Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

PRISM: Refracting Digital Influence into Ethical, Adaptive Advertising A Privacy-preserving, Real-time, Interoperable, Structure-aware and managerially governed framework for Social Network Analytics in e-commerce advertising

Linto Thomas
Christ Academy Institute for Advanced Studies
Ruchi Kumari Sharma
ruchikrisharma9525@gmail.com

Christ Academy Institute for Advanced Studies

Research Article

Keywords: social network analytics, influencer marketing, differential privacy, temporal centrality, multiplex networks, advertising governance

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11367043/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Social Network Analytics (SNA) promises influence-driven advertising, yet deployed systems still rank influencers once, on one platform, and treat privacy law as an afterthought. We propose PRISM, a five-layer framework that addresses these problems together. PRISM (i) computes influence only from noised node-level scores released under ε-differential privacy, aligned with the GDPR and India's DPDP Act 2023; (ii) keeps influence maps current through exponentially decayed temporal centrality; (iii) fuses platforms into a multiplex graph linked by a pseudonymous key rather than raw profiles; (iv) allocates seeds with awareness of community structure; and (v) embeds the analytics in a cross-functional governance loop.

We formalise PRISM, state three propositions, and test them in a reproducible simulation of a two-platform market (2,000 users, 6 periods, 15% edge rewiring per period, Independent Cascade diffusion, 5 replicates). PRISM reached 9.2% more adopters than static single-platform degree seeding on average, and 16% more by the final period. With ε = 1 it still beat both single-platform baselines and overtook every static baseline from period 5. Reach stayed within one standard deviation of the non-private level down to ε = 2. An ablation showed that community quotas slightly reduced reach, so diversity quotas are better justified by brand-safety goals than by reach. We close with managerial, regulatory and research implications.

1. Introduction

Influence in digital markets does not sit in individuals; it sits in the structure that connects them. A recommendation travels only as far as the ties that carry it, and those ties now span several platforms, shift weekly, and belong to people who have a legal right to decide how they are analysed. Advertising management has adopted SNA to exploit this structure, typically by ranking users on degree, betweenness or PageRank and seeding campaigns through the top of the list (Freeman, 1978; Page et al., 1999; Kempe et al., 2003).

That practice rests on three assumptions that rarely hold in e-commerce. It assumes the graph is one platform, when consumers split their attention across short-video, messaging and marketplace communities. It assumes the graph is static, when follower relations and conversational ties rewire continuously. And it assumes the analyst may see the raw graph, when the GDPR (European Parliament and Council of the European Union, 2016) and India's Digital Personal Data Protection (DPDP) Act 2023 (Government of India, 2023) make purpose limitation, minimisation and consent binding obligations. A fourth, organisational assumption sits underneath: that insight flows naturally from data scientists to marketers, when in practice the two groups seldom share a model of what the network means.

We argue these are not four independent problems. Privacy noise interacts with how often you can recompute; cross-platform fusion changes which communities you see; and none of it matters if the outputs cannot be governed. The field lacks a framework that designs for all of them at once and states what each component should buy in measurable reach.

Research questions.

1. How much advertising reach is lost when influence rankings are computed once, on one platform, while the underlying network drifts?
2. Can a privacy-preserving influence map retain competitive reach under budgets strict enough to satisfy data-protection regulators, and where does it break?
3. Which components of an integrated SNA pipeline contribute to reach, and which do not?

Contributions.

- A unifying framework. PRISM integrates privacy, temporal adaptation, cross-platform fusion, structure-aware seeding and governance into one pipeline with explicit interfaces between layers.
- A formal model and three propositions linking network drift, platform fusion and privacy budgets to expected campaign reach.
- A reproducible simulation study with a privacy–reach frontier and an ablation, including an honest negative result on community-quota seeding.
- A managerial and regulatory roadmap mapping each PRISM layer to DPDP/GDPR obligations and to the cross-functional roles that must own it.

The rest of the paper reviews related work (Section 2), presents PRISM (Section 3), describes the simulation (Section 4), reports results (Section 5), discusses implications (Section 6), and closes with limitations and an agenda (Section 7).

2. Literature review and research gaps

Five research streams bear on SNA-driven advertising. Each has advanced quickly since 2022, but they have advanced separately. This section reviews them in turn and then identifies the integration gap PRISM addresses.

2.1 Structural foundations of influence

SNA grew from sociology and graph theory. Granovetter (1973) showed that weak ties carry novel information across groups, and Katz and Lazarsfeld (1955) described a two-step flow in which opinion leaders relay media messages to their peers. Wasserman and Faust (1994) codified the methods. Centrality measures (Freeman, 1978; Page et al., 1999) and modularity-based community detection (Newman, 2006; Blondel et al., 2008) remain the working tools of marketing analytics. Zaefarian et al. (2022) offer a step-by-step guide to SNA constructs for marketing researchers, which shows that the methods are now mainstream in the discipline.

