

# Measuring The Unseen Hand: Methodological Challenges in Quantifying the Welfare Effects of Algorithmic Persuasion

Sharath Kumar M P<sup>1</sup> and Dr V Sumathi<sup>2</sup>

<sup>1</sup>Research Scholar, Department of Economics, PSG College of Arts and Science, Coimbatore.

<sup>2</sup>Assistant Professor, Department of Economics, PSG College of Arts and Science, Coimbatore.

**Abstract**—In today's India, our lives are increasingly intertwined with the digital world. From shopping on Flipkart and Amazon to ordering food on Zomato and managing investments on platforms like Zerodha and Groww, powerful Artificial Intelligence (AI) systems are constantly working behind the scenes.<sup>[1]</sup> These algorithms learn our habits and "nudge" us toward certain choices through personalised recommendations and curated information a process known as algorithmic persuasion.<sup>[2]</sup> This technology promises unprecedented convenience and efficiency, but it also raises a critical question: Are these digital nudges truly improving our well-being, or do they primarily serve the commercial interests of the companies that deploy them? This paper provides a comprehensive methodological review of the profound challenges researchers face in quantifying the real-world welfare effects of these persuasive systems. We identify and elaborate on five core, interlocking challenges: (1) the fundamental difficulty in measuring consumer surplus, especially its non-monetary components; (2) the ambiguity in defining a credible counterfactual against which to evaluate sophisticated algorithmic systems; (3) the complexity of establishing robust causal inference in dynamic, endogenous digital environments; (4) the severe data access limitations and algorithmic opacity that impede independent research and audits; and (5) the difficulty of capturing long-term, dynamic, and heterogeneous effects across a diverse population. By examining these challenges, this paper aims to create a detailed methodological roadmap to better understand and shape India's burgeoning digital economy for the equitable benefit of all its citizens.

## I. INTRODUCTION

The digital revolution has irrevocably transformed the Indian landscape. Propelled by some of the world's most affordable mobile data and widespread smartphone penetration, hundreds of millions of

Indians are now active participants in the online economy.<sup>[1]</sup> This has created a fertile ground for companies to deploy sophisticated AI to understand, predict, and influence consumer behaviour on a massive scale.<sup>[1]</sup> When Amazon India suggests a book, Flipkart shows you a new mobile phone, or a financial app like Groww recommends a specific mutual fund, it is not a random guess. It is the calculated output of a learning algorithm that has analysed your past clicks, purchases, search queries, and even the items you lingered on but ultimately did not buy.<sup>[1]</sup>

This is the essence of algorithmic persuasion: the use of data-driven, automated systems to gently but deliberately influence human beliefs, attitudes, and behaviours.<sup>[2]</sup> Unlike traditional advertising, which often relied on broad, one-to-many messaging, algorithmic persuasion operates at an individual level, creating what some have termed "persuasive architectures" digital environments meticulously designed to shape user choices.<sup>[3]</sup>

On the surface, this appears to be a win-win scenario. Consumers benefit from saved time, reduced effort, and the discovery of products that perfectly match their needs. Companies, in turn, see increased sales and user engagement. A case study of Flipkart, for instance, revealed that its AI-powered personalisation led to a remarkable 25% increase in sales and a 30% rise in user engagement.<sup>[1]</sup>

However, this technological prowess also grants companies immense power to influence our decisions, often in ways that are subtle and opaque.<sup>[4]</sup> The applied algorithmic decision-making is one of both great promise and significant peril.<sup>[4]</sup> While it holds the potential to promote economic justice by distributing opportunities more broadly, it also risks producing

biased, discriminatory, and otherwise problematic outcomes in critical areas like employment, credit, and housing.<sup>[4]</sup> This dual nature presents a formidable challenge for society: how do we rigorously measure whether this technology is, on balance, making our lives better? How do we quantify its impact on our overall well-being, or "welfare"?

This paper undertakes a deep dive into this complex question from a methodological perspective. We argue that a clear, evidence-based answer is currently elusive due to five fundamental and interconnected challenges that confront any researcher in this domain. First, we explore the Measurement Challenge, the inherent difficulty of putting a concrete value on welfare, especially its non-monetary aspects. Second, we examine the Counterfactual Challenge, the problem of establishing a realistic baseline for comparison. Third, we detail the Causal Inference Challenge, the difficulty of separating cause from effect in a complex web of digital interactions. Fourth, we address the Access Challenge, where corporate secrecy and algorithmic opacity create a "black box" that is nearly impenetrable to outside scrutiny. Finally, we discuss the Dynamics Challenge, the need to account for effects that vary across India's diverse population and evolve over the long term.

