, and advocates for the rights of minorities. Media preferences include independent journalism and progressive platforms such as Nexo Jornal, Brasil de Fato, as well as political debates on Instagram and Twitter.

Profile C simulates a politically ambivalent and distrustful voter who often gravitates toward populist and anti-establishment narratives. This user exhibits a mixture of ideological stances, expresses generalized distrust in institutions, and consumes content that ranges from traditional media to sensationalist sources. Profile C is characterized by erratic issue positions and a strong engagement with fringe YouTube channels, WhatsApp forwards, and clickbait aggregators.

Table 1. Key Features of the Simulated Political Profiles

Dimension	Profile A: Economic Liberal, Social Conservative	Profile B: Progressive Activist	Profile C: Populist, Distrustful Swing Voter
Economic Policy	Privatization, market liberalism	Redistribution, strong public services	Inconsistent; favors immediate economic protection
Social Values	Traditional family, anti-gender ideology	Pro-LGBTQ+, indigenous and civil rights	Ambiguous; responsive to moral panic and populist appeals
Institutional Trust	High trust in police, military, judiciary	Critical of law enforcement, pro-democracy movements	Generalized distrust of all institutions
Media Behavior	Jovem Pan, Gazeta do Povo, conservative YouTube channels	Nexo, Brasil de Fato, social media activism	Sensationalist news sites, WhatsApp forwards, fringe YouTube
Typical Searches	”School without party”, ”privatizations 2025”, ”censorship in agriculture”	”political violence”, ”indigenous rights”, ”fair tax reform”	”STF Corrpution”, ”new CPMF tax”, ”Telegram censorship”

While the number of simulated profiles in this study is limited to three, they were carefully constructed to reflect multidimensional ideological configurations relevant to the Brazilian political context. The goal is not to produce a statistically representative model of the electorate, but rather to empirically demonstrate how surveillance-based personalization mechanisms can adapt the delivery of disinformation to different behavioral signals. Each profile supports in-depth qualitative analysis of microtargeting effects, and the browsing routines were designed to expose subtle variations in content exposure. Future research may extend this methodology by incorporating a broader range of demographic, regional, and psychographic factors to explore generalizability and scale effects.

Each persona was instantiated in a dedicated browser environment configured with unique user agents, persistent cookies, and isolated session histories. Browsing behavior was automated using scripts that performed scheduled visits to a curated set of websites, including mainstream news portals, alternative media sources, political forums, and search engines. Interaction behaviors, such as scrolling, clicking, and keyword searches, were semi-randomized within the thematic constraints of each persona to simulate authentic yet consistent browsing routines over a two-week period. All the configuration files and scripts used in this study can be found at [Seabra 2025].

4.3. Web Traffic Interception and Data Logging

To monitor the data exchanged between the simulated users and online platforms, we deployed `mitmproxy`, a secure HTTPS interception and logging tool. This allowed us to capture all HTTP and HTTPS request-response cycles, including metadata such as cookies, session tokens, headers, redirect chains, third-party scripts, and advertisement payloads. These logs revealed the extent of data collection, the use of third-party trackers, and the structure of personalized content delivery.

The intercepted traffic was stored in structured formats such as JSON and CSV, timestamped, and tagged by profile identity to enable precise correlation between user behavior and received content. This phase provided the technical foundation for understanding how data generated by browsing activity propagates through the digital advertising and content recommendation ecosystem.

4.4. Content and Source Analysis

The captured content - including article bodies, headlines, advertisement copy, and embedded media - was analyzed using natural language processing (NLP) techniques to assess the information quality and the potential for manipulation in the material delivered to each persona. We employed a prompt-based classification approach using large language models (LLMs) in a zero-shot setting. Each content item was submitted to a structured prompt asking the model to categorize the text into one of three classes: *Informative* (fact-based and verifiable), *Opinion-Based* (editorial or subjective in tone), or *Disinformative* (misleading, false, or manipulative). In addition to returning a label, the model was instructed to provide a brief justification, enabling interpretability and downstream validation.

To ensure prompt calibration and increase the accuracy of the model, we adopted a few-shot strategy using representative examples drawn from the LIAR dataset, a publicly available corpus of fact-checked political claims annotated with varying levels of veracity [Wang 2017]. These examples served to contextualize the model's decision boundaries between factual content, opinionated language, and plain disinformation. Statements labeled as *true* or *mostly true* were aligned with our *Informative* class, while *half true* statements were treated as *Opinion-Based*, and those labeled as *false*, *barely true*, or *pants-on-fire* were associated with the *Disinformative* category. This mapping provided a consistent semantic foundation to guide LLM judgments.

