

INTRODUCTION

Ghosts in the Machine

The human labor powering many mobile phone apps, websites, and artificial intelligence systems can be hard to see—in fact, it’s often intentionally hidden. We call this opaque world of employment *ghost work*.¹ Think about the last time you searched for something on the web. Maybe you were looking for a trending news topic, an update on your favorite team, or fresh celebrity gossip. Ever wonder why the images and links that the search engine returned didn’t contain adult content or completely random results? After all, every business, illicit or legitimate, advertising online would love to have its site ranked higher in your web search. Or think about the last time you scrolled through your Facebook, Instagram, or Twitter feed. How do those sites enforce their no-graphic-violence and no-hate-speech policies? On the internet, anyone can say anything, and, given the chance, people certainly will. So how do we get such a sanitized view? The answer is people and software working together to deliver seemingly automated services to customers like you and me.

Beyond some basic decisions, today’s artificial intelligence can’t function without humans in the loop. Whether it’s delivering a relevant newsfeed or carrying out a complicated texted-in pizza order, when the artificial intelligence (AI) trips up or can’t finish the job, thousands of businesses call on people to quietly complete the project. This new digital assembly line aggregates the collective input of distributed workers, ships pieces of projects rather than products, and operates across a host of economic sectors at all times of the day and night. In fact, the rise of this shadow workforce is part of a larger, more profound reorganization of employment itself. This yet-to-be-classified form of employment done on demand is neither inherently good nor bad. But left without definition and veiled from consumers who benefit from it, these jobs can easily slip into ghost work.

Businesses can collect projects from thousands of workers, paid by the task. Now they can depend on internet access, cloud computing, sophisticated databases, and the engineering technique of human computation—people working in concert with AIs—to loop humans into

completing projects that are otherwise beyond the ability of software alone. This fusion of code and human smarts is growing fast. According to the Pew Research Center's 2016 report *Gig Work, Online Selling and Home Sharing*, roughly 20 million U.S. adults earned money completing tasks distributed on demand the previous year.² Professional, white-collar information service work, delivered through on-demand work platforms, is already projected to add \$2.7 trillion, or 2.0 percent, to global GDP by 2025.³ If trends continue at the current rate, economists estimate that by the early 2030s, tech innovation could dismantle and semi-automate roughly 38 percent of jobs in the U.S. alone.⁴ Left unchecked, the combination of ghost work's opaque employment practices and the shibboleth of an all-powerful artificial intelligence could render the labor of hundreds of millions of people invisible.

Who does this kind of work? People like Joan and Kala.

Joan works from the Houston home she shares with her 81-year-old mother. In 2012, Joan moved in to care for her mother after a knee surgery left her mom too frail to live on her own. A year later, Joan started picking up work online through MTurk—short for “Amazon Mechanical Turk,” a sprawling marketplace owned and operated by tech giant Amazon.com. Joan makes some of her best money doing “dollars for dick pics.” That’s how she describes labeling pictures flagged as “offensive” by social media users on platforms like Twitter and Match.com.

Companies can't automatically process every piece of content users flag for review, so some of the harder-to-evaluate materials are routed to workers like Joan. On the surface, her task seems simple: click on pictures and assess their content. Is that an X-rated penis selfie that should be removed, or some innocuous G-rated body part? She is paid for each task she completes and decides when she walks away from her computer. Joan, with years of practice, now knows how to piece together an average ten-hour day that will bring in roughly \$40* worth of such tasks.

Thousands of miles away in Bangalore, India, Kala works from her makeshift home office, tucked away in the corner of her bedroom.⁵ Joan and Kala do similar tasks, sorting and tagging words and images for internet companies, but Kala picks up work from an outsourcing company that supplies staff to the Universal Human Relevance System (UHRS), an MTurk-like platform used internally by its builder, Microsoft. Kala, a 43-year-old housewife and mother of two with a bachelor's degree in electrical engineering, calls her two teenage sons into the room, points to a word displayed inside a large text box on her LED monitor, and asks them, “Do you know what this word means? Is it something you shouldn't say?”

They giggle as she reads the text out loud to them. They make fun of her pronunciation of “chick flick.” Together they decide that, no, this sentence does not contain adult content. Kala clicks “no” on the screen, and the window refreshes with a new text phrase to read to her sons. “They are more qualified to recognize these words than me,” she says, laughing. “They help me keep the internet clean and safe for other families.” Though she’s typically unable to find enough tasks to fill more than 15 hours of work in a given week, Kala returns to UHRS almost every day to see if there are any new tasks that she feels qualified to do. Kala’s doggedness and luck in the past have paid off. Now that she’s learned how to browse and claim tasks quickly, Kala can make the time she has between making meals and checking her children’s homework feel, as she puts it, “fruitful” as she does web research for what she considers extra income.