2.2 Influencer marketing and seeding strategy

Kempe, Kleinberg and Tardos (2003) framed influencer selection as influence maximisation under Independent Cascade and Linear Threshold diffusion. Marketing studies then tested whom to seed. Hinz et al. (2011) found that seeding well-connected hubs and bridges outperforms random seeding. Watts and Dodds (2007) and Bakshy et al. (2011) showed that ordinary users can be more cost-effective than top influencers. Iyengar, Van den Bulte and Valente (2011) found that opinion leadership and social contagion jointly drive adoption. Aral and Walker (2012) separated influence from susceptibility experimentally, and Trusov, Bucklin and Pauwels (2009) measured word-of-mouth's long-run effect relative to traditional marketing.

Recent work has moved from who is central to what makes influence effective. Leung et al. (2022) modelled influencer marketing effectiveness across influencer, content and follower factors. Joshi et al. (2025) reviewed 214 articles and found the field dominated by parasocial relationships, sponsorship disclosure, authenticity and engagement, with far less attention to network structure. Benevento et al. (2025) restore that structural lens: in Twitter product launches, influencers in central hub positions grew brand networks most, and interaction frequency mattered more than content type.

Studies published in SN Business & Economics show the same tilt toward individual-level constructs. Shetu (2026) linked influencer posts to sustainable consumption through a stimulus–organism–response model. Babu et al. (2026) traced Gen Z brand advocacy to influencer credibility, attractiveness and relatability. Nyagadza et al. (2026) reviewed 316 articles on influencer and message credibility in emerging markets. All three measure perceptions of influencers through surveys; none models the network through which influence travels. PRISM is designed to supply that complementary structural layer.

2.3 Temporal dynamics of influence

Influence changes over time. Oufaida et al. (2024) modelled influence as a time-varying graph and showed that an influencer's impact rises and falls across topics and periods. Temporal graph networks (Rossi et al., 2020) learn node representations on evolving graphs. Wang et al. (2026) combined continuous-time graph neural networks with deep reinforcement learning to select seeds in dynamic networks and reported 15–28% higher spread than static baselines on six datasets. These results support our proposition that static rankings decay. They also show a cost: deep temporal models need training data, compute and specialist staff that many advertisers lack.

2.4 Machine learning on graphs

Graph neural networks (Kipf & Welling, 2017; Hamilton et al., 2017) learn node representations for prediction. Ling et al. (2023) applied them directly to influence maximisation, learning a latent space of seed sets jointly with a neural diffusion model. These models are accurate but opaque: they rarely explain why a user is ranked highly, which matters when a marketer must justify a seed list to a brand manager or a regulator.

2.5 Privacy, regulation and advertising

Differential privacy (Dwork et al., 2006) bounds what any output reveals about one individual, and federated learning (McMahan et al., 2017) keeps raw data on devices. Rahimian, Yu and Hurtado (2023) brought differential privacy to seeding itself. Their framework protects individuals' participation in observed cascades, with central and local variants, and makes the trade-off between privacy budget and seeding quality explicit. It is the closest precedent to PRISM's privacy layer, but it operates on a single, static network.

Regulation is raising the stakes. Johnson, Shriver and Goldberg (2023) found that, in the week after GDPR enforcement, websites' use of web technology vendors fell by 15% for EU residents while vendor-market concentration rose by 17%. Frey and Presidente (2024) estimated that firms targeting EU markets saw profits fall by about 2.1% after the GDPR, with smaller firms hit hardest. In India, the Digital Personal Data Protection (DPDP) Rules were notified on 14 November 2025 (Government of India, 2025), and the core consent, notice and security obligations take effect on 14 May 2027. Advertisers therefore have a short window to build privacy into their analytics rather than retrofit it.

2.6 The integration gap

Each stream solves one piece of the problem. Table 1 compares representative approaches against the five deployment requirements named in this special issue's call; none covers all five.

Table 1
Coverage of the five deployment requirements. "Partial" means the approach addresses the requirement indirectly or only in some variants; "Possible" means the method could be extended but the cited work does not do so.

Approach (representative work)	Privacy by design	Real-time adaptation	Cross-platform	Structure-aware seeding	Governance loop
Static centrality ranking (Freeman, 1978; Hinz et al., 2011)	No	No	No	Partial	No
Greedy influence maximisation (Kempe et al., 2003)	No	No	No	Yes	No
Deep learning for seeding (Ling et al., 2023)	No	Partial (retraining)	Possible	Yes	No
Temporal seeding (Wang et al., 2026)	No	Yes	No	Yes	No
Differentially private seeding (Rahimian et al., 2023)	Yes	No	No	Yes	No
Influencer perception studies (Leung et al., 2022; Babu et al., 2026)	No	No	No	No	No
PRISM (this paper)	Yes	Yes	Yes	Yes	Yes

The gap is therefore not a missing algorithm. It is a missing architecture that lets privacy, recency, cross-platform coverage and structure work together, and makes the result governable by the people who must act on it.