It is important to acknowledge the methodological limitations of using large language models for disinformation classification, even in a few-shot setting. While the inclusion of labeled examples from the LIAR helps anchor the model's judgments and provides useful guidance for distinguishing between factual, opinion-based, and disinformative content, the dataset consists of short political statements in English and is derived from the context of U.S. politics. As such, it does not directly represent the linguistic structures, cultural nuances, or misinformation strategies specific to the Brazilian political landscape. In addition, the model itself does not have access to external fact-checking databases and cannot assess the intentionality behind a message, an essential component of disinformation. For these reasons, the model output was treated as an informed heuristic, particularly useful for identifying linguistic and rhetorical patterns indicative of manipulation. To enhance reliability, flagged content was cross-referenced when possible with contextual metadata, such as domain reputation, recurrent targeting patterns, and consistency of message framing across sessions. This strategy enables scalable, interpretable annotation while maintaining awareness of the epistemic and cultural boundaries of LLM-based content analysis.

4.5. Cross-Profile Comparative Evaluation

In the final phase of the analysis, we conducted a comparative evaluation of content exposure across the three multidimensional user personas. This evaluation focused on measuring the frequency and thematic distribution of retrieved content, the prevalence of ideological bias, and the presence of disinformative narratives as classified through prompt-based LLM analysis. Content types were aggregated per persona and categorized by topic (e.g., economy, public safety, institutional trust), enabling a structured comparison of how different political identities experienced distinct informational environments.

Special attention was paid to assessing whether particular issues-based traits, such as support for law enforcement, skepticism toward public health institutions, or advocacy for minority rights, were correlated with the delivery of targeted or manipulative content. We analyzed not only the content itself but also the recurrence of specific domains, ad payloads, and third-party trackers to examine how underlying personalization systems shaped each user's content landscape. Temporal patterns were also considered, such as whether certain disinformative narratives intensified in proximity to specific events or topics trending in the Brazilian political agenda.

By correlating content patterns with the predefined ideological traits and simulated behavioral routines of each user profile, we were able to assess how microtargeting systems adaptively modulate content delivery strategies. The findings from this phase provide empirical support for our central hypothesis: that the infrastructure originally built for behavioral advertising and commercial personalization is now actively leveraged to shape political perception, amplify polarization, and facilitate the covert dissemination of disinformation at scale.

5. Experimental Results

The content collected over a two-week browsing period was analyzed to determine the prevalence of informative, opinion-based, and disinformative material encountered by each persona. Figure 2 shows the distribution of content types across profiles. Although all profiles were exposed to a mix of content, profile C, defined by political ambivalence and institutional distrust, received a substantially higher volume of disinformative content (100 items) compared to profiles A and B (60 and 30 respectively). In contrast, Profile B, the progressive activist, received the highest volume of informative content and the least disinformation, indicating a more fact-centered informational environment.

Profile A, which reflects economically liberal but socially conservative views, received a higher proportion of opinion-based content, often framed around themes like economic freedom, institutional integrity, and social order. The content of Profile B leaned heavily towards social justice narratives and government accountability. However, Profile C's content was marked by sensationalism, high repetition of questionable sources, and emotionally charged language that often aligned with populist messages.

