

The Attention Economy Is Dying

And What Comes After

Eleanor Whitfield

Kelford Press

First published in 2026 by Kelford Press

Copyright © 2026 Kelford Press. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written permission of

the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law.

ISBN 978-1-7394521-0-8 (Digital) ISBN 978-1-7394521-1-5 (Print) ISBN 978-1-7394521-2-2 (Audio)

Kelford Press Where Words Find Their Home Est. 2006

www.kelfordpress.com

For everyone who ever closed a tab and opened a book instead.

Contents

1. The Fifteen-Second Funeral
2. Peak Content
3. The Authenticity Paradox
4. The Collateral Damage
5. The Resistance
6. The Economics of What Comes Next

7. The Post-Attention World

Acknowledgements About the Author Also by Kelford
Press

Chapter 1: The Fifteen- Second Funeral

Adunni Okafor pressed her thumb against the glass and held it there for a beat longer than necessary, as though the act of posting deserved some small ceremony. It was 11:47 p.m. in Lagos, and the ceiling fan above her desk stirred the humid air without conviction. She had spent fourteen hours on this video — a seven-minute investigation into how Chinese textile imports were reshaping the markets of Balogun and Oshodi, filmed on two phones and edited on a cracked laptop running DaVinci Resolve. She had interviewed three traders, a customs broker, and an economics lecturer at the University of Lagos. She had colour-graded the footage to feel warm but serious. She had rewritten her script four times.

The upload bar filled. Instagram's servers, humming in data centres she would never see, ingested her work. Within ninety seconds, the platform's recommendation

engine had made its assessment: the video lacked the velocity markers — the first-frame hooks, the text overlays, the emotional triggering — that its classifiers associated with viral potential. It was served to a thin slice of her 11,000 followers. A handful watched. Most scrolled past. By midnight, the verdict was legible in the metrics: 200 views. No meaningful shares. The algorithm had already moved on.

Adunni is not a real person, though she is drawn from composites of several creators I interviewed in Lagos and Accra between 2024 and 2025. Her experience, however, is so common as to be banal. Every day, millions of people pour hours of skilled labour into content that is dead on arrival — not because it lacks quality, not because it fails to inform or move or provoke, but because a sorting mechanism designed to maximise advertising revenue has determined, in fractions of a second, that it will not generate sufficient engagement to warrant distribution. The funeral is over before most mourners know there was a death.

This book is about the system that made that possible, why it is failing, and what is already beginning to replace it.

But first, we need to understand how we arrived here — how human attention became the most traded commodity on Earth, and why the market for it now approaches something that looks very much like collapse.

The Economist Who Noticed

In 1971, Herbert Alexander Simon sat in his office at Carnegie Mellon University in Pittsburgh and wrote a sentence that would, decades later, be quoted so often and so carelessly that its original meaning would be almost entirely lost. "A wealth of information," he observed, "creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it."

The sentence appeared in a paper titled "Designing Organizations for an Information-Rich World," published in a collection edited by Martin Greenberger for Johns Hopkins University Press. The context matters, because the context is precisely what the Silicon Valley mythmakers tend to omit.

Simon was not writing about the internet. The internet, in 1971, was a handful of university mainframes connected by telephone lines; the first email would not be sent until later that year. He was not writing about social media, about advertising, about content creation, or about any of the phenomena later gathered under the banner of "the attention economy." He was writing about organisations — about how large bureaucracies and corporations could make better decisions when overwhelmed by data. His concern was managerial. How should a firm allocate its limited cognitive resources when available information exceeds the capacity of any individual or committee to process it?

Simon was a polymath — Nobel laureate in Economics (1978), pioneer of artificial intelligence, contributor to cognitive psychology and political science. He moved between disciplines not out of diletantism but because the problems he cared about refused to stay within departmental boundaries. His observation about attention was almost modest — an aside within a broader argument about information system design. The bottleneck in any information-rich environment, he argued, is not the information itself but the human capacity to attend to it. More reports, more memos, more

data do not automatically produce better decisions. They may produce worse ones, by scattering focus and rewarding whoever shouts loudest.

What Simon did not anticipate — what no one in 1971 could have anticipated — was that his observation would become the founding logic of an industry worth over seven hundred billion dollars. He described a problem of organisational cognition. The technology companies of the late twentieth and early twenty-first centuries built something different: a marketplace in which human attention itself was the product, harvested at scale and sold to the highest bidder.

The distance between Simon's insight and what it became is the distance between noticing that a river flows downhill and building a hydroelectric dam. The observation is neutral. The infrastructure is not.

The Long History of Selling Eyes

Tim Wu, the Columbia Law professor and later Biden White House competition adviser, published *The Attention Merchants* in 2016. The subtitle — *The Epic Scramble to Get Inside Our Heads* — gestured at melodrama, but the history was careful and genuinely illuminating.

Wu argued that the commodification of attention did not begin with Google or Facebook. It began in the 1830s, when Benjamin Day launched the *New York Sun* at one penny — far below production cost. Day's innovation was economic, not journalistic: a newspaper could be given away cheaply if advertisers paid to reach the large audience a cheap paper attracted. The reader was not the customer. The reader was the product.

This model — attract attention with content, then sell that attention to advertisers — proved so durable that it would underpin nearly every mass medium for the next two centuries. The patent medicine hawkers of the late nineteenth century perfected the attention-grabbing advertisement, promising miraculous cures in prose so lurid it could make a carnival barker blush. Radio arrived in the 1920s and, after a brief period in which it was imagined as a public utility, was colonised by commercial broadcasters who sold airtime to sponsors. Soap

manufacturers funded serialised dramas — hence "soap operas." The attention was the wife at home; the product was Palmolive.

Television intensified the exchange. By the 1960s, the American networks had refined a system in which programmes existed primarily as vehicles for advertising. Newton Minow, FCC chairman, famously described television in 1961 as "a vast wasteland" — but a vastly profitable one. The economics were simple and brutal: attract the largest possible audience, then sell thirty-second intervals of its attention to corporations. The Nielsen rating became the currency — a crude but effective measure of how many eyeballs were pointed at a screen at any given moment.

What changed with the internet was not the underlying logic but the precision and scale of its application.

The Machine Learns to Count

The year 2000 deserves a marker in this history. In October, Google launched AdWords, a system that allowed advertisers to bid on specific search terms and pay only when a user clicked on their advertisement. The implications were revolutionary, though they unfolded slowly.

For the first time, attention could be bought and measured with something approaching precision. A newspaper advertiser in 1950 knew roughly how many people subscribed to the paper, but had no reliable way of knowing how many had looked at his advertisement, let alone acted on it. A television sponsor knew the Nielsen ratings but could not distinguish between a viewer who watched the advert attentively and one who had gone to the kitchen for a cup of tea. The old line, attributed to John Wanamaker — "Half the money I spend on advertising is wasted; the trouble is I don't know which half" — had been the industry's rueful motto for a century.

Google abolished that ambiguity, or appeared to. With AdWords, an advertiser could track how many people searched for "running shoes," how many saw the ad, how many clicked, and how many bought a pair. The feedback

loop was immediate, granular, and intoxicating. Advertisers believed, for the first time, that they could measure the return on every dollar.

This was both the great promise and the great trap. Because once you could measure clicks, clicks became the thing that mattered. And once clicks became the thing that mattered, the entire apparatus of content production began to orient itself around generating them.

Six years later, Facebook introduced the News Feed. Before September 2006, Facebook had been a static collection of profile pages. The News Feed transformed it into a river — an algorithmically curated stream of updates, photographs, links, and advertisements in an endless scroll. The design borrowed from television's logic: keep the viewer watching; never let the screen go dark; fill every silence with something new.

The News Feed's algorithm — later refined, expanded, and replicated by every major platform — was the mechanism by which attention harvesting became industrialised. It learned, through billions of micro-interactions, what kept each individual user engaged: what they clicked, what they lingered on, what they shared, what provoked them to comment. It then served them more of whatever that was. The system was not

designed to inform, to educate, or to connect, though it sometimes did all three. It was designed to maximise time-on-site, because time-on-site was the inventory that Facebook sold to advertisers.

Mark Zuckerberg, in a 2007 presentation to advertisers, described Facebook's ambition with a candour that would later seem incriminating: "Once every hundred years, media changes. The last hundred years have been defined by the mass media... In the next hundred years, information won't just be pushed out to people, it will be shared among the millions of connections people have." The framing was democratic. The economics were extractive.

The Metrics That Ate the Internet

By the early 2010s, a vocabulary had emerged to describe the units of the attention economy, and each metric carried within it a set of incentives that would quietly reshape how information was produced and consumed across the planet.

Clicks were the original currency — the direct descendants of Google's pay-per-click model. A click indicated, at minimum, that a user had been sufficiently intrigued (or misled) to follow a link. The optimisation of content for clicks produced "clickbait," a term that entered common usage around 2014: headlines engineered for curiosity gaps ("You Won't Believe What Happened Next"), emotional provocation ("This Video Will Make You Cry"), or deliberate vagueness ("She Walked Into the Room. What Happened Next Changed Everything"). The form was perfected by outlets like Upworthy and BuzzFeed, though it spread rapidly to legacy media organisations desperate for digital traffic.

Views — or impressions — measured how many times content was displayed, regardless of engagement. On YouTube, a "view" counted after roughly thirty seconds of watch time; on Facebook, after just three seconds — so low that a user pausing while scrolling registered as a "viewer." In 2016, Facebook admitted it had overestimated average video watch time by 60 to 80 per cent — a discrepancy that drove a massive industry "pivot to video" in which publishers laid off writers and invested in video production on the basis of fraudulently inflated metrics.

Time-on-site measured how long a user remained on a platform or page. This metric incentivised infinite scrolling, autoplay videos, and the fragmentation of articles into multi-page slideshows — anything that kept the clock running. It also incentivised outrage, because outrage is sticky: a user arguing in the comments section of a divisive post is a user who has not left the platform.

Engagement rate — the ratio of interactions (likes, comments, shares) to impressions — became the metric by which social media algorithms ranked content. High engagement meant wider distribution; low engagement meant oblivion. This created brutal Darwinian pressure: content that generated strong emotional responses (anger, fear, amusement, moral indignation) was rewarded with reach, while content that was merely informative, carefully argued, or quietly beautiful was starved of oxygen.

Each metric was, at best, a proxy for human attention. None measured comprehension, satisfaction, trust, or value. A user who clicked an article, realised it was misleading, and left in disgust registered identically to one who read it carefully and found it useful. The metrics measured activity, not meaning — and the platforms, by

optimising relentlessly for activity, built a machine extraordinarily good at capturing attention and extraordinarily bad at doing anything worthwhile with it.

The result, by the mid-2020s, was an information environment in which the most successful content was not necessarily the most truthful, the most beautiful, or the most useful. It was the most engaging — where "engaging" was defined circularly as "whatever the algorithm decides to distribute."

The Cracks in the Cathedral

The system began to falter, as such systems tend to, not with a dramatic collapse but with a slow accumulation of contradictions.

The first and most measurable crack was the decline of organic reach. On Facebook, organic reach for brand pages averaged roughly 16 per cent in 2012, fell to 6 per cent by 2016, and by 2024 sat at approximately 2 per cent. A creator who had spent years building 100,000 followers could now expect 2,000 of them to see a given

post — unless they paid Facebook to distribute it. The platform had offered creators free tools to build communities, then throttled access to force the purchase of advertising. A bait-and-switch executed at civilisational scale, entirely legal.

The second crack was ad-blindness. By 2013, studies found that most internet users had learned to unconsciously filter out banner advertisements. Eye-tracking research showed gazes skipping over ad placements entirely. By the 2020s, the average user was exposed to an estimated 6,000 to 10,000 advertisements per day, and the cognitive response was not attention but perceptual anaesthesia. The noise had become so constant that the brain treated it as silence.

The third crack was economic. Between 2020 and 2025, the average CPM on Facebook rose by approximately 89 per cent, according to Revealbot. Over the same period, click-through rates fell by roughly a third. Advertisers paid more to reach people less likely to respond.

This is the arithmetic of a failing market. When the price of a commodity rises while its utility falls, something structural has changed. The attention economy was not

merely becoming more expensive. It was becoming less productive. The machine was running hotter and delivering less.

The Geography of Attention

There is a persistent distortion in how the attention economy is discussed in the Anglophone press: it is treated as though it were invented in Silicon Valley and experienced in San Francisco. The reality is almost precisely the reverse. The attention economy was designed in California, but it is lived most intensely in Lagos, Jakarta, Mumbai, Sao Paulo, Nairobi, and Manila.

In 2025, the global average daily time on social media was approximately 2 hours and 20 minutes (DataReportal/GWI). But this average concealed vast disparities: over 4 hours in Nigeria and the Philippines, 3 hours and 40 minutes in Brazil. In India, where Jio had brought half a billion people online in under a decade, the 18-to-34 demographic averaged over 4 hours daily.

These were not idle hours. For millions in the Global South, social media platforms served as essential economic infrastructure — shopfronts in Lagos, political channels in Jakarta, education systems and labour marketplaces in Mumbai. In countries where legacy institutions were inaccessible, unreliable, or expensive, the platforms had become the foundational layer of digital life.

The consequence was double extraction. Attention in Lagos or Manila was harvested by algorithms designed in Menlo Park and sold to advertisers paying in dollars and euros. The value flowed upward and outward. The content — millions of hours of creative labour keeping the platforms alive — was produced locally, often with great skill, but compensated (if at all) through revenue-share models paying fractions of a cent per view in the lowest-CPM markets.

Adunni's fourteen hours of work were not merely unrewarded by the algorithm. They were part of a system in which the attention of Nigerian users was worth, to advertisers, roughly one-tenth of what American attention was worth — while the cost of producing content (in time, in equipment, in creative energy) was not proportionally

lower. The attention economy did not merely commodify attention. It commodified it at rates that reflected and reinforced existing global inequalities.

This is not an incidental feature of the system. It is the system.

The Fifteen-Second Window

By the mid-2020s, the temporal logic of the attention economy had compressed to absurdity. TikTok, which had overtaken Instagram as the dominant platform for users under 25 by 2023, evaluated content with extraordinary speed. A new video's fate was typically sealed within fifteen seconds: the algorithm served it to a test audience, measured immediate engagement, and either amplified it or let it sink.

