THE PSYCHOLOGY OF INFLUENCER MARKETING

How the Subconscious Builds Trust and Trust Drives Sales

James Halder

Influencer Ink.

INFLUENCER INK.

Cover design by: Influencer Ink.
Printed in the United States of America

influencerink.com

FREE REVIEW COPY

The Psychology of Influencer Marketing

How the Subconscious Builds Trust and Trust Drives Sales

James Halder / Influencer Ink

This chapter is shared free for review purposes.

What the Full Book Covers

Beyond this opening chapter on algorithmic authority, the complete book explores the neuroscience of social proof, why parasocial relationships drive purchase intent, how micro-influencers outperform macro reach on engagement, the psychology of authenticity and disclosure, and the tactics that turn seeding campaigns into measurable revenue.

1. THE EVOLUTION OF AUTHORITY

Authority used to be held by people. Not just any people: elders whose survival meant they had seen things you had not, priests whose access to sacred texts placed them above ordinary knowledge, scholars who had spent decades inside a single subject, institutions whose longevity gave them a kind of weight. You trusted a doctor because of years of training, a judge because of legal standing, a newspaper because of the editorial process. That trust was imperfect and often abused, but at least you knew who was making the call.

Now, in most of your daily life, the calls are made by software. Not software designed by anyone you elected or could name. Software trained on your past behavior, your clicks, your pauses, your purchases, and the behavior of millions of people who look statistically similar to you. The system decides what news you see first, which products appear at the top of your search, which videos autoplay after the one you chose, and which job listings reach your screen. This is a fundamental redistribution of authority, and it happened fast enough that most people did not notice it happening.

This chapter traces that shift: where authority came from before the digital age, how it moved, what systems now hold it, and what that means for how you understand the world.

What Authority Actually Means

Before the argument can go anywhere useful, the word needs a working definition. Robert Dahl, the political scientist, defined power as the ability to get someone to do something they would not otherwise do. Authority is a specific form of that: the recognized right to make decisions that others accept as legitimate. The key word is recognized. A dictator has power but not necessarily authority. A judge has authority because the system grants it, and people accept the grant. A grandparent giving advice has authority because experience earns a kind of informal legitimacy.

Authority, in this sense, is always relational. It depends on what the person receiving the decision believes about the person or institution making it. That belief is shaped by culture, history, and evidence. When the evidence stops fitting the belief, authority erodes. When a new source of decisions starts producing outcomes that people accept, authority transfers.

Such a transfer happened recently.

Authority Before The Screen

For most of human history, authority over information was held by whoever controlled its scarcity. In oral cultures, that was the elder: the person who remembered. The elder who recalled which plant cured fever, which valley flooded in heavy rains, which alliances held, and which collapsed, held genuine power. Their authority derived from irreplaceable knowledge that could not be written down or looked up.

Religious institutions inherited and extended this model. The Catholic Church in medieval Europe didn't just provide spiritual guidance; it controlled literacy, maintained archives, translated texts, and decided which knowledge was valid. To challenge Church authority was to challenge the framework that organized all knowledge. Galileo's trial was about more than astronomy. It was about who gets to say what is true.

The printing press changed the equation in the fifteenth century. Gutenberg's press made it possible to copy text cheaply and at scale. That single development broke the Church's information monopoly within a century. Luther's theses spread because printing made distribution cheap. The Reformation was partly a media event: new technology, new authority structure.

Broadcast media in the twentieth century created a different kind of authority. Radio and television required expensive infrastructure and government licensing, which concentrated information production in a small number of hands. Three American television networks shaped what tens of millions of people understood about the world each evening. Walter Cronkite, the CBS news anchor, was repeatedly named in polls as the most trusted person in America during the 1960s and 1970s. His authority was reputational and structural: he appeared in living rooms every night, and the infrastructure required to do that credentialed him.

The model that came after the broadcast is the one you now live inside.

The Shift Online

The World Wide Web became publicly available in 1993. Its architect, Tim Berners-Lee, built it around the hyperlink: a mechanism that allowed any document to point to any other document, regardless of who owned either one. This was a deliberate political choice. Berners-Lee wanted to create a system with no central node, no single point of control. The hyperlink was a form of decentralized authority: you could follow knowledge wherever it led, and no one institution controlled the map.