3. The PRISM framework

PRISM is a pipeline of five layers, each with one job and one interface to the next (Fig. 1). The name reads as a metaphor: a prism does not create light, it separates what is already there into usable bands. PRISM likewise separates raw interaction data into privacy-safe, current, cross-platform and structure-aware signals that a manager can act on.

Data flow downward; campaign outcomes flow back into the temporal layer so the next period's influence map learns from observed adoption.

3.1 Formal setting

Let a market of N users be observed on L platforms over periods t = 1, …, T. Platform ℓ at period t yields a graph $G_\ell^t = (V, E_\ell^t)$. A campaign selects a seed set S_t of size k, and $\sigma_t(S)$ is the expected number

of adopters under Independent Cascade diffusion over the union of layers, where a user adopts if influenced on any platform. The manager's problem is to choose S_t maximising $\sigma_t(S_t)$ subject to a privacy budget ε per user per period and the constraint that no raw edge list leaves the platform.

3.2 Layer P — Privacy by design

Each platform computes a node-level influence statistic locally and releases only a noised version. Using the Laplace mechanism, the released score for user i is:

$$\tilde{s}_i = s_i + \mathrm{Lap}\left(\frac{\Delta}{\varepsilon}\right)$$

where Δ is the sensitivity of the statistic and ε the privacy budget. Smaller ε gives stronger protection and more noise. Consent state and a purpose tag travel with each record, so a user who withdraws consent under DPDP Act Section 6(4) drops out at the source rather than in a downstream audit.

3.3 Layer I — Interoperable fusion

Platforms differ in scale, density and API semantics, so raw scores are not comparable. PRISM keeps each platform as a layer of a multiplex graph, computes PageRank per layer, standardises it to z-scores, and averages across layers. Fusion needs only a pseudonymous shared user key (for example, a hashed login on the brand's own storefront), never a join of raw social profiles.

3.4 Layer R — Real-time temporal mapping

The influence map at period t blends past evidence with exponentially decaying weight λ ($0 < \lambda \leq 1$):

$$I_i^{t} = \frac{\sum_{a = 0}^{t-1} \lambda^{a}\, \bar{z}_i^{\,t-a}}{\sum_{a = 0}^{t-1} \lambda^{a}}$$

where $\bar{z}$ is the fused layer score. λ near 1 favours stability; λ near 0 favours recency. The update is incremental: each period adds one term and rescales, so the cost is one PageRank per layer per period rather than retraining a model.

3.5 Layer S — Structure-aware seeding

Communities are detected on the fused graph with the Louvain method (Blondel et al., 2008). Each user receives a bridge score, the standardised count of ties that cross community boundaries. Seeds are ranked by I_i^t plus 0.5 × the bridge score, subject to a quota per community proportional to its size. This quota-plus-bridge rule is PRISM's default configuration in all experiments; Section 5.3 tests whether it helps.

3.6 Layer M — Managerial governance

Analytics fail when no one owns the decision. Layer M assigns each decision to a role: the data scientist proposes seeds and reports expected reach with uncertainty; the marketer approves them against brand fit; the Data Protection Officer tracks cumulative ε spent per user and can veto. Every decision is logged with its inputs, which gives regulators an audit trail and gives the next campaign a labelled outcome.

3.7 Propositions

- P1 (Drift penalty). When the network rewires between periods, the reach of seeds chosen once at t = 1 declines over time relative to seeds recomputed each period.
- P2 (Fusion gain). When platforms carry partly different ties, fused multi-layer scores yield higher reach than single-platform scores of the same kind.
- P3 (Privacy frontier). Expected reach does not increase as ε shrinks, but stays close to the non-private level over a range of moderate ε, giving a practical region for compliant deployment.

4. Simulation study design

We test the propositions in a controlled synthetic market, where the true diffusion process is known and every strategy faces identical conditions. The code is short (Python, NetworkX 3.6, NumPy 2.4), seeded, and runs in under a minute; it is provided as supplementary material so reviewers can reproduce every number.

Market. 2,000 users are assigned uniformly to 8 latent communities. Two platforms are generated as stochastic block models over the same users with different structure, mirroring a dense, community-bound platform and a sparser, more cross-cutting one:

Parameter	Platform A	Platform B
Within-community tie probability	0.018	0.008
Between-community tie probability	0.0008	0.0012
Edge rewiring per period	15%	15%

Dynamics. Over 6 campaign periods, 15% of each platform's edges are rewired to random new endpoints every period, representing shifting attention and follow/unfollow churn.

Diffusion. Adoption follows Independent Cascade (Kempe et al., 2003) on the union of both layers with activation probability 0.05 per tie; a user adopts once, whichever platform reaches them. Each strategy seeds 40 users (2% of the market) per period, and reach is averaged over 60 Monte Carlo cascades per period.

