

Accounting for Market Frictions and Power Asymmetries in Online Labor Markets

Sara Constance Kingsley,* Mary L. Gray and Siddharth Suri

Amazon Mechanical Turk (AMT) is an online labor market that defines itself as “a marketplace for work that requires human intelligence.” Early advocates and developers of crowdsourcing platforms argued that crowdsourcing tasks are designed so people of any skill level can do this labor online. However, as the popularity of crowdsourcing work has grown, the crowdsourcing literature has identified a peculiar issue: that work quality of workers is not responsive to changes in price. This means that unlike what economic theory would predict, paying crowdworkers higher wages does not lead to higher quality work. This has led some to believe that platforms, like AMT, attract poor quality workers. This article examines different market dynamics that might, unwittingly, contribute to the inefficiencies in the market that generate poor work quality. We argue that the cultural logics and socioeconomic values embedded in AMT’s platform design generate a greater amount of market power for requesters (those posting tasks) than for individuals doing tasks for pay (crowdworkers). We attribute the uneven distribution of market power among participants to labor market frictions, primarily characterized by uncompetitive wage posting and incomplete information. Finally, recommendations are made for how to tackle these frictions when contemplating the design of an online labor market.

KEY WORDS: market friction, crowdsourcing, quality, Amazon Mechanical Turk, imperfect competition, online labor markets

Introduction

Crowdsourcing is task-orientated labor distributed online through an open call on the Internet (Brabham, 2013; Howe, 2006). Firms look to distribute tasks to crowdsourcing platforms to reduce labor and capital costs, increase the scale of production, and to reach large subject pools quickly. Firms and individuals may tap into “the crowd” to conduct usability testing, research surveys, medical studies (Ranard et al., 2014), and even investigate black market prices for street drugs (Dasgupta et al., 2013). Large firms like AOL,¹ Google,² Unilever,³ and Netflix⁴ all depend on the product of crowdsourcing labor done by hundreds of

*Most of this work was completed while visiting as a Ph.D. intern at Microsoft Research.

thousands of people across the globe to curate content, improve search results, survey potential consumers, and optimize the services they offer people who consume their services and products.

Amazon.com's Amazon Mechanical Turk (AMT), launched in 2005, is one of the earliest iterations of a publicly accessible crowdwork platform, making it an ideal focus for a case study of crowdsourcing market dynamics. Studies evaluating crowdsourcing labor on the AMT platform often recount a particular narrative about AMT's production levels and work quality. This narrative suggests that wages offered in this crowdsourcing market attract the worst quality of workers (Dow, Kulkarni, Klemmer, & Hartmann, 2012; Ipeirotis, 2010). Worker characteristics related to skill, educational attainment, motivation,⁵ and socioeconomic status are then said to explain the dearth of work quality or accuracy rather than the structural features or the design of the AMT labor market (for crowdsourcing literature see, Rogstadius et al., 2011; Shaw, Horton, & Chen, 2011; for economics literature see, Reich, Gordon, & Edwards, 1973).

Alternatively, this article argues that the structure of the platform produces market inefficiencies which channel the vast majority of market power to employers,⁶ which is, in turn, at least partially responsible for the poor quality work observed. Our key contribution is to demonstrate the features of the AMT platform that produce market inefficiencies and impact work quality. First, we show there is an information asymmetry between workers and requesters, which creates an uneven power dynamic between workers and requesters. Second, we evidence how the AMT platform is a highly concentrated labor market: a few requesters post the overwhelming majority of tasks. This limits the ability of workers to compete for tasks that best match their skills. Finally, we show that the platform API⁷ dictates that requesters ("employers") post wages, exacerbating market pressures that limit workers' ability to negotiate wages offered to them.

A key implication of these aspects of the AMT market is that employers wield far more market power than workers. As a result, worker characteristics do not adequately explain production quality on AMT. Instead market asymmetries mediated and reinforced through the API design work to degrade the quality of market outcomes, including employee–employer matches. For example, we argue the API design structures participant interactions in such a way that workers disproportionately absorb the cost of searching for tasks (e.g., labor recruitment costs).⁸

AMT as an Online Labor Market

Our analysis focuses on the AMT labor market, because AMT is "the crowdsourcing site with one of the largest subject pools" (Mason & Suri, 2012, p. 1). Employers, called requesters, post tasks to the AMT marketplace for individuals to do for pay. Individuals, who call themselves or are often referred to as "Turkers" do task-based labor in exchange for a wage set by requesters. Tasks posted to AMT are called Human Intelligence Tasks (HITs). A HIT Group

consists of similar micro-tasks or HITs posted by the same requester (Kittur, Smus, Khamkar, & Kraut, 2011).

The general workflow on AMT is as follows: A requester posts a group of tasks (HITs) for a set wage. Workers search for and do the tasks available to them. Requesters then review the work submitted by individual workers, accept the good work, and unilaterally reject any poor quality work. AMT delivers the requester's posted payment only for the work that the requester deems acceptable and approves through the API. The overall fraction of tasks that a worker has had approved over his or her lifetime is that worker's approval rating. A worker's approval rating serves as a type of reputation score that determines the jobs they will be able to access in the AMT marketplace. Two different types of accounts are available to workers on the AMT platform: a general account offered to all workers, and a Master's account that is, in principle, only offered to workers who maintain a reputation for high job performance. Amazon sets the parameters for establishing all accounts and standards for performance and reputation. Amazon has not disclosed how it defines "high job performance reputation" and does not make its parameters for Masters accounts and standards or criteria for general accounts public or transparent to requesters or workers. Since some tasks are only available to Masters accounts, in practice, AMT determines who may enter this marketplace and who is qualified to access different work opportunities on the platform.

