

GHOST WORK

How to Stop Silicon Valley
from Building a New
Global Underclass

Mary L. Gray and Siddharth Suri

Houghton Mifflin Harcourt

Boston New York

2019

Copyright © 2019 by Mary L. Gray and Siddharth Suri

For information about permission to reproduce selections from this book, write to trade.permissions@hnhco.com or to Permissions, Houghton Mifflin Harcourt Publishing Company, 3 Park Avenue, 19th Floor, New York, New York 10016.

Maps © Gregory T. Minton. Used with permission. All rights reserved.

Ming Yin, Mary L. Gray, Siddharth Suri, and Jennifer Wortman Vaughan. "The Communication Network Within the Crowd." In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, pages 1293–1303, Republic and Canton of Geneva, Switzerland, 2016. © International World Wide Web Conference Committee. Used with permission. All rights reserved.

Sara Constance Kingsley, Mary L. Gray, and Siddharth Suri. "Accounting for Market Frictions and Power Asymmetries in Online Labor Markets." *Policy & Internet*, 7(4):383–400, 2015. © 2015 Policy Studies Organization. Used with permission by Wiley Periodicals, Inc.

All rights reserved.

hnhco.com

Library of Congress Cataloging-in-Publication Data

Names: Gray, Mary L., author. | Suri, Siddharth, author.

Title: Ghost work : how to stop Silicon Valley from building a new global underclass / Mary L. Gray and Siddharth Suri.

Description: Boston : Houghton Mifflin Harcourt, 2019. |

Includes bibliographical references and index.

Identifiers: LCCN 2018042557 (print) | LCCN 2018044155 (ebook) |

ISBN 9781328566287 (ebook) | ISBN 9781328566249 (hardcover) |

ISBN 9780358120575 (int. ed.)

Subjects: LCSH: Labor supply — Effect of automation on. | Automation —

Economic aspects. | Artificial intelligence — Economic aspects. |

Technological unemployment.

Classification: LCC HD6331 (ebook) | LCC HD6331 .G826 2019 (print) |

DDC 331.1 — dc23

LC record available at <https://lcn.loc.gov/2018042557>

Book design by Chloe Foster

Printed in the United States of America

DOC 10 9 8 7 6 5 4 3 2 1

For Ila and George

— M.L.G.

To my family and in loving memory of my dad

— S.S.

Humans in the Loop

The Beginning of Ghost Work

In the early 2000s, Amazon.com was in a tight spot. A young startup at the time, it had a rapidly expanding marketplace of books just as “e-commerce” was taking off. To build out the online bookstall, Amazon electronically pulled data about millions of book titles directly from publishers’ catalogs and library listings. Much of it was rife with incorrect database entries. To build a loyal customer base, the company desperately needed to locate listing duplicates, typos, and outdated dust-jacket images linked to new editions before it could sell them to a still cautious segment of consumers willing to buy products online.¹ At first, Amazon hired temporary workers to clean up its databases. Following in the footsteps of neighboring tech companies, Amazon hired temps in both the U.S. and India, where it could get native English language fluency at low pay rates through staffing agencies that hired and managed contract workers. Not only did workers correct book titles, publication dates, and descriptions; they also made sure book covers matched the editions listed, and they embedded keywords in each book’s web page, which were designed to return as many relevant items as possible when customers used Amazon’s search bar.²

Then, in order to become the largest online retailer, Amazon would need to conquer a related challenge. Rather than stock millions of products, the company populated its site with small and medium-size busi-

nesses willing to list their inventory, from electronics and toys to cleaning supplies and niche foods, and sell their products through the Amazon site.³ As Amazon grew its online marketplace to include other booksellers and then other products beyond books, it faced the additional challenge of ensuring that every description of a product matched its accompanying image. All of these merchants, including Amazon, would need a small standing army of workers willing to handle these discrete, repetitive tasks — like verifying that product descriptions matched the photo and creating captions and keywords to help online shoppers browse the ever-growing catalog of goods. Amazon turned to its vendors to hire contract workers to fill the labor demand. And as Amazon.com grew, it also needed to polish customers' book reviews so that awkward or unclear wording and syntax didn't muddy the value of customers' comments on its site. Soon contract workers did that work, too.

In 2005, Amazon.com publicly debuted a website it had built to make it easier for anyone with a verified account to clean up product listings and customers' typo-ridden reviews. The company called it Amazon Mechanical Turk, a name that users quickly shortened to MTurk. MTurk was an online labor market where "requesters" could post various tasks they needed done and workers could do those tasks for pay. The platform would list tasks and pay rates, as easily as you might post a job on Craigslist. The platform also operated as a bank so that those with work could fill up their payment accounts and keep a tab that could be used to automatically pay workers once they turned in their projects.

Amazon added a percentage surcharge to whatever the requesters paid workers for each task, charging requesters extra for matching them with workers guaranteed to have certain qualifications. Anyone willing to share their bank account, credit card information, and verifiable mailing address with Amazon could sign up to work on Amazon Mechanical Turk and earn credit toward gift cards on the virtual superstore. Payouts for tasks ran anywhere from a penny, for adding keywords to a specific image, to \$25, for doing a single marketing survey. It's rumored that MTurk was a pet project of Jeff Bezos himself. Amazon's founder is said

to have created MTurk so that Amazon could not just offer a marketplace of books and other durable goods but make labor itself a service that anyone could find and pay for through the Amazon website.⁴

In its first two years — perhaps a signal that the global recession was beginning to build — more than 100,000 people created accounts to find work on the MTurk platform. MTurk automated the hiring and paying of these workers, allowing any company or individual with a set of tasks piling up to issue them to a crowd that was logged on and looking for work. Not long after the platform picked up business beyond Amazon's own product teams, MTurk ironed out the logistics of direct deposit for cash payments to workers with mailing addresses in the U.S. and, for workers based in India, cutting paper paychecks that converted U.S. dollars to Indian rupees, routed through the Bank of Singapore. Once MTurk could pay out in cash instead of gift cards, a vast pool of U.S. and Indian workers vying for posted jobs were pitted against other workers who were more casually topping off Amazon gift cards from other countries.⁵

MTurk filled a much-needed niche. Companies and individuals putting more and more of their products online needed some way to check the accuracy of their posted materials. Increasing numbers of people responsible for entering their receipts into expense reports for reimbursement at work could now turn to services like MTurk for quick help from a human able to take on the task, any time in the day, any day of the week. Startups at the time, like Yelp and those contractors hired to write and curate content for their databases, could offer accurate restaurant locations for neighborhoods that had never had such detailed listings before. Marketing and PR agencies could distribute short surveys to get hundreds of reactions to new product ideas, slogans, and word associations for less than the hourly pay of a full-time employee or temporary worker. Academics could now survey a broader population, sending out 1,000-person polls in a matter of an hour, with results comparable to what they might get from distributing a similar poll to an introductory undergraduate course.

Even better, MTurk's surveys were more likely to reach a broader age

and geographical demographic than the 18- to 22-year-olds found on most U.S. college campuses. While popular online classified listings like Craigslist in the United States had always included ads for work that people could do online, MTurk represented something entirely different. The platform offered contract work for tasks that required very little qualification or advanced computer experience. One needed time, attention to detail, and an internet connection. Whether for marketing, surveying, generating training data, or reviewing online content in real time, results from on-demand labor market platforms like MTurk could be obtained faster and more cheaply than anything collected from employees in the office. Soon, myriad new businesses formed, looking to profit from this mix of simple computer programming, a web interface, and unregulated hiring practices, producing a powerful new way to automate recruiting humans for ghost work.

Ghost work fueled a revolution in artificial intelligence through millions of people carrying out billions of small tasks hidden from technology consumers. Other companies soon figured out how to use ghost work to complete larger work projects, which we call “macro-tasks.” In either case, ghost work powers many popular websites and mobile phone apps today while keeping the workers hidden behind the APIs used to hire them.