Strategies compared.

Strategy	Platforms used	Recomputed each period	Privacy
Random	—	Yes	—
Degree, static	A only	No (period 1)	None
PageRank, static	A only	No (period 1)	None
PageRank, static fused	A + B	No (period 1)	None
PRISM, no privacy	A + B	Yes (λ = 0.5)	None
PRISM, ε = 1.0	A + B	Yes (λ = 0.5)	Laplace noise
PRISM, ε = 0.5	A + B	Yes (λ = 0.5)	Laplace noise

The static single-platform baselines represent the most common industry practice: rank influencers once, on the brand's main platform, and reuse the list across a campaign.

Experiments. (E1) the main comparison above; (E2) an ablation removing one PRISM layer at a time; (E3) a privacy sweep over $\varepsilon \in \{\infty, 4, 2, 1, 0.5, 0.25, 0.1\}$ on a single period. Every experiment is repeated on 5 independently generated markets; we report the mean and standard deviation across replicates.

Privacy calibration. In the simulation, noise is added to each user's fused score as Laplace noise with scale $(1/\varepsilon) \times$ the standard deviation of scores, a normalised stand-in for Δ/ε. Adding the noise once, after fusion, simplifies the per-platform release described in Section 3.2. It captures how the ranking degrades with ε but is not a full sensitivity analysis of PageRank; Section 7 returns to this.

5. Results

PRISM without privacy reached the most adopters overall, and its advantage grew every period while static rankings decayed. All figures below are from the synthetic study and should be read as evidence about mechanisms, not as forecasts for any real platform.

5.1 Main comparison (E1)

Table 2
reports mean adopters per period for each strategy, and Fig. 2 traces reach period by period.

Strategy	Mean adopters	SD	vs. static degree
PRISM (no privacy)	104.4	1.8	+ 9.2%
PageRank, static fused	99.7	1.0	+ 4.3%
PRISM (ε = 1.0)	96.9	1.5	+ 1.4%
Degree, static, platform A	95.6	2.0	—
PageRank, static, platform A	94.7	2.7	−0.9%
PRISM (ε = 0.5)	88.4	1.9	−7.5%
Random	77.8	1.3	−18.6%

Table 2. Mean adopters per period across 6 periods and 5 replicate markets.

P1 is supported. Every static strategy loses reach as the network rewires: static degree falls from 98.5 to 92.3 adopters and static fused PageRank from 105.2 to 95.8 (Fig. 2). PRISM starts slightly below fused PageRank in period 1, because decay-weighting has only one period of evidence, then rises to 107.2 by period 6, 16% above static degree. Even the private PRISM (ε = 1) overtakes every static baseline from period 5.

P2 is supported. Fusing both platforms lifts static PageRank from 94.7 to 99.7 adopters (+ 5.3%; Table 2), with no change to the algorithm, only to the data it sees.

5.2 Privacy–reach frontier (E3)

P3 is supported, with a clear knee. Reach stays within one standard deviation of the non-private level from no privacy down to ε = 2, drops about 7% at ε = 1, and reaches the random-seeding floor near ε = 0.1. E3 uses a single period and fresh markets, so its non-private level (97.8) sits below E1's six-period average (104.4), which includes the benefit of temporal updating. For practitioners this marks a usable compliance zone: ε between 2 and 4 costs almost nothing in reach, and ε = 1 costs about 7%, comparable to the 5% lost by ranking on one platform instead of two.

5.3 Ablation (E2)

Table 3
reports the ablation, in which one PRISM layer is removed at a time.

Variant	Mean adopters	SD	Change vs. full
Full PRISM	103.3	1.5	—
Without fusion (platform A only)	100.6	1.7	−2.6%
Without temporal update (period-1 scores)	100.9	1.8	−2.3%
Without community quota (plain top-k)	105.1	1.3	+ 1.7%

Table 3. Removing one layer at a time, 5 replicate markets.

The ablation uses its own five markets, so full PRISM scores 103.3 here rather than E1's 104.4. Fusion and temporal updating each contribute about 2–3% of reach. The community quota did the opposite: forcing seeds to be spread across communities lowered reach by 1.7%. A plausible mechanism, not tested directly here, is that a quota forces the model to skip higher-ranked users once their community's share is filled and pick lower-ranked users elsewhere. The plain top-k variant also dropped the bridge bonus, so the combined quota-plus-bridge step did not pay for itself here. We report this as a negative result: in a market like this one, the quota is better justified by brand-safety or fairness goals, such as avoiding over-exposure of one community, than by reach.

6. Discussion and implications

The central lesson is that an influencer list is a perishable asset. In our market, a list computed once lost 6–9% of its reach within six periods, while a list refreshed each period gained. Freshness, not algorithmic sophistication, drove most of the difference.

