

The Communication Network Within the Crowd

Ming Yin^{*}
Harvard University
mingyin@fas.harvard.edu

Siddharth Suri
Microsoft Research
suri@microsoft.com

Mary L. Gray
Microsoft Research
mlg@microsoft.com

Jennifer Wortman Vaughan
Microsoft Research
jenn@microsoft.com

ABSTRACT

Since its inception, crowdsourcing has been considered a black-box approach to solicit labor from a crowd of workers. Furthermore, the “crowd” has been viewed as a group of independent workers dispersed all over the world. Recent studies based on in-person interviews have opened up the black box and shown that the crowd is not a collection of independent workers, but instead that workers communicate and collaborate with each other. Put another way, prior work has shown the existence of edges between workers. We build on and extend this discovery by mapping the entire communication network of workers on Amazon Mechanical Turk, a leading crowdsourcing platform. We execute a task in which over 10,000 workers from across the globe self-report their communication links to other workers, thereby mapping the communication network among workers. Our results suggest that while a large percentage of workers indeed appear to be independent, there is a rich network topology over the rest of the population. That is, there is a substantial communication network *within* the crowd. We further examine how online forum usage relates to network topology, how workers communicate with each other via this network, how workers’ experience levels relate to their network positions, and how U.S. workers differ from international workers in their network characteristics. We conclude by discussing the implications of our findings for requesters, workers, and platform providers like Amazon.

Keywords

crowdsourcing, networks, online forums, Mechanical Turk

1. INTRODUCTION

Crowdsourcing is the act of outsourcing a job to an undefined group of people through an open call [11]. On a typical crowdsourcing platform like Amazon Mechanical Turk¹, re-

^{*}Most of this work done while visiting Microsoft Research.

¹<https://www.mturk.com/>

questers use an API provided by the platform to post tasks for workers to complete. Workers can then browse available tasks and choose which tasks to perform in exchange for prespecified payments. The platform’s API defines the communication protocol between a requester and the workers. By design, it hides personal attributes of workers, such as age, gender, and ethnicity, from the requester. Similarly, the API hides *social* characteristics of workers such as how many friends they have who also do crowdwork or if they are currently working on a task with other workers. In fact, Mechanical Turk does not ask workers about their personal characteristics and does not have access to social characteristics or any details of worker communication that takes place off the platform. Without this information, requesters may come to view workers as simply a black box method to accomplish tasks. In light of this, it is not surprising that crowds are often seen as groups of independent workers, with little attention paid to the connections between them.

This notion of crowds as independent workers was recently dispelled by Gray et al. [6], who opened up the black box and showed that workers are *not* independent but rather connected through social ties. Through a mix of ethnographic fieldwork, in-person interviews, surveys, and large scale data analyses of four different crowdsourcing platforms, they showed that workers collaborate with one another to meet social and technical needs left wanting by the crowdsourcing platforms studied. More specifically, they showed that workers collaborate on three fronts: 1) helping each other get through the administrative overhead involved in doing crowdwork (e.g., signing up for an account and getting paid, which can be especially challenging outside of the United States), 2) sharing information about lucrative tasks and reputable (or irreputable) requesters, and 3) completing work together. Thus, Gray et al. [6] showed the crowd is not a collection of independent workers, but that there exist edges between the workers.

While prior work showed that communication exists, it left open the problem of understanding the scale and structure of this communication. It showed that workers talk, but did not shed light on how widespread communication is, the topology of the communication network, or how worker attributes like geographic location correlate with network position. The central contribution of our paper is to map the entire communication network of workers on a leading crowdsourcing platform, Amazon Mechanical Turk (MTurk). We aim to understand the network’s properties and the implications that communication across the network has on the field of crowdsourcing. To do so, we designed a task that

encouraged workers to self-report their connections to other workers in a privacy-preserving way. The task was designed to provide value back to workers by allowing them to explore the network and learn about the workers they connect to and the greater network of crowdworkers. The edges that workers provide are self-reported and thus not perfectly accurate. However, they give us a close approximation of the true communication network underlying MTurk, and a sense of how widespread communication among workers is.

We analyze the structural features of the MTurk communication network. While a large segment of the population does, in fact, appear to be made up of isolated nodes, we show that there is a rich network topology over the subset of workers who report connections. That is, there is a substantial network *within* the crowd.

We show that online forums dedicated to working on MTurk play a key role in allowing workers to communicate across the network. Forums create overlapping subcommunities among workers. Forums differ from each other in terms of the topological structure of their subcommunities, the temporal nature of communication, and the content of discussions. Meanwhile, one-on-one channels are also used by some workers to communicate, yet they play a different role in fostering communication when compared to online forums. We also observe various types of homophily between workers. That is, we observe that workers are more likely to communicate with other workers who live in the same country, have worked on MTurk for a similar amount of time, and prefer the same types of MTurk tasks (e.g., classification or scientific experiments). By correlating topological features of the network with a number of worker properties, we find that workers' positions in the network are related to various aspects of their MTurk experiences, such as how long they have stayed on MTurk, whether they make use of online forums, how successful they are as MTurk workers, and how fast they can find interesting tasks on MTurk. Finally, as a case study of how workers with different properties participate in the network differently, we provide a comparison between workers who live in and out of the United States and show that these two populations hold different positions in the network, adopt different channels for communication, and focus on different topics in their communication.

Our findings have practical implications for all parties in crowdsourcing including requesters, workers, and platform providers. We discuss these in the conclusion.

2. RELATED WORK

The results of Gray et al. [6] are based on data gathered by a team of ethnographers who spent roughly 19 months in India interviewing over 100 crowdworkers, conducting repeat interviews with many of them over time to understand the longitudinal effects of crowdwork². Gray et al. [6] augmented their interviews with large scale surveys of the crowdworker population in both the U.S. and India and an analysis of a HIT designed to understand where MTurk workers are located and what resources they use to find HITs. Their key finding is that workers collaborate with each other, often to make up for technical or social shortcomings in the platform. The notion that some workers talk and collaborate with one another is also supported by the 35

²In ongoing work the authors have also conducted over 20 interviews with workers in the United States.

interviews of Indian workers that Gupta et al. [7] conducted, mostly via Skype. Both studies indicate that workers collaborate to share tasks, aid each other in doing tasks, and provide social interaction that is often missing in online labor. This notion inspired our goal of mapping the worker network. Our contribution above and beyond these studies is to scale up their findings and dig deeper into the structure of communication. While they find communication between 35 to over 100 interview subjects, we measure and analyze the communication network of over 10,000 MTurk workers.