Using APIs to Hire People

Anyone with access to MTurk’s application programming interface (API) could plug into MTurk’s growing pool of registered workers. Software developers, whether Amazon’s own or those working for another company, could now write software to place tasks on the MTurk platform, making it easy to hire workers, evaluate their work, collect their projects, and pay them — all within a matter of seconds.

Programmers had previously only written code for machines to execute. But MTurk’s innovation allowed humans to execute part of the

programmer’s code, as opposed to having only machines executing the code.⁶ The brilliant breakthrough of MTurk was that it took this basic technology — batching a series of tasks so that they could be completed by humans through an API — and turned it into a labor market where people could “buy” and “sell” human labor. Now the same piece of software could simultaneously draw on human creativity and a computer’s ability to crank through the same or similar tasks over and over. A programmer’s software and MTurk’s APIs effectively operated as managers of an on-demand contingent workforce. In the process, APIs and web-based platform interfaces, replicating MTurk’s business model, seemed to eliminate much of what many of us expect from our bosses — feedback, scheduling, workspace, payment, and affirmation that we’ve done a proper job and completed our tasks. In doing so, MTurk unraveled the role of “employers” and turned programmers and companies seeking immediate help with tasks into “requesters.”

The foremost implication of this new way of doing work is that the API determines the dialogue and communication between the programmer and the worker. For example, the API gives each individual requester and worker their own unique identifier, a string of seemingly random letters and numbers such as “A16HE9ETNPNONN.” From the programmer’s perspective, it makes the humans seem interchangeable, as each worker is represented by a worker ID, and everything that makes a human a person, such as their beliefs, attributes, and experiences, is stripped away from this identifier.

Computer scientists would say that all of your attributes are *abstracted* away. It’s as if you were hired by someone who knew only your social security number and absolutely nothing else about you. The abstraction resulting from the API makes it appear as if there is no need to figure out who the humans are. The same way that poker chips can make gamblers forget that they are gambling with actual money, representing people as unique identifiers can make programmers forget that they are hiring people and their code affects people’s lives. Amazon’s somewhat callous reference to the 19th-century “Mechanical Turk” chess automaton wasn’t

as arcane as it might seem.⁷ Amazon meant to draw parallels between its service and the enigmatic parlor game that toured for more than 80 years after its creation, in 1770. Ironically, the Mechanical Turk turned out to be a hoax, no more than a series of expert, small-statured chess players hiding inside the machine's wooden case. Humans in the loop — not machines — were the masterminds behind the automaton's chess moves. And, as the name suggests, it takes human intelligence to push the boundary of what machines can learn. APIs are the perfect taskmasters for teaching machines to advance AI.

Ghost Work, Machine Learning, and the Rise of AI

Computer scientist Kevin P. Murphy defines machine learning as “a set of methods that can automatically detect patterns in data, and then use the uncovered patterns to predict future data.”⁸

Recall the machine learning problem, from the introduction, of recognizing a camelback couch. A common machine learning approach would be to first gather what is called training data, in this case by gathering images of couches, say from furniture catalogs and social media posts, and having humans like Justin label them as “camelback” or “not camelback.” Then the machine learning algorithm would compare a new image of a couch against images in the training data. If it looks more like a camelback couch, the algorithm would classify the new image as such. Now say the lighting in the new image is bad, or the angle doesn't show its back clearly, or there are people sitting on the couch and obscuring its back, so that the machine learning algorithm might not know how to classify it. This is where yet more human help might come in.

ASSEMBLING IMAGENET

The overall goal of AI is to build computer systems with intelligence — the ability to evaluate and to act — comparable to what you might expect

from another human being. Understanding what objects are in an image is a part of the ambitious revolution reaching for general artificial intelligence. After all, even a one- or two-year-old can recognize an apple or a dog in an image. Fei-Fei Li, a computer science professor and co-director of the Stanford Human-Centered AI Institute, and her colleagues wanted to solve a much more general problem than how to train AIs to recognize a specific object, like a couch. They wanted to train machines to recognize the main object in an image, no matter what that object may be — a dog, a person, a car, or a mountain. To do it, they needed more training data than a single person could generate alone. *Much* more.

Li and her colleagues first wrote software to download millions of images from the World Wide Web. At first, they hired a team of undergraduates to label each image — the academic equivalent of hiring a temp worker. After testing this method, they could extrapolate how long it would take to complete — about 19 years. So they switched strategies. Next they tried developing machine learning algorithms to automatically guess labels for images and turn to human help for the ones that stumped the machines. This approach failed, because the machine learning algorithms made too many errors, and they were looking for highly accurate or “gold standard” data that other scientists could later reuse. Indeed, if this problem were easily solvable by machines, they wouldn't have needed this data set in the first place.

Shortly after, in 2007, Li and her colleagues found MTurk, and they realized that the MTurk API gave them a way to automatically distribute image-labeling tasks to people and pay them. They tried a few different workflows but were ultimately able to use about 49,000 workers from 167 countries to accurately label 3.2 million images.⁹ After two and a half years, their collective labor created a massive, gold-standard data set of high-resolution images, each with highly accurate labels of the objects in the image. Li called it ImageNet. Thanks to ImageNet competitions held annually since its creation, research teams use the data set to develop more sophisticated image recognition algorithms and to advance the state of the art. Having a gold-standard data set allowed researchers

to measure the accuracy of their new algorithms and to compare their algorithms with the current state of the art. This allowed researchers to make so much progress that some AIs can now do a better job than humans in recognizing images!¹⁰

The algorithmic and engineering advances that scientists achieved during the competition, between 2010 and 2017, fueled the recent “AI revolution,” which had an impact across a variety of fields and a variety of problem domains. The size and quality of the training data were vital to this endeavor. MTurk workers are the AI revolution’s unsung heroes. Without them generating and improving the size and quality of the training data, ImageNet would not exist.¹¹ ImageNet’s success is a noteworthy example of the paradox of automation’s last mile in action. Humans trained an AI, only to have the AI ultimately take over the task entirely. Researchers could then open up even harder problems. For example, after the ImageNet challenge finished, researchers turned their attention to finding *where* an object is in an image or video. These problems needed yet more training data, generating another wave of ghost work. But ImageNet is merely one of many examples of how computer programmers and business entrepreneurs use ghost work to create training data to develop better artificial intelligence.¹²

The Range of Ghost Work: From Micro-Tasks to Macro-Tasks

The platforms generating on-demand ghost work offer themselves up as gatekeepers helping employers-turned-requesters tackle problems that need a bit of human intelligence. Businesses no longer had to turn to contingent staffing agencies to access a global labor market. MTurk became known for “micro-tasks,” like the ones that Fei-Fei Li’s team generated, that could be done quickly but required many people. But a host of businesses have cropped up in recent years that match workers to larger projects called “macro-tasks.” These platforms, like Upwork and Fiverr,

where you can find people to copyedit a newsletter, develop webpages, or build a mobile app, use the same employment strategy: distribute tasks to a pool of internet-connected workers who are hired, scheduled, managed, and paid, at least in part, by AIs and APIs. All those paid to complete tasks do a variation of platform-based ghost work. And, as of today, everyone is operating outside the legally defined classifications we have for employment. That is, there are no laws governing who counts as an “employer” or “employee” in ghost work. And it is unclear where the platforms, which are the places the workers go to find work, stand. But it is clear that the platforms have become the de facto job sites of on-demand labor. It’s very difficult to see what it looks like to do this work until you meet the people behind the APIs.

MTURK: THE PUBLIC FACE OF MICRO-TASKS

Joan, whom we met in the introduction, wears her hair in a loose bun skewered by shiny black chopsticks to keep it out of her eyes while she’s working. She’s been living in Houston since 2011, when she returned to care for her 81-year-old mother. Joan cooks, keeps house, and drives her mother to doctor appointments. And for the past three years, she’s made most of her income working on Amazon’s Mechanical Turk.

