Monopsony and the Crowd: Labor for Lemons?

Sara C. Kingsley, University of Massachusetts Amherst[1]
Mary L. Gray, Microsoft Research
Siddharth Suri, Microsoft Research

Introduction

Crowdsourcing is task-orientated labor distributed by requesters ("employers") to crowdworkers ("employees") online through an open call on the Internet (Howe, 2006; Brabham, 2013). Crowdsourcing labor platforms are a mechanism to distribute requests from people and firms who want to outsource large numbers of micro-tasks requiring a large workforce, "the crowd." Firms look to outsource tasks to crowdsourcing platforms to reduce labor (Felstiner, 2011) and capital costs, increase the scale of production, and to reach large subject pools quickly. Firms want to tap into "the crowd" in order to conduct usability testing, research surveys, and even medical studies (Ranard, Ha, Meisel and et. al., 2013) and investigations of black market prices for street drugs (Nabarun, et. al., 2013), among other purposes. Services offered by large firms like AOL[2], Google[3], Unilever[4], and Netflix[5] all depend on the product of crowdsourcing labor done by hundreds of thousands of people across the globe. People who crowdsource for these companies are not legally their employee or even

[1] Sara C. Kingsley was a visiting researcher at Microsoft Research, New York during the time this paper was written. The authors would like to thank their colleagues at Microsoft Research and Professors Michael Ash, Nancy Folbre, and Gerald Friedman and Luke Pretz, PhD Candidate, at the University of Massachusetts, Amherst, for their generous insights and support.

[2] See: Amazon Mechanical Turk | Requester, Case studies. Available online: https://requester.mturk.com/case_studies

[3] Google Ireland posts HITs to the Amazon Mechanical Turk (AMT) platform. HIT details available on request from authors.

[4] Unilever contracts with the Jana.com crowdsourcing platform. Please see the "case studies" section on Jana's website. Available online: http://www.jana.com/case-studies/

[5] Netflix contracts with the Amara.org platform. Please see the "enterprise" section of Amara.org's website. Available online: http://about.amara.org/enterprise/. Please also see: Roettgers, J (July 30, 2012) Netflix experiments with crowd-sourced captioning. *Gigom*. Available online: https://gigaom.com/2012/07/30/netflix-amara-closed-captions-crowdsourcing/

the employee of crowdsourcing platforms like Amazon Mechanical Turk (AMT). Crowdworkers are either considered self-employed, freelancers, or are hired as independent contractors by third party vendors to work on private crowdsourcing platforms, internal to and used by most large tech companies.[6] By proxy, this means crowdworkers labor for multinational corporations with billions of dollars in revenues for pennies at a time, and labor at their own risk, without the affordance of job protections even the lowest paying occupations are expected to have.

Crowdwork is not legally defined or acknowledged by the U.S. Department of Labor or the U.S. Bureau of Labor Statistics, the government agency charged with defining occupational and employment categories.[7] Yet, federal government departments like the Central Intelligence Agency, the U.S. Department of Justice, DARPA, and the U.S. Army Research Lab[8] increasingly rely on crowdsourcing for translation, transcription, and other undisclosed activities. Crowdsourcing is also a growing segment of local and state government efforts to create data-driven and "smart" systems. Local police forces use crowdsourcing to identify license plates of traffic violators and send tickets by mail instead of paying police officers to stop people on the road. Hundreds of thousands of hours of surveillance tapes are outsourced to the crowd for similar purposes.

Despite this legal gray area, Amazon.com launched the first well-known, public crowdsourcing platform Amazon Mechanical Turk (AMT), on November 2, 2005, to streamline internal systems for

---

[6] Securities and Exchange Commission (S.E.C.) data available upon request from authors.

[7] Former Deputy Secretary of the U.S. Department of Labor, Seth Harris, provides an astute overview of the broader contracting issues with the "gig economy" on his personal blog post. He aptly notes that "The federal government's statistical agencies haven't counted independent contractors since 2005 when they found only 10 million such workers along with another 2.4 million 'on-call' workers." Harris further cites that "the Bureau of Labor Statistics counted more than 14 million workers identifying as self-employed in 2009." Please see: http://sethdharris.tumblr.com/post/99061916410/the-uberification-of-the-american-labor-force. Harris links to a BLS report citing the number of self-employed, please see: Hipple, S (September 2010) Self-Employment in the United States. *The Monthly Labor Review*: the U.S. Bureau of Labor Statistics (BLS). Available online: http://www.bls.gov/opub/mlr/2010/09/art2full.pdf

