The Crowd is a Collaborative Network

Mary L. Gray
Microsoft Research
Cambridge, USA
mlg@microsoft.com

Siddharth Suri
Microsoft Research
New York, USA
suri@microsoft.com

Syed Shoaib Ali
Independent
Nahan, India
syedshoaib@outlook.com

Deepti Kulkarni
Peepaldesign
Bangalore, India
deepti.d2005@gmail.com

ABSTRACT
The main goal of this paper is to show that crowdworkers collaborate to fulfill technical and social needs left by the platform they work on. That is, crowdworkers are not the independent, autonomous workers they are often assumed to be, but instead work within a social network of other crowdworkers. Crowdworkers collaborate with members of their networks to 1) manage the administrative overhead associated with crowdwork, 2) find lucrative tasks and reputable employers and 3) recreate the social connections and support often associated with brick and mortar-work environments. Our evidence combines ethnography, interviews, survey data and larger scale data analysis from four crowdsourcing platforms, emphasizing the qualitative data from the Amazon Mechanical Turk (MTurk) platform and Microsoft's proprietary crowdsourcing platform, the Universal Human Relevance System (UHRS). This paper draws from an ongoing, longitudinal study of crowdwork that uses a mixed methods approach to understand the cultural meaning, political implications, and ethical demands of crowdsourcing.

Author Keywords
crowdsourcing; social networks; collaboration; online labor

ACM Classification Keywords
J.4 Social and Behavioral Sciences: Sociology

INTRODUCTION
Crowdsourcing is the distribution of work through an open call [20]. Typically on crowdsourcing-for-pay sites (what we refer to throughout the paper as sites for "crowdwork"), such as Amazon Mechanical Turk, task creators use an API to place a task on the site. Workers then search the site for a task available to them that suits their interests and complete the chosen task. Finally, task creators review the work and either accept or reject it. If the work is accepted the workers are then compensated. Throughout this exchange, the API of the crowdsourcing site and the task itself mediate the interaction between the task creators and the task workers. As a result, the personal characteristics of the task worker are invisible to the task creator. For example, the task creator has no way of knowing if the task worker is male or female, young or old, religious or atheist, etc. Furthermore, the social network around the task worker is also hidden from the task creator. For example, the task creator has no way of knowing if the task worker has many contacts who also do crowdwork, receive help in doing a given task, or share information about tasks or task creators with other workers. Yet crowds are often thought of as a disaggregated, distributed set of *independent* workers.

The central research questions this work addresses are: do crowdworkers collaborate and, if so, why do they collaborate, how do they collaborate, and what do they collaborate on? Our ethnographic interviews, surveys and data analysis of four different crowdwork platforms, show that the presumed independent crowd of workers is actually a rich network of collaboration. Our evidence suggests that crowds are actually networks with edges hidden by the crowdsourcing platform and its API. When platforms do not natively support collaboration, workers create widespread yet invisible forms of collaboration that take place off-platform. Our paper argues that workers collaborate to address unmet social and technological needs posed by the crowdsourcing platform. These empirical findings underpin our theory that workers' investments in collaboration reflect needs for social relationships associated with the concept of employment that persist, even in the absence of a traditional workplace.

This paper expands our understanding of collaboration, including but also going beyond just the interactive practices of workers in the moment of task completion. Workers invested in making crowdwork a form of reliable employment engage in three types of collaboration to meet

CSCW '16, February 27-March 02, 2016, San Francisco, CA, USA

ACM 978-1-4503-3592-8/16/02…$15.00
DOI: http://dx.doi.org/10.1145/2818048.2819942

the social and technical needs. First, workers collaborate to manage the administrative overhead involved with doing crowdwork, such as signing up for an account, avoiding scams and receiving checks. Simplifying a complex signup process is, at first glance, a technical need. But no technical system gains an individual's trust on its own. Opening an account on a crowdsourcing site involves submitting intimate financial information. Thus having a friend help a worker open an account conveys a certain amount of trust, trust that this site is not a scam. It takes social endorsement through word of mouth and friends' presence to convey safety to us. The trust necessary to open an account, particularly associated with employment opportunities, requires social collaboration.

Second, we show that workers communicate via phone, forums, chat, Facebook—even in person—to share information about new tasks and good requesters. Displaying the expected hourly rate of each task and the reputation of each requester are technical fixes that would aid workers in searching for tasks. That said, a friend vouching for a requester or the time and effort it will take to complete a task provides the worker with a level of confidence that cannot be easily replicated with a purely technical solution.

Finally, we show that workers turn to each other to actually do the work itself, recreating social work environments to encourage each other's progress and development as crowdworkers. Social interaction is a basic human need. The crowdsourcing API which defines the interaction between one worker and one requester strips away any social interaction between the worker and the requester and, often, discourages or excludes any interaction between workers. Workers use online forums like the proverbial office break room or water cooler, to empathize, commiserate and confide with other workers. Like the previous two examples this is a need with both social and technical components. Taken as a whole, these three forms of collaboration among workers are more than an immediate, pragmatic attempt to compensate for a broken technical system. Collaboration represents the value that humans tenaciously assign to social connections in work environments.

We focus on two decentralized marketplaces for paid crowdwork, juxtaposing them to the experience of workers on two crowdwork platforms that explicitly build collaboration into their workflows. Our findings offer designers a mandate for considering the inevitability of collaboration among workers [25,35] and call out the value of incorporating collaboration into crowdsourcing workflows. Our research also advances the call of Kittur et al [29] to take seriously crowdsourcing's capacity beyond more efficient task output and design paid crowdwork systems that recognize the sociality of work and the shared identities produced through paid collaboration. Crowdsourcing must fully address and integrate both the technical and social needs of the workers to advance as systems for organizing productivity.

RELATED WORK

Research looking at computer-supported cooperative work has long recognized the value of collaboration for completing tasks distributed over a network [34,49]. Below, we review CSCW's investments in understanding the value of collaborative work and how work on collaboration has changed over the years. Throughout this section we will underscore the difference between systems that engineer collaboration through predetermined workflows and systems that recognize and value the collaboration that workers organically generate among themselves.

CSCW Investments in Collaborative Work

Computer-supported cooperative work has long grappled with the unique challenges posed by combining the social and technical demands that come with facilitating human interaction [15,21]. Take, for example, Mandivwalla and Olfman's [34] review of "groupware" for collaboration. Their analysis stretches back to systems built as early as 1968. The authors provide a thorough overview of the key challenges that accompany collaboration dynamics, from "multiple tasks and work methods" to group members' "multiple behaviors, permeable boundaries, and context" [34]. But much of the early CSCW literature examining workplace collaboration took a shared identity (e.g., "we are all employees of this firm") and shared goals (e.g., "we know we are trying to execute this shared project") for granted.

Building for Autonomous, Distributed Individuals through Crowdsourcing

Early work building collaboration for distributed systems assumed completing a multifaceted project required some form of collaboration between the organization or institution issuing the task and individuals doing the work at hand [15,21]. The turn to crowdsourcing suggested an escape from the intractable complexities of mediating human collaboration [28,46,49]. Individual actors could now contribute small bits of information or effort, working independently of each other, to achieve a larger (even opaque) collective goal [47], from editing large documents [27,28] and funding inspiring projects [16, 17] to incentivizing competition on a massive scale to surface the best idea [9]. The systems aggregated the individual efforts, redundancies and all, into a cohesive result [9,47]. Assembling independent results from workers is a simple way to engineer collaboration between them. This marks a decided shift in the approach to completing tasks through distributed systems.