The creator economy that emerged around this logic was, in its way, a marvel of adaptation. Millions of people learned to front-load their content with hooks — visual, verbal, or emotional — designed to survive the fifteen-second window. They learned to speak faster, cut more

aggressively, provoke more reliably. They learned the grammar of virality: the pattern-interrupt opening, the promise of revelation, the emotional payoff timed to arrive before the thumb could scroll onward. It was a new rhetoric, as formally demanding in its own way as the classical oratory of Cicero, and as ruthlessly optimised for its medium.

But the optimisation came at a structural cost. When the sorting mechanism rewards what is immediately arresting and punishes what is slow or complex, the information environment selects against depth. Investigative journalism cannot compete in a fifteen-second window with a video of someone falling off a skateboard. A careful analysis of trade policy cannot match a rage-bait clip. The algorithm does not suppress serious content through malice. It suppresses it through indifference — which is worse, because there is no one to appeal to, no editor to persuade, no decision that can be reversed.

This was what Adunni Okafor faced when she pressed "post": a system that processed her work in seconds, judged it against millions of competing signals, and assigned it a reach bearing no relationship to its quality or value. A system that had made its builders extraordinarily

wealthy by performing a single function at planetary scale: the conversion of human attention into advertising revenue.

And a system that was, by the time she posted that video, already dying — though it did not yet know it, and neither did she.

Economies do not die of single causes. They die of accumulated contradictions — of the growing distance between what a system promises and what it delivers. The attention economy is dying not because someone invented a better alternative (though alternatives are emerging), and not because regulators have acted (though they have begun to, haltingly). It is dying because the fundamental transaction — *give us your attention and we will give you content for free* — has degraded to the point where neither side of the bargain is being honoured.

The content is not free; it is paid for with data, with cognitive capacity, with hours that might have been spent otherwise. And the attention is no longer truly given; it is extracted, manipulated, and sold in a marketplace where buyer and seller are both, increasingly, aware that the goods are damaged.

What comes after is the subject of this book. But first, we need to understand what is breaking and why. The next chapter turns to the recommendation algorithms that govern what billions of people see, read, watch, and believe — and the mounting evidence that those algorithms now optimise for outcomes that serve no one.

The fifteen-second funeral is not just a metaphor for one creator's disappointment. It is the sound of a system consuming itself.

Chapter 2: Peak

Content

In January 2017, Reed Hastings stood before a room of investors in Los Gatos, California, and made a remark that would age poorly. Netflix's biggest competitor, he said, was not HBO, not Amazon, not any studio or streaming service. It was sleep. The line earned knowing laughter and approving nods from analysts who understood it as a declaration of dominance. Netflix had become so addictive it was competing with a biological imperative. The stock rose. The quote circulated.

Eight years on, the boast reads as a warning. Not because Netflix lost the war against sleep — by most measures, sleep deprivation in the industrialised world has worsened — but because Hastings failed to anticipate who else would join the battlefield. He imagined a contest between Netflix and the circadian rhythm. What arrived was a content avalanche so vast, so unrelenting, that it buried not only sleep but the economics of attention that Netflix's business model depended upon.

By 2025, Netflix was competing with five hundred hours of video uploaded to YouTube every minute. With four million podcast episodes published each month. With over five hundred million active blogs. With newsletters and live streams and TikTok clips and Reels and Shorts and X threads and Substack essays and Telegram channels — a planet-wide firehose no metaphor can contain. Hastings had been right about one thing: there was a finite resource at stake. He had catastrophically underestimated the number of claimants.

This is the story of what happens when content supply overwhelms attention supply. It is a story about economics — the oldest kind.

The Numbers That Broke the Model

To grasp the scale, anchor to a single platform and watch the curve steepen. In 2012, YouTube users uploaded roughly seventy-two hours of video per minute. By 2017: three hundred hours. By 2025: over five hundred. These are not marginal increases. They represent exponential

expansion in the volume of media competing for human eyes. And YouTube, for all its enormity, is merely one platform among dozens.

Consider podcasting. Apple's directory crossed one hundred thousand shows in 2013 — seemingly a mature market. By 2020: over a million. By 2025: four million, with four million new episodes each month. Most have audiences that could fit in a lift. Edison Research found that more than half of all podcasts launched after 2022 published fewer than ten episodes before going silent — the "podfade," a quiet epidemic of creators who discovered that speaking into a microphone was the easy part. Finding anyone to listen was the hard part.

The blogosphere never died. It metastasised. WordPress alone hosts over eight hundred million websites. Medium, Ghost, Hashnode, and dozens of smaller platforms add millions more. The volume of written content published online each day in 2025 exceeds what the entire human species produced per year in the 1990s. Often cited as a triumph of democratisation — but also a flood, and floods do not respect the quality of what they submerge.

Perhaps the most striking single statistic: the median YouTube video uploaded in 2025 receives fewer than fifty views. Not fifty thousand. Not five hundred. Fifty. The top

three per cent of creators capture ninety per cent of all viewing time. The remaining ninety-seven per cent share the scraps. This is not a long tail. It is a cliff face with a narrow ledge at the top.

Everyone Became a Publisher

The newsletter boom of the early 2020s illustrates the full arc of a content gold rush — from genuine innovation to saturated exhaustion — in compressed time.

Substack launched in 2017 with a seductive proposition: write what you want, own your audience, charge for your work. By 2018, it had roughly eleven thousand paying subscribers — modest but promising. Then the pandemic arrived, and with it a collapse in institutional trust, a surplus of homebound professionals, and a yearning for individual voices. By 2021, Substack's top writers earned seven-figure incomes. By 2023: two million paying subscriptions. By 2025: thirty-five million.

The imitators followed. Beehiiv, launched in 2021 by former Morning Brew engineers, grew to over a hundred thousand active newsletters within two years. Ghost attracted writers who wanted more control. ConvertKit

(now Kit) pivoted from email marketing to full creator platform. LinkedIn introduced its own newsletter feature, and by 2024 over a hundred and fifty thousand LinkedIn newsletters existed, many with subscriber counts in the hundreds of thousands.

The result was predictable. When everyone is a publisher, publication ceases to be remarkable. The inbox became a content battleground. The average professional in 2025 received over a hundred and twenty emails per day, roughly a third of them unread newsletters. Open rates declined from forty-two per cent in 2020 to twenty-six per cent by 2025. The medium that was supposed to rescue creators from algorithmic tyranny had, within five years, developed its own attention crisis.

Hamish McKenzie, Substack's co-founder, argued that the volume was a feature, not a bug — more voices, healthier ecosystem. This is true in the way that a forest fire is good for soil regeneration. The established organisms survive and find richer ground. The saplings burn.

When Supply Is Infinite

The economics here are not complicated. They are introductory-level supply and demand, playing out at global scale with unusual cruelty.

In any market where supply outpaces demand, prices fall. In the attention economy, the "price" is denominated not in currency but in audience — the number of people willing to give their time to a given piece of content. Total global attention is bounded: roughly eight billion humans, each awake for approximately sixteen hours a day. No technology can expand this. The human attention budget is, for all practical purposes, fixed.

Against this fixed demand, supply has grown by orders of magnitude. A creator publishing on YouTube in 2015 could, with good work and some luck, build a meaningful audience within a year. The competition was fierce but finite. By 2025, the same creator producing work of identical quality would face a content field so dense that discovery — not creation — became the central challenge. Getting seen required a parallel career in promotion, platform optimisation, and cross-posting.

A 2024 SignalFire report estimated over two hundred million people globally considered themselves creators. Of these, roughly two million — one per cent — earned enough to call it a livelihood. The rest occupied a grey

zone between hobby and hustle, subsidised by day jobs, savings, or family members who did not ask too many questions about when the audience would materialise.

This is not a failure of talent. It is a structural condition. When two hundred million people compete for a fixed supply of human attention, the mathematics are merciless.

The Streaming Plateau

If you want to understand how peak content corrodes even the most established creative industries, look at music.

In 2014, Spotify's average per-stream payout was approximately \$0.007. By 2020, as users grew faster than the payout pool, it fell to \$0.0045. By 2025: below \$0.003 — less than a third of a penny per play. An independent artist needed roughly one and a half million streams per month to earn minimum wage. Every month.

The paradox is excruciating. More people are making music than at any point in human history. DistroKid reported over a hundred thousand new track uploads per

day by 2024. Spotify's catalogue exceeded a hundred million tracks. Cheap software and global distribution meant barriers to entry had vanished. But barriers to livelihood had grown enormous.

The Music Industry Research Association found in 2024 that the median annual income from recorded music for "professional musicians" in the United States was \$3,400. Not monthly — annually. The mean was considerably higher, dragged upward by superstar outliers, which made the median all the more damning. The top one per cent of Spotify artists accounted for over ninety per cent of streams. The economist Alan Krueger called it a "winner-take-all" market — though "winner-take-almost-everything" would be more precise.

Daniel Ek, Spotify's chief executive, responded to criticism by arguing that artists needed to release music more frequently. "You can't record music once every three to four years and think that's going to be enough," he told an interviewer in 2023. The comment drew fury from musicians, but it was, in its grim way, an honest description of the treadmill. The platform's economics demanded volume. The algorithm rewarded recency. The artists who adapted — releasing singles every few weeks,

maintaining a constant presence in playlists — survived. Those who worked at the pace that art sometimes requires were algorithmically punished for their absence.

The Content Treadmill

The treadmill metaphor has become so common among creators that it risks losing its force. Let me be specific about what it means in practice.

A YouTuber in 2025 with two hundred thousand subscribers typically published twice weekly to avoid "algorithmic decay" — the reach reduction platforms impose on slower channels. Each ten-minute video required research, scripting, filming, editing, thumbnail design, metadata optimisation, and cross-platform promotion: easily thirty hours of labour. Two per week meant sixty — excluding analytics, sponsorships, and the administrative burden of running a one-person media company.

A 2024 study in the *Journal of Occupational Health Psychology*, surveying 1,200 full-time creators across six countries, found sixty-one per cent reported clinical burnout symptoms. Forty-three per cent reported anxiety

tied to algorithmic unpredictability. Twenty-seven per cent had sought professional help for mental health issues they attributed to content creation.

These are not dilettantes' complaints. These are the working conditions of a labour force the platforms call "partners" — a word implying equality that describes a piece-rate system where the rate is set unilaterally and changes without warning. A single algorithmic adjustment could halve a creator's income overnight. No appeal. No union. Only the next video.

The Algorithm Demands Novelty

Behind the content treadmill sits the machine that drives it: the recommendation algorithm. And the defining feature of the modern recommendation algorithm is its insatiable hunger for novelty.

TikTok, the dominant cultural platform for users under thirty by 2025, exemplified this most starkly. Flamingo Group found that the average "content lifecycle" on TikTok — the window during which a piece of content receives the majority of its views — had shrunk to approximately forty-eight hours. A video was either

amplified within two days or effectively dead, buried beneath newer material. Compare this to YouTube, where search-driven discovery gave content a half-life of roughly ninety days, or to a newspaper article that might once have shaped conversation for weeks.

On TikTok, consistency mattered less than frequency. A daily poster had seven times the algorithmic surface area of a weekly one — not because each video was likelier to succeed, but because each was a fresh lottery ticket. The engine prioritised novelty because novelty generates measurable interaction. A new video produces fresh data. An old video, however brilliant, produces nothing the algorithm can metabolise.

This created a paradox. Platforms needed quality content to retain users, but their algorithms punished the conditions under which quality is produced — time, reflection, revision, rest. A Gresham's Law of content: the quick and adequate drove out the slow and excellent. Exceptional work still found its audience. But the structural incentives pointed toward volume over depth, frequency over craft.

Beyond the Western Frame

The content explosion is often narrated as a Silicon Valley story. This framing misses most of the picture.

In India, the content explosion arrived with the force of a detonation, and on a specific date. On 5 September 2016, Mukesh Ambani launched Reliance Jio, a mobile network offering 4G data at prices so low they effectively made mobile internet free — less than 10 rupees per gigabyte, down from 250. Within six months: a hundred million subscribers. Within three years: six hundred million smartphone users, most online for the first time.

What followed was a second internet, built in Hindi and Tamil and Telugu and Bengali and Marathi and dozens of other languages. YouTube India became the platform's largest market by watch time. By 2025, India produced more digital content per day than any nation except the United States, and the gap was closing.

The pattern repeated across the Global South. Indonesian TikTok had, by mid-2024, more daily active users than the platform had in the entire United States. Creators in Jakarta produced shoppable content at a pace that made American influencer culture look sedate.

In Nigeria, the fintech revolution — Paystack, Flutterwave, mobile money — removed the payment friction that had blocked creator monetisation. Content exploded on YouTube, Instagram, and Audiomack. Lagos became, per capita, one of the most content-productive cities on earth.

In Brazil, Pix — the central bank's instant payment system, launched in November 2020 — made micropayments trivial. By 2024, a hundred and fifty million Brazilians used it. Creators on YouTube, Kwai, and TikTok operated in a market of two hundred and fifteen million people, most digitally connected, nearly all drowning in content.

The Western narrative of "too much content" dramatically understates the scale. In 2025, English-language content represented less than a quarter of global output. Five hundred hours per minute on YouTube is one platform's figure. The true volume across all platforms, all languages, all formats defies quantification. We are in the territory of "more than anyone can consume, by a margin so large the precise figure no longer matters."

The Silence at the Bottom

Return to that statistic: the median YouTube video in 2025 receives fewer than fifty views. Right now, someone is uploading a video they spent hours creating — a cooking tutorial, a guitar lesson, a short film, a travel vlog — and it will be seen by fewer people than could fit in a classroom. The creator will check analytics for days, watching the counter stall at thirty, forty-two, forty-seven. They will wonder what they did wrong. They will study the thumbnail, the title, the tags. Many will conclude they need to try harder, produce more, optimise better. Most will eventually stop.

People have always produced creative work that found no audience. The village poet, the Sunday painter, the garage band that never played beyond the local pub — obscurity is the historical norm for creative endeavour. What is new is the combination of accessibility and invisibility. The platforms promised that anyone could find an audience. They delivered a system in which everyone can publish but almost no one is heard. The gap between the promise and the reality is not a bug. It is the inevitable consequence of infinite supply meeting finite demand.

And the gap is widening. As generative artificial intelligence makes content creation faster and cheaper — a topic for later chapters — the supply curve will steepen. The top three per cent will capture an even larger share, because in a world of infinite content, curation becomes the scarce resource, and platforms will consolidate attention around creators already certified as reliable.

The Turning Point

Peak content is not a prediction. It is a description of the present. Content supply has overwhelmed attention supply so thoroughly that the economics of the attention economy have broken. The creator economy — which promised a new middle class of independent workers — has produced a power-law distribution so steep it resembles the pre-industrial aristocracy more than the democratic future its evangelists described.