6.1 Theoretical implications

- Influence is a flow, not a stock. Classic centrality treats influence as a property a node holds. PRISM's results support treating it as a property of recent structure, which links SNA to the temporal-network literature and cautions against cross-sectional influencer studies.
- Privacy and performance are not a simple trade-off. The frontier in Fig. 3 has a flat region followed by a knee, not a straight line. This reframes the ethics debate in advertising from "privacy versus accuracy" to "where is the knee for this market", an empirical question each firm can answer.
- Diversity is a design choice, not a free lunch. The quota result qualifies arguments that community-spread seeding always helps. Its value depends on how influence is distributed across communities, which future work should model directly.

6.2 Managerial implications

Table 4
translates each PRISM layer into an operating decision and names the role that owns it.

PRISM layer	What a manager does on Monday	Owner
P — Privacy	Set a privacy budget between ε = 2 and 4 as a starting point; record consent and purpose per user record	Data Protection Officer
I — Interoperability	Add a second platform before buying a more complex algorithm; fusion alone added 5% reach	Data engineering
R — Real-time	Recompute influencer rankings every campaign cycle, not once per quarter	Data science
S — Structure	Use community quotas for brand-safety or fairness goals, and measure their reach cost	Marketing strategy
M — Governance	Hold a short joint review of each seed list with expected reach, privacy spend and brand fit	Cross-functional lead

Table 4. Translating PRISM into operating decisions.

6.3 Regulatory alignment

PRISM maps onto core obligations in both regimes. Local noising and purpose tags support data minimisation and purpose limitation (GDPR Article 5; DPDP Act Sections 4–6 and 8). Consent flags that travel with each record make withdrawal effective at the source. The governance log in Layer M supplies the accountability evidence a Data Fiduciary or controller must be able to show. Timing matters: India's DPDP Rules, notified on 14 November 2025, bring the core consent, notice and security duties into force on 14 May 2027, so PRISM's layers can be phased in ahead of that date. PRISM does not by itself make a firm compliant; legal review of each deployment remains necessary.

7. Limitations, future work and conclusion

Limitations. The evidence is synthetic. Real platforms have heavy-tailed degree distributions, correlated rewiring, and adoption that depends on price and product fit, none of which a stochastic block model with uniform rewiring captures. The privacy mechanism adds calibrated noise to a normalised score rather than deriving the exact sensitivity of PageRank, so the ε values are indicative of shape, not of certified guarantees. Five replicates give stable means but modest power for small differences. Finally, the governance layer is specified, not tested; its value is organisational and needs field evidence.

Future work.

1. Field validation. Run PRISM against a static-ranking control in an A/B campaign with a partner retailer, measuring conversions rather than simulated adoption.
2. Formal privacy accounting. Derive sensitivity bounds for decayed multiplex PageRank and track cumulative ε across periods with composition theorems.

3. Learned decay. Replace the fixed λ with a rate learned from campaign outcomes fed back through Layer M.
4. Federated fusion. Test PRISM's interoperability layer under federated learning so no platform shares even pseudonymous scores centrally.
5. Organisational study. Use interviews or a multi-firm survey to measure whether a shared governance loop shortens the path from insight to campaign.

Conclusion. Social network analytics has taught advertisers where influence lives; it has been slower to admit that influence moves, spans platforms, and belongs to people with rights. PRISM treats these facts as design inputs. In a controlled test it outperformed the industry-default static ranking by 9% on average and 16% by the end of a campaign, kept most of that edge under meaningful privacy budgets, and exposed where a popular heuristic, community quotas, costs reach. The broader message for advertising management is practical: refresh often, look across platforms, spend privacy deliberately, and make one team own the decision.

Declarations

Ethical Statement

This study was conducted in accordance with established academic research ethics. All procedures involving human participants adhered to institutional guidelines for responsible research conduct. The research design, data collection and analysis ensured respect for participant autonomy, privacy and confidentiality.

Declaration of Technologies used

The authors declare that AI tools were used in a limited and responsible manner solely for language refinement, grammar, and stylistic clarity. No AI tools were used for data generation, analysis, interpretation or to develop original scholarly arguments. Full intellectual responsibility for the originality, accuracy, and integrity of the manuscript rests entirely with the authors.

Conflict of Interest Statement

The authors declare that there are no conflicts of interest, financial or non-financial, that could have influenced the research, authorship, or publication of this manuscript.

Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Author Contribution

Conceptualization: Dr. Linto Thomas and Dr. Ruchi Kumari Sharma ; Data Curation: Dr. Linto Thomas and Dr. Ruchi Kumari Sharma ; Formal Analysis: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas ; Investigation: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas ; Methodology: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas; Project Administration: Dr. Linto Thomas ; Resources: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas ; Software: Dr. Linto Thomas; Supervision: Dr. Ruchi Kumari Sharma ; Validation: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas ; Visualization: Dr. Linto Thomas ; Writing – Original Draft Preparation: Dr. Linto Thomas and Dr. Ruchi Kumari Sharma ; Writing – Review & Editing: Dr. Ruchi Kumari Sharma and Dr. Linto Thomas.

