

Introduction

The Hyperactive Hive Mind

In late 2010, Nish Acharya arrived in Washington, DC, ready to work. President Barack Obama had appointed Acharya to be his director of innovation and entrepreneurship, and a senior adviser to the secretary of commerce. Acharya was asked to coordinate with twenty-six different federal agencies and over five hundred universities to dispense \$100 million in funding, meaning that he was about to become the prototypical DC power player: smartphone always in hand, messages flying back and forth at all hours. But then the network broke.

On a Tuesday morning, just a couple of months into his new role, Acharya received an email from his CTO explaining that they had to temporarily shut down their office's network due to a computer virus. "We all expected that this would be fixed in a couple of days," Acharya told me when I later interviewed him about the incident. But this prediction proved wildly optimistic. The following week, an undersecretary of commerce convened a meeting. She explained that they suspected the virus infecting their network had come from a foreign power, and that

Homeland Security was recommending that the network stay down while they traced the attack. Just to be safe, they were also going to destroy all the computers, laptops, printers—anything with a chip—in the office.

One of the biggest impacts of this network shutdown was that the office lost the ability to send or receive emails. For security purposes, it was difficult for them to use personal email addresses to perform their government work, and bureaucratic hurdles kept them from setting up temporary accounts using other agencies' networks. Acharya and his team were effectively cut off from the frenetic ping-pong of digital chatter that defines most high-level work within the federal government. The blackout lasted six weeks. With a touch of gallows humor, they took to calling the fateful day when it all began "Dark Tuesday."

Not surprisingly, the sudden and unexpected loss of email made certain parts of Acharya's work "quite hellish." Because the rest of the government continued to rely heavily on this tool, he often worried about missing important meetings or requests. "There was an existing information pipeline," he explained, "and I was out of the loop." Another hardship was logistics. Acharya's job required him to set up many meetings, and this task was substantially more annoying without the ability to coordinate over email.

Perhaps less expected, however, was that Acharya's work didn't grind to a halt during these six weeks. He instead began to notice that he was actually getting *better* at his job. Lacking the ability to simply send a quick email when he had a question, he took to leaving his office to meet with people in person. Because these appointments were a pain to arrange, he scheduled longer blocks of time, allowing him to really get to know the people he was meeting and understand the nuances of their issues. As Acharya explained, these extended sessions proved "very valuable" for a new political appointee trying to learn the subtle dynamics of the federal government.

The lack of an inbox to check between these meetings opened up cognitive downtime—what Acharya took to calling “whitespace”—to dive more deeply into the research literature and legislation relevant to the topics handled by his office. This slower and more thoughtful approach to thinking yielded a pair of breakthrough ideas that ended up setting the agenda for Acharya’s agency for the entire year that followed. “In the Washington political environment, no one gives themselves that space,” he told me. “It’s all neurotic looking at your phone, checking email—it hurts ingenuity.”

As I talked to Acharya about Dark Tuesday and its aftermath, it occurred to me that many of the hardships that made the blackout “hellish” seemed solvable. Acharya admitted, for example, that his concern about being out of the loop was largely alleviated by the simple habit of calling the White House each day to learn if there were any meetings he needed to know about. Presumably, a dedicated assistant or junior team member could handle this call. The other issue was the annoyance of scheduling meetings, but this could also be handled by an assistant or some sort of automated scheduling system. It seemed, in other words, that it might be possible to preserve the profound benefits of the email blackout while avoiding many of the accompanying annoyances. “What would you think of this way of working?” I asked after explaining my proposed fixes. The phone line went silent for a moment. I had pitched an idea so preposterous—permanently working without email—that Acharya’s mind had temporarily frozen.

Acharya’s reaction was not surprising. A widely accepted premise of modern knowledge work is that email saved us: transforming stodgy, old-fashioned offices, filled with secretaries scribbling phone messages and paper memos delivered from mail carts, into something sleeker and more

efficient. According to this premise, if you feel overwhelmed by tools like email or instant messenger, it's because your personal habits are sloppy: you need to batch your inbox checks, and turn off your notifications, and write clearer subject lines! If inbox overload gets really bad, then maybe your organization as a whole needs to tweak their "norms" around issues like response time expectations. The underlying value of the constant electronic communication that defines modern work, however, is never questioned, as this would be hopelessly reactionary and nostalgic, like pining for the lost days of horse transport or the romance of candlelight.