Engineering Coordination on Crowdsourcing Platforms

Prior work has underscored the importance of explicitly engineering coordination in crowdwork to improve workflows and work output. As early as 2009 [33], Little et al. suggested a framework for crowdwork that serves to decrease mistakes by structuring an iterative process using their system, TurKit. Their study asked workers to attempt to type out handwritten messages, leaving blanks for words they did not understand. Their responses were then given to repeated workers in a chain until the passage was correctly transcribed. They proved that such an iterative method allows for decreased mistakes in completed work. Similarly, Ambati et al. [4] suggested a collaborative workflow model that would better support crowdwork for translation tasks. Their approach draws attention to the importance of collaboration between requesters and workers, and breaks up translation tasks such that different workers with different skillsets can be applied to various areas of the translation process, splitting workflow between different sets of workers.

Valuing Collaboration in Crowdsourcing

More recently, crowdsourcing research has turned its attention back to the value of coordinating human effort at scale. Some of this work is informed by the limits of leaving individuals to self-organize their contributions to larger projects [40]. Since the API, which governs interactions between workers and task creators, has no built-in way for workers to communicate with each other, task creators and crowdsourcing platform builders assume that workers do not communicate with each other as part of their work unless the platform is engineered to facilitate it. There is also a growing recognition of the tangible benefits of incorporating what Huang has recently dubbed "social facilitation" [13,23,24], and the value of group identity to collaboration [48]. "Friendsourcing," for example, [7] offers a valuable case study of incorporating one's social network into crowdsourcing processes. While Bernstein et al. [7] identify potential issues around friendsourcing and collaborative work (i.e. social loafing), they contend that the increase in output quality outweigh the potential challenges. Many of the studies examining the value of integrating social networks into crowdsourcing facilitate collaboration between sets of workers rather than making room for the organic collaborations that workers develop themselves [10].

Kulkarni et al [30] built an impressive system to test the value of a "scaffold" approach to task-based work. Specifically, the paper presents a system called 'Turkomatic,' built as a means of studying how effectively crowds can be used to support the execution of complex work. Through their research, they found that work undertaken by the crowd was improved when requesters were able to intervene and communicate with workers during the workflow process. Turkomatic effectively gives requesters the ability to, what we refer to as, "engineer" communication and collaboration between workers and requesters [30]. Furthermore, several researchers [6,19,18,35,23,39] have identified key ingredients to producing crowdwork environments better able to mentor workers resulting in skills-building and advancement opportunities that incorporate facilitating connections among workers.

Many, if not most, of the prior works have focused on how to harness collaboration either among workers or between workers and requesters by engineering different goal-oriented workflows. But, as Lee and Paine recently noted in their proposal for a Model of Coordinated Action (MoCA), CSCW frameworks for collaboration require much greater nuance, beyond time on task and shared, clearly-defined "goal-directedness" [32]. Their model is well-rooted in theories of work from the sociological literature such as Anselm Strauss' sociological concept of "articulation work"—actions individuals take to assemble the resources needed to accomplish something in a specific setting [45]. As noted sociology of work scholar Andrew Abbott suggests, studies of work and occupations are most fruitful when they examine "how exactly work is situated in the human experience" rather than narrowly focusing on the macro or micro mechanics of economic productivity and efficiencies [1,2]. Inspired by this sociological attention to how workers create meaning from their daily practices, we study how workers collaborate organically. That is, how they self-organize to do crowdwork. We consider not only the pervasiveness of collaboration that has organically grown among workers but also the variety and, arguably, overlooked value of organic collaboration to crowdsourcing as a new iteration of employment.

METHODS

Our data draws from a larger, mixed-method study, conducted from July 2013 through April 2015, that compares four crowdsourcing platforms: Amazon.com's publicly available Mechanical Turk (MTurk), Microsoft's proprietary Universal Human Relevance System (UHRS), the social entrepreneurial commercial start-up, LeadGenius, and Amara.org, a not-for-profit site dedicated to translating content for transnational audiences and the hearing impaired. UHRS and MTurk are similar platforms in that they both host an online, decentralized marketplace of microtasks, ranging from image-tagging to marketing surveys, posted daily. LeadGenius (formerly MobileWorks) focuses on business to business (B2B) services, helping its corporate clients identify and deliver "hard to find company information, key decision-makers, and contacts in new markets" (LeadGenius website). Amara.org provides captioning and translation services to a range of clients, most notably TED's Open Translation Project. We focused on four different crowdwork platforms in order to sample a range of platform approaches to worker support, from a completely "hands-off" decentralized market place approach, as seen in both MTurk and UHRS, to platforms invested in fostering worker interaction and task collaboration, in the case of LeadGenius and Amara.

The data sets analyzed for this paper include: 1364 completed surveys, collected from respondents living in India and United States, posted to the four crowdwork platforms studied for this project, between July 2013 and July 2014; results from posting a HIT to MTurk with 4,856 responses; analysis of 118 interviews and participant observations conducted in person from September 2013 to March 2015 in India.

Survey Data

We discuss our survey first as it provided data and a mechanism for recruiting interviewees. Our survey asked respondents doing paid crowdwork on our four platforms under study a range of questions, from inquiries about basic demographics to specifics concerning computer literacy and Internet skills. Questions focused on assessing the time and effort spent finding tasks, motivations for crowdsourcing, language skills, estimated yearly income, and venues to find tasks online, among other questions. Workers on all platforms were paid for doing our survey.

Merely posting the survey on MTurk, as is commonly done by those conducting surveys about crowdwork, may over-sample MTurk workers who typically do surveys as tasks for work. Thus, in addition to posting the survey to MTurk, we also embedded the survey into separate image-labeling tasks and email classification tasks. After a worker did 10 email classifications, for example, a link appeared asking if they would like to do our survey for an additional bonus payment. Since our survey also served as a vehicle to recruit interview participants, this methodological innovation allowed us to reach workers who might not typically do surveys on MTurk. The UHRS workforce is managed through vendor relationships rather than an open market place making it impossible to give workers a bonus payment as part of an additional, attached task. Thus we were unable to embed our survey into a collection of microtasks as done with MTurk. Work is centrally organized and distributed on both Amara and LeadGenius. As such, we relied on the cooperation of the platform owners and their willingness to circulate information and links to our survey task to all workers on these two platforms. Workers on these platforms received emails and saw announcements on worker newsletters informing them that participation was completely confidential and would not be shared with platform managers. We obtained a total of 451 survey respondents who use MTurk, 684 who use UHRS, 168 who use Amara, and 188 who use LeadGenius (note that some workers may use more than one platform).

Interviews and Ethnographic Fieldwork

To date, we have completed 118 in-person, open-ended, semi-structured interviews in India. The majority of our interview participants came from the largest of the four platforms (among 33 UHRS, 62 MTurk, 21 LeadGenius, and 2 Amara workers respectively) with hundreds of hours of informal follow up interviews and observations among research participants. We focused on interviewing workers in India as it is one of the hubs for crowd labor. Interview participants were recruited in the following ways: an invitation at the end of the crowdwork platform survey to participate in an in-person interview, scheduled at their convenience; worker referrals; and online contacts made on worker discussion forums. All names used in the discussion below are pseudonyms chosen by the research participants.

The ethnographic observations allowed us to understand people's experiences with crowdwork and how they come to their understandings of crowdwork and its relationship to their every day lives. The first author spent a total of six months with two research assistants who were present the entire 18 month period of the interview and ethnographic phase of the project. Following the leads from our surveys and our mapping HIT described below, interviewing and fieldwork focused on 3 major IT centers in South India, specifically Hyderabad, Bangalore, Chennai, as well as parts of Kerala, and Delhi in the North. Most interviews took place in people's homes, local cafés, or parks. Interviews lasted anywhere between one hour and three hours. The initial interviews included the equivalent of $15 USD cash gift in appreciation for the time that individuals gave us, recognizing that participants gave up time that could have been spent earning money crowdsourcing. The interviews included spending time with each participant and, in most cases, meeting them in their homes to see their work set up and have them demonstrate how they did their crowdwork. Fieldwork also included observing participants in their homes, with their families and friends, and joining them at events at cricket fields, shopping bazaars, mosques or temples that they signaled as important to them. The research team spent an average of 40 hours per week with a core group of 40 participants over the course of the India-based fieldwork.