One theme that appears prominently in Section 4 is the importance of online forums to the structure of the communication network. Prior research has shown the importance of these forums in the work and lives of MTurk workers. Martin et al. [10] spent hundreds of hours reading posts on TurkerNation, a popular online forum for MTurk workers, to understand this online community. They showed that workers primarily work on MTurk to augment their pay and that workers spend a lot of time talking about requesters and tasks in search of requesters with good reputations and tasks with high pay. Similarly, Zyskowski and Milland [19] conducted an ethnographic study of TurkerNation. They observed participants in chat rooms and interviewed them. They state that on TurkerNation, "common topics of discussion include the best jobs of the day, how to build one's reputation, how to earn more money, and how to make working more fun." Thus workers are using forums not just to find lucrative tasks but also to provide each other with social support. These qualitative studies inform our work. Our goal is to scale these studies up and see how big the communication network between MTurk workers is, what topology it has, and how workers use it.

Researchers have built at least two platforms that facilitate worker communication. First, TurkOpticon is a system developed by Irani and Silberman [8] used by workers to rate requesters in terms of their communicativity, fairness, generosity, and promptness. Second, Dynamo [15] is a community platform designed to aid MTurk workers with collective action problems such as "reining in problematic academic research practices" and gathering support for a letter-writing campaign. These works facilitate worker communication for focused goals. The purpose of our work is different in that we seek to understand the structure and scale of the overall communication network that has organically grown among the workers themselves.

3. EXPERIMENTAL DESIGN

Amazon Mechanical Turk (MTurk) is an online labor market in which requesters can post small jobs, referred to as *human intelligence tasks* or *HITs*, along with specified payments for completing each HIT. A typical HIT might involve translating a paragraph of text, labeling an image, or completing a survey. Workers can browse available HITs and choose HITs to work on. Once a worker has submitted her work for a given HIT, the HIT's requester can review this work, accepting it if it is high quality and rejecting it if not. If work is rejected, the worker receives no payment. The rejection is also reflected in the worker's *approval rate*, which is simply the fraction of HITs the worker has done that have been accepted. The approval rate serves as part of a *de facto* reputation system, and requesters often make HITs available only to workers with a high approval rate. Amazon additionally designates some workers as *Masters*.

Forum	URL	Registered Users	Posts	Start Date
Reddit HWTF	https://www.reddit.com/r/HITsWorthTurkingFor/	32,297	unknown	February, 2012
MTurkGrind	http://www.mturkgrind.com/	6,743	748,983	October, 2013
TurkerNation	http://turkernation.com/	15,411	311,816	August, 2011
MTurkForum	http://www.mturkforum.com/	53,883 (932 active)	1,354,249	January, 2009
CloudMeBaby	http://www.cloudmebaby.com/	4,180	32,072	July, 2012
Facebook (groups)	http://facebook.com/	unknown	unknown	unknown

Table 1: Statistics as of October 4, 2015 on the six online MTurk forums listed as options for the question on forums.

While Amazon does not disclose the criteria used to grant the Masters qualification, it is viewed as a sign of high quality, and requesters may choose to make HITs available only to workers who have received this qualification.

Amazon does not provide a platform for workers to interact with each other. However, MTurk workers have created a variety of forums focused on navigating MTurk. A brief overview of the most popular forums is given in Table 1. These forums differ somewhat in functionality. Reddit’s HITsWorthTurkingFor (HWTF) is a highly active subreddit primarily used by workers to share links to, and information on, good HITs. MTurkGrind, TurkerNation, and MTurkForum are post-driven discussion boards organized around a range of themes, much like USENET newsgroups. They each offer moderated areas and distinct but comparable conversation modes organized by discussion threads. Each of these forums has tens of thousands of threads dedicated to a wide range of topics. Registered members of these forums can participate in the discussion in any threads they are interested in, and they may also interact with each other in chat rooms or through private messaging systems provided by the forums. CloudMeBaby is a site devoted to helping navigate and improve cloud based workplaces including MTurk. In addition to these public forums, there are a number of both private and public MTurk-related Facebook groups, varying in size from tens of users to thousands of users.

Since the communication network among workers is not accessible from the API provided by MTurk—in fact, the network exists outside of and separate from the MTurk platform—we cannot simply download, crawl, or scrape this network. In this section, we describe a HIT that we designed to give workers incentive to self-report their connections with each other in the communication network underlying MTurk. We believe the approach described here is novel and could be of independent interest.

3.1 The Network Mapping HIT

We designed a five-step HIT to gather information from workers and allow them to self-report other workers with whom they communicate. In the first step of our HIT, each worker was asked to create a unique nickname for herself. This nickname had several intended purposes. First, it was used as a unique identifier for the worker, preserving the worker’s privacy since it was not based on the worker’s MTurk ID or other identifying information. (Workers were encouraged not to use their real name, though we had no way to enforce this.) Additionally, it was used as a way for other workers with whom this worker communicates to add edges to this worker and identify this worker in the network. This is described in more detail below.

In the second step, workers were asked nine survey questions about their demographics and MTurk usage:

- **Location:** Which country do you currently live in?
- **Age:** Which year were you born in?

- **Gender:** What is your gender?
- **Education:** What is the highest degree or level of school you have completed?
- **Master:** Are you a Mechanical Turk Master?
- **Approval Rate:** What’s your approval rate on Mechanical Turk?
- **Experience:** How long have you been Turking?
- **Tasks:** What types of MTurk tasks do you typically do?
- **Forums:** What online MTurk forums do you regularly use?

For the question on tasks, we provided a list of eleven common types of MTurk tasks such as data entry, survey, and scientific experiments, and allowed workers to choose one or more. For the question on forums, we enumerated the six popular MTurk forums in Table 1. Workers could choose any number of these forums, specify other forums they use, or say that they do not use any forums.

We allowed workers to set privacy preferences individually for each of the nine questions. For each question, a worker could choose whether to share her answer with all other workers who completed our HIT, to share her answer with only those workers connected to her in the communication network, or to keep her answer private.

In the third step, workers were asked to answer two free-form questions related to their experience on MTurk:

- Why did you start Turking?
- What motivates you to keep Turking?