Content moderation—from sifting through newsfeeds and search results to adjudicating disputes over appropriate content to help technology and media companies figure out what to leave up or take down—is just one example of a new type of work that depends on people like Joan and Kala. Reviewing content is a common, often time-sensitive task generated in the wake of social media companies’ attempts to identify family-friendly materials for the billions of people who use their sites every day. There are way too many webpages, photos, and tweets in every imaginable language for people like Joan and Kala to assess them all.

Companies like Google, Microsoft, Facebook, and Twitter use software to automatically remove as much “not safe for work” content as they can, wherever possible. But these software filtering systems, powered by machine learning and artificial intelligence, aren’t perfect. They can’t always tell the difference between a thumb and a penis, let alone hate speech and sarcasm. Remember that classic moment in the 2012 U.S. presidential campaign when Republican candidate Mitt Romney uttered the phrase “binders full of women!”? Twitter needed workers, doing the same type of work that Joan does, to figure out, in real time, why a hashtag attached to such an obtuse phrase was quickly soaring to the top of its trending topics. Was it a hack? A glitch? Bona fide, frenetic Twitter use? Current AI systems can’t reliably tell the difference. On-demand work offers the promise of blending the power of computation with the creativity and dynamism of human insight.

This book is the story of Joan, Kala, and the millions of workers like them who step in when AI falls short. They are the humans behind the seemingly automated systems that we all take for granted. But modern AI systems don’t just need humans to answer an unfamiliar or challenging

question; they also need humans to help them learn how to answer anything in the first place. For example, do an image search for “camelback couch” and you’ll get a whole bunch of pictures of couches with curved backs. Search engines like Bing and Google don’t see or understand images in the way we humans do. Furniture aficionados need no more than a second to recognize a swank piece of furniture with a curved back that multiple people can sit on as a camelback couch. The AI systems behind search engines must start with at least a few hundred images of curve-backed couches, each labeled “camelback couch.” Then, when the search engine encounters a new picture of a couch, it runs what is called a “classification algorithm,” which essentially checks to see if the couch in this new image matches the geometrical patterns of those labeled “camelback” more than those not labeled “camelback.” Now, where did the initial set of labeled images, called training data, come from? From people like Justin. With no more than a two-sentence task description as guidance, workers like Justin must claim a job within seconds or lose it to someone else willing to scoop up the job first. Justin’s a stay-at-home dad with two young sons, working around his kids’ preschool and nap schedules. He readily admits he had no idea what a camelback couch was at the start. “I had to spend an enormous amount of time on Google trying to look up these terms to figure out what they meant before I could answer the questions.”

TripAdvisor, Match.com, Google, Twitter, Facebook, and Microsoft are some of the better-known businesses that generate an array of projects that people like Justin are paid to do, task by task, 24 hours a day, seven days a week. New companies crop up every day with business models that depend on workers around the world who respond to open calls routed through software to do this behind-the-scenes work. Businesses that can contract out their day-to-day activities to independent workers instead of regular employees can use ghost work to answer a web-based customer chat query, edit a product review, or do just about any task that doesn’t require an employee’s full-time, physical presence.

How Does Ghost Work Work?

A computer program is no more than a list of instructions that tell a computer what to do. When two software programs (or a piece of software and a piece of hardware) need to communicate, they must first establish a

common language. They do so via an application programming interface, or API. The API determines the common language by defining the list of instructions that a program will accept and what will happen after each instruction is executed. One could say that the API specifies the computer program's "rules of engagement." For example, there are hundreds if not thousands of different kinds of computers on the market right now, so writing a custom version of a software system for each type would be impossibly complex. But when all (or at least significant fractions) of the machines available obey the same API, programmers can write code once for all of these kinds of machines, because the API ensures that all of the machines understand the same language. These types of APIs are limited to what a computer can do, but the MTurk API enabled software developers to write programs, using only a slightly different set of instructions, that automatically pay *humans* to do tasks.⁶

Normally, when a programmer wants to compute something, they interact with a CPU through an API defined by an operating system. But when a programmer uses ghost work to complete a task, they interact with a person working with them through the on-demand labor platform's API.⁷ The programmer issues a task to a human and relies on the person's creative capacity—and availability—to answer the call. Unlike CPUs, humans have agency: they make their own decisions. While CPUs just execute whatever instruction they are given, humans make spontaneous, creative decisions and bring their own interpretations to the mix. And they have needs, motivations, and biases beyond the moment of engagement with the API. Given the same input, a CPU will always output the same thing. On the other hand, if you send a hungry human into a grocery store, he or she will walk out with a dramatically different bag of groceries than if they were not hungry. In exchange for this impetuosity and spontaneity, humans bring something to work that CPUs lack: creativity and innovation. Joan, Kala, and Justin are members of a growing economy, hidden by APIs and fueled by ghost work.