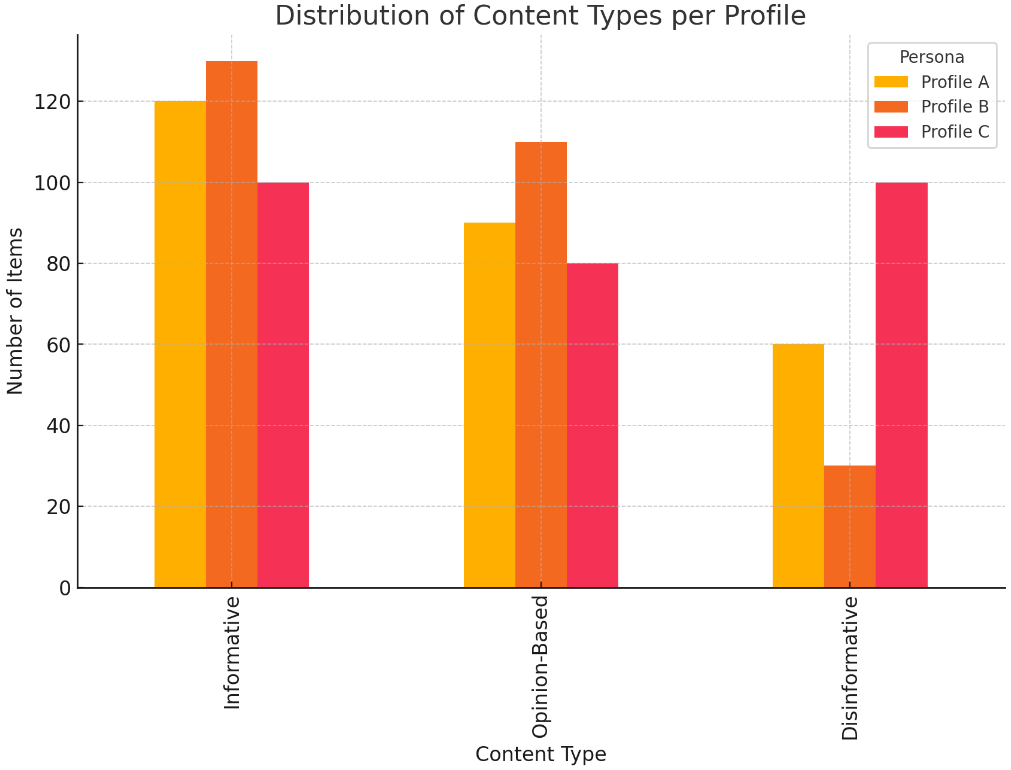

Figure 2. Distribution of Content Type per Profile. Source: Authors.

Figure 3 displays the number of unique content domains and third-party trackers encountered per profile. Profile C interacted with the lowest number of unique domains (30), but triggered the highest number of third-party trackers (75), suggesting a higher degree of exposure to ad-driven or commercially motivated information environments. This aligns with its behavioral routine, which prioritized less curated, more sensationalist content. Profile B, in contrast, exhibited both high domain diversity and relatively lower third-party tracking activity.

Temporal patterns of disinformation exposure are shown in Figure 3. In particular, Profile C experienced a spike in disinformative content during the second week of the simulation - coincident with increased engagement in political discussion forums and trending national topics. Profiles A and B maintained more stable, lower levels of disinformation exposure throughout the period, though Profile A showed occasional upticks around polarizing economic topics.

These results support the central hypothesis of the study: content delivery mechanisms adapt to behavioral traits, with disinformation disproportionately directed toward users who exhibit distrust, ambiguity, or ideological inconsistency. The underlying infrastructure of surveillance capitalism not only personalizes advertising but also conditions informational exposure in a way that may reinforce ideological segmentation and enable covert political manipulation.

5.1. Discussion

The experimental findings offer empirical support for the hypothesis that mechanisms originally designed for behavioral advertising are now repurposed to deliver ideologi-