We generally think of online labor markets as places where, "(1) labor is exchanged for money, (2) the product of that labor is delivered" online, "and (3) the allocation of labor and money is determined by a collection of buyers and sellers operating within a price system" (Horton, 2010, p. 516). Commercial crowdsourcing platforms like AMT "serve as the meeting place and market" where micro-task labor is exchanged online for pay (Mason & Suri, 2012, p. 2). This is a critical distinction because legally speaking crowdsourcing sites like AMT have no employment relationship with the people who supply labor on the platform. For this reason AMT differs from employers who typically hire ("buy") labor in offline labor markets. Amazon's primary role is to determine the boundaries of the online labor market. In practice this role does substantially affect labor market outcomes. Amazon's user agreement determines who and how people may participate. Only participants based in the United States may post work and only those registered as workers living in the United States and India may be paid in cash. Workers living in other countries are paid in Amazon.com credit. Thus, while Amazon is not legally an employer, the way Amazon designed the AMT platform does shape the market dynamics of this online labor market.

Research Methodology and Data Sources

We draw on several data sets from a longitudinal study of crowdwork to support our claims about crowdsourcing quality and labor market supply on AMT.⁹ Data sets include: (1) responses to a survey posted to the AMT platform

between July 2013 and 2014; (2) ethnographic data collected from 48 interviews and participant observations conducted in person from September 2013 to July 2014; and (3) results from a geographic mapping task (also called a HIT) posted to AMT.¹⁰

For this article, we examined a total of 317 survey responses from AMT workers (collected for the longitudinal study cited above), including 180 completed surveys from people living in the United States and 137 from people living in India. Since merely posting the survey on AMT may oversample workers who typically do surveys as tasks for work, the larger study embedded the survey into separate image-labeling tasks and email classification tasks. After a worker did 10 email classifications a link appeared asking if they would like to do our survey for additional pay. Since the survey was a vehicle to recruit interview participants this methodological innovation allowed us to reach workers who might not typically do surveys on AMT.

Complementing survey data and ethnographic interviews are measurements of the worker population gathered from tasks posted to AMT. Specifically, this article draws on data from a simple geographic mapping task. The mapping task paid participants to self-report their geo-location and then asked how they found out about the task. A map of the world (via the Bing maps API) was first displayed to workers who were then asked to place a pin where they were located. Our mapping task was posted to AMT for 5 weeks and collected 4,856 pins. Since Amazon does not publish statistics about the people who use and work on its platform, measurements like the mapping task allowed us to approximate the geographic distribution of workers on AMT.

Evidence of Market Power Imbalances

First, we discuss labor market frictions in regard to: (1) whether labor market information is equally available to all participants (e.g., perfect); and (2) the high degree of market concentration evidenced by the fact that a small minority of requesters post the majority of tasks to the platform (Ipeirotis, 2010). Then we turn to an assessment of AMT wage structures broken down by: (1) how wages are determined (*ex ante* vs *ex post* wage posting); (2) whether wage bargaining or negotiation occurs; and (3) if requesters pay the same wage to all workers who have equivalent job performance. All of these factors, mediated by the API, shape the labor market supply dynamics of this online labor market.

Labor Market Frictions

In labor economics, market frictions refer to different “transaction costs” that people and firms incur when participating in a given market (Coase, 1937; Williamson, 1979). For the AMT market we focus on the cost of labor market information. We will show that workers bear the brunt of the cost of the information in this market, which exacerbates the imbalance in market power.

Imperfect Information. When labor market information is costly rather than costless, labor economists call information “imperfect.” Imperfect information is a known source of market frictions (Autor, 2001). Imperfect information violates assumptions given by economic theories that depend on all parties having equal access to the information that people require to make decisions (Stigler & Sherwin, 1985). As we will describe below, AMT’s API distributes task and reputation information unevenly among participants to favor requesters at the expense of crowdworkers.

Amazon’s reputation system is one-sided in that it only signals to requesters how well crowdworkers have performed in the past, assigning each worker an approval rating based on previous requesters’ reported acceptance of workers’ tasks. AMT does not indicate to crowdworkers how well requesters have behaved as employers. Pertinent information not made available to workers on AMT include requesters’ past rejection rates, responsiveness to worker attempts to communicate, and payment history. This means crowdworkers lack mechanisms on the AMT platform to hold requesters accountable for the work they post, in the same way that requesters are able to hold crowdworkers accountable for the work they do.

So while requesters can penalize crowdworkers they deem to be bad actors by blocking workers from doing tasks, withholding payments, rejecting work without reason (though sometimes keeping the output submitted by workers), or reporting workers to Amazon.com (which can lead to suspension of worker accounts), crowdworkers have no mechanism to remedy concerns about requesters. For crowdworkers this makes the cost of finding good tasks to do on AMT higher than if sufficient information about requesters was made available (for relevant treatment of information costs in the economics literature, see Stigler, 1982).

In response to this information asymmetry researchers Irani and Silberman (2013) developed a browser extension (Turkopticon)¹¹ that allows crowdworkers to rate requesters and view ratings submitted by fellow workers. Turkopticon works by collecting and publicly releasing the labor market data workers share through their use of the plug-in that Amazon.com otherwise withholds from market participants. Our surveys and ethnographic data indicate that workers have widely adopted this tool. Several participants noted during interviews that Turkopticon was one of the first tools they read about in online worker forums and the one that they adopted early on to more efficiently identify the “good jobs.”