Before moving back to her hometown, Joan had a full-time job as a technical writer. She drafted and copyedited, among other things, manuals for filing for unemployment insurance in the state of Texas. At first, Joan lived off the money she’d cashed out of her 401(k) plan. But as her mom’s health worsened, Joan looked for work she could do from home. On-demand work seemed like a good fit. Joan turned a spare bedroom into a home office, crowding the small room with a weathered brown chair, computer desk, and large monitor. Then she started searching the internet for work that she could do online.

Joan can’t remember how she first found out about Amazon Mechanical Turk, but she suspects that she learned about it on Reddit. Reddit is one of several online communities where people doing ghost work share

tips on how to get started. As a 39-year-old white woman with a master's degree in communications, Joan is, in some respects, a typical MTurk worker. Almost 70 percent of workers have completed a bachelor's degree or higher educational attainment. In other ways, she stands out. MTurk workers skew young: 76.9 percent are between the ages of 18 and 37, that bracket of years when people are most actively seeking their first career-defining job.

Though Joan doesn't recall all the details, the process for setting up a worker account hasn't changed since she first signed on. She would have gone online, navigated to the main MTurk website, and clicked on the sign-up button. As a newcomer, Joan would have been asked to enter a verifiable name, email address, and password. From that point, she would have been given access to the behind-the-counter side of the site. Visible from Joan's "dashboard" would have been dozens of tasks. Tasks, or what Amazon refers to as HITs (Human Intelligence Tasks), are jobs for hire. If she had clicked on a task, it would have shown her a short description of what the task required, the deadline, and what it paid. She could click on and complete a task, but, as a new user, she'd need to wait until her account was validated to get paid. Before Amazon pays a worker, it verifies the person's physical mailing address, national identity, and bank account information. That's how easy it was for Joan to join the ghost workforce.

To a new worker like Joan, MTurk's dashboard can look chaotic. One sees multiple expandable menu tabs, including a tab to keep track of one's account, another to track the individual tasks, and a tab that lists the worker's "qualifications." That word doesn't align with skills. In the world of MTurk, qualifications can be things like a worker's age, gender, and location. People who post jobs on Amazon use "qualifications" to restrict the type of worker who can accept the job. For instance, if an advertising company is looking for a focus group to give it feedback on a product meant to appeal to women in their forties, it might add qualifications such as gender and age to the job. It can even pay Amazon an extra premium fee for workers who list qualifications like "smoker" (30 cents)

or "2016 voter" (10 cents). When Joan first looked at her MTurk dashboard, she remembers, she felt a bit overstimulated, but not deterred. "I thought, Okay, this is not going to pay out at the beginning, but if I do it for a while, it may become a decent source of side income," she says.

No one knows the exact number of people who use MTurk, but typically about 2,500 workers are actively either searching for tasks or completing tasks on the platform.¹³ Because no agency — like a labor union or the Department of Labor — tracks this information, big-picture numbers are equally hard to pin down. Amazon maintains that it has 500,000 registered MTurk workers. According to researchers, anywhere between 100,000 and 200,000 people are registered to work on MTurk.¹⁴ Panos Ipeirotis, a leading researcher most known for his work tracking the ebbs and flows of MTurk worker demographics, estimates that 2,000 to 5,000 workers can be found on the MTurk platform at any given moment. That is roughly the equivalent of a 10,000-to-25,000-person full-time workforce.¹⁵ If we apply this logic to every on-demand platform, there are potentially millions of full-time jobs in the shadows of ghost work. This, of course, assumes that people would want to do this work full-time. However, as will become clear, a sizable percentage of workers stick with on-demand ghost work precisely because it does not demand a full-time commitment.

We posted a task on the MTurk platform to understand how workers are distributed around the world. Upon accepting the task, workers were shown a Bing map of the world and told, "Just double click your location and submit the HIT — it's that simple." Over ten weeks, 8,763 workers across the globe self-reported their locations. Workers are distributed throughout the United States in both highly and sparsely populated regions, but Indian workers are concentrated in the southern part of the country, a point we'll return to in the next chapter (see figures 1A and 1B).¹⁶

Like most workers we met, Joan starts her day looking for tasks. One of the tasks she does the most is text categorization. She might read a snippet of text, perhaps a sentence or two from a news story, and either create a category for it or pick "politics" or "sport" from a list of options

presented to her. The first time we spoke with Joan, she was doing one such task. For every data point she categorized, she made two cents. She classifies tens of thousands of pieces of text every week.

Joan spent the first six months on MTurk finding her footing. In time, she learned that the trick to making decent money was to quickly find doable work and to evaluate the requester offering the job. She noticed that on MTurk, every second counted; a slow internet connection, time spent finding work, or any unplanned downtime was the equivalent of lost income. In her first year on MTurk, she made \$4,400. Some people might see that number as insignificant, she says, but "\$4,400 is a meaningful amount when your previous income was zero." Two years later, her MTurk earnings had almost quadrupled, to \$16,000. Joan is now among the 4 percent of MTurk workers who are skilled, practiced, and lucky enough to earn more than \$7.25 an hour completing tasks.¹⁷

Hypervigilance is a necessity for top earners. Those doing ghost work who make the most money spend hours monitoring their dashboards and scrolling through pages upon pages of job postings. Joan, like so many others who are trying to make MTurk a core source of income, turns to free software tools and workers' online forums to reduce some of the search costs that are an unpaid part of the job.¹⁸ They must be ready to snap up a well-paying or fast-and-easy task the second it pops on their screen, lest another worker click on the link and accept it first. "I've worked harder at this than I ever did at any office job," she says. To enhance her speed, Joan arranged her web browser's display of the MTurk dashboard to show 25 tasks or HITs at a time, and she uses keyboard shortcuts she created to flip through the pages quickly.

When Joan is in the zone, she can complete about 1,100 HITs an hour, netting roughly \$22 an hour. She knows that people are likely to assume that the work is mind-numbing, but she finds the variety of tasks intellectually stimulating. She especially enjoys work that involves editing, which plays to her strength as someone with a background in technical writing. "I'm good at it and it's easy to do," she says. When the work does feel mundane or repetitive, she stays alert by listening to techno music or

watching television. When we spoke, she was working her way through several seasons of *Top Gear*, a show for car lovers. "People talk about 'Netflix and chill,'" she says, "but I watch Netflix and MTurk."

MTurk set a minimum fee for work at one cent per task and, from there, requesters decide how much to offer MTurk workers for each assignment. On average, requesters price tasks to offer the equivalent of \$11 an hour, but low-paying requesters flood the market with minimum-fee work, which drags down the overall earning potential of workers, who must wade through lists of poor-paying tasks to find decent work. "It's a constant race to the bottom," says Joan. By some estimates, the total revenue of the requesters on MTurk and similar sites like CrowdFlower adds up to \$120 million per year.¹⁹ Workers keep what requesters pay out, but Amazon charges the requesters 20 percent of what MTurk calls "the reward" — a worker's paycheck, including any bonus amount (the equivalent of a tip) — as its fee for operating the platform. Amazon charges an additional 20 percent for HITs that require ten or more workers.²⁰

Unlike a traditional employer-employee relationship, MTurk workers are largely anonymous and mostly autonomous, meaning that a requester cannot specify the people who will carry out the work nor dictate exactly how the task is completed once it's been accepted by a worker. Workers alone are responsible for the taxes on their MTurk income. They are expected to file as independent contractors, the 1099 forms familiar to anyone in the freelance consulting world. The trade-off for the requester is that the work is done fast and without the associated costs of officially hiring an employee. The trade-off for the worker is that they don't have to stick with the same job any longer than it takes to complete the task. They can fit work around the demands of their lives rather than hand their lives over to the long commutes or hostile environments that come with some nine-to-five jobs. And they can stop working the second they've made the money they need to make. But the completion of a task does not always equal a payday. The work submitted by MTurk workers is reviewed by a human being or an algorithm that either deems the

work satisfactory or rejects it. If the work is rejected, the worker isn't paid. Each worker's approval rating, the fraction of tasks they have had accepted, serves as a reputation score on the site. Many tasks on MTurk require workers to have approval ratings of over 95 percent, so even one rejection can seriously affect a worker's ability to earn money by limiting their access to future tasks.