Less than two decades ago, software developers only wrote code for computers to execute. The MTurk API, and those that followed, allowed programmers to use humans to do tasks that are beyond a computer's capacity, like accurately making a quick judgment call, as Kala and Joan do when they determine what is and isn't adult content. In fact, anyone sitting in front of a web browser could now answer an automated request for help. Businesses call this mix of APIs, rote computation, and human ingenuity "crowdsourcing," "microwork," or "crowdwork." Computer scientists call it "human computation." Any project that can be broken

down into a series of discrete tasks can be solved using human computation. Software can use these APIs to manage the workflow and process the output of computers and individuals and even pay people for their contributions once they have completed the task. These people power modern AI systems, websites, and apps that we all use and take for granted.

Imagine a woman in her early twenties—let’s call her Emily—standing on a curb in Chicago. Emily opens the Uber app on her smartphone and an Uber driver responds. Neither Emily nor the driver knows that their meeting hinges on another woman, two oceans away—perhaps her name is Ayesha.⁸

Emily and her driver have no idea that Uber’s software just flagged his account. The driver—let’s say his name is Sam—shaved off his beard last night for his girlfriend’s birthday. Now the selfie he took this morning—part of Uber’s Real-Time ID Check, rolled out in 2016 to authenticate drivers—doesn’t match his photo ID on record. It didn’t occur to Sam that a discrepancy between the two photos—one showing him with a beard, one without—would automatically suspend his account. But suddenly, and unbeknownst to him, his livelihood hangs in the balance.

Meanwhile, overseas in Hyderabad, the Silicon Valley of India, Ayesha sits at her kitchen table, squinting at her laptop. She just accepted a job routed from Uber to CrowdFlower’s software, and now she is an invisible yet integral part of the ride. CrowdFlower and its competitors with similarly hip-techy names, like CloudFactory, Playment, and Clickworker, offer their platform’s software as a service to anyone who needs quick access to a ready crowd of workers. Tens of thousands of people like Ayesha log on to crowdsourcing platforms like CrowdFlower every day, looking for task-based work. Now Ayesha—and any other invisible workers who happen to have responded to CrowdFlower’s request—will determine whether Sam picks up Emily.

Uber and CrowdFlower are two links in a growing supply chain of services that use APIs and human computation to put people to work. Uber uses CrowdFlower’s API to pay someone to review the results of Ayesha’s work, and, if it passes muster, it will process Uber’s payment to her within minutes. If it doesn’t meet the preprogrammed bar, Ayesha won’t get paid for her efforts, nor will she have any meaningful opportunity to lodge a complaint. The API isn’t designed to listen to Ayesha.

Ayesha compares the two photos of the driver side by side. A timer in the top right-hand corner of CrowdFlower’s webpage winds down,

prompting her to speed up. If she doesn't submit a response before the timer runs out, CrowdFlower won't process Uber's payment for the task. Ayesha blinks, glances at the timer, and squints at the thumbnail-size photos: *Yes, those are the same brown eyes. The same dimpled cheeks.* She clicks "okay."

Sam's account is authorized to pick up Emily just as he pulls up to the curb. Emily stops scanning the congested Chicago traffic and climbs into his car. By the time the car door closes, Ayesha has moved on to the next task. She hopes to net a few more rupees before she ends her workday.

Neither Uber's passengers nor their drivers realize that a person, working far away or perhaps just down the road, might vet their transaction in real time. Imperceptible exchanges like this one determine one out of every 100 Uber pickups in the United States, which means they happen roughly 13,000 times a day. We never saw the ghost work that Ayesha could do for CrowdFlower, but, having spent time with her and workers like her, we can imagine the fleeting market exchanges that consumers like Emily and drivers like Sam will never see. Ayesha is the only artifact of ghost work's presence and, as such, the only one who can help us recover the experience of ghost work after Emily and Sam are long gone.

Billions of people consume website content, search engine queries, tweets, posts, and mobile-app-enabled services every day. They assume that their purchases are made possible by the magic of technology alone. But, in reality, they are being served by an international staff, quietly laboring in the background. These jobs, dominated by freelance and contingent work arrangements rather than full-time or even hourly wage positions, have no established, legal status. Sometimes these jobs are given heft as harbingers of the "Second Machine Age" or the "Fourth Industrial Revolution" or part of a larger digital or platform economy. Other times, they're simply, glibly called *gigs*.⁹

No employment laws capture the on-demand gig economy's odd mix of independence from any single employer and dependency on a web-based platform. As the taskmasters of the gig economy, on-demand platforms make their money by matching those buying and selling human labor online, generating a two-sided market of myriad businesses and anonymous crowds of workers. And, importantly, as media scholar and sociologist Tarleton Gillespie points out, platforms may not create the content that they host, "but they do make important choices about it."¹⁰ On-demand work platforms can easily become silent business partners

more aligned with the interests of those willing to pay a fee to find workers than with the workers searching for jobs.