All of the surveys and interviews analyzed for this paper were conducted in English. Half of the India-based interviews were conducted in person in English by the first author or jointly with one of two research team members fluent in the interview participant's primary language (mother tongue), while the remaining interviews were conducted one-on-one by the India-based research team members (who are also 3rd and 4th author on this paper).

Data Analysis Techniques

Because we have not completed analysis of United States-based fieldwork and interviews, we focus our analysis on the interview transcripts and research fieldnotes collected in India. Our survey data does, however, indicate more similarities than differences between India-based and U.S.-based workforces. While recording was not feasible in some cases (particularly in loud café settings), we were able to record interviews with 72 participants. All interview participants also completed a project survey, either before or after our interview with them, allowing us to compare and contrast participants' responses to each format.

Because of the longitudinal nature of the fieldwork and interview process, some of the early qualitative data shaped later interviews and field observations. We used a semi-structured interview protocol to ensure that the interviews were as consistent across researchers as possible and that the interviews covered specific subject areas. For example, we asked: how workers first found the platforms they currently or most recently used; motivations for doing this work;

experience with tasks; information-sharing practices; other work and educational interests; and their aspirations for the future. We used a qualitative interpretative approach for analysis of the fieldnotes and interview materials. Specifically, we prioritized participants' accounts, then explored themes in relation to insights from the survey results and mapping HIT that we conducted (which is further discussed in the next section).

Critical to our approach was not to impose categories of collaboration on participants' reflections a priori, but to document how participants articulated their understanding of crowdwork activities both through their recollections and live demonstrations of their work. As such, the categories of collaboration that we use here come from both the descriptions of activities and language from the participants themselves, corroborated by quantitative data. To begin the analysis of the qualitative materials (interviews and fieldnotes), the lead author read through transcripts and fieldnotes closely, making notes on emerging themes that related to our research questions. After this initial analysis, the co-authors used several data sessions to discuss observations from the fieldnotes and transcripts and how the quantitative further illustrated or countered the themes found in the qualitative materials. In coding qualitative data and identifying forms of collaboration, we built on the previously established presence and value of collaboration [35,42], analyzing the material both in detail and comparing the resulting themes to one another.

Geographic Mapping HIT

We combined our ethnographic work with a large scale online task which had two goals: measure the geographic location of the population of active workers and, more importantly for this paper, measure how workers share tasks. The task and its instructions were quite simple. Upon accepting the task, workers were shown a Bing map of the world and told, "Just double click your location and submit the HIT—It's that simple." This allowed us to achieve the first goal of this task. We intentionally asked workers to self-report their location so they could report their location down to any level of granularity they were comfortable sharing. Since the map allowed workers to search, zoom and pan as they saw fit, they were free to put a pin on their house, their neighborhood, their county, city etc. Moreover, workers were also told, "We will not reveal your location to anyone. Instead, we will randomly move everyone's location a short distance," to protect their privacy.

After placing a pin on their location and clicking save, workers were then shown a Bing map of the world with the pins of the last 500 workers to do this HIT. Again, each pin was randomly perturbed to protect worker privacy and they were reminded of that on this page. We only showed the last 500 pins because testing revealed that showing more pins resulted in browser lag which caused a poor user experience. We showed them this map after they placed their pin to avoid biasing the results and so they would see first-hand that they are part of a global community of workers. On this page we also asked them "How did you find out about this HIT?" along with a pull-down menu with the following choices: searching the MTurk site, from an online forum (e.g. Turker Nation, MTurk Grind etc.), referred by a friend, following a requester, decline to answer, and other. This allowed us to achieve the second goal of this task, which is the most relevant for this paper, to understand how workers share task information and how widespread of a phenomenon task sharing is.

This HIT could easily be done in under one minute and we paid $0.25 for completing it. This is a much higher wage rate compared to other tasks on MTurk but seemed fair compensation for a task that did not offer practical experience or outcomes for workers [11]. This HIT received almost 5,000 pins and ran from April 23 - May 28, 2014.

DATA ANALYSIS/RESULTS

The pertinent survey data, ethnographic data, and interviews suggest three, widespread practices of collaboration among crowdworkers on the platforms we studied: 1) sharing administrative overhead to reduce costs of managing the work process, 2) sharing task information such as employment opportunities and 3) helping each other complete individual tasks. Below we offer examples from the ethnographic data of crowdworkers collaborating in each of these three ways.

Sharing Administrative Overhead to Reduce Costs of Managing the Work Process

Joseph, a 22 year old Christian student living in the south Indian state of Kerala, has worked on Amazon Mechanical Turk (MTurk) since 2012. He recently completed his Bachelors in Computer Applications. While Joseph planned to pursue a Masters degree in design at the time we met on January 12, 2014, his passion is music. He plays guitar as part of a local band, performing at small events. Joseph joined MTurk to earn money from home, so that he could leave his schedule as flexible as possible to pick up paid musical jobs. Joseph tried other online job websites but they all turned out to be fake. He found out about MTurk through an online jobs community that he found on Facebook. He wanted to open his own MTurk account but his first request for an account was rejected because he now lived in part of a modest patchwork of houses built by local residents, along the banks of the Karuvannur River and his identification papers listed an old home address. Unable to generate a valid postal address in MTurk's accounting system, Joseph bought an account from an agency called STS Education Institution. But, soon after, the account got suspended. Joseph, now knowing what to search for, found MTurk account sellers on Facebook. As Joseph put it "I found a person on Facebook who was from Thrissur [2 hour drive from Joseph's hometown] but he does not work on it [MTurk] anymore. I work using his account and give him 20% of my salary." The arrangement allowed Joseph to work on MTurk, keeping 80% of his earnings, approximately 20,000 Indian Rupees ($330 USD) per month. The person who sold Joseph the