Like all people doing ghost work, Joan must weather fluctuating income streams. Requesters can bring booming business one day and disappear the next. Not long after she'd signed up with MTurk, Joan got a string of decent-paying tasks posted by Taste of the World, a pseudonym widely rumored among workers to belong to the popular travel site TripAdvisor. Taste of the World posted hundreds of thousands of tasks on MTurk, jobs like removing duplicate hotel listings, validating website links, writing descriptions of top travel destinations, creating city-specific lists of the best places to eat, and cleaning up typos. The average Taste of the World task could net an experienced worker the equivalent of \$10 an hour, and working for the requester had other perks, too. "The work was available just about every day . . . and it was posted hours at a time," says Joan, meaning she didn't have to jump on it before it disappeared. She could step away from her computer to make dinner and, when she returned, the Taste of the World tasks were still available, because the sheer volume of work was so large. But just as abruptly as the jobs arrived, they dried up. Joan told us that less than a year into using MTurk, Taste of the World posted to MTurk Forum that "we have enough people." Joan flatly added, "And that was the end."

When it comes to paying for ghost work, each platform operates a bit differently. Amazon, in some ways, operates like both an ATM and a company store. New workers on MTurk must wait out the initial ten-day holding period before they can claim any money made doing tasks. After successfully submitting ten days of requester-approved work on the platform, U.S. workers have a choice: they can receive the full value of their earnings in the form of an Amazon.com gift card or they can transfer their paycheck into an Amazon Pay account. From an Amazon Pay

account, workers can then transfer their earnings to a personal bank account, but they have to pay a transfer fee to Amazon for the privilege. International workers, with the exception of citizens of India, can only convert their earnings into an Amazon.com gift card.

Indian citizens are the only international workers who can also earn cash for their MTurk ghost work. There's no reason for this other than the fact that Amazon's multinational holdings allow it to operate and transfer money between its U.S. and India office locations. India's MTurk workers could opt to fill up an Amazon gift card, though the company doesn't reliably deliver to many of India's sprawling, informal neighborhoods. If a worker in India wants to move their money to a personal bank account, they must first hand over their birth date and a scanned copy of their permanent account number (PAN) card, the equivalent of a U.S. social security number. It takes Amazon a week or more to verify PAN card information. Once that is done, workers in India have one more bar to clear: they must send their bank account information to Amazon for verification. Once it's verified, Amazon can, for an additional fee, cut paper checks or offer direct deposits to India's MTurk workers.

For her part, Joan didn't plan to turn MTurk into a full-time job. It just happened. Now she has settled into the life of an independent worker, and her long-term goal is to create financial stability by knitting together several sources of income. This was a common theme among workers we met. And, indeed, 75 percent of workers on MTurk report having at least one other source of income. In Joan's case, in addition to her work on MTurk, she spins her own wool and sells knitted crafts at a local market. She's also ramping up her technical writing skills, with the goal of having a more competitive freelancing profile on the macro-task site Upwork. And she thinks about getting a part-time telecommuting job, like online customer service work, but she hasn't figured out how to do that while taking care of her mom. And, like 75 percent of MTurk workers, Joan does ghost work on other platforms, including Microsoft's UHRS, even though she notes that her primary source of income, for more than a year, was MTurk.

UNIVERSAL HUMAN RELEVANCE SYSTEM: MICRO-TASK
GHOST WORK BEHIND CORPORATE FIREWALLS

Before she had children, Kala, 43, whom you met in the introduction, worked as an electrical engineer. She stepped out of the workforce when her second child came along. But after staying home for several years, she missed the sense of community and purpose she'd enjoyed at work. She broached the subject of going back to work, even part-time, with her husband, but he was skeptical of her ability to juggle family and work. "He worried that it would be too much," she says. But Kala pressed. Finally, she and her husband landed on a compromise: she'd work from home. Now she works on Microsoft's proprietary platform, UHRS.

Microsoft, like several of the large tech companies, has its own internal micro-task ghost work platform, modeled on the mechanics of MTurk. Companies pushing technological innovation need legions of workers to beta-test products and check code. They also rely on people doing ghost work to improve their services' algorithms and artificial intelligence by cleaning up training data from large stores of proprietary data. Tech companies collect and archive information about how people use their sites. Data such as top search-query terms, popular song choices, and mouse cursor movements can be harvested to fuel product development. If customer data is the new oil, the people doing ghost work operate the rigs.

The biggest difference between MTurk and tech companies' internal platforms, like UHRS, is that MTurk recruits and sells labor as well as the platform work site itself, while, on big tech companies' platforms, a third party — a vendor management system (VMS) — recruits and supplies ghost work labor. All of this is to say that vendor management systems create yet another layer of opaqueness, acting as a broker finding people willing to do ghost work on contract, under nondisclosure agreements.²¹ For example, Google used vendor management systems to populate its enigmatic ghost work platform, EWOQ. People hired through contract staffing companies like Leapforce worked on EWOQ, identifying and ranking new webpages to fine-tune links between ads and users'

search queries.²² According to press accounts and people we interviewed and surveyed who work on multiple on-demand platforms, Twitter and Facebook use internal tool kits that function much like MTurk and plug VMS-provided workers into their platforms to monitor and review content.²³ Kala does similar tasks on UHRS for Microsoft.

Before Kala started doing ghost work on UHRS, she worked for a small company that processed back office files from a U.S. business. Her company, one of many so-called business process outsourcing (BPO) shops handling work from the United States, was located in the heart of Bangalore's Electronic City neighborhood, not far from tourist-choked Cubbon Park.²⁴ Perhaps it is ironic that the company's biggest contract was with one of the oldest and largest labor organizations based in the United States. Kala and three other women sat shoulder to shoulder in the BPO shop's cramped four-desk office and did small tasks like deleting duplicates and fixing typos in data entries and updating the labor organization's contact database. She enjoyed going on a virtual scavenger hunt for information, like the correct spellings and zip codes of cities like Chattanooga and Hoboken that are scattered across the web. She liked tracking vital clues, entering the correct search terms, and finding information pertinent to the task at hand.

That's how, three years ago, Kala found her way to Microsoft's internal platform, UHRS. She started by searching Glassdoor, a job review site, where she found a link to a vendor management system. "I clicked that link. It seemed like it offered good ways to keep up my web search and other computer skills."

Kala qualified to work on UHRS with ease. Each vendor management system has its own procedures for vetting potential hires for ghost work, which often come down to language proficiency (the ability to write in English is often privileged) and prowess in finding materials online. On UHRS, for example, if an applicant passes a short quiz testing their language and web search skills, within a matter of minutes they're assigned a unique account. (For UHRS, a worker account is no different than a Microsoft account used for logging on to Xbox.) Once the account is ac-

tive, workers can complete the platform-specific training, learning the ins and outs of UHRS's equivalent to the MTurk dashboard, and can start looking for projects, called "HitApps," on UHRS.

Which HitApps are available depends on the worker's current IP address and the country and language on file with their VMS. Once a worker on UHRS has completed at least one job in the HitApp (called a "Judgment") and shown they can do it correctly, the remaining HitApps are transferred from the Marketplace area into the worker's My HitApps area, the home screen where they can see micro-tasks available to them. Kala set up a desk in the corner of her bedroom and got to work.