[8] Please see: Amazon Mechanical Turk | Requester, Case studies. Available online: https://requester.mturk.com/case_studies

its growing online consumer marketplace. AMT is an online labor market strictly speaking (Horton, 2010a), and defines itself as "a marketplace for work that requires human intelligence."[9] Early advocates and developers of crowdsourcing platforms argued 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 identifies a peculiar issue: work quality expressed as job performance among workers is not responsive to changes in price (Mason and Watts, 2010; Shaw, Horton and Chen, 2011; Buhrmester, Kwang and Gosling, 2011; Rogstadius et al. 2011). In effect, paying crowdworkers higher wages does not lead to higher quality work, as economic theory would predict. This observation leads some to believe crowdsourcing platforms like Amazon Mechanical Turk (AMT) are labor markets for lemons. Since quality is unresponsive to changes in price, this implies crowdsourcing is also not skill-neutral in terms of labor quality. This paper examines the different market dynamics that might, unwittingly, lead to unfair compensation of workers and contribute to inefficiencies in the market that, to date, have been attributed to the quality of workers available on crowdsourcing platforms.

The economics of crowdsourcing is dominated by studies evaluating the responsiveness of quantity (the number of workers who do a task) and speed (the rate workers complete tasks) to changes in price (Horton and Chilton, 2010; Mason and Suri, 2012). The behavior of work quality in relation to price is given little attention, especially in terms of labor market supply. It is important to study work quality, price, and labor market supply because crowdsourcing labor markets require a significant "crowd" of laborers at very little cost (price or wage), and firms or employers (requesters) require the satisfactory completion of work in order to turn what workers produce into a sellable commodity.

---

[9] Amazon Mechanical Turk (AMT), FAQ, online: https://www.mturk.com/mturk/help?helpPage=overview

Computer scientists, engineers, and economists care to study how the quantity of workers doing a particular task increases or decreases in response to price, because crowdsourcing, by definition, is a type of labor reliant on rapidly amassing a "crowd" labor force on demand. Researchers study how quickly crowdworkers complete tasks in response to changes in price, because people and firms (requesters) who post tasks to platforms want those tasks done expediently. Requesters need to consider the quality of work crowdworkers do, because they cannot turn subpar labor into monetizable results. When quality is subpar, total labor costs also tend to be higher, because employers receive less for the wages they pay, and for this reason, they tend to pay workers lower wages.

Economists who study price in relation to quantity, rates of production (speed), and quality of work are concerned with labor market supply, and how wages affect the number of persons willing to sell their labor to a particular firm. Market competition is one primary factor labor economists must evaluate in order to understand how labor market supply and wages interact. Labor market competition is a critical factor, because the degree of market competitiveness largely determines when and at what rate workers will participate in a labor market when employers or firms change wage rates (prices). A labor market with only one firm is not competitive, for example, because it is the only employer purchasing labor (Robinson, 1959; Manning, 2003). For this reason, workers in a labor market with only one firm are less likely to quit their current job if the firm lowers wages, because they have no other employment opportunities (Manning, 2003). This example is extreme, and unlikely because most labor markets have more than one employer. However, employers and firms are known to collude, and act like a single employer rather than competing entities. In these cases, similar labor market dynamics are observed as with the extreme case of a labor market with only one firm (Manning, 2003). Uncompetitive situations like these describe monopsony and oligopsony, or labor markets characterized by imperfect competition (Robinson, 1959).

Imperfect competition is harmful to workers, because it limits their bargaining power. Bargaining power or the ability to negotiate with employers and firms over wages is curtailed, because fewer employment options are available to workers, since the number of firms competing with each other is less than what would occur under more competitive circumstances. When imperfect competition—monopsony or oligopsony—prevails, employers and firms are said to have market power (Manning, 2003). Conceivably this power grants employers the latitude to behave unscrupulously without the same repercussions economic theory predicts firms would face under perfectly competitive conditions. Imperfect competition also has important repercussions for work quality and the behavior of price in labor markets, which is discussed in latter sections of this paper.

Economics utilizes models to determine the level of competition in a market, and how various factors like quality responds to changes in price. Perfect competition and imperfect competition describe the dominate models economists use to study labor market supply (Robinson, 1959; Dixon, 1987). Perfectly competitive labor market supply assumes employers pay wages according to work quality[10] that reflects employees' job performance (Cahuc and Zylberberg, 2004: 246). This assumption implies that if employers lower or raise wages, the quality of goods and services workers produce, changes proportionately. Crowdsourcing labor markets are not thought to follow this normative economic story. Instead, paying people who crowdsource (known as crowdworkers) marginally higher wages does not clearly produce higher work quality.

