

The Author as Customer

Vanity Publishing, Advertising Platforms, and the Sale of Entry into Imagined Opportunity

Flyxion

Independent Researcher

October 21, 2026

Abstract

This essay argues that vanity publishing is not best understood as a quality problem but as an economic inversion, and that the same inversion appears wherever a platform is paid for access to an attempt and not for its outcome. In conventional publishing, revenue flows from readers to publishers, which aligns the publisher's financial interest with the public circulation of work. In vanity publishing, the author replaces the reader as the primary revenue source, decoupling editorial selection from audience demand and transferring risk from the institution to the creator. The essay develops a structural account of this inversion, distinguishes it from legitimate self-publishing and service provision, and situates it within a broader platform pattern in which participation itself becomes the billable surface. It then applies the account to advertising platforms. These can earn revenue by distributing promises of economic advancement, including artificial-intelligence-based income opportunities, under payment terms that are not contingent on whether purchasers ever benefit, which sets the business of selling advertising against the function of discovery. Indirect constraints such as refunds, litigation and reputation may weaken the conflict, and how far they do is an empirical question. The claim concerns incentive structure. What a platform knew about a particular advertisement is a separate question of fact, and the essay sets out the evidence that bears on it and what that evidence cannot show. It proposes that clearly separated advertising can coexist with discovery only if ranking, verification and enforcement are insulated from the payer, and it examines the disposable aspirational identities that such an economy leaves behind. It also considers whether generative AI services add a further possibility, that feedback from paying users becomes useful training material, and states what documentation would be needed to say so.

1. Introduction

Vanity publishing is usually criticized as a marginal embarrassment within the literary economy: a domain of weak editorial standards, inflated promises, and authors misled by the symbolic appeal of seeing their names in print. This criticism is not wrong, but it is incomplete. It treats vanity publishing as a degraded version of publishing rather than as a distinct economic structure. The central issue is not merely that some books are poorly produced or that some publishers exaggerate what they can deliver. The deeper issue is that the economic relation between publisher, author, and reader has been reversed.

This essay argues that vanity publishing should be understood as an inversion of the press. In conventional publishing, the publisher assumes risk by selecting, editing, producing, and distributing work in the hope that readers will buy it. In vanity publishing, the author becomes the primary source of revenue. The publisher is compensated before readership is established and often independently of whether readership ever materializes. Publication itself, rather than public circulation, becomes the product.

The significance of this inversion extends beyond publishing. Vanity publishing belongs to a broader class of platform systems that sell access to participation under conditions where success is rare. These systems do not primarily sell outcomes. They sell entry into environments where outcomes might, in principle, occur. The participant pays to try, while the platform captures value from the attempt whether or not the attempt succeeds.

The aim is not to defend traditional publishing as pure or democratic. Traditional publishing has its own exclusions, inequities, and failures. The point is more specific: when an institution claims the symbolic authority of publishing while shifting cost and risk onto creators, it changes the meaning of publication itself. What appears as expanded access may instead become a system for extracting value from the desire to participate in culture.

The second part of the essay applies this structure to advertising-funded platforms. Its thesis is that advertising platforms can earn revenue by distributing promises of economic advancement under terms that do not make payment depend on whether the people who accept them ever benefit, and that this sets the business of selling advertising against the function of discovery. The relevant offers are those that sell entry into an imagined opportunity: tools, courses and subscriptions that promise income or productivity, including through artificial

intelligence. In its most circular form, buyers purchase instruction in earning money with artificial intelligence and then sell products or instruction of the same kind to the next group of hopeful buyers, so that much of the income in the chain may be the fee collected at the entrance.

The claim concerns incentive structure. It does not assert that any platform knowingly accepted any particular harmful advertisement, because knowledge is a separate matter that depends on evidence about specific advertisements and the enforcement history attached to them. Keeping the two claims apart is the discipline of the second part, since an argument from incentives, however sound, cannot by itself establish what an institution knew.

2. The Inversion of the Press

Criticism of vanity publishing is often dismissed as snobbery, and the objections most commonly voiced concern quality: the books are poorly edited, the prose is undistinguished, the covers derivative. These observations may be accurate, but they identify symptoms and not mechanisms. A publisher can produce a handsome, well-written book and still operate an exploitative structure, and a publisher can produce work of modest merit while performing a genuine service. Quality is not the diagnostic. The real problem concerns who pays, who bears risk, and in whose interest selection decisions are made.

In the conventional model the publisher combines the roles of investor, curator and distributor [25]. It receives manuscripts, declines most of them, and concentrates resources on those it selects, investing in editing, design, production and distribution and recovering that investment through sales. Revenue flows from readers, through booksellers, to publishers, and if the book fails to find an audience the publisher bears the loss. The model guarantees neither good judgment nor equitable treatment of authors, and consolidation, gatekeeping, genre conservatism and inequitable contracts are documented features of it. Its one structurally important property is that the publisher's revenue depends at least partly on reader demand, which creates a weak but real alignment between institutional incentive and public value, and which makes selection an economic necessity and not a cruelty. Because the publisher invests its own capital, it cannot accept everything.

In vanity publishing the relation is reversed. The publisher's revenue comes from selling publication to authors, who pay for packages that may include edi-

torial services, cover design, formatting, distribution listings and marketing support, and the publisher is compensated whether or not the book finds an audience. The consequences cascade. Rejection means declining a paying customer, so the editorial selection function weakens or disappears. The reader ceases to act as an external discipline, since the question that governs acceptance is whether the author wants to publish and not whether anyone will want to read. The worse form of the model needs hopeful authors more than successful books, and its logic points toward volume. The visible result is the outward apparatus of publishing, with ISBNs, covers, retail listings and press-style language, produced by a service workflow in which nobody has wagered capital on the work's reception.

Risk is not eliminated by this arrangement but relocated. The author pays upfront, sometimes substantially, for services that may or may not result in sales, and bears the loss if the book finds no readers. This is not inherently illegitimate when it is transparent and the author understands that they are purchasing services and not gaining an institutional partner. It becomes a problem when risk is supposed to travel with curatorial authority and does not. A publisher that claims the authority to select, endorse and legitimize a work implicitly accepts responsibility for that judgment, and when the same authority is claimed while risk is transferred entirely to the author, the publisher speaks the language of selective investment while operating the economics of a service vendor.