From the largest firms to the smallest startups, companies rely on this shared pool of on-demand workers amassed by on-demand platforms. They use this assembly of workers to satisfy customers who have grown to expect responses to their requests within seconds. Businesses turn to this pool, instead of traditional temporary staffing agencies, to fill last-minute gaps on their teams. They draw from it to spin up new projects, from testing a new software privacy setting to vetting descriptions of culturally attuned mac-and-cheese flavors. Such ventures are too speculative or loosely understood to justify hiring a full-time employee or the expense of recruiting, even through a temp service. No business wants to invest in launching a new service or product without gauging how consumers will respond. Service industries, driven by the ever-shifting winds of customer taste and satisfaction, can try out ideas generated by ghost work and iterate on responses from other workers, standing in for the average consumer.

Robots Might Be Coming, but They Aren't Here Yet

Every week, another breathless headline proclaims the end of work. Soon, we are warned, the robots will rise up against us. Automation and its handmaiden, artificial intelligence, are widely understood as processes making human labor obsolete. Robotic arms can move sheets of metal across the factory floor. Software bots can take texted pizza orders. Drones can deliver packages to our doorsteps. These intelligent systems, now hitched to many traditional employment sites, are said to herald the rapid disappearance of humans in the workplace. The inevitable triumph of AI, so the story goes, will make all but the most uniquely qualified workers redundant. We all need to skill up. *Now.*

Tesla and SpaceX founder Elon Musk, renowned physicist Stephen Hawking, and Google co-founder Larry Page are just a few of the prominent voices in this chorus.¹¹ Either they express panic about “summoning the demon” of AI or wax nostalgic about a time before AI, when humans supposedly controlled their own destiny.¹² But arresting headlines obscure a messier reality. While it’s undeniably true that robots are on the rise, most automated jobs still require humans to work around

the clock, often part-time or on a contract basis, fine-tuning and caring for automated processes when the machines get stuck or break down, as technical systems, like humans, are apt to do.

It's also true that the long march toward automation has historically created new needs and different types of human labor to fill those needs. In this respect, the new, software-managed work world shares features of the factory jobs that assembled cars by placing workers on a production line where and when they were needed most. It also resembles the so-called piecework that women and children did on farms in the 19th century, assembling matchstick boxes for pennies a pop. And it overlaps in obvious ways with the outsourcing of medical transcription and call center work to the Global South that boomed with the expansion of the internet in the late 1990s.

Factory work, piecework, and outsourcing were all precursors to tasks distributed online insofar as they involved jobs that were small, repetitive, and removed from the bigger picture. These jobs came with little stability or support. They were done, most often, by people whom economists might consider expendable or "low skill." The market calls this, unironically, "human capital." Clicking "dog" or "cat" to label an image that will eventually enable an iPhone to recognize a family pet is not that different from turning a screw on what will eventually become a Ford truck. But that's where the job similarities end.

Blue-collar manufacturing jobs have been the most visible targets of AI's advance. The Foxconn factories that make iPhones allegedly replaced 60,000 humans with robots in 2016. Amazon's 20 fulfillment centers reportedly deployed 45,000 robots to work alongside 230,000 people that same year. Yet these numbers confound how many jobs are *created* by automation. And the media coverage of AI's impact on full-time blue-collar work can distract us from the rapid growth of a new category of human workers to complement or tend to automated manufacturing systems when AI hits its limits.

In the past 20 years, the most profitable companies have slowly transitioned from ones that mass-manufacture durable goods, like furniture and clothing, to businesses that sell services, like healthcare, consumer analytics, and retail. There's more money to be made in selling consumers an experience, from sipping a latte to watching a bit of infotainment, than building a television set.¹³ Businesses of all types manage costs by tapping into and maintaining control of a pool of contingent workers. Having who you want, when you want them, is now a half-century-old strategy for

avoiding negotiations with full-time employees and the classification and employment laws that protect them.

This hybrid of humans and AI reconfiguring manufacturing, retail, marketing, and customer service has outstripped familiar employment categories. Unlike the repetitious lockstep of factory-controlled, full-time manufacturing shift work, these task-based services, such as correctly amending a client's tax return or translating and captioning a video in real time, depend on endless iterations of human discernment and divination that don't fit neatly into a traditional 40-hour workweek. The tasks are dynamic, not merely mechanical, which is why it is difficult to eliminate humans from the task at hand.