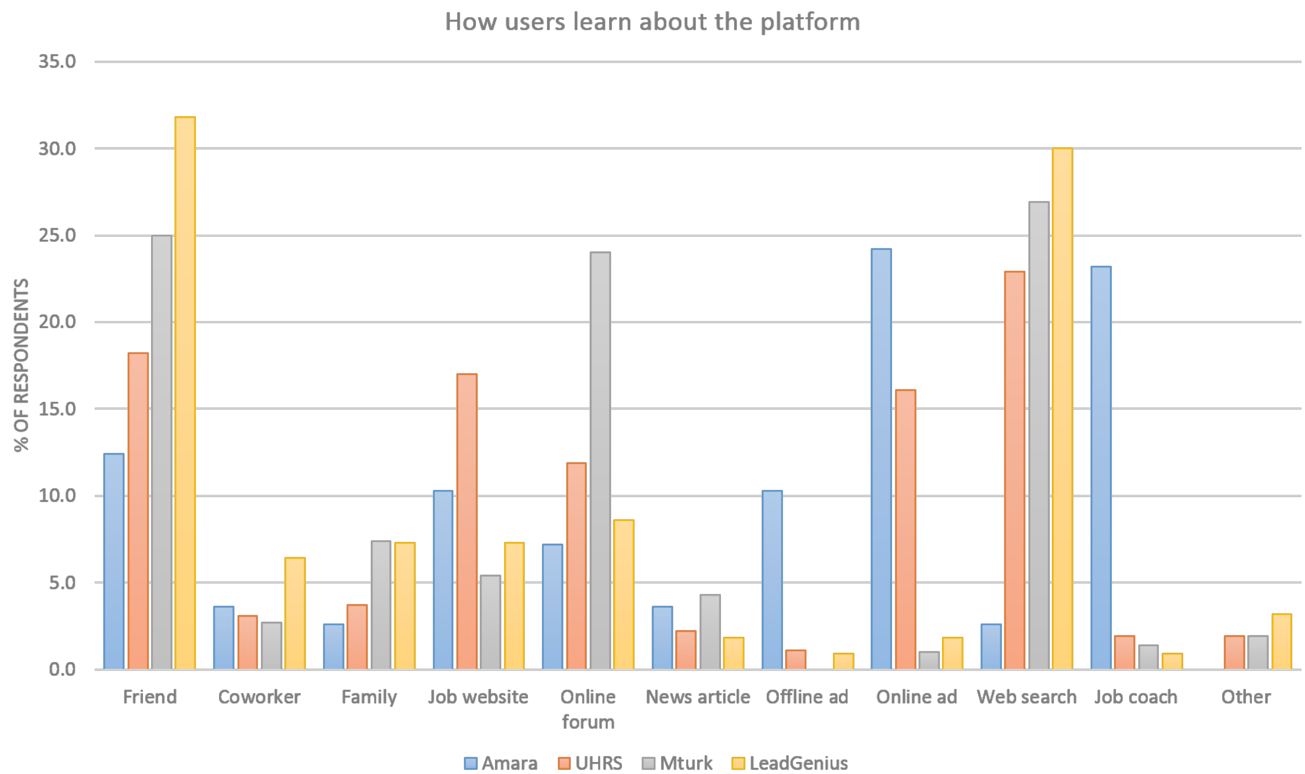

Figure 1. How users learned about each of the platforms studied.

account also accepts the paper checks, deposits them in his account, processing the checks and electronically transferring the funds to Joseph's online bank account. They have brokered this arrangement for more than 4 years now.

Joseph uses the money earned on MTurk to contribute to his family's household income, investing in his dad's stationary shop, buying his mother her first washing machine on her birthday, and a motorcycle for himself. Without the cooperation of this individual, who Joseph first found online and now considers a "work friend," Joseph would not be able to manage the administrative requirements of having a validated physical address and access to a bank for processing his pay. Several other participants described the very real barrier of being unable to deposit or cash the paper checks mailed by Amazon in their small towns as there were no banks where they lived. Others described using the postal addresses of friends and family members' established businesses because physical mail could be more reliably delivered to those addresses than one's home residence.

Like Joseph, Mohasin is also a student, 24 years old, working on MTurk part-time while studying for his Masters of Computer Applications. He lives in the southern city of Kochi with his mother; his father passed away two years ago. His two sisters are married and living in the United States. He came to know about MTurk from a friend and also from a newspaper ad by an institute, which helped individuals sign-up for their own MTurk accounts and provided basic training on navigating the site. Moshasin paid Rs. 1000 for his MTurk account and training. He is in daily contact with friends he has met on social networking sites and forums, predominantly closed Facebook groups started by fellow workers. Moshasin argues that "If we are not here, MTurk is also not here" arguing that workers' helping others join the site is what keeps the platform refreshed with new workers. He feels that "if we help others, others will also help us." He felt that "MTurk is thriving" because of workers' commitment to exchanging information about how the platform functions.

Kumuda, a 34 year old Hindu woman and computer trainer working on MTurk, articulated another commonly cited reason that workers relied on each other to identify and navigate the sign-up process for platforms like MTurk: many wanted to help relatives and friends in their hometowns avoid the business process outsourcing (BPO) scams they had experienced already. In the absence of a physical place of employment or any clear system for vetting the authenticity of BPOs, individuals came to rely on each other to help sort through the legitimate businesses and the ones simply looking to glean emails or other personal information from people looking for work opportunities. As Kumuda put it, "I actually started with outsourcing. I and a friend of mine were searching for job offers. We searched a lot but we faced great losses. Everywhere it was a scam. I am the first person in my area to find about MTurk. My friends have come to know about MTurk through me [so they know it is safe]."

As the examples above illustrate, a core form of collaboration workers discussed was helping each other identify reliable platform work and, in some cases, sign up for accounts. Figure 1 shows how widespread referring friends is. Roughly 25% of those surveyed in both the United States and India were referred to MTurk by a friend. LeadGenius had even higher rates of referrals from friends showing this phenomenon cuts across platforms. In addition, we found that workers accept and process payments as well as distinguish legitimate work opportunities from the ubiquitous scams that are part of the backdrop of decades of fly-by-night business process outsourcing (BPO) industries now flooding the world of online crowdwork. Having a friend recommend a crowdsourcing site helped workers avoid online scams by conveying trust. This would be difficult to achieve with a purely technical solution hence the need for interpersonal worker collaboration.

Sharing Task Information as Employment Opportunities

A second form of collaboration that we identified was finding and sharing information about tasks and specific requesters posting to the platforms. Workers created and circulated phone lists of task types and called each other when task creators posted good jobs to the platform.

For example, Sanjeev is a 22 year old student working on MTurk. He is an active blogger and has blogs on love and friendship, and earns much of his money through Google Ad Words. He learned about MTurk from his friend, a fellow classmate in a Masters of Computer Applications course. He felt MTurk was a good part-time job because he could keep an eye out for tasks late in the night when he was studying. In a joint interview with his college friend, Sanjeev said, “if I am working and find a good HIT then I call him and tell him about it.”

Fareed supports his family through his work on MTurk. He is a devout Muslim in his late 20s, and the eldest brother in his family. A native of Hyderabad, his uncles and father are increasingly pressuring him to join them as a driver for hire in the Arab Emirate states, a common employment opportunity for young Muslim men with only general high school educations and few employment opportunities in Hyderabad. He has worked on MTurk since 2011. He tried to get worker accounts on other crowdworking sites but has not been successful so far. Fareed was most concerned with managing his worker reputation on MTurk. As he noted, “rejections used to happen more [when I first signed up] as I didn’t know of requesters and the given instructions for tasks.” Fareed, dependent on MTurk for his primary income, was deeply invested in his reputation score, which is the fraction of tasks a worker submitted that have been approved. Fareed regularly turned to his childhood neighbor and good friend, Zafar, who also introduced him to the platform, to identify requesters who had a good reputation for responding to workers’ queries for clarifications about instructions. Fareed also asked for Zafar’s guidance on how to do certain tasks, as Fareed familiarized himself with the novel world of image-tagging and looking for physical addresses (location verification tasks) of places he had never lived himself organized by streets, names, and postal codes completely unfamiliar to him. The stakes of workers’ reputations on MTurk are high. A low reputation rating, even a rating in the low 90th percentile, locks most tasks out of the reach of