The only people able to submit a request for ghost work on UHRS are Microsoft's full-time employees and authorized partners working with Microsoft to develop new products for the company. That means that Microsoft's more than 120,000 full-time employees can, at any time, become requesters in the loop commissioning workers in the ghost economy to help them with their tasks.

In the same way that Joan shared some qualities with her peers on MTurk, Kala has some similarities and some differences with her peers on UHRS. On UHRS, nearly 80 percent of workers are between the ages of 18 and 37, and more than 70 percent are male. But, like Kala, more than 85 percent of workers have a bachelor's degree or higher.

Much like the case of MTurk, there are no clear labor laws governing who can sign up and do paid work on UHRS. But ghost work for any multinational, like Microsoft, is by necessity a global enterprise. The availability of micro-tasks to workers on UHRS around the world largely revolves around Microsoft's immediate needs to support a range of products and services delivered to more than 20 countries in 70 languages.

The types of micro-tasks available to workers on UHRS are not surprising if you think about the products that Microsoft sells. Workers review voice recordings, rating the sound quality of the recorded clip. They check written text to ensure it's not peppered with adult content. Another popular task is translation. Microsoft's strength in speech rec-

ognition and machine translation comes from the ghost work of people training algorithms with accurate data sets. They create them by listening to short audio recordings of one sentence in one language, typically English, and entering the translation of the sentence in their mother tongue in an Excel file.

Other common types of work on UHRS are market surveys — often restricted by demographics like age, gender, and location — and a task called "sentiment analysis." In sentiment analysis, workers may look at a series of words, selfies, videos, or audio files and add a word to each data point that describes their sense of the mood of the word, person, action, or sound in front of them. These human insights become the training data for algorithms later shown the same materials.

Back at home, Kala often turns to her sons for help completing categorization tasks, especially ones that require knowledge of American colloquialisms. The boys help her categorize and sort the best terms for finding common websites. (For example, if a person wants to find an expensive wedding gift, would they enter "fine china" or "fancy dinnerware"?) Kala's children also help when she picks up tasks identifying "adult content," a common job that information studies scholar Sarah T. Roberts refers to as "commercial content moderation."²⁵

This kind of content moderation requires someone like Kala in the loop precisely because words, as plain as they may seem, can mean many different things depending on who is reading and writing them. Artificial intelligence can learn and model some human deliberation like that between Kala and her sons, but it must be constantly updated to account for new slang or unexpected word use.

One of the points of friction when dealing with a vendor management system is that they are limited in the amount of technical support they can offer when someone has a question or something goes wrong. The most a worker can do when a problem or question about tasks comes up is click on the "Report a technical issue with this HitApp" link in their MyHitApps area, enter their issue in a text box that would be familiar to

anyone who's typed in a service complaint online, then wait for someone to reply to their email with help.

Kala echoed what other UHRS workers talked about in the online forums: the MyHitApps area might be down at any time, with no explanation posted, locking workers out of their tasks. Because the VMS doesn't operate the platform itself, it can neither fix the glitch nor explain what's going on. And when this happens, UHRS engineers are always too busy getting the labor platform back up to respond to queries from people doing ghost work. Even though the engineers are sympathetic to the workers' plight, it's contractually not their job to respond to the worker side of the equation. Dealing with worker frustration becomes a giant game of "not it" between the VMS and the tech company using the VMS's ghost workforce.

Likewise, if there's a dispute about the quality of a worker's submitted tasks, VMS agencies don't intervene. Written into most ghost work contracts is that the workers bear the full responsibility of arbitrating disputes. Vendor management systems do typically provide forums to workers, though they are siloed by nondisclosure agreements. Workers signed up for tasks on Google through EWOQ cannot see or talk with workers picking up tasks for Microsoft through UHRS. But if workers have questions or problems specific to UHRS, they can turn to an active group of people doing ghost work talking about the jobs on UHRS.

Finally, new workers on UHRS must wait three weeks before their first paychecks arrive. Tech companies typically pay the VMS, which, in turn, pays the workers. Like MTurk, UHRS uses the first few weeks of a worker's presence on the site verifying their account information, checking their work, and setting up the money transfers that will go to the worker's VMS user account. After the initial waiting period, workers receive payment every two weeks.

Kala's in-laws disapprove of how much time she spends on the family's computer. They'd prefer if she spent more time with them, she says. But she enjoys the independence and having a little extra money in her pocket. For Kala, most work on UHRS means constantly learning new

things. It also means staying connected to the IT sector and a particular job. Working on UHRS gives Kala confidence that she knows the latest software and strategies for finding information online.

Kala can also use her work on UHRS to fill the years on her résumé that might otherwise suggest that she left the workforce altogether. "It is hard for women my age to return to work or break into a new job. Everyone assumes that you only know how to chase after babies or you won't be able to keep up with what you need to know." And her husband has come to appreciate her work and finds ways to support her. She smiles recounting how he brings her tea and snacks in the evening when they are both home. "I love that he sees me working and does little things for me, like I do for him."

Although none of Kala's old office mates at the BPO followed her into ghost work, she does talk with them about how to improve their search query skills. The women trade tips once a week when Kala travels downtown to meet up with her former colleagues so she can feel, as she puts it, "part of the working world."

MTurk and UHRS illustrate how much micro-tasks still depend on human creativity. But the same approach to breaking down work can also be applied to larger projects. As the next two cases show, figuring out where and when to place humans in the loop opens up its own business opportunities. These companies blur the lines between micro-tasks and macro-tasks. In doing so, they also open up the alarming question: What job can't be turned into ghost work?

LEADGENIUS: BLURRING THE LINES BETWEEN MICRO AND MACRO GHOST WORK

Zaffar, 26, has spent his life in the heart of Hyderabad's one-hundred-square-mile "Old City," a walled neighborhood along the banks of the Musi River. Built five centuries ago, it is one of the largest, oldest Muslim neighborhoods in all of India. His father helped build the freeway that rings the edges of the densely populated HITEC City, an IT industrial

park built in the 1990s to accommodate the flood of outsourced work sent to India's shores. Because the first boom in IT work opportunities typically flowed to the upper castes of Hyderabad's Hindu majority, many of Zaffar's uncles and cousins joined part of a wave of Muslim Indian men emigrating to the United Arab Emirates, on the Persian Gulf, to find higher-paying work as drivers, line cooks, or salespeople in the shops crowding the emirates' beaches.

But Zaffar's father wanted a different life for his two sons. He pushed them to get college degrees and white-collar jobs. Zaffar's brother studied finance and works as a teller at a bank in a wealthy suburb of Hyderabad. Zaffar studied IT, getting the equivalent of a bachelor's in engineering from a local technical college.

Tech company jobs in Hyderabad go to candidates with polished English, especially if they've mastered British or U.S. English intonations. Young Muslim men who have few chances to practice their spoken English skills and who have limited engineering training, like Zaffar, are at a disadvantage. He spent about a year applying to call center positions and tech support positions at several of the large tech companies based in Hyderabad but never made it past a second interview. So when he saw a news article about the on-demand platform LeadGenius, he applied.

LeadGenius is a business-to-business service selling leads to salespeople. Other companies sell sales leads, but the genius in LeadGenius is the creativity and insights of its workers. Think about it this way: a basic web search can return contact information for potential new customers and clients, but artificial intelligence can't determine whether the information is useful. That's where people come in. A person can look at the information about any two businesses — for example, how long they've been open, whether they have other store locations, or if either business owner is in the middle of a lawsuit — to help a salesperson decide which of those businesses is more or less likely to be a good target for sales. By cultivating a workforce that feels seen and valued in ways that are hard to come by in on-demand work, LeadGenius bridged the chasm between micro-tasks (small, repetitive, and somewhat mindless

work) and macro-tasks (work that requires thoughtful sleuthing).²⁶ Zaffar found that signing up with LeadGenius worked much like it did on other ghost work sites. As a newcomer, he would have clicked on a button that said "Apply" and created a "candidate account." He would have taken a typing and proofreading quiz and answered a few demographic questions — age, gender, location. He would have submitted a résumé. LeadGenius makes a point of letting applicants know that informal work experience, like working for one's family, is legitimate — a small gesture of recognition that many of its recruits live in places where work is less than formal.

