

Quantifying the Invisible Labor in Crowd Work

CARLOS TOXTLI, Northeastern University, United States

SIDDHARTH SURI, Microsoft Research, United States

SAIPH SAVAGE, Northeastern University

Crowdsourcing markets provide workers with a centralized place to find paid work. What may not be obvious at first glance is that, in addition to the work they do for pay, crowd workers also have to shoulder a variety of unpaid invisible labor in these markets, which ultimately reduces workers' hourly wages. Invisible labor includes finding good tasks, messaging requesters, or managing payments. However, we currently know little about how much time crowd workers actually spend on invisible labor or how much it costs them economically. To ensure a fair and equitable future for crowd work, we need to be certain that workers are being paid fairly for *all* of the work they do. In this paper, we conduct a field study to quantify the invisible labor in crowd work. We build a plugin to record the amount of time that 100 workers on Amazon Mechanical Turk dedicate to invisible labor while completing 40,903 tasks. If we ignore the time workers spent on invisible labor, workers' median hourly wage was \$3.76. But, we estimated that crowd workers in our study spent 33% of their time daily on invisible labor, dropping their median hourly wage to \$2.83. We found that the invisible labor differentially impacts workers depending on their skill level and workers' demographics. The invisible labor category that took the most time and that was also the most common revolved around workers having to manage their payments. The second most time-consuming invisible labor category involved hyper-vigilance, where workers vigilantly watched over requesters' profiles for newly posted work or vigilantly searched for labor. We hope that through our paper, the invisible labor in crowdsourcing becomes more visible, and our results help to reveal the larger implications of the continuing invisibility of labor in crowdsourcing.

CCS Concepts: • Human-centered computing → Collaborative and social computing systems and tools.

Additional Key Words and Phrases: Crowdsourcing, Amazon Mechanical Turk, Worker training, Peer review, Peer advice, Future of work

ACM Reference Format:

Carlos Toxtli, Siddharth Suri, and Saiph Savage. 2021. Quantifying the Invisible Labor in Crowd Work. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 37 (November 2021), 26 pages. <https://doi.org/10.1145/3274306>

1 INTRODUCTION

Crowdsourcing markets, and their APIs, can help absorb some of the costs associated with crowd work [51, 82]. From the requesters' perspective, these platforms provide an always-available pool of workers and an easy-to-use payment API to contract workers and start getting work done [5, 63, 75]. From the workers' perspective, these markets provide a central place to find work and offer them the flexibility of working from wherever they desire [1, 87].

Authors' addresses: Carlos Toxtli, Northeastern University, Boston, United States, carlos.toxtli@mail.wvu.edu; Siddharth Suri, Microsoft Research, Redmond, United States, suri@microsoft.com; Saiph Savage, Northeastern University & Universidad Nacional Autonoma de Mexico (UNAM), s.savage@northeastern.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2021/11-ART37

<https://doi.org/10.1145/3274306>

However, recent research has identified that some of these costs do not actually get absorbed by the crowdsourcing platform, but rather, they are passed onto the workers in the form of invisible labor [15, 51]. Invisible labor is defined as “unpaid activities that occur within the context of paid employment that workers perform in response to requirements (either implicit or explicit) from employers and that are crucial for workers to generate income, to obtain or retain their jobs, and to further their careers, yet are often overlooked, ignored, and/or devalued by employers, consumers, workers, and ultimately the legal system itself [15].”

Invisible labor is also present in crowd work and it includes activities such as: the unpaid time workers have to invest in finding work, figuring out on their own how to complete the job at hand, or managing their payments [51, 67]. The problem is that crowd workers are forced to engage in these unpaid activities just to be able to complete the labor for which they are paid [91, 107]. If we are aiming to create a future where crowd work is fair and equitable to workers, we need to ensure that workers receive a fair wage for *all* of the labor they do, whether it is the actual tasks for which they get paid, or the unpaid invisible work they do above and beyond that work.

The central question this work addresses is how much time do workers actually spend on invisible work, and how does this affect their overall hourly wages? This is an important question not only to ensure that workers receive a fair wage now but also to ensure that workers receive a fair wage in the future. Notice that our research is addressing a critical problem because a common use case for crowd work is to train machine learning algorithms, or to provide a human-in-the-loop approach when A.I. fails [13, 44, 100, 106]. Since we are in the midst of an “A.I. revolution,” it is plausible that we will see dramatic growth in the use of crowd labor [10, 45, 64, 65, 99]. In addition, post-COVID-19, there will likely be a large increase in people who need to work from home, whether that is for safety reasons or because of the massive number of worldwide layoffs [29, 31, 104]. Measuring invisible labor in crowd work will only grow in importance going forward.

To start to quantify the invisible labor in crowd work, we develop a web plugin¹ that allows us to detect when a worker is performing invisible labor and quantitatively measure the amount of time the worker spends on such efforts on Amazon Mechanical Turk (MTurk), one of the most popular crowdsourcing markets [73]. We conduct a field study with our plugin to measure in the wild the amount of time crowd workers invest in invisible labor, and contrast with the amount of time workers spent on paid labor. Armed with our plugin, we had 100 crowd workers use our plugin for one week. Workers completed 40,903 human intelligence tasks (HITs). Through our plugin, we uncovered that crowd workers spent 33% of their time on MTurk doing unpaid work. Relatively similar to prior work, we found that workers’ median hourly wage considering only paid labor was \$3.76 [57]. But, if we consider the time workers spent on invisible labor, we calculated that workers’ median hourly wage dropped to \$2.83. We also found that the amount of time that workers dedicated to invisible labor varied across workers’ skill level and demographics. We found that master workers spent 23% less time on invisible work than regular workers. We also observed that the time spent in invisible work appears to be heavily correlated with demographic factors.

The invisible labor in which crowd workers spent the greatest portion of their time revolved around payments. In particular, this most time consuming activity involved doing tasks for which workers were not paid because they experienced a “time out” (and hence they did not receive any payment for any of the labor they conducted for the task). Workers spent a median of 4.5 minutes daily on this activity. Overall, invisible labor around “payments” was the most time-consuming for workers; it was also among the most common. In fact, 97% of the workers in our study practiced invisible labor around visiting the earnings section on their workers’ dashboard (perhaps to ensure they had gotten paid fairly [112]). The second most time-consuming category of invisible labor

¹<https://github.com/anonym-research/invisible-labor>

involved hyper-vigilance where workers were “on-call” vigilantly watching over requesters’ profiles ready to do, at all hours of the day, the labor that certain requesters posted, as well as vigilantly searching for work on Amazon Mechanical Turk [24, 51, 114]. Understanding invisible labor is key to creating positive change in crowd work [15]; however, it has remained so far understudied. Bettering our understanding of invisible labor will allow us to design fairer crowdsourcing markets.

2 RELATED WORK

Our research builds on two main pieces of literature: (1) research on invisible labor, especially within digital labor markets [15]; and (2) methodologies to quantitatively measure the time crowd workers spend completing paid labor on MTurk [96, 98], i.e., HITs.

2.1 Invisible Work

The literature has traditionally characterized invisible labor as work that is “economically devalued through cultural, legal, and/or spatial dynamics” [59]. Under this definition, invisible labor is usually done in private rather than public [18, 50]. Usually, housework is one of the most commonly cited examples of invisible labor [14, 30], and it involves both physical labor and mental labor (e.g., planning what types of household chores should be done and in what manner.)

In their book *Invisible Labor* [15], Crain et al. describe the concept of invisible work as the “*activities that occur within the context of paid employment that workers perform in response to requirements (either implicit or explicit) from employers.*” They explain how this concept has existed in different offline settings before, but nowadays, technology has enabled a large part of all invisible labor. In particular, a number of technology companies are passing several aspects of digitization labor to consumers and workers, e.g., consumers are expected to install all the required Internet infrastructure at their homes. This labor is typically presented as something that is mundane, flexible, and part of the “do-it-yourself” culture [48]. However, this dynamic also reduces the meaning of this type of labor, making it invisible, and something for which people are not paid.

2.2 Invisible Labor in Crowd Work

Gray and Suri [51] explored these concepts of invisible labor with a particular angle toward crowd workers. Through in-depth interviews with crowd workers, their book “*Ghost Work*” unveils the current conditions to which crowd workers are exposed and explains how companies have placed on the shoulders of workers a great portion of the invisible labor that companies themselves would traditionally do. The book also went a step further and started to describe the different types of invisible labor present in crowd work. Within this setting of describing invisible labor in crowd work, it is important to consider that crowd work does not emerge only from the requesters’ side of the market; it is also something that crowdsourcing platforms, such as Amazon Mechanical Turk facilitate [51] and could, with different design choices, help alleviate. For instance, crowdsourcing platforms could help match workers to tasks to reduce search time. Also, they could potentially pay workers for the time they spend searching for tasks or reading messages from requesters, which is something that companies have traditionally covered [20, 51].

Furthermore, when thinking about invisible labor in crowd work, we have to consider that much of the labor that crowd workers complete is fed into machine learning models that power the A.I. industry [72, 106]. For instance, crowd workers might label content so that Facebook’s News Feed algorithm will not recommend posts that are filled with hate speech or pedophilia [17, 46]. Crowd workers might also transcribe audio to help Amazon’s Alexa better understand the user [6]. Given that most end-users are unaware that there are humans helping to power the A.I. services they access [51], the work done by workers and their possible unfair labor conditions, are hidden from

sight. Notice that here the invisibility of crowd work is again not just due to requesters and their HIT design choices, but rather it is an issue within the A.I. industry as a whole. In this particular research, we focus on measuring the different categories of invisible work that the book of Ghost Work identified that exist within the context of crowd work [51]. We believe that by quantifying the different costs that invisible work has in this setting, we can design better solutions to improve crowd workers' conditions. Notice that invisible labor in crowd work includes activities that go unnoticed while doing paid work, such as finding HITs and communicating with requesters to resolve conflicts [42, 55, 57, 97]. Invisible labor in crowd work has recently gained more attention because it has become clearer that the independent nature of crowd work has led workers to now have to assume invisible labor that was traditionally taken by companies and employers [20, 95]. In this work, we present computational mechanisms for quantifying for the first time the invisible labor that exists in crowd work and bring much needed light to a critical topic.

2.3 Quantifying Working Time

Saito et al. [96] studied different ways to measure the time crowd workers spent completing HITs using their system called TurkScanner. They found that through web plugins, they could quantify how much time workers spent completing HITs. We built upon their methodology to properly measure the working time on HITs and expand their method to now also measure time spent in unpaid work. Hara et al. [57] also used plugins to measure the wages of MTurk workers. Ignoring unpaid work, they estimated an average hourly wage of \$3.18, which roughly coincides with the average hourly estimate of \$3.76 that our study uncovered when we also ignore unpaid work. This shows that our measurement instrument is roughly calibrated to theirs (we likely had higher hourly wages because we considered more conservative measurements). We build on their work and provide a more detailed accounting and classification of the unpaid work that workers endure.