3. Symbolic Claims Without Symbolic Substance

Three distinctions locate the problem more exactly. The first separates vanity publishing from self-publishing, which in its defensible forms is transparent about what it is. The author knows that they are purchasing discrete services from editors, designers and printers and distributing through platforms built for direct publication, and they are not under the impression that a press has selected them from a field of candidates. The second separates vanity publishing from honest service provision. Editors charge for editing, designers for covers and printers for printing, and author-facing payment is not the problem. A service provider sells a service and does not claim to be investing in the work, selecting it or conferring legitimacy on it. The vanity model becomes distinct when it adopts the symbolic apparatus of institutional publishing, speaking of editorial review and implying selection and endorsement, while operating on the author-as-customer logic. The problem is mimicry, and the ambiguity it

creates is often the product. Authors pay more when they believe they are receiving institutional endorsement than when they believe they are buying a service package, and if the process were described honestly as service purchasing, many would decline. The argument is not about legality and is not reducible to fraud in any simple sense, since many of the claims are ambiguous enough to survive legal scrutiny. The issue is the gap between the symbolic claims and the economic reality, and the degree to which the gap is productive for the seller.

The third distinction concerns what is actually sold. Publishing carries symbolic weight beyond its economic function, since to be published in the traditional sense is to have been selected and to have crossed a threshold that not everyone can cross. Vanity publishing encloses that symbolic dimension and resells access to it. The ISBN, the distributor listing, the press-release language and the cover resembling those of established houses are the tokens of having been published, and the author pays for the experience of arriving on the other side of a gate that resembles the gates of traditional publishing. This is why the model can be emotionally powerful when it is commercially negligible: the result in sales and readers may be close to zero while the symbolic result, the ability to say that one has published a book, feels complete. The platform has delivered the experience of recognition without its substance.

Vanity publishers often defend themselves by appeal to access. Traditional publishing is slow, opaque and exclusionary, favouring established authors, agents and networks, and the critique of it contained in this defence is substantially correct. But democratization and extraction are not the same thing. A democratic model would give creators more control, more transparency about economic arrangements, more access to readers and a more equitable distribution of revenue. Direct-to-reader tools, open-access infrastructure, cooperative presses and platforms that charge transparent fees for discrete services without claiming curatorial authority all expand access by reducing the cost of reaching readers. The vanity model expands access to the symbol of endorsement and charges for the symbol.

4. Participation as Product

The conceptual structure beneath this extends well beyond the publishing industry. The author is not simply buying a service but access to a process in which their own aspiration is the raw material, and the platform sells the creator's desire to be published back to them in packaged form. The creator is simul-

taneously producer and customer, generating the work, supplying the energy and bringing the aspiration, and then paying to have the aspiration formally recognized. Publication ceases to be a bridge between writer and audience and becomes a passage through institutional-looking gates whose function is to complete a transaction. Institutions have always been capable of selling proximity to prestige in place of prestige itself, and what is distinctive here is the degree to which a model built on author payment imitates the symbolic language of one built on audience value.

Shared infrastructure sharpens the asymmetry. Once a platform has templates, editorial pipelines, design workflows and distribution agreements, the marginal cost of processing another manuscript is low, and the cost of the fixed infrastructure is amortized across the whole author pool. Each author may be charged as though receiving a bespoke arrangement while receiving a templated one. Operating at scale is not dishonest in itself, and all large publishing operations standardize. The issue is representation. The author's belief that the manuscript is receiving special institutional attention is part of what is being paid for, and a platform that presents individualized investment while running a high-volume workflow lets the gap between description and reality do important work in the transaction.

5. Chokepoints and Compatibility with Failure

Vanity publishing is one instance of a pattern that appears wherever a platform controls a chokepoint, a position of mediated access between participants and the outcomes they seek [6, 24]. Publishing platforms control access to the apparatus of authorship, professional networks to visibility and opportunity, application marketplaces to distribution, and cloud infrastructure to the tools needed to operate in a digital economy. The position acquires its full power under oversupply. The number of aspiring authors vastly exceeds the number of commercially successful books, the number of applications uploaded exceeds the number meaningfully downloaded, and the number of people attempting to build businesses or audiences through platform channels exceeds the number who will achieve meaningful returns. In such environments failure is not an anomaly but the statistical baseline.

A particular economic logic then becomes available. Payment can be collected at the point of entry, before outcomes are determined, so that the supplier is not dependent at that point on any given participant succeeding. Authors pay

to be published regardless of whether the books are read, developers invest in applications that may never be installed, and professionals buy tools and subscriptions in pursuit of opportunities that may never materialize. The systems do not sell success. They sell access to the possibility of success. Aspiration is the input, and the costs of failure are borne individually while the benefits of scale are captured by the supplier, since infrastructure is shared while pricing is individualized [7].

A qualification is needed here to avoid attributing too much to the structure. Charging for access does not require that most customers fail, and widespread success would not force the supplier to become an investor or partner. A useful software service can charge subscriptions while most customers benefit. What the structure establishes is narrower: because payment is collected at entry, the supplier can remain profitable even where most participants do not obtain what they sought, provided that acquisition and retention sustain revenue. The structure is compatible with failure as the baseline and does not depend on it, and the worry developed in the later sections concerns the cases where failure is in fact the baseline and the supplier's revenue is largely indifferent to it. Whether a given market is such a case is an empirical question about success rates, refunds and retention.

Where failure is the baseline, success stories function as proof that the system works. They justify continued participation by showing that the outcome, while unlikely, is possible, and the many unsuccessful attempts remain largely invisible or are reframed as part of the process of striving. The platform need not deny the difficulty of success. It can incorporate the difficulty into its narrative of legitimacy, so that the failure rate is absorbed into the story and not disclosed as a feature of the business model. Participants are not necessarily misled about the existence of risk; they are invited to accept it as the price of access, and the asymmetry lies in who bears it. The participant profits only by succeeding, whereas the supplier has already been paid.

A complementary dynamic can appear at the other end of the distribution. The few participants who achieve traction often do so within or adjacent to the systems that mediated their participation, and such success is frequently absorbed through acquisition, integration or talent capture. The field of participation can then function as a distributed search process, generating many low-probability attempts from which a small number of viable outcomes emerge, and the platform may capture those outcomes as well. The cost of trying is distributed and

the reward of succeeding can be centralized.