From this perspective, Acharya's Dark Tuesday experience was a disaster. But what if we have this exactly backward? What if email didn't save knowledge work but instead accidentally traded minor conveniences for a major drag on real productivity (not frantic busyness, but actual results), leading to *slower* economic growth over the past two decades? What if our problems with these tools don't come from easily fixable bad habits and loose norms, but instead from the way they dramatically and unexpectedly changed the very nature of how we work? What if Dark Tuesday, in other words, was not a disaster, but instead a preview of how the most innovative executives and entrepreneurs will be organizing their work in the very near future?

I've been obsessed with studying how email broke work for at least the past half decade. An important inflection point in this journey was in 2016, when I published a book titled *Deep Work*, which went on to become a surprise hit. This book argued that the knowledge sector was undervaluing concentration. While the ability to rapidly communicate using digital messages is useful, the frequent disruptions created by this behavior also make it hard to focus, which has a bigger impact on our

ability to produce valuable output than we may have realized. I didn't spend much time in *Deep Work* trying to understand how we ended up drowning in our inboxes, or suggesting systemic changes. I thought this problem was largely one of insufficient information. Once organizations realized the importance of focus, I reasoned, they could easily correct their operations to make it a priority.

I discovered that I was overly optimistic. As I toured the country talking about my book, meeting with both executives and employees, and writing more about these topics on my blog, as well as in the pages of publications like *The New York Times* and *The New Yorker*, I encountered a grimmer and more nuanced understanding of the current state of the knowledge sector. Constant communication is not something that gets in the way of real work; it has instead become totally intertwined in how this work actually gets done—preventing easy efforts to reduce distractions through better habits or short-lived management stunts like email-free Fridays. Real improvement, it became clear, would require fundamental change to how we organize our professional efforts. It also became clear that these changes can't come too soon: whereas email overload emerged as a fashionable annoyance in the early 2000s, it has recently advanced into a much more serious problem, reaching a saturation point for many in which their actual productive output gets squeezed into the early morning, or evenings and weekends, while their workdays devolve into Sisyphean battles against their inboxes—a uniquely misery-inducing approach to getting things done.

This book is my attempt to tackle this crisis. To pull together—for the first time—everything we now know about how we ended up in a culture of constant communication, and the effects it's having on both our productivity and our mental health, as well as to explore our most compelling visions for what alternative forms of work might look like. The idea of a world without email was radical enough to catch Nish Acharya

off guard. But I've come to believe it's not only possible, but actually inevitable, and my goal with this book is to provide a blueprint for this coming revolution. Before I can better summarize what to expect in the pages ahead, we must start with a clearer understanding of the problem we currently face.

As email spread through the professional world in the 1980s and 1990s it introduced something novel: low-friction communication at scale. With this new tool, the cost in terms of time and social capital to communicate with anyone related to your job plummeted from significant to almost nothing. As the writer Chris Anderson notes in his 2009 book, *Free*, the dynamics of reducing a cost to zero can be “deeply mysterious,”¹ which helps explain why few predicted the changes unleashed by this arrival of free communication. We didn't just shift our existing volume of voicemails, faxes, and memos to this new, more convenient electronic medium; we completely transformed the underlying *work-flow* that determines how our daily efforts unfold. We began to talk back and forth much more than we ever had before, smoothing out the once coarse sequence of discrete work activities that defined our day into a more continuous spread of ongoing chatter, blending with and softening the edges of what we used to think of as our actual work.

One study estimates that by 2019 the average worker was sending and receiving 126 business emails per day, which works out to about one message every four minutes.² A software company called RescueTime recently measured this behavior directly using time-tracking software and calculated that its users were checking email or instant messenger tools like Slack once every six minutes on average.³ A team from the University of California, Irvine, ran a similar experiment, tracking the computer behavior of forty employees at a large company over twelve

workdays. They found that the workers checked their inboxes an average of seventy-seven times a day, with the heaviest user checking more than four hundred times daily.⁴ A survey conducted by Adobe revealed that knowledge workers self-report spending more than three hours a day sending and receiving business email.⁵

The issue, then, is not the tool but the new way of working it introduced. To help us better understand this new workflow, I'll give it a name and definition:

The Hyperactive Hive Mind

A workflow centered around ongoing conversation fueled by unstructured and unscheduled messages delivered through digital communication tools like email and instant messenger services.