3 METHODS

The goal of our IRB-approved field study is to measure and contrast the time that crowd workers spent on invisible labor and the time they spent on regular paid labor (i.e., completing HITs.) Since this data is not part of the official MTurk API, and prior work has not been able to measure invisible labor at the level of detail in which we were interested, we build computational mechanisms to measure these variables. Armed with these computational mechanisms, we conduct a field study to investigate in the wild how much time workers on MTurk dedicate to invisible labor. In the following, we describe how we measured these activities through the computational mechanisms that we designed and detail how we conducted our field study. It is important to highlight that our computational methods for measuring invisible labor focus on measuring invisible labor in a conservative manner. We consider it is best to err in underestimating the amount of time that workers spend in invisible labor than to overestimate. We make this design decision because quantifying invisible labor can potentially call attention to the structural issues surrounding crowdsourcing markets and the conditions they provide workers. Operating in a conservative manner helps us to avoid being labeled as “exaggerated activists” and allows us to present the study in a scientific, objective way. This approach helps us to bring much-needed attention to understanding invisible labor in crowd work. As we will see, even with erring on the side of underestimating invisible labor, it is still a sizeable overhead for the workers.

3.1 Computational Mechanisms to Measure Invisible & Paid Labor

For our study, we need computational mechanisms for: (1) detecting when a worker is doing invisible labor or when she is doing paid work; and (2) measuring how much time a worker invests in each of these two activities. To address these two points, we created a Chrome Extension (plugin).

3.1.1 Methods for Quantifying Paid Labor. Our plugin builds on prior research that was able to detect and measure with plugins when a crowd worker was completing a HIT, the amount of time the worker invested in completing the HIT, and the daily earnings that workers made from the HITs (notice that this value is important as it can help us to quantify the monetary costs of invisible labor)[96]. In particular, building on prior work, we developed a plugin that can: 1) automatically record the exact times when a worker accepts a HIT and when she finishes and submits the HIT; 2) track when a tab about a HIT is in focus and automatically record the time period in which the worker is active on the HIT page tab by checking whether there were any type of interactions from the worker (e.g., mouse movements, typing) under a given time window; and also 3) measure the daily income that each worker makes from these HITs by querying the information from their workers' dashboard on MTurk. In summary, as a starting point, we developed our own plugin that mimics prior work and quantifies the amount of time that a given worker dedicates to completing HITs and the earnings that the worker is making.

3.1.2 Methods for Quantifying Invisible Labor. Next, we expand the plugin to now provide new functionality through which we can also track and measure the time workers spend on invisible labor. Notice that we focus on quantifying invisible labor in a conservative manner, which means that we prefer to err on the side of under measuring the invisible labor (we took this methodological decision based on the reasoning stated above). Our conservative approach to the measurement of invisible labor comes in especially when we consider cases where there is disagreement in the literature on whether an activity is invisible work or not [47, 94]. In such cases, we prefer not to label the activity as invisible labor. We prefer to underreport so that the invisible labor we measure will be at least as large as we quantify here, if not larger. Some of the discussions around what is and what is not invisible labor especially arise for the activities of "reading instructions of the HITs", and "taking breaks" [25, 36, 43, 47, 51, 76]. Gray and Suri [51] label "reading instructions" and "taking breaks" as examples of invisible labor activities. However, we decided not to categorize these activities as invisible labor because there is research that considers these two activities as part of paid work [25, 36, 43, 47, 76]. Now, given that workers are not actually paid for either of these two activities, we designed our computational methods to detect when workers take breaks or read instructions; but, we do not count these activities as either paid nor invisible labor. It is important to highlight that because workers are not paid any wages for reading instructions, it is incorrect to categorize the work as being paid.

Our plugin, therefore, in addition to what prior work had already developed, provides now the novelty of being able to detect and quantify all other activities that workers do aside from completing HITs. For this purpose, we developed new computational mechanisms to detect when a worker is visiting other parts of the MTurk platform that are different from the HIT page tab²(e.g., perhaps the worker entered the MTurk page to search for HITs³ or the worker entered the MTurk page for sending messages to requesters⁴). Our plugin tracks the exact time when a worker enters one of these other MTurk domain pages and then scrapes and parses the HTML of the page to understand how the worker interacted with the page and identifies the intervals of time in which the worker is active on these other pages. We consider a worker to be active on a page when the worker has the page in focus and does any type of user interaction on that page, e.g., mouse movements, scrolls, clicks, keyboard typing. Notice that we do not track what a worker does on these pages (e.g., we do not track what they type). We simply detect that they are active on a particular MTurk page. To accomplish all of this, we developed two new components into our plugin: a page crawler and a

²<https://worker.mturk.com/>

³<https://worker.mturk.com/?filters>

⁴https://worker.mturk.com/contact_requester

time-driven background process that detects the different browser events that happen on MTurk (e.g., that the worker visited another page on MTurk, or that she started typing, or began a new HIT). The page crawler detects the current MTurk domain page that the worker is on, as well as the status of the page (e.g., that the page is loaded, active, inactive, or closed). The background process focuses on detecting the HITs that the worker is currently doing and identifying which she has finished. In order to accomplish this, the background process polls workers' task queues on MTurk every 30 seconds. From the task queue, the background process obtains the metadata and status of all the HITs the worker has accepted to do. Notice that the page crawler is the primary element that we use to detect whether the worker is completing paid labor or invisible labor. The background process helps our plugin to be able to better detect when the worker is completing HITs (some of them reside outside the MTurk platform) and also when the worker is multi-tasking (doing multiple HITs at the same time.) Through this, we create a plugin that automatically detects when a worker is doing invisible or paid labor and the amount of time the worker invested in each of these two activities. Our plugin is available here: <https://github.com/anonym-research/invisible-labor>.

3.1.3 Quantifying Types of Invisible Labor. We were not only interested in detecting whether or not a worker was doing invisible work; we also wanted to understand what type of invisible labor was the most taxing and contextualize our results with prior interview work that started to document the invisible labor that workers perceived by conducting interview studies with them [51]. In the following, we present the different types of invisible labor we consider (i.e., broad categories) and how we detected their related individual activities. The categories and activities we study are based on prior interview research that studied invisible labor [51, 113]. Note that for most cases, we detect that a worker started a new activity when they loaded, focused, or changed their browser tab to a page on MTurk related to that particular type of invisible labor (below, we mention which pages relate to specific invisible labor activities). Similarly, our plugin considers that a worker paused or finished an activity when the worker changed to another tab, unloaded, blurred, or closed the MTurk page related to that particular activity. The categories and activities we consider are:

(a) Category: Hypervigilance. This category involves workers spending time in: (1) identifying good work, e.g., *“wading and sorting through spam or suspicious offers for at-home-work.”*[22]; and (2) being “on-call,” ready to do HITs for requesters at any time. Invisible activities include:

- *Watching over requesters' profile:* Notice that this activity relates to Hypervigilance because workers are visiting requesters' profiles to be ready to do any HIT that requesters post. In other words, workers are “on-call.” To detect this activity, our page crawler detects when a worker is on a requester's profile page.
- *Searching for general HITs (unfiltered):* To detect this activity, our page crawler identifies that a worker is on the main page where HITs are posted.
- *Searching for filtered HITs:* Our page crawler detects when the search URL for the main page of HITs has a query in it to filter HIT results. This activity relates to hypervigilance as it involves “wading and sorting” through HITs.
- *Managing their queued HITs:* this activity relates to Hypervigilance as it involves workers filtering out fraudulent HITs and focusing on HITs from specific requesters (i.e., being “on call”). To detect this, our crawler identifies when a worker is visiting her tasks queue.
- *Checking their own qualifications:* This activity relates to Hypervigilance as prior work has identified that workers watch over their own qualifications to vigilantly identify whether they could now access certain HITs and thus more effectively find and access quality labor [51]. In this case, our crawler detects when the worker is viewing her earned qualifications.

(b) Category: Lack of Guidance. Crowd workers are generally left on their own to figure out how to do jobs as fast and accurately as possible [86]. Activities related involve: • *Starting HITs but then returning them*: This activity relates to “Lack of Guidance” as it usually occurs because workers believe the HITs will be different than what they actually end up being [49, 88] (e.g., less complex or of another nature.) The lack of guidance leads workers to have to return HITs they already started. In this case, our crawler detects when workers click the return HIT button on MTurk.

- *Sending messages*: Workers send messages to requesters to ask them questions about a HIT and better understand what the requester wants. To detect this type of invisible labor, our crawler detects when a worker opens MTurk’s messaging form to send a message.
- *Reading HIT information*: Page crawler detects when a worker clicks the “More Info” option while previewing or working on a HIT. Notice that this activity is different from reading HIT instructions, as reading HIT information helps workers get a preview of what a HIT is about. It is an activity that workers have to do in order to obtain guidance.
- *Previewing HITs*: Page crawler detects when the page of a HIT is open in preview mode. Notice that here we could potentially say that workers are previewing HITs in order to “vigilantly” find tasks from certain requesters (and hence this activity could be labeled as being from the category of Hypervigilance). However, the search filtering option allows workers to do that more effectively, and that is also not the main purpose of the preview [70]. We, therefore, decided to categorize this activity as Lack of Guidance. Additionally, prior work has labeled this activity as related to guidance [83, 116].
- *Reading platform help*: Page crawler detects when workers are in MTurk support sections. Notice that within this category, we could have considered the activity of reading instructions as part of the invisible labor that a worker has to do related to the lack of guidance. However, as mentioned before, we opted to just detect the activity but not label it as invisible nor paid labor. To detect the activity of “reading instructions,” the page crawler detects the time that passes from when a worker accepted a HIT until the worker has her first interaction with the HIT (e.g., she presses a key, or she opens another tab related to the HIT, etc.) We assume that this time-lapse corresponds to when the worker is reading instructions.

(c) Category: Payments. In crowd work, even after workers have vigilantly identified legitimate labor and they have been able to figure out how to complete the work, they still run the risk that they will not get paid for their efforts. The broad category of “Payments” encompasses the invisible labor that workers do to ensure payment and also instances where they worked on HITs but were not paid. This category of invisible labor includes:

- *Visiting their worker’s dashboard*: workers visit their dashboard to oversee if requesters have paid them and ensure they made a certain amount of daily income. To detect this activity, our crawler identifies when workers are visiting their general MTurk dashboard.
- *Doing HITs that eventually timeout*: Some HITs have an expiration time on them. If workers take longer to complete the HIT than the allowed expiration time, the HIT times out. In these cases, workers are not paid for any of the labor they have done on the HIT, and thus we consider this activity within the broader category of Payments. To detect these instances, the background process of our plugin identifies when a HIT has an end time equal to or higher to the HIT expiration time. Our plugin also checks in the worker’s dashboard that the worker was never paid for those HITs.
- *Viewing their earnings*: Page crawler detects when workers are in earning sections on MTurk.

(d) Category: General Logistics. The last category we detect relates to MTurk logistics. We focus on the activities of logging into MTurk. Our crawler detects when workers log into MTurk.