6. Persistence and Performance

If failure is the baseline condition of participation, a further question arises: why does participation persist? Why do individuals continue to invest time, money, and attention in systems where the probability of meaningful success remains low?

Part of the answer lies in the construction of necessity. Platforms do not present themselves merely as optional tools, but as components of what it means to participate seriously in a domain. To publish is to exist as an author within recognized channels. To maintain a professional presence is to be visible within networked systems. To build a business is to operate through accepted infrastructures. These claims are not entirely false, but they are generalized far beyond the subset of participants for whom they are materially justified. Tools that are indispensable at scale are presented as prerequisites at all scales.

The result is a condition of latent participation. Individuals remain enrolled in systems not because those systems are producing commensurate outcomes, but because disengagement appears to foreclose possibility. Accounts are maintained, subscriptions renewed, profiles updated, and projects sustained in a state of partial activation. The participant does not exit the system upon failing to achieve success. They remain aligned with the possibility of it. The platform does not need to deliver results. It needs to remain legible as a pathway through which results might, in principle, be obtained.

This persistence is reinforced by the introduction of metrics that render participation visible and comparable. Activity is quantified: tokens consumed, interactions performed, content produced, visibility maintained. These measures begin as proxies for productive engagement, but once established, they invite direct optimization. Participants do not merely pursue outcomes; they pursue the signals that stand in for them. The system provides a means of acting and a means of displaying that action, and the latter can come to dominate the former.

Under these conditions, the distinction between work and its representation becomes unstable. The engineer who maximizes interaction with a system appears productive. The author who achieves publication appears legitimate. The professional who maintains constant visibility appears engaged. These appearances are not meaningless, but they are only indirectly related to the outcomes

they imply. The metric measures participation, and participation becomes the target.

This dynamic aligns closely with the economic structure described in earlier sections. Platforms benefit when participation is both persistent and measurable. Persistent participation ensures a continuous flow of engagement, while measurable participation provides the basis for comparison, evaluation, and, where applicable, pricing. The system does not need to ensure that participants succeed. It needs to ensure that their activity remains legible, optimizable, and ongoing.

In this sense, the system does more than extract value from attempts under conditions of failure. It stabilizes those attempts by embedding them within structures that encourage continuation and reward display. Participants are not only drawn into the system by aspiration; they are held within it by necessity and guided within it by metrics. The result is a self-reinforcing loop in which the pursuit of success is sustained, measured, and converted into value regardless of its outcome.

A further extension of this dynamic can be observed in the integration of agent-based systems into platform infrastructures. These systems promise to move beyond discrete tools toward continuous, action-oriented workflows: coordinating communication, managing tasks, and executing operations across domains. At the level of user experience, this appears as an expansion of capability. Structurally, it is an expansion of the platform boundary. Activities that might previously have occurred across disparate tools and contexts are drawn into a single environment, increasing both dependence on and exposure to the system.

This expansion is often tied to subscription tiers that differentiate between basic access and so-called power usage. The participant is encouraged to see themselves not merely as a user, but as someone operating at a higher level of integration, leveraging the system across all aspects of their work. What remains unchanged is the absence of any corresponding mechanism for evaluating whether the actions being taken are viable. The system facilitates execution, but does not impose constraints derived from physical, logistical, or financial reality. It expands what can be done within the platform without expanding what can be successfully done in the world. Agent-based systems do not depart from the structural logic described above. They extend it. The surface of interaction expands from discrete inputs to continuous workflows, increasing the volume and persistence of activity that can be measured, optimized, and monetized. The

participant experiences this as increased capability. The platform experiences it as increased engagement within a controlled environment. The underlying asymmetry remains intact.

At its limit, this structure does not merely extract value from participation. It reorganizes who performs the work that makes the system itself more capable.

7. Feedback from Paying Users as Possible Training Material

Generative platforms raise a further question about who performs the work that improves the system. In conventional software development the cost of improvement flows outward from the institution: engineers are employed, annotators are contracted and testers are paid. Annotation pipelines for large language models, including preference ranking, output evaluation and the identification of edge cases, require large volumes of human judgment applied to system outputs, and the people who perform it have been documented as often poorly paid [18, 23]. The possibility examined here is that a service sold by subscription could obtain some of that judgment from its paying users, in the form of reprompts that reveal failure modes, corrections that carry preference information, and iterative refinements that resemble comparison data, so that customers would be supplying material that conventional institutions purchase.

The possibility should be separated from what is established. User interactions can furnish potential training or evaluation material, but whether they do so depends on the service, the account type, the settings chosen, the contractual terms, the procedures by which interactions are selected and the training that follows. A correction is not automatically a preference comparison of the kind used in reinforcement learning from human feedback, and most interactions are likely to be unusable without substantial processing. Nor does the use of such material show that the arrangement is exploitative, since users receive a working service in exchange and may benefit from improvements that result. The question is whether an additional appropriation of value occurs and on what terms, which is an empirical and contractual question.

Public descriptions of two consumer services show that practices differ and change. According to press and consumer-guide accounts, OpenAI's consumer ChatGPT service has treated conversations as eligible to improve its models unless the user turns off a data-control setting [26], and Anthropic announced in 2025 that consumer Claude plans (Free, Pro and Max) would ask users to choose

whether chats and coding sessions may be used for training, with retention of up to five years for those who agree and thirty days for those who decline, while commercial, government and education offerings kept existing protections [1]. These accounts are secondary, dates and defaults have changed over time, and the services' own terms should be consulted before any claim about a particular period. They do show that the presence of a paid tier does not settle whether interactions are used, and that user choice and plan type matter.

The argument that survives is therefore conditional. If a service uses interactions of paying users as training or evaluation material, and if that use is not disclosed clearly or not priced into the arrangement, then the customer and the contributor have been merged without the customer being treated as a contributor, and the symbolic substitution and the labour substitution described for vanity publishing coincide in a single transaction. If the use is disclosed, optional and compensated through a lower price or better service, the arrangement is closer to an ordinary exchange. Establishing which description fits a given service requires documentation of its terms, settings and practices, and this paper does not supply it beyond the examples above.

8. Symptom and Structure

When a specific platform is criticized as exemplifying vanity publishing, the criticism often focuses on surface features: the marketing language is hyperbolic, the books appear unremarkable, the promotional claims seem disproportionate to commercial results. These observations may be accurate, but they miss the structural point. The ridiculousness, if that is the right word, is not accidental. It is the visible symptom of an economic model that has detached publication from readership.