Scholars have hypothesized quality is unresponsive to price, because of the low-wage nature of the market. Crowdsourcing labor markets are low-wage by design, because this type of labor is marketed as a virtual, cloud-based version of outsourcing. Outsourcing is when firms hire contract labor to service or provide a business segments previously produced in-house. Firms outsource labor

[10] Sometimes labor economists control for work quality when analyzing wage differentials (see Manning, 2003: 220).

to low-wage markets for the purpose of reducing labor costs. Call centers are a well-known example of outsourcing. Another example is the infamous Nike sweatshops abroad. Notably, in the 1990s, Nike's outsourcing practices sparked the ire of labor and human rights activists around the globe. Some have compared crowdsourcing labor markets like AMT to sweatshops (Horton, 2010b), but the utility of this comparison is questionable for the hundreds of thousands of workers who rely on the income crowdsourcing provides, and equally, for the computer scientists and engineers laboring to make the design of crowdsourcing platforms better for workers.

Outsourcing and crowdsourcing share common features. As mentioned, firms are incentivized to outsource or crowdsource to reduce labor costs (Felstiner, 2011). In order to do so, they seek opportunities to hire contingent labor in low-wage markets, or markets and workplaces where workers are made vulnerable by a lack of legal enforcement and protection. Previously this meant firms in the global north outsourced business segments to the global south, but crowdsourcing allows firms to outsource and hire anywhere workers with a computer and Internet connection are willing to labor for pennies a task.

Analyzing the historical distinctions and connections between these labor practices is beyond the scope of this paper, but the low-wages characteristically paid to workers under these circumstances speak to a narrative about crowdworkers commonly held in some circles. That is, some argue the low-wage nature of commercial crowdsourcing markets attracts 'the worst quality of worker,' because these workers are presumed not hirable for other positions, and are thought unable to produce higher work quality for this reason—no matter the marginal increase in wage rates. Worker characteristics are then said to explain why crowdsourcing platforms cannot improve work quality by increasing wages: these platforms are simply labor markets for lemons. This view is informed and bolstered by the way neoclassical economics thinks about the decisions firms and workers make in regards to the supply of labor. Neoclassical economics often but not always starts from the presumption that labor markets are perfectly competitive (Moore, 1906; Robinson, 1959), and workers chose when to supply their labor and when to consume more leisure time (Camerer, Babcock, Loewenstein and Thaler, 1997). Simply, this translates to the misleading and ill-informed idea that unemployment is partially a choice workers make in response to wages the "free market" sets (Mortensen, 2011).

An alternative interpretation is that competitive models of labor market supply do not provide the best predictions (Robinson, 1959) for pricing, motivations for laboring, or the quality of work produced in some labor markets, including crowdsourcing platforms. In fact, models that assume labor markets are perfectly competitive arguably limit our ability to understand online, crowdsourcing labor because the assumptions made by these models are not met by the pricing and work quality trends observed in commercial crowdsourcing markets. Models of perfect competition also do not permit a comprehensive, and vigorous discussion of inequality in labor markets (Manning, 2003). Since perfect competition assumes unemployment is a choice made by workers (Mortensen, 2011), and wages are set by markets rather than firms, there is not much room for a debate about what the societal implications are when workers are paid below subsistence wages while firms take home revenues in excess of billions per year. If the free market determines what wages firms will offer, and workers decide whether or not to work in relation, what agency is left to civic society to ask if wealth and resources should be distributed alternatively? What accountability is left to hold governments responsible for enforcing the economic distribution of resources a society considers fair? (see Piketty)

In order to ask these questions, and interpret their implications for low-wage markets like Amazon Mechanical Turk (AMT), alternative models of labor market supply beg consideration. This paper looks to models of imperfect competition for this reason. Specifically, the dynamic model of monopsony articulated by Alan Manning (2003) is used as a vehicle to ask whether requesters who pay crowdworkers for their labor on Amazon Mechanical Turk hold sufficient market power to determine the nature of employment relations on the platform. From observations made and datasets collected directly from the AMT market, this paper concludes that monopsony (power) best describes labor market supply, employment relations, and the outcomes for work quality on crowdsourcing platforms. This conclusion is articulated by data gathered from a longitudinal, multi-modal study[11] of Amazon Mechanical Turk ("AMT") that we analyzed to test how competitive labor market supply is, and whether monopsony prevails, allowing requesters a consequential degree of market power. Again, based on our study's findings to-date, we believe the monopsony framework (see Manning, 2003) offers a more comprehensive explanation of observed trends, particularly why varying wages does not alter the quality of job performance.