The hyperactive hive mind workflow has become ubiquitous in the knowledge sector. Whether you're a computer programmer, marketing consultant, manager, newspaper editor, or professor, your day is now largely structured around tending your organization's ongoing hive mind conversation. It's this workflow that causes us to spend over a third of our working hours in our inbox, checking for new messages every six minutes. We're used to this now, but when viewed in the context of even recent history, it represents a shift in our work culture that's so radical it would be absurd to allow it to escape closer scrutiny.

To be fair, the hyperactive hive mind is not obviously a bad idea. Among the benefits of this workflow is the fact that it's simple and incredibly adaptive. As one researcher explained to me, part of email's appeal was that this one easy tool could be applied to almost every type

of knowledge work—a much smaller learning curve than needing to master a separate bespoke digital system for each type of work. Unstructured conversation is also an effective method for identifying unexpected challenges and quickly coordinating responses.

But as I'll argue in part 1 of this book, the hyperactive hive mind workflow enabled by email—although natural—has turned out to be spectacularly ineffective. The explanation for this failure can be found in our psychology. Beyond the very small scale (say, two or three people), this style of unstructured collaboration simply doesn't mesh well with the way the human brain has evolved to operate. If your organization depends on the hive mind, then you cannot neglect your inbox or chat channels for long without slowing down the entire operation. This constant interaction with the hive mind, however, requires that you frequently switch your attention from your work to talking about work, and then back again. As I'll detail, pioneering research in psychology and neuroscience reveals that these context switches, even if brief, induce a heavy cost in terms of mental energy—reducing cognitive performance and creating a sense of exhaustion and reduced efficacy. In the moment, the ability to quickly delegate tasks or solicit feedback might seem like an act of streamlining, but as I'll show, in the long run, it's likely *reducing* productivity, requiring more time and more expenses to get the same total amount of work accomplished.

In this first part of the book, I'll also detail how the social element of the hive mind workflow clashes with the social circuits in our brains. Rationally, you know that the six hundred unread messages in your inbox are not crucial, and you remind yourself that the senders of these messages have better things to do than wait expectantly, staring at their screens and cursing the latency of your response. But a deeper part of your brain, evolved to tend the careful dance of social dynamics that has allowed our species to thrive so spectacularly since the Paleolithic,

remains concerned by what it perceives to be neglected social obligations. As far as these social circuits are concerned, members of your tribe are trying to get your attention and you're ignoring them: an event that registers as an emergency. The result of this constant state of unease is a low-grade background hum of anxiety that many inbox-bound knowledge workers have come to assume is unavoidable, but is actually an artifact of this unfortunate mismatch between our modern tools and ancient brains.

The obvious question is why we would ever adopt a workflow that comes with so many negative features. As I explain at the end of part 1, the story behind the rise of the hyperactive hive mind is complicated. No one really *decided* that it was a good idea; it instead arose, in some sense, of its own volition. Our belief that frenetic communication is somehow synonymous with work is largely a backfilled narrative we tell ourselves to make sense of sudden changes driven by complex dynamics.

Understanding the arbitrariness behind how we currently work, perhaps more than anything else, should motivate us to seek better options. This is exactly the goal I take on in part 2 of the book. In this second part, I introduce a framework I call *attention capital theory* that argues for creating workflows built around processes specifically designed to help us get the most out of our human brains while minimizing unnecessary miseries. This might sound obvious, but it actually contradicts the standard way of thinking about knowledge work management. As I'll show, driven by the ideas of the immensely influential business thinker Peter Drucker, we tend to think of knowledge workers as autonomous black boxes—ignoring the details of how they get their work done and focusing instead on providing them with clear objectives and motivational leadership. This is a mistake. There is massive potential productivity currently latent in the knowledge sector. To unlock it will

require much more systematic thinking about how best to organize the fundamental objective of getting a collection of human brains hooked together in networks to produce the most possible value in the most sustainable way. Hint: the right answer is unlikely to involve checking email once every six minutes.