Hastings understood that attention was finite and that capturing it was the game. What he did not understand — what almost no one in Silicon Valley understood — was that platforms designed to maximise content production

will eventually choke on their own output. The competitor was never sleep. The competitor was everything, all at once, forever.

But peak content is only half the crisis. An ocean of content is manageable — if you trust the compass. The deeper problem is that the compass has broken. In the rush to produce more, faster, for more platforms, something essential was lost: the ability to distinguish what is real from what is performed, what is honest from what is optimised, what is worth your attention from what merely captured it.

That is the subject of the next chapter: what happens when an entire civilisation's information ecosystem loses its claim to authenticity — and what it means that the loss was not accidental but designed.

Chapter 3: The Authenticity Paradox

It is four o'clock in the morning in Gangnam, and Yuna Park is staging her spontaneity.

The ring light — a twenty-inch Neewer, the preferred instrument of Seoul's creator class — casts its telltale circle across her coffee. She adjusts the mug three centimetres to the left. The reflection persists. She switches to a matte-finish ceramic cup, pours again, wipes the counter. The apartment is immaculate, though her 1.3 million TikTok followers will never know this, because the camera captures only a narrow band of countertop, a window showing pre-dawn blue, and her face, bare of makeup, still puffy from what is supposed to look like sleep.

"My mornings are so boring," she will type in the caption, followed by a shrug emoji. "Just me and my coffee."

She will reshoot this seven-second clip eleven times before breakfast. Two takes are discarded because she blinked too naturally. Three because the lighting made her skin look uneven. One because the coffee had gone cold and the steam — a crucial signifier of candour, of a moment just seized — had dissipated. In the final cut, there is a half-second of apparent camera fumble, as if she had propped her phone against a stack of books and hit record as an afterthought. That fumble took four attempts to get right.

This is not a story about Yuna Park. She is talented, shrewd, and doing precisely what the market demands. This is a story about what happens when an entire civilisation decides it craves authenticity and then builds an industrial apparatus to mass-produce it.

I. The Authenticity-Industrial Complex

The word "authentic" derives from the Greek *authentikos* — "original, genuine, principal." For most of human history, it was not something one performed. Authenticity was what remained when performance stopped. But in the attention economy, where every human behaviour is potential content, the cessation of performance became the most valuable performance of all.

The progression was swift and grimly predictable.

Between 2015 and 2018, Instagram rewarded careful curation: flat-lay brunches, golden-hour travel shots, immaculate apartments styled by Scandinavian design magazines. Influencers operated as living catalogues, each image processed through the same VSCO preset, calibrated to suggest effortless refinement. The unspoken contract was simple: *I show you beauty; you give me your attention; brands pay me for that attention.*

Then the market shifted. By 2019, Later's research showed that Instagram posts featuring "unfiltered" or "no makeup" in their captions received 24 per cent higher engagement than polished equivalents. The audience, saturated by manufactured perfection, had developed what Sianne Ngai might call an aesthetic allergy.

Perfection had become boring — worse, suspicious. If everyone's life looked like an advert, the advert was no longer persuasive.

What followed was Instagram's "authentic era."

Influencers began posting imperfect photos. They shared acne. They photographed unmade beds. They wrote long, confessional captions about anxiety, imposter syndrome, and the pressure of maintaining an online presence. The aesthetic shifted from aspirational to relatable. This was not, in itself, cynical; many creators were genuinely exhausted by the performance of perfection, and their admissions were sincere.

But sincerity, under the attention economy, is a resource. It gets mined.

Within months, the "authentic" aesthetic had its own playbook. Agencies coached clients on how to appear vulnerable. Templates emerged for confessional captions. A 2020 guide from the influencer management firm Fohr advised creators to "schedule your breakdowns" — not to fabricate emotional crises, but to ensure that genuine difficulty was captured and posted at optimal engagement windows (Tuesdays and Thursdays, between 11 a.m. and 1 p.m. EST).

Then TikTok detonated the entire framework. By 2022, the platform's algorithm had trained its users to prize a new register: raw, chaotic, apparently unmediated. The dominant format — the "Get Ready With Me," the "day in my life," the mid-sentence ramble — was built on the premise that the viewer received unfiltered access to a real person's real existence. Shaky camera work. Bad lighting. Mid-sentence corrections. The aesthetic grammar of authenticity was now fully codified.

This is the authenticity-industrial complex: the systematic, market-driven manufacture of the appearance of genuineness. Not every creator is dishonest. But the system cannot distinguish between genuine vulnerability and its simulation — and does not need to. The algorithm measures engagement, not sincerity. A staged moment of rawness and an actual moment of rawness produce identical metrics. The incentive to manufacture authenticity is structurally inevitable.

II. The Deinfluencing Mirage

No episode illustrates the paradox more precisely than the rise and instantaneous co-option of "deinfluencing."

In January 2023, a TikTok creator named Michelle Skidelsky posted a video telling followers which popular products were not worth buying — a consumer looking into the camera and saying, essentially, *save your money*. Within two weeks, the hashtag #deinfluencing had accumulated over 300 million views. Hundreds of creators joined in, cataloguing the overhyped and overpriced. For a brief moment, it looked like genuine counter-programming: the attention economy producing antibodies against its own excesses.

The moment lasted approximately three weeks.

By mid-February, brands had grasped the opportunity. Deinfluencing was not a threat to the influence economy; it was a new channel within it. Creators who told audiences to skip a forty-dollar mascara would, in the same breath, recommend a twelve-dollar alternative — often with an affiliate link. The products recommended *against* frequently saw sales spikes, driven by the Streisand effect of public dismissal. CeraVe, the skincare brand, reportedly experienced a 19 per cent increase in online searches during the peak of deinfluencing discourse, much of it driven by creators citing it as the

sensible, unglamorous alternative to luxury products. CeraVe is owned by L'Oreal, the largest cosmetics company on Earth.

Marketing analytics firm CreatorIQ tracked the phenomenon closely. By March 2023, sponsored deinfluencing content — deinfluencing videos that were themselves paid brand partnerships — had appeared on the platform. The revolution had been monetised before its first month was out.

This pattern recurs with metronomic regularity. Every counter-movement within the attention economy is reabsorbed by it. Minimalism becomes a content category. Anti-consumption becomes a consumption niche. "Logging off" becomes a video series about logging off. The system is not fragile; it is antifragile, in Nassim Nicholas Taleb's sense. It feeds on its own disruptions.

III. The Burnout Machine

The human cost of this system is not metaphorical.

In 2023, the creator economy platform Kajabi surveyed 1,200 full-time content creators across the United States, the United Kingdom, and Brazil. The findings were stark: 61 per cent reported burnout. Forty-two per cent said their mental health had worsened since becoming full-time creators. Twenty-eight per cent had considered quitting in the previous six months but felt financially unable to do so.

Broader research confirms this. A 2024 study in the *Journal of Occupational Health Psychology* by researchers at the University of Zurich examined 840 professional creators across Europe and found that algorithmic platforms produced stress profiles comparable to emergency medical workers — not because of physical danger, but because of chronic absence of recovery time. The algorithm punishes absence. A creator who takes a two-week holiday can expect a 30 to 50 per cent decline in engagement upon return — "algorithmic decay," the study called it. The platform does not pause when you do. It replaces you.

The architecture matters. TikTok, Instagram, and YouTube do not employ creators. No contract, no minimum wage, no sick leave. Creators are, in gig-economy parlance, independent contractors — except that

even gig workers have some contractual relationship with the platform distributing their work. Creators have none. They are suppliers in a marketplace that can alter its terms unilaterally. When TikTok adjusted its algorithm in late 2023, reportedly deprioritising videos over sixty seconds, thousands of creators whose format relied on longer content saw their income collapse overnight. No announcement, no transition period, no appeal.

The psychological result is a population of workers simultaneously self-employed and utterly dependent on systems they neither control nor understand. The British psychotherapist Josh Cohen, writing in *Granta* in 2024, described this as "entrepreneurial helplessness" — the paradox of being nominally free while functionally captive.

And yet the supply of new creators grows. The allure is obvious: the handful of spectacular successes, the mythology of the ordinary person who becomes extraordinary through sheer creative will. What is less visible is the base rate. According to Goldman Sachs' 2024 report on the creator economy, fewer than 4 per cent of full-time creators earn more than the median household income in their respective countries. The median annual

earnings of a full-time creator in the United States were approximately \$15,000 — below the federal poverty line for a single-person household.

IV. Drowning in Sponsored Content

On the demand side — among audiences — a parallel exhaustion is setting in.

In 2007, the marketing research firm Yankelovich estimated that the average American encountered roughly 5,000 advertisements per day. By 2025, according to Red Crow Marketing, that figure had risen to between 6,000 and 10,000 — depending on how one defines "advertisement" when sponsored stories, product placement, and algorithmically inserted promotions blur the boundary between advertising and organic material until the distinction collapses.

The human perceptual system responds the only way it can: by tuning out. Banner blindness — the tendency to ignore banner advertisements — was first identified at Rice University in 1998. It has since generalised into

something the marketing literature calls "ad fatigue" but which is better described as comprehensive numbness to persuasion. A 2025 Hootsuite study found that 78 per cent of social media users scrolled past sponsored content without registering its message. Among users aged 18 to 24, the figure was 86 per cent.

This creates a doom loop. As audiences grow more resistant to advertising, advertisers demand more immersive, less detectable formats. Native advertising. Sponsored storytelling. Influencer partnerships designed to be indistinguishable from organic endorsements. Each escalation temporarily restores effectiveness — and then, as audiences adapt, accelerates the erosion of trust that caused the problem in the first place.

V. The Trust Deficit

Trust, in this context, is not a sentiment. It is an economic variable, and it is in freefall.

Morning Consult's January 2025 trust survey found that 67 per cent of Gen Z consumers (born 1997–2012) trusted user-generated content *less* than two years prior. User-generated content was supposed to be the antidote to corporate messaging — the genuine voice of real people, uncorrupted by institutional interests. If even that has lost credibility, the implications for attention-based advertising are severe.

The distrust is not irrational. Years of undisclosed sponsorships, manufactured authenticity, and algorithmically amplified misinformation have trained audiences to approach all online content with suspicion. The attention economy has educated its consumers — just not as intended. It has taught them that nothing online can be taken at face value.

In India, the tension between influencer culture and traditional media trust has produced a striking inversion. The country had an estimated 80 million Instagram creators in 2025, according to the marketing platform Koflence — the platform's largest market by creator count. The rise of "nano-influencers" — creators with fewer than 10,000 followers, often operating in regional languages on Reels and YouTube Shorts — has defined

Indian digital culture. They are valued because they appear local, relatable, and unpolished: the authenticity premium at its most granular.

And yet print newspaper subscriptions in India have been *rising*. The Registrar of Newspapers reported in 2024 that registered print publications had increased for the sixth consecutive year, reaching over 150,000 titles. Daily print readership stood at approximately 300 million — unimaginable in Western markets. The standard explanation is literacy growth and rising incomes in semi-urban and rural areas. But in a country where WhatsApp forwards have become synonymous with misinformation — a problem linked to lynchings and communal violence — the physical newspaper, with its masthead and editorial accountability, has acquired renewed authority. It is trusted not despite being old-fashioned, but *because* of it.

VI. When Attention Turns Dangerous

If manufactured authenticity is the attention economy's chronic illness, the pursuit of attention at any cost is its acute crisis.

South Korea's "molka" epidemic represents the darkest endpoint. *Molka* — from *mollae* (secret) and *camera* — refers to the hidden filming of people, overwhelmingly women, in intimate situations. Between 2015 and 2023, police investigated over 50,000 cases, according to the Korean Institute of Criminology. The footage is shared on forums, sold on encrypted platforms, and uploaded to pornographic websites where it generates advertising revenue.

The connection to the attention economy is structural, not incidental. A society organised around the capture and monetisation of human experience will, at its margins, produce actors who extend that logic beyond consent. The crisis prompted massive protests — in 2018, approximately 70,000 women marched in Seoul under the slogan "My life is not your porn" — and led to legislative reform, including increased penalties and a dedicated police unit. But the underlying incentive structure — attention as currency, content as commodity — remains intact.

VII. The Virtual Escape

Japan has produced the most philosophically interesting response to the authenticity crisis.

The country's VTuber (Virtual YouTuber) industry — performers who stream through animated digital avatars — generated an estimated \$1.5 billion in revenue in 2024, according to User Local. The largest agency, Hololive Production, manages over fifty performers whose real identities are unknown to their audiences. Gawr Gura, a Hololive English performer represented by an animated shark-girl avatar, surpassed four million YouTube subscribers before her 2024 hiatus — more than many television presenters.

The obvious reading is escapism, a retreat into cartoon fantasy. This is not wrong, but it is incomplete. VTubers offer, paradoxically, a form of honesty that human creators increasingly cannot. A VTuber makes no pretence of showing you their real life. No staged morning coffee, no "unfiltered" selfie run through three editing apps. The artifice is announced, upfront, non-negotiable. Within that declared artifice, the performer is free to be genuine. Freed from performing authenticity, VTubers often

display a spontaneity and directness that their human counterparts, trapped in the authenticity-industrial complex, cannot risk.

The audience grasps this. A 2024 Cross Marketing survey found that 38 per cent of regular VTuber viewers cited "honesty" as a primary reason for engagement — striking, given that the performers are hidden behind fictional characters. What they mean is not factual transparency but the absence of pretence. They would rather engage with a cartoon shark openly playing a character than a human influencer covertly playing one.

When audiences find animated avatars more trustworthy than human beings, something fundamental has shifted between performance and reality.

VIII. The Simulacra Have Won

Jean Baudrillard, in *Simulacra and Simulation* (1981), proposed a four-stage model for the relationship between representation and reality. First, the image reflects reality. Second, it masks and distorts reality. Third, it masks the

absence of reality. Fourth — "pure simulation" — the image has no relation to reality at all. It is its own referent.

Baudrillard was writing about Disneyland and television. He could not have imagined TikTok. But his framework maps onto the attention economy's trajectory with uncomfortable precision.

The polished Instagram era was Stage Two: images that distorted reality by presenting an idealised version of it. The authenticity era was Stage Three: images that masked the absence of genuine spontaneity, performing a rawness that concealed the fact that rawness had become impossible within the platform's incentive structure. And the current moment — in which AI-generated influencers coexist with human creators manufacturing authenticity coexist with VTubers offering honesty through acknowledged fiction — is arguably Stage Four. The copy of authenticity has not merely replaced the original; it has rendered the concept of an "original" incoherent.