Before proceeding with study results, it is important to consider that evaluating model assumptions associated with traditional, offline labor markets in new, online labor exchanges allows us to rethink the dominant assumptions of labor market supply theory under new employment contexts. Determining which economic model best describes crowd labor is a critical first step to informing our understanding of crowdsourcing markets (Horton, 2010a). Understanding how work quality, among other factors, functions in relation to price, and whether requesters hold substantive monopsony (market) power, will enable us so to design better tools and technologies future digital workforces and employers will use to exchange labor for pay.

More broadly, questioning the prevalence of monopsony power is a vehicle to explore how empirical studies of labor market supply (see Ashenfelter, Farber and Ransom, 2008; Manning, 2003: 360-367; Schmitt, 2013) could inform governments and other legal entities about establishing the labor policies that will shape the future of a globally distributed, crowdsourced workforce. Reckoning with monopsony power in online commercial crowdsourcing labor markets, we argue, could prove critical to deciding how we think about this growing sector and the public policies this space requires.

---

[11] See Microsoft Research (ongoing) *Face in the Crowd*: http://research.microsoft.com/en-us/projects/crowdwork/

Finally, while we argue that imperfect labor market competition best describes crowdsourcing sites facilitating the exchange of labor for pay, these models are not a perfect fit. A noteworthy takeaway from our analyses is that new economic models are required to fully understand online labor markets. In the concluding section, we discuss the shortcomings of models for imperfect competition, and provide recommendations for future directions in empirical studies.

Section II: Research Methodology and Data Sources

We draw on several data sets to argue that crowdsourcing labor markets exhibit features of monopsony power. The data sets include: (1) responses to a survey posted to the Amazon Mechanical Turk (AMT) platform between July 2013 and July 2014 (2) ethnographic data collected from 48 interviews and participant observations conducted in person from September 2013 to July 2014, (3) results from a geographic mapping task (HIT) posted to AMT, and (4) data gathered from SEC filings and other regulatory agencies.

Our analysis focuses on the AMT labor market, because AMT is "the crowdsourcing site with one of the largest subject pools" (Mason and Suri, 2012). 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 set wage. Wage rates are unilaterally determined by requesters. Tasks posted to AMT are called Human Intelligent Tasks or HITs. A HIT group consists of similar micro-tasks or HITs posted by the same requester. As we argue above, the growth of the information services economy increasingly depends on the types of tasks posted as HITs on AMT. Small firms digitalizing business cards collected at a marketing event, large technology companies farming out beta-testing of a new software product, and researchers running behavioral experiments online rather than a college Psychology 101 course all rely on crowdsourcing marketplaces like AMT.

The general workflow on AMT is as follows. A requester posts a group of HITs. Workers do those HITs. Requesters then review the work, accept the good work and reject any poor quality work. Requesters deliver payment only for the work that they deem acceptable. The overall fraction of HITs that a worker has had approved over his or her lifetime is that worker's approval rating. A worker's approval rating serves as type of reputation score that determines the jobs they will be able to access in the AMT Marketplace. Finally, 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 high job performance reputation. Amazon sets the parameters for establishing all accounts and standards for performance and reputation, but Amazon does not make these parameters and standards public or transparent to requesters or workers.

Longitudinal Survey of Amazon Mechanical Turk (AMT)

Our survey asked respondents doing paid crowdsourcing work ("crowdwork") on AMT a range of questions, from inquiries about basic demographics to specifics concerning computer literacy and Internet skills. A set of questions focused on assessing the time and effort spent finding tasks, motivations for crowdsourcing, language skills, estimated yearly income, and venues to find tasks online, among other questions. Early ethnographic evidence suggested that workers focus on finding and doing specific types of tasks on AMT to optimize their time on the platform, specializing in certain task types that fit their availability and skill sets. Merely posting the survey on AMT, as is commonly done by those conducting surveys on crowdsourcing platforms, may over-sample workers who typically do surveys as tasks for work. Thus, in addition to posting the survey to AMT, we also embedded the survey into separate image-labeling tasks and email classification tasks. After a worker did 10 email classifications, for example, a link appeared asking if they would like to do our survey for additional pay. Since our survey also served as a vehicle to recruit interview participants, this methodological innovation allowed us to reach workers who might not typically do surveys on AMT. We obtained a total of 317 responses from the AMT survey, to date, including 180 completed surveys from people living in the United States, and 137 from people living in India.

Ethnographic Fieldwork on Amazon Mechanical Turk (AMT) Workers

We integrate qualitative data, gathered from 9 months of ethnographic fieldwork in India, for this paper. To date, we have completed 48 in-person, open-ended, semi-structured interviews, and hundreds of hours of follow up interviews and observations with research participants met through the AMT surveys, worker referrals, and online contacts made on worker discussion forums.