The bulk of part 2 explores a collection of principles for applying attention capital theory to rebuild the workflows that drive organizational, team, and individual work in this direction—moving us away from the hyperactive hive mind and toward more structured approaches that avoid the problems of constant communication detailed in part 1. Some of the ideas supporting these principles come from cutting-edge examples of organizations experimenting with novel workflows that minimize unscheduled communication. Other ideas are drawn from the practices that enabled complex knowledge organizations to function effectively in an age before digital networks.

The principles described in part 2 don't insist that you banish messaging technologies like email and instant messenger. These tools remain a very useful way to communicate, and it would be reactionary to return to older and less convenient technologies just to make a point. But these principles will push you to reduce digital messaging from a constant presence to something that occurs more occasionally. The world without email referenced in the title of this book, therefore, is not a place in which protocols like SMTP and POP3 are banished. It is, however, a place where you spend most of your day actually working on hard things instead of talking about this work, or endlessly bouncing small tasks back and forth in messages.

This advice is designed to apply to many audiences. This includes business leaders looking to overhaul their company's operation, teams looking to function more efficiently, solo entrepreneurs and freelancers looking to maximize their value production, and even individual em-

ployees looking to get more out of their individual communication habits by viewing them from the perspective of attention capital. Accordingly, my examples span from the large scale, such as CEOs making drastic changes to their company's culture, to the small scale, such as my own experiments with using systems borrowed from software development to move my academic administrative tasks out of my inbox and into a more organized format.

Not every suggestion in part 2 applies to every situation. If you're an employee of a company that still worships at the altar of the hyperactive hive mind, for example, there are only so many changes you can make on your own without infuriating your coworkers. Some care will therefore be needed in picking and choosing the strategies you implement. (I attempt to help you in this selection by highlighting examples of how the various principles have applied in the individual context.) Similarly, if you're a start-up entrepreneur, you're better able to experiment with radical new work processes than if you're the CEO of a large company.

But I firmly believe that any individual or organization who starts to think critically about the hyperactive hive mind workflow, then systematically replaces elements of it with processes that are more compatible with the realities of the human brain, will generate a substantial competitive edge. The future of work is increasingly cognitive. This means that the sooner we take seriously how human brains actually function, and seek out strategies that best complement these realities, the sooner we'll realize that the hyperactive hive mind, though convenient, is a disastrously ineffective way to organize our efforts.

This book, therefore, should not be understood as reactionary or anti-technology. To the contrary, its message is profoundly future-oriented. It recognizes that if we want to extract the full potential of digital networks in professional settings, we must continually and

aggressively try to optimize how we use them. Attacking the flaws of the hyperactive hive mind is decidedly not an act of Luddism—if anything, the true obstruction to progress is giving in to the simplistic comforts of this blunt workflow at the expense of further refinement.

In this formulation, a world without email is not a step backward but a step forward into an exciting technological future we're only just beginning to understand. Knowledge work does not yet have its Henry Ford, but workflow innovations with impact on the same scale as the assembly line are inevitable. I can't predict all the details of this future, but I'm convinced it will not involve checking an inbox every six minutes. This world without email is coming, and I hope this book will get you as excited about its potential as I am.

Part 1

The Case Against Email



Chapter 1

Email Reduces Productivity

The Hidden Costs of the Hyperactive Hive Mind

When I first met Sean, he told me a familiar story about communication in his workplace. Sean was the cofounder of a small technology firm that designed internal-facing applications for large organizations. His company had seven employees working out of a London office, and they were, as Sean described it, enthusiastic practitioners of the hyperactive hive mind workflow. “We used to have Gmail opened constantly,” he told me. “Everything was handled in email.” Sean would start sending and receiving messages immediately on waking up and continue into the night. One employee even asked Sean to stop sending emails so late, as the knowledge of messages from the boss piling up while he slept was stressing him out.

Then the hyperactivity shifted into a new gear. “There was all this hype about Slack, so we decided to try it,” Sean remembered. The rate of back-and-forth communication intensified, especially after a demanding

client was provided access to their channels, allowing them to check in and ask questions whenever they felt like it: “Constant interruptions, every day.” Sean could feel the whiplash attention swings from messages to work to messages and back again wearing down his ability to think clearly. He grew to despise his phone’s notification pings. “I hated it—the sound still gives me the shivers,” he said. Sean worried that the mental grind of managing all this communication was reducing his company’s effectiveness. “I would work until one a.m. every night,” he said, “because that was the only time I felt free from distractions.” He also began to doubt that all this incessant chatter was mission critical. When he ran a review of his team’s Slack usage, he found that the most popular feature was a plug-in that inserts animated GIFs into the chat conversations. Sean reached a new low when two of his project supervisors suddenly quit. “They were burnt out.”