3.1.4 Detecting and Processing Multi-Tasking. When measuring the time workers spent in completing HITs, it is important to properly detect when workers are multi-tasking and properly measure and account for the time they spent doing so [74]. In our study, we refer to multi-tasking as when a worker accepts multiple HITs or batches of HITs around the same time and then starts completing these multiple HITs. The background process of our plugin checks the workers' tasks queue to detect workers completing HITs via multi-tasking. To account for this time, we adopt an approach similar to prior work [96]. A common feature of working in this manner is that the HITs are chained in succession. This means that the start and end times may overlap with one or more HITs in the batch. Also, similar to prior work [96], our study does not consider batch HITs that take more than one day to be completed (0.6% of our sample). We filtered out all the multi-day batches and HITs since these imply computing the effective working schedule of each worker.

3.2 Field Study

The purpose of our field study was to have workers use our plugin and through it measure in the wild the amount of time workers dedicate daily to invisible labor. Armed with our plugin that could detect and measure the amount of time workers dedicated to different types of invisible labor, as well as the time they dedicated to paid labor, we conducted a field study that lasted a week. Note that we included weekends in our analysis as MTurk presents itself as a platform that offers workers the flexibility to work whenever workers want (weekends included). Similar to prior work [58, 98], we did not see changes in the days workers completed tasks.

3.2.1 Field Study Logistics. We recruited workers from MTurk by posting a HIT inviting workers to our study. We also used mailing lists of Turkers (workers on MTurk) who had participated previously in studies with us. For our study, we first surveyed participants on their perceptions of how much time they estimated that they spent on invisible labor. We asked workers to report how much time they felt they invested on MTurk: searching for work; looking over their worker dashboard; sending messages to requesters; and doing HITs that eventually timeout. This helped us understand workers' prior beliefs and awareness of invisible labor and how much time they believed they spent on it. We also asked workers about how COVID-19 had affected them (none of our participants expressed any work disruptions). Our initial survey also asked workers about their basic demographic information, such as current location, gender, disabilities, etc. Overall, we based our survey on prior work [37, 58].

After the initial survey, we asked participants to: (1) install and use our plugin for a week; (2) work on MTurk as normal; (3) visit the plugin dashboard, which showed to each worker graphs of how much time the plugin detected that they invested in different MTurk activities for a given day. At the end of the field study, workers completed a short survey evaluating the accuracy of the plugin in detecting and measuring the amount of time they spent on different activities on MTurk. In general, workers in our study stated that they felt that our plugin was able to adequately track the time they spent daily on MTurk completing HITs and doing different invisible labor activities (the median score for the plugin's accuracy was 4 on a 5 point Likert scale). We paid each participant \$10 USD for taking part in our study. Notice that this accounts for the US federal minimum wage (\$7.25/hour) as our initial survey took 5-8 minutes to complete, the installation of our plugin took less than 4 minutes, and the end survey we gave participants took 5-8 minutes.

4 RESULTS

We had 100 MTurk workers install and use our plugin for a week. We allowed all types of workers to participate in our study. This resulted in us recruiting 21 "master workers" and 79 "non-master" workers. Note that we considered that a worker was a master worker if we

detected that they had

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:9

Description of the Statistics	Value
Total number of workers in our study	100
Total number of HITs workers did in a week	40,903
Minimum number of HITs a worker did in a week	1
Median number of HITs a worker did in a week	185
Maximum number of HITs a worker did in a week	3,168
Minimum number of HITs a worker did per day	0
Median number of HITs a worker did per day	30
Maximum number of HITs a worker did per day	1,149
Minimum time a worker invested in completing HITs per day	0 min
Median time a worker invested in completing HITs per day	1:07 hrs
Maximum time a worker invested in completing HITs per day	7:36 hrs
Minimum time a worker invested in invisible labor per day	0 min
Median time a worker invested in invisible labor per day	33 min
Maximum time a worker invested in invisible labor per day	5:31 hrs
Minimum earnings made by a worker in a week	0.92
Median earnings made by a worker in a week	\$55.39
Maximum earnings made by a worker in a week	\$542.06
Minimum earnings made by a worker per day	0.01
Median earnings made by a worker per day	\$8.07
Maximum earnings made by a worker per day	\$178.62
Median hourly wage with invisible labor	\$2.83
Median hourly wage without invisible labor	\$3.76
Percentage of workers who multi-task	96%
Minimum number of batches a worker did in multi-tasking	1
Median number of batches a worker did in multi-tasking	32
Maximum number of batches a worker did in multi-tasking	333
Minimum number of HITs a worker did in a batch	2
Median number of HITs a worker did in a batch	3
Maximum number of HITs a worker did in a batch	689

Table 1. Summary statistics of the workers in our study with regard to: HITs workers did, the time they invested in working, workers' earnings, and their multi-tasking information.

completed at least one HIT with master qualifications. Table 1 presents the statistics of the workers in our study and their general labor patterns. We had 73 men and 27 women, who had a median age of 30 years old. 41 participants were from the United States, 45 from India, five from Brazil, three from Italy, and the remaining six from Venezuela, Spain, Mexico, United Kingdom, United States Virgin Islands, and Thailand.

Through our plugin, we identified that workers did a median of 30 HITs each day. The median daily earnings of each worker were \$8.07 US dollars. Figure 1 presents the median amount of time that each worker invested in completing HITs during our one-week study. Each bar represents a worker, and the bars are sorted along the X-axis based on the median amount of time they worked daily on MTurk. The Y-axis shows the amount of time each worker dedicated to completing HITs or doing invisible labor. The light gray part of each bar shows how much time the worker spent doing HITs, and the dark gray part shows how much time they spend doing invisible labor. Observe that invisible labor occupied a substantial amount of workers' overall time. The median time that



Figure 1. Overview of the labor patterns of each worker in our study.

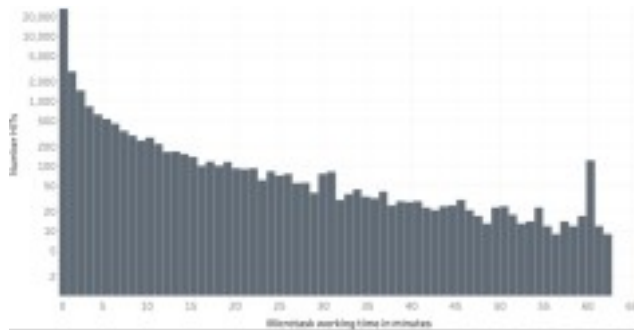


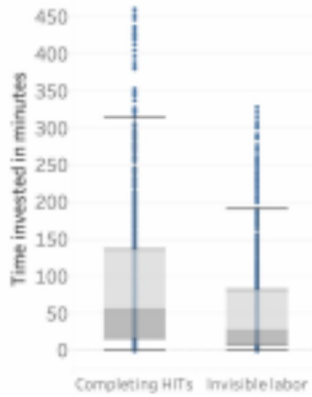
Figure 2. Working time distribution of the HITs that crowd workers in our study completed.

workers invested daily in completing HITs was 1 hour 7 minutes, and the median time that workers invested in invisible labor was an additional 33 minutes, with some workers spending a maximum of 5 hours and 31 minutes daily. Notice that to calculate this value, we summed up all of the time that workers invested in the different invisible labor activities that our plugin detected. Workers spent a median of 33% of their daily time on MTurk doing invisible labor.

We also graphed a histogram of the amount of time that workers dedicated to completing HITs (see Fig. 2). This graph helps to calibrate whether our plugin is measuring paid labor adequately as we can compare our findings to prior work [98]. Note that we used a log scale on the y-axis so that the distribution was easier to visualize. From here, we observe that similar to prior work [96], the distribution of the time that workers invested in completing HITs had a long tail that was heavily weighted towards shorter tasks, meaning workers usually did HITs that took under a minute.

Next, we were interested in studying whether there was a significant correlation between the time workers spent working and the time they spent conducting invisible labor (as this can help us to better understand the phenomena of invisible labor). For this purpose, we computed the Spearman's correlation and obtained 0.283 (p-value 0.004) for the time workers spent working and time doing invisible labor, and 0.517 (p-value 0.000) for the percentage of time working and time in invisible labor. Given these values, for both cases, we reject the null hypothesis that the samples are uncorrelated, i.e., we identified that there is correlation between the time workers' spent working and the time they spent completing invisible labor. Future work could thus study the type of paid labor that might minimize the amount of time a worker has to dedicate to invisible labor.

4.0.1 Quantifying Invisible Labor and its Economic Costs. We aimed here to understand the economic costs that invisible labor has on workers' wages. For this purpose, we first visualized in greater depth the median time workers spent daily in invisible labor and in completing HITs (see Fig. 3). We also aimed to understand the distribution of payments of the completed HITs (see



Quantifying the Invisible Labor in Crowd Work 37:11

Figure 3. Box plots showing the time that workers dedicate to completing HITs and doing invisible labor

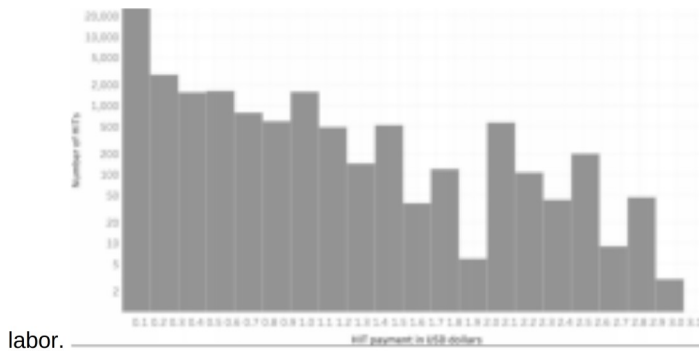


Figure 4. Payment distribution of the HITs that crowd workers in our study completed.

Fig. 4). Armed with this information, we calculated the median hourly wage of workers. We used an approach similar to prior work [57, 98]. We first calculated the total hours a worker spent completing HITs on a given day D . We call this the worker's $\sum_{i \in D} \text{Time}_{i, D}$, and it is the sum of all the time series ($\text{Time}_{i, D}$), measured in hours, that the worker dedicated to doing HITs on day D within the time period D :

$$\sum_{i \in D} \text{Time}_{i, D} = \sum_{i \in D} \text{Time}_{i, D} \quad (1)$$

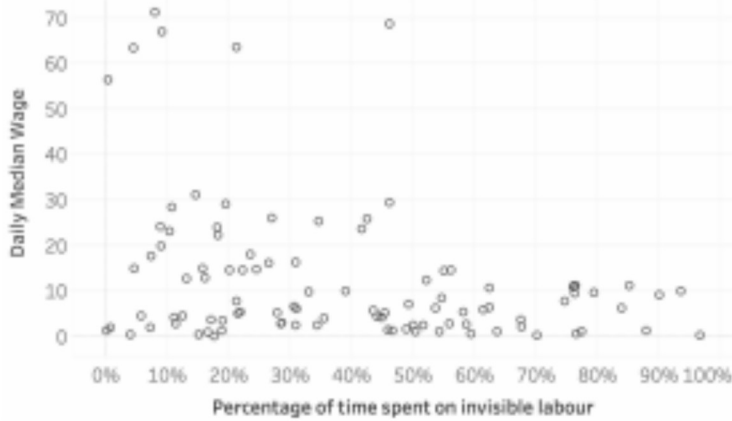
After this, we obtain the total $\sum_{i \in D} \text{Time}_{i, D}$ the worker made on day D . We take this value from the rewards and bonuses logged on the worker's "Daily Income" on

her MTurk dashboard. For worker w_i , her overall hourly wage for day d is:

$$w_i = \frac{\text{Income}_i}{\text{Time}_i} \quad (2)$$

With this, we calculate for each worker her hourly wage for each day of our study. We then use that information to calculate the median hourly wage of the 100 workers participating in our study. Excluding invisible labor, we calculated that workers earned a median hourly wage of \$3.76, which roughly coincides with prior work, which calculated \$3.18 [57]. Notice that it is likely that we calculate a slightly higher salary because we utilize a slightly more conservative approach for our measurements, with the purpose of limiting the overreporting of invisible labor that workers do.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.