Where LeadGenius's hiring practices diverge from those of other on-demand platforms is what comes next — a rigorous interview round conducted by the workers who've advanced up the ladder. The entire process can take up to three weeks. A job invitation is offered only when an applicant passes the interview, as well as some additional tests.

If the job offer is accepted, the new hire takes part in a paid video orientation and post-video quiz. Paying workers for an orientation also sets LeadGenius apart from many of the other vendors and open ghost work platforms. All new hires start out on a 90-day trial, but if they make it through the first 90 days, keep up their requirements by logging in and staying connected to teams for at least 20 hours a week, and make it to their shifts on time, they get an automatic bump of 8 percent in their hourly pay.

In exchange for a 20-hour commitment, workers must provide their own computer and internet connection, be able to work with office software like Microsoft Word, Excel, and Google Docs, and be comfortable using instant-message and voice-chat software, like Skype. Zaffar works on a laptop he bought himself. He likes to move to different parts of his house, to break up his work shift, rather than sit anchored to the desktop computer that he set up in the foyer that serves as his home office.

Like Zaffar, 85 percent of workers on LeadGenius are between the ages of 18 and 37.²⁷ Slightly more than 70 percent of LeadGenius's ghost workforce — called researchers — have at least a bachelor's degree. Globally, women make up 49 percent of the platform's workforce, although

among our surveyed India workers there were 10 percent more men than women. Almost 75 percent of workers on the platform use LeadGenius and at least one other platform to do on-demand work. When Zaffar first applied to LeadGenius, he was still working full-time on MTurk but hadn't made his daily goal of \$20 in more than a month. According to LeadGenius, one out of every three researchers supports a household of three or more people. And we found that more than 60 percent of workers on LeadGenius rely on the platform, in addition to at least one other income stream, to meet their basic needs.

LeadGenius has a global workforce. The company's largest research teams are based in India and the Philippines. Wages depend on the business markets for the sales leads that LeadGenius is hired to find. As with UHRS, then, workers' pay depends on a mix of two things: first, what's called labor arbitrage — how cheaply a business can get work of comparable quality out of workers connected to global trade around the world but living in a country where the wages are lower — and, just as important, the value of “localization” to a business's product or service. As more and more companies try to sell what they make to a global market, they create a demand for workers who know the local language, idioms, and quirks of their corner of the globe.²⁸

LeadGenius structures its research teams like a traditional workforce. Positions with stair-stepped levels of responsibility — trainer, junior manager, and project manager — create a company-wide trajectory for career growth. Workers at higher levels assemble teams to tackle specific client requests, which often looks like waves of web searches to collect, sort, and refine sales leads.

Teams are organized so that all the members are working in the same time zone, typically in the same country. Project managers also reside in the same time zone, so they can answer questions. Workers must be able to devote between 20 and 40 hours a week. Once they join a project team, they are asked to stay with it for a minimum of 30 days, unless an emergency arises. Workers are evaluated by project managers. There's room for a learning curve, within reason. They receive a strike if they return

work that's incomplete or done incorrectly. After three strikes in three months, workers are removed from the platform.

Workers have a dashboard for most projects, where the lead requests are preloaded, just as they would be for MTurk or UHRS. In some cases, LeadGenius's business clients have workers come to their internal web-sites for privacy and security of the data. The dashboards are also the hub for all communication among workers, managers, and those clients connecting directly with the projects. If questions pop up, workers can turn to a manager for help, and workers have live-chat software running in the background, enabling them to talk to one another, just as they might if they were working in a retail store. Zaffar spends hour after hour making tough, quick decisions about which leads to send to his team.

Since it's a business-to-business company, the daily work can be hard to picture, so here is one example. Law offices in the United States pay platforms like LeadGenius to collect the names of people found in public records. Workers like Zaffar take a city, like Cambridge, Massachusetts, then go online and search local newspapers for postings of people who broke the law. Maybe they were arrested for driving under the influence or maybe they defaulted on their alimony. Those are the kinds of sordid details that may show up when a potential employer looks up a job applicant online. Workers for LeadGenius compile an exhaustive list of people who've been publicly listed for allegedly breaking the law, and then they hand over this list to a lawyer, who will begin calling the people on it, offering to expunge their record from search engine results in exchange for a fee. As one worker told us, this is the kind of lead generation that is constantly providing new work, because “people will always commit crimes that they want to cover up later.” You might say this is a new form of ambulance chasing that comes with the internet. It requires smart searching techniques that a computer can't do on its own.

A worker's day ends when the time they've listed for their availability runs down — what LeadGenius labels as a shift. And, much like we found among other workers, LeadGenius's team members end their shifts and meet up with fellow workers who live in their local communities.

LeadGenius pays workers every other Tuesday, by noon California time (that's where the startup's headquarters are based), using digital payment services PayPal, Payoneer, and even Bitcoin. But, as with MTurk and UHRS, getting paid requires direct deposit, so workers must take a leap of faith and connect their bank account to accept a global transfer from a company headquartered on the other side of the planet, operated by people they may never meet or speak to in person.

Working for LeadGenius, Zaffar saved enough money for his marriage ceremony and to take almost a month off from work. He also was able to step out for more than three weeks to nurse his mother back to good health after a debilitating auto rickshaw accident, and, just as quickly, he returned to LeadGenius. They had guaranteed him a spot if he wanted to come back. Zaffar turned down a junior manager position with LeadGenius right before he got married, because his fiancée did not want them to start out married life with Zaffar working the night shifts or 30 hours of work required of the position.

AMARA: TRANSLATING LANGUAGE INTO GHOST WORK

Karen, 37, has a bachelor's in comparative literature. She lives in Portland, Oregon, with her husband, three-year-old son, and ten-month-old daughter. Amara, a web-based interface for linking video to captions, wasn't Karen's first online job. In years prior, she'd taken on work through Lionbridge, a vendor management system, doing search engine evaluation for (she suspects) Google. She also worked for Fancy Hands, an on-demand virtual assistant service. The work involved fielding live webchat and text-based requests from people asking for help with anything from scheduling a flight to purchasing furniture. But both of those jobs, which demanded constant searching for tasks, were short-lived.

After her second child arrived, Karen started looking for work that was more creative. She began writing and editing "how to" articles for a media company (what she called a "content farm") with clients like eHow and Livestrong.com. Karen says she was "fired" from the media company

after several tense email exchanges with supervisors who marked her work as "unsatisfactory" but didn't offer feedback on what she should do differently with the next assignment. One cannot, technically, be "fired" as a freelance worker, but, as Karen says, "no matter its legal status, it felt like getting fired." Then, searching Craigslist for other copyediting jobs, she came across an ad for Amara.

Amara is a translation and video-captioning service that blends the mechanisms of ghost work with automated features that manage the slicing and recombination of translated video content. Amara blurs the boundaries between repetitive micro-tasks and larger macro-tasks that call on a worker to bring creative insights to a project. The work available through Amara also challenges the clear lines that most in society would draw between paid work and volunteerism, or a labor of love.

The idea for Amara dates back to 2006. It was the brainchild of friends Nicholas Reville, Tiffiny Cheng, Holmes Wilson, and Dean Jansen, all then working with the nonprofit Participatory Culture Foundation (PCF).²⁹ The group had a small amount of grant money to build tools that helped people more easily share videos and creative work on the internet with no gatekeepers or advertisers controlling content. At the time, Real Player and Windows Media Player were the only choices for distributing video online. Building on its early software tools, PCF launched Amara in 2011, a web-based platform that helps people collectively add language translations to dialogue and scenes in videos playing on their screens.

