CHAPTER ONE

The Six-Minute Mind

From Stolen Minds: The Engineering of Human Obsolescence

It's 2004, and Gloria Mark is sitting in a windowless office in the informatics department at UC Irvine, watching a software developer through a one-way observation mirror. The developer is trying to write a function. Something she's done a thousand times. But Mark isn't watching the code. She's watching the clock. Counting the seconds between interruptions. Not the interruptions themselves. The spaces between them.

Mark had started this research with a simple question: how long can a knowledge worker sustain focus on a single task before something pulls them away? She'd assumed the answer would be measured in hours. She thought she was studying productivity. She didn't yet know she was documenting the early stages of a cognitive catastrophe.

The developer lasts eleven minutes before she checks her email. Then she opens a browser tab. Then she responds to a colleague's message. Then she tries to return to the code. But the thread of thought is gone. She has to reconstruct it from scratch. Mark watches her eyes scan back through lines she'd already written, re-reading, re-processing, rebuilding the mental model that the interruption had demolished.

Twenty-three minutes later, the developer is back where she started. If she's lucky.

Mark recorded eight years of this. Forty thousand hours of observed computer activity across hundreds of knowledge workers. The pattern was consistent and the trajectory was terrifying: in 2004, the average time on a single screen before switching was two and a half minutes. By 2012, it had dropped to seventy-five seconds. By the late 2010s, forty-seven seconds. The decline tracked precisely with the rise of smartphone notifications, social media platforms, and algorithmic content feeds.

The title figure of this chapter — the six-minute mind — is not Mark's number, and I want to be clear about that before going any further. Mark measures the inside half of attention: the average duration of focus on any single screen before the next switch, the increasingly rapid micro-cycle that her data has been documenting since 2004. The six-minute figure is the outside half: roughly the duration of sustained focus most adults can now muster before some larger interruption — a phone, an alert, an internal urgency — breaks the work entirely. I have arrived at the rough six-minute estimate from interview rather than from a clean longitudinal study. It is consistent with what teachers describe in classrooms, what knowledge workers describe at their desks, and what I have measured on myself in the work that the rest of this chapter is built around. The screen-switch numbers are the

cleanest measurement of the substrate. The six-minute mind is what that substrate, accumulated over twenty years, has produced.

But the number that stopped Mark in her tracks was the recovery time. Twenty-three minutes and fifteen seconds, on average, to fully return to a task after interruption. Not because people were slow. Because the brain doesn't have a save function. When you switch away from a complex task, the mental model you've been building doesn't pause. It collapses. Rebuilding it from scratch takes cognitive effort that most people, by the time they've been interrupted four or five times, no longer have the resources to expend.

The mathematics are brutal: interrupted every eleven minutes, twenty-three minutes to recover. You never actually achieve sustained attention. You exist in a permanent state of partial focus, accomplishing bugger all of substance whilst feeling exhausted by the effort of trying.

Mark herself, in the work she has continued to publish across the past decade — gathered for a general audience in her 2023 book *Attention Span* — has been blunter about what the trajectory implies than her cautious early papers were. Her most recent published figure for the average screen-switch interval in the office settings she tracks is forty-seven seconds. The recovery time figure has not changed. What has changed, on her account, is the proportion of the working day in which sustained focus is happening at all: across the twenty years of her observation programme, the continuous stretches in which a worker holds attention on a single complex task have collapsed. The work the workers in those offices were paid to do was the same work. The conditions under which they could now do it were not.

That's what the platforms have done to you. And it wasn't an accident.

What This Did to Everything Else

Mark's recovery time figure — twenty-three minutes and fifteen seconds — does not measure attention by itself. It measures everything attention is the precondition for. You cannot encode a memory you didn't hold long enough to encode. You cannot construct an opinion you didn't hold an argument in your head long enough to weigh. You cannot maintain a sense of who you are across contexts if your sense of who you are has to be rebuilt from scratch every eleven minutes. The cognitive theft that the rest of this book describes — outsourced memory, fragmented identity, replaced judgement, deskilled work — happens against a backdrop in which the older capacity for sustained attention has already been removed.