Systems-Level Measurements of Crowdsourcing Platforms

Complementing the AMT survey and ethnographic interviews are data from systems-level measurements which we obtained from posting tasks posted to AMT. In particular, we draw heavily on a simple geographic mapping task, posted to AMT, which paid participants to self-report their geo-location and asked how they found out about the task itself. More specifically, workers were shown a map of the world (via the Bing maps API) and asked to place a pin where they are located. After clicking save, they were shown a map of the 500 previously reported pins, each of which was randomly perturbed to protect worker privacy. This was done to display the global community of crowdworkers to fellow participants completing mapping task. The HIT also asked workers to identify how they found out about the HIT. This HIT ran for 5 weeks and collected 4,856 pins. Notably, AMT does not publish statistics about the people who use and work on its platform. Systems-level measurements, like our mapping task, allow us to extrapolate the geographic distribution of workers on AMT, a key data point for our study.[12]

Government & Corporate Financial Records

Lastly, we looked at data from third-party companies that largely depend on crowdsourcing platforms like AMT for the bulk of their revenue stream. This data includes primary source documents obtained from public filings, as submitted to regulatory agencies like Securities and Exchange Commission (SEC), and equivalent international regulatory agencies. We obtained investor relation and financial information from company websites. These data help us more fully develop a picture of crowdsourcing labor markets as entwined employment landscapes and ecosystems rather than separate job opportunities operating independently and equally available to all workers.

Section III: Key Definitions

Market Definitions

We define the crowdsourcing labor market as the outsourcing of jobs to, "an undefined group of people in the form of an open call" that is issued to online, commercial labor exchanges (Mason and Suri, 2012). Notably, commercial crowdsourcing platforms like Amazon Mechanical Turk ("AMT"), "serve as the meeting place and market" where micro-task labor is exchanged online for pay (Mason and Suri, 2012). Crowdsourcing platforms serve as both the *location* and *marketplace* for labor exchange. This is a critical distinction because, legally speaking, crowdsourcing sites like AMT have no employment relationship with the people who exchange labor on their platform. AMT is different

[12] Literature reviewing experiments in labor economics likewise emphasize the importance of examining the "representativeness of the laboratory population relative to the field population of interest." The literature suggests "conducting lab experiments directly on the field population […] and studying the selection process itself." See: Ashenfelter, O., and Card, D. (2010). "Lab Labor: what can labor economists learn from the lab?" *Handbook of Labor Economics*, 4(A), North Holland/Elsevier: London.

from employers who typically hire ("buy") labor in offline labor markets. Instead, Amazon's primary role is to define the online boundaries and *geography* of the online labor market. This role, however, does substantially effect labor market outcomes. Amazon's user agreement, for example, determines who and how people may participate in the market. 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. Finally, 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, 2010a).

Model Definitions

The dominant model of labor market supply presumes conditions that allow for perfect competition (Robinson, 1959). In this section, we first describe this model's core assumptions and components. We then describe alternative models of imperfect competition. The central tenet of competitive labor market supply theory is the *law of one price* (Stigler and Sherwin, 1985). The *law of one price* assumes that firms will pay all workers who are the same in terms of their ability, skills and occupation the same wage (Cahuc and Zylberberg, 2004: 246). The *law of one price* only applies to workers who are employed to do the same job. This rule is not expected to hold across occupational groups. So, for example, the *law of one price* assumes a hospital employing doctors will pay doctors a wage rate that is different than the rate they offer to nurses. However, doctors employed in the same specialization, who have equivalent education and work experience, should be offered identical wages by the hospital. The *law of one price* mitigates wage dispersion, or the power of firms to pay differential wages to workers whose job performance (productivity) is identical. In other words, the *law of one price*, in principle, shores up the promise of equal pay for equal work in a perfectly competitive labor market.

Competitive models of labor market supply make a number of other important assumptions. For example, *ex post* wage bargaining where employers and employees negotiate wages after meeting face to face is considered normative.[13] Also, price discrimination, based on non-productive worker attributes, should not occur (Manning, 2003).[14] This means employers do not offer different pay to similar workers doing the same job, because of their socio-demographic characteristics like age, gender, ethnicity, nationality, and religious affiliation. As Cahuc and Zylberberg (2004: 261) strongly emphasize, "employer discrimination cannot exist under perfect competition" Intertemporal substitution is also imagined to directly affect labor market supply, meaning that workers may freely choose when to work and when not to work in direct response to changes in wage rates (Robinson, 1959; Camerer, Babcock, Loewenstein and Thaler, 1997).[15] Finally, in some models, competition also depends on the ability of "workers" to "sort across tasks on the basis of their comparative advantage"[16] (Costinot, Arnaud and Vogel, 2010).