By systematically dissecting these issues, this paper aims to move beyond a simple acknowledgment of the problem. Our goal is to provide a structured, detailed framework for understanding *why* measurement is so hard, and to propose a concrete path forward for developing the tools and institutions necessary for effective oversight in the Indian context.

## II. CONCEPTUAL FRAMEWORK: DEFINING AND MEASURING WELFARE IN THE DIGITAL AGE

To diagnose the challenges in measuring welfare, we must first establish a clear and comprehensive conceptual framework. This requires carefully defining what constitutes algorithmic persuasion and operationalising the multifaceted concept of "welfare" in a way that is relevant to the digital experiences of Indian consumers.

**A. From Advertising to Algorithmic Persuasion**  
The economic literature has long grappled with the role of advertising, traditionally distinguishing

between two primary functions: informative and persuasive. Informative advertising provides consumers with objective facts about a product's existence, price, or features, thereby reducing information asymmetry and, in theory, intensifying price competition.<sup>[5]</sup> In contrast, persuasive advertising seeks to alter consumer tastes and build brand loyalty, often through emotional appeals, which can soften price competition and create market power.<sup>[5]</sup>

Algorithmic persuasion, however, masterfully blurs this distinction. It leverages vast datasets to deliver messages that are simultaneously informative (i.e., highly relevant to a user's inferred preferences) and persuasive (i.e., framed and timed to appeal to known cognitive biases and emotional triggers).<sup>[6]</sup> The goal is to change behaviour without overt coercion, often by subtly redefining the "choice architecture"-the environment in which we make decisions.<sup>[7]</sup> This might involve changing the default option, highlighting a particular choice, or creating a sense of urgency.

### B. A Multi-dimensional View of Welfare

A robust analysis of these systems requires a definition of welfare that extends beyond simple monetary outcomes. We conceptualise welfare as having three critical dimensions: consumer surplus, non-monetary factors, and distributional effects.

- **Consumer Surplus (Value for Money):** The traditional cornerstone of welfare economics, consumer surplus is the difference between what a consumer is willing to pay for a good and the price they actually pay. It represents the net economic gain to the consumer from a transaction. While conceptually simple, its empirical measurement is a significant challenge, as it requires knowledge of the entire demand curve, not just the observed market price.
- **Non-Monetary Factors (Time, Effort, and Match Quality):** A significant portion of the value provided by algorithmic systems is non-monetary. Recommender systems, in particular, can dramatically reduce search costs, both in terms of the time spent browsing and the cognitive effort required to evaluate options.<sup>[8]</sup> They can also mitigate the uncertainty associated with purchasing unfamiliar products. Perhaps most importantly, they can improve "match quality" by

helping consumers discover niche products that better align with their specific tastes the so called "long tail effect."<sup>[9]</sup> Any complete welfare measure must account for these crucial, albeit hard to quantify, benefits and costs.

- **Distributional Effects and the Bayes Welfare Set:** An analysis focused solely on aggregate or average welfare is dangerously incomplete, especially in a country as diverse as India. Algorithmic personalisation, by its very nature, creates heterogeneous outcomes, producing winners and losers.<sup>[10]</sup> A powerful theoretical tool for analysing these distributional trade-offs is the Bayes welfare set, a framework developed by economists Laura Doval and Alex Smolin.<sup>[11]</sup> This framework models recommendation algorithms and other information policies as mechanisms that shape the welfare of different types of individuals within a population.<sup>[11]</sup> The Bayes welfare set is the complete set of all possible "welfare profiles" vectors describing the expected utility for each group (e.g., urban vs. rural consumers, high-income vs. low-income)- that are achievable under some feasible algorithmic policy.<sup>[11]</sup> It is the information-age equivalent of the classic utility possibility set from welfare economics.