I want to make this concrete, because it is the kind of finding that loses its bite when it stays at the level of statistics. The Tuesday I sat down to write this opening section, I gave myself ninety minutes. I closed my email. I put my phone in another room. I opened a clean document and started typing. By the eighteen-minute mark, I noticed I had walked from my desk to the kitchen to refill a kettle that didn't need refilling. By the thirty-minute mark, I

had reread the same paragraph four times because each time I had finished it I had got up to do something — adjust the curtain, find a different pen, check whether I had remembered to put the bins out — and lost the thread of the argument I had been building. At forty-two minutes I stopped, sat in the chair without moving, and tried to identify what was happening. What was happening was that the eleven-minute interval Mark had measured in 2004 had become, in me, in 2025, an interval of about three minutes. The interruption was no longer external. It was being generated by my own brain, which had been trained over twenty years to expect a small dopamine event roughly every three minutes and was supplying its own when none arrived.

This is the part of Mark's research that took me longest to internalise. The figure that gets quoted in the press is the eleven-minute interval, because eleven minutes sounds short and shocking and fits in a headline. The figure that matters is the recovery time. Eleven minutes of work followed by twenty-three minutes of recovery is not a productivity problem. It is the absence of a productivity problem in the older sense, because productivity in the older sense — sustained engagement with a task long enough to do it well — is no longer happening. Most knowledge workers I have interviewed in the past three years describe a working day that feels exhausting and produces little. They are describing the modal experience of a brain doing constant interruption-recovery cycles without ever actually arriving at the work the interruptions are interrupting. The exhaustion is real. It is the cost of the cycle. The work the cycle is preventing is invisible because, by definition, it is not getting done.

Forty-four per cent of those interruptions, Mark's longitudinal data found, are self-initiated. By 2016, the average user was reaching for an interruption before any external interruption arrived. The platforms had, by that point, been training users in interruption for ten years; the trained behaviour outlasted the training context. People reached for the phone in queues, in lifts, in the gap between two thoughts at the kitchen table. The interruption was no longer something the platform did to you. It was something you did, on the platform's behalf, because the dopamine system had been rewired to expect the next hit.

The other research that has accumulated in the years since Mark's original study consistently confirms the direction. Stanford's Memory Laboratory found that heavy media multitaskers — the people who are best, on self-report, at managing several streams of information simultaneously — were measurably worse on objective tests of filtering irrelevant information, controlling working memory, and switching between tasks. The very practice the platforms had habituated them to had degraded the underlying capacity. The literature on the dopamine mechanism that drives this — Anna Lembke's clinical work at Stanford in particular — I will come back to in detail in Chapter 11. The literature on what fragmented attention does to memory, opinion formation, and identity has its own chapters later in this book. The point of this opening chapter is to establish what the rest of the book is built on: the broken substrate.

The neurochemistry is worth naming once here, briefly, because it is the same machinery that the rest of the book's arguments rest on. Every notification arrives with a small dopamine spike — not at the moment of the notification, but at the moment of anticipation. The variable schedule that makes you reach for the phone before any specific notification has arrived is the same variable schedule that runs slot machines, and the dopamine response is the same response. The reason willpower keeps failing against it is not that willpower is weak. It is that willpower is not the right tool. The architecture of consumer technology has been calibrated, over twenty years and tens of billions of dollars of A/B testing, against the dopamine mechanism specifically. There is no realistic version of the contest in which a single user, exercising conscious self-discipline against an industrially-scaled engineering programme calibrated against their own neurology, comes out ahead. The contest is mismatched at the design level. We will look at the architecture of the mismatch in Chapter 4. For now, the relevant point is its consequence: a permanently distracted population that experiences its own distraction as an individual failing.

There is one further finding from the broader attention literature that I want to put on the table here, because it is the one that surprised me most when I came across it during the research for this book. The phenomenon is called *attention residue*. The term was coined by Sophie Leroy, the organisational psychologist at the University of Washington Bothell, in a 2009 paper that has held up remarkably well. The finding, simplified: when you switch from one task to another before you have finished the first, the first task does not actually leave your head. A piece of it stays — the unresolved problem, the incomplete sentence, the email you almost wrote — and it interferes with your ability to engage fully with the second task. The interference is measurable in the lab. It is also, on Leroy's evidence and the subsequent literature, larger than people realise from inside it.