This is not an abstract philosophical problem. It has concrete economic consequences.

When audiences cannot distinguish genuine recommendation from paid promotion, they stop trusting recommendations. When consumers cannot tell organic content from sponsored content, they discount all content. When voters cannot differentiate authentic political expression from coordinated campaigns, they disengage from political discourse. The collapse of authenticity is not merely an aesthetic concern. It is a collapse of the informational infrastructure on which markets, democracies, and institutions depend.

IX. The Exhaustion of Everything

There is a particular tiredness that belongs to our era. It is not the tiredness of labour — the aching muscles, the spent body that at least testifies to something accomplished. It is the tiredness of sorting. Of parsing. Of trying to determine, with every piece of content that appears on one's screen, whether it is genuine or manufactured, organic or sponsored, human or generated, sincere or strategic.

The philosopher Byung-Chul Han, in *The Burnout Society*, argues that the contemporary subject is exhausted not by external compulsion but by the imperative to perform. The attention economy extends this analysis: we are exhausted not only by our own performance but by the performance of others, and by the cognitive labour of evaluating those performances in an environment where the signals of authenticity have been systematically degraded.

This is the condition that Yuna Park's followers inhabit, though most could not name it. They watch her morning coffee video. Some believe it is real. Some suspect it is staged. Most exist in an intermediate state — neither believing nor disbelieving, simply consuming, scrolling onward, their attention extracted and monetised in the fraction of a second before the next video loads.

And Yuna herself? She is exhausted too. After the eleventh take, she eats breakfast — real breakfast, off camera, at the kitchen table she has pushed out of frame. The coffee has gone cold. She does not reheat it. She has, by 5 a.m., produced seven seconds of content, maintained the illusion of effortlessness, and contributed her small quantum of manufactured authenticity to a system that will process it, distribute it, and forget it by noon.

She has fourteen more videos to record today.

The authenticity crisis might be survivable if it were confined to consumer culture — to makeup tutorials and morning routines, to the frothy surface of digital entertainment. But the same dynamics that have corrupted the relationship between creators and audiences have infiltrated the institutions societies depend upon for shared truth: journalism, science, politics, education. The performance of credibility has begun to replace credibility itself — and unlike a staged coffee video, the consequences of that substitution cannot be scrolled past.

That institutional damage is where we turn next.

Chapter 4: The Collateral Damage

On the morning of 14 March 2022, Dr Katrin Meissner sat in her office at the University of Hamburg's Centre for Molecular Neurobiology, refreshing her browser with the compulsion of someone who knows the numbers will keep climbing but cannot look away. Her team's paper — on epigenetic markers in aggressive breast cancers, published the previous week in *Nature* — had crossed 1.2 million views on Twitter. By that afternoon, it would reach 2.3 million. *The Guardian* had run a feature. *Le Monde* had phoned. *The Hindu* had requested a video interview.

Meissner, who had spent eleven years in molecular biology publishing papers that typically reached a few hundred specialists, was suddenly a public figure. Her university's press office issued three statements in a week.

A German television crew filmed in her laboratory. Colleagues in adjacent fields began citing the work before they had finished reading it.

Seven months later, a postdoctoral researcher at the Karolinska Institute in Stockholm identified a methodological error in the statistical analysis — a mis-specified covariate that, once corrected, dissolved the paper's central finding. *Nature* published a retraction notice. It was, by the journal's standards, a thorough and transparent process.

The retraction received 23,000 views. One per cent of the original audience learned the truth.

By the time I spoke with Meissner in early 2025, she had tracked forty-seven subsequent papers that cited her retracted work as though it remained valid. "The internet has a magnificent memory for sensation," she told me, "and almost none for correction."

This chapter is about that asymmetry — and the institutional wreckage it leaves behind.

The Replication Crisis, Accelerated

Science was in trouble before the attention economy found it. In 2005, John Ioannidis published what would become one of the most-cited papers in medical history: "Why Most Published Research Findings Are False." His argument was statistical and structural. Most studies were underpowered. Researchers had perverse incentives to find positive results. Journals had perverse incentives to publish them. The entire apparatus of scientific publishing was optimised to produce confident-sounding claims from uncertain data.

Ioannidis described a system that rewarded novelty over rigour. He could not have anticipated a technology that would amplify that reward a thousandfold.

The trouble is not merely that attention metrics have entered the academy — they have become proxies for value. Journal impact factors, a crude measure of citation frequency, began to function as academic credit scores. Hiring committees used them. Funding bodies used them. Tenure decisions rested on them. Social media added a second, more volatile layer: altmetrics. A paper's online mentions, shares, and media coverage became part of its

permanent record, visible on Altmetric.com, where a colourful badge was appended to each article like a digital medal.

The incentive structure was clear. A surprising finding attracts journalists, who write accessible articles, which generate social media shares, which boost altmetric scores, which influence hiring committees. A replication study — even a successful one — does none of these things. By the attention economy's reckoning, it is worthless.

A 2023 *PLOS Biology* analysis found that papers with high altmetric scores were significantly more likely to report surprising or counterintuitive results — and more likely to be subsequently challenged or corrected. The correlation was not subtle. The attention economy's selection mechanism was filtering for unreliability.

Here is the fact that should unsettle everyone. Retracted papers continue to be cited an average of 4.3 years after retraction. Not by cranks. By working scientists in mainstream publications who either do not know the paper has been retracted or do not check. A 2022 *Scientometrics* study traced over 10,000 retracted papers and found that more than half continued accumulating

citations long after withdrawal. The attention economy has no correction mechanism. It has an amplification mechanism. These are not the same thing.

The COVID-19 pandemic made the problem impossible to ignore. Preprint servers — where researchers post findings before peer review — became the primary channels for urgent scientific communication. A novel virus demanded rapid sharing; in practice, unvetted claims reached the public at the speed of Twitter. A French microbiologist's small, poorly controlled hydroxychloroquine study was viewed millions of times before its flaws were identified. By then, a sitting American president had endorsed the drug on national television. The preprint had done its work. The correction arrived, as corrections always do, to an empty room.

The Newsroom Collapses

If science's vulnerability to the attention economy is structural, journalism's is existential. Both depend on public trust and a business model the internet has

dismantled. But journalism's collapse has been faster, more visible, and more consequential for democratic life.

Between 2005 and 2025, the United States lost more than 2,900 newspapers. Not journalists — entire publications. The *Rocky Mountain News* in Denver, founded in 1859, closed in 2009. The *Seattle Post-Intelligencer* abandoned print the same year. Hundreds of smaller papers — the *Youngstown Vindicator*, the *Paducah Sun*, the *Trentonian* — either folded or were hollowed out by hedge-fund owners who bought them for their real estate and presses.

The United Kingdom followed suit. More than 300 local newspapers closed between 2005 and 2024, according to Press Gazette. Entire regions — the Midlands, the Welsh Valleys, the north-east — became "news deserts": areas with no local journalist covering council meetings, court proceedings, or school board decisions. A 2023 Reuters Institute study found residents of these areas were less likely to vote locally and more likely to report feeling disconnected from their communities.

The mechanism was straightforward. Newspapers had sustained themselves on classified advertising — property, jobs, cars, personals. Between 2000 and 2010, that revenue migrated to Craigslist, Rightmove, Indeed, and

their equivalents. Display advertising increasingly flowed to Google and Facebook, whose targeting no local paper could match. Between 2010 and 2020, the two platforms captured an estimated 60 per cent of all US digital advertising revenue. Newspapers were not competing; they were being bypassed.

The survivors adapted by chasing attention metrics. Click-driven newsrooms — a phrase that would have been meaningless in 2000 — became the norm. Editors measured success in pageviews. Headlines were A/B tested, optimised for engagement rather than accuracy. The distinction between what mattered and what performed became a source of daily institutional anguish.

India's experience offers a revealing counterpoint. The country gained a generation of impressive digital-native outlets — *The Wire*, *Scroll.in*, *Newslaundry*, *The News Minute* — staffed by journalists who rejected the compromised practices of legacy media. Their reporting was often superb. But they faced the same structural problem: advertising revenue flowed to platforms, not publishers. *Scroll.in* relied heavily on institutional grants. *Newslaundry* pioneered a subscription model that was, by its own admission, a constant struggle. *The Wire*, which broke several significant investigative stories, spent

substantial resources defending itself against legal challenges it might have weathered more easily with a larger revenue base.

The pattern was global. Quality journalism persisted — it always does, through the sheer stubbornness of people who believe it matters — but the economic scaffolding that had supported it was gone.

The Attention Arbitrage of Political Actors

No group understood the attention economy's incentives faster than political operatives. While journalists and scientists debated whether Twitter was a useful professional tool, political actors had already reverse-engineered its logic and turned it into a governing strategy.

The principle was embarrassingly simple: outrage generates more attention than governance. A provocative statement at 6 a.m. could dominate the news cycle all day, displacing coverage of policy decisions or institutional

corruption. The algorithms, designed to maximise engagement, treated political outrage as premium content.

Donald Trump's Twitter strategy was the most analysed example, though neither the first nor the most sophisticated. His team discovered that the platform rewarded escalation. Each tweet generating fury or mockery became the centre of a media vortex that consumed the oxygen for any other story. His opponents spent four years responding to provocations rather than advancing alternatives. Trump received an estimated \$5 billion in free media coverage during the 2016 campaign alone, according to *mediaQuant*.

India's Bharatiya Janata Party built a more systematic apparatus. Its IT cell, overseen by Amit Malviya, ran what researchers described as the most extensive political social media operation in any democracy. Thousands of volunteers and paid operatives amplified messaging across WhatsApp, Twitter, Facebook, and ShareChat, coordinating to trend specific hashtags at scheduled times. The strategy was attention dominance — ensuring the party's framing reached citizens before any competing account could.

In Brazil, Bolsonaro's network operated primarily on WhatsApp, exploiting end-to-end encryption to circulate claims that could not be fact-checked or traced. A 2019 *Folha de S.Paulo* investigation revealed that sympathetic companies had spent millions of reais on bulk messaging — a violation of Brazilian election law that was almost impossible to prosecute given WhatsApp's architecture.

Saudi Arabia's troll farms, documented by the *New York Times* and the Citizen Lab at Toronto, represented state-sponsored attention manipulation designed not to win arguments but to drown out dissent. Critics of Mohammed bin Salman were met with coordinated floods of abuse and harassment designed to make public speech so unpleasant that silence became the rational choice.

Each of these operations discovered the same structural truth: in an attention economy, you do not need to be right. You need to be loud.

The Myth of the Shrinking Attention Span

In 2015, a report attributed to Microsoft Canada claimed that the average human attention span had fallen to eight seconds — less than a goldfish's. The statistic appeared in thousands of articles, conference presentations, and marketing decks, cited as proof that the internet was rewiring our brains.

It was nonsense.

The BBC's *More or Less* programme traced the claim and found it rested on no credible research. The goldfish figure derived from a misreading of a dubious study. The human figure was unsupported by any methodology. Microsoft's own researchers could not identify the data's origin. The statistic had, fittingly, gone viral because it was surprising and shareable — the very dynamics it purported to describe.

The reality is more mundane and more interesting. Humans have not lost the capacity to focus. A reader who cannot concentrate on a 3,000-word article can watch a four-hour film or play *Elden Ring* for six hours without a bathroom break. A student who struggles with a textbook chapter can follow the intricate threads of a 70-hour television series. The issue is not neurological degradation. It is competition.

In 1980, a person sitting down to read a book competed, for attention, with a few television channels, a radio, and whatever was visible through the window. In 2025, that same person competes with the entire output of human civilisation, algorithmically sorted and personally targeted, available through a device in their pocket. The book has not changed. The environment has.

This distinction matters for education. The "shrinking attention span" narrative has justified shorter readings, more video, "micro-learning" modules, gamified assessments — changes that may be precisely wrong. If the problem is competition rather than capacity, the solution is not shorter education but more compelling education, or less competition. Schools in France and Denmark that banned smartphones during hours reported improvements in focus and performance — nothing to do with attention spans, everything to do with removing a device engineered to be maximally distracting.

India's WhatsApp Crisis

In the summer of 2018, a series of mob lynchings swept across India. In Dhule, Maharashtra. In Bidar, Karnataka. In Rainpada, Maharashtra. In Karbi Anglong, Assam. The pattern was consistent and horrifying: strangers — often migrant workers, mendicants, or simply people who looked unfamiliar — were attacked by mobs who had received WhatsApp messages warning of child kidnappers in the area. At least thirty-three people were killed between May 2017 and July 2019, according to *IndiaSpend*.

The messages were fabrications. Many included grainy videos, repurposed from unrelated incidents, showing children being led away by adults. Others were text-only, written in the breathless style of urgent warning: "BEWARE! Kidnappers operating in this area! Share with all groups!" The forwarding chains were long. The messages arrived from trusted contacts — family members, neighbours, members of local WhatsApp groups formed around temples, schools, or housing societies. Each forward carried an implicit endorsement.

WhatsApp responded by limiting forwarding to five groups at a time and marking frequently forwarded messages with a double-arrow label. These measures reduced virality but did not address the fundamental

problem. The platform had become India's primary information infrastructure — used by more than 500 million people for family coordination, business, and news — without the editorial structures, fact-checking, or institutional accountability that had traditionally mediated public information.

The government demanded traceability — which would have meant breaking end-to-end encryption. WhatsApp resisted, citing privacy. The impasse persisted for years, a knot of competing values — privacy, safety, free expression, state power — that no single actor could untangle.

What the WhatsApp crisis revealed, at the cost of human life, was that the attention economy's damage was not confined to wealthy nations with strong institutional buffers. In countries where traditional media infrastructure was thinner, where digital literacy was unevenly distributed, and where communal tensions required only a spark, the consequences were not measured in declining newspaper revenues or distorted academic incentives. They were measured in bodies.

China's Split Attention

China presents the most instructive case of what happens when a state attempts to manage the attention economy rather than simply participate in it.

The country operates two dominant platforms serving fundamentally different functions. Douyin — TikTok's Chinese original — is an entertainment engine: short-form video, algorithmically curated, designed to maximise time spent. It is, by most measures, the most effective attention-capture technology ever built. Its algorithm, reportedly evaluating more than a thousand signals per interaction, can identify a viewer's interests within minutes.

WeChat is an information and communication platform: messaging, long-form articles, payments, government services, and an ecosystem of "mini-programs" functioning as apps within the app. Its content distributes through personal networks rather than algorithmic recommendation, giving it a fundamentally different attention profile.