Workers, particularly newer platform participants with less experience or connections to other workers, are left most vulnerable to incurring added costs exacted by the lack of equally available information. While tools like Turkopticon have helped mitigate some information asymmetries generated by the AMT platform’s “algorithmic authority” (see Gillespie, 2014; Lustig & Nardi, 2015), Turkopticon does not solve the asymmetry of market power. Requesters can still unilaterally reject work and block workers. Turkopticon also does not provide any method for workers to have their reputation repaired or to regain lost wages (Irani & Silberman, 2014).

Online forums frequented by crowdworkers have noticeably proliferated in recent years (Martin, Hanrahan, O’Neill, & Gupta, 2014). Crowdworkers realize the cost of market information asymmetries or frictions when they use online forums to share information about the quality of tasks available on AMT, and swap recommendations about requesters. As in the case with Turkopticon, online forums have helped mitigate the information asymmetry between workers and requesters generated by the AMT API but these forums do not solve the market power asymmetry. Indeed, the online forum traffic specifically illustrates the recruitment costs that requesters shift to workers searching for tasks to do on AMT. Since AMT does not post labor market information directly on the platform, crowdworkers spend substantial amounts of time searching for information on the Internet (Yuen, King, & Leung, 2012).¹² Most economists will agree that time spent searching for work is a cost or rent job seekers bear in order to secure future employment. Normally wages and salaries are considered the return workers receive for investments made when searching for jobs (Diamond, 2011; Mortensen, 2011; Pissarides, 2011).

We wanted to assess the value of online forums as a resource that could offset a lack of perfect information about AMT for this reason. We posted an experimental mapping task to AMT to evaluate what fraction of workers came to our task via online forums and what fraction came to our task via other methods, such as by searching the AMT website using AMT’s built-in search functionality.

Figure 1 displays the traffic flows to our experimental mapping task by the information source workers used to discover it. The x-axis denotes time starting at April 23, 2014, which is when the mapping HIT was launched, and indicates each 8 hour-period until May 28, 2014, which is when the mapping HIT was taken down. Thus the HIT ran for 5 weeks in total. The y-axis shows how many




Figure 1. Online Forum Use by Crowdworkers to Find HITs on AMT.

Note: Just under half of the workers came to this HIT via an online forum. Based on 4,856 responses from April 23, 2014 to May 28, 2014. Graph credits: Gregory Minton.

workers did our task during a given day. So the height of each bar indicates the total number of workers who did our HIT in a given 8-hour period. The coloring of the bar indicates how the workers who did the HIT learned about the HIT. The length of the black portion of each bar represents the number of workers who were directed to the task from online forums and the length of the blue portion indicates the number of workers who found the task by using the search bar tool on the AMT platform. These were the two predominant channels by which workers found our HIT. Figure 1 shows that many of the bars, especially the higher ones, have large black portions. This indicates that a large fraction of traffic, in fact almost half, came to our experimental mapping task from online forums. If sufficient information about AMT was available to market participants we would expect the overwhelming majority of workers to have arrived via searching the AMT site. That tens of thousands of registered AMT users generate traffic on multiple forums dedicated to finding good tasks highlights the scale of the inefficiency that an information-starved market creates. As such, we argue that online crowdsourcing forums signal the transfer of labor recruitment costs to crowdworkers.

Economic theory posits that searching for a job is a rent that workers pay now in order to find income earning opportunities in the future (Manning, 2003). However, the nature of crowdsourcing labor is iterative, meaning workers must constantly search for new tasks to do because no long-term employment contract exists between workers and requesters. Arguably this means workers continue to pay search or labor market entry costs without receiving a full return from the wages they earn. Data about forum traffic are indicators of market entry costs incurred by crowdworkers for this reason. Ultimately for crowdsourcing labor markets it is critical to understand that workers' search efforts are far from free. Search time is a cost or rent borne by those who are actively looking to find decent work online, and these workers' costs should be factored into the valuation of this market.

Search activity within labor markets marred by imperfect information produce another telling outcome: poor employer–employee matches (Benson, 2013; Priest, 2008). Matches in labor markets with patchy information about past performances or human capital are typically bad because parties exchanging pay for labor do not have the information needed to make optimal choices about who to hire and what jobs to accept. Put another way, prices are not set at a level where “the quantity that the buyers want to purchase is exactly that which sellers are willing to provide” (Mortensen, 2011, p. 1074). This might explain why AMT's reported work quality is not responsive to changes in price (Mason & Watts, 2010).

While it may be fair to suggest that forums register as an aberration in a well-functioning market, we do recognize that forums also serve a range of other important functions, which go beyond providing market information. They also offer social cohesion, mentorship, help workers build a sense of identity, and offer them entertaining breaks from work routines (Martin et al., 2014). Forums represent the trade-offs between efficiency and social interaction found in a

marketplace (Lehdonvirta & Castronova, 2014, p. 130, figure 7.2) and crowdsourcing platforms are no exception. However, not all workers interviewed were aware of this resource. And without a mechanism on the platform itself to right the imbalance of market power produced by AMT's reputation system, the costs associated with AMT market participation remains a significant impediment for many workers.