AI is simply not as smart as most people hope or fear. Take, for example, the celebrated accomplishments of the AI powering AlphaGo, most recently chronicled in technologist Scott Hartley's book *The Fuzzy and the Techie*.¹⁴ In May 2017, AlphaGo became the first computer program to beat Ke Jie, the reigning world champion of the ancient Chinese board game go. Five months later, AlphaGo fell to its progeny, AlphaGo Zero. But, lest we be too impressed, it's important to keep in mind that the rules of go are fixed and fully formalized and it is played in a closed environment where only the two players' actions determine the outcome. AlphaGo and AlphaGo Zero's human programmers at the Google-backed company DeepMind gave the programs clear definitions of winning versus losing. Winning go is about foreseeing the long-term consequences of one's actions as one plays them out against those of an opponent.¹⁵ So AlphaGo was trained on billions of board positions using a large database of games between human experts, as well as games against itself, allowing it to learn what constitutes a better move or a stronger board position.¹⁶ AlphaGo Zero was then steeped in all of those prior experiences by playing against AlphaGo, a mirror image of self. But, as Tom Dietterich, a noted expert in artificial intelligence research, suggests, "we must rely on humans to backfill with their broad knowledge of the world" to accomplish most day-to-day tasks. Real life is more complicated than a game of go.

The new online work platforms that channel jobs to Joan, Kala, Justin, and Ayesha upend the mediagenic stories about AI's boundless wisdom and the inexorable rise of robots. Real-world tasks, from identifying hate speech or categorizing a rental as a great springtime wedding venue to correctly amending a tax return, require human discernment. Formalizing the singular, best choice, as you might in a game of go, won't work. For example, it would be difficult, if not impossible, to enumerate every

attribute of a wedding venue that would make it the “best.” Even if this were possible, people would have different preferences when it came to the attributes of the venue. Moreover, the training data to teach AI to recognize what counts as the “best choice” does not exist. In addition, an endless set of external factors, from vernacular slang and climate-change-induced hurricanes to haphazard tax reform legislation, can intrude and influence the outcome. In many cases, there are too many unknowns to train current AIs to be aware enough or gain enough experience to intelligently respond to all cases of the unexpected. This is why AI must return to humans to backfill decision-making with their broad knowledge of the world.

Anyone who scrutinizes the shadows of AI, as we have done, will find a new world of work in which software manages people doing jobs that computers can’t do. As builders create systems to transfer tasks from humans to machines, they surface new problems to solve through automation. For example, it was only after the web became mainstream that companies like Facebook, Twitter, and Instagram faced growing demand to moderate their online content, outstripping the limited capacity of automated moderation tools. At the same time, as novel systems are brought online, they typically face unanticipated problems and fall short of their promise, hence the need for Kala’s and Joan’s work. Thanks to workers like them, automated moderation software is better, but it is far from perfect. The inevitable glitches that automated processes encounter along the way to perfection generate temporary work for people. Once they have successfully trained artificial intelligence to perform like humans, workers move on to the next tasks engineers assign them that push the boundaries of automation. Since the finish line moves as people dream of new applications for AI, we can’t be sure if the “last mile” of the journey toward full automation will ever be completed. We call this the “paradox of automation’s last mile.”

As AI advances, it creates temporary labor markets for unforeseen and unpredictable types of tasks.¹⁷ The great paradox of automation is that the desire to *eliminate* human labor always *generates* new tasks for humans. What we call “the last mile” is *the gap between what a person can do and what a computer can do*. Without a doubt, software developers will use ghost work to perform the tasks at hand and push AI to its limits. And it is just as likely that as more companies aspire to give us AI-enabled “smart” digital assistants to manage our calendars and book our flights, we’ll need more and more people to step in when AI falls short of our increasingly exacting and extensive demands. In fact, dependency on temporary human

labor has always been a part of the history of technology's long march toward automation. Today's engineers aiming to solve problems through algorithms and AI are the latest iteration of the paradox of automation's last mile. On this frontier, the peaks and valleys of temporary work shift constantly, redefining relationships between humans and machines in the process.

The rise of on-demand labor platforms signals the allure of using APIs to organize, route, and schedule work. As the examples in this book suggest, this reorientation to use contingent labor to develop new technologies fueled the recent "AI revolution." When an AI system that powers a phone app or online service isn't confident about what to do next for a customer, it needs human help, and it needs it fast. End users expect software running search engines and social media to respond in milliseconds. Traditional methods of hiring won't do here. So if an AI needs a human in the loop, to make sense of a spike in search terms tied to, say, a sudden natural disaster, it needs to get human input immediately. The disaster will fade into history. The software will have learned what it needed from the momentary flood of human input. That is exactly what an always-on labor pool, plugged into APIs, provides. Software developers can write code that automatically hires someone to solve an immediate problem, checks their work, and pays them for doing the job. Similarly, scientists and researchers using modern machine learning systems depend on training data that's clear and error-free. They need an automated method to get help generating and cleaning up that data, and they rely on many people around the world to do it. On-demand labor platforms offer today's online businesses a combination of human labor and AI, creating a massive, hidden pool of people available for ghost work. Delivering services and jobs on demand could be an integral part of the future of work. It could also have unintended, potentially disastrous consequences if not designed and managed with care and attention to how it is restructuring the experience and meaning that people attach to their day jobs.