When publication is no longer constrained by the need to find an audience, the language surrounding it need no longer be constrained by the need to describe something real. The press releases can become untethered from market performance because market performance is not what the platform is measuring. The claims about reach, visibility, and impact can be expansive because no one's revenue depends on them being accurate. The symbolic apparatus can be elaborate precisely because the economic substance beneath it is thin.

A structurally grounded critique of such a platform would not primarily note that it seems ridiculous. It would ask the diagnostic questions: where does

the platform's revenue come from? How is editorial selection conducted, and what are the consequences of rejection? What does distribution actually deliver? What claims does the platform make about its own institutional function, and how do those claims relate to its economic operation? If the answers reveal a system in which authors are the primary payers, editorial rejection is rare or absent, distribution is limited, and institutional language exceeds institutional reality, then what is being described is the inversion this essay has outlined.

9. Payment Without Contingent Outcome

The structural account developed so far describes a relation among a seller, a participant and an outcome. The remaining sections ask what follows when the seller is an advertiser and an advertising platform is paid for bringing the promise to a receptive audience. The first step is to state the payment structure precisely, since the statement shows how far it extends beyond publishing and how far it does not. Let a promise m be offered to a population of prospective buyers, and let Y denote the outcome a buyer realises after acting on it, such as income earned or a skill acquired.

Definition 1 (Outcome-Contingent and Non-Contingent Payment). *A payment for distributing or purchasing m is non-contingent on outcome if its amount is determined by exposure, click, purchase or subscription events and no term of the transaction makes it depend on Y .*

The definition applies to the initial transaction and makes no claim about a party's revenue over time. In vanity publishing the buyer is the author, the promise is publication, and the seller is the press. In the income-opportunity market the buyer is the aspirant, the promise is entry into a business or a source of income, and the seller is the advertiser, while the advertising platform is paid per impression or per click for the exposure it supplies. The platform's payment is non-contingent on outcome with respect to every promise it carries, and the advertiser's entry fee is collected before Y is observed. The buyer bears the most direct dependence on Y . In the circular case, buyers become sellers, and the entry fees of each tier of sellers are again collected from entrants below.

Remark 1. *Non-contingency at the point of payment is not independence over time. Refunds and chargebacks, subscription retention, repeat purchase, reputation, litigation and regulatory action, and platform enforcement all link a seller's or platform's later revenue to buyer outcomes, and each can weaken or strengthen the incentive effect. The*

structure also does not imply that any particular offer is deceptive, since some promises are fulfilled. It implies only that the initial payment does not reward fulfilment, and it leaves open how strong the indirect constraints are in a given market. That strength is what the evidence below is meant to bear on, and where the indirect constraints are strong the conflict described in the next section is correspondingly weak.

For the circular case a statement about total flows would require an accounting model that specifies external customer revenue, transfers among participants, costs, and what counts as an outcome. No such model is offered here, and the case is described only as a possibility: when fees are collected from entrants at each tier, income can be sustained by continued entry and not by outcomes realised outside the chain.

9.1. A Historical Note

The financial history of the Royal Society's journal is included to mark a distinction and not an equivalence. According to the St Andrews project on the *Philosophical Transactions* [27], the journal's editors before 1752 were generally wealthy enough to finance it themselves. Henry Oldenburg appears to have been the only one who made money from it, and less than he hoped, while Hans Sloane reported spending £1,500 of his own money during his time as secretary. From 1752 the Royal Society bore the costs, and through the nineteenth century specialist journals were widely assumed to be unprofitable. Publishing produced regular surpluses only from the 1970s, when generating income became formal policy alongside the dissemination of knowledge, and by 2010 the Society's publishing surplus was about £1.4 million, roughly a tenth of its annual income. The sequence shows a long period in which the party paying for the journal was a patron or society and not the reader or the author, followed by a period in which publication became a source of income. Subsidized scholarly publication is not itself the vanity inversion. Third-party support for publication, whether from editors, patrons or learned societies, does not involve selling an author the appearance of endorsement. The vanity inversion begins where the author pays and the seller represents the result as selection. The history is useful for separating these cases, and it carries no claim about the motives of any editor.

10. The Discovery Conflict

A search or discovery service claims to surface the most relevant result for a query. Its advertising revenue, however, is greater when a paid result receives attention that the best organic result would otherwise have received. In their 1998 description of a prototype search engine, Brin and Page [4] argued that advertising-funded search engines would be inherently biased toward advertisers and away from the needs of users. The conflict is structural: the revenue model rewards a departure from the product's stated purpose, and the departure is larger the more the paid result differs from the best one.

The conflict is a gradient, and the essay does not treat it as a flat contradiction. For many commercial queries an advertisement for the thing sought is relevant, and in those cases the paid and organic interests coincide. The conflict is sharpest for offers whose value to the buyer is uncertain and whose seller's incentive to promote is high, which describes income opportunities well. The seller may expect the entry fee to exceed the cost of reaching the buyer, the buyer often cannot assess the offer before purchase, and the discovery service is paid in proportion to the reach it supplies. Refunds, complaints and enforcement can push against this, and the strength of those forces is again an empirical matter. The structure is that of the chokepoint described above: the platform stands between the aspirant and the outcome sought, and it is paid for the standing.

11. What the Incentive Argument Establishes

The argument of the preceding sections establishes that the revenue structure gives platforms a standing incentive to distribute promises of advancement without regard to their outcome, and that this incentive opposes the protective function of discovery. It does not establish what a platform knew about a given advertiser, whether enforcement was applied consistently, or whether the harm in a particular case was foreseeable. Those are questions of fact.

Evidence bearing on them falls into identifiable tiers, each with limits that should be stated alongside the findings. Advertising libraries and transparency centres maintained by the platforms show which offers were distributed, in what volume, and over what period, but they do not show internal assessments. Regulatory actions against advertisers of income opportunities show that authorities found specific conduct unlawful, and they bear on what a platform could have recognised when the same advertiser continued to run offers after-

wards. Complaint and refund data indicate the rate of buyer dissatisfaction, subject to selection effects in who complains. Public records of enforcement, including removals, policy changes and the timing of each, show how the platform responded. A claim about knowledge in a specific case would need to cite items from several of these tiers, and a claim that rested on one would be correspondingly weaker.