37:12 Toxtli et al.

Figure 5. Overview of the percentage of time each worker spent daily on invisible labor (X-axis) and their median daily wage (Y-axis).

Now, if we include invisible work into the calculation of the hourly wage, the median hourly wage of workers drops to \$2.83. Next, we were interested in better understanding the dynamics around invisible labor and wages. Figure 5 presents a scatter plot where each point represents a worker. The X-axis represents the median percentage of time a worker invested in invisible labor daily, and the Y-axis the worker’s median daily wage. From Fig. 5, we observe that the highest-earning workers, in general, all invested less than 50% of their time in invisible labor. Given this result, there might be value in exploring coaching systems that teach workers how to best manage their invisible labor to ensure high wages.

4.0.2 Invisible Work for Different Segments of Workers. In this section, we provide a breakdown of the different demographics of workers in our study (segments) and study the type of invisible labor they presented in their work practices. This analysis is important as research has started to showcase how workers’ different demographics can impact how they approach work on MTurk [37, 58, 90, 102]. We were thus interested in further studying and understanding this aspect, but now for invisible labor. In Table 2, we present an overview of the amount of paid labor and invisible labor that different population segments conducted. Notice that in the table, we also calculate the “unpaid percentage ratio,” which denotes the percentage of the total working time that is unpaid. We calculated the unpaid percentage rate as follows, where $\text{Time}_{\text{invisible}}$ is the median time workers in a particular segment spent on invisible labor, and $\text{Time}_{\text{paid}}$ the median time workers in that segment spent on paid work.

$$\begin{aligned}
 & \text{[Symbolic representation of a mathematical equation with various symbols and subscripts]} \\
 & = \text{[Symbolic representation of a mathematical equation with various symbols and subscripts]} \\
 & + \text{[Symbolic representation of a mathematical equation with various symbols and subscripts]} \\
 & \text{[Symbolic representation of a mathematical equation with various symbols and subscripts]}(3)
 \end{aligned}$$

Armed with these measurements, we next conducted statistical analysis to study whether there were significant differences between how invisible labor impacted the different segments of workers. First, over each worker segment we performed the Shapiro-Wilk test, which allows us to identify whether our distribution is normal or not. We found that for all the segments, the p-value was less than .05, so we rejected the null hypothesis (i.e., our distribution is not normal). Given that we do not have a normal distribution, we proceeded to use a non-parametric analysis of variance. We performed a Kruskal-Wallis H Test as a non-parametric alternative to the parametric one-way between-groups analysis of variance for independent groups. We found that there was a significant

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:13

Workers' segment	Working Time	Invisible Work	Unpaid rate	Payment #	workers	Masters	Workers	1hr	37min	24min	19.8%	\$13.8	21						
Non-Masters	Workers	58min	43min	42.5%	\$5.5	79													
Workers based in United States (English Speaking)	1hr	28min	27min	23.4%	\$11.9	41	Workers based in India (English Speaking)												
42min	35min	45.4%	\$4.0	45	Workers based in Brazil (non-English Speaking)	18min	1hr	15min	80.6%	\$1.9	5	Workers based in Italy (non-English Speaking)	1hr	11min	1hr	34min	56.9%	\$9.5	3
					Women	1hr	02min	42min	40.3%	\$8.2	27								
					Men	53min	28min	34.5%	\$5.5	73									
					18-24 years old	26min	33min	55.9%	\$3.3	9									
					25-34 years old	1hr	01min	45min	42.4%	\$6.3	52								
					35-44 years old	59min	22min	27.1%	\$5.9	22									
					45-54 years old	55min	20min	26.6%	\$6.6	9									
					55-64 years old	1hr	12min	36min	33.3%	\$9.8	6								
					65-74 years old	21min	19min	47.5%	\$1.3	2									
No impairment declared	54min	31min	36.4%	\$5.7	93														
Mobility impairment	1hr	19min	19min	19.3%	\$15.9	5													
Mental disorder	1hr	16min	1hr	44.1%	\$9.2	2													
Frequently multi-task	1hr	04min	28min	30.4%	\$6.7	42													
Rarely multi-task	48min	32min	40.0%	\$5.4	58														
					Use tools	1hr	03min	37min	37.0%	\$6.8	79								
					Not use tools	33min	21min	38.8%	\$3.1	21									

Table 2. Median of times and payments per segment. The unpaid rate shows the percentage of the total working time that is unpaid (invisible work). The payment amount represents the median daily payment.

difference (p-value < 0.05) in the invisible labor time between the workers who were: Masters and Non-masters (p-value 0.00), male and female (p-value 0.04), tool users and non-tool users (p-value 0.01), from English speaking countries and non-English speaking countries (p-value 0.02). We did not find a significant difference among the following groups: workers without disabilities and workers with some disabilities (p-value 0.64); workers who do multi-tasking and workers who do no multi-tasking (p-value 0.32); workers in the U.S. and workers in India (p-value 0.07).

Next, we dug deeper into several of these results to better understand the dynamics

behind invisible labor and workers' demographics. Table 2 shows that from the 100 Turkers who participated in our study, 21 of these were MTurk Masters, and 79 were not. We found that the median amount of time that master workers invested in completing HITs daily was 1 hour and 37 min and the median amount of time they invested in invisible work daily was just 24 minutes, as shown in Table 2. Non-master workers worked slightly less time on HITs and spent more time on invisible labor than master workers. Non-master workers spent a median of 58 minutes daily completing HITs and a median of 43 minutes on invisible labor (almost double the time to what master workers invested.) Thus, workers with the Masters distinction spent more time working and less time doing invisible work than non-masters. Overall, a key takeaway from Table 2 is that Master workers perform 23 percentage points less invisible work than non-Masters workers (20% vs. 43%) and earn a median of \$8.3 more a day. Naturally, 21 Masters is not a huge sample, so one should view this result as suggestive and follow up with future work to confirm. There are also a variety of explanations for this finding. It could be that the experience and know-how of the Masters workers help them minimize the amount of time they spent doing invisible work. Similarly, it could also be

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.

37:14 Toxtli et al.

Invisible Labor Activity	Mean [min]	Median [min]	Std [min]	% workers
Doing HITs that eventually timeout (Payments)	32.3	4.5	1.5	37%
Starting HITs but then return (Lack of Guidance)	11.2	4.2	12.1	92%
Viewing their worker's dashboard (Payments)	10.6	2.8	16.3	97%
Sending messages (Lack of Guidance)	2.4	1.9	0.7	51%
Watching over requesters' profiles (Hypervigilance)	15.0	1.1	12.9	69%
Searching for general HITs (Hypervigilance)	3.6	0.9	5.6	96%
Managing queued HITs (Hypervigilance)	3.2	0.7	4.6	93%
Previewing HITs (Lack of Guidance)	1.5	0.6	1.0	66%
Viewing their earnings (Payments)	0.9	0.5	0.3	85%
Searching for filtered HITs (Hypervigilance)	3.9	0.5	0.6	46%
Checking Worker's qualifications (Hypervigilance)	0.4	0.2	0.0	27%
Login to MTurk (General Logistics)	0.3	0.1	0.1	64%
Reading HIT information (Lack of Guidance)	0.1	0.0	0.0	63%
Reading Platform Help (Lack of Guidance)	0.0	0.0	0.0	0%

Table 3. Overview of the invisible labor activities that workers did, the amount of time they dedicated daily to each activity per day, and the percentage of workers who engaged in the activity. Doing HITs that eventually time out was the median most time consuming activity; viewing their earnings was the most common activity.

that Masters workers have more experience using tools. 86% of our Masters worker participants reported using tools, while only 57% of the non-masters workers reported tool use. Additionally, these workers might not be using these tools as effectively as the master workers. Prior work had identified that there are differences in how experts and non-experts use tools [67, 98].

However, it is important to highlight that Table 2 does show that workers who used tools spent more time doing paid work (30 minutes more) and earned substantially higher wages (\$1.3 USD more daily, when measuring workers' median wages.) Notice that these results might be emerging because most tools focus on increasing the wages that workers receive for their paid labor [67]. But, given our results, we believe there is value in exploring mechanisms through which workers learn how to better navigate crowdsourcing markets to focus primarily on paid work.

Within this study of worker segments, we also studied the relationship between

adopting particular strategies and invisible labor. Prior work has shown that experienced workers often use strategies to boost their performance [54, 98]. This can include using different tools or multi-tasking. Our study identified that workers who completed HITs in batches did 9.6% less invisible work than workers who did not (see Table 2). The reasoning behind this finding is likely that within batches, the same type of tasks is continuously presented to workers (one after the other). Therefore, workers do not have to search for new tasks (thus reducing their invisible labor). Batch tasks are also usually similar, so workers do not have to spend time context switching [74].

4.0.3 Quantifying Categories of Invisible Labor. We were also interested in understanding the type of invisible labor that was the most taxing for workers. Table 3 presents an overview of the different invisible labor activities that our plugin detected that workers did and the percentage of workers who engaged in each activity. For each activity, we also present in parenthesis the main categories to which the activity belongs. In Table 4, we present a summary of the time workers invested in each of these main categories. From Tables 3 and 4, we observe that the invisible labor category of “Payments” was the most time-consuming category (especially when taking the median value) and was also highly common among workers. For example, Table 3 shows how 97% of all workers in our study engaged in the Payments related activity of checking their daily earnings on their

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:15

Main Category of Invisible Labor	Mean	Median	Std
Payments	14 min	13 min	23.8
Hypervigilance	28 min	11 min	56.8
Lack of Guidance	16 min	6 min	62.1
Breaks	3 min	3 min	12.6
General Logistics	1 min	1 min	0.1

Table 4. Overview of the categories of invisible labor that workers did and the median amount of time they dedicated to it daily. The category of Payments was the one workers invested the most median time daily.

worker dashboard. Similarly, the most time-consuming activity was “doing HITs that eventually timeout,” which took a median of 4.5 minutes. Luckily, timeouts were not as common (only 37% of workers engaged in this activity). It is important to mention that timeouts relate to “Payments” because requesters on MTurk have to specify the amount of time that workers have for completing their tasks; if workers take longer than that time, the HIT is timed out, and workers do not get paid for any of the labor that they did for the HIT. We calculated timeouts only if the worker was actually working on the HIT (had any current mouse or keyboard-related activity on the HIT). The timeouts we detected were, therefore, cases where the worker was actively doing labor but at the end did not get paid for it.