For roughly a decade, that architecture held. Early web use was genuinely decentralized. Personal homepages, forums, mailing lists, and later blogs created a fragmented but pluralistic information environment. Authority online in this period was earned by reputation within specific communities. If you ran a credible technology forum or a well-sourced political blog, people returned because you were reliable, not because an algorithm sent them.

Then, in the mid-2000s, platforms arrived. Facebook launched to the general public in 2006. Twitter launched the same year. YouTube had been acquired by Google in 2006 as well. These new websites were architecturally different from what came before. Instead of a decentralized web of linked documents, they were centralized spaces where the platform itself controlled what users saw. The feed replaced the hyperlink as the primary navigation mechanism. And the feed was not neutral: it was sorted.

This is the precise moment when authority moved from creators to curators, and the curators were not human.

Tim Wu, writing in 2016, described social media platforms as technological prostheses: extensions of human cognition so integrated into daily life that removing them felt like amputation. That description captures why algorithmic authority is difficult to escape. You are not choosing to consult an algorithm the way you might choose to read a newspaper. The algorithm is the environment. It is the water you swim in.

What An Algorithm Is And What It Does

In this context, an algorithm, stripped of technical language, is a set of rules a system follows to decide what to show you, based on data about your past behavior and the behavior of users who resemble you. That is the whole thing. The sophistication varies enormously, but the logic is consistent: observe behavior, find patterns, predict what action will produce the desired outcome, serve that content.

Modern algorithms are more complex than simple rule sets. Machine learning systems identify patterns in large datasets without being explicitly programmed with rules. Deep learning systems use layered neural networks that can detect features in images, audio, and text without human-defined categories. Natural language processing systems analyze and generate human language. These are not interchangeable terms, but they share a common feature relevant to authority: none of them operate through logic you can read or audit. They produce outputs from processes that their own designers cannot fully explain after the fact.

That opacity matters. When a judge delivers a verdict, you can read the reasoning. When a doctor recommends treatment, you can ask why. When an algorithm ranks your search results or selects your news feed, there is no explanation to request. The decision is made, and the mechanism is proprietary, legally protected, and in many cases technically inexplicable even to the engineers who built it.

You are expected to accept the output without understanding the process. That is exactly what Dahl's definition of authority describes: decisions that others accept as legitimate. The difference is that traditional authority sources, at least in principle, could be held accountable. They had faces, offices, and reputations. The algorithm has none of those.

Scale And Daily Exposure

The scale of this exposure makes it qualitatively different from earlier information environments. Netflix has stated that more than 80 percent of content watched on its platform is discovered through its recommendation engine[1] rather than through user search. Amazon's recommendation algorithm is credited with driving approximately 35 percent of the company's total revenue[2]. Google processes over 8.5 billion searches per day, and the ranking decisions made in those results shape what knowledge billions of people encounter as authoritative.

You are not exposed to algorithms occasionally. On any given day, if you check a social media feed, watch a streaming service, search for information, shop online, or read a news aggregator, you have received algorithmically curated information multiple times before lunch. The systems doing this curation are not neutral. The question is: optimized for what?

The answer, in most cases, is engagement. Platforms generate revenue through advertising. Advertising revenue depends on time-on-platform. Algorithms are therefore tuned to maximize the amount of time you spend interacting with content. It's a business model. But the side effects of that optimization are not neutral.

Bias Built In

Algorithms are biased because they are trained on historical data, and historical data reflects historical inequalities. This distinction matters: the problem is structural, not individual. Fixing it requires more than good intentions.

Joy Buolamwini, a researcher at the MIT Media Lab, demonstrated this concretely. Working on facial recognition systems built by major technology companies, she found that the systems performed significantly worse on darker-skinned faces, particularly darker-skinned women[3]. In her tests, published in 2018, error rates for lighter-skinned men ran as low as 1 percent. For darker-skinned women, error rates reached 35 percent. The systems were not programmed to discriminate. They were trained on datasets that contained many more lighter-skinned male faces, so they became optimized for that demographic and failed on others.

Her documentary, *Coded Bias*, extended this finding beyond the lab. She documented how facial recognition tools were being used in real-world applications: law enforcement, housing assessments, and school monitoring. In each context, the same bias translated directly into consequential decisions. A system that misidentifies darker-skinned faces at 35 times the error rate of lighter-skinned faces is not a neutral tool when police departments use it to identify suspects.