This framework provides several profound insights. First, it offers a geometric characterisation of the feasible welfare trade-offs, allowing a society's choice of an algorithm to be framed as a choice of a specific welfare distribution from this set.<sup>[11]</sup> Second, it formally demonstrates that fairness (e.g., ensuring equal outcomes across groups) can be in direct conflict with Pareto efficiency (where no single group can be made better off without making another group worse off).<sup>[11]</sup> Third, and crucially for practical application, the framework shows that the Pareto frontier of the welfare set - the set of all efficient outcomes- can be traced by solving a series of standard Bayesian persuasion problems, where a hypothetical social planner acts as the information designer.<sup>[11]</sup> This elegantly connects the societal problem of choosing a fair and efficient algorithm to a well-understood economic model, providing a path for both theoretical analysis and computational exploration.<sup>[12]</sup>

### III. CORE METHODOLOGICAL CHALLENGES

While a clear conceptual framework for welfare exists, its empirical measurement is fraught with difficulty. This section provides a detailed examination of the five core methodological challenges that confront any researcher attempting to quantify the real-world welfare effects of algorithmic persuasion.

#### A. The Measurement Challenge: The Difficulty of Quantifying Well-Being

The first and most fundamental hurdle is the difficulty of measuring consumer surplus itself. Estimating consumer surplus requires identifying the entire demand curve for a product, yet researchers typically only observe a single point on that curve: the equilibrium price and quantity where supply meets demand.<sup>[13]</sup> To trace the full curve, one must isolate exogenous shifts in supply while holding demand constant, a notoriously difficult task in most market settings.

A landmark study by Cohen et al. on the ride-sharing service Uber provides a rare exception.<sup>[13]</sup> By leveraging the richness of Uber's data and its "surge" pricing algorithm which effectively runs thousands of natural prices experiments every day the researchers were able to estimate demand elasticities at multiple points along the demand curve and construct a robust estimate of consumer surplus. However, this approach is the exception that proves the rule; such granular, high-frequency pricing data is almost never available to independent researchers, making this method impossible to replicate in most contexts.

The challenge is even greater for non-price factors. Traditional economic measures like the Consumer Price Index (CPI) struggle to account for welfare changes that come from the introduction of new goods, quality improvements in existing goods, or shifts in shopping patterns to lower-priced online stores.<sup>[14]</sup> Correctly incorporating these effects requires not just price data, but also detailed quantity or expenditure data, which is often unavailable.<sup>[14]</sup> While sophisticated structural models of consumer search can be used to estimate the value of non-monetary factors like reduced search costs, these models are themselves complex, data-intensive, and rely on strong assumptions about consumer behaviour.<sup>[15]</sup>

#### B. The Counterfactual Challenge: What Is the Proper Baseline?

A causal estimate of an algorithm's welfare effect requires a credible counterfactual: what would have happened in its absence?<sup>[16]</sup> Defining this baseline is far from straightforward. Should a sophisticated personalised recommender system be compared to a world with no recommendations at all, or to one with a simpler, non-personalised algorithm, such as a "most sold" or "highest rated" list?

The choice of this baseline dramatically alters the conclusion. This gives rise to what has been termed the "personalization paradox."<sup>[9]</sup> An algorithm can be designed to maximise firm profit - for instance, by exploiting the fact that consumers are less price-sensitive to products placed in more salient positions - and thereby reduce consumer surplus relative to a theoretical social optimum.<sup>[17]</sup> However, that same profit-maximising algorithm may still generate substantially higher consumer surplus than the simpler, less efficient systems it replaced.<sup>[18]</sup>

Large-scale field experiments confirm this dynamic. One study with a major online retailer found that its personalised ranking algorithm, while incorporating profitability, still substantially increased consumer surplus compared to the uniform bestseller list it replaced.<sup>[18]</sup> Another study simulated a profit-maximising recommender and found that while it reduced surplus by 8.4% compared to the social optimum, it still yielded higher surplus than systems based on "most sold" or "highest rated" products.<sup>[17]</sup> Consumers, observing a clear improvement over a poor baseline, are unlikely to be aware of the further surplus being extracted by the firm and are thus unlikely to complain.<sup>[17]</sup> This creates a "good enough" trap, where firms have little incentive to move toward a truly consumer-centric algorithm, making it incredibly difficult for researchers to quantify the "unseen" welfare loss relative to an unobserved, optimal counterfactual.

#### C. The Causal Inference Challenge: Disentangling Cause and Effect

Establishing a clear causal link between an algorithm and a welfare outcome is exceptionally difficult due to the complex, dynamic, and interconnected nature of digital platforms. The central challenge of causal inference is estimating the unobservable counterfactual, and this problem is particularly acute in the context of algorithmic persuasion for several reasons.<sup>[16]</sup>

First, these systems are characterised by endogeneity and simultaneity. For example, a recommender system has a positive effect on sales, but sales data simultaneously feeds back into the algorithm to update future recommendations.<sup>[18]</sup> This feedback loop makes it difficult to isolate the causal effect of the recommendation itself from the sales patterns that are shaping it.

