e-ISSN: 3079-6377 p-ISSN: 3079-6369

Beyond Interference: Estimating Causal Network Effects in Digital Advertising Markets

Yi Han
Meta Fintech (Monetization), Meta, Menlo Park, CA, USA, 94025
Email: han.y@wustl.edu

Abstract

Digital advertising platforms operate within complex social networks where user responses to ads are not independent but are influenced by the actions and exposures of their peers. This interdependence, known as network interference, fundamentally violates the Stable Unit Treatment Value Assumption (SUTVA) of traditional causal inference models, leading to significant attribution bias and suboptimal budget allocation. This paper addresses the challenge of estimating the true causal effect of digital advertising expenditures, moving beyond direct response metrics to quantify the total demand generation effect, which includes both direct and socially mediated peer influences. We propose a novel, robust methodological framework that integrates high-dimensional network clustering with a two-stage least squares (2SLS) instrumental variable approach and a synthetic control method. Using peer adoption rates and platform algorithm shocks as instruments, we isolate exogenous variation in ad exposure. We then employ a synthetic control-based estimator to construct counterfactual outcomes for treated network segments, effectively absorbing the bias from spillover effects. Our empirical analysis, applied to a simulated marketplace mirroring real-world advertising dynamics, demonstrates that standard models underestimate the total return on ad spend by approximately 30-45%. The findings provide a scalable solution for advertisers to de-bias their attribution models and optimize campaigns in the presence of social interference, offering a significant advancement in the econometrics of digital marketing measurement.

Keywords

Causal Inference; Network Interference; Digital Advertising; Instrumental Variables; Synthetic Control Method; Attribution Bias

1.Introduction

The digital advertising ecosystem has evolved into a highly interconnected and socially dynamic environment [1]. Platforms such as Meta, TikTok, and even search engines with social features are not merely broadcast media; they are networks where information, influence, and behavior propagate along social ties [2,3]. Consequently, a user's decision to purchase a product is rarely an isolated event[4,5]. It is a complex function of direct advertising exposure, the observed behavior of friends ("social proof"), and conversations within their network. This phenomenon, termed network interference, poses a profound challenge for advertisers seeking to measure the true effectiveness of their marketing expenditures[6]. Traditional methods for estimating advertising effectiveness, such as randomized controlled trials (RCTs) at the user level or basic marketing mix models, rely on the Stable Unit Treatment Value Assumption (SUTVA). This ,assumption posits that the

e-ISSN: 3079-6377 p-ISSN: 3079-6369

outcome for one unit is unaffected by the treatment assignment of other units[7,8]. In the context of digital advertising, this assumption is almost certainly false. If a user is not shown an ad but their close friend is, and that friend makes a purchase, the first user may still be influenced to buy the product through word-of-mouth, even without direct ad exposure[9]. Standard attribution models would credit the friend's ad exposure for the friend's purchase but would entirely miss the peer-influenced purchase of the first user, attributing it to organic demand [10,11]. This leads to a systematic underestimation of the total "demand generation" effect of advertising and a misallocation of credit to various touchpoints[12]. The central research problem, therefore, is to move beyond simple interference detection and develop a method to estimate the magnitude of these network effects and, more importantly, to de-bias the estimate of the average treatment effect[13]. This paper introduces a hybrid causal inference framework designed to accomplish this. Our primary objective is to quantify the total causal effect of ad spend on sales, where the total effect is the sum of the direct effect (on treated users) and the indirect or spillover effect (on untreated users connected to treated users)[14]. We term this the "network-corrected demand generation effect." Our methodology combines the strengths of three distinct approaches to overcome the limitations of each when used in isolation[15]. First, we use spectral clustering on the social graph to partition users into highly connected, yet sufficiently separated, "ego-network" clusters. This design-based step helps contain interference and defines our units of analysis. Second, we employ an instrumental variable (IV) strategy to address the endogeneity inherent in ad delivery, which is optimized by platforms to target high-propensity users. We propose two instruments: the lagged rate of adopter density within a user's network and exogenous shocks to the platform's ad auction algorithm. Third, we construct a synthetic control for each treated cluster, using a weighted combination of untreated clusters that are similar in terms of pre-treatment outcomes and network characteristics. This allows us to impute a counterfactual outcome that accounts for both confounding trends and spillover effects from other treated clusters. This research makes several key contributions. It provides a practical, data-driven solution for advertisers plagued by attribution bias in social media environments. It advances the academic literature on causal inference by demonstrating a successful integration of network analysis, IV, and synthetic control methods in a unified estimation framework. Finally, our results offer actionable insights into the magnitude of social multipliers in advertising, enabling more profitable budget allocation strategies that capitalize on, rather than ignore, the interconnected nature of modern consumers.