Boulamwini's work is precise, and it exposes a general principle: the appearance of objectivity can be more dangerous than acknowledged subjectivity. A human police officer who acts on racial bias can be challenged, trained, and prosecuted. An algorithm that produces racially biased outputs carries an implicit claim of mathematical neutrality that makes it much harder to contest.

Diversity in algorithm development teams compounds this problem. When the people designing and testing these systems come from narrow demographic ranges, they are less likely to test for failure modes outside their own experience. This is a matter of perspective. If you have never experienced being misidentified or excluded by a system, you are less likely to check for that failure mode when building one.

YouTube's advertising algorithm offers a related example. Because ad placement is automated and based on contextual signals from the content being viewed, advertisers who explicitly did not want their brands associated with extremist content found their ads appearing alongside it anyway. The algorithm was optimizing for relevance and engagement without human review of each placement decision. The business consequence was a significant advertiser boycott in 2017. The social consequence was that extremist content had been, for years, partially subsidized by mainstream brands that had no idea this was happening.

Filter Bubbles And What They Cost

In 2011, the internet activist Eli Pariser introduced the term filter bubble to describe what happens when personalization algorithms remove disagreement from your information diet. The logic is simple: if you consistently engage with content from one political perspective and skip content from another, the algorithm learns to stop showing you the latter. Over time, your feed becomes a reflection of what you already believe, minus the friction of contradiction.

Pariser's main concern was about the social function of shared information. Democratic deliberation requires, at a minimum, that people have overlapping pictures of reality. When algorithms partition those pictures along engagement lines, the overlap shrinks. You and someone with different political views are not just reaching different conclusions from the same facts; you are receiving different facts.

The research on filter bubbles is genuinely contested. Some studies suggest the effect is smaller than Pariser claimed, because most people still encounter diverse sources. Other studies find that algorithmic curation measurably reduces exposure to cross-cutting political content. The debate is ongoing, but a narrower version of the claim holds regardless of its scale: recommendation algorithms are not designed to provide exposure to diverse perspectives. They are designed to maximize engagement. Whatever breadth of perspective they provide is incidental, not structural.

For you, as a reader navigating daily information, the practical implication is that the fact that something appeared in your feed is not evidence that it is representative. It appeared because the system predicted you would engage with it. That prediction is based on your past behavior, which is itself a product of the choices the algorithm made for you previously. The circularity is the problem. Your preferences and the algorithm's predictions of your preferences are not independent of each other.

Transparency And The Black Box

The opacity of algorithmic systems is not accidental. It has two sources: commercial interest and technical complexity.

On the commercial side, recommendation algorithms are core intellectual property. Facebook's News Feed ranking system, Google's PageRank and its successors, Amazon's product recommendation engine: these are proprietary assets that companies protect through trade secret law and patent protections. Publishing the full logic of these systems would allow competitors to replicate them and would allow bad actors to game them. The confidentiality has a business rationale, but the effect is that the systems making daily authority decisions about your information environment cannot be inspected.

On the technical side, modern deep learning systems are not transparent even to their creators. When a neural network produces an output, the output is the product of millions of weighted connections adjusted during training. There is no single rule you can point to that explains why the system ranked this article first and that one fifth. This is called the explainability problem, and it is an active area of research. The EU has pushed for explainable AI as a regulatory standard, but the technical capacity to explain most production systems at scale does not yet exist.

The combination of commercial opacity and technical inscrutability creates what is routinely called the black box problem. You receive a decision. You cannot see how it was made. You cannot easily contest it. And because the decision arrives embedded in a user interface designed for frictionless use, you may not register that a decision was made at all. You scrolled, the content appeared, and you engaged. The selection process was invisible.

This is authority without accountability. It is decisions made at scale, affecting billions of people, by systems whose logic is unavailable for public inspection.

Regulation And Its Limits

Governments have begun to respond, though the responses are uneven and generally behind the pace of deployment.

The European Union's AI Act, finalized in 2024, is the most comprehensive attempt to regulate algorithmic systems by risk category. High-risk applications, including AI used in employment decisions, credit scoring, education assessment, law enforcement, and critical infrastructure, face mandatory transparency requirements, human oversight obligations, and pre-market conformity assessments. The Act prohibits certain practices outright, including social scoring by governments and most uses of real-time facial recognition in public spaces. It is a serious piece of legislation, and it sets a global benchmark by virtue of the EU's market size.