The result is a bifurcated attention culture unlike anything in the West. Chinese users spend their entertainment time on Douyin and their informational time on WeChat, and the two modes rarely overlap. A 2024 study by the China Internet Network Information Center found that average daily Douyin usage had reached 118 minutes — nearly two hours — while WeChat article reading averaged 23 minutes per day. The entertainment layer consumed five times as much attention as the informational layer.

The Chinese government has intervened actively. In 2021, it imposed time limits on Douyin for users under fourteen — forty minutes per day — and mandated that the platform's youth mode show only "educational" content. The regulations were part of a broader campaign against what state media called "spiritual opium," a phrase linking digital entertainment to a history of foreign-imposed addiction. Whether these interventions were motivated by genuine concern for public welfare or by the party's desire to control the narrative — or, most likely, both — China was the only major country attempting to structurally reshape its attention economy through state power.

The experiment's long-term results remain uncertain. But its existence demonstrates that the attention economy is not a natural law. It is an engineered system. And engineered systems can, in principle, be re-engineered.

Africa's Mobile-First Reckoning

In much of sub-Saharan Africa, the attention economy arrived without preamble.

There was no era of desktop internet, no decade of blogging, no gradual migration from print to digital. For hundreds of millions, the internet arrived fully formed as an affordable smartphone and a data plan. An extraordinary liberation — access to information, commerce, and communication that previous generations could not have imagined. Also an arrival without context.

Kenya's M-Pesa, launched by Safaricom in 2007, became the world's most successful mobile money platform, processing transactions equivalent to roughly half the country's GDP. It demonstrated that mobile technology could leapfrog decades of institutional development. But

it also established a precedent: in African mobile markets, attention and commerce were intertwined from the start. The phone was not merely a device for consuming content; it was a wallet, a bank, a marketplace. Every moment of attention had immediate transactional potential.

Nigeria's experience was shaped by MTN, whose aggressive expansion made it the dominant carrier in Africa's most populous nation. MTN's pricing — offering cheap or free access to platforms like Facebook Free Basics — determined what the attention economy looked like for millions of first-time users. If your only affordable internet was Facebook, then Facebook was the internet. The implications for information diversity and political discourse were profound.

In Ethiopia, the government's repeated internet shutdowns — during political crises in 2016, 2019, and 2020 — demonstrated the fragility of attention economies built on mobile infrastructure. When the state controlled the telecommunications network, it could sever the entire information ecosystem with a single decision. The attention economy, for all its decentralised rhetoric, depended on centralised infrastructure.

Across the continent, the pattern held: mobile technology delivered the attention economy's benefits and pathologies simultaneously, without the transitional period that allowed Western societies — however imperfectly — to develop some defences. Social media arrived in countries with young populations, thin institutional structures, and recent histories of communal conflict. The combination was volatile.

The Common Thread

The examples in this chapter span continents, languages, political systems, and levels of economic development. A molecular biologist in Hamburg. Mob violence in rural Maharashtra. A hollowed-out newspaper in Kentucky. A state-mandated screen time limit in Beijing. A mobile money transaction in Nairobi. They appear, at first glance, to share nothing.

But they share everything. Each is a story of an institution — science, journalism, democracy, education, public safety — built for a world where information was scarce and attention the default. Libraries existed because books

were expensive. Newspapers because reporting required resources most people lacked. Peer review because scientific claims needed gatekeeping before reaching the public. These institutions were not perfect — often exclusionary, slow, resistant to change. But they mediated between information and attention, ensuring that not everything that could be said would be heard with equal force.

The attention economy removed that mediation. It did not replace it with something better. It simply removed it, and then monetised the chaos.

The damage documented here is not hypothetical, nor confined to edge cases. It is structural, ongoing, and compounding. Retracted papers cited for years. News deserts spreading across democracies. Political actors who have discovered that governing is optional if you can dominate the feed. Students blamed for neurological deficits they do not have. Mobs mobilised by fabrications. Entire countries arriving at the digital commons without a map.

And yet the damage has not gone unresisted. In Vermont and Seoul, in Bangalore and Copenhagen, in converted barns and government ministries, a counter-movement has been forming. It is uncoordinated, ideologically

diverse, and often internally contradictory. It includes Luddites and technologists, monks and venture capitalists, parents who ban screens and parents who teach their children to code.

Whether it will amount to anything more than a footnote in the attention economy's long triumph is the question we turn to next.

Chapter 5: The Resistance

The barn smells of woodsmoke and beeswax. Built in 1847 to store hay for a dairy operation that supplied Stowe village through Vermont's brutal winters, it retains the proportions of that era — a cathedral ceiling of hand-hewn beams, walls of mortise-and-tenon joinery that predate the telegraph. On this Friday evening in October 2025, twelve people sit in a rough circle on cushions and low benches, their faces lit by hurricane lanterns and the amber glow of a wood-burning stove. Outside, the sugar maples are turning the colour of rust and blood.

They have come from six countries. A venture capitalist from San Francisco who manages a \$340 million fund. A product manager from Bangalore who works on recommendation algorithms for one of India's largest e-commerce platforms. A secondary school teacher from Manchester who has not gone a waking hour without checking her phone in four years. A marketing executive from São Paulo whose agency runs social media

campaigns across Latin America. A retired psychiatrist from Montreal, a documentary filmmaker from Lagos, a sommelier from Lyon, a junior doctor from Seoul. Each has paid \$2,400 for three days in this barn, plus accommodation in a neighbouring farmhouse with aggressively plain rooms — no televisions, no Wi-Fi, no clocks.

Their phones are sealed in grey neoprene pouches locked with magnetic clasps, the kind used at comedy shows to prevent recording. The organisers — a married couple, both former Google employees — will unlock the pouches on Sunday afternoon. In the meantime, there are books, conversation, walks, cooking, and silence. No agenda beyond the absence of one.

The organisers call it a "cognitive retreat." The participants, when they describe it to friends afterwards, use a different word. They call it sanity.

The resistance to the attention economy has become, with the inevitability that attends all things in a market society, its own economy.

The global digital detox industry was valued at roughly \$1.6 billion in 2025, encompassing luxury retreats like the one in Vermont, smartphone-blocking apps, distraction-free writing tools, analogue watches marketed to former smartwatch users, and "dumb phone" handsets stripped of browsers and social media. Camp Grounded, which began in 2013 as a summer camp for adults in Anderson Valley, California — no phones, no real names, no discussion of work — was regarded at the time as charmingly eccentric. By 2025, it had spawned dozens of imitators across four continents, and the mockery had given way to envy.

The trajectory is instructive. What begins as fringe becomes niche, what is niche becomes mainstream, and what is mainstream becomes an industry with its own trade conferences, certification programmes, and venture capital. The attention economy produces the disease; the detox economy sells the cure. Neither has much incentive to address the underlying condition.

But to dismiss the resistance as merely commercial would be dishonest. Behind the retreats and the apps and the dumb phones lies a genuine intellectual tradition — a body of thought that has been building, quietly and then not quietly at all, for the better part of a decade.

In 2019, Jenny Odell published *How to Do Nothing: Resisting the Attention Economy*. The title was slightly misleading; the book was not an instruction manual for idleness but a philosophical argument for redirecting attention away from the extractive logic of social media and towards the immediate, the local, the ecological. Odell, an artist and Stanford lecturer, drew on birdwatching, public parks, bioregionalism, and Hannah Arendt to argue that doing "nothing" — by the metrics of the attention economy — was in fact doing the most important thing: reclaiming one's capacity to notice the world.

The book became a bestseller. It spent weeks on the *New York Times* list and was translated into more than twenty languages. It was discussed in the very ecosystems it critiqued, shared in Instagram stories, summarised in Twitter threads, excerpted in newsletters optimised for open rates. The irony was not lost on Odell, who acknowledged it with the weary grace of someone who understood that the attention economy's most insidious quality is its ability to absorb its own critiques. You cannot promote a book about resisting the attention economy without participating in it.

What Odell achieved, however, was a shift in register. Before *How to Do Nothing*, the dominant language around technology overuse was therapeutic: addiction, compulsion, dopamine, withdrawal. Odell reframed it as political. The problem was not that individuals were weak; an economic system had been built to exploit a cognitive vulnerability for profit. This was not a matter for self-help. It was a matter for collective action.

The distinction mattered enormously, even if the collective action proved difficult to organise.

Cal Newport arrived at similar conclusions from a different direction. A computer science professor at Georgetown, Newport had been writing about focus and productivity since 2012, but *Deep Work: Rules for Focused Success in a Distracted World* (2016) established him as the attention economy's most rigorous domestic critic. His central argument was deceptively simple: the ability to concentrate without distraction on a cognitively demanding task is becoming both rarer and more valuable. In an economy where most knowledge workers spend their days switching between email, Slack, spreadsheets, and social media — Newport cited research suggesting the average office worker checks their inbox

every six minutes — the capacity for sustained, uninterrupted thought is a competitive advantage of the first order.

Digital Minimalism (2019) extended the argument from the professional to the personal. Newport proposed a philosophy of technology use based on intention rather than default: choose your tools deliberately, optimise them for the things you value, discard everything else. He advocated long walks, physical hobbies, face-to-face conversation, and extended periods away from screens — not as Luddism, but as strategy.

Newport himself did not use social media. He had never had a Facebook, Twitter, or Instagram account. For years, this was treated as an eccentricity, the kind of thing interviewers raised with gentle scepticism, as though he had announced he did not own shoes. By 2024, it looked less like eccentricity and more like prescience. He had not predicted the specific pathologies — the adolescent mental health crisis, the epistemic collapse, the political polarisation — but he had identified the underlying mechanism. Attention is a skill. Like any skill, it atrophies without practice. And the attention economy is designed, with extraordinary precision, to prevent that practice.

His phrase "deep work" entered the lexicon in a way that few academic concepts do. Managers invoked it to block out meeting-free mornings. Writers cited it when they deleted Twitter. It was, in its way, as influential as any Silicon Valley product launch — and considerably more durable.

If Odell and Newport represented the intellectual wing of the resistance, the technology companies offered something more ambiguous: tools to manage the very problem they had created.

In September 2018, Apple introduced Screen Time as part of iOS 12, allowing users to monitor daily usage, set time limits for specific apps, and schedule "Downtime" periods. Google followed almost simultaneously with Digital Wellbeing for Android, offering a usage dashboard and a "Wind Down" mode that drained the screen of colour as bedtime approached.

The announcements drew a mixture of applause and suspicion. Here were the architects of the attention economy acknowledging, implicitly, that the edifice had a

problem. It was, as more than one commentator observed, the foxes offering the hens a pamphlet on predator awareness.

The suspicion proved well-founded. Studies conducted between 2019 and 2024 consistently found that Screen Time and Digital Wellbeing had negligible impact on usage patterns. A 2022 analysis in *Nature Human Behaviour*, examining data from over 10,000 participants across eight countries, concluded that digital wellbeing tools reduced average daily screen time by less than four minutes. Users set limits, then overrode them. They reviewed their dashboards, felt a momentary pang, and carried on. The tools assumed the problem was informational — that people didn't *know* how much they used their phones. But the problem was structural. The apps on the other side of those time limits had been engineered by some of the most talented designers in the world, spending billions annually, to be as difficult to put down as possible. A countdown timer was not a fair fight.

Apple and Google continued to refine the tools. But neither company made the obvious move: redesigning the products themselves to be less compulsive. The business

model would not permit it. Engagement drives advertising revenue. Advertising revenue drives share price. Share price drives everything.

Other countries took the problem more seriously — or at least more literally.

South Korea, which has one of the highest rates of internet penetration in the world, began treating internet overuse as a public health issue in the early 2010s. The Ministry of Science and ICT has spent roughly \$18 million since 2019 funding digital detox camps for young people — one- to two-week programmes combining outdoor activities, group therapy, art workshops, and structured periods without devices. Over 10,000 young people have attended.

The figure bears repeating: a national government, one of the most technologically advanced in the world, spending tens of millions of dollars to help its young citizens *stop using technology*. South Korea frames this not as individual weakness but as an environmental health issue, no different from air pollution or contaminated water. The technology is the environment. The overuse is the contamination. The camps are the treatment.

Results have been mixed, as public health interventions usually are. Participants report short-term improvements in sleep, mood, and concentration. Longer-term follow-up data is sparser, and what exists suggests a familiar pattern: gains that fade as participants return to the environment that produced the problem. A teenager who spends two weeks in the mountains without a smartphone will sleep better. A teenager who then returns to a school system where homework is assigned via app and social standing is mediated by Instagram will, in most cases, revert. The intervention addresses the individual. The environment remains unchanged.

Japan's response has been characteristically different — less programmatic, more cultural. *Shinrin-yoku*, or forest bathing, dates formally to the 1980s, when the Ministry of Agriculture, Forestry and Fisheries coined the term. Peer-reviewed studies have since documented reductions in cortisol, blood pressure, and heart rate among forest bathers, along with increases in natural killer cell activity. But for decades, *shinrin-yoku* was simply what Japanese people — particularly older Japanese people — did. It was not a product. It was a practice.

What changed was the clientele. By the mid-2020s, specialised *shinrin-yoku* retreats had emerged in the mountains around Tokyo and Kyoto, marketed to technology workers. The programmes combined forest walking with digital fasting, mindfulness instruction, and elaborate multi-course meals served in near-total silence. The waiting lists were long. Here was a peculiarly modern absurdity: young engineers from the most advanced technology companies in the world paying substantial sums for an experience their grandparents had enjoyed for free, in the same forests, without needing to call it anything at all.

In India, the resistance took a more communal form. Bangalore's technology corridor — home to Infosys, Wipro, Tata Consultancy Services, and the Indian operations of virtually every major American technology firm — saw, from around 2022 onwards, a growing movement of IT professionals organising "digital sabbaths": screen-free weekends structured around group activities. Some groups gathered in parks for conversation and board games. Others organised cycling trips into the

countryside south of the city. A few, influenced by India's tradition of silent retreat, spent their sabbaths at *ashrams* near Mysore.

The Wipro Foundation's "Unplug" programme, launched in 2023, formalised some of these impulses. It offered employees structured weekend retreats, subsidised at company expense, combining digital fasting with workshops on attention management, sleep hygiene, and meditation in the Vipassana tradition. The programme was small, reaching perhaps a few thousand employees in its first two years, but its significance lay less in scale than in provenance. Here was one of India's largest technology services companies — a company whose revenues depend entirely on the digital infrastructure of the attention economy — quietly acknowledging that its own workers needed periodic rescue from the thing they were building.