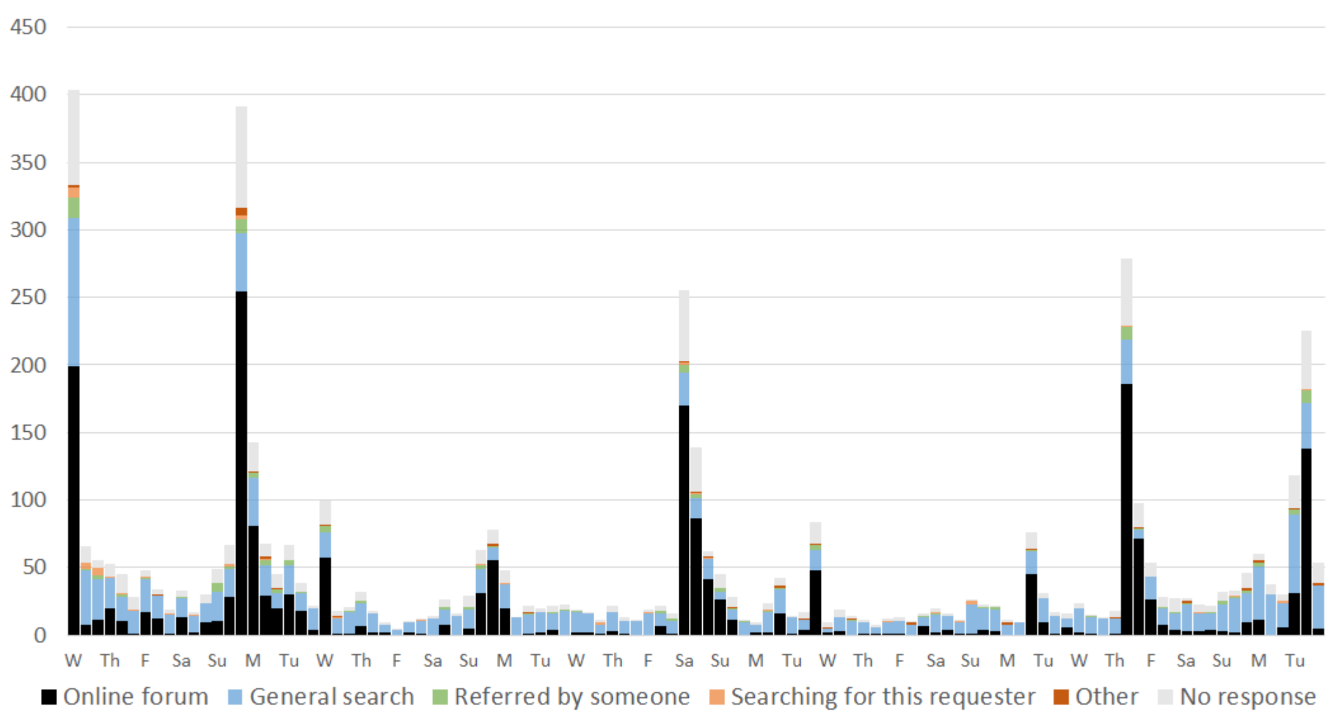

Figure 2. How workers found our mapping HIT (n=4,856) which ran from April 23 - May 28, 2014

workers. After about six months of working on MTurk, Fareed, at Zafar's urging, joined online forums created by other Indian workers. As Fareed noted, "members who are workers share each other's experiences." From the forums, Fareed has made close connections and shared that, "there are also some close ones who, when there is some good work posted they would give a 'miss call'—hanging up before the call connected to save on the costs of phone conversation. We hurry to open the system then and look for the work. I check on phone straight away at times. Anyone who sees work posted calls and tells everyone. There is no fixed timing. Whosoever is alert and sees informs everyone and in this way everyone helps everyone else. Around 150 friends (on Facebook)."

Akbar, an Android phones enthusiast, hoping to migrate to Australia is another Hyderabad native, 19 year-old Muslim and cricket enthusiast. He has worked on MTurk for 2 years. When he and his friends working on MTurk meet at their mosque, "we keep discussing about MTurk for five to ten minutes. 'I worked on this, he worked on that, this was good, this wasn't…and when we get back to work (MTurk) we keep chatting on Skype/Facebook. We talk and even do video chat." Even though they are neighbors within a few kilometers of each others' houses, the young friends work from their own broadband connections and cellular data plans because they worried that sharing the same internet connection while doing their work might lead to them having their accounts suspended. Many crowdwork platform companies do not provide an explicitly technical list of criteria explaining reasons for account suspensions. MTurk and UHRS workers circulate a variety of folk theories about what might prompt the banning of a worker. Systems operators feel the need to keep details opaque to discourage bad actors from gaming the system. However the lack of clarity takes its toll on well-intentioned workers trying to keep their accounts alive.

Akbar talked most with his friend Mohsin because, "we both use Aircel [mobile phone service] and it's free to call between Aircel numbers. So we talk a lot. Among us, I and Mohsin chat the most." Akbar describes another reason that workers connect with each other: to keep each other motivated and awake through the long evenings of shift work that come with participating in an industry driven by U.S. and British time zones. "If you have to work throughout night" Akbar tells us, "you plug in earphones, put the phone to charging and talk all night while working."

As Fareed noted, "we turn to Facebook to ask, 'did any of you work for this requester?' and if a friend says that he worked and had all bulk HITs approved then we also work more for that requester. Requesters' replies don't come immediately…asking friends is easier." Collaborating with friends, both made online or known through offline connections, reduced the costs of spending time finding tasks and reliable requesters. It also helped workers find ways to cut the tedium and challenges of working alone in their homes, across multiple time zones without other physical sources of support we might associate with an office environment.

Figure 2 (previous page, bottom) shows how widespread the phenomenon of referring tasks is. It shows the results of the data gathered from asking workers how they were referred to our mapping HIT. Recall that our mapping HIT ran for 35 days and simply asked users to put a pin on a map wherever they are. Afterwards we asked them how they discovered our HIT. We broke the 35 weeks into consecutive 8 hour periods and laid them out one after another on the x-axis from beginning to end. The y-axis counts how many workers did our HIT each 8 hour period. The bars are colored to indicate how many workers come to the HIT via a forum, via searching the MTurk site or via other methods.

Figure 2 shows the fraction of traffic that came via the forums (the black part of each bar) versus the fraction of traffic that came via searching the MTurk site (the light blue part of each bar). Overall, 41.3% of the traffic came from an online forum whereas only 36.0% of the traffic came from searching the MTurk site. These numbers might depend on the specific attributes of our HIT, but they do suggest that workers collaborating on sharing HITs is widespread.

Figure 2 also shows that the traffic to our HIT was extremely bursty. Most of the 8 hour periods had 50 or fewer workers do our task, however, a handful of 8 hour periods had over 200 workers do our HIT. If we restrict our attention to days with over 100 workers doing our HIT, we see that 55.4% of the traffic came from an online forum as compared to 41.3% of the traffic coming from forums over all days (as mentioned in the previous paragraph). On the other hand, if we look at days with less than 100 workers doing our task we see that only 21.8% of traffic came from the forums.

Since this data is observational it is difficult to make causal claims, but we can say that posting to the forums is correlated with the bursts of traffic shown in Figure 2. We found that most of the major spikes correspond to online posts by searching the archives of the forums for mentions of our HIT. We found posts referring specifically to our HIT on MTurk Grind on 4/23/15, 4/27/15, and 5/13/15, and on Reddit (HITsWorthTurkingFor) on 4/30, 5/4, 5/10, 5/22 and 5/27. There were six posts referring to our HIT to MTurk Forum on 4/23. These correspond to many of the spikes in Figure 2.

Our ethnographic data showed that workers collaborate by referring tasks to each other. Here again workers use their social connections to convey trust, this time to convey trust that a task or requester is legitimate and will pay. Figure 2 shows that collaboration on this information sharing is widespread. Furthermore, it shows that collaborations on the level of individuals can scale up to an overall burstiness of traffic.