Acknowledgement

The authors gratefully acknowledge the institutional support and scholarly resources that facilitated this study. They also extend their appreciation to all participants and reviewers whose valuable insights contributed to the refinement of this work.

Data Availability Statement

All results can be reproduced with the simulation script provided as Supplementary Material (Python 3, NetworkX 3.6, NumPy 2.4; fixed random seeds).

References

1. Aral S, Walker D (2012) Identifying influential and susceptible members of social networks. Science 337(6092):337–341. https://doi.org/10.1126/science.1215842
2. Babu MA, Sabuj MMI, Jaman SMS, Nila SS, Hossain MF (2026) Influencing the advocates: Exploring how social media influencers (SMIs) drive brand advocacy in Gen Z. SN Bus Econ 6:234. https://doi.org/10.1007/s43546-026-01255-w
3. Bakshy E, Hofman JM, Mason WA, Watts DJ (2011) Everyone's an influencer: Quantifying influence on Twitter. In Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (WSDM '11) (pp. 65–74). ACM
4. Benevento E, Aloini D, Roma P, Bellino D (2025) The impact of influencers on brand social network growth: Insights from new product launch events on Twitter. J Bus Res 189:115123. https://doi.org/10.1016/j.jbusres.2024.115123
5. Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. J Stat Mech: Theory Exp 2008(10):P10008
6. Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography Conference (TCC 2006), LNCS 3876 (pp. 265–284). Springer

7. European Parliament and Council of the European Union (2016) Regulation (EU) 2016/679 (General Data Protection Regulation). Official J Eur Union L 119:1–88
8. Freeman LC (1978) Centrality in social networks: Conceptual clarification. Social Networks 1(3):215–239
9. Frey CB, Presidente G (2024) Privacy regulation and firm performance: Estimating the GDPR effect globally. Econ Inq 62(3):1074–1089. https://doi.org/10.1111/ecin.13213
10. Government of India (2023) The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023). Ministry of Law and Justice
11. Government of India (2025) The Digital Personal Data Protection Rules, 2025. Ministry of Electronics and Information Technology. Notified 14 November 2025.
12. Granovetter MS (1973) The strength of weak ties. Am J Sociol 78(6):1360–1380
13. Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. In Advances in Neural Information Processing Systems 30 (NeurIPS 2017)
14. Hinz O, Skiera B, Barrot C, Becker JU (2011) Seeding strategies for viral marketing: An empirical comparison. J Mark 75(6):55–71
15. Iyengar R, Van den Bulte C, Valente TW (2011) Opinion leadership and social contagion in new product diffusion. Mark Sci 30(2):195–212
16. Johnson GA, Shriver SK, Goldberg SG (2023) Privacy and market concentration: Intended and unintended consequences of the GDPR. Manage Sci 69(10):5695–5721. https://doi.org/10.1287/mnsc.2023.4709
17. Joshi Y, Lim WM, Jagani K, Kumar S (2025) Social media influencer marketing: Foundations, trends, and ways forward. Electron Commer Res 25(2):1199–1253. https://doi.org/10.1007/s10660-023-09719-z
18. Katz E, Lazarsfeld PF (1955) Personal influence: The part played by people in the flow of mass communications. Free
19. Kempe D, Kleinberg J, Tardos É (2003) Maximizing the spread of influence through a social network. In Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '03) (pp. 137–146). ACM
20. Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In International Conference on Learning Representations (ICLR 2017)
21. Leung FF, Gu FF, Li Y, Zhang JZ, Palmatier RW (2022) Influencer marketing effectiveness. J Mark 86(6):93–115. https://doi.org/10.1177/00222429221102889
22. Ling C, Jiang J, Wang J, Thai MT, Xue L, Song J, Qiu M, Zhao L (2023) Deep graph representation learning and optimization for influence maximization. In Proceedings of the 40th International Conference on Machine Learning (ICML 2023), PMLR 202
23. McMahan B, Moore E, Ramage D, Hampson S, Agüera B (2017) Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on

Artificial Intelligence and Statistics (AISTATS 2017), PMLR 54 (pp. 1273–1282)