Ghost Work and the Future of Employment

The dismantling of employment is a deep, fundamental transformation of the nature of work. Traditional full-time employment is no longer the rule in the United States. It used to be that a worker could spend decades showing up day after day to the same office, building a career, with the

expectation of getting steady pay, healthcare, sick leave, and retirement benefits in return. Now, centuries of global reforms, from child labor laws to workplace safety guidelines, are being unraveled. In fact, according to the U.S. Department of Labor's Bureau of Labor Statistics, only 52 percent of today's employers sponsor workplace benefits of any kind. In the wake of the Great Recession, Americans have come to realize that the best alternatives to serving food, providing healthcare, or selling goods in brick-and-mortar shops are the growing number of jobs that can be found in the on-demand gig economy. Because this work doesn't fit any ready-made classification in employment law, the terms-of-service agreements for platforms like MTurk and CrowdFlower are almost indistinguishable from the boilerplate dialogue boxes that we all click to update our software, erasing the protections that traditional workers enjoy.

While the Pew Center's best estimate puts the number of individuals involved in ghost work today at around 20 million, there is no corroborating tally of how many people like Joan, Kala, Justin, and Ayesha cobble together contract-based ghost work gigs to make ends meet. When the Bureau of Labor Statistics added a supplemental survey of Contingent and Alternative Employment Arrangements to the U.S. Census Bureau's May 2017 Current Population Survey (CPS), a monthly snapshot of 60,000 eligible households that provides the nation's employment and unemployment data for the U.S. Bureau of Labor Statistics (BLS), it was the first time it had tried to gauge the growth of contingent jobs in more than a decade.¹⁸ According to the BLS's estimates, 10.1 percent of U.S. workers work without an explicit or implicit long-term employment contract. But this survey counts only people who hold an alternative employment arrangement as their primary or stand-alone job. So if a person does ghost work while also holding down a nine-to-five job with a single employer for a set salary or hourly wage—a very common trend among the most active workers we met—they are even harder to identify, let alone count.

The Bureau of Labor Statistics' 2017 Contingent and Alternative Employment Arrangements supplement to the Current Population Survey poses two hurdles for measuring the rise of ghost work. It is hard to really understand what "long-term employment" means to workers in a multiple-choice survey. It might be as hard to know what "primary job" means when so many people hold down multiple jobs to make their rent. The confusion over how to think about old work categories, like "long-term" or "primary job," is reflected in a head count from the Government Accountability Office that diverges with the BLS's numbers. It reported,

just two years earlier, that at least 31 percent of the U.S. workforce claims that it does some form of alternative work arrangement that includes freelancing or independent contract work for hire.¹⁹ Labor economists Lawrence Katz and Alan Krueger estimate that temporary and alternative contract-driven work delivered through self-employed workers or those temporarily employed by staffing agencies—the so-called casualization of the workforce—rose from 10 to 16 percent, accounting for all net employment growth in the U.S. economy in the past decade.²⁰ The closest we might come to understanding the size and growth of ghost work comes from independent think tanks rather than governmental data.

The most conservative estimates of on-demand gig labor markets come from the Economic Policy Institute. Economist Lawrence Mishel and his research team estimate that between 0.5 and 1 percent of working adults in the U.S., or 1.25 to 2.5 million people, participate in the on-demand gig economy. But they come to that number through a very specific study of Uber drivers and the assumption that Uber and other ride-hailing mobile apps make up the bulk of gig work. A study produced by the JPMorgan Chase Institute found that 4.3 percent of U.S. adults, or 10.73 million people, had worked an online-platform-economy job at least once between 2015 and 2016.²¹ A revolving door of temporary tasks defines this job market. No obvious professional title. No ladder. No bonuses. No guarantees. Tasks are finite, built to disappear once a firm has reached its specific target and the people hired to hit it have moved on to other projects.