Market Concentration on AMT. Ipeirotis (2010, p. 17) found that the "top requesters" generate "more than 30 percent of the overall activity in the market" yet comprise a sparse "0.1 percent of [...] total requesters" on AMT. This trend is consistent with monopsonistic or oligopsonistic features of labor markets with imperfect competition (Manning, 2003; Ransom & Sims, 2008). We evaluated market power among AMT requesters in our sample by looking to see if the number of income-earning opportunities or tasks posted to the platform are limited to a small number of requesters. We evaluated a data set produced by scraping information about HIT Groups posted to AMT on August 8, 2014.¹³ Market level information for AMT in this data set includes the wages requesters set for tasks, how many tasks each HIT Group contained, and a description of each task that requesters were asking workers to do. From this data set we calculated the distribution of tasks posted by the top 10 percent of requesters. We defined the top 10 percent of requesters by the frequency at which requesters posted tasks to the market. We then estimated the distribution of tasks by HIT Group size. For our sample, HIT Group size ranged from those HIT Groups with only one task to a maximum size of 39,588 tasks posted to a single HIT Group by an individual requester.

Findings from our data analysis of requester and task distribution on AMT are illustrated by Figures 2 and 3. Figure 2 shows that 10 percent of all requesters post approximately 98–99 percent of all tasks to the AMT platform. Figure 3 shows that approximately 98–99 percent of all tasks on AMT are posted to only 10 percent of HIT Groups. A HIT Group is a collection of similar tasks posted by the same requester each for the same pay.

Although the total number of requesters posting tasks to AMT seems large, our data demonstrates that the fraction of requesters posting tasks most frequently to the platform is highly concentrated to a small percent of all requesters, supporting claims found in the existent crowdsourcing literature. Ipeirotis (2010, p. 17), for example, notes that this "high concentration is not unusual for any online community" and that "there is always a long tail of participants" who have "significantly lower activity than the top contributors." While it may be unsurprising to see such a concentration in this market, since this happens in an online labor market this concentration has an important effect. We argue that this concentration of requesters could hold a significant degree of market power. Requester concentration implies that workers have fewer wage earning opportunities outside this top percent of requesters. In economic terms, this means that a small fraction of the total number of requesters likely capture the majority of labor supply on AMT.




Figure 2. Market Concentration by Requesters.




Figure 3. Market Concentration by HIT Groups.

AMT Wage Structures

One of the strongest indications of prevailing imbalances in market power comes from analysis of AMT's site design itself (Ipeirotis, 2010; Khanna, Ratan, Davis, & Thies, 2010; Silberman, Ross, Irani, & Tomlinson, 2010). As Figure 4 shows, AMT's design requires requesters to post tasks to the AMT market and list the wage for each task (HIT) that they wish to offer crowdworkers on the platform before any work is done. Requesters log onto the platform and post tasks within a group of similar tasks (a so-called HIT Group), predetermining and unilaterally setting wage rates through the API. In other words, Figure 4 clearly shows that requesters set wages *ex ante*. This means wages for tasks are determined before requesters interact with the crowdworkers who end up taking on the requesters' posted tasks. In labor economics wage determinations set by

The screenshot shows the Amazon Mechanical Turk interface. At the top, there's a navigation bar with 'amazonmechanicalturk' logo, 'Artificial Intelligence', and buttons for 'Your Account', 'HITS', and 'Qualifications'. A notification indicates '363,450 HITs available now'. Below this is a search bar with 'Find HITs' and filters for 'that pay at least \$ 0.00' and 'require Master Qualification'. The main content area displays 'All HITs' and '1-10 of 2119 Results'. The results are sorted by 'HITS Available (most first)'. Five HITs are listed in a table format:

Requester	HIT Expiration Date	Reward	HITS Available
Fiber Data Co.	Sep 18, 2014 (4 weeks)	\$0.03	67823
Jon Breilig	Aug 28, 2014 (6 days 23 hours)	\$0.03	44405
Jon Breilig	Aug 27, 2014 (6 days 3 hours)	\$0.08	24256
rohzt0d	Sep 12, 2014 (3 weeks 1 day)	\$0.00	19441
Jon Breilig	Aug 28, 2014 (6 days 23 hours)	\$0.09	17139

Figure 4. Screenshot of the Amazon Mechanical Turk Marketplace.

firms who do not negotiate with workers beforehand are called *ex ante* wage postings (Manning, 2003).

That requesters set wage rates *ex ante* and do so uniformly for each individual task they post suggests that crowdworkers have little to no ability to negotiate their wages with requesters on AMT. Economic theory about imperfect competition assumes market distortions (“frictions”), which limit the ability of workers to negotiate wages, produce market power that favors employers since employers rather than the market set wage rates (Ashenfelter, Farber, & Ransom, 2010a, 2010b; Manning, 2003; Staiger, Spetz, & Phibbs, 2008). Since we cannot provide evidence that employers are systematically posting wages below the expected marginal product of a worker, we cannot conclusively state that *ex ante* wage posting causes employers to have market power. However, we do believe market concentration, discussed previously, suggests that *ex ante* posting gives requesters a greater degree of market power than they might have had otherwise.

Ex ante wage posting has additional market implications. First, since wages are predetermined and posted by requesters, crowdworkers have few alternatives to accepting the wages requesters offer, aside from searching for other tasks to do on AMT and incurring added search costs or walking away from the AMT platform altogether (for a related discussion in the economics literature, see Dixon, 1987). Once a worker has spent time learning how to search and do specific tasks on the AMT platform, they are hard-pressed to find alternatives, which will allow them to capture surplus profits or extract rent from their experience. This limits the ability of workers to transfer their skills to another workplace.

Second, *ex ante* wage posting also means workers with higher productivity rates cannot be proportionately compensated for their work when the wage rate is set ahead of time. That requesters do not pay crowdworkers according to their

job performance for each individual task they do or the investments they make in obtaining higher skills and education evidences the noncompetitive nature of wage structures on AMT.