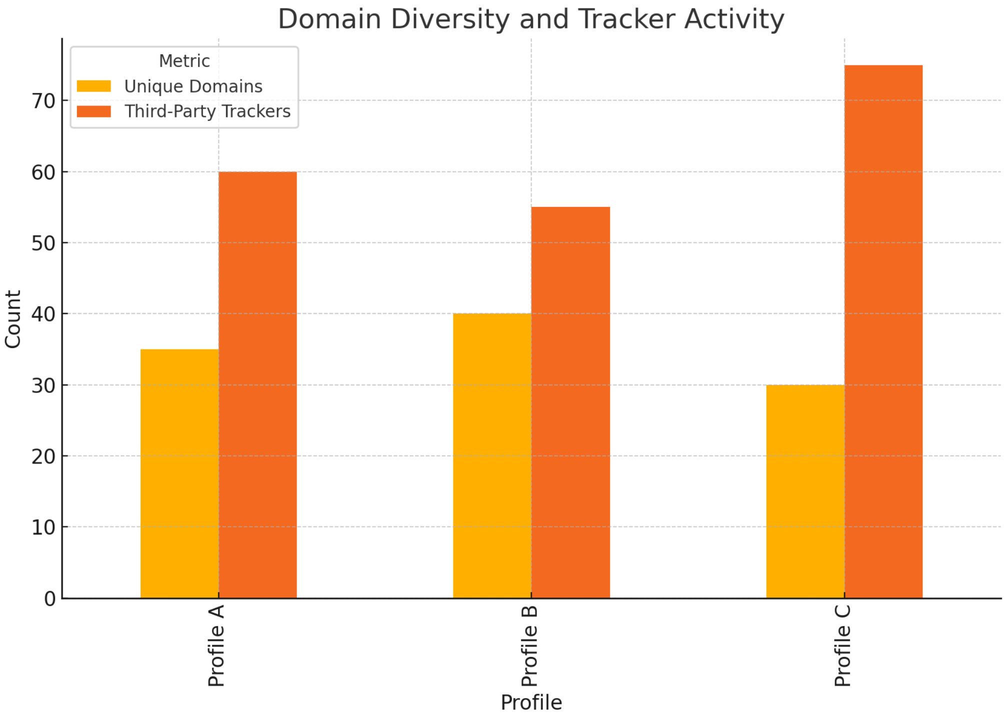

Figure 3. Domain Diversity and Tracker Activity. Source: Authors.

cally segmented and, in some cases, disinformative content. The differential exposure observed in the three multidimensional personas highlights how algorithmic content delivery systems exploit behavioral signals, not just for commercial purposes but also to adapt political messaging to user-specific vulnerabilities and beliefs.

Profile C, characterized by institutional distrust and ideological ambiguity, was disproportionately targeted with disinformative content. The browsing routine of this persona, which included interactions with low-credibility domains, YouTube recommendation loops, and viral clickbait, triggered a higher frequency of follow-up activity and content repetition. The pattern suggests a self-reinforcing cycle: behavioral ambiguity invites data-driven targeting, which increases exposure to manipulative content, which in turn may deepen distrust and polarization. The temporal spike in disinformation observed during politically salient periods further underscores the reactivity of these systems to current events, precisely when public attention is most fragile.

In contrast, Profiles A and B, both ideologically anchored in distinct worldviews, received a greater share of opinion-based or informative content, with relatively lower levels of disinformation. This may reflect the fact that established ideological commitments produce more stable engagement patterns, making these users less vulnerable or less valuable as targets for manipulative content streams.

These findings have several implications. First, they demonstrate that disinformation cannot be addressed solely as a content moderation problem; it must also be recognized as a systemic outcome of surveillance-driven personalization. Second, the study illustrates how the commodification of attention and behavioral data creates fertile ground for political influence operations, particularly among users whose browsing patterns re-

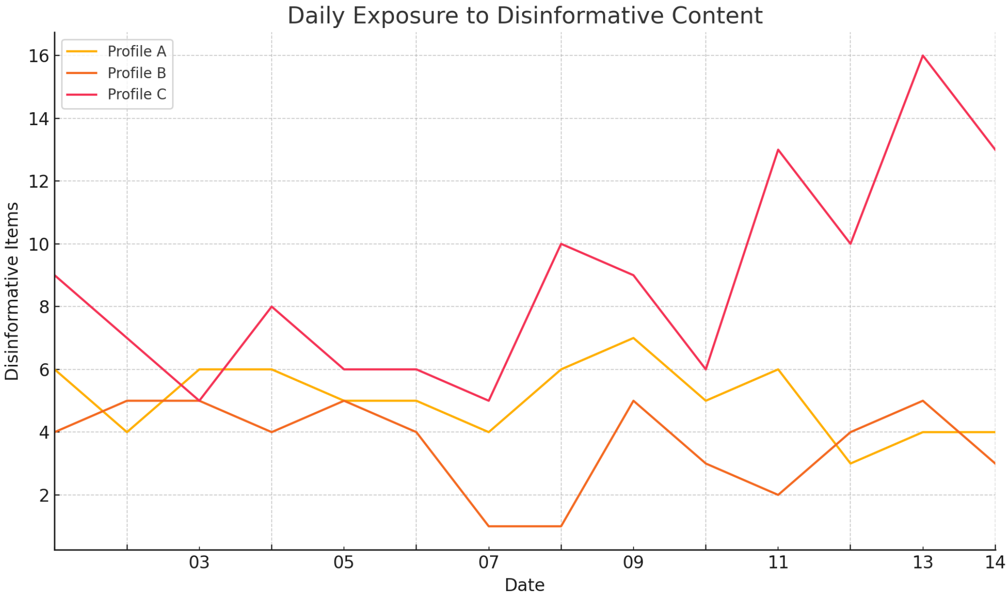

Figure 4. Daily Exposure to Disinformation Context. Source: Authors.

flect indecision or disillusionment. Finally, the results emphasize the need for greater transparency and accountability in the design and operation of content recommendation systems, particularly in contexts where public trust and democratic institutions are under strain.