When the *law of one price* is shown inconsistent with labor market dynamics, and labor market information is not perfect, then we must consider models of imperfect competition (Robinson, 1959; Manning, 2003). And emphatically, if the nature of a labor market calls upon us to ask questions about inequality, then alternative models to perfect competition are required (Manning, 2003).

---

[13] Employers and workers meet a priori to wage rate determination, allowing for a degree of wage negotiation.

[14] McAfee, Mialon and Mialon (2006) argue price discrimination does not strongly correlate with market power, but their study concerns consumer markets. When labor economists speak of price discrimination they refer to wage differentials that are not explained by transportation costs or compensating wages (hazard pay for dangerous work). By law, at least in the United States, employers are not allowed to offer workers different wages based upon socio-demographic characteristics. Thus, while McAfee, Mialon and Mialon are concerned with the power of monopolists to price discriminate in consumer markets with regard to anti-trust laws, labor economists are concerned with price (wage) discrimination against workers with regard to employment or labor laws upheld by the Equal Employment Opportunity Commission (EEOC), the U.S. Department of Labor (DOL), and the court systems.

[15] Economists usually describe this decision as the trade-off between consumption and leisure. This means consumers freely choose to participate in the labor market, and trade-off consumption and leisure in relation to wage rates.

[16] David Autor explains, "Comparative advantage in production means that the factor with the lowest economic cost of performing a task is assigned that task," while, "economic costs in turn reflects both a factor's technological capability and its opportunity cost" (see Autor, 2013).

Monopsony is a model used to analyze imperfect competition, and permits an analysis of social equality, especially in regards to work (Manning, 2003). Monopsony traditionally describes a situation where employers set wage rates, instead of the market, and bargaining between employers and employees does not occur (Manning, 2003). Models of imperfect competition, such as monopsony, differ from competitive models of labor market supply in a number of important ways. First, employers have wage-setting power and *ex ante* wage posting occurs, that is, employers set wage rates before meeting workers (Manning, 2003). Second, substantive labor market frictions exist, such as asymmetric or imperfect labor market information (Robinson, 1959; Isard, 1977; Manning, 2003; Ehrenberg and Smith, 2012). Imperfect labor market information means employers and workers do not have the information they require to make optimal decisions about hiring workers and accepting jobs online. This often leads to poor employer-employee matches and results in the provision of low quality goods and services (Priest, 2008). An extreme example of this would be: a hospital hires a nurse to do the job of a cardiac surgeon, because the hospital did not have sufficient information to know that the nurse was not a qualified cardiac surgeon, while at the same time, the nurse did not have sufficient information to understand they were being hired to do the job of a cardiac surgeon. Outcomes from such an employer-employee match would obviously not be ideal (for the hospital, nurse, or unfortunate patient).

In labor markets with imperfect competition, workers with different skills, education, and work experience are not necessarily paid according to their degree of productivity, or according to their job performance (Manning, 2003). Alternatively, this means different workers, with ranging abilities, are paid the same rate by firms. This fact directly counters standard models of perfect competition, which assume only workers with identical job performance or of equivalent ability, will be paid the same rate (Manning, 2003).[17]

Oligopsony describes a situation where the number of employers in a labor market is highly concentrated to the extent that competition is imperfect (Manning, 2003). Moving forward, we use the terms monopsony and oligopsony interchangeably for a number of reasons. First, monopsony in the strict sense, refers to a situation where only one employer buys labor from numerous workers ("sellers of labor"). However, as Alan Manning aptly notes, and we agree, the term monopsony is not meant to be interpreted literally, as mono or one, employer (Manning, 2003: 16; 360). Instead, the model for monopsony should be viewed conceptually along a spectrum, where the concentration of employers in a particular labor market varies from one or more (monopsony)[18] to a few or more (oligopsony) to an infinite number (perfect competition).

Section IV: Empirical Evidence

In this section we discuss evidence of monopsony in the Amazon Mechanical Turk ("AMT") labor market. We do so by evaluating the competitiveness of wage structures on AMT and the extent that labor market frictions prevail. Elements of wage structure competitiveness and market frictions serve as diagnostics to test competitive as opposed to monopsonistic dynamics of crowdsourcing labor markets.

---

[17] A notable consequence of paying the same wage rate to all workers is that firms lose important sources of revenue, because wages are not utilized as a mechanism to sort workers according to their ability, or job performance. Offering different wages to different workers typically serves as an incentive to retain desirable workers, and increase the quality of goods and services produced by these workers (Manning, 2003).