Sean’s frustrated sense that all this digital back-and-forth is making us less productive turns out to be a common sentiment. In the fall of 2019, as part of the research for this book, I invited my readers to participate in a survey about the role of email (and related tools like Slack) in their professional lives. More than 1,500 people responded, and many of them echoed Sean’s frustration—not with the tools themselves, which are self-evidently efficient ways to communicate, but with the hyperactive hive mind-style workflows they enable.

One thread of these responses concerned the sheer volume of communication generated by this workflow. “Every day it’s a barrage of emails regarding scheduling, deadlines, and they’re not used very effectively,” wrote a lawyer named Art. George, also a lawyer, described his inbox as containing “an avalanche of messages” in which important things get lost.

Another thread focused on the inefficiency of stretching out conversations into endless back-and-forth messaging. “The asynchronous nature is both a blessing and a curse,” wrote a financial analyst named Rebecca. “It is a blessing in that I can ask a question or delegate a task without having to find the person. It is a curse in that there is an implicit expectation that we are checking email all the time and will respond quickly.” An IT project manager similarly complained: “Simple conversations (that could have been dealt with within a matter of hours) can end up beginning a drawn-out email thread being read by an ever-increasing list of recipients.” A public services administrator noted that moving these interactions to digital messages also makes them “overly formal” and “less creative or on-point.” As she elaborated: “A project or task that could be relatively simply completed with a group working together in person becomes far more complicated by trying to manage all of the back-and-forth communication via email.”

Another common argument for email diminishing productivity centered on its ability to increase the amount of irrelevant information it suddenly forces you to process. “I’m frustrated that I receive so many updates . . . that have nothing to do with my position,” wrote a teacher named Jay. “People now confuse answering emails with real work,” wrote an editor named Stephanie. “There is a performative dimension to writing emails and cc’ing everybody, like ‘Look at all the work I’m doing.’ It’s annoying.” As an HR consultant named Andrea put it: “In at least 50% of messages you still have open questions. . . . You get the feeling that the person just shot off an email without caring about how I could answer it.”

As in Sean’s story, instant messenger tools like Slack weren’t let off the hook by my respondents, as they were described by many readers as simply email with faster response expectations. “Slack is just a string of

messages. It invites people to post almost without limitations,” wrote an executive coach named Mark. “It’s awful.”

The above stories, of course, are anecdotal. But as I’ll elaborate in the following pages, when you turn your attention to the relevant research literature, it becomes clear that the problems the respondents hinted at are even worse than most probably realize. Email might have made certain specific actions much more efficient, but as the science will make clear, the hyperactive hive mind workflow this technology enabled has been a disaster for overall productivity.

Constant, Constant Multitasking Crazyiness

In the late 1990s, Gloria Mark enjoyed an enviable professional setup. Mark’s research focused on a field known as computer-supported collaborative work (CSCW), which, as the name suggests, looks at ways that emerging technology can help people work together more productively. Though CSCW had been around since at least the 1970s, when it began with a focus on dry topics like management information systems and process automation, it received a jolt of energy in the 1990s as computer networks and the internet enabled innovative new approaches to work.

At this time, Mark was a researcher at the German National Research Center for Information Technology in Bonn, where she could, as she told me, “work on whatever I wanted.” Practically, this translated to her “going deep” on a small number of projects at a time, most of which focused on novel collaboration software. Among other projects, Mark worked on a hypermedia system named DOLPHIN, meant to make meetings more effective, and a digital document-handling system named PoliTeam, meant to simplify paperwork within a government ministry. As was the custom in Germany, lunch was the main meal of the day. As

Mark explained, she would enjoy long meals with her colleagues followed by long walks around the campus—they called these “rounds”—to digest their food and work through interesting thoughts. “It was beautiful,” she told me. “The campus had a castle on it.”