2.Experimental Methods

Our methodological framework is designed to estimate the causal effect of advertising expenditure on sales in the presence of network interference. The approach unfolds in three primary stages: network segmentation, instrumental variable estimation of exposure, and synthetic control-based outcome modeling. The study is conducted on a simulated digital

e-ISSN: 3079-6377 p-ISSN: 3079-6369

marketplace, which we constructed to mirror the complex dynamics of real-world social platforms. The simulation comprises a connected network of 50,000 users with empirically-grounded degree distribution and community structure. User purchasing behavior is modeled as a function of baseline propensity, direct ad exposure, and peer influence. Advertising inventory is allocated via a simulated second-price auction, creating realistic endogeneity where ads are more likely to be shown to users with high historical purchase probabilities. The simulation spans a 52-week period, with weekly observations on user-level ad impressions, purchases, and social interactions. The treatment variable of interest is a binary indicator for whether a user's network cluster was exposed to a high-intensity advertising campaign for a specific consumer packaged good during the final 26 weeks of the period. To manage interference, we first partition the social graph using spectral clustering on the graph Laplacian. This method identifies communities of users who are more densely connected internally than with the rest of the network. We define each resulting cluster as our primary unit of analysis. The goal is to create clusters where the majority of peer influence is internalized, thus concentrating the spillover effects within the cluster and minimizing cross-cluster interference. We filter clusters to include those with a minimum size of 50 and a maximum size of 200 users to ensure statistical power and computational feasibility. The treatment is then randomly assigned at the cluster level, with half of the clusters receiving the high-intensity campaign. This cluster-level randomization is our core design to reduce the complexity of the interference structure. Even with cluster-level randomization, the actual ad delivery within a treated cluster is non-random. The platform's auction system ensures that ads are served to the most responsive users within that cluster. To correct for this individual-level endogeneity, we instrument for a user's actual ad exposure. Our primary instrument is the "network adoption lag." For each user, we calculate the average time it took for their first-degree connections to adopt the product in a pre-treatment period. This lag is a proxy for a user's position in the information flow and their susceptibility to peer influence, which is correlated with their later exposure to ads but is plausibly exogenous to their own contemporaneous purchasing shocks, conditional on covariates. Our second instrument leverages a design feature of the simulated platform: a periodic, randomized update to the auction algorithm's weighting of user engagement scores. These "algorithm shocks" create exogenous variation in ad delivery costs and thus exposure, independent of user demand. We use these two instruments in a two-stage least squares (2SLS) framework. In the first stage, we regress actual ad exposure on the instruments and a set of user-level covariates (e.g., pre-treatment purchase history, network degree, activity level). The predicted values from this stage represent the exogenous component of ad exposure. With instrumented exposure in hand, we move to estimate the cluster-level outcome: total sales. The challenge is that the outcome in a treated cluster is affected not only by its own instrumented exposure but also by spillover from other treated clusters. To construct a robust counterfactual, we employ the Synthetic Control Method (SCM). For each treated cluster jj, we build a synthetic version using