Implications of Market Power Imbalance

We have established that requesters have far more market power due to the presence of asymmetric reputation systems, the concentration of tasks supplied by small number of firms or individuals providing work, and the ability to set wages and unilaterally reject work. Next we describe some implications of these phenomena. Since there is no version of AMT or a competitor available for comparison that is free from these market power imbalances, we turn to economic theory to illustrate the effects of the market power imbalances. Because there is no reputation mechanism for requesters that is a default on the AMT system, workers are forced to either pay the cost of exploring different requesters by installing third party browser plug-ins like Turkopticon and browsing online forums or take on the risk of working for requesters of unknown reputation. In either case, workers bear the costs of the information asymmetry.

Since requesters have to uniformly set wages across the worker population, they cannot preferentially set higher wages for the good workers and lower wages for the bad workers. This makes it more difficult for good workers to find the good requesters and requesters to pay according to the quality that they might expect from a worker, which contributes to poor worker–requester matches. The concentration of work among a small fraction of all requesters means less diversity of tasks are available to workers. This makes it harder for a new worker to find a task suited to them. Again, this contributes to poor worker–requester matches. Finally, we give some evidence that *requesters* are already experiencing poor matches with workers. Mason and Suri (2012) advocate recruiting a panel of trusted workers to participate in research studies, which has been used to great effect in follow up work (Mason & Watts, 2012; Suri & Watts, 2011; Wang, Suri, & Watts, 2012). Furthermore, Crowdfunder and SpeechInk (now called SpeechPad), companies that once posted tasks on AMT on behalf of other companies, implemented their own reputation systems for its workers. One can surmise that Crowdfunder and SpeechInk did this because they found the AMT reputation system insufficient for finding the optimal workers for sets of tasks. In summary, we see that the asymmetrical distribution of market power between workers and requesters created by the AMT API has a real and felt impact on the quality of worker–requester matches. For this reason, we next consider how to remedy asymmetrical market power in online, commercial crowdsourcing markets like AMT.

Conclusion: Remedies for Online Market Power

This article primarily seeks to understand the balance of market power in the AMT crowdsourcing labor platform through a detailed, multimodal analysis of

the labor market supply dynamics on AMT. We scrutinize AMT as a longstanding player and point of reference for others designing crowdsourcing platforms to understand what features might produce a consequential degree of unchecked (and unrecognized) market power. A few different technical remedies for these market power imbalances are possible for crowdsourcing labor markets. The solutions presented are given according to the most probable causes of unchecked market power in the AMT market; that is, noncompetitive wage structures, and labor market frictions, specifically imperfect information.

Solutions for Noncompetitive Wage Structures

First, wage structures on AMT could be made more competitive by instituting mechanisms that allow for wage bargaining—negotiations—between requesters and crowdworkers. One solution might be to allow a double auction between workers and requesters, which would encode the bargaining over wages and jobs into a well-studied mechanism.¹⁴ A second and preliminary example of how this could work is illustrated by Dynamo, a platform built by researchers to support crowdworkers' efforts to share information, collaborate, and determine guidelines for academic requesters setting wages and task design.¹⁵ The guidelines speak specifically to the issue of fair payment for crowdwork on Dynamo. Workers are given a central role (voice) on Dynamo in deciding what constitutes fair pay.

That said, the Dynamo platform does not enable face-to-face or real-time wage negotiation. In online labor markets where speed is an essential factor, information and communication delays are costly, and real-time negotiation mechanisms become critical to correcting skewed market power between those paying for labor and the people supplying it. Crowdsourcing platforms could incorporate online chat services directly into the platform, permitting requesters to talk directly to crowdworkers in real-time. Scalability, however, may limit the feasibility of this solution. Requesters typically need large subject pools to complete their tasks, and for this reason, it is hard to imagine requesters chatting with each individual crowdworker they need to hire. Alternative tools can communicate information quickly to all parties working in a virtual system. Answers to a prompt about what constitutes fair pay for a particular task could rapidly circulate opinions among participants. Some researchers have started to explore the role that systems-level visualizations can play in this regard.¹⁶ Either way, innovations like Dynamo offer examples of what it could look like to explore and create spaces for requesters and crowdworkers to work together to determine wage rates, reduce market frictions and correct imbalances of power between AMT participants.¹⁷

Solutions for Imperfect Labor Market Information

On the surface it might seem as though Amazon's internal platform reputation system offers an effective quality control mechanism, accurately signaling to requesters what they can expect from a worker's job performance.

However, AMT's reputation system creates more information asymmetries than clarity. As argued above, this is strongly evidenced by the widespread awareness and adoption of external remedies like Turkopticon, the plug-in that seeks to provide crowdworkers with information about the quality of requesters and tasks being posted to the market. However, even Turkopticon is not a sufficient fix for the information asymmetries on AMT, as data from online forums for workers who crowdsource make abundantly clear. As mentioned, crowdworkers on forums frequently discuss the unscrupulous behavior of many requesters. This is also why remedies like Dynamo have sought to, first and foremost, provide basic guidelines to particular categories of requesters.¹⁸ As helpful as these additional tools are, we argue that the dynamics of a healthy market demand that information be made readily available to participants through information tools directly embedded in the infrastructure of the marketplace to ensure equal access for all participants.