These questions were carefully chosen to obtain information that other workers would find valuable and interesting as a way of providing value back to workers who completed our HIT. We ran a pilot survey in which we asked workers what they would most like to know about other workers in the MTurk community and extracted the most popular questions. Our hope was that presenting information that workers found valuable would encourage workers to explore the communication network through the visualization and to truthfully report their connections. Workers were informed that their answers to these two questions would be shared with all other workers who completed our HIT as part of the network visualization which they would view in Step 5.

In Step 4, each worker was asked to pause and take a moment to exchange nicknames with other MTurk workers that she knows. Workers were told that they could do this in any way they wanted and given several examples including exchanging nicknames in person, over the phone, through instant messaging, or through text messaging.

In the final step, workers were shown a visualization of the current state of the communication network. Each worker in the network was represented by a node displaying the national flag of her country (if her privacy settings allowed), and edges were shown between pairs of connected workers. The worker was able to locate herself, zoom in and out of the network, and click on any worker in the network to view

his information. When a worker clicked on another worker to whom she was connected, she would see his nickname as well as all information that he had chosen to share with his connections. Crucially, when a worker clicked on a worker to whom she was *not* connected, she would *not* see his nickname and would see only information he chose to share with all workers. Thus such workers were effectively anonymous.

At this point, workers could add an edge to any other worker by providing his nickname. When adding an edge, the worker was asked an additional two questions:

- How do you usually talk to this worker?
- What do you usually talk about with this worker?

For the first question, the worker was provided with a list of communication channels such as forums, phone calls, email, and instant messaging, and allowed to choose one or more. For the second, the worker was given a list of topics such as sharing HITs, discussing requesters, sharing Turking tools/scripts, and chatting about day-to-day life, and could choose one or more. After entering this information, an undirected edge between the two workers was immediately added to the network. Workers were also able to remove edges to other workers.

Before submitting the HIT, the worker was given a unique URL that would allow her to return to the visualization to add or remove additional edges and continue to explore.

Note that by design, an edge between two workers could only be added if one of the workers knew the other's nickname, which could only occur if the workers had communicated³. Thus we believe that the vast majority of the edges in the network represent a true exchange of information, or in other words, a communication between workers. Of course there are likely pairs of workers who communicate but did not choose to exchange nicknames. However, exchanging nicknames allowed workers to learn interesting information about each other and better understand their own place in the MTurk community. We believe this design nudged workers towards reporting many of their true connections, though the true communication network is perhaps even more dense and vast than we show here.

We cannot rule out the possibility that the very existence of our HIT caused communication between pairs of workers who had not previously communicated with each other. This is unavoidable; in general, every new HIT has the potential to provoke new communication and the communication network is always evolving. We attempted to minimize this effect by intentionally deciding not to pay workers per edge added, as this would result in workers adding edges to those they do not regularly communicate with.

3.2 Experimental Procedure

We posted our HIT to MTurk. Workers who accepted the HIT read through a description of the task, signed a consent form stating that they were voluntarily participating in our experiment, and then completed the HIT as described above. The payment for the HIT was fixed at \$1 USD and the average completion time was roughly 10 minutes. The HIT was open to all workers on MTurk. Each worker was allowed to complete the HIT only once, but could return to their personalized URL to further explore the network and

³A worker could potentially guess another worker's nickname, but we do not believe this frequently occurred. If it did, the second worker could remove the unwanted edge.

add or delete edges as often as they liked. Our experiment was approved by the Microsoft Research IRB.

To ensure our HIT was well-functioning and scalable, we intentionally launched our experiment in phases during August and September of 2015. We first launched two small batches on August 11 (60 HITs) and August 12 (200 HITs). We notified workers on TurkerNation ahead of time about these two test launches. Next, to test the scalability we launched two larger batches on August 17 (596 HITs) and August 20-21 (1594 HITs). Satisfied with these initial tests, we finally left our HIT up for 2 weeks straight from August 28 to September 11, with the exception of 2 days (September 3-4) during which our requester account accidentally ran out of money due to the unexpected popularity of our HIT. After our HIT was taken down, workers continued to update the network via their private URLs. We report on data collected on September 13 once the addition and removal of edges had greatly slowed.

4. RESULTS

A total of 10,354 workers completed our HIT. Stewart et al. [16] estimated that when conducting behavioral research on MTurk, one laboratory is sampling from a pool of roughly 7,300 workers, and that the seven laboratories they studied sampled from an overall pool of roughly 11,800 workers. This suggests that our HIT was approximately a census of the active workers at the time. Of the workers who did our HIT, 1,389 (13.4%) either added at least one edge to another worker or had an edge added to them by another worker. We refer to these workers as *connected*. Among connected workers, a total of 5,268 edges were added, resulting in a mean degree of 7.6, median degree of 2, and maximum degree of 321. The largest connected component of the communication network consisted of 994 workers (71.6% of all connected workers), while the next largest consisted of just 49 workers (3.5% of all connected workers). Of the remaining connected components 117 were made up of a single edge between a pair of workers.

The communication network is shown in Figure 1a. Examining the network visually, it appears that the largest connected component is made up of several densely connected clusters of workers. Below we show that this structure largely coincides with workers' use of different online forums dedicated to Mechanical Turk work.

4.1 A Network Enabled by Forums

Forum use is extremely wide-spread among workers who completed our HIT, with 59.1% of all workers and 83.0% of connected workers reporting that they use at least one forum⁴. The overwhelming majority of edges involved communication through a forum as 89.9% of the edges added were between pairs of workers that communicate via forums, and 86.2% between pairs that communicate *exclusively* through forums. Since the vast majority of communication between

⁴The Forums question was added to Step 2 of our HIT on August 20 when we first realized the prevalence of forum usage. We asked the 856 workers who completed our mapping HIT before August 20 which forums they regularly use in a separate, one-question follow-up HIT and 659 responded. As a result, 98.1% of all workers answered the question. Whenever we report statistics related to forum usage, we restrict attention to workers who answered the Forums question.