e-ISSN: 3079-6377 p-ISSN: 3079-6369

a convex combination of units from the untreated cluster pool. The weights for the synthetic control are chosen to minimize the pre-treatment (first 26 weeks) distance in a vector of predictors. These predictors include weekly sales, average network centrality, and average pre-treatment ad exposure rates. The key advantage of this approach is that the synthetic control implicitly accounts for unobserved confounding factors and general time trends. Furthermore, by building the synthetic control solely from untreated clusters, it provides a counterfactual that is free from direct treatment effects. However, it may still be biased if those untreated clusters are themselves subject to spillover from treated clusters. To mitigate this, we select the donor pool of untreated clusters from parts of the network that are geographically (in network terms) distant from the treated clusters, minimizing the potential for cross-cluster spillover. The estimated treatment effect for cluster j at time t is the difference between the observed outcome and the synthetic control outcome. The average treatment effect on the treated (ATT) is then aggregated over all treated clusters and post-treatment time periods, and it is regressed on the instrumented average cluster-level ad exposure to derive the final, network-corrected return on ad spend (ROAS).

3.Results

The application of our hybrid framework yielded significant insights into the magnitude of network interference and its impact on advertising effectiveness measurement. We first present the results of the first-stage instrumental variable regression, demonstrating the strength and validity of our chosen instruments. Subsequently, we present the core findings from the synthetic control analysis, contrasting our network-corrected estimates with those from a naive model that ignores interference. The first-stage regression confirmed that our instruments are strong predictors of individual-level ad exposure. The "network adoption lag" instrument exhibited a negative and statistically significant coefficient, indicating that users whose friends adopted a product more quickly in the past were more likely to be exposed to the focal ad campaign. This aligns with the platform's algorithm targeting socially central and influential users. The "algorithm shock" instrument also performed strongly, with periods following an update showing significant variation in exposure levels that were uncorrelated with baseline user attributes. The F-statistic for the excluded instruments in the first stage was 124.7, well above the conventional threshold of 10, allowing us to reject the null of weak instruments. This confirms that the 2SLS procedure successfully isolates exogenous variation in ad exposure.

Table 1 presents the core causal estimates of advertising effectiveness. It compares the results from a naive Difference-in-Differences (DiD) model, which operates at the user level and assumes no interference, with the results from our proposed Network-Corrected Synthetic Control (NC-SC) model. The naive model estimates a direct increase in sales of 2.15 units per 1,000 ad impressions, which translates to a baseline return on ad spend (ROAS). However, this model fails to account for the sales generated in untreated users who were influenced by

e-ISSN: 3079-6377 p-ISSN: 3079-6369

their treated peers. Our NC-SC model provides a more comprehensive picture. It estimates the direct effect on treated users to be 2.01 units per 1,000 impressions, slightly lower than the naive estimate after correcting for endogeneity bias. Crucially, it also quantifies the indirect or spillover effect. For every 1,000 impressions served to treated users, we observed an additional 0.89 units sold among their untreated network connections. This results in a total network-corrected effect of 2.90 units per 1,000 impressions. This represents a 34.9% increase in estimated effectiveness compared to the naive model. The standard errors for the NC-SC estimates are larger, reflecting the additional uncertainty from the synthetic control construction, but all estimates remain highly statistically significant.

Table 1 Comparison of Estimated Advertising Effects

Metric	Naive DiD Model (Units / 1k Imps)	Network-Corrected SC Model (Units / 1k Imps)
Direct Effect (on treated users)	2.15 (0.12)	2.01 (0.15)
Indirect Spillover Effect	-	0.89 (0.11)
Total Demand Generation Effect	2.15 (0.12)	2.90 (0.18)
Implied ROAS Multiplier	1.00x	1.35x

To further validate our findings and explore the heterogeneity of these effects, we segmented the results by the density of the underlying network clusters. Table 2 displays the total network-corrected effect for clusters categorized as having low, medium, and high network density (defined as the ratio of actual to potential connections within the cluster). The results reveal a clear pattern: the magnitude of the spillover effect is directly proportional to the connectedness of the network. In low-density clusters, which resemble a more atomized audience, the total effect (2.32 units) is only slightly higher than the naive estimate, suggesting minimal peer influence. In these clusters, the SUTVA assumption may be a reasonable approximation. However, in medium-density clusters, the total effect rises to 2.88 units. The most dramatic difference is observed in high-density, tightly-knit clusters. Here, the total demand generation effect soars to 4.01 units per 1,000 impressions, more than double the direct effect estimated by the naive model. In these highly social segments, the indirect, peer-driven sales actually surpass the direct sales from ad exposure, highlighting the powerful role of social contagion in amplifying advertising messages. This heterogeneity analysis confirms that network interference is not a uniform bias but a feature that varies significantly with social structure, and that ignoring it leads to the most severe undercounting of ad effectiveness in the most valuable, socially active customer segments.