Second, the risk of unobserved confounding is high. It is difficult to distinguish whether a change in user behaviour is caused by the algorithm or by the user's own evolving preferences, which the algorithm is merely reflecting. This requires moving beyond identifying mere statistical correlations to uncovering true cause-and-effect relationships, a task complicated by the presence of latent (unobserved) variables that influence both user behaviour and algorithmic output. Third, modern consumer journeys suffer from data fragmentation. A consumer may see an ad on a social media app on their phone, browse on a tablet, and complete the purchase later on a desktop computer. Tracking users across these multiple devices and channels is a significant technical challenge, making it hard to accurately attribute conversions and revenue to specific persuasive interventions. These issues mean that standard statistical methods are often insufficient, requiring more sophisticated approaches from the growing field of "causal machine learning" designed to handle such complexity.<sup>[19]</sup>

#### D. The Access Challenge: Opaque Algorithms and Proprietary Data

Perhaps the most significant practical barrier to independent research is the combination of algorithmic opacity and restricted data access. Many modern machine learning systems are effectively "black boxes," meaning their internal workings are opaque even to the firms that deploy them, let alone to outside stakeholders.<sup>[20]</sup> These forces independent researchers, journalists, and civil society groups to conduct "algorithm audits," where they systematically probe a system with inputs and observe the outputs to infer its behavior.<sup>[21]</sup>

However, these audits are severely hampered by a lack of data access.<sup>[22]</sup> Firms are often reluctant to disclose the granular, user-level data required for rigorous welfare analysis due to a combination of commercial secrecy and privacy regulations like the EU's General Data Protection Regulation (GDPR).<sup>[23]</sup> This lack of

access is a primary impediment to effective, independent research and oversight.<sup>[23]</sup> Auditors must often resort to creative but limited workarounds, such as using publicly available data or attempting to replicate the algorithm themselves, which can expose them to legal retaliation and limit the robustness of their findings.<sup>[23]</sup>

While privacy-preserving techniques like differential privacy and synthetic data have been proposed as a solution, they are not a panacea. Recent research simulating audits under various data access conditions found that these methods can mask or understate the very biases and disparities that auditors seek to uncover.<sup>[23]</sup> For example, synthetic datasets tend to "invisibilise" disparities, leading to underestimations of algorithmic bias.<sup>[23]</sup> This suggests that while privacy is a legitimate concern, it should not be used as an excuse to prevent necessary oversight. Institutional solutions are needed to provide secure access to real, granular data for vetted researchers.<sup>[23]</sup>

#### E. The Dynamics Challenge: Capturing Long-Term and Heterogeneous Effects

Finally, a comprehensive welfare analysis must account for effects that vary over time and across people.

Most empirical studies measure the short-term effects of an intervention. However, the long-term consequences could be very different. Behavioural changes from nudges might decay once the intervention stops, or they might persist and become habitual. Encouragingly, some recent long-term field experiments suggest that effects can be sustainable. One study on cancer patients found that machine learning-triggered nudges to doctors led to a significant and lasting increase in important end-of-life conversations.<sup>[24]</sup> Another month-long experiment on YouTube found that an "algorithmic nudge" had a significant and sustainable effect on news consumption that persisted even after the intervention ended.<sup>[25]</sup> Capturing these long-term dynamics is methodologically demanding but crucial for a full welfare accounting.

Furthermore, average treatment effects can be dangerously misleading in a country as diverse as India. Algorithmic personalisation, by its very nature, creates a distribution of outcomes. A field experiment on personalised pricing found that while total consumer surplus declined, over 60% of individual

consumers actually benefited from the personalisation.<sup>[10]</sup> This heterogeneity means that any policy evaluation must move beyond averages to understand the full distributional consequences, identifying who wins, who loses, and by how much.

#### IV. THE INDIAN EXPERIENCE: PROMISES, PROBLEMS, AND POLICY IMPLICATIONS

These methodological challenges are universal, but they take on a special significance in the unique context of India's rapidly growing and incredibly diverse digital economy.