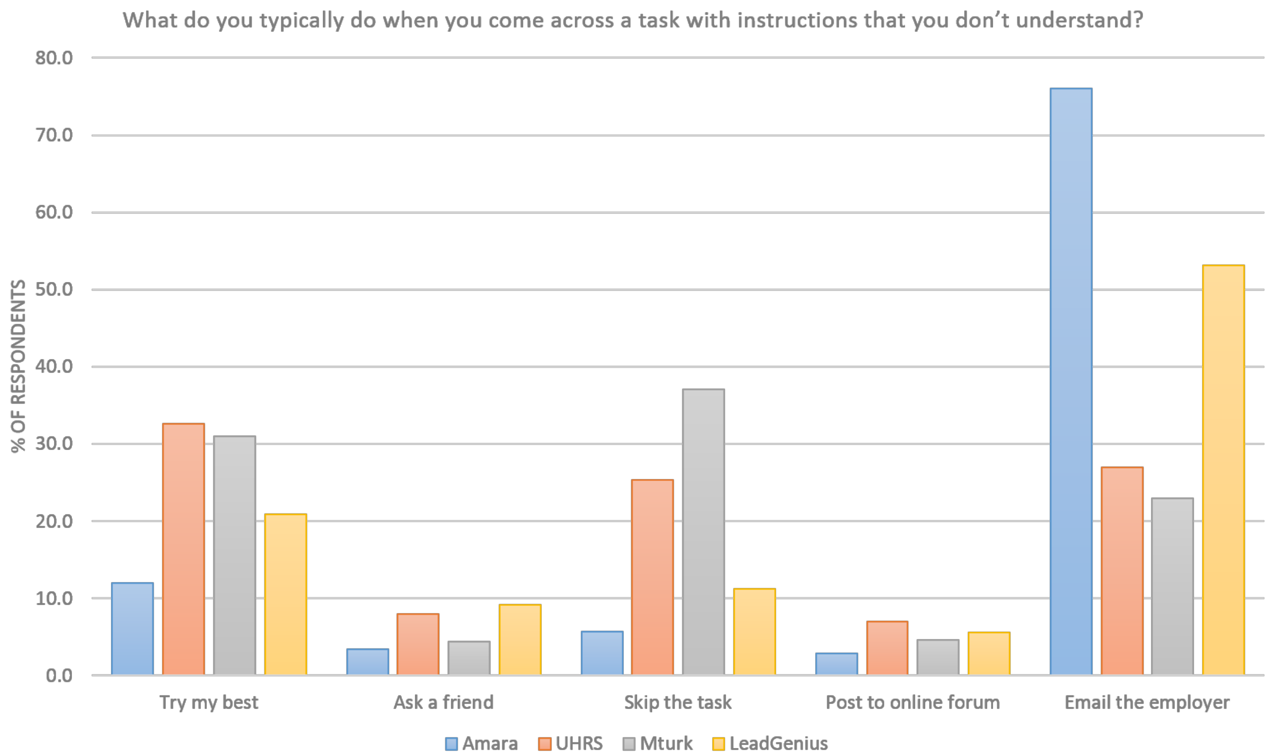

Figure 3. How workers respond to instructions they don't understand across 4 platforms.

Helping Others Complete Tasks and Advance as Crowdworkers

The last type of collaboration that workers discussed in our ethnographic fieldwork and interviews was helping each other complete the actual work itself. They used Facebook and other forums, chat, and in-person meetings to describe how to manage one's time completing tasks and how to do search queries or execute basic scripts or computing techniques, like copying and pasting, to get tasks done.

Poonam, a housewife in her early 20s, and her husband Sanjay, hope to start a family soon. Poonam left her work at a BPO to work from home after signing up for UHRS and finding that it allowed her to forego a tiring commute to the center of Chandigarh and still make enough to money to justify the switch to home-based work. Sanjay noted that his crowdsourcing earnings were important to making ends meet but that he has a hard time managing UHRS work and his own work in his office at a graphic design and print shop. Poonam and Sanjay, somewhat sheepishly, acknowledged that they hand tasks off to each other, if they do not feel confident in their ability to complete them. While Sanjay is more comfortable using spoken English, Poonam excels at the language-based tasks, like search query evaluation. Sanjay takes on any task that involves visualization skills that allow him to take advantage of his advanced knowledge of design.

Anand, a 24 year old student working on MTurk, with a Masters in Embedded Systems, lives in Chennai. Anand uses MTurk for personal expenses not covered by pocket money. His parents do not understand what he is doing with his time though Anand hopes that it will lead to a formal position with Amazon.com. Anand learned about MTurk from his friend Raja and says, "because of Raja all of us have learnt MTurk". Noting the long hours that Raja spent with him, Anand showed us the hand-written computer shortcuts and key commands that Raja gave him. He keeps the tattered note taped to the wall to the left of his desk, showing him the commands for saving screen shots, common search queries, and describing how to download and search excel spreadsheets for the names of U.S. states and cities that Anand could use to answer basic questions that he comes across in tasks.

Lalitha, a Christian mother of two living in Hyderabad described Riyaz as her "guru in MTurk" noting, as many of the people we interviewed did, that crowdwork didn't provide, "someone who could guide one clearly on how to work, how to increase the approval rate and move steadily forward." Kali, a 43 years old, housewife and mother of 2, quit her job as an electrical engineer at a manufacturing plant when she had her children. While her husband now supports her work on UHRS, her in-laws still dislike the amount of time that Kali spends on the computer rather than with them. Before finding UHRS, Kali worked several hours a week, for close to seven years, at an office with four other women in downtown Bangalore doing data entry and database management for a small BPO. While none of the women at Kali's previous workplace followed her to UHRS, Kali still travels to the small office to meet with her former colleagues at least weekly. Kali describes the time in the office as her chance to "feel a part of the working world" and discuss strategies for how to improve search queries—a skill common to all the women in her office. While Kali did not directly collaborate on UHRS to complete tasks with her former coworkers, she did find social connection and support through the shared space of the downtown office environment and opportunities to talk strategy. To complete tasks, she turned to her sons, who routinely help their mother categorize and sort search terms and "adult content" joking, "they are more qualified to recognize these words than me!...I need their help to keep the internet clean and safe for other families." When we asked Kali what improvements she would recommend to crowdsourcing platform designers, she immediately exclaimed "open an office so that I can meet my co-workers!" Kali's enthusiasm for a shared office suggests that the experience of sharing space provides opportunities for collaboration of a different kind: that of finding mutual support and chances to advance one's skills through the kind of "shop talk" that defines most tech environments.

Workers consistently described relying on finding someone—a mentor or group of friends online or off—willing and able to walk a new worker through the disorienting world of survey questions, culturally-specific knowledge like "twerking," and the dizzying range of appliances and other consumer goods relatively unknown to or uncommon among India-based crowdworkers to complete the most mundane tasks posted to the platforms. These collaborations provided help with the task at hand, but also provided social interactions in an online labor setting otherwise devoid of any human contact.

Our survey asked, "What do you typically do when you come across a task with instructions you do not understand?" The responses were that roughly 5% of workers across all platforms "Asked for help on a discussion forum or a blog," and roughly 5% of workers across all platforms, "Asked a friend or relative who crowdsources for help". Thus, roughly 10% of workers tapped their social network, whether it be through a forum or directly, to ask for help. Figure 3 shows these numbers broken down across the platforms that we studied.

Limitations and Qualifications

While our interview and ethnographic data demonstrate specific forms of collaboration among the India-based workers, our survey responses and mapping HIT, which were open to workers from all countries, illustrate similar forms of collaboration. To be sure, some of the India-based workers' collaborations are tied to their limited opportunities to sign up for accounts and additional administrative overhead that comes with being in a developing nation. India-based workers are also less versed in and often explicitly ostracized from U.S. web-based discussion forums. Indeed, the specific geographical and structural barriers of different crowdwork labor pools suggests a pressing need to understand the population of workers solicited for projects, particularly for research that may be affected by sampling bias. But the broader categories of collaboration that we found in our ethnographic data among India-based workers are also discussed in the extant literature in studies of U.S.-specific forums [35,42]. More recently, a 2014 study of TurkerNation, a popular forum dominated by U.S.-based MTurk users, found, through thematic coding of discussion threads and 29 open-ended, semi-structured interviews with forum participants, that U.S. MTurk workers also widely practice the forms of collaboration that we saw among India-based workers (Zyskowski and Milland, not published).

CONCLUSION

We note that our mixed methods approach was essential to grounding our findings. The interviews showed that collaboration happened routinely among workers. Without this finding, we would have never thought to measure this at all. Thus, the ethnographic data gave us hypotheses to then test via surveys and data analysis. While ethnography can unearth unexpected phenomenon, it is difficult to tell how widespread a phenomenon is just using interviews or field observations. We used our surveys and data analysis to fulfill this need. Moreover, our surveys gave us a method to compare our findings across platforms. Thus the strengths of each of our methods could compensate for the weakness of our other methods.