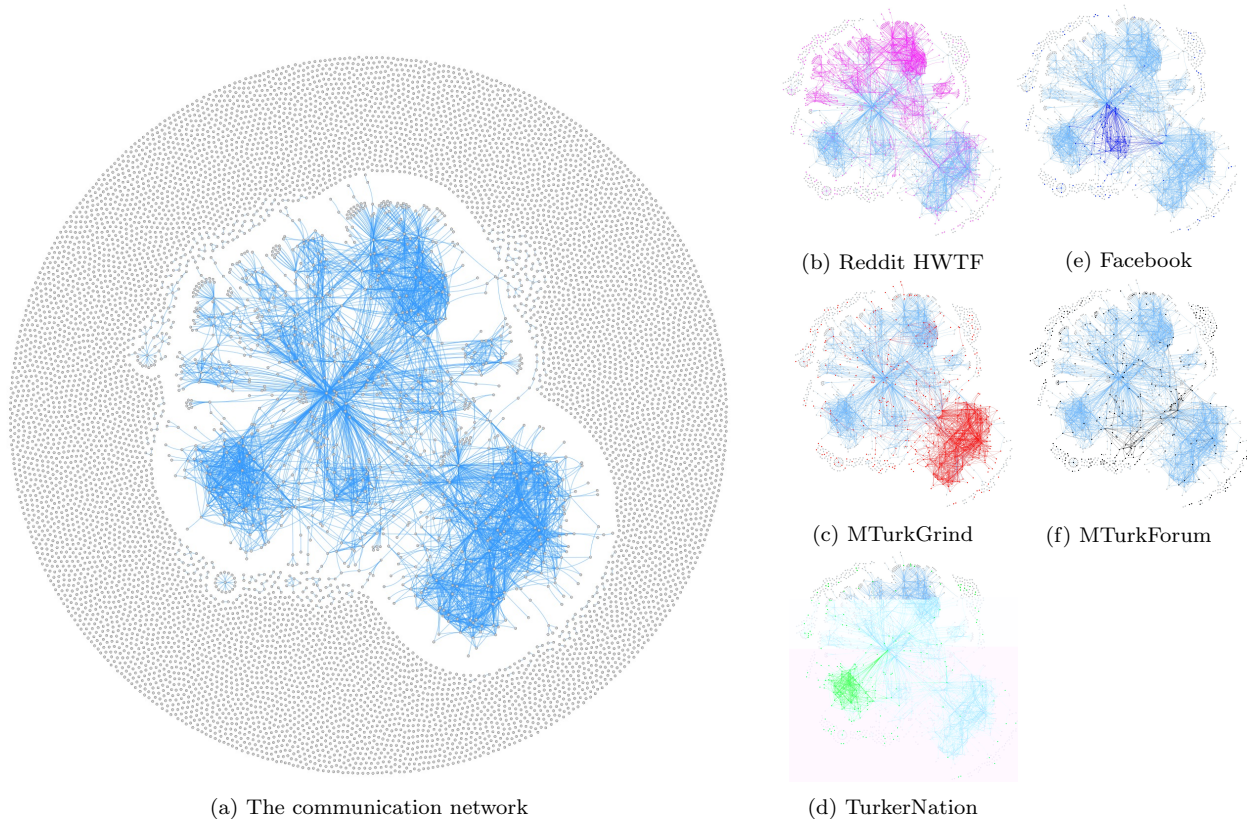


Figure 1: 1a: The communication network among Amazon Mechanical Turk workers. 1b-1f: Subnetworks for Reddit HWTF (magenta; 660 workers, 1837 edges), MTurkGrind (red; 392 workers, 1331 edges), TurkerNation (green; 200 workers, 740 edges), Facebook (blue; 133 workers, 357 edges), and MTurkForum (black; 312 workers, 244 edges).

workers occurs on online forums, we next analyze the structure of the subnetworks defined by each of the forums.

We extract the subnetwork corresponding to each forum by keeping only *connected* workers who use that forum and only edges between pairs of these workers. As a sanity check, pairs of workers reported communicating with each other through forums for the vast majority of these edges (93% averaged over all subnetworks). Figures 1b-1f illustrate the subnetworks for Reddit HWTF, MTurkGrind, TurkerNation, Facebook, and MTurkForum, respectively. We omit CloudMeBaby as only 0.9% of all workers reported using it. As is visually apparent from the figures, users of different forums make up distinct but overlapping subcommunities, which explains much of the structure in the network.

To quantify our visual intuition, we measure whether or not workers who use the same forum are more likely to connect to each other than to other workers. The sociological phenomenon that contact between similar people occurs at a higher rate than among dissimilar people is called *homophily* [12]. Thus, we are interested in understanding the extent to which homophily exists with respect to forum use.

One standard approach to quantifying homophily is the *homophily test* described in Easley and Kleinberg [5]. Consider a binary property \mathcal{C} that a node may or may not satisfy. In our case, satisfying \mathcal{C} might mean using a particular forum like MTurkGrind. Let q denote the fraction of the population who satisfy \mathcal{C} , S denote the set of all nodes that satisfy \mathcal{C} , and T denote the set of all nodes that do not. If there is no homophily with respect to \mathcal{C} , edges would be equally likely to form between all pairs of nodes in the network in-

dependent of whether those nodes satisfy \mathcal{C} . So, in the case of no homophily each node on an edge would satisfy \mathcal{C} independently with probability q , and the probability that any edge would be between one node in S and one node in T would be $2q(1 - q)$. We refer to this quantity as the *expected cross-group ratio (ECGR)* of \mathcal{C} . If, on the other hand, nodes in S were more likely to connect to other nodes in S , and nodes in T to other nodes in T , then the actual fraction of edges that would be between nodes in S and T , or the *actual cross-group ratio (ACGR)*, would be significantly lower. The homophily test compares these ratios.

Table 2 (left section) reports the results of homophily tests run separately for each forum⁵, limited only to connected workers. For each of the five forums, we find that the actual cross-group ratio is lower than the expected cross-group ratio. This provides evidence for homophily with respect to the use of each forum, confirming the visual intuition given by Figure 1. To check whether the differences are statistically significant, we keep the network structure fixed and simulate a random assignment of node property values (that is, whether or not a node uses a particular forum) by assigning each node to use the forum with probability equal to the fraction q of users who use the forum in the real worker population. We repeat this process 1,000 times, calculating the cross-group ratio for each of the 1,000 resulting networks. An empirical p-value can then be computed as the fraction

⁵Note that while the test of Easley and Kleinberg [5] easily extends beyond binary properties, we must run it separately for each forum since workers may select multiple forums.