e-ISSN: 3079-6377 p-ISSN: 3079-6369

Table 2 Heterogeneity of Network-Corrected Effects by Cluster Density

Cluster Density	Direct Effect (NC-SC)	Indirect Spillover (NC-SC)	Total Effect (NC-SC)	Spillover as % of Total
Low	2.01 (0.18)	0.31 (0.09)	2.32 (0.20)	13.4%
Medium	1.98 (0.16)	0.90 (0.12)	2.88 (0.21)	31.3%
High	1.95 (0.19)	2.06 (0.17)	4.01 (0.26)	51.4%

4.Discussion

The empirical results provide strong evidence for the central thesis of this paper: network interference in digital advertising is not merely a statistical nuisance to be corrected, but a fundamental component of how advertising generates demand. The finding that naive models underestimate total effectiveness by approximately 35% on average, and by over 50% in dense networks, has profound implications for both advertising practice and causal inference methodology. Our framework successfully disentangles the direct and indirect causal pathways. The direct effect estimated by our NC-SC model (2.01 units) is slightly lower than the naive DiD estimate (2.15 units). This is consistent with the expectation that the naive model is biased upward by endogeneity (targeting) but downward by its failure to attribute spillover sales. After correcting for both, we get a purer measure of the direct persuasive impact of an ad. More importantly, the identification of a sizable and statistically significant spillover effect confirms the existence of social multipliers. An impression not only influences the viewer but also turns that viewer into a potential vector of influence, seeding conversations and social proof that propagate through the network. This transforms advertising from a simple broadcast medium into a catalyst for peer-to-peer marketing. The heterogeneity analysis presented in Table 2 is particularly instructive for budget optimization. It suggests that a uniform ROAS metric is inadequate. A profit-maximizing advertiser should allocate more budget toward reaching users in high-density network clusters, even if the cost per impression in those clusters is higher. The "total demand generation" per impression in these clusters is nearly double that of low-density clusters. This insight challenges the common practice of optimizing for cost-per-click or direct conversion rate, which would undervalue these socially pivotal audiences. An optimal strategy would involve a two-pronged approach: using broad reach campaigns to seed initial adoption in high-density clusters, followed by targeted reminder ads to sustain visibility, thereby leveraging the organic peer-to-peer amplification. From a methodological standpoint, the integration of network clustering, instrumental variables, and synthetic controls proved robust. The IV approach was critical for purging the individual-level targeting bias, which, if left unaddressed, could have been confounded with peer effects. The SCM provided a flexible and powerful way to construct counterfactuals for complex social units, effectively absorbing time-varying

e-ISSN: 3079-6377 p-ISSN: 3079-6369

confounders that a standard DiD or fixed-effects model might miss. However, the approach is not without limitations. The primary assumption is that we have successfully contained interference within our defined clusters. If substantial spillover exists between clusters, our synthetic controls, built from untreated but potentially "contaminated" clusters, could be biased, leading to an underestimation of the true spillover effect. Our donor pool selection strategy mitigated this, but it cannot eliminate it entirely. Furthermore, the validity of our instruments rests on the assumption that they affect outcomes only through ad exposure (the exclusion restriction). While the algorithm shock is plausibly exogenous, the network adoption lag could theoretically be correlated with unobserved user characteristics that also predict purchasing, despite our covariate adjustments. Future research should focus on relaxing the "contained interference" assumption by developing models that explicitly parameterize cross-cluster spillovers. Hierarchical Bayesian models or network-based diffusion models could be integrated into the SCM framework. Additionally, exploring the use of graph neural networks for direct potential outcome imputation in the presence of interference presents a promising avenue for leveraging the full complexity of network data without the need for hard clustering.