[18] Employers sometimes collude, and act together rather than behave as individual, autonomous employers. In this sense, colluding firms have monopsony power, because they limit the degree of competition in the labor market. A recent example of this is the Google/Apple/Intel wage-setting scandal, where the firms agreed not to hire each other's' employees in order to keep wages lower than they would be if perfect competition prevailed. See: Konczal M (February 14, 2014) The Silicon Valley Labor Scandals Prove Minimum Wage Hikes Don't Cost Jobs. *The New Republic*: http://www.newrepublic.com/article/116608/silicon-valley-labor-scandals-prove-minimum-wage-hikes-dont-cost-jobs

Our assessment of AMT wage structures is broken down by: (1) how wages are determined (*ex ante* versus *ex post* wage posting), (2) whether wage bargaining occurs, (3) if requesters pay the same wage to all workers who have equivalent job performance, work experience, and education, and finally (4) if requesters have the ability to price (wage) discriminate against workers.

We discuss labor market frictions in regards to: (1) whether information is perfect, and (2) if labor market entry is free or costly. Additional examples of how models for perfect competition do not accurately describe AMT labor market dynamics, include: (1) workers likely use of daily income targets to determine when to stop crowdsourcing, as opposed to the use of intertemporal substitution, or the trade-off between work and leisure in relation to changes in wage rates (Mason and Watts, 2009), and (2) the high degree of concentration among requesters posting tasks to the AMT platform, where the concentration of requesters indicates a lack of competition in the AMT marketplace (Ipeirotis, 2010).

AMT Wage Structures

Some of our strongest evidence of monopsony power in AMT's Marketplace come from the design of the site itself (Ipeirotis, 2010; Silberman, Ross, Irani and Tomlinson, 2010; Khanna, Ratan, Davis and Thies, 2010). Requesters posting tasks to the AMT market list the wage for each task (HIT) made available to crowdworkers before any work is done. Figure 1 shows an image of the AMT marketplace, which clearly demonstrates how requesters unilaterally set wages *ex ante*. This meets the first basic criteria for monopsony power among employers as economic theory assumes only employers in labor markets with imperfect competition are able to unilaterally set rates (Manning, 2003). That requesters set wage rates, and do so uniformly for each individual HIT, prior to meeting workers (ex-ante), implies crowdworkers have little to no bargaining power. Models for imperfect competition likewise assume workers have little to no bargaining power when employers have monopsony power, because employers rather than the market set prices (Manning, 2003; Cahuc and Zylberberg, 2004; Ashenfelter, Farber and Ransom, 2008; Staiger, Spetz and Phibbs, 2008; Ashenfelter, Farber and Ransom, 2010). This means workers have no alternative options for earning income apart from what employers pre-determinately offer them.