24. Newman MEJ (2006) Modularity and community structure in networks. Proceedings of the National Academy of Sciences, 103(23), 8577–8582
25. Nyagadza B, Bashar A, Dangaiso P, Maziriri ET, Chuchu T (2026) Sustainable consumption government policy influencer and message credibility impact on Generation Z's social media behavioural engagement in emerging markets. SN Bus Econ 6:150. https://doi.org/10.1007/s43546-026-01150-4
26. Oufaida H, Lhadj LS, Oufaida S, Grine A, Smaili N (2024) Influence lifetime: Modeling the temporal variation of social influence through domains. Social Netw Anal Min 14:174. https://doi.org/10.1007/s13278-024-01333-7
27. Page L, Brin S, Motwani R, Winograd T (1999) The PageRank citation ranking: Bringing order to the web (Technical Report 1999-66). Stanford InfoLab
28. Rahimian MA, Yu F-Y, Hurtado C (2023) Differentially private network data collection for influence maximization (Extended abstract). In Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023) (pp. 2795–2797)
29. Rossi E, Chamberlain B, Frasca F, Eynard D, Monti F, Bronstein M (2020) Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637.
30. Shetu SN (2026) Exploring the role of social media influencers' posts on sustainable consumption behavior: The moderating role of environmental knowledge. SN Bus Econ 6(1):12. https://doi.org/10.1007/s43546-025-01017-0
31. Trusov M, Bucklin RE, Pauwels K (2009) Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. J Mark 73(5):90–102
32. Wang Y, Alawad MA, Alfilh RHC, Singh SS, N (2026) Temporal influence maximization via continuous-time graph neural networks and deep reinforcement learning. Sci Rep 16:8987. https://doi.org/10.1038/s41598-026-37640-6
33. Wasserman S, Faust K (1994) Social network analysis: Methods and applications. Cambridge University Press
34. Watts DJ, Dodds PS (2007) Influentials, networks, and public opinion formation. J Consum Res 34(4):441–458
35. Zaefarian G, Misra S, Koval M, Iurkov V (2022) Editorial: Social network analysis in marketing: A step-by-step guide for researchers. Ind Mark Manage 107:A11–A24. https://doi.org/10.1016/j.indmarman.2022.10.003

Figures

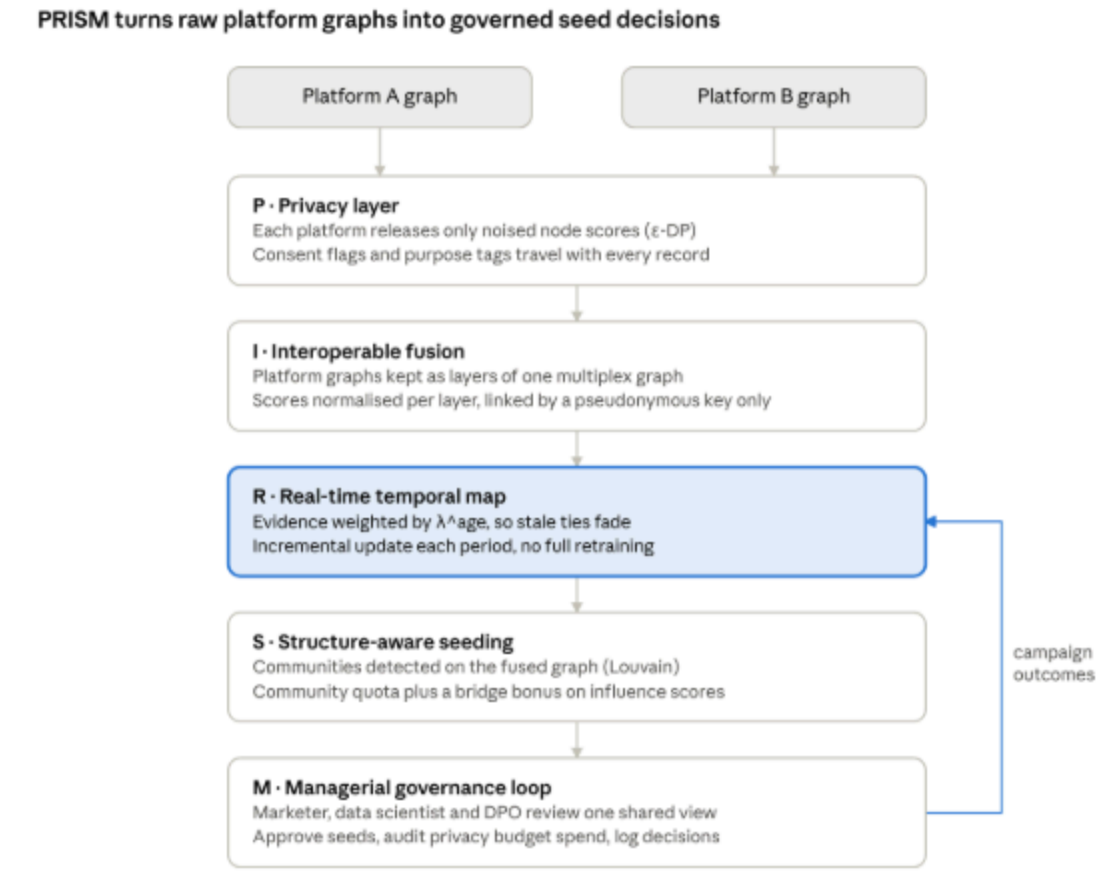
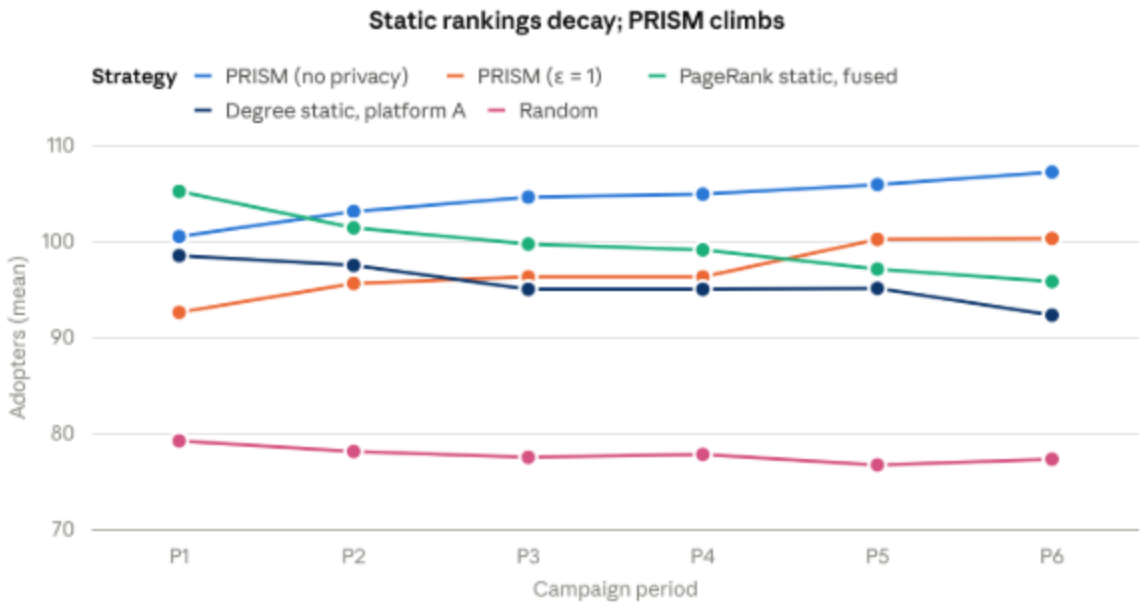