We saw, triangulating our interviews, surveys and ethnographic data analysis, that workers collaborate in three main ways. First, a significant number of MTurk and UHRS workers collaborated on the administrative overhead involved with doing crowdwork, such as creating accounts on the crowdsourcing site, avoiding employment scams, and collecting checks. Second, we saw workers collaborate by notifying each other when high quality tasks are posted or when a trusted requester posts a task. Third, we saw that workers actually help each other by working on tasks together. All three of these modes of collaboration involve workers addressing both social and technical needs. While these collaborations address technical needs they also fairly clearly address social needs as well. They broker, convey and circulate trust for a platform, a requester or a task and generate the social interaction otherwise stripped away by the crowdsourcing API.

With a more expansive definition of collaboration in mind, one that includes building and consulting social networks and mentoring others in different aspects of crowdwork, we argue that workers organically generate forms of collaboration between themselves. Indeed, workers often, seamlessly, interweave these modes of collaboration. Though hard to detect, the articulation work that accompanies collaboration is fundamental to crowdwork [32]. While focusing on the specific interaction patterns of collaborators in work settings is valuable, our interests here are to understand how worker's organic collaboration contribute to the ability of people to complete crowdwork and make it a meaningful experience of employment. We denote the idea that workers self-organize without external help, what we refer to as "organic collaboration", and contrast it to systems were experimenters build workflows for workers to use which we call "engineered collaboration". Furthermore, our research shows that crowds are not collections of independent workers. Instead they are dynamic, self-organizing networks of people. Researchers and engineers alike may need to account for this in their research designs and when framing research questions.

DISCUSSION

Crowdsourcing platforms like MTurk and UHRS remove almost everything from the work process, leaving only the labor, the payment and little else. In fact, many routine uses of crowdsourcing presume humans as independent sources of information and computation. Behavioral experiments, surveys, and polls rely on one person's response with the assumption that they are not sharing that response with others. Also, task creators often de-bias worker output by taking the majority vote of judgments [44], but if the crowds' responses are correlated, the results can be skewed [36]. However, our study shows that workers themselves are putting collaboration back into such systems. Much as the research on the sociology of work and contingent labor suggests, workers deploy a range of strategies to cope with the instability of no longer having the "standard" 40-hour workweek updating [43, 26, 14]. Thus, any reorganization of labor that shifts workflow management, job searching and matching, mentoring, and acknowledgment of labor from formal employers and firms to platforms and APIs will need to consider how, when, and why workers attempt to put particular activities back into the work process. After all, if workers spend so much effort putting, for example, different forms of collaboration back into the work process, they clearly consider these activities important. As such, our paper's findings offer key insights for both crowdwork and the larger ecosystem of platform economies rapidly growing in this global market of on-demand services.

Currently, many platforms do not supply the tools to connect workers, leaving some workers with less English fluency or familiarity with discussion boards unable to fully take advantage of the opportunities to collaborate. In constrast, we found that structural support for collaboration, illustrated by the two counter examples of Amara and LeadGenius, discussed below, appears to lower the costs to workers to connect, coordinate and teach one another. These counter examples suggest that there are substantial benefits to workers in providing the means to interact on the platform itself and fostering other opportunities for articulation work among workers [32].

In many ways, workers on MTurk and UHRS recreated the type of collaboration built into the design of the other two platforms in our larger study. Amara.org, for example, draws translators both volunteering their time to TED's Open Translation Project and serving as a paid, on-demand team for a range of translation and captioning projects [3]. Amara matches translators of similar abilities through smaller captioning exercises, gauging the translation skills level of individuals before putting them in explicitly organized small teams, that include a lead translator. The LeadGenius platform [31] includes built in, real-time chat tools that allow groups to speak directly with other crowdworkers assigned to the same tasks. They can ask each other for help, keep each other company, and reach a junior manager to respond to their questions anytime during their scheduled work shift. Team leaders and junior managers are paid for the time that they spend checking the quality of a crowdworker's tasks as well as for time spent responding a crowdworkers' questions. These management responsibilities are handled and compensated for as discrete tasks.

The use of vendors to recruit and manage labor pools on UHRS may help explain the platform's relatively low levels of collaboration. Beyond an asynchronous internal discussion board associated with each vendor, there are few ways for individuals to see the larger universe of UHRS workers. One upside of the siloing that comes from using specific vending services to manage and curate crowdworkers is that it may make it easier to monitor and track tasks that require individual contributions while taking advantage of the institutional structures folded into vendor systems to build in mentoring, skills-building and a sense of shared identity that may not be possible under current regulatory restrictions for contract labor.

Design Implications

Our findings suggest that crowdsourcing systems cannot sidestep the demands of collaboration and human-driven coordination. APIs can efficiently distribute disaggregated tasks at scale, minimizing the need for hands-on management. Better matching algorithms can reduce the search costs that come with labor markets (for both buyers and sellers), making it easier for workers and employers to execute discrete tasks. These systems cannot, however, eliminate the desire to invest in work as something more than a single payment transaction. Nor can these systems eliminate the very human need for social connection, validation, recognition, and feedback of one's efforts that currently accompany the experience of employment. The widespread and varied reliance on collaboration, from information-sharing to sitting with someone else while doing one's tasks, speak to the value of recognizing the social systems that are part of any employment relationship, even the most contingent or ephemeral forms of employment like paid crowdwork.

Rather than resist the tenacious presence and organic nature of collaboration among crowdworkers, we turn to the collaborative strategies of our research participants to inform several design recommendations. The first recommendation is creating two clearly defined streams of crowdwork: one explicitly available for group collaboration and the other requiring independent, subjective results. There are many tasks, from sales lead generation to location verification that do not, by design, require independent responses. There are cases, however, such as generating training data for machine learning algorithms or survey responses where independent results are required for validity. Where collaboration works against the desired outcomes, we could focus on explicit directions banning collaboration, building in more instructional transparency where the task quality depends on it. We could also focus our efforts to ferret out breaches of the terms of participation where a workers' desire to collaborate really hurts the final outcomes.

In the short run, we could also develop systems of "task-ifying" management that turn affirmation and encouragement into doable, paid tasks. Several respondents suggested meet ups and issuing certificates of achievement that would validate contributions to the highest quality work. We currently associate managing or curating a workforce, through recognition and coordination of collaborative teamwork, with full-time employment. To make such recognition and validation possible, let alone training and other forms of formalized mentorship, we will need to redefine "independent/freelance" workers to better incorporate and value managing and curating digital workforces. While autonomous workers are able to recreate the systems of collaboration needed to find and complete tasks and build the social bonds that make work manageable, the next iteration of crowdwork may be best served by returning to the challenges of mediating collaboration. Work, after all, remains a sociotechnical system that requires as much attention to the cultural needs and values we attach to our labor as the technical tools to do our jobs efficiently.

Acknowledgments
The authors would like to thank the crowdworksers and platform engineers who have generously contributed their insights to this project; members of the research team who have supported different facets of the research project, specifically Andrea Alarcón, Wei-Chu Chen, Rebecca Hoffman, Sara Kingsley, Kate Miltner, and Greg Minton; and anonymous reviewers for CSCW.