Vital pieces of information, from requester reputations to a real-time list of jobs and workers in the system, remain scattered across the Internet. Crowdwork labor market participants must therefore absorb the costly scavenger hunt to make informed decisions about their participation on AMT. Crowdsourcing labor markets are doomed to reproduce labor market inequalities and generate market frictions if they fail to supply all market participants with the same information needed to fairly compete. To correct this problem, communication and reputation systems on platforms like AMT need to be made more transparent and inclusive of key constituents if they are to incorporate the "voice" of workers (see Freeman, 1980). This could take the form of crowdworkers being able to rate requesters directly on the AMT platform without needing to install additional software, while also allowing crowdworkers to determine the metrics or standards by which these ratings are constructed. Then, hopefully, crowdworkers would have equal opportunity to hold requesters accountable for their on-platform behavior, and the quality of tasks they design, as requesters are already able to hold crowdworkers to account for the work that they do.

Other solutions could include real-time communication tools made directly available to both requesters and crowdworkers on the AMT platform. This would reduce the amount of time crowdworkers spend in online forums, searching for good tasks to do on the AMT platform. Some platforms already try to implement in-platform communication tools. The Lead Genius (formerly MobileWorks) platform, for example, provides a chat service to crowdworkers so that they may communicate with each other in real-time when doing projects together. Legal implications, however, might currently prevent AMT from adopting similar measures. For this reason, in the next section we discuss policy and legal concerns relevant to online, crowdsourcing labor markets.

Policy Consideration for the AMT Labor Market

Technical remedies for imbalance in market power are severely limited by policy and legal frameworks. As much as Amazon's technological systems shape

the kind of information exchanged directly on the platform, legal systems shape the parameters and rules for permissible activities and actors in labor markets. For this reason, the interplay between technological choices, and how our legal institutions either broaden or limit those options, should not be ignored. For instance, if platform providers like AMT are defined as employers, legally speaking, many platform providers would likely opt out of the market, and no longer provide the environment necessary for the online exchange of labor to occur. Conversely, if the people who post tasks to online, crowdsourcing labor markets are defined as employers, they will face prohibitive costs associated with the legal obligations of being an employer. This outcome would not only harm people posting tasks, but the hundreds of thousands of people who rely on the income they earn from the work they do online. How, then, might we imagine a future that expands the opportunity to earn money through flexible, short-term contracts while still offering fair payment for quality work?

Today, platform providers are incentivized to minimize the risk of being deemed to be an employer under the law. Most platform providers will not integrate technical fixes to their systems that support workers through training, collaboration, and information sharing, as such enhancements may suggest that the platform curates a workforce. As the class action lawsuit¹⁹ brought against the editorial crowdsourcing site, Crowdfunder, suggests, we have yet to legally decide what kind of employment crowdwork, technically, is (see NewScientist, 2013). Additionally, most platforms do not directly set wage rates, and instead leave wage setting to people posting tasks. Horton (2010, p. 517) suggests, however, and we agree that, “the influence of the market creator is so pervasive that their role in the market is closer to that of a government...they determine the space of permissible actions within the market, such as what contractual forms are allowed and who is allocated decision rights.” The AMT platform is the *location* where online labor takes place. And today, at least in the United States, it is hard to think of many workplace environments that are not at least minimally regulated to ensure the well-being and safety of both employers and their employees.

Many platforms strive to thoroughly integrate computational infrastructure and “humans-as-a-service,” potentially eliminating the interactivity that we traditionally associate with labor markets. We would argue that the cultural logics that currently orient most of society’s members to employment are permeated with expectations of professional and personal relationships that we all associate with work environments. These expectations require us to rethink what constitutes meaningful employment rather than assume that interactivity is no longer necessary. Delegating the management of workflows through an API does not eliminate workers’ needs for these relationships. Indeed, the expansive use of forums amplifies the value of these relationships as workers and requesters seek other means to communicate and collaborate with each other. In consideration of these points, we argue that the effort that platform providers make to avoid costly legal responsibilities contributes to the market power imbalances we observe. Therefore, the following policy fixes are recommended. First, treat crowdsourcing labor markets according to their needs, and not those of traditional, offline

markets. Doing so requires policymakers to enact new rules, which will define employment relationships in crowdsourcing labor markets, and protect crowdworkers who exchange their labor for pay online. Second, institute enforcement mechanisms to hold bad actors on platforms like AMT accountable for their actions, especially those requesters who commit cybercrimes, and violate best practices, such as researchers not abiding by ethical standards of universities and Institutional Review Boards. Finally, consider mechanisms to make the role of platform providers similar to those of a fiduciary, in that they should act in the best interest of all parties on the platform, and not the select interests of a few.

Numerous crowdsourcing studies, particularly those focused on AMT, offer preliminary evidence that online labor markets feature power imbalances between workers and requesters. We argue that power imbalances are a result of platform design, specifically how the AMT API influences: (1) employer-based wage setting; (2) the number of requesters posting tasks on the platform (market concentration); and (3) reputation, and other costly market frictions caused by asymmetric information problems. We emphasize that as new online markets emerge, the discipline of economics will need to consider the implications of online labor market design for the standard models employed by economics. In doing so, economics, as a whole, will require a new framework to understand the contexts in which people exchange labor online for pay.

Sara Constance Kingsley, B.A., University of Massachusetts Amherst, Amherst, Massachusetts [saraki@microsoft.com].

Mary L. Gray, Ph.D., Microsoft Research, Cambridge, Massachusetts, and Indiana University, Indiana.

Siddharth Suri, Ph.D., Microsoft Research, New York City, New York.