11.1. Evidence Assembled for AI Income Opportunities

The offer category examined here is the sale of artificial-intelligence-based income opportunities, and the evidence assembled so far comes from four sources. The descriptions below are drawn from agency announcements, complaints, public letters and press reports, and the allegations they report are allegations unless a court has found otherwise.

Regulatory actions. In September 2024 the Federal Trade Commission announced a set of cases under the name Operation AI Comply [8]. Three of them concerned income opportunities. The Commission alleged that Ascend Ecom promised that AI-powered tools would generate thousands of dollars a month in passive income, with at least \$25 million in consumer losses, and that Ecommerce Empire Builders claimed buyers could earn millions, with social media advertisements alleged to promise \$10,000 per month without substantiation. It alleged that FBA Machine promised guaranteed income through AI-powered pricing tools, with more than \$15.9 million in consumer losses. Courts temporarily halted the schemes and placed them under receivers. In March 2025 the Commission brought a case against Click Profit, which allegedly promised passive income from AI-powered online stores and took at least \$14 million, and final orders followed in August 2025 [14]. In August 2025 it sued Air AI for allegedly false earnings and refund claims, with consumers alleged to lose as much as \$250,000 [12], and in March 2026 it announced a proposed settlement that includes an \$18 million judgment, largely suspended because the defendants cannot pay it, and a permanent ban on marketing business opportunities [16]. These items show that the agency found conduct in this category sufficiently serious to seek court intervention, and in one case that social media advertising carried the earnings claim. They do not show that any platform carried a particular advertisement with knowledge of its character. The announcements do not say which channels carried the offers, a

question taken up below from the complaints themselves.

Internal platform documents. Reuters reported in November 2025, on the basis of internal Meta documents, that the company estimated about ten percent of its annual revenue, roughly \$16 billion, came from advertisements for scams and banned goods, that a December 2024 document estimated about 15 billion higher-risk scam advertisements shown to users each day, that the company acted against likely scammers only when it was highly certain of the classification and charged some suspected advertisers higher rates, and that a guardrail limited enforcement actions whose cost exceeded 0.15 percent of revenue. These characterisations are taken from a Consumer Reports letter to the Commission and state attorneys general [5] and a letter from two senators to the Commission and the Securities and Exchange Commission [2], both of which cite the reporting and neither of which independently verified it. A trade-press summary records that Meta disputed the reporting, stating that the documents were not meant for public release and may not reflect the full scope of the issue, and that its improved detection reduced user reports of scam advertisements globally by 58 percent in 2025 [21]. These documents bear directly on the incentive argument, because they describe enforcement thresholds shaped by revenue considerations. They concern scam and banned-goods advertisements in general and do not identify any income-opportunity advertiser, so they support a claim about general awareness and policy and not a claim about knowledge of any particular offer.

Complaint data. The Commission estimated that consumers lost \$2.7 billion to social media fraud between January 2021 and June 2023, with commerce scams making up 44 percent of reports but 8 percent of losses and investment scams 20 percent of reports but 53 percent of losses [22]. The data do not isolate income opportunities, and the Consumer Reports letter cites third-party research that fewer than one percent of scams are reported to government agencies, so the reports are an incomplete and unevenly selected record of harm.

Transparency tools. Meta’s advertising library shows non-political advertisements and their active status, with filters for language, advertiser, platform and media type, but according to a third-party guide it does not show spending, audience or impressions for such advertisements [20]. Google’s Ads

Transparency Center shows an advertiser’s name and location and the advertisements it served over a period, including those of advertisers that have not completed verification [17]. These tools make it possible to document which offers circulated and from whom, which is the necessary first step in testing whether the offers named in regulatory actions were carried by a given platform. They do not reveal internal assessments or enforcement decisions, and their retention periods for non-political advertisements are not established in the sources reviewed.

Channels named in the complaints. The advertising libraries could not be searched for the period in which these offers ran, since the retention of non-political advertisements is not established in the sources reviewed and the schemes were halted in 2024 and 2025. The match between offers and channels was therefore made at the level of the Commission’s complaints, which describe where each defendant advertised. Table 1 records what each complaint names.

Defendant	Channels named in the complaint	Notes
Ascend Ecom	Instagram, Facebook, X, YouTube videos, podcasts, websites; TikTok videos	Complaint states the defendants advertise on these channels; Google not mentioned
Ecommerce Empire Builders	Facebook, Instagram, TikTok, YouTube, email, website	A Facebook advertisement of 6 December 2023 from the owner’s account is described; Facebook, Instagram and TikTok advertisements linked to a free training video
FBA Machine and Passive Scaling Click Profit	Instagram, Facebook, YouTube, TikTok, email, websites Google search advertisements, TikTok video advertisements, a third-party listing site	Facebook advertisements described; the owner’s Instagram account used for earnings claims Facebook, Instagram and YouTube not mentioned; complaint filed 3 March 2025
Air AI	Social media advertising, online videos, email, text and telephone calls	No specific platform named

Table 1: Advertising channels named in Federal Trade Commission complaints against AI income-opportunity sellers, from Federal Trade Commission [9, 10, 11, 15, 13]. Entries record allegations.

Four of the five complaints name at least one specific platform as a channel for the offer. Meta’s Facebook and Instagram appear in the complaints against As-

cend Ecom, Ecommerce Empire Builders and FBA Machine, Google search advertising in the complaint against Click Profit, and TikTok in the complaints against Ascend Ecom, Ecommerce Empire Builders, FBA Machine and Click Profit. YouTube, a Google property, is named in three of them, although those complaints describe it as a place where marketing videos appeared and do not say whether paid advertising was used there. The Air AI complaint refers to social media advertising without naming a platform. Platforms are not defendants in any of these actions, and the complaints make no allegation about what a platform knew. What the match establishes is narrower and still useful: for these offers the question of whether the major advertising platforms distributed them can be answered, at the level of allegation, in the affirmative, so that the open question is now what the platforms knew and did and not whether the offers reached users through their services. The specific advertisement of 6 December 2023 described in the Ecommerce Empire Builders complaint is the kind of item against which a platform's own records could in principle be compared.