In the United States, the Algorithmic Accountability Act has been proposed in Congress in various forms since 2019. As of this writing, it has not passed. The proposal would require companies above a certain size to audit automated decision systems for accuracy, fairness, bias, and privacy implications. The Federal Trade Commission has taken enforcement actions against companies using algorithmic systems in deceptive ways, but the US lacks a comprehensive federal framework equivalent to the EU Act.

The UK's Information Commissioner's Office has issued guidance on automated decision-making under data protection law, and sector-specific regulators in finance and healthcare have developed their own standards. Australia, Canada, and Brazil have active legislative processes. The global picture is one of accelerating regulatory attention but significant variation in speed and ambition.

Regulation faces a fundamental challenge beyond political will: the systems being regulated change faster than legislative processes can track. A law drafted to address the algorithmic systems of 2022 may be poorly fitted to the systems of 2027. Regulatory frameworks that focus on specific techniques rather than outcomes tend to become obsolete. The EU Act's risk-based approach, which focuses on application domains rather than specific algorithms, is partly designed to address this problem, but it remains untested at scale.

There is also a jurisdictional problem. The companies deploying the most influential algorithmic systems are headquartered in the United States, though they operate globally. A European user of an American platform is subject to EU rules, but enforcement depends on cooperation that is not always forthcoming. The global reach of these systems does not match any single government's regulatory authority.

What Algorithmic Literacy Actually Means

Much of the public debate about algorithms ends with a call for algorithmic literacy, which is a useful concept that is rarely given specific content. Literacy does not mean learning to code or understanding neural network architecture. It means developing a practical, accurate model of how these systems work and what they are optimizing for.

At the most basic level, algorithmic literacy means understanding that your feed is a selection from the world, made by a system with specific optimization targets that do not include your long-term health. That knowledge changes how you read what you see. An article that appears at the top of your search results is there because Google's ranking system assessed it as the best match for its model of your intent, filtered through its model of page quality, influenced by the link structures of the broader web.

At a practical level, algorithmic literacy means deliberately introducing friction into your information consumption. This does not require dramatic gestures. It means occasionally searching for the same topic using different platforms and comparing results. It means following sources that routinely challenge your existing views, not to be contrarian but to maintain calibration. It means treating the personalized nature of your information environment as a standing condition to be managed, not a neutral background.

Tools exist to support this. Browser extensions that reveal the factors driving search rankings, media literacy curricula that train students to identify the difference between editorial and algorithmic curation, and journalistic organizations that publish their editorial standards explicitly. None of these solves the structural problem. But they shift you from passive recipient to active reader, which changes your relationship to the authority these systems claim.

The larger point is institutional as much as individual. Algorithmic literacy at a societal level requires that schools teach the basics of how recommendation systems work alongside the basics of how media organizations work. It requires that regulators have the technical staff to understand the systems they are overseeing. It requires that journalists develop enough fluency in algorithmic systems to cover them accurately rather than either uncritically or hysterically.

The Accountability Gap

The deepest problem is the accountability gap between the authority these systems exercise and the mechanisms that exist to contest or correct them.

When a newspaper publishes a false story, there are mechanisms: corrections, retractions, press council complaints, defamation law, and reputational consequences. When an algorithm systematically suppresses certain voices, demotes certain content, or amplifies certain political tendencies, the mechanisms are far weaker. You can appeal a specific content moderation decision on most platforms. You cannot appeal a ranking decision. You cannot request an explanation for why your content reached fewer people this month than last. The decision is made, it has effects, and there is no address to write to.

This gap is what makes the shift from human to algorithmic authority meaningful rather than merely technical. The traditional sources of authority that algorithms have displaced were imperfect, biased, exclusionary, and frequently corrupt. But they existed within accountability structures that, over time and through struggle, could be reformed. The press could be pressured into better practices. Courts could be challenged and overturned. Institutions could be legislated into different behaviours. The accountability structures for algorithmic authority are nascent, underfunded, and outpaced.

This does not mean algorithmic authority is inherently illegitimate or that the previous order was preferable. Algorithmic systems also surface knowledge that would never have reached you through traditional gatekeepers, connect communities that geography would have kept apart, and process information at scales that genuinely exceed human capacity. The claim is narrower: authority without accountability is dangerous regardless of its form, and the current pace of algorithmic deployment has outrun the accountability frameworks that would make it safe.

The question worth sitting with is not whether algorithms should influence your information environment. They already do, at a scale and depth that is not reversible in any practical sense. The question is what accountability structures are appropriate for systems that now perform functions once performed by editors, librarians, teachers, and institutions.