As one Wipro engineer told me during a reporting trip to Bangalore in early 2025: "We spend Monday to Friday making the systems more engaging, more sticky, more impossible to leave. Then on Saturday we go to a retreat to recover from engagement." He paused. "I don't think we talk about this contradiction enough."

The corporate resistance, where it existed, was dramatic in gesture if uncertain in impact.

Basecamp — the project management company founded by Jason Fried and David Heinemeier Hansson — had long positioned itself as a counterweight to Silicon Valley's growth-at-all-costs ethos. In 2021, the company removed social media metrics from its internal dashboards. Employees could no longer see follower counts, post likes, or content "performance" by the usual measures. "We were making decisions based on what got attention," Fried wrote, "not what was good."

The move attracted admiration, some scepticism, and little imitation.

Patagonia's decision was bolder. In 2020, the outdoor clothing company pulled all paid advertising from Facebook, citing concerns about hate speech and misinformation. The boycott, initially framed as temporary, became permanent. Patagonia never returned. Revenue continued to grow.

Lush Cosmetics went further still. In November 2021, the British retailer shut down its accounts on Instagram, Facebook, TikTok, and Snapchat entirely. Its chief digital officer, Jack Constantine, described the platforms as "not

a safe place for our customers" and compared the decision to refusing to open a shop on a street where customers were routinely harassed.

Whether these gestures constituted genuine resistance or savvy brand positioning — Patagonia and Lush both serve demographics that reward conspicuous ethical consumption — admits no clean answer. Probably both. The attention economy is capacious enough to accommodate even its own rejection as a marketable stance.

And this is the central difficulty. The resistance, in all its forms — intellectual, personal, corporate, governmental — confronts a problem that individual action cannot solve. You can delete Instagram. You can lock your phone in a magnetic pouch. You can spend a weekend in an *ashram* or a converted barn in Vermont. But your child's school still communicates via a WhatsApp group you cannot leave without missing essential information about closures, homework, and head lice. Your employer still expects you on Slack, where a green dot beside your name signals availability and its absence signals something close to dereliction. Your elderly parent still shares family news on Facebook because that is where the family is.

The attention economy is no longer an optional layer atop ordinary life. It is infrastructure. Opting out in 2026 is not like cancelling a magazine subscription. It is like opting out of the postal system in 1926, or the telephone network in 1966. It can be done, but the cost is measured in information missed, relationships strained, and opportunities quietly foreclosed.

This is why the most honest participants in the resistance — Odell, Newport, the Vermont retreat organisers — stop short of recommending wholesale withdrawal. Their counsel is not *leave* but *renegotiate*. Not *reject* but *choose*. They advocate for consciousness, intentionality, the examined digital life. These are valuable things, genuinely so. But they are available primarily to people with the education to analyse their own technology use, the financial security to risk missing a work email, and the social capital to explain, without penalty, why they are unreachable on a Saturday afternoon. The venture capitalist can afford a \$2,400 weekend without screens. The delivery driver whose income depends on a smartphone app cannot afford an hour.

The resistance has a class problem. And class problems, left unaddressed, have a way of becoming the entire problem.

There is a historical parallel that clarifies, if it does not resolve, the dilemma.

The temperance movement of the nineteenth century began as individual moral reform: people pledging, one by one, to abstain from alcohol. The early societies emphasised personal responsibility, self-improvement, the cultivation of willpower. They held meetings. They signed pledges. They supported one another in sobriety. And for a time, this seemed sufficient. If the problem was drink, the solution was not drinking.

But as the movement matured, its more perceptive advocates recognised that individual abstinence could not address the structural conditions that produced widespread abuse: poverty, industrial working conditions, the economic power of the brewing industry, the absence of alternative social institutions. A factory worker in Manchester in 1850 did not drink because he lacked willpower. He drank because the pub was the only warm room available to him, the only place of social connection after twelve hours of labour, and because cheap gin was more accessible than clean water.

The temperance movement's lasting achievements were structural: licensing laws, public health campaigns, regulation of the alcohol industry, the creation of alternative social spaces. The pledges mattered less than the policies. The personal was necessary but insufficient. What was required was a renegotiation of the economic and social arrangements that made the problem endemic.

The attention economy's resisters stand today roughly where the temperance advocates stood in the 1870s. They have identified the problem. They have developed personal practices of mitigation. They have built communities of mutual support. They have even achieved some corporate and governmental recognition. What they have not yet produced is a viable alternative to the economic model that makes the problem structural.

Individual resistance is the beginning of the story. It is not the end.

Jenny Odell knew this. In the final pages of *How to Do Nothing*, she acknowledged that the book's argument could be read as a counsel of privilege — a tenured professor telling people to go birdwatching. She resisted that reading, arguing that attention was a political act, that choosing where to direct one's notice was a form of power available to everyone. She was right, in theory. In

practice, the attention economy does not distribute its burdens equally, and the capacity to resist is not evenly shared.

Cal Newport knew it too. His later work increasingly emphasised institutional reform: the redesign of workplaces, the restructuring of communication norms, the creation of environments in which deep work was the default rather than the exception. He had begun, by the mid-2020s, to sound less like a self-help author and more like a labour organiser — albeit one with tenure and a bestseller list.

The twelve people in the Vermont barn will leave on Sunday afternoon. Their phones will be unlocked. The magnetic clasps will release with a soft click. Most will check their messages immediately, unable to resist even after seventy-two hours of practice in resisting. They will feel a brief, vertiginous guilt, and then the guilt will dissolve into the familiar current of notifications, and Monday will arrive, and the experiment will become an anecdote — something to mention at dinner parties, a credential of consciousness.

But something will have shifted. The venture capitalist will return to San Francisco and begin asking different questions of the startups that pitch him. The teacher from Manchester will institute a phone-free first period in her classroom — not because she believes it will change the world, but because she has seen what her own mind does when given thirty minutes of quiet. The product manager from Bangalore will return to his recommendation algorithms with a discomfort he cannot quite articulate, a sense that the metrics he optimises all day might be measuring the wrong things.

These are small changes. They are insufficient. They do not alter the incentive structures, the revenue models, the fiduciary obligations, the quarterly earnings calls, the vast machinery of extraction that turns human attention into advertising dollars. The barn in Vermont is a beautiful place, and the silence within it is real, and it is not enough.

The question is not whether resistance is possible. It plainly is. Millions of people, in dozens of countries, are resisting in ways large and small, personal and political, commercial and sincere. The question is whether resistance alone — even at scale — can produce the structural transformation that the moment requires.

Whether the temperance pledge can become the licensing law. Whether the individual "no" can become the collective "instead."

That question — what economic model replaces the attention economy when the attention economy fails — is not one the resistance movement has yet answered. But others are beginning to try.

Chapter 6: The Economics of What Comes Next

Wanjiru Kamau writes a weekly newsletter from her flat in Nairobi about African technology policy. She has five hundred paying subscribers, each contributing eight dollars a month. That is four thousand dollars monthly — roughly six times the median Kenyan salary, and considerably more than she ever earned from five hundred thousand social media followers spread across three platforms. Her Instagram once reached a quarter of a million people with a single post about Kenya's data protection bill. She made nothing from it. Her YouTube channel, where she posted explainer videos with careful motion graphics, generated eleven dollars in ad revenue over eight months. "I was performing for robots," she told me, "and the robots weren't even paying."

When she launched her newsletter in 2024 — first on Substack, later on Ghost, where she could own her subscriber list — she expected fifty readers. Within six months, she had two hundred. Within a year, five hundred. The growth was slow by the standards of virality. It was also, by every measure that matters, transformative. She quit her communications job at a fintech startup. She hired a research assistant. She bought her mother a new water heater.

When I asked about her "strategy," she laughed. "I stopped performing for algorithms and started writing for people who care."

Kamau's story is not exceptional. It is becoming ordinary. And in that ordinariness lies the most significant economic reorientation the internet has seen since Google sold its first advertisement in 2000.

From Attention to Intention

For twenty-five years, the dominant economic logic of the internet rested on a single transaction: someone sees something, and an advertiser pays for the seeing. Attention was the commodity. It could be measured in impressions, clicks, time-on-page, scroll depth, video completions. The entire architecture of the commercial web — from Facebook's News Feed to YouTube's recommendation engine to Google's search results page — was engineered to harvest this commodity as efficiently as possible.

The shift now underway is not a refinement of this model. It is its inversion.

Attention is passive. You did not ask to see that banner advertisement; it appeared in your peripheral vision as you read about the weather. You did not seek out that pre-roll video; it stood between you and the recipe you wanted. The attention economy was built on interception — inserting commercial messages into the gaps between what people actually wanted to do.

Intention is active. You sought something out. You chose to subscribe. You decided to pay. The intentional consumer is not a pair of eyeballs to be harvested. She is a person making a decision.

This distinction — passive attention versus active intention — might seem semantic. It is not. It rewires the entire economic relationship between creators and audiences. In the attention economy, the customer was the advertiser; the audience was the product. In the intention economy, the audience is the customer. Full stop. When the audience is the customer, the incentives shift radically: instead of maximising time spent, you maximise value delivered. Instead of engineering compulsive engagement, you earn deliberate trust.

The advertising industry generated an estimated \$626 billion globally in 2025, built almost entirely on the first model. What is emerging is not yet its equal in raw revenue. But it is growing in precisely the places where advertising is failing.

The Paywall Pivot

Consider the transformation of *The New York Times*. In 2012, the institution was closer to bankruptcy than most of its readers realised. Print advertising revenue had

collapsed. Digital advertising paid pennies on the dollar. The paper's market capitalisation had fallen below that of a mid-tier technology startup.

Then, under Mark Thompson, the *Times* did something the industry considered suicidal: it erected a hard paywall. Not a metered sample. Not a porous suggestion. A wall.

The industry watched with the morbid curiosity one reserves for bridge jumpers. The consensus was that readers, accustomed to free online news, would go elsewhere. Some did. But enough stayed — and enough new subscribers arrived, drawn by the signal that the *Times* considered its journalism worth paying for — that by 2023, the paper had surpassed ten million digital subscribers. Subscription revenue overtook advertising revenue for the first time in the paper's history. The share price more than tripled.

The *Times* was not alone, merely the most visible. *The Information*, a technology publication launched in 2013 by former *Wall Street Journal* reporter Jessica Lessin, charged \$400 per year from day one. It has never run advertisements. It has never chased virality. It has, however, built a readership of senior technology executives who consider it indispensable. Lessin once told

an interviewer that the most important decision she made was not about journalism but about economics: "We decided our reader was our customer, not our product."

In India, *The Ken* followed a similar path, launching in 2016 with a subscription model many considered untenable in a market where free content was the overwhelming norm. Priced at roughly 2,500 rupees per year — not trivial in India, where a cup of chai from a street vendor costs ten rupees — *The Ken* bet that a small audience willing to pay would prove more sustainable than a large audience that advertisers undervalued. By 2024, it had expanded from Bangalore to Southeast Asia, covering technology, business, healthcare, and education with a reporting depth that ad-supported competitors could not match.

Ben Thompson's *Stratechery*, perhaps the purest expression of the model, has operated since 2014 as essentially one person writing technology strategy analysis for a paying audience. No newsroom. No sales team. No venture capital. Thompson has never disclosed his subscriber count, but industry estimates place his annual revenue north of \$3 million. He writes four articles a week. He answers to no one but his readers.

What these examples share is not merely a business model but a philosophical reorientation. The paywall is not a barrier. It is a quality signal. You pay because you trust, and the payment reinforces the trust. When a reader commits eight dollars a month, she pays attention differently than when she stumbles upon a link in a social media feed. The transaction creates a covenant: I will be worth your money, and you will hold me to that standard.

One Thousand True Fans, Revisited

In 2008, Kevin Kelly — the founding editor of *Wired* — published a brief essay titled "1,000 True Fans." His argument was elegant: any creator could earn a sustainable living if they cultivated one thousand people willing to pay \$100 per year. One thousand multiplied by one hundred equals \$100,000 — a comfortable living in most of the world.

The essay circulated widely but was treated as aspirational rather than practical. The infrastructure did not exist. Accepting payments from one thousand individuals scattered across multiple countries was, in 2008, a logistical nightmare of PayPal disputes, currency conversion fees, and homemade mailing lists in spreadsheets.

Kelly was early. He was not wrong.

The infrastructure now exists, and it is formidable. Substack has made paid newsletters frictionless. Patreon hosts over a quarter of a million active creators and has processed more than \$3.5 billion in payments. Gumroad enables digital product sales. Buy Me a Coffee enables micro-patronage. Ko-fi. Memberful. Ghost. The tooling has become so commoditised that the question is no longer *how* a creator collects payment but *whether* they can build a devoted audience.

Li Jin, the venture capitalist who coined the phrase "passion economy," has argued persuasively that in many niches, one hundred true fans willing to pay \$1,000 per year may be more achievable — and more stable — than one thousand fans at \$100. A boutique consultancy. A

specialised research service. A private investment newsletter. The numbers shrink while the per-capita value rises.

Kamau does not have one thousand true fans. She has five hundred. But at \$96 per year each, she earns \$48,000 annually — life-changing money in Nairobi, respectable money in New York. She needs no venture capital. She needs no advertising sales team. She needs no algorithm's permission.

The Economics of Intimacy

What Kamau has, and what the five-hundred-thousand-follower influencer typically does not, is a relationship. Not the parasocial simulacrum that social media manufactures — where a creator broadcasts to an undifferentiated mass, each member of which feels privately addressed — but something closer to correspondence. Her subscribers reply to her emails. She replies back. They suggest stories. They correct her errors,

sometimes before she has noticed them herself. One subscriber, a policy researcher at the African Union, has become an informal adviser.

This is the economics of intimacy, and it inverts nearly every assumption of the attention economy.

In the attention model, scale is everything. A creator with one million followers is, by definition, more valuable than one with one thousand. But "valuable" to whom? To advertisers, certainly, who pay for reach. To the creator? Often not. The creator with a million followers earns perhaps \$3,000–\$5,000 per sponsored Instagram post — which sounds generous until one considers that the next post might earn nothing, that the algorithm might suppress her reach tomorrow, that her audience is rented, not owned, and that the platform's terms of service can change without notice.

The creator with one thousand devoted subscribers — people who have actively chosen to pay, who open every email, who would notice and complain if the quality declined — earns less in a single transaction but more in aggregate, more consistently, and with far less anxiety. She does not check her follower count compulsively. She does not game the algorithm. She has, in the language of the old economy, a customer base.

The platforms facilitating this shift are, notably, not social media platforms. They are infrastructure. Substack, Ghost, Patreon, Discord, Circle, Geneva — these services do not stand between the creator and her audience, extracting attention rents. They provide plumbing and take a modest cut. The creator owns her subscriber list. She can leave at any time and take her audience with her. This portability is not a minor technical detail. It is the foundation of creative independence.

Private podcasts — audio feeds accessible only to paying subscribers — have grown significantly since 2022. Discord servers, once associated primarily with gaming, now host tens of thousands of professional and creative micro-communities: writers' rooms, investment clubs, language-learning cohorts, academic reading groups. These are not broadcasts. They are gatherings. The distinction matters enormously.

The Machine That Cannot See Advertisements

Here is the most disruptive shift, and the one least discussed in polite company among advertising executives.

By 2028, AI agents are projected to handle thirty per cent of routine web browsing — not all internet usage, but the functional, transactional browsing that generates the majority of display advertising revenue: searching for products, comparing prices, reading reviews, scanning news summaries, booking travel, researching purchases.

If an AI agent reads articles on your behalf, compares products, summarises research, and recommends purchases, who sees the advertisements? The global display advertising market — worth roughly \$200 billion in 2025 — rests on a single assumption: that a human being will look at a screen. AI agents do not look at screens. They cannot be persuaded by a banner advertisement for running shoes or enticed by a pre-roll video for a meal delivery service. They process information; they do not experience it.

Google's search advertising business generated \$175 billion in 2024, all of it dependent on humans viewing a results page that mixes organic results with paid placements. When an AI agent performs that search instead — retrieving information directly, without loading

the page, without rendering the advertisements — the impression simply does not occur. The human never sees the page at all.

This is not speculative. It is already happening. OpenAI's browsing capabilities, Google's AI Overviews, Anthropic's tool-using models, and a proliferating ecosystem of AI assistants are already reducing traditional web page loads for a growing segment of users. Early data from web analytics firms suggests that AI-driven traffic increased by roughly forty per cent year-on-year in 2025. These visits generate zero advertising revenue.

The implications cascade. If display advertising revenue declines — not because of ad-blockers, but because the human is no longer in the loop — then every business model built on advertising subsidies comes under pressure. Free-to-read news sites. Free social media platforms. Free email services. The "free" internet was never free; it was paid for by advertising. If advertising's foundation erodes, the edifice must find new foundations or crumble.

This is, paradoxically, good news for creators like Kamau. Her income does not depend on whether her readers see advertisements. It depends on whether they value her work enough to pay directly. The intention economy is, by

its nature, immune to the disruption that AI agents pose to the attention economy. You cannot automate trust. You cannot delegate taste to a machine. You can ask an AI agent to summarise the news, but you cannot ask it to replicate the feeling of reading a writer whose judgement you have learned, over months and years, to rely upon.

The Coffeehouse Principle

In the late seventeenth century, London was home to roughly three thousand coffeehouses. They were not merely places to drink coffee — which was, in any case, often terrible. They were information exchanges. Lloyd's of London began as Edward Lloyd's coffeehouse, where ship captains and merchants gathered to discuss maritime insurance. Jonathan's Coffee-House in Exchange Alley became the London Stock Exchange. *The Tatler* and *The Spectator* were written in and distributed from coffeehouses.

What made these spaces powerful was not their scale but their specificity. Each attracted a particular clientele — merchants, scientists, writers, politicians. The

conversations were substantive because the participants shared context. Trust was high because the community was small enough that reputation mattered. A merchant who lied about a cargo's contents would find himself unwelcome by Thursday.

The micro-communities now forming across Discord, Geneva, Circle, and Mighty Networks bear a striking resemblance to those coffeehouses. They are small, specific, and — crucially — bounded. Entry often requires payment, invitation, or both. A Discord server for five hundred people discussing climate technology policy is not trying to be Twitter. It is trying to be a coffeehouse.

The economics of these communities are modest but sustainable. Five hundred members paying ten dollars per month generates \$60,000 annually — enough to support a full-time moderator-curator, occasional guest speakers, and a shared resource library. It is not venture-scale revenue. It is human-scale revenue. And that, increasingly, is the point.

The shift from broadcast to gathering represents a fundamental change in how information circulates. In the attention economy, information flowed from one to many: a publisher broadcast, an audience received. In the emerging model, information circulates within small

groups: a community discusses, debates, refines, and creates knowledge collectively. The value lies not in the content alone but in the context — who is saying it, to whom, and within what shared frame of reference.

Trust as the Scarce Resource

Herbert Simon, whose 1971 observation about attention scarcity launched the conceptual framework examined throughout this book, understood something his intellectual descendants often missed. He wrote: "A wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it." The key phrase is "allocate efficiently." Simon was not describing a problem of willpower. He was describing a problem of trust. How do you decide what deserves your attention? You trust something — a brand, a person, an institution, a recommendation from a friend.

In the attention economy, trust was unnecessary. The algorithm decided what you saw. Engagement metrics — likes, shares, comments — served as crude proxies for

quality, but everyone understood that virality and truthfulness were orthogonal qualities. A conspiracy theory could go viral precisely because it was false, therefore novel, therefore engaging.

In the intention economy, trust is the entire game. When you pay for a subscription, you purchase trust. When you join a community, you invest in a trust network. When you follow a curator's recommendation, you borrow their credibility. Trust becomes the scarce resource — not attention, which is merely the downstream consequence of trust well placed.

This has profound implications for how reputations are built. In the attention economy, reputation was a function of visibility: you were credible because you were everywhere. In the intention economy, reputation is a function of reliability: you are credible because you have been consistently right, consistently honest, consistently useful, over a sustained period. This favours depth over breadth, consistency over novelty, and — uncomfortably for those who have built careers on self-promotion — substance over performance.

The Curation Premium

If artificial intelligence can generate an effectively infinite quantity of competent prose, passable images, serviceable code, and plausible analysis — and by 2026, it demonstrably can — then creating content is no longer a competitive advantage. What becomes valuable is the ability to select.

This is the curation premium, and it is rising fast.

A human editor who reads one hundred articles and tells you, "Read these three; skip the rest," provides immense value — not because the other ninety-seven are bad, but because your time is finite and your trust in your own ability to sort signal from noise is, quite reasonably, declining. The volume of available content now so vastly exceeds any individual's capacity to evaluate it that the curator — the person or institution whose taste and judgement you trust — becomes the essential intermediary.

This is why independent bookshops are experiencing a quiet renaissance, even as overall book sales remain flat. The bookshop owner who arranges a table of twelve carefully chosen titles, with handwritten recommendation

cards, performs a curatorial act that Amazon's algorithm, for all its sophistication, cannot replicate. The algorithm knows what people like you have bought. The bookseller knows what you should read next. The difference is the difference between correlation and judgement.

Newsletter editors, podcast hosts who conduct interviews rather than read scripts, and critics who stake their reputation on their opinions are all, in different ways, selling curation. They sell the promise that they have done the reading so you do not have to, and that their judgement — imperfect, personal, occasionally idiosyncratic — is more useful than an algorithmic ranking optimised for engagement.

Payment Rails and Local Revolutions

The intention economy cannot function without infrastructure for small payments. This is obvious but frequently overlooked by commentators in San Francisco

and London, where credit cards are ubiquitous and Stripe integration is trivial.

Globally, the picture is more varied and more interesting.

India's Unified Payments Interface — UPI — processed over fourteen billion transactions in a single month in late 2025. The system, developed by the National Payments Corporation of India and launched in 2016, enables instant bank-to-bank transfers at zero cost. Razorpay, built atop UPI, allows Indian creators to accept payments of as little as ten rupees — roughly twelve US cents — without the transaction fees that would make such micro-payments uneconomical on Western platforms. An Indian newsletter writer charging forty-nine rupees per month can build a sustainable audience of a few thousand readers. The per-unit economics differ from the Western subscription model, but the aggregate mathematics work.

Brazil's Pix, launched in November 2020, has achieved comparable ubiquity. By 2025, over 150 million Brazilians — more than seventy per cent of the adult population — had used Pix at least once. Brazilian creators on platforms like Hotmart and Eduzz use Pix to sell courses, memberships, and digital products with a frictionlessness that Stripe is only now approximating in its core markets.

In East and Southern Africa, M-Pesa — the mobile money system launched by Safaricom in Kenya in 2007 — has created a payment infrastructure that predates and, in several respects, surpasses the card-based systems of wealthier nations. Over fifty million active users across Kenya, Tanzania, Mozambique, and other markets conduct everything from grocery purchases to creator subscriptions via their mobile phones. Kamau's subscribers in Kenya pay her through M-Pesa integration with her Ghost instance. No bank account required. No credit card. Just a phone number and a willingness to pay.

These local payment rails are not footnotes to a Western story. They are, in many cases, enabling the intention economy to develop faster and more inclusively than it is developing in the United States or Europe. When a student in Dar es Salaam pays 300 Tanzanian shillings via M-Pesa to access a weekly agriculture technology briefing, she participates in precisely the same economic reorientation as a hedge fund manager in Connecticut paying \$400 per year for *The Information*. The currency differs. The principle is identical.

The Uneven Transition

It would be dishonest to suggest that this transition is clean, complete, or universally beneficial.

The intention economy favours creators with existing credibility, education, and social capital. Kamau succeeded in part because she had a communications degree, years of professional experience, and fluency in the policy language her subscribers valued. A first-generation creator without those advantages faces a steeper climb. The democratisation is real but incomplete.

Subscription fatigue is a genuine concern. The average American household now subscribes to an estimated eleven digital services. Each additional subscription competes not just with other subscriptions but with the deeply ingrained expectation that online content should be free. There is a ceiling, and we do not yet know where it sits.

Nor will advertising vanish. For all its pathologies, advertising subsidises an enormous quantity of useful content and services that most people are unwilling to pay for directly. The weather forecast. The search engine. The mapping application. These will continue to be ad-

supported, and the advertising industry will adapt — as it has adapted before, from print to radio to television to digital. The transition is not from advertising to subscriptions. It is from an economy in which advertising dominated to one in which it is one model among several, and perhaps not the most important one.

But the direction is unmistakable. The infrastructure is built. The payment rails are laid. The creators are migrating. The audiences are following — slowly, selectively, and with the deliberateness that characterises intention rather than the passivity that characterised attention.

Wanjiru Kamau, when I spoke to her last, was considering hiring a second researcher. She had been approached by two venture capital firms offering investment and a traditional media company offering a column. She declined all three. "The moment I take their money," she said, "I answer to them instead of my readers. And my readers are the whole point."

She paused, then added: "Besides, five hundred people who trust me is worth more than five million who don't. I think the whole internet is about to learn that."

She may be right. But knowing what comes next — the shape of the post-attention world, its possibilities and its dangers, its winners and its casualties — requires more than economics. It requires a reckoning with what we want the internet to be, and what we are willing to build.

That reckoning is the subject of our final chapter.

Chapter 7: The Post-Attention World

Yuki Tanaka wakes at 6:14 a.m. in her flat in Kita-ku, Osaka, to the sound of rain against the window and nothing else. Her phone does not buzz. No notification banner slides across the lock screen. No red badge demands acknowledgement. On the low wooden table beside her futon, a small e-ink display — her *jouhou koukan*, her information companion — shows a single line of text: **7 items prepared. 12-minute read.**

She rises, brews hojicha, and sits. The device has read 340 articles overnight, drawn from eighty-seven sources across four languages, and surfaced seven that match the profile she has spent two years refining. A report on polymer recycling from *Nature Materials*. A municipal notice about changes to Osaka Metro's Midosuji line. A long essay from a Portuguese urbanist whose thinking she has followed since 2027. She reads two pieces over breakfast — the polymer report and the urbanist — then

sets the device face-down. The other five will wait, or they won't. There are no ads. There is no feed. There is no algorithm competing for her gaze, because the business model does not profit from her lingering. She pays Saisei Labs ¥1,800 per month. That is the entire transaction.

Eight thousand miles west, Mateus Oliveira da Silva opens his phone in his apartment in Vila Madalena, São Paulo. It is 7:40 a.m. A single notification awaits him — not from Instagram, not from X, not from any platform whose engineers once optimised for exactly this moment, the vulnerable first reach for the screen. The notification is from the neighbourhood cooperative's message board, a federated service hosted on servers in Campinas: *Ana Paula, Rua Harmonia 412, has surplus açaí from her garden. First come, first served. Bring a container.*

Mateus smiles, types a reply — *Chego em 20 minutos* — and sets the phone on the counter. A small widget reads: **Atenção: 3 / 45 min.** This is his attention budget, a setting he configured six months ago when he installed Claridade, a Brazilian-made operating system fork that treats screen time the way a budgeting app treats money. He has allocated forty-five minutes of non-essential screen time for the day. He has used three. The system

will not lock him out when he reaches the limit — it is not a parent; he is not a child — but it will ask, plainly and without manipulation, whether he wishes to continue. Most days, the question is enough.

Astrid Lindqvist sits in the number 7 tram as it rattles along Kungsgatan in Stockholm, reading a printed newsletter. It arrives every Thursday in a cream envelope with her name typeset in Garamond — not "Hey {first_name}!" generated by a marketing automation platform, but her actual name, set by an actual person who runs an actual small publishing operation from a shared office in Södermalm. The newsletter is called *Tråden* — "The Thread" — and it is written by Per Jansson, a former culture editor who left *Dagens Nyheter* in 2025 after the paper's third round of AI-driven redundancies.

Astrid has subscribed for four years. She pays 89 SEK per month — roughly eight pounds — and she has never seen an advertisement in it. She does not know how many other people subscribe. Per has never published the number. She suspects it is fewer than two thousand, because he once mentioned he could still, just about, reply to every letter he received. The newsletter covers Swedish

culture, Nordic design, European politics, and occasionally Per Jansson's own unpredictable enthusiasms — last month, a 2,000-word digression on the engineering of canal locks. Astrid read every word. She cannot remember the last time she read every word of anything served to her by an algorithm.

These three mornings are speculative, but not fantastical. Every technology they describe either exists in prototype or is the logical extension of trends already underway. Yuki's AI briefing agent is a more refined descendant of tools such as Artifact, Google's NotebookLM, and the summarisation layers already embedded in Arc Browser. Mateus's attention-budgeting operating system is a natural evolution of Screen Time and Digital Wellbeing, stripped of their performative half-measures and redesigned by people who actually want users to spend less time on their phones. Astrid's printed newsletter is not futuristic at all. It is arguably the oldest information technology in the vignette.