Notes

1. See AMT requester case studies. https://requester.mturk.com/case_studies.
2. Google posts HITs to the AMT platform. HIT details available on request from authors.
3. Unilever contracts with the Jana.com crowdsourcing platform. Please see the “case studies” section on Jana’s website. <http://www.jana.com/case-studies/>.
4. Netflix contracts with the Amara.org platform. Please see the “enterprise” section of Amara.org’s website. <http://about.amara.org/enterprise/>. Please also see Roettgers, J. (July 30, 2012) Netflix experiments with crowd-sourced captioning. *Gigom*. <https://gigaom.com/2012/07/30/netflix-amara-closed-captions-crowdsourcing/>.
5. Often this narrative is framed in terms of the intrinsic and extrinsic motivations of workers (Rogstadius et al., 2011) and the incentives of workers to “game the system” (Ipeirotis, 2010).
6. Employers on AMT refer to those who post tasks for workers to do for pay. We strongly emphasize, however, that employers or requesters on AMT are not defined as employers under the law in the United States. Please see *Otey v. Crowdflower, Inc. et al.* for details on a legal case making its way through the court system; determining who is an employer (if anyone) on AMT under the law is a question central to the legal argument.
7. For an informative discussion about APIs, please see Evans, Hagiu, and Schmalensee (2011).
8. Task recruitment refers analogously to one type of labor recruitment cost.
9. See Gray M.L., and Suri S. *On-Demand: Crowds, Platform Economies, and the Future of Work in Precarious Times* (in progress), which studies four different crowdsourcing platforms, across two

- continents, over a period of 18 months. For more information, see <http://research.microsoft.com/en-us/projects/crowdwork/>.
10. For this article, we integrate some qualitative data gathered during 10 months of ethnographic fieldwork in India undertaken as part of the larger research project. The larger project's survey asked respondents doing paid crowdwork on AMT a range of questions, from inquiries about basic demographics to specifics concerning computer literacy and Internet skills. We did open coding analysis of a subset of survey 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. Qualitative data includes responses and data gathered from 48 in-person, open-ended, semi-structured interviews and hundreds of hours of follow up observations with research participants. Interview participants were identified through worker referrals, contacts made in online worker discussion forums, and the survey itself.
 11. See Turkoicon: <http://turkoicon.ucsd.edu/>.
 12. For a discussion of increased income from better labor market information, please see Agrawal, Horton, Lacetera, and Lyons (2015).
 13. Data was gathered by Chien-Ju Ho.
 14. We thank an anonymous reviewer for this suggestion.
 15. See <http://www.wearedynamo.org/>.
 16. These living guidelines are collaboratively prepared by crowdworkers who are active on the Dynamo platform. Currently, participation is only open to active crowdworkers on AMT.
 17. See <http://www.wearedynamo.org/>.
 18. See http://wiki.wearedynamo.org/index.php/Guidelines_for_Academic_Requesters.
 19. See *Otey v. Crowdflower, Inc. et al.* <http://law.justia.com/cases/federal/district-courts/california/candce/3:2012cv05524/260287/124>.

References

- Agrawal, A., J. Horton, N. Lacetera, and E. Lyons. 2015. "Digitization and the Contract Labor Market: A Research Agenda." In *Economic Analysis of the Digital Economy*, eds. A. Goldfarb, S. Greenstein, and C. Tucker. Chicago: The University of Chicago Press, 219–50.
- Ashenfelter, O., H. Farber, and M. Ransom. 2010a. "Labor Market Monopsony." *Journal of Labor Economics* 28 (2): 203–10.
- Ashenfelter O., H. Farber, and M. Ransom. 2010b. "Modern Models of Monopsony in Labor Markets: A Brief Survey. Institute for the Study of Labor (IZA)." Discussion Paper No. 4915, Bonn, Germany.
- Autor, D. 2001. "Wiring the Labor Market." *The Journal of Economic Perspectives* 15 (1): 25–40.
- Benson, A. 2013. "Firm-Sponsored General Education and Mobility Frictions: Evidence From Hospital Sponsorship of Nursing Schools and Faculty." *Journal of Health Economics* 32 (1): 149–59.
- Brabham, D. 2013. *Crowdsourcing*. Cambridge, MA: MIT Press.
- Coase, R.H. 1937. "The Nature of the Firm." *Economica* 4 (16): 386–405.
- Dasgupta, N., C. Freifeld, J.S. Brownstein, C.M. Menone, H.L. Surratt, L. Poppish, J.L. Green, E.J. Lavonas, and R.C. Dart. 2013. "Crowdsourcing Black Market Prices for Prescription Opioids." *Journal of Medical Internet Research* 15 (8): e178.
- Diamond, P. 2011. "Unemployment, Vacancies, and Wages." *American Economic Review* 101: 1045–72.
- Dixon, H. 1987. "A Simple Model of Imperfect Competition With Walrasian Features." *Oxford Economic Papers* 31 (1): 134–60.
- Dow, S.P., A. Kulkarni, S.R. Klemmer, and B. Hartmann. 2012. "Shepherding the Crowd Yields Better Work." In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*. New York: ACM, 1013–22.
- Evans, D.S, A. Hagiu, and R. Schmalensee. 2011. *Invisible Engines*. Cambridge, MA: MIT Press.
- Freeman, R.B. 1980. "The Exit-Voice Tradeoff in the Labor Market: Unionism, Job Tenure, Quits, and Separations." *The Quarterly Journal of Economics* XCIV (4): 643–73.
- Gillespie, T. 2014. "The Relevance of Algorithms." In *Media Technologies: Essays on Communication, Materiality, and Society*, eds. T. Gillespie, P. Boczkowski, and K. Foot. Cambridge, MA: MIT Press, 167–94.