In 1999, Mark decided it was time to return to her native United States. Both she and her husband had secured academic jobs at the University of California, Irvine, so they packed up, said goodbye to the long stretches of deep work interspersed with leisurely meals and afternoon rounds by the castle, and headed west. Arriving in an American academic job, Mark was immediately struck by how *busy* everyone seemed. “I had a very difficult time focusing,” she said. “I had all of these projects to work on.” The long lunches she enjoyed in Germany became a distant memory. “I barely had time to grab a sandwich or salad for lunch,” she said, “and when I returned, I could see my colleagues in their offices doing the same thing, eating in front of their computer screens.” Curious to figure out how general these work habits had become, Mark persuaded a local knowledge sector company to allow her research team to shadow a group of fourteen employees over three workdays, looking over their shoulders and precisely recording how they spent their time. The result was a now famous paper—or infamous, depending on your perspective—presented at a 2004 computer-human interaction conference, with a provocative title that quotes a research subject’s description of her typical workday: “Constant, Constant, Multi-tasking Crazyiness.”¹

“Our study confirms what many of our colleagues and ourselves have been informally observing for some time: that information work is very fragmented,” Mark and her co-author, Victor González, write in the paper’s discussion section. “What surprised us was exactly how fragmented the work is.” The core finding of the paper is that once you eliminate formally scheduled meetings, the employees they followed

shifted their attention to a new task *once every three minutes* on average. Mark's experience of suddenly being pulled in many different directions when she arrived in California was not unique to her—it instead seemed a more universal property beginning to emerge in knowledge work.

When I asked Mark what caused this fragmentation, she replied quickly: “Email.” She came to this conclusion, in part, by diving back into the relevant literature. Since at least the 1960s, researchers have been measuring how managers spend their time in the workplace. Though the different categories they tracked have changed over the years, there are two key types of effort that show up consistently: “scheduled meetings” and “desk work.” Mark pulled out the findings on these two categories from a series of papers beginning in 1965 and ending with a 2006 follow-up to her original multitasking craziness study.

When Mark tabulated these results into a single data table, a clear trend emerged. From 1965 to 1984, the employees studied spent around 20 percent of their day engaged in desk work and around 40 percent in scheduled meetings. In the studies since 2002, these percentages roughly swap. What explains this change? As Mark points out, in the gap between the 1984 and 2002 studies, “email became widespread.”²

When email arrived in the modern workplace, people no longer needed to sit in the same room as their colleagues to discuss their work, as they could now simply trade electronic messages when convenient. Because email counts as “desk work” in these studies, we see time spent on desk work grow as time spent in scheduled meetings falls. Unlike scheduled meetings, however, conversations held through email unfold asynchronously—there's usually a gap between when a message is sent and ultimately read—meaning that the compacted interactions that once defined synchronous meetings are now spread out into a shattered

rhythm of quick checks of inboxes throughout the day. In Mark and González's study, the average scheduled meeting took close to forty-two minutes. By contrast, the average time spent in an email inbox before switching to something else was only two minutes and twenty-two seconds. Interaction now occurs in small chunks, fragmenting the other efforts that make up the typical knowledge worker's day.

It's here, therefore, in these nondescript data tables from CSCW papers published over a decade ago, that we find some of the first empirical evidence for the hyperactive hive mind hypothesis I outlined in this book's introduction. We shouldn't, however, place too much emphasis on just a single study. Fortunately for our purposes, around the time Gloria Mark began studying how communication technologies were transforming knowledge work, other researchers began asking similar questions.

A 2011 paper appearing in the journal *Organization Studies* replicated Mark and González's pioneering work by shadowing a group of fourteen employees in an Australian telecommunications firm. The researchers found that, on average, the employees they followed divided their workday into eighty-eight distinct "episodes," sixty of which were dedicated to communication.³ As they summarize: "These data . . . seem to lend support to the notion that knowledge workers experience very fragmented workdays." In 2016, in another paper co-authored by Gloria Mark, her team used tracking software to monitor the habits of employees in a research division at a large corporation and found that they checked email, on average, over seventy-seven times per day.⁴

Papers measuring the average number of email messages sent and received per day also show a trend toward increasing communication: from fifty emails per day in 2005,⁵ to sixty-nine in 2006,⁶ to ninety-two by 2011.⁷ A recent report by a technology research firm called the