While the use of LLM-based classification and external datasets like LIAR provided scalable methods for labeling, it is essential to reiterate the methodological limits: language mismatches, lack of contextual fact-checking, and the difficulty of detecting communicative intent. Nonetheless, the classification outputs, combined with metadata patterns and cross-profile analysis, reveal consistent and concerning trends in how digital ecosystems shape political information flows. Future research could expand the scope to include additional personas, multilingual datasets, and experimental interventions designed to disrupt or counteract algorithmic disinformation loops.

6. Conclusions and Future Work

This study investigated how the infrastructure of surveillance capitalism, originally developed to optimize commercial personalization, is now repurposed to support the targeted dissemination of disinformation within digital ecosystems. By simulating three multidimensional user personas aligned with diverse political orientations in the Brazilian context, and analyzing their content exposure over time, we demonstrated that algorithmic content delivery systems adaptively modulate information streams in response to behavioral traits. Our findings show that users who exhibit ideological ambiguity or generalized distrust, such as the simulated Profile C, are more likely to receive disinformative content, especially around politically charged events.

The results reinforce the notion that disinformation is not solely a problem of malicious content, but of systemic design. The same mechanisms that predict consumer preferences are being used to segment, polarize, and manipulate citizens at scale. As such, tackling this challenge requires going beyond content moderation and toward a critical examination of how data-driven personalization architectures shape informational autonomy and democratic deliberation.

While our methodology offers a scalable and interpretable framework for disinformation analysis using LLM-based classification and profile-driven simulation, it also has limitations. The LIAR dataset used for prompt calibration is in English and U.S.-centric, which introduces cultural and linguistic mismatches when applied to Brazilian political content. Additionally, the classification of disinformation through large language models remains a heuristic approach, limited by the absence of external fact validation and intent inference. Nevertheless, the integration of behavioral data, content analysis, and cross-profile comparison provides a valuable foundation for further investigation.

Future work may expand this framework by incorporating a wider range of simulated personas that vary not only in ideological dimensions but also in geography, age, and media literacy. The development or curation of labeled datasets in Portuguese, focused on Brazilian misinformation campaigns, would significantly enhance both accuracy and cultural relevance. Moreover, real-time packet interception combined with source credibility scoring could offer deeper insights into how disinformation campaigns evolve in response to user interaction. Finally, policy-focused research is needed to explore regulatory and technical interventions that can safeguard democratic discourse in the age of algorithmic manipulation.

References

Brunton, F. and Nissenbaum, H. (2015). *Obfuscation: A User's Guide for Privacy and Protest*. MIT Press.

Christl, W. (2017). Corporate surveillance in everyday life. Technical report, Cracked Labs. Accessed online: `https://crackedlabs.org/en/corporate-surveillance`.

Gonzalez, R. (2021). Algorithmic surveillance and the corporate world. *Surveillance & Society*, 19(1):5–14.

Isaak, J. and Hanna, M. J. (2018). User data privacy: Facebook, cambridge analytica, and privacy protection. *Computer*, 51(8):56–59.

Libert, T. (2015). Exposing the invisible web: An analysis of third-party http requests on 1 million websites. *International Journal of Communication*, 9:3544–3561.

Seabra, A. (2025). Github repository. `https://github.com/`. Accessed: 2025-05-05.

Seabra, A., Junior, L. G., and Lifschitz, S. (2024). Surveillance capitalism revealed: Tracing the hidden world of web data collection. *arXiv preprint arXiv:2412.17944*.

Tufekci, Z. (2014). Engineering the public: Big data, surveillance and computational politics. *First Monday*, 19(7).

Wang, W. Y. (2017). " liar, liar pants on fire": A new benchmark dataset for fake news detection. *arXiv preprint arXiv:1705.00648*.

Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier.