

THE ATTENTION RESET

Focus in a World of AI, Apps
and Algorithms

*A research-grounded guide to reclaiming
deliberate attention*

FIRST EDITION

2026

Contents

Preface	xv
How to Use This Book	xix
The Architecture of This Book	xxi
 I The Architecture of Distraction	 1
1 The Scarcest Resource	3
1.1 Simon's prophecy	4
1.2 From eyeballs to engagement	6
1.3 The machine that learns you	7
1.4 What the market actually buys	10
1.5 The generative turn	11
1.6 Why this is not a story about willpower	12
 2 The Engineered Pull	 15
2.1 The pull and the schedule	17
2.2 Dopamine, told correctly	18
2.3 From finding to method	22
2.4 Deceptive design, measured	24
2.5 The social fuel	27
 3 Death by a Thousand Pings	 31
3.1 Eleven minutes, then something else	32
3.2 The notification as a designed interrupt	37
3.3 Why we interrupt ourselves	40
3.4 A pace not dictated by electrons	41
3.5 The room, and then the norm	43

II	How Attention Works—and What AI Does to It	47
4	The Machinery of Focus	49
4.1	Not one faculty but several	50
4.2	The circuit-breaker	53
4.3	Four things at once, at most	55
4.4	The lapse	57
4.5	The mind that leaves	59
4.6	Load, and what it decides	62
5	The Myth of Multitasking	67
5.1	The ledger in the laboratory	69
5.2	Residue	71
5.3	Why the experience misleads	74
5.4	The media multitasking literature, told honestly . . .	75
5.5	Classrooms, screens and the shallowing of reading . .	78
6	Thinking with Machines	83
6.1	The mind has always leaked	84
6.2	What the generative turn actually changes	88
6.3	Overreliance is a measurable phenomenon, not a moral failing	92
6.4	What the productivity evidence actually shows	93
6.5	Cognitive debt, and the case for chosen difficulty . . .	95
III	The Attention Reset Playbook	101
7	Layer One—Physiology	103
7.1	The architecture of Part III	104
7.2	Sleep: the first casualty	106
7.3	Movement: small effects, reliably available	109
7.4	Restoration: what attention recovers from	112
7.5	Training the substrate	115
7.6	The effort economy	117
8	Layer Two—Environment: Making Distraction Expensive	123
8.1	The situational turn	124
8.2	A ladder, not a leap	127

8.3	What the rungs are worth	130
8.4	Why the built-in tools do not work	133
8.5	Rooms, desks and pockets	135
8.6	Binding the future self	137
9	Layer Three—Protocol	141
9.1	Why intention is the wrong unit	142
9.2	The deep-work block	143
9.3	Ready to resume: the highest-value four minutes in this book	148
9.4	Batching the channel	150
9.5	The AI co-working boundary	151
9.6	Keeping a protocol alive	156
10	Layer Four—Purpose: What Attention Is For	161
10.1	A commons, not a virtue	162
10.2	The norms that do the work	164
10.3	The assistant in the room	166
10.4	Growing up inside it	168
10.5	Design, policy and the ethics of capture	171
10.6	What attention is for	172
A	Worksheets and Instruments	179
A.1	The Attention Audit	179
A.2	The Trigger Inventory	181
A.3	The Ready-to-Resume Note	181
A.4	The Personal AI Charter	183
A.5	The Weekly Review	183
A.6	The Deep-Work Block Specification	185
A.7	Baseline Measures	186
B	The Twelve-Week Attention Reset	189
B.1	Weeks 1–2: Measure	190
B.2	Weeks 3–5: Subtract	191
B.3	Weeks 6–9: Build	193
B.4	Weeks 10–12: Consolidate	194
B.5	What twelve weeks can and cannot do	196

C Adaptations by Context	199
C.1 Students and researchers	199
C.2 Knowledge workers and managers	201
C.3 Software and technical workers	202
C.4 Parents and families	203
C.5 Clinicians and high-stakes professionals	205
C.6 Creative workers	206
 Glossary	 209
 Bibliography	 217
 Index	 237

The Scarcest Resource

What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention.

HERBERT A. SIMON, 1971

Consider the first ninety seconds of a fairly ordinary morning.

An alarm sounds from a device that is also a newspaper, a post-box, a casino, a shopping arcade and a telescope pointed at several thousand acquaintances. The sleeper silences it. In the two seconds during which a thumb is already in contact with the glass, a notification badge resolves into focus: eleven messages, three of which are work, one of which contains a word that will occupy a corner of the mind until roughly eleven o'clock. Somewhere in the same gesture a video begins to play, because the application has learned that autoplay converts idle thumbs into watched minutes. By the time the feet reach the floor, the day's first cognitive act has already been performed—not chosen, but performed—and it was performed on behalf of a system that will be paid for it.

None of this is an accident, and none of it is a moral failing. It is the predictable output of an economic arrangement that has been assembling itself for roughly a century and a half, and that has, in the last fifteen years, acquired machinery of a power that its architects did not originally possess. This chapter describes that arrangement: how human attention came to be treated as a commodity, how it came to be priced, and how the arrival of learning algorithms—and now generative models—changed the terms of the exchange. The purpose is not to induce despair. It is to establish the first premise of everything that follows: that the difficulty you have in holding your own attention is a rational response to an environment engineered

by extremely capable people to make holding it difficult, and that environments, unlike character, can be redesigned.




Figure 1.1. Five eras in the industrialisation of attention. Each transition did not replace its predecessor so much as absorb it: the portal inherited the logic of the broadcast advertisement, the feed inherited the logic of the portal, and the generative layer now inherits the logic of the feed while operating inside the tools we use to think.

1.1 Simon's prophecy

In 1971, the economist and cognitive scientist Herbert Simon contributed an essay to a volume on the future of computing in public life [170]. His subject was organisational design, but the observation that made the essay famous was almost incidental. Information, Simon pointed out, is not free to consume. It costs something to attend to, and what it costs is attention. Therefore the abundance of information, which everyone at the time regarded as an unambiguous good, would generate a corresponding scarcity—a poverty of attention—and the central design problem of an information-rich world would become the allocation of that scarce resource among the overabundant sources competing to consume it.

This was, in 1971, a remarkably counter-intuitive claim. The prevailing anxiety of the era concerned the difficulty of obtaining information, not the difficulty of escaping it. Simon's insight was that the binding constraint would migrate. Once information becomes effectively free to produce and distribute, the bottleneck moves downstream to the only input that cannot be manufactured: the finite number of seconds during which a human being can consciously process anything at all. That number is stubbornly fixed. It does not scale with bandwidth, storage or computation. It is, in the most literal economic sense, the scarcest resource in a digital economy—and it is the one input that every digital business must obtain from its users before it can obtain anything else.

Economists eventually formalised what Simon had described in prose. Falkinger modelled the consequences of limited attention in information-rich economies and showed that when the supply of signals grows faster than the capacity to process them, competition for attention systematically distorts what gets produced: senders invest not in the quality of their message but in its capacity to be noticed [57]. This is a precise statement of a phenomenon that anyone who has seen a headline written for a social feed will recognise. In an attention-scarce environment, the return on being interesting exceeds the return on being right.

EVIDENCE BOX: Attention as an economic constraint

Falkinger's formal analysis [57] establishes the key asymmetry: in an information-rich economy, the incentive to compete on *salience* rises relative to the incentive to compete on *substance*, because salience determines whether a message is processed at all. The mechanism does not require any actor to behave badly. It is a structural property of a market in which supply of information vastly exceeds the processing capacity of its recipients. Applied to digital media, it predicts precisely what has been observed: escalating stimulus intensity, shortening formats, and increasing investment in the science of capture.

1.2 From eyeballs to engagement

Attention became a commodity long before it became a data product. Tim Wu's history of the trade locates the origin in the nineteenth-century penny press, which discovered that a newspaper need not be sold at a profit if the assembled readership could be resold to advertisers [206]. The essential move—giving away the content in order to sell the audience—has been re-executed in every subsequent medium: radio, television, the web portal, the social network. What changes across eras is not the business model but the resolution at which attention can be measured and the precision with which it can be targeted.

By 2001, when Davenport and Beck published their study of what they called the attention economy, the language of attention as corporate currency had entered management practice [44]. Organisations were advised to think of their employees' and customers' attention as a budget to be allocated rather than an inexhaustible resource to be assumed. The advice was sound and almost universally ignored, in part because the metrics available at the time were crude. A newspaper could count circulation; a television network could estimate ratings. Neither could observe an individual's attention second by second, and neither could adjust its content in response.

The internet dissolved that limitation. When the medium is also the measurement instrument, every scroll, hover, pause and abandonment becomes an observation. Shoshana Zuboff's account of what she terms surveillance capitalism traces the consequences of this shift: behavioural data generated as a by-product of use was reconceived as a raw material for predicting, and eventually shaping, future behaviour [211]. Whether or not one accepts the full weight of her framing, the operational fact is not seriously disputed. The dominant platforms of the last two decades optimise explicitly for engagement metrics, and engagement metrics are, definitionally, measurements of captured attention.

James Williams, who worked on advertising systems before turning to philosophy, offers the sharpest statement of what this optimisation implies for the individual [200]. The technologies are not merely distracting us from our goals, he argues; at scale they interfere with the faculties by which we form goals at all. He distin-

guishes distractions from our immediate tasks, from our broader life projects, and from our capacity to reflect on what we want—and observes that the third is the most consequential and the least discussed. The reader who has ever concluded a two-hour session on a feed unable to reconstruct a single thing they learned, and equally unable to explain why they continued, has encountered the phenomenon Williams describes.

1.3 The machine that learns you

For most of the history of the attention trade, the seller had to guess. An editor chose a headline; a programmer scheduled a broadcast; the audience's response arrived days later in aggregate and could not be traced to any individual. The decisive change of the 2010s was the replacement of that guesswork with closed-loop optimisation. A recommender system observes what each specific person does, updates a model of that person, selects the next item, observes the response, and repeats—millions of times per second, across billions of people. What the system is optimising is not your satisfaction, your learning or your wellbeing, unless those happen to be measurable and included in the objective. What it is optimising is whatever engagement signal the operator has chosen to maximise.

The empirical literature on the consequences is more nuanced than either the enthusiasts or the alarmists usually concede, and it is worth being precise, because Part III's recommendations depend on getting this right.

There is good evidence that engagement-optimised ranking systematically changes what people see. In a large-scale study of Twitter's algorithmic timeline, Huszár and colleagues found that the algorithmic feed amplified political content relative to the reverse-chronological baseline, with the effect differing by political grouping in ways that were not attributable to any deliberate ideological choice [86]. Milli and colleagues went further, comparing engagement-optimised ranking with alternatives and finding that optimising for engagement signals increased exposure to divisive and low-quality content *relative to what users themselves reported preferring* [124]. That last clause is the important one. The gap between

revealed preference (what you click) and reflective preference (what you say you want more of) is exactly the gap in which the attention economy operates.

At the same time, several of the strongest claims in public circulation have not survived careful testing. The large field experiments conducted with Meta during the 2020 United States election found that switching users from an algorithmic feed to a chronological one changed consumption substantially but had minimal effects on political attitudes and polarisation over the three-month study period [76], and that exposure to like-minded sources, while prevalent, did not show the polarising effect the popular echo-chamber narrative predicts [137]. Earlier work had already shown that individual choice, rather than algorithmic ranking alone, accounted for a large share of ideological filtering on social platforms [10]. Studies of YouTube consumption similarly find that the consumption of radical content is concentrated among users with pre-existing preferences, rather than being manufactured by the recommender in otherwise uninterested viewers [84, 102]. A subsequent counterfactual study, which compared real viewing histories against bots following the recommender's suggestions, found that the recommendations themselves drifted towards more moderate material and that user preference, not the algorithm, was the dominant force [85]. The recommender is a weaker ideological actor than the popular account supposes.

A WORD OF CAUTION: The algorithm is not a mind-control ray

The evidence supports a specific, bounded claim: engagement-optimised ranking systematically shifts *what* is placed in front of us, and shifts it toward material that is politically activating [86] and divisive relative to what users themselves say they want [124]. The evidence does not support the stronger popular claim that feeds straightforwardly rewire political beliefs in the average user; large field experiments have repeatedly failed to find such effects on attitudes over periods of weeks to months [10, 76, 137]. It is worth being equally careful in the other direction: what these studies establish about *exposure* does not by itself establish how much time the systems consume, which is a separate quantity requiring separate

evidence. This distinction matters practically. If the problem were belief manipulation, the remedy would be information hygiene. Because the problem is *time and attentional capture*—a claim the next section defends on its own evidence—the remedy is architectural: changing how much access the system has to your attention in the first place.

The honest summary is therefore narrower and, for the purposes of this book, more useful. The algorithmic feed is not primarily an instrument of persuasion. The proposition this book will defend is that it is an instrument of **duration**—that what these systems reliably purchase is not belief but time, in the form of extended sessions, increased return frequency and the filling of interstitial moments.

That proposition is an inference rather than a single finding, and it is worth saying so plainly, because it is the premise on which Part III's whole strategy rests. It is assembled from four independent bodies of evidence, none of which is decisive alone. The first is the objective function itself: engagement-optimised ranking is, by construction, tuned to a behavioural quantity, and the divergence between what such systems surface and what users reflectively endorse has been measured directly [124]. The second is the checking habit documented in field studies of device use, in which brief inspections recur dozens of times a day and are frequently not recalled as having been initiated at all [143, 148]. The third is the interruption literature examined in Chapter 3, which establishes that these brief episodes are not cheap. The fourth, and the most direct, is the experimental evidence from removal: when access is deliberately restricted—by deactivating an account for a month [2], by blocking mobile internet for a period of weeks [30], or by reducing use in a structured way [19]—participants recover substantial quantities of time and reallocate it to offline activity. That last class of study is the strongest evidence available, because it manipulates access rather than merely observing it, and what it recovers is not accuracy of belief. It is hours.

Those interstitial moments—the ninety seconds in a lift, the four minutes before a meeting, the gap between putting down a book and picking it up again—were, for most of human history, where

reflection happened. Their systematic colonisation is, on this argument, the central attentional fact of our era.

1.4 What the market actually buys

It is tempting to describe the cost of all this in hours, and the hours are indeed striking. But hours are a misleading unit, because attention is not consumed linearly. What the attention market actually purchases is something more specific: the *onset* of your engagement, repeatedly, at moments you did not select. As Chapters 3 and 5 will show in detail, the cognitive cost of an interruption is concentrated in its transitions, not its duration. Thirty seconds of checking a notification does not cost thirty seconds. It costs thirty seconds plus a reconfiguration of your task set, plus the residue of the interrupting task, plus the elevated probability of a further self-interruption in the following minutes [4, 106, 116].

Economists have attempted to price the exchange from the other direction, by asking what people would accept to give up access. Allcott and colleagues ran a randomised experiment in which participants were paid to deactivate Facebook for four weeks [2]. The results were genuinely mixed—deactivation increased offline activity and subjective wellbeing, reduced political polarisation, and reduced post-experiment use—and, critically, participants still valued the service substantially, requiring meaningful compensation to give it up. The platforms deliver real consumer surplus. The point is not that these products are worthless. It is that their value is delivered alongside a cost that is invisible on any individual occasion and cumulative over a lifetime.

EVIDENCE BOX: The deactivation experiment

In a randomised controlled trial with over 2,700 participants, four weeks of Facebook deactivation produced small but significant improvements in subjective wellbeing, substantial substitution toward offline activity and television, reduced factual news knowledge, reduced political polarisation, and persistently lower Facebook use after the experiment ended [2]. Participants nonetheless placed a considerable monetary value on continued access. The study is

the strongest available evidence that these products deliver genuine value *and* impose genuine costs, and that users systematically under-appreciate the second until they experience its absence.

1.5 The generative turn

Until recently, the attention economy operated principally in the domain of leisure. The feed competed with the novel, the television and the conversation. Work, for the most part, took place in a different set of tools, and the boundary—while porous—was recognisable.

The generative turn dissolves that boundary. Since the deployment of large language models at scale, the same underlying technology that ranks a feed now sits inside the document, the inbox, the code editor, the search bar and the meeting. Foundation models, as Bommasani and colleagues named them, are general-purpose systems adapted to an enormous range of downstream tasks [16], and their commercial logic pushes them toward becoming the default interface to all cognitive work. The architecture that makes this possible is, with a pleasing irony, itself built on a mechanism its inventors named attention: the transformer's capacity to weight the relevance of every element of an input against every other [192].

The productivity evidence is real and should not be minimised. In a controlled experiment with professional writers, access to a generative assistant reduced task completion time by roughly forty per cent and improved output quality, with the largest gains accruing to initially lower-performing participants [136]. A field experiment with consultants found substantial gains on tasks within the model's competence, alongside a striking failure mode: on tasks that fell just outside that competence—what the authors call the jagged technological frontier—assisted participants performed *worse* than unassisted ones, because they accepted plausible but incorrect output [47]. A study of customer support agents in a real firm found productivity increases concentrated among novices, consistent with the model encoding and distributing the tacit knowledge of experts [23].

So the generative layer is not simply a new distraction. It is a genuine amplifier of certain kinds of output. But it changes the attentional situation in three specific ways that Part II will examine in depth. It moves the point of temptation inside the work itself, so that the alternative to difficult thinking is no longer a different application but a button within the same one. It makes fluent output available at near-zero cognitive cost, which alters what feels worth thinking through [67, 103]. And it introduces a new form of interruption—the suggestion, the completion, the proactive summary—that arrives precisely when concentration is deepest, because that is when you are typing.

PRACTICE: The ninety-second audit

Before reading further, reconstruct the first ninety seconds of your own morning as literally as you can. Write down, in order, every act of attention between waking and standing up: what you touched, what appeared, what you read, what you replied to. Do not judge the list and do not change anything yet. The purpose is calibration. Almost everyone discovers at least one event they had no memory of initiating, and that event is usually the most informative item on the page. Keep this list; it is the first entry in the Attention Audit that Appendix A formalises.

1.6 Why this is not a story about willpower

There is a version of this book that would end here with an exhortation: be more disciplined, want it more, try harder. That version would fail, and the research explains why with unusual clarity.

Hofmann and colleagues used experience sampling to track desires and their resistance across a week of ordinary life, and found that people spend a considerable proportion of their waking hours actively resisting some impulse, with media and communication among the most frequently resisted and most frequently yielded-to categories [83]. Resistance, in other words, is not a rare emergency measure; it is a near-constant background tax. And the people who achieve their goals are, on the evidence, not the people who win more of those confrontations. Milyavskaya and Inzlicht found that goal attainment was better predicted by experiencing

fewer temptations than by successfully resisting more of them [125]. Duckworth and colleagues drew the practical conclusion directly: the effective strategies for self-control are situational rather than heroic—changing the environment so that the confrontation does not occur—and they are systematically underrated by the people who need them most [51].

This is the hinge on which the entire book turns. If attention were principally a matter of character, the appropriate response to an engineered environment would be shame, and shame is famously ineffective as an intervention. If attention is principally a matter of architecture—physiological, environmental, procedural and motivational—then the appropriate response is engineering, and engineering is something one can actually do on a Tuesday.

RESET DRILL: Name the exchange

For one day, keep a tally of every occasion on which you open an application without having decided to. A mark on paper is sufficient; do not attempt to stop yourself, and do not use an app to do the counting. At the end of the day, look at the tally and ask a single question: what did I receive in exchange? The answer is sometimes genuinely valuable—a message that mattered, a piece of news you needed. Frequently it is nothing you can name. That gap, between the transactions you would repeat and those you would not, is the territory this book is about. Expect a number between thirty and one hundred; the count itself typically reduces the behaviour by a noticeable margin within a day, an effect we will exploit deliberately in Chapter 7.

KEY TAKEAWAYS

- Attention is the binding constraint of an information-rich economy. Simon predicted in 1971 that abundant information would produce scarce attention [170]; the formal economics confirms that such markets reward salience over substance [57].
- The business model is not new—selling an assembled audience predates the internet by more than a century [44, 206]—but the

measurement resolution is. When the medium is also the instrument, optimisation becomes closed-loop and personalised [211].

- Engagement-optimised ranking systematically shifts *what* we are shown, toward politically activating [86] and divisive [124] material, but large field experiments do not support the stronger claim that feeds straightforwardly change political attitudes [76, 85, 137]. The product of the feed is better understood as *duration* than as persuasion—an inference drawn from the removal experiments, in which restricting access returns substantial quantities of time [2, 30].
- These services deliver real value alongside real costs; a randomised deactivation experiment found improved wellbeing and reduced polarisation, yet participants still required substantial compensation to give up access [2].
- Generative AI relocates the competition for attention from leisure into the centre of cognitive work, delivering genuine productivity gains [23, 136] while introducing a characteristic failure at the edge of the model’s competence [47].
- Willpower is the wrong instrument. Goal attainment tracks the number of temptations encountered more closely than the number resisted [83, 125], which is why the remainder of this book is about architecture rather than effort [51].