In Spring 2011, not long after PCF released Amara online, activists turned to it to translate videos documenting human rights crises, most notably during the Arab Spring and the Fukushima reactor meltdown. This launched Amara into the limelight. Filmmakers and the nonprofit Technology, Entertainment, and Design, the creators behind TED Talks, approached PCF looking for ways to offer "rush captioning" to media creators and TED presenters who want to caption video for a global audience.³⁰ By mid-2013, PCF Executive Director Nicholas Reville and seasoned technology strategist Aleli Alcala co-founded Amara On Demand (AOD) to fill this niche. Amara represents two realities folded into the growth

of ghost work. First, ghost work's simple existence and persistence belies claims that it's ever easy to completely jettison humans from workflows that require creativity. Second, Amara speaks to a nascent desire among some businesses to explicitly acknowledge that people, rather than software, are the more valuable component behind ghost work.

Karen started out on Amara as a volunteer, subtitling YouTube videos, short documentaries, and college lectures for deaf and hard-of-hearing communities. The macro-tasks appealed to her for reasons that we heard from other Amara workers as well. She enjoyed being a stay-at-home mother but was eager to find avenues for adult interaction with her co-workers. When she moved to working with Amara On Demand, the work was already familiar, but now she was getting paid.

Amara's pay rates vary according to the demand for a specific language. Captioning and translating content in languages that are more commonly spoken, particularly those that map to a wealthier country, fetches a higher premium. So, for example, Amara can pay Karen around \$68 per hour of video to caption what she sees and hears into written English. The first time Karen tried to caption a video, it took her an hour to complete a mere minute of material. But she practiced and got faster, and the captions started to flow more quickly, bringing her pay up to the average.

To that end, Amara's workers captioning English content make about a dollar for every minute of video they caption. Now Karen works on about 15 minutes of video at a time, but that earns her a little over \$15 an hour, almost twice as much as she would earn making coffee drinks at her local Starbucks. And Karen doesn't have to commit to a full-time job to get that pay rate. But the best part, according to Karen, is the teamwork.

In her earlier forays into online work, Karen worked alone. Now she collaborates on video-captioning projects with a team. Amara On Demand organizes teams working on a single video into small groups, the size being dependent on the language and size of the translation project. If it is a feature-length film heading to an international film festival, a job Karen recently took on, Amara will assign a few team leaders to supervise the project, send out job requests, and assemble the team.

Once in place, everyone receives an official invitation to work on the team. Then interested team members send back the times they're available to work on the video assignment. Once a team member accepts the invitation, they can choose any of the videos available and share clips or notes with other people on the team. Amara's on-demand teams work mostly as equals, both producing original subtitles and editing the translations that other team members produce.

Nearly 75 percent of the people working on Amara are between the ages of 18 and 37. More than 60 percent of Amara's on-demand team are women (the reverse of all the other platforms that we studied).³¹ More than 78 percent of Amara members hold a bachelor's degree or higher (and more than 40 percent working on Amara hold a master's degree or higher). Eighty percent of Amara workers rely on the platform, alongside at least one other source of income, to meet their needs. And for nearly 70 percent of workers, Amara is the only ghost work platform that they use to access on-demand work.

Amara still operates most of its team communication via email or a live-chat channel. Team members can create profiles so that their name, picture, and bio can be shared with those with whom they're working on a project. The overall tone among team members is amicable; some people even go so far as to email friendly notes to each other, she says. All and all, it's a very different vibe from her earlier experience working on demand.

But there is another big difference between working for Amara On Demand and working for other platforms. Amara makes it easy for team members to return a task (a video) if they start it and then realize it's either too demanding or not interesting to them. Karen recalls rejecting a task only once after accepting it. It was a video recording of a Samuel Beckett play. "Amara needed subtitles for it. The work was so daunting! The characters talked really fast, and the dialogue was, literally, 'Blah, blah, blah, blah, blah.' Somebody had to caption it, but I mean, I can't imagine." Karen was paid for the amount of time she spent trying to figure out how to do the subtitles: exactly five minutes of reviewing the

video. The structure and ease of picking up as well as dropping projects encourages team members to experiment with topics that might otherwise feel too intimidating.

Amara took pains to build software that would not only be easy for workers to use, but fun. After all, it needed to first appeal to volunteers. Amara On Demand team members are supplied with all the software they need to do the captioning and translation of videos. Jansen and Wilson used the popular dance video game *Dance Dance Revolution* as a model for adding captions to videos and adding features like big, colorful buttons and a simple interface. The software makes it easy to use the tab key to browse the video materials. As a team member watches a video clip, they type translations of the dialogue or descriptions of the action into a text window on the screen. They can then click and drag their subtitles into the Amara Editor, attaching their captions to that video segment. They can also stop and start a video as often as they need. "You're just basically hitting two keys to either start or stop your subtitle to sync it up with the video. It's a great program." You could say that it's as easy as playing a video game.

Automatically recognizing and translating language looks easy in some ways because people are accustomed to the everyday nature of tools like Siri, Cortana, and Alexa. Automating human speech recognition and translation is a fundamental part of artificial intelligence that grew into a field called natural language processing. Natural language processing was helped immensely by the internet's capacity to amass tons of examples of people writing and speaking in various languages. Yet capturing dialogue in video, particularly action scenes that change the mood and meaning of an actor's words, remains a difficult task for a computer program to understand, let alone translate into different languages. In fact, for the most part, because of the lack of intelligence to the computers, it takes a team of people to achieve this, too.

Amara, as a nonprofit, challenges the unbridled venture capital enthusiasm for companies like Uber and the wave of startups in their wake that have pitched themselves as part of the "Uberization" of their market. Together, LeadGenius and Amara represent business models willing

to concede that they don't just sell matching software. They are in the business of banking on people's creativity. They also use ghost work to deliver on tasks much larger than tagging an image. These macro-tasks are, at least today, out of automation's reach.

UPLOADING WORK: WHEN FULL-TIME EMPLOYEES MANAGE MACRO-TASKS

In matching workers to another company's tasks, LeadGenius's and Amara's business practices can't help but draw attention to the other humans in the loop here, who are also hard to see. They are the requesters — small sole proprietors of small businesses or full-time employees at larger companies — managing their own workloads by finding and hiring help on ghost work platforms. They are yet another set of humans entangled in ghost work's loop, working to clear time-sensitive or new projects from their desks.

Commercial platforms like Upwork and a growing crowd of competitors take a hybrid approach to managing ghost work. They allow API access, which enables automated hiring, evaluation, and payment, not unlike MTurk. But they also allow individuals or companies to place tasks on the site manually and interact more with those contracting workers to complete a chore that looks more like something that might happen a few cubicles down than highly atomized micro-tasks. Common tasks placed on such sites include graphic design, video production, and content creation but also more engineering-focused tasks like website production and software engineering, some paying \$100 per hour or more. Larger, more complex tasks require more interaction, so Upwork allows requesters to chat in real time with and send email to workers, which is useful for more complex tasks. Nonetheless, the platform still brokers the interaction between the worker and the requester. In doing so, it also distances the two parties. This can dehumanize the worker in the requester's eyes, as though the worker is simply part of the platform's software. Platforms like Upwork show that automated and semi-automated

processes now distribute and manage the full range of work, from simple to complex.