- Horton J.J. 2010. Online Labor Markets. In *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE)*, ed. A. Saberi, Lecture Notes in Computer Science 6484. Berlin, Heidelberg: Springer, pp. 515–22.
- Howe, J.J. 2006. "The Rise of Crowdsourcing." *Wired Magazine* 14: 1–4.
- Ipeirotis, P.G. 2010. "Analyzing the Amazon Mechanical Turk Marketplace." *XRDS: Crossroads, The ACM Magazine for Students* 17 (2): 16–21.
- Irani, L., and S.M. Silberman. 2013. "Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York: ACM, 611–20.
- Irani, L., and S.M. Silberman. 2014. "From Critical Design to Critical Infrastructure." *Interactions* 21 (4): 32–5.
- Khanna, S., A. Ratan, J. Davis, and W. Thies. 2010. "Evaluating and Improving the Usability of Mechanical Turk for Low-Income Workers in India." In *Proceedings of the First ACM Symposium on Computing for Development*. New York: ACM, Article No. 12.
- Kittur, A., B. Smus, S. Khamkar, and R.E. Kraut. 2011. "CrowdForge: Crowdsourcing Complex Work." In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. New York: ACM, 43–52.
- Lehdonvirta, V., and E. Castronova. 2014. *Virtual Economies: Design and Analysis*. Cambridge, MA: MIT Press.
- Lustig, C., and B. Nardi. 2015. "Algorithmic Authority: The Case of Bitcoin." Paper presented at the IEEE Computer Society Proceedings of HICSS-48, January 5–8, Kauai, HI.
- Manning, A. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.
- Martin, D., B.V. Hanrahan, J. O'Neill, and N. Gupta. 2014. "Being a Turker." In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. New York: ACM, 224–35.
- Mason, W., and S. Suri. 2012. "Conducting Behavioral Research on Amazon's Mechanical Turk." *Behavior Research Methods* 44 (1): 1–23.
- Mason, W., and D. Watts. 2010. "Financial Incentives and the Performance of Crowds." *ACM SigKDD Explorations Newsletter* 11 (2): 100–8.
- Mason, W., and D.J. Watts. 2012. "Collaborative Learning in Networks." *Proceedings of the National Academy of Sciences* 109 (3): 764–69.
- Mortensen, D. 2011. "Markets With Search Friction and the DMP Model." *American Economic Review* 101: 1073–91.
- NewScientist. 2013. *Time to Focus on the Welfare of Online Workers*. <http://www.newscientist.com/article/mg21729033.400-time-to-focus-on-the-welfare-of-online-workers.html#.VATftknJaQ>.
- Pissarides, C. 2011. "Equilibrium in Labor Markets With Search Frictions." *American Economic Review* 101: 1092–105.
- Priest, G. 2008. "Timing 'Disturbances' in Labor Market Contracting: Roth's Findings and the Effects of Labor Market Monopsony." *Journal of Labor Economics* 28 (2): 447–72.
- Ranard, B.L., Y.P. Ha, Z.F. Meisel, D.A. Asch, S.S. Hill, L.B. Becker, A.K. Seymour, and R.M. Merchant. 2014. "Crowdsourcing—Harnessing the Masses to Advance Health and Medicine, a Systemic Review." *Journal of General Internal Medicine* 29 (1): 187–203.
- Ransom, M.R., and D.P. Sims. 2008. "Estimating the Firm's Labor Supply Curve in a 'New Monopsony' Framework: Schoolteachers in Missouri." *Journal of Labor Economics* 28 (2): 331–55.
- Reich, M., D.M. Gordon, and R.C. Edwards. 1973. "Dual Labor Markets: A Theory of Labor Market Segmentation." *American Economic Review* 63 (2): 359–65.
- Rogstadius, J., V. Kostakos, A. Kittur, B. Smus, J. Laredo, and M. Vukovic. 2011. "An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets." Presented at the International AAAI Conference on Web and Social Media, North America, July.
- Shaw, A., J. Horton, and D. Chen. 2011. "Designing Incentives for Inexpert Human Raters." In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*. New York: ACM, 275–84.

- Silberman, M., J. Ross, L. Irani, and B. Tomlinson. 2010. "Sellers' Problems in Human Computation Markets." In *Proceedings of the ACM SIGKDD Workshop on Human Computation*. New York: ACM, 18–21.
- Staiger, D., J. Spetz, and C. Phibbs. 2008. "Is There Monopsony in the Labor Market? Evidence From a Natural Experiment." *Journal of Labor Economics* 28 (2): 211–36.
- Stigler, G.J. 1982. *The Process and Progress of Economics. The Nobel Prize Committee*. Graduate School of Business, University of Chicago, Nobel Memorial Lecture.
- Stigler, G.J., and R.A. Sherwin. 1985. "The Extent of the Market." *Journal of Law and Economics* 28 (3): 555–85.
- Suri, S., and D.J. Watts. 2011. "Cooperation and Contagion in Web-Based, Networked Public Goods Experiments." *ACM SIGecom Exchanges* 10 (2): 3–8.
- Wang J., S. Suri and D. Watts. 2012. "Cooperation and Assortativity With Dynamic Partner Updating." *Proceedings of the National Academy of Sciences* 109 (36): 14363–68.
- Williamson, O.E. 1979. "Transaction-Cost Economics: The Governance of Contractual Relations." *Journal of Law and Economics* 22 (2): 233–61.
- Yuen, M.-C., I. King, and K.-S. Leung. 2012. "Task Recommendation in Crowdsourcing Systems." In *Proceedings of the First International Workshop on Crowdsourcing and Data Mining*. New York: ACM, 22–6.