Forum Name	ECGR	ACGR	p-val.	User			Density (d)	Transitivity (t)	Diam.	Avg. Shortest Distance
				Frac. (q)	R	H				
Reddit HWTF	0.50	0.30	<0.001	0.48	0.69	0.70	0.008	0.30	9	8.36
MTurkGrind	0.41	0.23	<0.001	0.28	0.53	0.65	0.017	0.38	13	7.15
TurkerNation	0.25	0.18	0.005	0.14	0.56	0.62	0.037	0.48	5	4.55
Facebook	0.18	0.17	0.362	0.10	0.53	0.45	0.041	0.38	6	4.37
MTurkForum	0.36	0.25	<0.001	0.23	0.39	0.28	0.005	0.11	10	7.85

Table 2: Left section: Expected cross-group ratio and actual cross-group ratio for the usage of each forum. Middle section: One-sided homophily measures for each forum. Right section: Density, transitivity and distance metrics for each subcommunity.

of these simulated networks with a cross-group ratio smaller than the ACGR we measure. As reported in Table 2 (left section), the differences are significant for almost all forums.

The results of the homophily tests may, in fact, underestimate the amount of homophily in the network. This is because while we might expect workers who use Facebook forums, for example, to be more likely to connect with other workers who use Facebook forums, it is unclear if workers who do not use Facebook forums are much more likely to connect with other workers who do not. To address this, we look at two alternative measures of such “one-sided” homophily. For a given node i , let n_i be the total number of edges incident on i , and $n_{i,S}$ be the number of edges incident on i that connect to nodes in S . Intuitively, it is a sign of homophily if, on average, the fraction of the edges that are incident on some node in S that to connect to other nodes in S is higher than the fraction of nodes in the total population that are in S , i.e., if $R \equiv (1/|S|) \sum_{i \in S} (n_{i,S}/n_i) > q$. The measure R treats all nodes equally. The homophily index of Currarini et al. [4], defined as $H \equiv \sum_{i \in S} n_{i,S} / \sum_{i \in S} n_i$, is similar but effectively gives more weight to nodes with higher degree. Again, if $H > q$, there is evidence of homophily.

Table 2 (middle section) shows both R and H for each forum along with the fraction of workers who reported using that forum, again limited to connected workers. As expected, these measures show a clear and striking tendency for workers to connect to other workers who use the same forums.

Given that workers are more likely to communicate with others from the same forums, information should flow easily within subcommunities. One may wonder how information spreads *between* subcommunities. Are there “connectors” in the network who bridge subcommunities [2, 3]? In fact, 32.4% of connected workers reported using more than one forum regularly, providing ample opportunities for information to flow from one forum to another through these individuals. Furthermore, among all edges connecting a pair of workers that both reported using forums, 71.8% are between pairs in which at least one worker uses a forum that the other does not. This provides another route for information to spread between subcommunities. This observation supports the theoretical prediction of Kleinberg et al. [9] that if there are informational benefits to bridging communities, many people will take a position in the network to earn, share, and ultimately dilute these benefits.

4.2 Differences Between Subcommunities

We next highlight three major differences across these subnetworks in terms of topological structure, temporal communication patterns, and content of communication, and then discuss implications. As before, we extract the subnetwork corresponding to a forum by taking all connected workers who use the forum and all edges between these workers.

Topological Differences

We first examine how tightly connected each subcommunity is using two metrics: *density* and *transitivity*. Given a network with n nodes and m edges, the density of the network is defined as $d \equiv \frac{2m}{n(n-1)}$, which is the ratio between the actual number of edges in the network and the maximum number of edges that could exist in any network with n nodes [17]. Transitivity measures the degree to which triangles in the network are closed. Let $n_{triangle}$ be the number of triangles in a network (i.e., sets of three nodes with edges between each pair) and n_{triple} be the number of connected triples (i.e., nodes x , y , and z with an edge between x and y and another between y and z ; a set of three nodes can create up to three triples). The network’s transitivity is then $t \equiv 3n_{triangle}/n_{triple}$, which measures the ratio between the actual number of triangles and the maximum number of triangles that could occur in any network with n_{triple} triples.

Intuitively, higher density and higher transitivity both imply a more densely connected network. Table 2 (right section) reports the density and transitivity for each of the five subcommunities. The degree of connectivity varies a lot between subcommunities, with TurkerNation and Facebook being the most tightly connected and MTurkForum the least tightly connected.

To further understand how densely connected the subcommunities are, we measure the *diameter* and the *average shortest distance* between two nodes for the largest connected component of each subcommunity. With the exception of MTurkForum, the largest connected component contains the majority of nodes in the subnetwork for each forum. Table 2 (right section) summarizes these results. TurkerNation and Facebook have the smallest diameter and average shortest distance respectively, suggesting that workers in the largest connected components of these two subcommunities are closer to each other. This echoes our previous observation that the TurkerNation and Facebook subcommunities are more highly interconnected. Despite the largest connected component in the MTurkForum subnetwork containing only 35.3% of workers who use the forum (110 workers), the diameter is still large, further evidence that the MTurkForum community is not tightly connected.

Individual subcommunities are not uniformly dense, but composed of a mixture of tight-knit groups and “star structures,” consistent with the core-periphery structure of social networks [1]. Within tight-knit groups, most workers communicate with each other, forming cliques in the extreme. The sizes of the largest cliques in Reddit HWTF, MTurkGrind, TurkerNation, Facebook, and MTurkForum are 11, 16, 16, 12, and 6, respectively, and these largest cliques account for 1.67%, 4.08%, 8.00%, 9.02% and 1.92% of all workers in each subcommunity. In contrast, *star structures* occur when a large number of workers connect to a common cen-

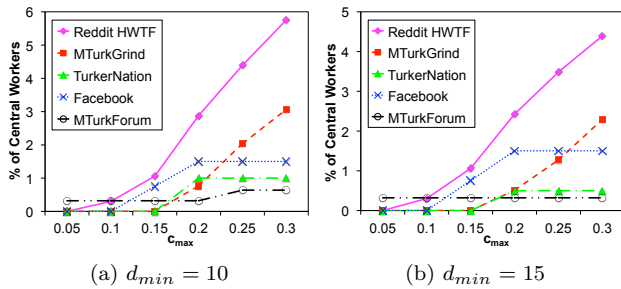


Figure 2: The percentage of central workers of star structures in each subcommunity.

tral worker but not much to each other. To identify star structures, we formally define a “central worker” to be any node with degree at least some value d_{min} and clustering coefficient⁶ at most some value c_{max} . We use the number of central workers identified in a network as a proxy for the number of star structures in it. Figure 2 shows the fraction of workers who are central workers in each subcommunity when we vary d_{min} and c_{max} . By this measure, there exist many more star structures in the Reddit HWTF subcommunity than in any others, a phenomenon that can be observed by a visual inspection of Figures 1b – 1f. This suggests that workers may be using Reddit HWTF in a different way than the other forums. We provide more evidence of this below.