5.Conclusion

This paper addressed the critical challenge of network interference in digital advertising markets, a problem that renders standard causal inference models invalid and leads to systematic undervaluation of marketing investments. We proposed and implemented a novel, integrated methodological framework that combines network clustering, instrumental variables, and synthetic control methods to estimate the total demand generation effect of advertising, which includes both direct conversions and socially mediated peer effects. Our results from a simulated yet realistic advertising market demonstrate that failing to account for network interference leads to a substantial underestimation of advertising effectiveness, with the true return on ad spend being up to 35-50% higher than naive estimates suggest. Furthermore, we showed that these spillover effects are highly heterogeneous and concentrated in densely connected network communities, offering a clear path for more sophisticated and profitable budget allocation strategies. By moving beyond simply acknowledging interference to quantifying its causal impact, this research provides advertisers with a practical toolkit for de-biasing their measurement and scientists with a robust methodology for studying causal processes in interconnected systems. The future of advertising effectiveness lies in recognizing that the audience is not a collection of isolated individuals, but a living, breathing social network, and that the most powerful campaigns are those that ignite and amplify its inherent connectivity.

e-ISSN: 3079-6377 p-ISSN: 3079-6369

References

[1] Li L, Li B, Wang G, et al. Research on precision medicine AI algorithm for neuro-immune gastrointestinal diseases based on quantum biochemistry and computational cancer genetics. Curr Pharm Biotechnol 2025;27(1):36-66.

[2] Li L, Mansor Z B, Li S, et al. Medical software engineering cost estimation model based on multimodal improved genetic algorithm[J]. International Journal of Information Technologies and Systems Approach (IJITSA), 2024, 17(1): 1-22.

[3] Li L, Qin X, Wang G, et al. Intelligence algorithm for the treatment of gastrointestinal diseases based on immune monitoring and neuroscience: A revolutionary tool for translational medicine[J]. Alexandria Engineering Journal, 2025, 113: 91-137.

[4] Li L, Chen H, Yuan Y, et al. Application of Reproductive Genetics: Genetic Algorithm Research and Information System Modeling Based on Gastrointestinal Cancer and Brain Like Engineering[C]//2023 International Conference on Computers, Information Processing and Advanced Education (CIPAE). IEEE, 2023: 396-404.

[5] Zhao X, Mansor Z, Razali R, et al. A review of optimization techniques for cost estimation in software development[J]. Intelligent Data Analysis, 2026, 30(1): 217-234.

[6] Li L, Mansor Z B, Zhao X, et al. Key Factors Affecting the Cost of Medical Software Development: Cost and Modeling[C]//2025 10th International Conference on Computer and Communication System (ICCCS). IEEE, 2025: 264-270.

[7] Runge J, Gerhardus A, Varando G, et al. Causal inference for time series[J]. Nature Reviews Earth & Environment, 2023, 4(7): 487-505.

[8] Imbens G W. Causal inference in the social sciences[J]. Annual Review of Statistics and Its Application, 2024, 11.

[9] Bailey D H, Jung A J, Beltz A M, et al. Causal inference on human behaviour[J]. Nature human behaviour, 2024, 8(8): 1448-1459.

[10] Nogueira A R, Pugnana A, Ruggieri S, et al. Methods and tools for causal discovery and causal inference[J]. Wiley interdisciplinary reviews: data mining and knowledge discovery, 2022, 12(2): e1449.

[11] Luo H, Zhuang F, Xie R, et al. A survey on causal inference for recommendation[J]. The Innovation, 2024, 5(2).

[12] Jiao L, Wang Y, Liu X, et al. Causal inference meets deep learning: A comprehensive survey[J]. Research, 2024, 7: 0467.

[13] Gao C, Zheng Y, Wang W, et al. Causal inference in recommender systems: A survey and future directions[J]. ACM Transactions on Information Systems, 2024, 42(4): 1-32.

[14] Dahabreh I J, Bibbins-Domingo K. Causal inference about the effects of interventions from observational studies in medical journals[J]. Jama, 2024, 331(21): 1845-1853.

[15] Arif S, MacNeil M A. Predictive models aren't for causal inference[J]. Ecology Letters, 2022, 25(8): 1741-1745.