What can be claimed now. Taken together, the assembled material supports five statements. The revenue structure gives platforms an incentive to distribute outcome-independent promises. Federal authorities have alleged that AI-based income opportunities were deceptive, and their complaints name Facebook and Instagram, Google search advertising, TikTok and YouTube among the channels that carried them. Internal documents reported by Reuters and characterised by public letters describe awareness of scam advertising at scale and enforcement thresholds that account for revenue, a characterisation Meta contests. These documents concern scam and banned-goods advertisements in general. And none of this yet establishes that a platform knowingly accepted a specific income-opportunity advertisement. Closing that gap would require the platform's own records of the advertisers named in the complaints, including when each advertiser's account was opened, what review it received, what complaints about it were logged, and when and why its advertisements were removed. Those records are not public, and the tools that are public could not be used to reconstruct them for this period. No evidence about Google's enforcement history for these advertisers was assembled. The inference from scam advertisements in general to income-opportunity advertisements in particular is not drawn.

12. Constraints on Payment and Judgment

The constructive principle is that payment should not purchase favourable judgment or exemption from scrutiny. Two readings of it must be kept apart, because they lead to different institutions. One would prohibit purchased placement within discovery altogether. The other would permit advertising that is clearly separated and labelled, while insulating the functions that judge and enforce from commercial influence. The reading adopted here is the second. A platform may sell labelled placement, and a user who sees such placement knows what it is, but the ranking of organic results, the recommendation of content to people who did not ask for it, the verification of claims and the enforcement of rules should be governed by processes that an advertiser's payment cannot reach.

The separation concerns control over those four functions. A party that pays for exposure should have no influence over how its offer is verified, how organic results are ranked, whether it is recommended outside the labelled placement it bought, or how violations of the rules are detected and sanctioned. Where a platform's revenue depends on the advertiser, these functions are performed by the party with a financial interest in the advertiser's continued presence, and the principle asks for an arrangement in which that interest cannot be exercised through them. Even then the separation is partial, since the same organization pays the enforcers, and the institutional forms below are candidates and not solutions.

The principle does not condemn every arrangement in which a party is paid for evaluation, since auditors, testing laboratories and certifiers are paid for judgments and are often trustworthy. The relevant issue is whether the evaluator's compensation or continued engagement depends on rendering a favourable judgment, and what institutions constrain that dependence. Candidates include independence rules and fixed terms, disclosure of fees and conflicts, liability for negligent or false certification, external audit, ownership structures that bind a service to its users, charters whose terms cannot be amended by later management, and revenue drawn only from the people the service claims to serve. Which of them works best, and at what cost, is not settled here.

12.1. Physical Businesses as a Comparison

Businesses whose output is physical often face a tighter link between payment and delivery, because the buyer can inspect what was delivered. A recycling operation is paid for material it processes, and an analytical laboratory is paid for results that can be checked. The comparison is useful only where payment is meaningfully tied to inspectable delivery or outcome. Physicality does not by itself ensure alignment, since a physical product can be misrepresented and a service can be unverifiable, and the comparison should be applied case by case.

13. Capacities and Accountable Objectives

The promises examined here direct capable people toward objectives that were selected for them by whoever pays for distribution. Three functions are worth separating in any institution that puts skilled people to work: selecting an objective, translating it into bounded tasks, and executing those tasks competently. Competent execution is a capacity that can be cultivated and rewarded, and much of the discussion of education, training and reward concerns it. The selection of the objective is a different function, and it is the one that income-opportunity advertising may capture. Execution quality can sometimes supply evidence about an objective's feasibility or consequences, since an objective that cannot be pursued competently by capable people is suspect on that ground. What it cannot do by itself is establish the objective's social worth, and the point matters because a capable person handed a poorly chosen objective may pursue it well and thereby make it look better chosen than it was.

Three criteria can be used to assess whether the selection of an objective is accountable to the people it affects. Its justification should be open to examination (justification), the people who bear its costs should be consulted and able to contest it (consultation), and it should be revisable by a party who does not gain from its continuation (revisability). They are offered as criteria against which arrangements should be assessed and not as a verdict, and whether a given advertising arrangement meets them has to be shown case by case, for example by examining what an advertiser disclosed, whether buyers could contest the offer, and who decided whether it continued to run. The pattern described in the preceding sections gives reasons to expect failures where the party that profits from persistence controls whether the offer continues to circulate, but the reasons are not demonstrations. Agent-based workflows that organise execution

in fine detail make the distinction more pressing, since they make execution cheaper while leaving the choice of ends with whoever buys the machinery. It follows that results showing that a reward scheme cultivates a capacity within a fixed project do not establish that the project deserves the effort, and that a distribution channel's success in directing people to an objective is not evidence of the objective's worth.

14. Aspirational Exhaust: Disposable Identities

The argument so far concerns what is sold to aspirants and who is paid for delivering them. A further question concerns the supply side of attention: what kinds of entities populate the platforms on which these offers are advertised and discussed. Large social platforms host a great many minimally active Pages whose names announce a profession, a lifestyle or a state of self-realization and whose activity does not sustain the claim. A Page may call itself a studio, an academy or a path to some kind of freedom and carry a handful of posts, a stock image and a few followers, with no sign of a person or group whose continuing practice the name describes. The phenomenon is easy to notice and hard to measure from outside, because the platforms do not publish the distribution of activity across Pages, and this section therefore makes no claim about how many such Pages exist. It offers an account of the conditions that make them cheap to create and cheap to keep, and a discussion of what they might be for. The residue they leave is called *aspirational exhaust* here, and the term is used descriptively.

The account is structural. It does not hold that these Pages are all spam or all innocent, and it treats the question of what any given Page is for as an empirical matter that the structure leaves open. The claim is that a platform which gives identity no persistent cost and no binding history, and which allocates visibility according to measured interaction, will accumulate disposable identities as a matter of course, whatever any individual operator intends.

14.1. From Expression to Probing

Early social platforms leaned on identities that accumulated a social record. A profile gained standing through reciprocal recognition, and abandoning it meant giving up something built over time. Where creation is cheap and the record does not bind, a different use of identity becomes rational, which is to

instantiate a Page in order to see what the system does with it. Each Page then works as a small experiment sampling the recommendation surface for the minimum conditions of visibility, and most experiments are expected to fail. Failure is not a pathology of the individual Page, because what matters to an operator is the yield across a large population of cheap attempts, and nothing about the platform requires that any single attempt mean anything. This is the logic of failure as baseline described earlier for publishing and for generative platforms: when attempts are cheap and failure is the expected outcome, the aggregate looks less like a set of failed expressions than like a sampling process, and its residue is the population of attempts that never became anything.