From software engineering and legal services to commercial media and healthcare, a wide range of businesses now turn to on-demand labor platforms to convert white-collar careers into bundles of projects. Such all-digital information services and knowledge work convert the creative expertise required to think with and massage data into the consumable services delivered online by industries from tech and law to finance and entertainment. Because of these seismic shifts, the days of large enterprises with full-time employees working on-site are numbered. A crowded field of companies compete to sell information services that pair computers and smart devices with artificial intelligence. Companies like Catalant (formerly HourlyNerd), Popexpert, and Upwork use APIs to deliver the larger “macro-tasks” of knowledge work, on demand, to other businesses or individuals. The future of employment wrought by automation will undoubtedly be far more disjointed than traditional nine-to-five work. Some labor economists argue that a new reality of “fissured workplaces” is the ultimate result of turning long-term employment into a series of short-

term contracts throughout the 1980s and '90s.²² And yet this newly unpredictable reality hasn't dissuaded millions of digital workers around the world from sitting down at their keyboards day and night and performing the countless behind-the-scenes tasks that make our apps seem smarter than they are. This means that the future of business and employment will more likely resemble today's on-demand economy than a dystopian sci-fi film in which humans disappear and robots rule. It will require people to navigate layers of software interfaces and learn to labor in the shadow of AI. It will contain an ecosystem of independent contractors like Joan, typing away in spare bedrooms, cafés, and cinder-block homes in rural India, Knoxville, Tennessee, and Portland, Oregon—or anywhere else a person with an internet connection, a computer, ambition, or financial need can get online. When little attention is paid to the workers behind these jobs, on-demand labor can quickly become alienating, debasing, precarious, and isolating ghost work.

All of the workers we interviewed have something unexpected in common: hope. They *hope* to use on-demand jobs to control when they work, who they work with, and what tasks they take on. They *hope* to stay close to their families. They *hope* to avoid long commutes and hostile work environments. And they *hope* to gain experience that refreshes their résumé or opens a door to new possibilities. Also true is that many saw few other options for themselves or their families. Full-time employment in their towns often meant an hourly wage at a big-box store, working a fixed shift, adapting to unpredictable work schedules, and without meaningful opportunities to advance. On-demand jobs gave them real-world experiences scheduling meetings, testing and debugging websites, developing computer expertise, finding sales leads, and managing full-time employees' HR files. What worker doesn't hope to one day fully control both the schedule and the purpose of their workdays?

Ghost Work draws on a five-year study in which we—an anthropologist and a computer scientist and the research team we mustered—investigated this booming yet still largely hidden sector of the economy.²³ It is the culmination of more than 200 interviews and tens of thousands of survey responses collected from workers across the United States and India; dozens of behavioral experiments and social network analyses of on-demand work platforms; and unique studies of this labor market's other key players, namely the people turning platforms into businesses and those hiring workers on them. It exposes a world in which steady work and salaries are being replaced by a chaotic string of small projects and micropayments, and human bosses are being replaced by automated

processes that are programmed to oversee a far-flung workforce of anonymous independent contractors. *Ghost Work* departs from the well-known story about the rise of robots by documenting a more complicated future that is already emerging. It shows how ghost work platforms foster our belief in the magical promise of technology.

As an anthropologist, Mary had her interest sparked by the specter of an atomized world of workers earning money by sorting and annotating thousands of pictures of pointy-eared dogs, hairless cats, and “dick pics.” When Mary asked those hiring workers what they knew about the people picking up their tasks, the responses ranged from “I don’t know” to “Why would I want to know that!?” As a computer scientist, Siddharth had used on-demand platforms for years to conduct online behavioral experiments, but he knew little about the workers, as the API kept them hidden from him.²⁴ Who were the people offering themselves up for hire? What motivated them to do what many consider “mindless tasks,” and how did they make this ill-defined form of employment pay off? What did this work mean to them? How many tasks flow online through these on-demand platforms? What are the business models that produce the demand for task-based work? What are the overall workings of this task-based economy?

When our research team started asking these questions in 2013, the only people in the conversation were economists, computer scientists, and businesspeople. All three groups evaluated the on-demand labor market on the basis of its ability to enhance efficiency and maximize a company’s bottom line. When humans did happen to come up in the discussion, it was in reference to the consumer. What was the quality of the consumer’s experience? The engineers and computer scientists building APIs, for companies or for their own experiments to advance AI, wanted to design systems that eliminated what they assumed were costly, superfluous operations that annoyed end users. They were in the business of building smarter, faster software that could automatically match people to services, whether it was a ride, a meal, or tax advice, with an end goal of using the data from each iteration to train future software to automate even more. Few people were tracking what this approach to productivity would mean for the people who vied to do task-based work for hire. They operated from the assumption that the workers needed to generate training data and improve software would disappear once the AI got things right. Companies were building software, after all, not temp jobs.

For the next five years, we did something our respective research fields had not: we learned about the range of ghost work and the lives of people

doing it by conducting one of the most comprehensive studies of its kind. *Ghost Work* is the first book to illuminate ghost work's role in building artificial intelligence and the lives of workers who are invisible yet central to the functioning of the internet and the future of automation. It offers an intimate, detailed look at the experience of workers in this new economy. We focus on workers living in India and the United States, the two countries with the largest on-demand labor pools, both with a long, entwined history of technological advancement. Our team interviewed and observed hundreds of people, in their homes and other makeshift workspaces, as they did everything from flag tweets to transcribe doctors' visits. We surveyed thousands more to establish a baseline to help us gauge which practices were typical and which were exceptional. We then scaled up the findings from our interview data by conducting dozens of behavioral experiments and "big data"-style analyses, each with thousands of participants. Throughout *Ghost Work*, the reader will see us toggle between these two types of analysis, combining their strengths to shed more light on those who work in the on-demand economy.