What this means for the world Mark's twenty-three-minute number describes is that the cost of an interruption is not just the recovery time. It is also the residue of every previous interruption you have not yet fully cleared. By mid-afternoon on a normal working day, the average knowledge worker is carrying residue from twelve, fifteen, twenty unfinished cognitive episodes. Each of them is taking a small amount of working-memory capacity. The cumulative drag on the day's thinking is the sum of all of them. The exhaustion that most knowledge workers report by the end of an ordinary office day is not the cost of the work they did. It is the cost of the work they did not get to do, plus the residue of every interruption they accumulated trying. The system is producing, at population scale, a workforce that arrives home in the evening drained by a day spent recovering from interruptions to work that, in many cases, never quite happened.

Internal documents revealed through Frances Haugen's whistleblower disclosures in 2021 confirmed what users had been describing for years. From 2017 onwards, Facebook's ranking algorithm weighted emoji reactions — including the angry reaction — five times more heavily than the 'like'. The result was that angry, divisive posts spread further than ordinary 'liked' posts. The integrity teams flagged this internally. The weighting was kept. I

will not dwell on the Haugen disclosures here because the same documents will recur, in service of different points, in three of the next eight chapters. What matters for the present chapter is the simple fact: the company knew the engineered amplification was producing the cognitive harms its own researchers were documenting, and the company chose the engagement metric over the cognitive harm. That choice — engagement over wellbeing, every time — is the choice the architecture of control was built around.

The six-minute mind is not the end of the story. It is the substrate on which the next layer is being built. The cognitive infrastructure described in this chapter — fragmented attention, anticipation-driven reaching, depleted working memory — is precisely the substrate that the second wave of platform technology, the generative AI products, requires in order to take root. A mind that has been trained over a decade to delegate concentration to a feed is the mind most ready to delegate reasoning to a chatbot. The first wave produced the consumer that the second wave needs. We will return to this throughout the book; for now, hold the thought that the cognitive theft you have just read about is not the destination. It is the on-ramp.

What to Do Tonight

I'd been reading Mark's research for months, writing about it, explaining it to other people. I understood the mechanism intellectually. I could draw the dopamine cycle on a whiteboard. And I still couldn't read for more than twelve minutes without reaching for my phone.

I noticed it on a Sunday afternoon, sitting in my garden with a book I'd been looking forward to. Page thirty-seven. My hand moved to the table where my phone was sitting, face-down, as every digital wellness article had advised. I picked it up, unlocked it, opened Twitter, scrolled for ninety seconds, put it down, and tried to find my place.

I couldn't remember what had happened on page thirty-six. The twenty-three minutes started again.

I did this four more times in the next hour. Each time fully aware of what I was doing. Each time understanding the neuroscience. Each time doing it anyway. The knowledge made no difference. The system was better engineered than my willpower was strong.

That's not a personal failing. That's the design working exactly as intended.

What I'd tell you to do tonight, and I'm being specific because vague advice is useless, is this.

Set a timer for fifteen minutes. Sit with a physical book. Not a Kindle, a paper book. Read for fifteen minutes without touching your phone. When the impulse hits, and it will, probably within three minutes, notice it. Note the physical sensation: the restlessness, the itch in your fingers, the sudden conviction that something important might be happening in your pocket.

That sensation is the feeling of a system working on you. The discomfort is the diagnostic.

If you make it to fifteen minutes, you're doing better than most. If you can't, you've got your answer about how far the fragmentation has gone. Tomorrow, try sixteen minutes.

The second thing: count your pickups tomorrow. Not screen time. Pickups. The number of times you lift your phone without a specific purpose. Most people guess thirty or forty. The tracking studies that have measured it put the real figure at double that or more. The gap between what you think you're doing and what you're actually doing is the measure of how invisible the system has become.

I'm not going to give you a four-week programme. I've watched enough people try those. What I've seen work is smaller and more stubborn: the repeated, daily choice to use your own brain for things you've been letting your phone handle.

Read the physical book. Count the pickups. Notice the itch. That's enough for tonight.

Attention destruction is just the foundation. With your focus shattered, you're vulnerable to what comes next: the rewiring of your brain itself. Layer One is complete. Layer Two is already underway.