Temporal Communication Differences

HIT completion timestamps can be used to understand the temporal nature of communication in each subcommunity. For this analysis, we coarsely divide the edges in the network into three categories: edges between workers who completed the HIT on the same day, edges added by a worker to another worker who completed the HIT on an earlier day, and edges added by a worker to another worker who completed the HIT on a later day. Note that the third type of edge can only occur when a worker returns to the network visualization another day via their private URL.

Figure 3 shows the fraction of edges that are of each type for each of the subcommunities. More than half of the edges in the Reddit HWTF, MTurkGrind, and MTurkForum subcommunities are between workers who took the HIT on the same day. On the contrary, workers who use TurkerNation and Facebook are much more likely to communicate with other workers who took the HIT on different days. Strikingly, at least 15%-20% of the edges in the TurkerNation and Facebook subcommunities were created by workers who had submitted the HIT on a previous day, but returned to the network to add additional edges later.

To further understand the temporal nature of communication, we calculate two additional quantities for each subcommunity: the empirical probability of a worker in the subcommunity adding an edge to another worker in the subcommunity conditioned on that worker arriving the same day, and the empirical probability of a worker adding an edge to another worker conditioned on that worker arriving a differ-

⁶The clustering coefficient of a node is $c \equiv 2 \times |\{e_{j,k} : e_{j,k} \in E, j, k \in N\}| / (d(d-1))$, where d is the node’s degree, N is the set of the node’s neighbors, E is the set of edges among nodes in N , and $e_{j,k}$ is the edge connecting nodes j and k [18]. This is the ratio between the number of edges between the node’s neighbors and the maximum number of edges between d nodes.




Figure 3: Temporal communication in each subcommunity.

ent day. Specifically, for each worker in a subcommunity we calculate the fraction of all workers who arrived the same day with whom the worker shares an edge and the fraction of all workers who arrived on different days with whom the worker shares an edge, and we average these empirical probabilities across workers. The results, given in Table 3, show that an average worker who uses Reddit HWTF or MTurkForum is an order of magnitude more likely to connect to a worker who accepted the HIT on the same day as opposed to a different day. This effect is dramatically smaller for workers using TurkerNation or Facebook.

Forum Name	Same Day	Different Day
Reddit HWTF	0.049	0.005
MTurkGrind	0.077	0.010
TurkerNation	0.081	0.032
Facebook	0.074	0.035
MTurkForum	0.030	0.002

Table 3: Mean probability of connecting to a worker who took the HIT on the same day or a different day.

Taken together, these results suggest that workers might use Reddit HWTF and MTurkForum to broadcast or obtain information that is immediately actionable, communicating primarily with other workers who happen to be online at the same time. This is in contrast with workers on TurkerNation and Facebook, perhaps indicating that workers on the latter forums form longer lasting relationships.

Communication Content Differences

We turn to a comparison of the topics discussed in different subcommunities. Figure 4 shows the fraction of connected pairs that report communicating about each of five topics: HITs, requesters, Turking scripts and tools, day-to-day-life, and other things. Consistent with the previous literature on forum usage [6, 7, 19], we find intensive discussion about HITs in all subcommunities. Workers in Reddit HWTF almost exclusively discuss HITs. Workers in other subcommunities are more likely to share information about requesters, provide technical support, and recreate the social environment otherwise missing from online work. TurkerNation has the most communication on day-to-day life and Facebook has the most communication on other topics, suggesting that workers use these forums in a more social manner.

Comparing the Subcommunities

Next we put all the differences we have observed together to help us understand how these subcommunities are similar and how they are different. On the one hand, TurkerNation and Facebook might be more socially oriented than other




Figure 4: Topics discussed in each subcommunity.

forums, leading to more tightly connected subcommunities, workers who felt the urge to add edges to other workers they know even if those workers took the HIT on a different day, and more discussions not directly related to MTurk work. In comparison, Reddit HWTF, MTurkGrind, and MTurkForum appear to be mostly dedicated to discussions about details of MTurk work. Reddit HWTF in particular displays a variety of features (e.g., prevalence of star structures and discussions almost exclusively about HITs) which suggest that workers treat it as a platform for broadcasting good HITs above all else. MTurkGrind appears to be something in between a social community and a broadcasting platform, which may be related to the fact that 51.3% of all connected workers who use MTurkGrind also reported using Reddit HWTF. One might conjecture either that MTurkGrind has developed into an independent, more socialized community partly from a pool of Reddit HWTF users, or that MTurkGrind has started to attract users from Reddit HWTF who seek more social interactions. Finally, as we discuss in Section 4.6, MTurkForum accounts for a significant amount of the communication that occurs between workers outside of the United States. This might explain why it seems less connected than other subcommunities.

4.3 The Role of One-on-One Communication

While the majority of communication occurs over forums, workers also report communicating one-on-one via in-person discussions, phone calls, emails, text messages, instant messages, video chatting, and other channels. Overall 13.8% of connected pairs communicate at least partially through one-on-one channels, and 10.1% communicate exclusively through one-on-one channels. Among those pairs that communicate at least partially one-on-one, the three most popular communication channels are instant messaging (27.3%), in-person discussion (18.0%), and email (15.8%).

The role of one-on-one communication is different from that of communication via forums. While forum use is responsible for enabling much of the communication within the largest connected component, one-on-one communication is much more common in the smaller components. Inside the largest component, only 10.7% of connected pairs communicate at least partially through one-on-one channels, and 7.29% exclusively so. Outside of this component, the story is very different: 74.0% of pairs communicate at least partially through one-on-one channels and 63.6% exclusively so. Thus, one-on-one communication accounts for the majority of edges outside of the largest component.

In addition, the distribution of topics discussed by pairs of workers who communicate one-on-one differs substantially from that of workers who communicate over forums. Fig-




Figure 5: Comparison of topics discussed through forums vs. one-on-one communication.

ure 5 compares the amount of discussion for each topic (i.e., the percentage of pairs that communicate on the topic) among pairs who communicate one-on-one vs. in forums. Workers primarily use forums to discuss HITs, while workers who communicate one-on-one communicate much less about HITs and more about day-to-day life and other topics.