Engineering, 112, 42–46. https://doi.org/10.54254/2755-2721/2024.17905

[2] MacKenzie, I., Meyer, C., & Noble, S. (2013, October 1). How retailers can keep up with consumers. McKinsey & Company. https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers#/

[3] Buolamwini, J. (n.d.). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification – MIT Media Lab. MIT Media Lab. https://www.media.mit.edu/publications/gender-shades-intersectional-accuracy-disparities-in-commercial-gender-classification/

[4] Meshi, D., Morawetz, C., & Heekeren, H. R. (2013). Nucleus accumbens response to gains in reputation for the self relative to gains for others predicts social media use. Frontiers in Human Neuroscience, 7, 439. https://doi.org/10.3389/fnhum.2013.00439

[5] Montag, C., & Hegelich, S. (2020). Understanding detrimental aspects of social media use: Will the real culprits please stand up? ResearchGate. https://doi.org/10.3389/fsoc.2020.599270

[6] Djafarova, E., & Rushworth, C. (2016). Exploring the credibility of online celebrities' Instagram profiles in influencing the purchase decisions of young female users. Computers in Human Behavior, 68, 1–7. https://doi.org/10.1016/j.chb.2016.11.009

[7] Berryman, R., & Kavka, M. (2017). 'I Guess A Lot of People See Me as a Big Sister or a Friend': the role of intimacy in the celebrification of beauty vloggers. Journal of Gender Studies, 26(3), 307–320. https://doi.org/10.1080/09589236.2017.1288611

[8] Salla, K. T. (2025). Little fish, big pond: Are micro influencers worth the catch? Analyzing the difference and effectiveness between micro and macro influencers. Undergraduate Review, 19(1), 291–304. https://vc.bridgew.edu/undergrad_rev/vol19/iss1/25

[9] Wilt, J. A., Thomas, S., & McAdams, D. P. (2019). Authenticity and inauthenticity in narrative identity. Heliyon, 5(7), e02178. https://doi.org/10.1016/j.heliyon.2019.e02178

[10] Cassidy, P. (2017, November 21). Survey finds consumers crave authenticity - and User-Generated Content delivers. Social Media Today. https://www.socialmediatoday.com/news/survey-finds-consumers-crave-authenticity-and-user-generated-content-deli/511360/

[11] Geyser, W. (2024). The state of influencer marketing 2024: Benchmark report. Influencer Marketing Hub. https://influencermarketinghub.com/ebooks/Influencer_Marketing_Benchmark_Report_2024.pdf

[12] 2022 Bad Bot Report | Evasive Bots drive online fraud | Imperva. (2025, February 10). Resource Library. https://www.imperva.com/resources/resource-library/reports/bad-bot-report/

[13] Apostol, A., & Asiminei, R. (2025). De-Influencing as a means of preventing overconsumption. Encyclopedia, 5(4), 202. https://doi.org/10.3390/encyclopedia5040202

[14] Rogers, R. (2026, January 14). TikTok Shop showed me search suggestions for products with nazi symbolism. WIRED. https://www.wired.com/story/tiktok-shop-nazi-symbolism/

[15] Naceva, N., & Naceva, N. (2025, August 1). Elasticity Modeling: How CPM changes with Creator Tier. Influencer Marketing Hub. https://influencermarketinghub.com/indicatesinfluencer-cpm-elasticity/

[16] Moore, C. (2026, January 19). The seeding strategy that's replacing paid ads. Branding With Benefits. https://artofthebrand.substack.com/p/the-seeding-strategy-thats-replacing

About Influencer Ink

Influencer Ink is an independent publishing house dedicated to the business, culture, and psychology of influence. We publish nonfiction for marketers, brand strategists, creators, and the curious minds trying to understand how modern persuasion really works.
Every title we release goes deeper than the trend cycle — because the most valuable insights are the ones that hold up.

Our Titles

Browse our full catalog at influencerink.com.

Connect With Us

Website: influencerink.com
Submissions & Rights: submissions@influencerink.com
Media & Publicity: press@influencerink.com
General Inquiries: hello@influencerink.com

For Booksellers & Librarians

Influencer Ink titles are available through major distributors. For bulk orders, educator pricing, and institutional licensing, visit influencerink.com or contact your regional sales representative.

Printed and distributed in partnership with leading printing and fulfillment partners worldwide.