We examined four different ghost work platforms: Amazon.com's Mechanical Turk (MTurk); Microsoft's internal Universal Human Relevance System (UHRS); the socially minded startup LeadGenius; and Amara.org, a nonprofit site dedicated to translating and captioning content for transnational audiences and people with hearing disabilities. Each of these four platforms offers different products and business models. Investigating them alongside one another helped show us that our observations and conclusions hold broadly across the on-demand economy, as opposed to being specific to one category of ghost work. MTurk, as one of the first commercially available ghost work platforms, set the norms for how others would apply human computation to business solutions. UHRS stands in for the internal platforms that every large tech company maintains to meet its own ghost work demands. LeadGenius and Amara illustrate just how complex and sophisticated ghost work can be, as well as how much companies can play a role in designing better conditions for ghost work.

And then there were the workers. Among those working on these platforms, we met people stringing together on-demand projects to recreate the work hours, pay rates, and career development associated with full-time employment. We also met college-educated, stay-at-home parents staving off boredom; first-generation college students working 50 hours a week to save money for a wedding or fund a younger sibling's degree; and people, disabled or retired, looking for alternative routes to

employment or extra money to pad their social security checks. We also met engineers and entrepreneurs who founded, designed, and built ghost work platforms.

When we started, we wondered: Who are these people, and how does their work differ from traditional nine-to-five jobs? On many on-demand labor platforms, a requester like Siddharth sees no personal information about a worker—gender, location, age, and prior work experience are all unknown. And workers have no information about the requester beyond the task description. The range of tasks can be endless and can change from one day to the next. APIs can be used to have a human tag a cat photo or run a research experiment, and similar APIs can be used to hire someone to deliver a meal, send a car, or design a website. The moment that the API is called and the work is produced looks automated to both consumers and requesters. But who benefits from this veneer of automation? And who might be harmed?

By the time we finished our study, we understood that people doing ghost work were no different from our friends and family making a living through freelance writing, research, software development, or adjunct teaching. Their work lives were often vulnerable and insecure. Yet the anonymity and remote access of on-demand platforms also made it easier for those marginalized in formal employment—because of where they lived, a perceived disability, or their belonging to a stigmatized minority—to earn an income.

The more closely we looked at the nascent edges of on-demand work, the more we saw people using familiar strategies to stay afloat and create meaningful employment for themselves and their peers. Sometimes these workers succeed by collaborating with one another. They share strategies for making difficult tasks easier, they swap intel about those with tasks for sale, and they help one another stay awake as they wait for new tasks to come online. We met workers who learned to move forward after their failed forays. Who learned to thwart exploitative business models, labor laws, and APIs designed to be indifferent to their interests. And we noted that businesses have no clue how much they profit from the presence of workers' networks. This book describes the thoughtless processing of human effort through APIs as *algorithmic cruelty*—literally, computation incapable of thought, let alone empathy. People doing ghost work understand the perils and potential of on-demand work better than any engineer, tech company CEO, policy maker, or labor advocate. They live it every day. And they are the most invested—economically and psychologically—in making it better.

Just as we need companies to be accountable for the labor practices that produce our food, clothes, and computers, so should the producers of digital content be accountable to their consumers and workers. We should demand truth in advertising in cases where humans have been brought in to benefit us—whether it is to curate our news or field complaints about what some troll just posted to our favorite social media site.

Along with a call for transparency, *Ghost Work* holds lessons for tech entrepreneurs who want a productive workforce, engineers who are building the labor platforms of the future, and policy makers charged with shaping this new commercial landscape. But the still untold story of the invisible workers who power the apps on our phones and the websites we look at should interest a wide range of general readers who've seen some coverage of "gigging it" or "Turk work," not to mention "crowdsourcing" and "microwork," and heard *a lot* about the rise of robots but want a deeper look at how, exactly, AI reshapes the working world and what, precisely, people do in the shadow of it. We offer a textured, nuanced, and ultimately hopeful account. Among other things, we show how moving beyond the full-time-freelance divide alone could go a long way toward sharing the wealth generated by the internet with those tasked to grapple with the paradox of automation's last mile. We hope, too, that the lessons we learned from the many workers we interviewed in the U.S. and India will help the millions of people who already, or will soon, do this work make the most of it. More than anything, *Ghost Work* is for anyone who works and wants to see what their future holds.