4.4 Homophily in the Network

We have seen that there is a communication network within the crowd and that workers communicate across the network both via forums and one-on-one channels. It is natural to ask who it is that workers are most likely to communicate with. In Section 4.1, we showed that there is homophily in the network in terms of forum usage. We now examine whether there is homophily in the network with respect to other worker characteristics.

To answer this question, we follow the same approach used in Section 4.1. First, we apply (generalized, non-binary) homophily tests to examine and compare cross-group ratios. Next, we compare the one-sided homophily measures R and H with the fraction q of workers who share the same property among all connected workers. Using this approach, we do not see strong, consistent evidence for homophily along characteristics such as worker age, gender, education, approval rate, or if a worker is an MTurk Master.

We did, however, find that there is homophily in the network for two other worker characteristics: location and length of time on MTurk. For a worker’s location (limited to just U.S. and Indian workers, $ECGR = 0.249$, $ACGR = 0.107$, $p < 0.001$), it is observed that U.S. workers are much more likely to connect to other U.S. workers ($q = 0.857$, $R = 0.906$, $H = 0.943$), and the tendency for Indian workers to connect with other Indian workers is even more substantial ($q = 0.130$, $R = 0.781$, $H = 0.580$). For the length of time on MTurk ($ECGR = 0.844$, $ACGR = 0.809$, $p < 0.001$), the values for both one-sided homophily measures are also larger than the fraction of workers for almost all groups (“less than 1 year”, “1-2 years”, “2-3 years”, “more than 4 years”) and close for the remaining “3-4 years” group ($q = 0.0914$, $R = 0.1694 > q$, yet $H = 0.0907$ is just slightly less than q). This implies that experienced workers are likely to connect to experienced workers while inexperienced workers tend to communicate with inexperienced workers.

Finally, we analyze homophily around the types of tasks workers regularly do. We could not conduct a single unified homophily test for task type since the vast majority of workers regularly work on more than one type of task. We also did not conduct homophily tests on the binary property of whether or not a worker does a particular type of task (as we did with forum usage) because two workers who




Figure 6: Comparisons of worker fraction and one-sided homophily measures for each type of task.

do not have a particular task type in common may still be very likely to connect to each other because of their shared interest on one or more other types of task. This would make interpreting the ACGR difficult. Figure 6 shows that both one-sided homophily measures are larger than the corresponding fraction of workers who do that type of task for almost all task types, with only the exception of transcription. This indicates that workers tend to communicate with others who work on similar tasks.

4.5 Correlates of Network Position

Next we report our findings on the relationship between network position and various worker properties such as length of time on MTurk, success on MTurk, and access to information. Note that the relationships we report are correlations only. It is impossible to determine whether there is a causal relationship between network position and worker properties from our data.

First, we examine whether workers’ positions in the network have any relationship with how long they have been on MTurk. According to Table 4, the percentage of workers that have been on MTurk for more than 1 year is higher among connected workers than unconnected workers. Consistent with our understanding that the network within the crowd is largely conducted over forums, Table 4 shows that connected workers are also more likely to use forums than unconnected workers.

Property	Connected	Unconnected
Be active >1 year	54.9%	45.9%
Use forums	83.0%	55.5%
Have Master status	11.4%	6.9%
Mean approval rate	98.6%	97.4%

Table 4: Relationship between whether a worker is connected and various worker properties.

Next, we attempt to understand whether workers’ network positions relate to how successful they are. While “success” on MTurk is hard to measure, we can use as a proxy a worker’s approval rate and whether or not the worker has been granted Masters status. These capture how successful a worker has been at getting her own work approved. As Table 4 suggests, by both of these measures, connected workers are more successful than unconnected workers as they are more likely to be MTurk Masters and have higher approval rate on average. At first glance these two effects may seem small, but a 1% increase in approval rate or a Masters qualification allows a worker access to many more HITs




Figure 7: The percentage of connected workers in each bin of 200 workers, ordered by time. The red dashed line is a linear regression trend line.




Figure 8: Degrees for connected workers who did our HIT by day. We omit four outliers with degree over 100.

which could dramatically affect her income. Thus these are very important outcomes for workers.

Finally, we investigate the connection between workers’ network positions and how fast they learn about HITs. We analyze how the network characteristics of workers who accepted our own network mapping HIT changed over time. Specifically, we sort all workers according to the time that they took our HIT and bin them into groups of 200. Figure 7 shows the percentage of connected workers in each bin⁷. There is a clear decreasing trend over time: connected workers were likely to learn about our HIT earlier than unconnected workers. Figure 8 shows a box plot of the degrees of connected workers who took our HIT on different days. Since our data was collected two days after we took down the HIT when few new edges were being added, we believe we gave workers ample time to connect to those workers who took our HIT late, reducing the chance that the low degrees of these workers are an artifact of our data. Here we see that workers who found our HIT earlier also seem to have larger degrees. If this phenomenon generalizes across HITs, this dynamic might result in connected workers starving out isolated workers from high paying tasks.

These results suggest that there are potential benefits to crowdworkers associated with their positions in the network. Being connected is correlated with longevity on the site, higher probability of getting work accepted, and the ability to learn about HITs faster than unconnected workers.

4.6 U.S. vs. International Workers

Finally, we study the differences between workers who

⁷As mentioned in Section 3.2, we notified TurkerNation workers about our test batches of HITs on August 11–12 before launch. Hence we exclude workers who took the HIT on these 2 days in this analysis to minimize possible bias.

are located inside and outside of the United States. Overall, 9,108 workers (88.0%) reported being located in the U.S., while the remaining 1,246 workers (12.0%) reported being located in other countries. While international workers are more likely to be connected than U.S. workers (13.1% U.S. vs. 16.0% international), connected U.S. workers have higher degree on average than connected international workers (8.19 vs. 3.96). This coincides with the finding that a higher percentage of U.S. workers (59.9%) reported using forums than international workers (53.2%), and this comparison is even sharper when we restrict to connected workers (85.8% U.S. vs. 66.5% international).

U.S. workers rely heavily on forums to communicate with each other (91.1% of connected pairs of U.S. workers communicate with each other on forums, and 88.1% exclusively so). International workers tend to use one-on-one channels dramatically more often (76.7% of connected pairs of international workers communicate through one-on-one channels, and 56.9% exclusively so). Interestingly, the most popular forum among U.S. workers is Reddit HWTF while international workers are most likely to use MTurkForum.