What makes these scenes feel foreign is not the technology. It is the economics. In each case, the person consuming information is the customer — not the product. The transaction is legible: money for service,

attention for value. There is no hidden third party harvesting behavioural data. There is no engagement optimisation loop. The relationship is bilateral, not trilateral.

This is what the post-attention world looks like. Not the abolition of attention — we shall always attend to things; that is what consciousness does — but the end of attention as an extractive commodity. The dismantling of an economic model that treated human focus as raw material to be mined, refined, and sold.

It will not arrive all at once. It will not arrive cleanly. And in several important domains, it will not fully arrive at all.

The Incomplete Transition

One of the persistent failures of technology writing — a failure I have committed more than once — is the prediction of total replacement. The internet will kill television. E-books will kill print. Streaming will kill cinema. These predictions are alluring because they are

narratively satisfying: the old world falls, the new world rises, and the writer stands at the hinge of history to narrate the turn.

Reality is less accommodating. Television adapted. Print contracted but survived. Cinema reinvented itself around spectacle and communal experience. The attention economy, too, will not die uniformly. It will recede from certain domains while persisting — perhaps intensifying — in others.

Entertainment will remain attention-driven for the foreseeable future. The competitive logic of streaming, gaming, and social video is fundamentally attention-based, and audiences are, by and large, willing participants. When someone opens TikTok at eleven o'clock at night, they are not being deceived. They may later regret the time spent, as one might regret a third glass of wine, but the transaction is broadly understood. The same holds for live sport, reality television, and multiplayer gaming. These are arenas where attention is the *point*, not a byproduct being siphoned.

Breaking news will retain attention economics, because urgency is genuinely scarce and genuinely valuable. When an earthquake strikes or a government falls, the race to inform is not an artificial competition manufactured by

engagement metrics — it is a real public need. The problem was never that attention economics exist. The problem is that they metastasised into domains where they did not belong: education, healthcare information, civic discourse, professional knowledge, publishing.

It is in these domains — where the goal is understanding rather than entertainment, where the stakes are truth rather than engagement — that the attention model is already failing and will continue to retreat. Several things will replace it, simultaneously and unevenly. But three successor economies are already visible.

The Trust Economy

In January 2024, *The New York Times* reported that Wirecutter, its product-recommendation subsidiary, generated more than \$100 million in annual revenue through affiliate commissions. This figure is remarkable not for its size but for what it represents. Wirecutter earns money because people trust its recommendations enough to buy on the basis of them. The entire value chain rests on a single, fragile, non-fungible asset: trust.

Trust economies are not new. The solicitor you hire because your neighbour recommended her. The restaurant you visit because a friend whose palate you respect told you to go. These are trust transactions, and they predate the internet by millennia. What is new is the scale at which trust can now operate as an explicit economic model — and the premium it commands in a world drowning in unreliable information.

Consider *The Information*, which charges \$399 per year and has never run an advertisement. Its subscriber base — estimated at fewer than 30,000 — is small by mass-media standards, but its readers are disproportionately senior executives, investors, and policymakers whose decisions move capital. Lessin's editorial model rests on a simple premise: if you trust us, you pay us. If you stop trusting us, you stop paying. There is no attention flywheel to disguise the loss of trust behind a wall of impressions.

Or consider Narayana Health, the Indian hospital network founded by the cardiac surgeon Devi Shetty. In a country where medical misinformation circulates freely on WhatsApp and trust in healthcare institutions varies enormously, Narayana has built its expansion not on advertising but on outcomes. Its cardiac surgery mortality rates are published. Its pricing is transparent. Patients

travel from across South Asia not because they saw a sponsored post, but because someone they trust — a relative, a referring doctor, a community health worker — told them: these people will take care of you.

This non-transferability is what makes trust economies fundamentally different from attention economies. Attention is liquid. It flows wherever the algorithm directs it. A viral post by a teenager in Manila can command more attention than a Pulitzer-winning investigation. Trust is solid. It accumulates. It cannot be gamed at scale — or rather, when it is gamed, the fraud eventually surfaces and the collapse is catastrophic. Ask the cryptocurrency exchanges. Ask the wellness influencers.

The Curation Economy

In 2025, Spotify reported that its algorithmic playlists accounted for roughly a third of all listening. The remaining two-thirds came from user-curated playlists, editorial playlists built by Spotify's in-house team, and direct artist searches. The algorithm was useful, but not dominant — and its influence was plateauing. Listeners

were gravitating towards human-curated collections: playlists built by DJs, by record shops, by friends with good taste.

This is the curation economy in miniature. When content is infinite — when generative AI can produce a competent article, a plausible image, a passable melody in seconds — creation loses its scarcity premium. What becomes scarce is judgement. The ability to say: *of the ten thousand things you could read this morning, these seven are worth your time.*

Curation is what Yuki's AI briefing agent performs, but it is also what Per Jansson performs for Astrid in Stockholm. One is algorithmic, the other human, and the interesting tension of the next decade will be their interplay. AI curation will handle volume — no human can read 340 articles overnight — but human curation will handle meaning. The algorithm can learn that Yuki is interested in polymer science. It cannot tell her why the Portuguese urbanist's essay matters, what it connects to in the broader discourse, why this argument at this moment carries weight. That requires a point of view. A sensibility. A mind willing to say, with the full weight of its accumulated judgement: *pay attention to this.*

The independent bookshop is a curation economy in physical form. So is the restaurant critic, the gallery owner, the film programmer at a repertory cinema. These roles have always existed, but they acquire new economic significance as AI-generated content makes undifferentiated abundance the default state of every medium. The curator's value proposition is not "I made this." It is "I chose this, and here is why."

The Intention Economy

In 2006, Doc Searls published an essay arguing that the future of commerce lay not in companies broadcasting products to consumers, but in consumers broadcasting their needs to companies. He called this the "intention economy" and spent the next decade developing the concept through ProjectVRM (Vendor Relationship Management) at Harvard's Berkman Klein Centre.

For most of those years, the idea remained stubbornly theoretical. The tools did not exist. The platforms had no incentive to build them. The attention economy was too profitable to permit its opposite.

But the infrastructure is now assembling. Personal AI agents — descendants of Siri and Alexa, but far more capable and far less beholden to their manufacturers' advertising ambitions — are beginning to mediate the relationship between consumers and commerce. When Yuki's AI assistant identifies that she needs a new portable battery charger, it does not show her advertisements. It queries a product database, cross-references specifications with her requirements, checks reviews from pre-approved sources, and presents three options ranked by suitability. The battery manufacturers cannot buy placement. They can only make a better product.

This is the intention economy made operational. If a significant portion of consumer decision-making is mediated by AI agents loyal to the user rather than the advertiser, then the entire apparatus of brand awareness — the billboards, the sponsored posts, the influencer partnerships, the pre-roll ads — loses its point of contact. You cannot emotionally manipulate an algorithm. You cannot create brand loyalty in a machine-learning model. You can only compete on merit, which is how markets were supposed to work in the first place.

The Regulatory Ratchet

None of these successor economies will emerge solely from market forces. Regulation — uneven, imperfect, frequently captured by the industries it purports to govern — is nonetheless applying directional pressure.

The European Union's Digital Services Act, which took full effect in 2024, requires large platforms to offer chronological feeds as an alternative to algorithmic ones, to provide transparency about recommendation algorithms, and to grant researchers access to platform data. It does not ban the attention economy. But it establishes a principle: that users have a right to understand and control the mechanisms that compete for their attention.

China's algorithmic recommendation regulations, enacted in 2022, went further in one respect: they required platforms to offer an explicit opt-out for algorithmic recommendations. The motivations were not libertarian — the Chinese Communist Party has its own reasons for controlling what algorithms promote — but the

mechanism is instructive. When Douyin was required to let users turn off its recommendation engine, a meaningful minority did so.

India's Information Technology Rules, amended in 2023, focused on content traceability — the ability to identify the originator of viral messages. The rules were controversial, rightly criticised for their potential to chill free speech, and unevenly enforced. But they reflected a genuine problem: in a country where WhatsApp forwards have been linked to mob violence, the frictionless virality of the attention economy carries costs that are not merely economic.

Brazil's Marco Civil da Internet, often called the country's "internet constitution," established net neutrality and user privacy protections as early as 2014 — prescient legislation that has since been tested and, in some cases, weakened by subsequent political pressures. Yet the framework endures, and it has shaped a Brazilian digital culture more conscious of platform power than its North American equivalent.

The regulatory picture is messy. But the direction is consistent across jurisdictions that otherwise agree on very little: towards transparency, towards user agency,

towards the principle that the design of information systems is a matter of public concern.

A Map, Not a Manifesto

I am suspicious of books that end with action plans. The twelve-step programme, the bullet-pointed "takeaways," the breathless call to arms — these are the trappings of airport business books, and this is not one. But I would be abdicating responsibility if I described the post-attention world without saying anything about how to move through it.

For individuals: Curate ruthlessly. Every subscription you maintain, every notification you permit, every app on your home screen is a claim on your attention, and claims left unexamined accumulate like debts. Pay for the publications and services you value — not as charity, but as the most direct way to ensure they survive and remain accountable to you rather than to advertisers. Build direct relationships with the writers and creators whose work matters to you. An email subscription is worth more than a follow. A paid membership is worth more than a like.

For creators: Build what you own. An email list is yours; a follower count belongs to the platform and is worth precisely what the platform's algorithm decides on any given Tuesday. Focus on depth over breadth — a thousand readers who trust you are more valuable, by every meaningful measure, than a million who encountered you once in a feed they were already scrolling past. Diversify your revenue away from advertising. This is not ideological purity; it is risk management.

For companies: Prepare for a world in which AI agents increasingly mediate consumer attention. Brand awareness — the traditional goal of advertising — will matter less than brand trust, because an AI agent comparing products will weight verified quality, transparent pricing, and genuine reviews over emotional resonance and visual identity. Impressions tell you how many eyeballs passed over your message. Conversions tell you how many people acted. But relationships — repeat customers, direct subscribers, community members — tell you how many people trust you enough to return. That is the metric that will matter, and most companies are not yet measuring it.

The Gift

I want to return, as this book must, to Herbert Simon.

His full passage, from the 1971 paper "Designing Organizations for an Information-Rich World," deserves reading in full, because it is so frequently truncated that its meaning has been distorted:

"In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it."

The passage has been quoted in hundreds of books, thousands of articles, and innumerable conference keynotes. It is the founding text of the attention economy. But notice the word almost always glossed over, the word doing the heaviest lifting: *efficiently*.

Simon did not say that attention should be captured. He did not say it should be extracted, monetised, or optimised. He said it should be *allocated* — and allocated *efficiently*, which is to say, in a manner that serves the person whose attention it is. The attention economy took Simon's diagnosis — that attention is scarce — and built a prescription he never endorsed: that because attention is scarce, it is valuable, and because it is valuable, it should be treated as a commodity to be harvested by those with the most sophisticated tools for harvesting it.

This was the fundamental error. Not a technical error, not a business error, but a moral one. It confused the scarcity of a human capacity with the exploitability of a natural resource. It treated the finite hours of a person's conscious life as an input to someone else's revenue model. And it built, over thirty years, an infrastructure of extraction so pervasive that we stopped noticing it, the way a fish does not notice water.

What comes after the attention economy is not the end of attention. We shall go on attending to things — to people we love, to work that matters, to art that moves us, to the ten thousand daily negotiations between our minds and the world. Attention is not going away. It is the substrate

of conscious experience. It is, if one is inclined to be philosophical, the closest thing we have to a secular definition of being alive.

What is ending — unevenly, incompletely, but irrevocably — is the treatment of attention as a commodity. The notion that it is acceptable to design systems whose explicit purpose is to capture and hold human attention beyond the point of value to the person attending. The business model that made surveillance the price of participation and distraction the cost of being informed.


In its place, something older is re-emerging. Not a new technology, not a new platform, not a breathless announcement at a conference in San Francisco. Something older and less glamorous: the idea that if you want someone's attention, you must deserve it. That attention is not a resource to be mined but a relationship to be honoured. That the appropriate response to someone giving you their attention is not to extract maximum value from the moment, but to make the moment worth their while.

Yuki in Osaka finishes her hojicha and sets down her reading device. Mateus in São Paulo walks to Ana Paula's house with an empty container, his phone quiet in his

pocket. Astrid in Stockholm folds Per Jansson's newsletter and tucks it into her bag as the tram reaches her stop.

None of them is thinking about the attention economy. That is rather the point. The structures that mediate their mornings have been designed not to demand thought about themselves, not to draw attention to their own mechanisms, not to keep their users engaged a moment longer than the content warrants. They have been designed around an old and simple principle that the attention economy spent three decades trying to make us forget.

Attention was never a resource to be mined. It was always a gift to be earned.



Acknowledgements

This book began as a series of observations scribbled in margins — of academic papers, of conference programmes, of the quiet moments between notifications.

It became something larger because of the people who insisted those observations mattered.

My editors at Kelford Press gave this project the rarest gift in publishing: time. Time to think, time to report, time to be wrong before being right. David Ashworth's editorial instinct shaped the argument from its earliest drafts.

The technologists, creators, researchers, and policy-makers who gave their time for interviews — in Lagos, Nairobi, Seoul, São Paulo, Bangalore, Hamburg, Tokyo, and dozens of other cities — are the spine of this book. Their generosity with their expertise and their candour about their struggles made every chapter possible.

Any errors that remain are, of course, my own.

About the Author

Eleanor Whitfield is the Senior Technology Correspondent at Kelford Press, where she covers artificial intelligence, enterprise software, and the intersection of technology and society. Her reporting spans six continents and draws on two decades of

covering how technology reshapes the way we live, work, and think. She is based in London and can usually be found in a library.

Also by Kelford Press

Kelford Press publishes original journalism and books across technology, business, science, culture, fashion, sport, and world affairs. Our Daily Papers arrive fresh each morning. Our Weekly Digest lands every Friday. Our Monthly Letters offer a considered view of the world.

Visit **kelfordpress.com** to explore our full catalogue.

Become a Member

Kelford Press readers are thinkers, not scrollers. If this book resonated with you, consider joining our community:

Reader (Free) — 5 articles per month, digest previews
Subscriber (£4.99/mo) — Unlimited papers, full digest,
monthly letter **Patron** (£9.99/mo) — Everything plus
digital books, audibles, and an ad-free experience **Fellow**
(£79.99/yr) — Patron benefits plus quarterly signed books
and exclusive events

Join at kelfordpress.com/subscribe

Kelford Press — Where Words Find Their Home Est.
2006