14.2. Aspirational Naming

The uniformity of the names follows from the probing logic. An aspirational title is interpretable by many users, offensive to few and hard to classify, so it retains broad appeal while committing the Page to nothing. A Page named for a vague profession or virtue can be repurposed without contradiction, and the name does not constrain what the Page later posts or sells. Moderation systems that work by mapping labels to functions find such Pages difficult to act on for the same reason, since a vague label supplies no stable function to enforce against.

14.3. The Costlessness of Failure

The enabling condition is that identity creation and abandonment carry no meaningful cost. Pages can be generated at volume, sometimes with automated tools, and their dormancy or deletion brings no penalty. History does not accumulate as liability. In the loose vocabulary of thermodynamics one could say that the system lacks dissipation at the level of identity, so that abandoned structures are not reclaimed and gradually thicken the medium through which legitimate signals must travel. This is a metaphor with no quantitative content: no entropy is measured, and the point is only that nothing in the architecture removes dormant identities or charges their creators for the clutter.

14.4. Engagement Optimization as the Selection Pressure

Recommendation systems trained on dwell time, reaction speed and return visits favour content that provokes quick responses. They do not ask whether an

entity is sincere or grounded in a practice, only whether it produces measurable changes in user behaviour, and in that field persons, collectives, brands and synthetic actors are flattened into sources of engagement signals. The systems are proprietary and their objectives are known only in outline, so the argument needs only the weak premise that visibility responds to interaction and not to continuity of identity. That premise is consistent with what the platforms have said publicly about ranking, and anyone willing to compare the reach of new and long-established Pages carrying similar content could test it.

14.5. What the Pages Might Be

At least four explanations deserve separate treatment, because they predict different things and call for different responses. The first is probing, in which operators create large numbers of Pages to find cheap routes to visibility, and which predicts clusters of Pages created in bursts, with similar structure and with content copied or generated at volume. The second is the funnel, in which an aspirational Page serves as bait for an audience that holds the aspiration and the destination is a paid offer, often an offer of income or productivity. The payment structure behind that case is the one analysed above, and the explanation predicts links, forms or messages leading to a sale, and perhaps paid promotion of the Page or its offers. The third is the abandoned personal experiment, in which someone created a Page in hope, did not return, and left it to sit, and which predicts a single creation date, no coordination with other Pages and a name that reflects the creator's own interest. The fourth is the placeholder, in which a name is held in case it becomes valuable, and which predicts little or no content at all.

These explanations are not exclusive and the proportions among them are unknown. The funnel explanation is the one with a documented commercial mechanism behind it, since the complaints cited in Section 11.1 describe income-opportunity schemes advertised through social platforms, but the existence of the mechanism does not show that it accounts for any particular share of the Pages. A study able to distinguish the four would need creation timing, network structure, content provenance and outbound destinations for a defined sample, and none is attempted here.

14.6. Identity as Infrastructure

If identity were treated as a namespace in which each entity carries an irreducible history that conditions its future, transformations of that entity would be append-only and abandonment would leave a durable mark. A Page could not change silently from one function to another without that change being recorded and, within limits, costed. Dormancy would likewise have consequences, such as decay of visibility or reclamation of the name, so that inert identities did not accumulate without limit. A platform built on that footing would treat discoverability as derived from demonstrated continuity of purpose, so that visibility accrued to entities with provenance and a track record and not to the winners of a series of cheap probes. It would also apply the principle of Section 12, so that paid promotion could not substitute for the history a Page lacks. The economics of probing would then change, because each failure would carry a cost and each success would have to be earned by continuity, and the platform would tend toward fewer and more legible entities.

The proposal has costs that should be acknowledged. Binding history to identity penalizes people who legitimately change direction, and it makes anonymity and pseudonymity harder to protect. A workable design would have to separate the persistence that supports accountability from the exposure that endangers vulnerable users. How many fewer entities, and at what cost to open participation, is an empirical question this paper cannot answer, and no design is specified here.

15. Limits of the Argument

The first part of the essay is structural and offers no measurements. Its statements about failure as the statistical baseline, about platforms consolidating value from successful attempts, and about the possible use of paying users' feedback describe the logic of arrangements that are compatible with those outcomes and are not estimates of how often each occurs. The second part argues from payment structure and presents the evidence of Section 11.1 as allegations and reported documents, with the knowledge question explicitly deferred to evidence not yet assembled. The strength of indirect links between outcomes and revenue, through refunds, retention, reputation and enforcement, is not measured, and neither are success rates in any market discussed. The account of feedback from paying users is conditional and rests on secondary descriptions

of two services. The discussion of aspirational Pages is a set of competing hypotheses with the evidence that would separate them, and no count of such Pages is offered. The treatment of physical businesses is a comparison and not an analysis of any particular enterprise, and the institutional forms proposed in Section 12 are candidates and have not been tested.

16. Creation Without Reciprocal Risk

The conclusion toward which this analysis points is not primarily about any specific company or platform. It is about what happens to a cultural institution when its economic structure inverts. Publishing, in its traditional form, is a system in which institutions claim authority over what is worth making available and accept financial risk in proportion to that claim. When that system inverts, when the institution claims authority without accepting risk, and when the creator accepts risk without receiving the institutional investment that risk is supposed to purchase, something important has been broken.

What has been broken is the reciprocal relationship between creator and institution. In the conventional model, the publisher and the author are jointly committed to the work's public reception: the publisher because it has financial exposure, the author because the work is theirs. This alignment is imperfect and is frequently violated, but it structures the relationship as one of shared interest. When the publisher has already been paid by the author, the alignment dissolves. The publisher's interest in the work's success is now only incidental, a matter of reputation rather than revenue.

The deeper consequence is what this does to the cultural function of publishing. Publishing is not merely a commercial arrangement. It is, or has been, a mechanism for bringing work before audiences, for organizing cultural attention, for maintaining the connection between creators and readers that makes a literary culture possible. When the publisher's revenue is decoupled from reader demand, this function is not abolished but attenuated. Publication continues to occur in its administrative sense, books appear, ISBNs are registered, listings go up, but the institutional investment in connecting those books with readers weakens in proportion to the weakening of the financial incentive that once drove it.