##### A. The Promise of Digital India:

- **A Catalyst for Economic Growth:** AI is a major engine of India's e-commerce sector. As noted, a case study of Flipkart revealed that AI-powered personalisation led to a 25% increase in sales and a 30% rise in user engagement, demonstrating a clear link between this technology and economic activity.<sup>[1]</sup>
- **Deepening Financial Inclusion:** In the financial technology (fintech) space, AI is playing a transformative role. Robo-advisors from companies like Zerodha, Paytm Money, and Groww provide personalised investment advice to millions of Indians who may not have access to traditional financial advisors.<sup>[1]</sup> One survey found that 65% of users on these platforms relied on AI-driven insights for their stock market investments, suggesting AI can be a powerful tool for improving financial literacy AND decision-making across a wider segment of the population.<sup>[1]</sup>
- **Enhancing Efficiency:** At an organisational level, AI-driven nudges are being used to improve employee performance and operational efficiency, which can lead to better services and lower costs for consumers.

##### B. The Problems and Perils

- **The Digital Divide and Algorithmic Bias:** The benefits of AI are not being shared equally. Algorithms are trained on data, and in India, this data is often skewed towards urban, English-speaking, and higher-income populations. This can lead to significant biases. One study noted that consumers with smaller "digital footprints" (less online history) received fewer personalised

offers, creating a new form of digital inequality.<sup>[1]</sup> This is a critical issue in a country with vast linguistic, cultural, and economic diversity. An algorithm that works well for a user in Mumbai may be ineffective or even discriminatory for a user in rural Bihar.

- Erosion of Privacy and Autonomy: There is a growing and palpable concern among Indians about how their personal data is being used. A survey revealed that 58% of users were worried about AI tracking their online behaviour without their clear permission.<sup>[1]</sup> This is not just a privacy issue; it is an autonomy issue. As firms gain the power to shape user behaviour and preferences by "tweaking their code," it threatens to rob individuals of their ability to make independent choices.<sup>[4]</sup>
- The "Black Box" and Lack of Trust: The problem of algorithmic opacity is deeply felt by Indian consumers. The same study found that only 24% of respondents felt they understood how AI algorithms made recommendations for them.<sup>[1]</sup> This lack of transparency makes it difficult for people to trust the system or know if they are being treated fairly.
- Harmful "Dark Nudges": The power of persuasion can be used for ill. Techniques like creating false urgency ("Only 1 item left!"), showing misleading social proof, or exploiting cognitive biases can push people into making hasty or poor financial decisions they later regret.

#### V. A PATH FORWARD: A METHODOLOGICAL ROADMAP FOR INDIA

Measuring the true welfare impact of algorithmic persuasion is a complex but not insurmountable task. To make progress, India needs a concerted, multi-pronged approach that combines methodological innovation with institutional reform.

1. Investing in Advanced Econometric Modeling: To overcome issues of endogeneity and unobserved confounding, researchers should make greater use of structural models that explicitly map the underlying economic environment. In the context of online choice, this involves building models of consumer

search and firm behaviour that can disentangle the various channels through which algorithms affect welfare. For instance, structural models can separate the effect of a recommender system on consumer search costs from its effect on consumer beliefs about hidden product quality.<sup>[26]</sup> By estimating the primitive parameters of consumer tastes and search costs, these models allow for counterfactual simulations of alternative algorithmic designs.

2. Promoting Large-Scale, Context-Specific Field Experiments: Randomised Controlled Trials (RCTs), or field experiments, remain the gold standard for establishing causality. The literature on algorithmic persuasion would benefit immensely from more large-scale, long-term experiments conducted in the Indian context. These experiments should be designed to test not only the average effect of an intervention but also the mechanisms behind it. This requires greater collaboration between academic researchers and digital platforms, which possess both the user base and the technical infrastructure to run such experiments.
3. Leveraging Causal Machine Learning: The emerging field of causal machine learning (causal ML) offers a powerful toolkit for tackling the complexity of digital environments.<sup>[16]</sup> While traditional ML focuses on prediction and correlation, causal ML adapts these techniques for causal inference.<sup>[19]</sup> These methods are particularly well-suited for handling the high-dimensional data and complex, non-linear interactions that characterise algorithmic systems. A key strength of causal ML is its ability to systematically uncover heterogeneous treatment effects, moving beyond average effects to identify for whom an intervention is most (or least) effective, a crucial capability for policymaking in a diverse nation like India.
4. Building Frameworks for Secure Data Access and Audits: Methodological advances will be of limited use without a solution to the data access problem. Ensuring that independent, third-party researchers can conduct meaningful audits is critical for public accountability. This requires building new institutional frameworks for data

access that balance legitimate privacy concerns with the public interest. Potential solutions include the development of secure data environments or "data clean rooms" where vetted researchers can analyze sensitive data remotely, and regulatory mandates that require platforms to provide such access.[22]

## VI. CONCLUSION

The "unseen hand" of algorithmic persuasion is already a powerful force shaping the Indian economy and society. It is more than just a marketing tool; it is a technology that is fundamentally re-architecting our choices, our markets, and our access to information. While it offers incredible benefits in terms of convenience, efficiency, and discovery, it also carries significant risks of manipulation, bias, and unfairness, particularly in a country as diverse as India.