Figure 1

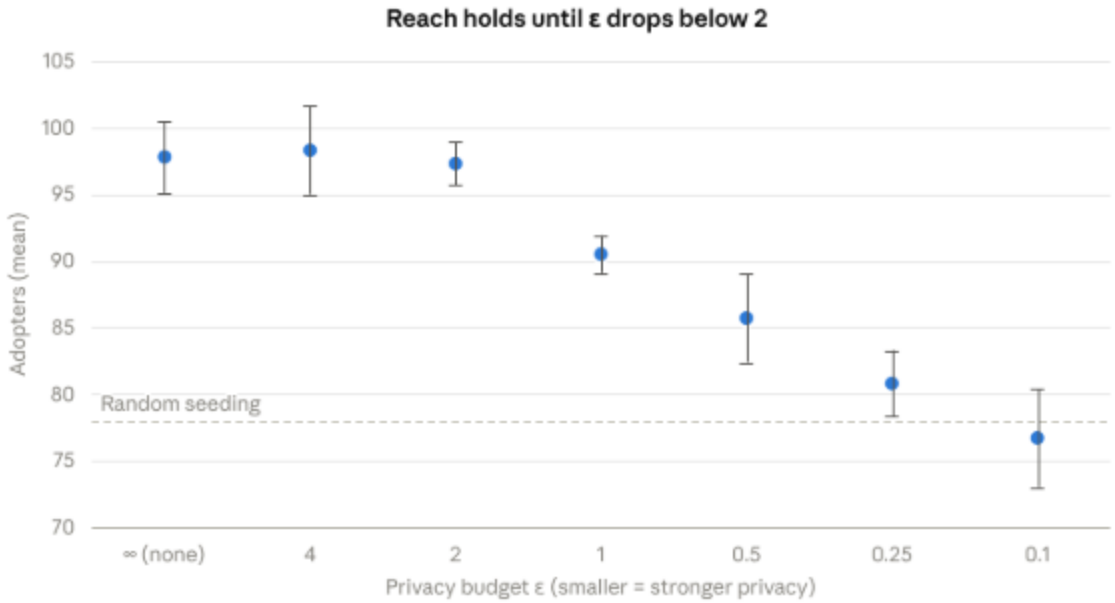
PRISM architecture · 5 layers, 1 feedback loop



Mean adopters per period from 40 seeds, 5 synthetic markets, 60 cascades each. The network rewires 15% per period.

Figure 2

simulation output (sim.py) · 6 periods × 5 strategies



Dot = mean adopters over 5 markets, whisker = ±1 SD. Dashed line = random seeding (78.0).

Figure 3

simulation output (sim.py) · 7 privacy budgets