Vanity publishing matters, then, not because it is embarrassing or low-status, and not because it produces books that are inferior to those of traditional pub-

lishers. It matters because it reveals what cultural institutions become when they stop bearing reciprocal risk. Once the creator's desire to be published is the commodity, and once the platform's revenue depends on that desire rather than on the public value of what is created, the institution has ceased to serve the culture it nominally represents. It has become a system for extracting value from the act of trying to participate in that culture.

The argument applies unchanged where the payer is an advertiser and the product is reach. An advertising platform paid on distribution of promises of advancement, with no part of its revenue depending on what purchasers later achieve, stands to the aspirant as the vanity press stands to the author, and the conflict it creates with discovery is the same inversion applied to the function of finding what is worth finding. The remedy is correspondingly the restoration of reciprocal risk: arrangements under which scrutiny cannot be bought, under which identities carry the history of what they have done, and under which those who select objectives answer to the people who bear their consequences.

References

- [1] Bitdefender. Anthropic shifts privacy stance, lets users share data for AI training. <https://www.bitdefender.com/en-us/blog/hotforsecurity/anthropic-shifts-privacy-stance-lets-users-share-data-for-ai-training>, 2025.
- [2] Richard Blumenthal and Josh Hawley. Blumenthal and hawley call for federal investigation of Meta's profiting from scams and fraud. Press release, 24 November 2025 (letter dated 22 November 2025). https://www.blumenthal.senate.gov/newsroom/press/release/blumenthal-and-hawley-call-for-federal-investigation-of-metas-profiting-from-scams_fraud, 2025.
- [3] Rishi Bommasani et al. On the opportunities and risks of foundation models. arXiv:2108.07258, 2021.
- [4] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117, 1998.
- [5] Consumer Reports. Letter to the FTC and state attorneys general regarding scam advertisements on Facebook and Instagram. Letter, 13 November

2025. <https://advocacy.consumerreports.org/wp-content/uploads/2025/11/Letter-to-FTC-and-state-AGs-re-FB-scams.pdf>, 2025.
- [6] Cory Doctorow and Rebecca Giblin. *Chokepoint Capitalism: How Big Tech and Big Content Captured Creative Labor Markets and How We'll Win Them Back*. Beacon Press, 2022.
- [7] Thomas Eisenmann, Geoffrey Parker, and Marshall Van Alstyne. Strategies for two-sided markets. *Harvard Business Review*, 84(10):92–101, 2006.
- [8] Federal Trade Commission. FTC announces crackdown on deceptive AI claims and schemes. Press release, 25 September 2024. <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>, 2024.
- [9] Federal Trade Commission. Complaint, federal trade commission v. Ascend Ecom. https://www.ftc.gov/system/files/ftc_gov/pdf/2423023ascendecomcomplaint.pdf, 2024.
- [10] Federal Trade Commission. Complaint, federal trade commission v. Ecommerce Empire Builders. https://www.ftc.gov/system/files/ftc_gov/pdf/2423027ecommercebuilderscomplaint.pdf, 2024.
- [11] Federal Trade Commission. Complaint, federal trade commission v. FBA Machine et al. https://www.ftc.gov/system/files/ftc_gov/pdf/Complaint-FTCvFBAMachine-et-al.pdf, 2024.
- [12] Federal Trade Commission. FTC sues to stop Air AI from using deceptive claims about business growth, earnings potential, and refund guarantees. Press release, 25 August 2025. <https://search.ftc.gov/news-events/news/press-releases/2025/08/ftc-sues-stop-air-ai-using-deceptive-claims-about-business-growth-earnings-potential-refund>, 2025.
- [13] Federal Trade Commission. Complaint, federal trade commission v. Air AI Technologies. https://www.ftc.gov/system/files/ftc_gov/pdf/Air.ai-Complaint.pdf, 2025.
- [14] Federal Trade Commission. Click profit, LLC. Case page, docket 1:25-cv-20973-DSL. <https://www.ftc.gov/legal-library/browse/cases-proceedings/click-profit-llc>, 2025.
- [15] Federal Trade Commission. Complaint for permanent injunction, monetary judgment and other relief, federal trade commission v. Click Profit.

- Filed 3 March 2025. https://www.ftc.gov/system/files/ftc_gov/pdf/ClickProfit-20250303-Complaint.pdf, 2025.
- [16] Federal Trade Commission. Air AI and its owners will be banned from marketing business opportunities to settle FTC charges that the company misled entrepreneurs. Press release, 24 March 2026. <https://ftc.gov/news-events/news/press-releases/2026/03/air-ai-its-owners-will-be-banned-marketing-business-opportunities-settle-ftc-charges-company-misled>, 2026.
- [17] Google. Ads transparency center. Google Ads Policy Help. <https://support.google.com/adspolicy/answer/13614690?hl=en-GB>, n.d.
- [18] Mary L. Gray and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt, 2019.
- [19] Jaron Lanier. *Who Owns the Future?* Simon & Schuster, 2013.
- [20] Jon Loomer. Meta ad library. Third-party guide, Jon Loomer Digital. <https://www.jonloomer.com/meta-ad-library/>, n.d.
- [21] Marketing Dive. Report claims Meta is earning billions from scam ads. Trade-press summary of Reuters reporting of 6 November 2025. <https://www.marketingdive.com/news/report-claims-meta-earning-billions-scam-ads/804945/>, 2025.
- [22] MediaPost. FTC links social media scams to targeted ads. <https://www.mediapost.com/publications/article/389964/ftc-links-social-media-scams-to-targeted-ads.html>, 2023.
- [23] Long Ouyang et al. Training language models to follow instructions with human feedback. arXiv:2203.02155, 2022.
- [24] Nick Srnicek. *Platform Capitalism*. Polity Press, 2017.
- [25] John B. Thompson. *Merchants of Culture: The Publishing Business in the Twenty-First Century*. Polity Press, 2012.
- [26] Tom's Guide. Keep your ChatGPT data private by opting out of training: here's how. <https://tomsguide.com/ai/keep-your-chatgpt-data-private-by-opting-out-of-training-heres-how>, 2025.

- [27] University of St Andrews, Philosophical Transactions project. How much money did the Philosophical Transactions make? <https://arts.st-andrews.ac.uk/philosophicaltransactions/how-much-profit-did-the-philosophical-transactions-make/>, n.d. Web page, accessed October 2026.
- [28] Shoshana Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.