As this paper has shown, understanding the true balance of these benefits and risks is a major methodological challenge. Our traditional tools for measuring economic welfare are struggling to keep up with the pace of technological change. However, by acknowledging these hurdles and working towards innovative solutions through more sophisticated research methods, greater transparency, and smarter regulation we can begin to measure the unseen hand of the algorithm. For a country like India, which stands to gain so much from the digital economy, ensuring that this technology truly serves the well-being of all its citizens is one of the most important tasks of our time.

## REFERENCES

- [1] Upadhyay, A. K., & Khandelwal, K. (2024). AI in Behavioral Economics and Decision-Making. *Journal for Research in Applied Sciences and Biotechnology*, 3(4).
- [2] Fogg, B. J. (2009). A Behavior Model for Persuasive Design; and Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*.
- [3] Rose, J. (2020). The Architecture of Algorithm-driven Persuasion. *Journal of Information Architecture*.
- [4] Slaughter, R. K. (2021). Algorithms and Economic Justice: A Taxonomy of Harms and a Path Forward for the Federal Trade Commission. *Yale Journal of Law & Technology*.
- [5] Schmalensee, R. (1982). Advertising and Market Structure.
- [6] Fogg, B. J. (2009). A Behavior Model for Persuasive Design.
- [7] Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*.
- [8] De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2012). Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior. *American Economic Review*.
- [9] Kaye, A. P. (2025). The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets. *Toulouse School of Economics*.
- [10] Dubé, J. P., & Misra, S. (2023). Personalized Pricing and Consumer Welfare.
- [11] Doval, L., & Smolin, A. (2024). The Bayes Welfare Set. *Journal of Political Economy*.
- [12] Kamenica, E., & Gentzkow, M. (2011). Bayesian Persuasion. *American Economic Review*, 101(6), 2590-2615.
- [13] Cohen, P., et al. (2016). Using Big Data to Estimate Consumer Surplus: The Case of Uber. *NBER Working Paper*.
- [14] Hausman, J. A. (2003). Sources of Bias and Solutions to Bias in the Consumer Price Index. *Journal of Economic Perspectives*.
- [15] De los Santos, B., et al. (2012); Harrison, G. W., & Swarthout, J. T. (2022). Behavioral welfare economics and risk preferences: a Bayesian approach; and Hong, J. C., & Kim, J. (2025). Pretraining Estimators for Structural Econometric Models.
- [16] Varian, H. (2016). Causal inference in economics and marketing.
- [17] de Clippel, G., Eliaz, K., & Rozen, K. (2022). Welfare properties of profit-maximizing recommender systems.
- [18] Donnelly, R., Kanodia, A., & Morozov, I. (2024). Welfare Effects of Personalized Rankings. *Marketing Science*, 43(1), 92-113.

- [19] Athey, S., & Imbens, G. W. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*.
- [20] Broussard, M., et al. (2022). Using Algorithm Audits to Understand AI. *Stanford HAI*.
- [21] Broussard, M., et al. (2022) and Sandvig, C., et al. (2014). Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms.
- [22] Broussard, M., et al. (2022) and Djolonga, J., et al. (2023). Access Denied: A Triptych on the State of Public Interest Research on Online Platforms.
- [23] Wernerfelt, N., Johnson, G., & Gummadi, K. P. (2024). The Effect of Privacy Regulation on the Digital Advertising Market. *NBER Working Paper*.
- [24] Manz, C. R., et al. (2023). Long-term Effect of Machine Learning-Triggered Behavioral Nudges on Serious Illness Conversations and End-of-Life Outcomes Among Patients with Cancer. *JAMA Oncology*.
- [25] Yu, X., et al. (2024). Nudging recommendation algorithms increases news consumption and diversity on YouTube. *PNAS Nexus*.
- [26] De los Santos, B., et al. (2012); and Hong, J. C., & Kim, J. (2025). Pretraining Estimators for Structural Econometric Models.