To understand the details of the workers who engaged the most in this type of highly taxing invisible labor, we first identified the workers who were outliers (i.e., invested the most time in this activity) and then conducted a manual inspection of their digital traces. We considered outliers to be the workers whose time invested for this particular activity was above the 95th percentile (typical method to calculate outliers [56]). We observed that in this case, the outliers tended to be workers who accepted a high number of HITs within a given time window (likely to avoid having other workers take the HIT before them). However, the problem was that it would sometimes take workers significant time to get to some of the HITs they had “reserved” for themselves, and hence they experienced timeouts. We thus believe there is value in exploring tools [81], that based on workers’ log data, can automatically learn the best amount of time that should be allocated for a given task and then

recommend to requesters to use a significantly higher time window than that time to avoid timeouts and also be sympathetic with the labor practices of some workers.

Table 4 also shows that the second most time-consuming category was that of Hypervigilance, taking workers' a median of 11 minutes daily. The Hypervigilance activity that took the most time was watching over requesters' profiles. It is likely that workers engaged in this activity because through this they could more easily grab the HITs that their favorite requesters posted [51]. Upon manual inspection of workers' digital traces, we identified that the workers who invested the most time in this activity (i.e., the outliers, which we calculated with a similar method as stated above), were the workers who appeared to hunt the profiles of multiple requesters ready to be "on-call". (In specific, these workers opened the profile pages of multiple requesters and then iterated through the list of profile pages, likely inspecting if the requesters had posted anything new.)

Finally, the third most time-consuming category was "Lack of Guidance," which took a median of six minutes daily. The most time-consuming activity here were cases when workers started a HIT but then decided to return it. There are several reasons why workers might engage in this behavior; for example: workers realize that the HIT is more time-consuming than they expected; or the HIT involves skills that the worker lacks; or the HIT consists of activities that the worker does not enjoy. In general, these are instances where the HIT instructions likely did not correctly guide the worker on the type of labor to expect, and hence the worker had to return the HIT. Prior

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.

37:16 Toxtli et al.

Perception of Time in Invisible Labor Percentage of Workers Perceived Time Actual Time

Far too much time	25%	3 hrs 4 min	2 hrs 23 min
Too much time	38%	2 hrs 3 min	1 hr 40 min
An adequate amount of time	29%	1 hr 50 min	1 hr 32 min
Too little time	5%	1 hr 15 min	56 min
Far too little time	3%	1 hr 40 min	

Table 5. Summary statistics of workers' perceptions of how much time they felt they invested in invisible labor. Notice that the perceived and actual times are the medians for each perception group.

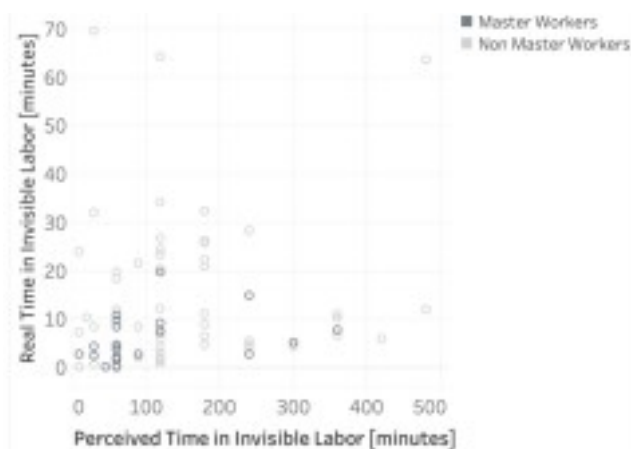


Figure 6. Comparisons of the estimated and actual time that workers invested in invisible labor. Most workers overestimated how much time they dedicated to invisible labor.

work has already reported how the lack of guidance can lead to these types of dynamics [43, 83]. From Table 3, we note that the activities related to the Lack of Guidance were

actually some of the most commonplace for workers and also some of the most time consuming (e.g., 92% engaged in starting HITs but then returning them; and this was also the second most time consuming activity.) It was surprising to see the large percentage who returned HITs. Upon manual inspection of the outliers, we observed that they appeared to primarily follow a discard-by-doing labor pattern [68].

4.1 Perceptions of Invisible Labor

Workers from digital labor platforms typically underestimate the actual amount of time and effort they dedicate to invisible labor [113]. However, workers' perceptions of invisible labor can play a strong role in how they feel about their work. In this section, we investigate the amount of time that crowd workers believe they invested in invisible labor and their satisfaction. For this part, we use the initial survey that we gave workers, which was inspired by prior work [4]. Through this, we found that workers in our study estimated that they spent a median of 2 hours daily on invisible labor on MTurk (with the minimum time that some workers' estimated as 0 and a maximum of 8 hours.) Figure 6 plots the actual time workers invested in invisible labor against perceived time. Notice that each point represents a worker in our study, and workers are color-coded based on whether workers are master-workers (dark gray) or non-master workers (light gray). We made this distinction given that prior work has identified that there are differences in how more experienced

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:17

workers operate [54, 98], and our results in the previous sections were also highlighting these differences. Notice that workers in Figure 6 who were able to accurately guess the amount of time they spent on invisible labor are located on the diagonal line, as that is when the actual time is equal to the perceived time. The cluster of points that we observe above the diagonal line close to the Y-axis showcases that the majority of crowd workers in our study overestimated the amount of time that they thought they invested in invisible labor. Notice that this overestimation occurred for both master and non-master workers. Next, we quantify the relative error of workers in estimating how much time they invested daily in invisible labor:



Through this, we identified that the median relative error was -0.14 . Notice that the negative value highlights that workers are overestimating how much time they dedicate to invisible labor, but our conservative metrics used to quantify invisible labor might also contribute to the error. Next, we analyze workers' satisfaction with the time they perceived they invested in invisible labor. Table 5 presents a summary of these statistics. Only 10% felt they invested too little time on invisible labor (8% felt they spent "too little time" and 2% "far too little"), while slightly more than half (63%) felt they invested too much in invisible labor (38% "too much time," and 25% "far too much time"). Lastly, 27% of workers considered they invested the right amount of time in invisible labor. Future work could study the type of labor dynamics that might lead workers to feel more satisfied with the amount of invisible labor that they do, and also what circumstances might lead them to feel the most dissatisfied.

5 DISCUSSION

The core result from our study is that crowd workers spent a median of 33 minutes of their

daily time on MTurk doing invisible labor, and this labor leads workers to drop their median hourly wage from \$3.76 to \$2.83. Notice that because we used conservative methods to measure invisible labor, we are obtaining a lower bound of the amount of invisible labor that exists on MTurk. However, this lower bound is still highlighting and providing quantitative support to the literature's qualitative claim that invisible work makes up a substantial fraction of the work done in crowdsourcing markets and, therefore, dramatically reduces workers' hourly wages [51]. Considering that the median hourly wage of workers is just \$3.76 (without considering invisible labor), it is clear that crowd workers still need a dramatic increase in their wages before we can consider this labor fair. However, this is not only something for requesters to consider, but also something for platforms, workers, and even policy makers. In this section, we discuss: the details of the most taxing categories of invisible labor that our study uncovered; design and policy solutions to mitigate invisible labor on crowdsourcing platforms. Additionally, we make an effort to connect with invisible labor in other workplaces, as well as with critical theory, to have a broader discussion on the implications of our research.

5.0.1 Most Common and Most Time-Consuming Invisible Labor. The invisible labor that the overwhelming majority of workers in our study practiced was around Payments. In fact, 97% of the workers in our study visited the earnings section on their worker's dashboard at least once daily. Crowd workers are likely visiting their earnings dashboard to ensure that they: (1) were paid for their labor; and (2) made a certain daily income amount [67]. For the first point, it is important to note that crowd workers typically have to deal with faceless requesters, machines that are outdated, unreliable internet connections, and have nowhere to report when things go wrong (e.g., report that a requester decided to unjustly withhold payment, or report that due to technical issues they can no longer access their MTurk account and earnings.) Pew Research reported that 30% of on-demand

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
37:18 Toxtli et al.

gig workers experienced situations where they were not paid for their labor [61, 103]. Similarly, the US Freelancers Union found that 71% of freelancers have struggled to collect payment for their work. As we note, ensuring payment is a critical and stressful aspect of crowd work [62]. For the second point, we have to be aware that most crowd workers struggle to make a minimum wage [57]. Therefore, another likely reason why workers were visiting their earnings dashboard was to see if they had made sufficient wages. The stress of not receiving payment for their labor or not receiving enough appears to be very present and real in crowd work.

Invisible labor around payments was actually also the most time-consuming, and one of the most critical, as it relates to workers' livelihood. To address this problem, designers could explore interfaces where workers are constantly informed of their current earnings. However, seeing their earnings constantly could also create stress on workers. Future work could explore optimal settings for displaying wages in crowd work. We also believe there is value in further exploring interfaces where requesters, platforms, and workers agree to fair wages [101]. Offering workers a space where they know they will be treated fairly could reduce repeated payment checking [112].

It is important to mention that this type of invisible labor is also present in other digital workspaces [53, 92]. For example, Uber had reports of drivers and passengers organizing to check how much a passenger was actually charged for a ride vs. how much the driver received. This dynamic emerged after Uber changed its pricing algorithm and did not provide transparency on how it functioned [12]. The lack of transparency not only led drivers and passengers to have to engage in this type of invisible labor, it also led them to feel cheated and betrayed by the platform [92]. Here it is important to highlight that this invisible labor does not only emerge due to the fault of requesters (passengers). But rather, platforms can play a key part in the promotion of this invisible labor. Here it can be important for platforms to see that this type of invisible labor is likely emerging out of mistrust and has

the potential to alienate people from their platforms.

5.0.2 Second Most Time-Consuming Category of Invisible Labor. Our study uncovered that the invisible labor category of Hypervigilance was the second most time-consuming for crowd workers. Workers spent a median of 11 minutes daily on this category of invisible labor. Crowd work has been championed as offering people the unique flexibility of working anytime and from anywhere [21, 117]. However, our work highlights how this flexibility is likely more of a myth. Crowd workers have to dedicate significant time daily to search for work and be on-call for requesters. Intuitively, this suggests that there are more workers on the site than there is work to be done. (If there were lots of requesters constantly posting lots of high-paying jobs, workers would not feel the need to be on call to get the good work.) This connects with prior work that shows that requesters have the majority of the power in this market partly due to the fact that there is an extreme concentration of a few requesters who post the majority of the tasks [28, 69]. Thus, workers are forced to take whatever jobs at whatever pay these few requesters post.