Perhaps it's ironic that we met people hiring workers on sites like Upwork who turned out to be full-time employees themselves. They turned to ghost work platforms to subcontract out a variety of macro-tasks driven, in most cases, by four reasons familiar to anyone who's ever felt understaffed or overwhelmed at work. First, just as they might post a temp job ad on Craigslist or Monster.com, full-time employees used on-demand platforms because their own firms did not have the expertise that they needed for a project in-house. This expertise ran the gamut, from writing copy to analyzing vibrations in engines. As one senior manager at an engineering firm explained, he typically hires on-demand workers to do detailed engineering designs. "I was looking for a vibration specialist on a large-size induction motor. This is where I was looking for someone specializing in designing as well as engineering." A marketing manager at an online education company who uses on-demand workers primarily for content design, writing, and animation said, "If we can't make it work in-house, we look for a freelancer to complete the job. But there are also instances where we just don't have the skills, so we have no choice but to hire a freelancer."

Second, hiring on-demand workers can be done much faster and with lower costs and lower overhead than hiring via a traditional staffing agency, which is the most attractive feature for companies trying to maximize their profit and increase their bottom line. A communications specialist at a marketing company specifically called out and estimated the savings in hiring on-demand workers: "The most significant aspect of working with a freelancer is that they deliver value at a minimal cost. We can save up to 40 percent by not paying benefits or allocating office space." The marketing manager could also estimate the savings right off the top of their head: "If you go to [a staffing] agency regarding a project, they will charge, say, \$2,500, but if you go to Upwork, you can finish that project for, say, \$700 to \$800."

Third, one of the most common reasons cited for using on-demand work is an unexpected spike in workload. In these scenarios, full-time employees of the company may be busy with other work when a new task arrives, so the firm hires an on-demand worker as an extra pair of hands. The marketing manager said, "We hire freelancers when we have a quick turnaround project. If our in-house team isn't available or we're getting slammed with a lot of work at once, we bring freelancers on board." Full-time employees hiring out ghost work reported a variety of reasons behind the increased workload, including seasonality or a rapidly approaching deadline, and in some cases both. A project manager at a direct mail company said, "We're a cyclical business, and even seasonal, so during peak times everyone is working at full capacity. That's when we raise our hand and say we need to hire a freelancer."

The final reason our interviewees said they hire on-demand workers is that they produce higher-quality work than contractors hired from a staffing agency or, in some cases, even full-time employees. A marketing manager at an advertising company said, "I found in a lot of cases that their performance was above an in-house employee." There are a few reasons for the higher-quality work. First, on-demand workers work for themselves and want to be called back for repeat work. We will see that a common technique among hiring managers for overcoming some of the shortcomings of hiring on-demand work is to maintain a trusted pool of on-demand workers they can repeatedly draw from. On-demand workers realize this, so they do high-quality work to make it into these trusted pools. It's as if every task from a new requester is a type of tryout to make it onto their callback list. A second reason for the higher-quality work output is the competition for jobs in the on-demand labor market. Jobs will go to the workers with the most up-to-date skills, whereas full-time employees can let their skills go stale without serious immediate consequences. A project manager at a healthcare company said, "I think some of the finest technical people come from the area of [freelancing]. It takes a special kind of person, who is diverse, who is more into learning more

systems, new processes. I think [freelancing] sharpens your skill set, it keeps you sharp and on your toes . . . Their skill set makes them more marketable to a company.”

Despite the sincere appreciation for workers' sharp skills and work ethic, there was also a tacit recognition among requesters, sometimes voiced, that they were outsourcing tasks that, under other circumstances, they might have done themselves. In plenty of cases, the distinctions among those generating macro-tasks and those completing them seemed categorically arbitrary. If there wasn't much difference between what a full-time employee did and what they might hand off to a worker, why did the full-time employee seem to have all of the perks and the risks of carrying out that assignment?

Weaponized Ignorance

As ghost work swiftly and stealthily displaces a full-time workforce, it is upending a century of efforts to keep full-time, long-term employment as the cultural default and stabilizing cornerstone of middle-class life.

Humans in the loop appear interchangeable. As noted earlier, thanks to application programming interfaces (APIs), workers are represented as a string of letters and numbers instead of a name and a face. In this dehumanized zone, few companies that sell ghost work have any idea who makes up their workforce ranks. Some of this erasure can be chalked up to logistics. The crowd is too big to see individual faces, one might say. But it's important not to gloss over the fact that erasure is a purposeful feature rather than an errant bug of the ghost work economy.

There are legal reasons why an on-demand platform might not want to know or care too much about its workers. But as the range of ghost work above suggests, it is hard to ignore how necessary humans have become in the growing shadow of AI.

CrowdFlower, a crowdsourcing and data-mining business founded in

2007 by Lukas Biewald and Chris Van Pelt, is the company behind Uber's "selfie security," discussed in the book's introduction. The company has other large business clients, like eBay, Mozilla, Twitter, and Facebook. People doing ghost work on CrowdFlower complete micro-tasks like approving photos, customer support, and content moderation. In 2012, a CrowdFlower worker, Christopher Otey, filed suit against the company, calling out its labor practices. Otey and a second worker, Mary Greth, were named as the suit's plaintiffs, but before it was over, an estimated 19,992 CrowdFlower workers had signed on to the lawsuit. The basis of Otey's lawsuit was that, according to Otey, CrowdFlower set expectations accordant with a full-time employee, but the pay and benefits were commensurate with those of an independent contractor. Otey says of his time working for CrowdFlower, "I didn't have control over the work I did. It was all done on their platform. I couldn't choose my own hours. I had to work when they provided the work. They pretty much controlled all the aspects of the work that was being offered." Given the degree to which CrowdFlower set the terms of employment, argued Otey, CrowdFlower owed him and other workers minimum wage, per the Fair Labor Standards Act. The company's legal team countered that, because CrowdFlower's workers were "free contractors," the FLSA didn't apply. Ultimately, in 2015, CrowdFlower paid \$585,507 to settle the lawsuit, which left the question of the employment status of its workers unanswered.

Since 2015, companies buying and selling on-demand work have tiptoed cautiously around any activity that makes them look as if they are doing anything other than providing an online meeting place and matching service between people with jobs to be filled and workers willing and able to work. On-demand ghost work platforms see themselves as neutral parties, arguing that they are the software serving as the middlemen managing what economists call a two-sided market.

They connect requesters seeking workers, on one side of the platform's marketplace, and workers seeking jobs, on the other side. And in

the absence of set hours, work sites, or agreement about who's in charge, it is difficult to gauge how much ghost work across this burgeoning industry, who's paying for it, and which are completing the tasks. The transaction costs that economist laureate Ronald Coase so long ago identified as the reason very existence of firms seemed to melt away with the new on-line systems. Platforms could keep themselves and requesters at arm's length from workers, shielded from a formal employer's legal responsibilities.

WHEN YOUR WORK HAS NO CATEGORY

The paradox of automation's last mile suggests that the shift to ghost work to deliver services is just heating up.³² As of today, hundreds of companies offering on-demand ghost work services are sorting, annotating, and refining the terabytes of "big data" that companies produce every moment they spend online, and an explosion of cloud services hosting larger tasks that are, at least in part, managed by AI. Treating ghost work as a consumable good drains ghost work job protections.

It can be difficult for those involved to fully see and value the wide range of ghost work. Complicating the issue is that workers themselves don't know how to categorize their work or their status, making it harder to figure out what people performing ghost work might want or need. Without a shared workplace, hours, or professional identity to orient them, those doing ghost work form casual, fluid communities and social circles made up of diverse interests. A common feature of online environments where individuals drift out of their networks, drawn to different people and projects, is the lack of time they have in their day. Organizing formal employment to equally prioritize and support each worker's ability to choose their work, who they work with, and what projects they take on is unprecedented. Usually that privilege is reserved for the most

Everyone else must cram their lives into the constraints of the grind or step aside for someone who will. Using a seemingly expendable, fungible contingent labor force to do out miscellaneous tasks outside the scope or purview of the employee is not new. Ghost work is arguably just the most recent iteration of a well-established historical trend.