REFERENCES

1. Andrew Abbott. 2005. The sociology of work and occupations. In Handbook on Economic Sociology, Neil J. Smelser and Richard Swedberg (eds.). Russell Sage Foundation and Princeton University Press, New York and Princedon, 307-330.
2. Andrew Abbott. 1993. The sociology of work and occupations. Annu Rev Sociol 19: 187-209.
3. Amara – Caption, Translate, Subtitle and Transcribe Video. 2015. Retrieved May 21, 2015 from http://www.amara.org
4. Vashi Ambati, Stephan Vogel, Jamie Carbonell. 2012. Collaborative workflow for crowdsourcing translation. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '12), 1191-1194.
5. Morgan G. Ames, Janet Go, Joseph Kaye, Mirjana Spasojevic. 2010. Making love in the network closet: The benefits and work of family videochat. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '10), 145-154.
6. Benjamin B. Bederson, Alexander J. Quinn. 2011. Web workers unite! Addressing challenges of online laborers. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11), 97-105.
7. Michael S. Bernstein, Desney Tan, Greg Smith, Mary Czerwinski, Eric Horvitz. 2010. Personalization via friendsourcing. ACM Trans Comput-Hum Interact 17, 2: 1-28.
8. Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, Tom Yeh. 2010. Nearly Real-time answers to visual questions. In Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '10), 1-10.
9. Daren C. Brabham. 2014. Moving the crowd at Threadless. Information, Communication, & Society 13 (8): 1122-1145.
10. Jon Chamberlain. 2014. Harnessing the intelligence of the crowd for problem solving and knowledge discovery In Proceedings on the AAAI Conference on Human Computation and Crowdsourcing (HCOMP '14), 1-5.
11. Lydia B. Chilton, John J. Horton, Robert C. Miller, Shiri Azenkot. 2010. Task search in a human computation market. In Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KIGKDD '10), 1-9.
12. Lydia B. Chilton, Juho Kim, Paul André, Felicia Cordeiro, James A. Landay, Daniel S. Weld, Steven P. Dow, Robert C. Miller, Haoqi Zhang. 2014. Frenzy: collaborative data organization for creating conference sessions. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14), 1-10.
13. Franco Curmi, Maria Angela Ferrario, John Whittle, Floyd Mueller. 2015. Crowdsourcing synchronous spectator support: (go on, go on, you're the best). In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15), 757-766.
14. Alison Davis-Blake, Joseph P. Broschak. 2009. Outsourcing and the changing nature of work. Annu Rev Socio 35: 321-340.
15. Paul Dourish. 1995. Developing a reflective model of collaborative systems. ACM Trans Comput-Hum Interact 2, 1: 40-63.
16. Elizabeth Gerber, Julie Hui. 2013. . Crowdfunding: Motivations and deterrents for participation. ACM Trans Comput-Hum Interact 20, 6: 1-32.

17. Elizabeth Gerber, Michael Muller, Rick Wash, Elizabeth F. Churchill, Lily C. Arani, Amanda Williams. 2014. Crowdfunding: An emerging field of research. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14), 1093-1098.
18. Jonathan Grudin. 1994. Computer-supported cooperative work: History and focus. Comput 27, 5: 19-26.
19. Neha Gupta, David Martin, Ben Hanrahan, Jacki O'Neill. 2014. Turk-life in India. In Proceedings on the ACM International Conference on Supporting Group Work (GROUP '14), 1-11.
20. Vaughn Hester, Aaron Shaw, Lucas Biewald. 2010. Scalable crisis relief: Crowdsourced SMS translation and categorization with Mission 4636. In Proceedings of the Symposium on Computing for Development (DEV '10), 1-7.
21. James Hollan, Edwin Hutchins, David Kirsh. Distributed cognition: toward a new foundation for human-computer interaction research. ACM Trans Comput-Hum Interact 7, 2: 174-196.
22. Jeff Howe. 2006. The rise of crowdsourcing. Wired Magazine 14: 1-4.
23. Shih-Wen Huang, Wan-Tat Fu. 2013. Don't hide in the crowd! Increasing social transparency between peer workers improves crowdsourcing outcomes. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13), 621-630.
24. Shih-Wen Huang, Wan-Tat Fu. 2013. Motivating crowds using social facilitation and social transparency. In CSCW '13 Companion, 149-152.
25. Lily Irani, M. Six Silberman. 2013. Turkopticon: Interrupting worker invisibility in Amazon Mechanical Turk. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13), 1-10.
26. Arne L. Kalleberg. Nonstandard employment relations: Part-time, temporary and contract work. Annu Rev Sociol 26: 341-365.
27. Aniket Kittur, Ed Chi, Bryan A. Pendleton, Bongwon Suh, Todd Mytkowicz. 2009. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09), 1-9.
28. Aniket Kittur, Robert E. Kraut. Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '08), 37-46.
29. Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, John Horton. 2013. The future of crowd work. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '13), 1301-1317.
30. Anand Kulkarni, Matthew Can, Björn Hartmann. 2012. Collaboratively crowdsourcing workflows with Turkomatic. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '12), 1003-1012.
31. LeadGenius – B2B Lead Generation and Custom Data. 2015. Retrieved May 21, 2015 from http://www.leadgenius.com
32. Charlotte P. Lee, Drew Paine. 2015. From The Matrix to a Model of Coordinated Action (MoCA): A Conceptual Framework of and for CSCW. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '15), 179-194.
33. Greg Little, Lydia B. Chilton, Robert C. Miller, Max Goldman. 2009. TurKit: Tools for iteratives on Mechanical Turk. In Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP '09), 29-30.
34. Munir Mandviwalla, Lorne Olfman. 1994. What do groups need? A proposed set of generic groupware requirements. ACM Trans Comput-Hum Interact 1, 3: 245-268.
35. David Martin, Benjamin V. Hanrahan, Jacki O'Neill, Neha Gupta. 2014. Being a turker. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '14), 1-12.
36. Lev Muchnik, Sinan Aral, Sean J. Taylor. 2013. Social influence bias: A randomized experiment. Science 341, 6146: 647-651.
37. Jacki O'Neill, David Martin. 2013. Relationship-based business process crowdsourcing. In Proceedings of the IFIP conference on human-computer interaction (INTERACT '13), 429-446.
38. Cheong Ha Park, KyongHee Son, Joon Hyub Lee, Seok-Hyung Bae. 2013. Crowd vs. crowd: Large-scale cooperative design through open team competition. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '13).
39. Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, Clickhappier. 2015. We are dynamo: Overcoming stalling and friction in collective action for crowd workers. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15), 1-10.
40. Aaron Shaw, Benjamin Mako Hill. 2014. Laboratories of oligarchy? How the iron law extends to peer production. J Commun 64, 2: 214-238.
41. Aaron Shaw, Haoqi Zhang, Andrés Monroy-Hernández, Sean Munson, Benjamin Hill, Elizabeth Gerber, Peter Kinnaird, Patrick Minder. 2014.

Computer supported collective action. ACM Interactions 24, 2: 74-77.

42. M. Six Silberman, Joel Ross, Lily Irani, Bill Tomlinson. 2010. Sellers' problems in human computation markets. In Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP '10), 18-21.
43. Vicki Smith. 1997. New forms of work organization. Annu Rev Sociol 23: 315-339.
44. Rion Snow, Brendan O'Connor, Daniel Jurafsky, Andrew Y. Ng. 2008. Cheap and fast—but is it good? Evaluating non-expert annotations for natural language tasks. In Proceedings of the ACL Conference on Empirical Methods in Natural Language Processing (EMNLP '08), 254-263.
45. Anslem Strauss. 1988. The articulation of project work: An organizational process. Sociological Quarterly 29, 2: 163-178.
46. James Surowiecki. 2005. The wisdom of crowds. Random House.
47. Haoqi Zhang, Eric Horvitz, Rob C. Miller, and David C. Parkes. 2011. Crowdsourcing general computation. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11), 1-4.
48. Haiyi Zhu, Robert Kraut, Aniket Kittur. 2012. Group identification, goal setting, and social modeling in directing online production. In Proceedings of the ACM conference on Computer supported cooperative work (CSCW '12), 935-944.
49. Vinko Zlatić, Miran Božičević, Hrvoje Štefančić, Mladen Domazet. 2006. Wikipedias: Collaborative web-based encyclopedias as complex networks. Phys Rev E 74, 1.