A way to start addressing this problem could be to build off the different tools and computational methods that have been developed to achieve fair compensation [112]. Potentially these computational methods could be extended by incorporating an invisible labor component. For instance, workers could be computationally guided to cooperate with each other to ensure fairer wages and minimize the amount of invisible labor in which they engage [32], such approach could be extended to potentially lead to reduced invisible labor. Similarly, we could also consider how algorithms that facilitate automatic task assignment and recommendations [60], could be helpful in reducing invisible labor by minimizing the task search time.

When thinking about the invisible labor around Hypervigilance, it is also important to notice that this type of invisible labor is one that promises workers high returns (especially as by being vigilant, workers can potentially earn high wages). Here, it can be important to identify that other

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:19

digital labor platforms have started to weaponize this type of invisible labor to manipulate workers to stay longer on their platforms [52, 110]. For example, Uber sends drivers messages to motivate them to keep being vigilant of surge pricing [12]. The following is an example message that Uber sends drivers to motivate them to remain vigilant of surge pricing: *“The weekend is here, and demand is on the rise in Lehigh Valley! Plan to go online tonight, and keep an eye out for surge around the area, where you can earn over 3X on fares! Stay online through midnight to take advantage of the highest fares. Uber on!”* [92]. In this context, we believe there is value in providing workers with tools that can help them to visualize how digital labor platforms might be manipulating them to engage them in free labor. Related, there is likely also value in tools that can inform workers of the likelihood of achieving specific wages if they engage in hypervigilance within particular time windows.

5.0.3 Invisible Labor in Other Workplaces and Policy. Researchers have argued that within our “capitalist societies”, there is a propensity to manage the workforce in ways that will profit the “capitalists” (who in this context could be considered to be Amazon or the requesters.) [26, 33]. Such “workforce management” can include defining what labor is counted and what labor is turned invisible [18, 19]. Labor visibility (what is counted) is considered to be especially important in this societal context because the cultural worth of a piece of labor is directly connected to how much the labor costs [89]. Work that is done for free (invisible) usually will fail to be valued [3, 33, 108]. Several labor collectives, researchers, practitioners, and individual citizens have therefore fought to empower workers to gain visibility and recognition for their work [26, 41]. For instance, the International Feminist Collective has been fighting for decades to give more visibility to the housework that

women perform [33]. The collective has argued that housework has been undervalued, underpaid, and its invisibility has been used as a means to empower primarily “white middle-class men to do lucrative waged jobs,” e.g., office work [19, 40]. This in return has profited companies and factories as they now have a more specialized and dedicated workforce [38, 39, 78].

In 2013, several of these collectives had a breakthrough when labor statisticians agreed internationally to begin measuring in official workforce surveys both paid and unpaid labor, such as housework [9, 11]. This inclusion influenced the development of new policies around invisible labor [9, 11]. Historically, policymakers had overlooked unpaid labor simply because the work was not included in the official statistics that they used to define policy [9, 111]. Its exclusion also meant that policymakers did not understand why the labor was problematic or the number of citizens who were impacted. But, by now counting and including the labor within the official stats that policymakers used for their decision-making, they were able to more easily pay attention to this type of labor, grasp its problematic, and design policy to address its problematics.

Inspired by the impact that the quantification of invisible labor has had in transforming policy within other industries and workplaces, our hope is that our plugin tool, study, and anonymized worker data, can in the future also be used to motivate new policies to improve the labor conditions of crowdworkers. However, given that the use of data in policymaking is usually an organic, political process [23] (which might not be obvious to outsiders, e.g., workers and their advocates), we believe there is value in designing socio-technical mechanisms that guide citizens on how they can best use the data from our plugin to drive policy innovation [16]. This could include tools that guide citizens on the time in which they should release the data on invisible labor to match the political cycle. Being in tune with the political cycle could help citizens to have a better chance at influencing policymakers [111]. Similarly, other tools could focus on helping citizens to easily visualize which policymakers might be most influenced by seeing the stats from our plugin on invisible labor. There is likely also value in tools that can guide citizens on how to use our plugin’s data to gather the public’s support and create pressure on policymakers [7, 105].

5.0.4 Design Implications & Future Work. Future work could explore mechanisms to help workers manage the time overheads from invisible labor. Notice that here there are still numerous aspects of invisible labor that need to be further investigated. For instance, are more experienced crowd workers able to reduce the amount of time they spent in invisible labor in comparison with novices? Our results highlight that at least master and non-master workers have similar perceptions of the amount of invisible labor they do. But more analysis in this space is necessary. Especially because there might be a benefit in designing tools that help novice workers adopt some of the strategies from more experienced workers [54, 98]. Other questions we are interested in exploring in this space are: How does the way that workers manage their invisible labor relate to their wages? How exactly does multi-tasking and context switching relate to invisible labor? Is a worker’s invisible labor increased when workers have to switch between HITs? Are there certain HITs or requesters that magnify workers’ invisible labor? Our hope is that by releasing our plugin, we will enable the scientific community to study this.

Notice that our plugin tool can be easily extrapolated to other digital labor platforms to help workers quantify the amount of invisible labor that they spend on those other workplaces (the only main piece that needs to be changed is the mapping between the websites the workers use and the work done on each platform; primarily if it is paid or unpaid labor). Our hope is that our tool will inspire cross-platform studies on invisible labor and will help the scientific community to derive principles around how invisible labor looks like across digital workplaces. As we described above, our plugin and study could also help to motivate action from policymakers. Facilitating tools for cross-platform auditing can be extremely

important as digital labor platforms have traditionally been black boxes. But, to design better platforms or drive policy change, it is crucial to understand what happens inside these platforms. Our hope is that our research will be a step forward to better understand and address the dynamics existing in these online spaces.

We believe there is likely value in exploring data visualizations that could help to better showcase the different types of invisible labor that crowd workers have to do. Here, we could take inspiration from the visualizations that Github has developed to showcase the labor surrounding the writing of collaborative code [26, 77]. Github has made great strides to provide visualizations that help people to rapidly understand the quantity, frequency, and duration of the contributions made by each individual to a codebase. Such visualizations in this context could help requesters to better grasp the amount and type of invisible labor that their tasks are forcing workers to do and potentially lead requesters to better compensate workers for their effort and time [112]. It is important here to consider how to design such visualizations to also not incite unhealthy competition between workers or enable abuse and surveillance from requesters [71, 93].

5.0.5 Critical Theory and Design to Address Invisible Labor in Different Digital Workplaces. An important question in CSCW is whether a new design truly engages with the root cause of a societal problem or if it is primarily dealing with the symptoms of a problem [2]. For example, a design could make a societal problem bearable. However, this might lead people to no longer have a need for addressing the root problem. In this setting, the design could provide enjoyable experiences to end-users; but it could also reinforce the structural issues that are harming end-users. Within this context, Herbert Marcuse, a theorist from the Frankfurt School of Critical Theory [66], introduced the concept of “one-dimensional” people who have a conformist understanding of society that does not allow them to critique or question how society could be different [84]. Marcuse argues that the one-dimensional person has lost her ability to critique society because consumerism has tricked her into having false needs and wishes (notice that consumerism is considered to be “a social and economic order that encourages the acquisition of goods and services in ever-increasing amounts” [8]). As a result, the person focuses on fulfilling those “fake needs” instead of questioning

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:21

the problematic societal structures in which she is immersed. According to Marcuse, this dynamic leads us to be imprisoned into one-dimensional thinking, and that makes it extremely challenging to critically question the structures and processes that exist in our society.

As CSCW researchers, we believe it is crucial that we question to what extent we are falling into one-dimensional thinking and possibly strengthening the structural issues that are already in place. This is especially important when designing interfaces that aim to address the problem of invisible labor in crowdwork and also within other digital labor platforms. Without this critical analysis, we might fall into designing interfaces that make the problem of invisible labor bearable; but we never address the systematic problems surrounding workers, requesters, and digital platform owners. Notice that engaging in such critical analysis is an ambitious, complex, and difficult undertaking, but as Marcuse discusses, it is very much necessary [2, 84].

Marcuse argues that a way to engage in such critical analysis and challenge our one-dimensional thought is by participating in artistic creativity that allows us to leave the reality that has been defined by society [2, 84, 85]. Artistic creativity facilitates developing new designs that are not confined by the current reality of what is possible and allows us to consider designs we might have been blind to consider otherwise. Based on this, we believe there is value in engaging with workers in “creative artistic co-design sessions.” These sessions would allow workers to creatively define the type of digital labor platforms that they would like to see and how they would design to address invisible labor [109].

Similarly, we believe there is value in drawing on scholarship that has studied the link between fiction and design [27, 35, 79, 115]. Here we envision we could engage researchers, workers, platform owners, and practitioners to use fictional narratives to design “alternative realities” to contemporary digital labor platforms and tools [34, 80].

5.0.6 Limitations. The insights from our research are limited by the methodology and population we studied. Our study also focused on breadth instead of depth to start to shed needed light on the quantification of invisible labor in crowd work. Notice that we had to develop specific tools in order to do our field study, which is not simple. However, these types of studies are important, especially given the lack of transparency that MTurk or other crowdsourcing platforms provide around invisible labor. Upon publication, we will open-source our plugin and anonymous worker data so that the scientific community can conduct longitudinal studies around invisible labor, as well as study other principles surrounding invisible labor.

6 CONCLUSION

We developed a new computational tool to be able to quantify and study the invisible labor of crowd workers on MTurk. We have demonstrated that the invisible labor that workers do can take a toll on their wages. Particularly, we saw that if we consider the amount of time that crowd workers invest in invisible labor, their hourly wages go down to \$2.83 from \$3.76. We also identified that the two most time-consuming categories of invisible labor revolved around payments and hyper-vigilance. Additionally, our study identified that workers tended to overestimate the amount of invisible labor that they believed they did. Our results also suggest there is a wide range of dynamics that influence the amount of invisible labor that a particular worker conducts. These different dynamics deserve more investigation.

Finally, we hope that our plugin tool inspires the auditing of different digital labor platforms and helps to potentially generate a range of positive policy innovations in digital work. Our paper has provided much-needed light to the invisible labor of crowd workers.

Acknowledgments. Special thanks to all the anonymous reviewers who helped us to strengthen the paper. This work was partially supported by NSF grant FW-HTF-19541.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
37:22 Toxtli et al.