The topics discussed among these workers also differ. In particular, a larger fraction of U.S. pairs communicate about HITs (83.2% U.S. vs. 61.2% international), while international workers are much more likely to chat about day-to-day life (12.5% U.S. vs. 30.2% international). This finding coincides with the finding in Figure 5 that forum communication is more likely to focus on HITs while one-on-one communication is more likely to focus on day-to-day life.

In terms of network position, most of the connected U.S. workers (79.8%) are part of the largest connected component, while the majority of connected international workers (77.9%) are in smaller components.

Taken as a whole, this analysis resolves a question left open from Section 4.1: Who are the connected workers who lie outside the largest connected component? These are largely international workers who mostly communicate one-on-one on topics not limited to MTurk work only and are most likely to use MTurkForum if they use a forum at all.

5. CONCLUSION AND DISCUSSION

We designed and executed a HIT to map the network of workers on Amazon Mechanical Turk. Our main result, and the main contribution of this paper, is that there is a substantial communication network within the crowd. Put another way, the crowd is not a collection of independent workers. It is a network. The largest connected component of this network is made up mostly of U.S. workers communicating on various online MTurk forums on which discussion is mostly focused on aspects of MTurk work such as sharing HITs. The network additionally contains many smaller components composed largely of international workers talking with each other through one-on-one channels in which conversations focus on topics like the workers' day-to-day lives in addition to MTurk work. Workers who are part of the network tend to communicate with other workers who are similar to themselves in terms of geographic location, worker experience, and the types of tasks they prefer. Being part of the network may confer some informational advantages to workers allowing them to hear about HITs before workers who are not part of the network. Overall, connected workers tend to be experienced and of high quality.

The existence of the network within the crowd has im-

plications for requesters, workers, and platform designers. Requesters should be aware that the workers they recruit are not an independent sample from the community of active workers. Instead, workers are effectively sampled from a network of workers bound together by the online forums they use or the type of tasks they prefer to do. Since there is homophily among workers, if one worker does a HIT she is more likely to recruit a fellow worker who is similar to her to do the HIT next. If a requester who is using Mechanical Turk to conduct behavioral experiments [11, 13, 14] randomly assigns workers to the treatment and control groups, both groups are still statistically equivalent in all aspects. However, such requesters should carefully consider if the treatment itself would be artificially increased or decreased depending on the characteristics of the population sampled. This is especially true for characteristics like location and experience with MTurk, for which we have shown homophily. Additionally, since workers frequently communicate with one another about HITs, it is natural to ask whether the work that they submit is generated independently or whether they may, for example, share answers with one another. Any discussion among workers of the contents of HITs could bias results.

Our results show that many workers share lucrative tasks and information about reputable requesters with their network connections. With access to this extra information, connected workers might be able to start on high quality tasks before other workers hear about them. In the extreme, this might lead to connected workers using up all of the high paying tasks before isolated workers have had a chance to find them, effectively starving out the isolated workers. Thus, we speculate that being a part of the network may confer an advantage to workers.

All of the forums discussed in this paper were built by workers and exist outside of the Mechanical Turk platform and website. We offer two explanations as to why workers would spend their time building and using these forums. First, it could be the case that participation in forums results in higher pay for workers since they gain access to information about lucrative tasks, as discussed above. Beyond that, workers might inherently value the social interactions that these forums provide. A quote from Zyskowski and Milland [19] indicates that some workers value online forums for both of these reasons: "If I had not found TurkerNation, I would not have made as much money for sure. And the fun we have when things are slow: priceless." Platform designers should be aware that some functionality of their site is missing, so much so that workers felt the need to build that functionality on their own, at their own expense. Crowdsourcing platforms should perhaps consider whether there are ways to make their sites more social and provide workers with the interaction they clearly value.

6. ACKNOWLEDGMENTS

We thank Andrew Mao for helpful technical advice, Kati London for helpful design advice, and the anonymous WWW 2016 reviewers for many helpful comments.

References

- [1] S. P. Borgatti and M. G. Everett. Models of core/periphery structures. *Social Networks*, 21:375–395, 1999.
- [2] R. S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–99, September 2004.
- [3] R. S. Burt. Second-hand brokerage: Evidence on the importance of local structure for managers, bankers, and analysts. *Academy of Management Journal*, 50:119–148, 2007.
- [4] S. Currarini, M. O. Jackson, and P. Pin. An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77:1003–1045, 2009.
- [5] D. Easley and J. Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, 2010.
- [6] M. L. Gray, S. Suri, S. S. Ali, and D. Kulkarni. The crowd is a collaborative network. In *The 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2016.
- [7] N. Gupta, D. Martin, B. V. Hanrahan, and J. O’Neil. Turk-life in India. In *The International Conference on Supporting Groupwork (Group)*, 2014.
- [8] L. C. Irani and M. S. Silberman. Turkoption: Interrupting worker invisibility in Amazon Mechanical Turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2013.
- [9] J. Kleinberg, S. Suri, Éva Tardos, and T. Wexler. Strategic network formation with structural holes. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, 2008.
- [10] D. Martin, B. V. Hanrahan, J. O’Neil, and N. Gupta. Being a turker. In *The 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*, 2014.
- [11] W. Mason and S. Suri. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior Research Methods*, 44(1):1–23, March 2012.
- [12] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.
- [13] G. Paolacci, J. Chandler, and P. G. Ipeirotis. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5:411–419, 2010.
- [14] D. G. Rand. The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299:172–179, April 2011.
- [15] N. Salehi, L. C. Irani, M. S. Bernstein, A. Alkhatib, E. Ogbe, K. Milland, and Clickhappier. We are Dynamo: Overcoming stalling and friction in collective action for crowd workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*, 2015.
- [16] N. Stewart, C. Ungemach, A. J. L. Harris, D. M. Bartels, B. R. Newell, G. Paolacci, and J. Chandler. The average laboratory samples a population of 7,300 Amazon Mechanical Turk workers. *Judgment and Decision Making*, September 2015.
- [17] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge University Press, 1994.
- [18] D. J. Watts and S. H. Strogatz. Collective dynamics of “small-world” networks. *Nature*, 393(6684):440–442, 1998.
- [19] K. Zyskowski and K. Milland. Crowdsourcing visibility: An ethnography of a discussion board for digital labor. Unpublished Manuscript, 2015.