REFERENCES

- [1] Ali Alkhatib, Michael S Bernstein, and Margaret Levi. 2017. Examining crowd work and gig work through the historical lens of piecework. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4599–4616.
- [2] Jeffrey Bardzell, Shaowen Bardzell, and Mark A Blythe. 2018. *Critical theory and interaction design*. MIT Press.
- [3] Veronica Beechey. 1977. Some notes on female wage labour in capitalist production. *Capital & Class*, 3 (1977), 45–66.
- [4] Michael S Bernstein, Eytan Bakshy, Moira Burke, and Brian Karrer. 2013. Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 21–30. [5] Michael S Bernstein, Joel Brandt, Robert C Miller, and David R Karger. 2011. Crowds in two seconds: Enabling realtime crowd-powered interfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. 33–42.
- [6] Jeffrey P Bigham, Raja Kushalnagar, Ting-Hao Kenneth Huang, Juan Pablo Flores, and Saiph Savage. 2017. On how deaf people might use speech to control devices. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. 383–384.
- [7] Monica Blagescu and John Young. 2006. *Capacity development for policy advocacy: current thinking and approaches among agencies supporting civil society organisations*. Overseas Development Institute London. [8] Encyclopaedia Britannica et al. 2008. *Britannica concise encyclopedia*. Encyclopaedia Britannica, Inc. [9] Mayra Buvinic and Elizabeth M King. 2018. Invisible No More? A Methodology and Policy Review of How Time Use Surveys Measure Unpaid Work. *United Nations Foundation, Data2X*. Available at https://data2x.org/wp-content/uploads/2019/05/Data2X-Invisible-No-More_1_2018.
- [10] Erran Carmel and Steve Sawyer. 2020. The Five-Dimensional Space of the Futures of Work: A View to 2030. In

- IFIP Joint Working Conference on the Future of Digital Work: The Challenge of Inequality. Springer, 295–309. [11] NU CEPAL. 2015. Classification of time-use activities for Latin America (CAUTAL). (2015). [12] Le Chen, Alan Mislove, and Christo Wilson. 2015. Peeking beneath the hood of uber. In Proceedings of the 2015 internet measurement conference. 495–508.
- [13] Justin Cheng and Michael S Bernstein. 2015. Flock: Hybrid crowd-machine learning classifiers. In Proceedings of the 18th ACM conference on computer supported cooperative work & social computing. 600–611. [14] Lucia Ciciolla and Suniya S Luthar. 2019. Invisible household labor and ramifications for adjustment: Mothers as captains of households. Sex Roles 81, 7-8 (2019), 467–486.
- [15] Marion Crain, Winifred Poster, and Miriam Cherry. 2016. Invisible labor: Hidden work in the contemporary world. Univ of California Press.
- [16] Emma Crewe and Mr John Young. 2002. Bridging research and policy: context, evidence and links. (2002). [17] Brandon Dang, Martin J Riedl, and Matthew Lease. 2018. But who protects the moderators? the case of crowdsourced image moderation. arXiv preprint arXiv:1804.10999 (2018).
- [18] Arlene Kaplan Daniels. 1987. Invisible work. Social problems 34, 5 (1987), 403–415.
- [19] Angela Davis. 1983. The approaching obsolescence of housework: A working-class perspective. Women, Race, and Class (1983), 222–44.
- [20] Valerio De Stefano. 2015. The rise of the just-in-time workforce: On-demand work, crowdwork, and labor protection in the gig-economy. Comp. Lab. L. & Pol'y J. 37 (2015), 471.
- [21] Xuefei Nancy Deng and KD Joshi. 2013. Is crowdsourcing a source of worker empowerment or exploitation? Understanding crowd workers' perceptions of crowdsourcing career. (2013).
- [22] Xuefei Nancy Deng and KD Joshi. 2016. Why individuals participate in micro-task crowdsourcing work environment: Revealing crowdworkers' perceptions. Journal of the Association for Information Systems 17, 10 (2016), 3. [23] Iqbal Dhaliwal and Caitlin Tulloch. 2012. From research to policy: using evidence from impact evaluations to inform development policy. Journal of Development Effectiveness 4, 4 (2012), 515–536.
- [24] Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, Panagiotis G Ipeirotis, and Philippe Cudré-Mauroux. 2015. The dynamics of micro-task crowdsourcing: The case of amazon mturk. In Proceedings of the 24th international conference on world wide web. 238–247.
- [25] Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, Panagiotis G. Ipeirotis, and Philippe Cudré-Mauroux. 2015. The Dynamics of Micro-Task Crowdsourcing: The Case of Amazon MTurk. In Proceedings of the 24th International Conference on World Wide Web (Florence, Italy) (WWW '15 Companion). Association for Computing Machinery, New York, NY, USA, 617. <https://doi.org/10.1145/2740908.2744109>
- [26] Catherine D'Ignazio and Lauren F Klein. 2020. Data feminism. Mit Press.
- [27] Carl DiSalvo, Tom Jenkins, and Thomas Lodato. 2016. Designing speculative civics. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. 4979–4990.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:23

- [28] Arindrajit Dube, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2020. Monopsony in Online Labor Markets. American Economic Review: Insights 2, 1 (March 2020), 33–46. <https://doi.org/10.1257/aeri.20180150>
- [29] Benjamin Duke and UK April. 2020. The effects of the COVID-19 crisis on the gig economy and zero hour contracts. Group 24 (2020), 8.
- [30] Margrit Eichler and Patrizia Albanese. 2007. What is household work? A critique of assumptions underlying empirical studies of housework and an alternative approach. The Canadian Journal of Sociology/Cahiers Canadiens de Sociologie (2007), 227–258.
- [31] Fairwork. 2020. The Gig Economy and Covid-19: Fairwork Report on Platform Policies. Fairwork Report on Platform Policies, Oxford, United Kingdom 24 (2020), 8.
- [32] Shaoyang Fan, Ujwal Gadiraju, Alessandro Checco, and Gianluca Demartini. 2020. CrowdCO-OP: Sharing Risks and Rewards in Crowdsourcing. Proceedings of the ACM on Human-Computer Interaction 4, CSCW2 (2020), 1–24.
- [33] Silvia Federici. 1975. Wages against housework. Falling Wall Press Bristol.
- [34] Casey Fiesler. 2018. Owing the Servers: A Design Fiction Exploring the Transformation of Fandom into 'Our Own'. Transformative Works and Cultures 28 (2018).
- [35] Casey Fiesler. 2021. Innovating Like an Optimist, Preparing Like a Pessimist: Ethical Speculation and the Legal Imagination. Colorado Technology Law Journal 19, 1 (2021).
- [36] Ailbhe Finnerty, Pavel Kucherbaev, Stefano Tranquillini, and Gregorio Convertino. 2013. Keep It Simple: Reward and Task Design in Crowdsourcing. In Proceedings of the Biannual Conference of the Italian Chapter of SIGCHI (Trento, Italy) (CHIItaly '13). Association for Computing Machinery, New York, NY, USA, Article 14, 4 pages. <https://doi.org/10.1145/2499149.2499168>
- [37] Claudia Flores-Saviaga, Yuwen Li, Benjamin Hanrahan, Jeffrey Bigham, and Saiph Savage. 2020. The Challenges of Crowd Workers in Rural and Urban America. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing. Vol. 8. 159–162.
- [38] Maria S Floro and Hitomi Komatsu. 2011. Gender and work in South Africa: what can time-use data reveal? Feminist Economics 17, 4 (2011), 33–66.
- [39] Maria S Floro and Anant Pichetpongsa. 2010. Gender, work intensity, and well-being of Thai home-based workers. Feminist Economics 16, 3 (2010), 5–44.

- [40] Nancy Folbre and Jayoung Yoon. 2007. What is child care? Lessons from time-use surveys of major English-speaking countries. *Review of Economics of the Household* 5, 3 (2007), 223–248.
- [41] Harley Frazis and Jay Stewart. 2012. How to think about time-use data: What inferences can we make about long-and short-run time use from time diaries? *Annals of Economics and Statistics/Annales d'économie et de statistique* (2012), 231–245.
- [42] Ujwal Gadiraju, Alessandro Checco, Neha Gupta, and Gianluca Demartini. 2017. Modus operandi of crowd workers: The invisible role of microtask work environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–29.
- [43] Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. 2017. Clarity is a worthwhile quality: On the role of task clarity in microtask crowdsourcing. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media*. 5–14. [44] Ruijiang Gao, Maytal Saar-Tsechansky, Maria De-Arteaga, Ligong Han, Min Kyung Lee, and Matthew Lease. 2021. Human-AI Collaboration with Bandit Feedback. *arXiv preprint arXiv:2105.10614* (2021).
- [45] Martha Garcia-Murillo and Ian MacInnes. 2019. The impact of AI on employment: a historical account of its evolution. (2019).
- [46] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- [47] Thomas Gillier, Cédric Chaffois, Mustapha Belkhouja, Yannig Roth, and Barry L Bayus. 2018. The effects of task instructions in crowdsourcing innovative ideas. *Technological Forecasting and Social Change* 134 (2018), 35–44. [48] Miriam Glucksmann. 2016. Completing and complementing: The work of consumers in the division of labour. *Sociology* 50, 5 (2016), 878–895.
- [49] Elham Shafei Gol, Mari-Klara Stein, and Michel Avital. 2018. Why Take the Risk?: Motivations of Highly Skilled Workers to Participate in Crowdfunding Platforms. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. Association for Information Systems. AIS Electronic Library (AISeL).
- [50] Patricia A Gouthro. 2009. Neoliberalism, lifelong learning, and the homeplace: problematizing the boundaries of 'public' and 'private' to explore women's learning experiences. *Studies in Continuing Education* 31, 2 (2009), 157–172. [51] Mary L Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.
- [52] Harish Guda and Upender Subramanian. 2019. Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science* 65, 5 (2019), 1995–2014.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
37:24 Toxtli et al.

- [53] Jonathan Hall, Cory Kendrick, and Chris Nosko. 2015. The effects of Uber's surge pricing: A case study. *The University of Chicago Booth School of Business* (2015).
- [54] Lei Han, Eddy Maddalena, Alessandro Checco, Cristina Sarasua, Ujwal Gadiraju, Kevin Roitero, and Gianluca Demartini. 2020. Crowd worker strategies in relevance judgment tasks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 241–249.
- [55] Lei Han, Kevin Roitero, Ujwal Gadiraju, Cristina Sarasua, Alessandro Checco, Eddy Maddalena, and Gianluca Demartini. 2019. All those wasted hours: On task abandonment in crowdsourcing. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 321–329.
- [56] Lingxin Hao, Daniel Q Naiman, and Daniel Q Naiman. 2007. *Quantile regression*. Number 149. Sage. [57] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P Bigham. 2018. A data-driven analysis of workers' earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [58] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Benjamin V Hanrahan, Jeffrey P Bigham, and Chris Callison-Burch. 2019. Worker demographics and earnings on amazon mechanical turk: An exploratory analysis. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [59] Erin Hatton. 2017. Mechanisms of invisibility: rethinking the concept of invisible work. *Work, employment and society* 31, 2 (2017), 336–351.
- [60] Danula Hettiachchi, Niels Van Berkel, Vassilis Kostakos, and Jorge Goncalves. 2020. CrowdCog: A Cognitive skill based system for heterogeneous task assignment and recommendation in crowdsourcing. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–22.
- [61] Paul Hitlin. 2016. *Research in the Crowdsourcing Age. a Case Study: How Scholars, Companies and Workers are Using Mechanical Turk, a "Gig Economy" Platform, for Tasks Computers Can't Handle*. Pew Research Center.
- [62] Sara Horowitz. 2015. The Costs of Nonpayment. Retrieved June 26, 2019 from <https://blog.freelancersunion.org/2015/12/10/costs-nonpayment/>
- [63] Ting-Hao Kenneth Huang and Jeffrey P Bigham. 2017. A 10-month-long deployment study of on-demand recruiting for low-latency crowdsourcing. In *Fifth AAAI Conference on Human Computation and Crowdsourcing*. [64] Mohammad Hossein Jarrahi. 2018. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons* 61, 4 (2018), 577–586.
- [65] Mohammad Hossein Jarrahi, Gemma Newlands, Brian Butler, Saiph Savage, Christoph Lutz, Michael Dunn, and Steve Sawyer. 2021. Flexible work and personal digital infrastructures. *Commun. ACM* 64, 7 (2021), 72–79. [66] Martin Jay. 1996. *The dialectical imagination: A history of the Frankfurt School and the Institute of Social Research*.

1923-1950. Vol. 10. Univ of California Press.

- [67] Toni Kaplan, Susumu Saito, Kotaro Hara, and Jeffrey P Bigham. 2018. Striving to earn more: a survey of work strategies and tool use among crowd workers. In Sixth AAAI Conference on Human Computation and Crowdsourcing. [68] Patrick Gage Kelley. 2010. Conducting usable privacy & security studies with amazon's mechanical turk. In Symposium on Usable Privacy and Security (SOUPS)(Redmond, WA). [69] Sara Constance Kingsley, Mary L. Gray, and Siddharth Suri. 2015. Accounting for Market Frictions and Power Asymmetries in Online Labor Markets. Policy & Internet 7, 4 (2015), 383–400. <https://doi.org/10.1002/poi3.111> [70] Aniket Kittur, Ed H Chi, and Bongwon Suh. 2008. Crowdsourcing user studies with Mechanical Turk. In Proceedings of the SIGCHI conference on human factors in computing systems. 453–456. [71] Aniket Kittur, Susheel Khamkar, Paul André, and Robert Kraut. 2012. CrowdWeaver: visually managing complex crowd work. In Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work. 1033–1036. [72] Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In Proceedings of the 2013 conference on Computer supported cooperative work. 1301–1318. [73] Siou Chew Kuek, Cecilia Paradi-Guilford, Toks Fayomi, Saori Imaizumi, Panos Ipeirotis, Patricia Pina, and Manpreet Singh. 2015. The global opportunity in online outsourcing. (2015). [74] Laura Lascau, Sandy JJ Gould, Anna L Cox, Elizaveta Karmannaya, and Duncan P Brumby. 2019. Monotasking or Multitasking: Designing for Crowdworkers' Preferences. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. 1–14. [75] Walter S Lasecki, Christopher D Miller, and Jeffrey P Bigham. 2013. Warping time for more effective real-time crowdsourcing. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. 2033–2036. [76] Walter S Lasecki, Jeffrey M Rzeszotarski, Adam Marcus, and Jeffrey P Bigham. 2015. The effects of sequence and delay on crowd work. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. 1375–1378. [77] Zhifang Liao, Dayu He, Zhijie Chen, Xiaoping Fan, Yan Zhang, and Shengzong Liu. 2018. Exploring the characteristics of issue-related behaviors in github using visualization techniques. IEEE Access 6 (2018), 24003–24015.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
Quantifying the Invisible Labor in Crowd Work 37:25

- [78] Linda Lim. 1983. Capitalism, imperialism, and patriarchy: the dilemma of third-world women workers in multinational factories. Women, men and the international division of labor (1983), 70–91. [79] Joseph Lindley and Paul Coulton. 2015. Back to the future: 10 years of design fiction. In Proceedings of the 2015 British HCI Conference. 210–211. [80] Conor Linehan, Ben J Kirman, Stuart Reeves, Mark A Blythe, Theresa Jean Tanenbaum, Audrey Desjardins, and Ron Wakkary. 2014. Alternate endings: using fiction to explore design futures. In CHI'14 Extended Abstracts on Human Factors in Computing Systems. 45–48. [81] Anna Lioznova, Alexey Drutsa, Vladimir Kukushkin, and Anastasia Bezzubtseva. 2020. Prediction of Hourly Earnings and Completion Time on a Crowdsourcing Platform. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 3172–3182. [82] Caitlin Lustig, Sean Rintel, Liane Scult, and Siddharth Suri. 2020. Stuck in the middle with you: The Transaction Costs of Corporate Employees Hiring Freelancers. Proceedings of the ACM on Human-Computer Interaction 4, CSCW1 (2020), 1–28. [83] VK Manam and Alexander Quinn. 2018. Wingit: Efficient refinement of unclear task instructions. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Vol. 6. [84] Herbert Marcuse. 1964. One-dimensional man: Studies in the ideology of advanced industrial society. Beacon Press. [85] Herbert Marcuse. 1966. The individual in the great society. Collected Papers of Herbert Marcuse Vol. 2: Towards a Critical Theory of Society (1966), 59–80. [86] Anoush Margaryan. 2016. Understanding crowdworkers' learning practices. (2016). [87] David Martin, Benjamin V Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a turker. In Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing. 224–235. [88] Brian McInnis, Dan Cosley, Chaebong Nam, and Gilly Leshed. 2016. Taking a HIT: Designing around rejection, mistrust, risk, and workers' experiences in Amazon Mechanical Turk. In Proceedings of the 2016 CHI conference on human factors in computing systems. 2271–2282. [89] Yann Moulrier-Boutang. 2011. Cognitive capitalism. Polity. [90] Gemma Newlands and Christoph Lutz. 2020. Crowdwork and the mobile underclass: Barriers to participation in India and the United States. new media & society (2020), 1461444820901847. [91] Sihang Qiu, Ujwal Gadiraju, and Alessandro Bozzon. 2020. Improving worker engagement through conversational microtask crowdsourcing. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. 1–12. [92] Alex Rosenblat. 2018. Uberland: How algorithms are rewriting the rules of work. Univ of California Press. [93] Jeffrey Rzeszotarski and Aniket Kittur. 2012. CrowdScope: interactively visualizing user behavior and output. In Proceedings of the 25th annual ACM symposium on User interface software and technology. 55–62. [94] Jeffrey M Rzeszotarski, Ed H Chi, Praveen Paritosh, and Peng Dai. 2013. Inserting micro-breaks into crowdsourcing workflows.

(2013).

- [95] Jeffrey M Rzeszutowski and Aniket Kittur. 2011. Instrumenting the crowd: using implicit behavioral measures to predict task performance. In Proceedings of the 24th annual ACM symposium on User interface software and technology. 13–22.
- [96] Susumu Saito, Chun-Wei Chiang, Saiph Savage, Teppei Nakano, Tetsunori Kobayashi, and Jeffrey Bigham. 2019. Turkscanner: Predicting the hourly wage of microtasks. In The World Wide Web Conference, 3187–3193. [97] Shruti Sannon and Dan Cosley. 2019. Privacy, Power, and Invisible Labor on Amazon Mechanical Turk. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. 1–12.
- [98] Saiph Savage, Chun Wei Chiang, Susumu Saito, Carlos Toxtli, and Jeffrey Bigham. 2020. Becoming the Super Turker: Increasing Wages via a Strategy from High Earning Workers. In Proceedings of The Web Conference 2020. 1241–1252. [99] Saiph Savage and Mohammad H Jarrahi. 2020. Solidarity and AI for Transitioning to Crowd Work during COVID-19. In The New Future of Work Symposium 2020.
- [100] Nigel R Shadbolt, Daniel A Smith, Elena Simperl, Max Van Kleek, Yang Yang, and Wendy Hall. 2013. Towards a classification framework for social machines. In Proceedings of the 22nd International Conference on World Wide Web. 905–912.
- [101] M Six Silberman, Bill Tomlinson, Rochelle LaPlante, Joel Ross, Lilly Irani, and Andrew Zaldivar. 2018. Responsible research with crowds: pay crowdworkers at least minimum wage. Commun. ACM 61, 3 (2018), 39–41.
- [102] RACHEL N SIMONS, DANNA GURARI, and KENNETH R FLEISCHMANN. 2020. "I Hope This Is Helpful": Understanding Crowdworkers' Challenges and Motivations for an Image Description Task. ACM CSCW (2020). [103] Aaron Smith. 2016. Labor platforms: Technology-enabled 'gig work'. Pew Research Center, Internet and Technology. November 17 (2016).
- [104] Daniel Spurk and Caroline Straub. 2020. Flexible employment relationships and careers in times of the COVID-19 pandemic.

Proc. ACM Hum.-Comput. Interact., Vol. 2, No. CSCW, Article 37. Publication date: November 2021.
37:26 Toxtli et al.

- [105] Daniel Start and Ingje Hovland. 2004. Tools for policy impact: a handbook for researchers. Overseas Development Institute London.
- [106] Hao Su, Jia Deng, and Li Fei-Fei. 2012. Crowdsourcing annotations for visual object detection. In Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence.
- [107] Will Sutherland and Mohammad Hossein Jarrahi. 2018. The sharing economy and digital platforms: A review and research agenda. International Journal of Information Management 43 (2018), 328–341.
- [108] Tiziana Terranova. 2000. Free labor: Producing culture for the digital economy. Social text 18, 2 (2000), 33–58.
- [109] Carlos Toxtli and Saiph Savage. 2020. Meta-Gig: Empowering anyone to create crowd marketplaces. Avances en Interacción Humano-Computadora 1 (2020), 11–19.
- [110] Hai Wang and Hai Yang. 2019. Ridesourcing systems: A framework and review. Transportation Research Part B: Methodological 129 (2019), 122–155.
- [111] Vanesa Weyrauch and Gala Díaz Langou. 2011. Sound expectations: from impact evaluations to policy change. International Initiative for Impact Evaluation Working Paper 12 (2011).
- [112] Mark E Whiting, Grant Hugh, and Michael S Bernstein. 2019. Fair Work: Crowd Work Minimum Wage with One Line of Code. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Vol. 7. 197–206.
- [113] Rebecca Whiting and Gillian Symon. 2020. Digi-housekeeping: the invisible work of flexibility. Work, Employment and Society (2020), 0950017020916192.
- [114] Alex C Williams, Gloria Mark, Kristy Milland, Edward Lank, and Edith Law. 2019. The Perpetual Work Life of Crowdworkers: How Tooling Practices Increase Fragmentation in Crowdwork. Proceedings of the ACM on Human-Computer Interaction 3, CSCW (2019), 1–28.
- [115] Richmond Y Wong, Ellen Van Wyk, and James Pierce. 2017. Real-fictional entanglements: Using science fiction and design fiction to interrogate sensing technologies. In Proceedings of the 2017 Conference on Designing Interactive Systems. 567–579.
- [116] Meng-Han Wu and Alexander Quinn. 2017. Confusing the crowd: Task instruction quality on amazon mechanical turk. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Vol. 5. [117] Ming Yin, Siddharth Suri, and Mary L. Gray. 2018. Running Out of Time: The Impact and Value of Flexibility in On-Demand Crowdwork. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). ACM, New York, NY, USA.

Received April 2018; revised July 2018; accepted September 2018