Figure 1 Screenshot of the Amazon Mechanical Turk marketplace

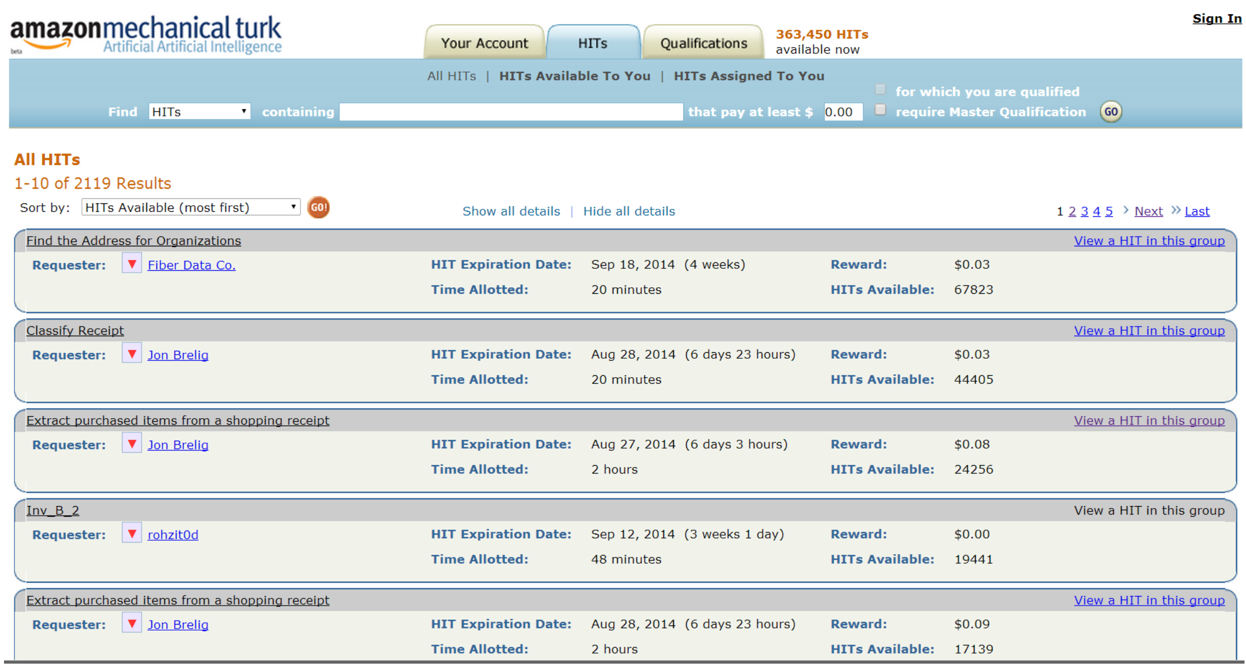

Informal interviews with requesters, from university-based researchers learning how to post work on AMT to those software engineers A/B testing new user interfaces, indicate that it is a common practice for new requesters to ask their more experienced colleagues the prices they should set for their tasks. This is an understandable action, considering there are no obvious channels of communication that would allow a requester to "price check" a task with experienced workers. However, this observation strongly suggests that the normative wage structures on AMT afford requesters enough monopsony (market) power to determine wages in a non-competitive fashion, signaling that the requester, not the market, is the price-setter in the AMT labor market. Again, because wages are predeterminately set by requesters, crowdworkers have no alternative to accepting the wages requesters offer, except for unemployment or earning no crowdsourcing income from the platform. While crowdsourcing income may seem an insignificant amount to some, our survey data indicates that at least 14% of crowdworkers in the United States are living under the federal poverty line, and a majority report that they rely on crowdsourcing money to supplement their income, if they even have other source income (many do not).

That requesters unilaterally determine wages before meeting crowdworkers, and often share or determine wages in coordination with their peer requesters, also means requesters are not using wage rates as a mechanism to compete with other requesters. Normative economic theory expects firms or employers to use wages as a competitive mechanism to attract desirable workers, as the model of perfect competition presumes they would do. That requesters do not pay crowdworkers according to their job performance for each individual HIT, and the investments they've made in their education and job skills, is further evidence of the non-competitive nature of wage structures on AMT. Crowdworkers who are, likely, qualitatively different in terms of their skills, work experience, and education—what economists call human capital—are paid the same rate when doing the same task as there is no transparent mechanism for sorting workers according to their human capital. This clearly violates the *law of one price*, which assumes ***only*** workers with a similar degree of human capital will receive the same wage rate for doing the same job (Robinson, 1959; Manning, 2003). This is, by design, impossible on platforms like AMT.

According to standard economic theory, for competition to prevail, the *law of one price* must hold for a given labor market. Labor economist, David Autor (2013), extends this to the allocation of tasks, and argues that, "Competitive labor markets require that the Law of One Price" must also apply "for skill[s]." Despite this, the law of one price is the exception not the rule (i.e., the fallacy of) Amazon Mechanical Turk (AMT).

Evidence of Price Discrimination

Price discrimination occurs when an employer or firm pays different wages to workers based on their socio-demographic characteristics, such as race, gender, age, sexual orientation, or religious affiliation (Varian, 1987; Manning, 2003; Cahuc and Zylberberg, 2004). Problematically, it is often hard to find evidence of price discrimination in traditional labor markets, because private firms usually keep the wages they offer to workers secret.[19] Crowdwork poses a particularly challenging setting for teasing out the presence and impact of price discrimination.

For AMT, evidence of price discrimination is equally hard to find, because the prejudices of requesters are not made immediately apparent through the wages they post. However, this does not mean requesters do not price discriminate. Requesters, for example, can bar certain crowdworkers from doing HITs based on their geographic location. In practice, this could reflect an actual need (i.e., targeting a specific population for a survey) but it could also signal employer bias (i.e., a requester discriminating against crowdworkers from India based upon unfounded prejudices). In fact, it is common practice among requesters to restrict their HITs to U.S.-based workers because U.S.-based workers are perceived to do higher quality work on average (Chandler and Kapelner, 2013). Thus high quality India-based workers face wage discrimination simply because they live in India. That requesters can, and do, discriminate against workers based on socio-demographic identities suggests that the AMT labor market is not perfectly competitive, and that requesters hold a certain amount of monopsony power in the market, because they can price discriminate.

---

[19] A well-known example of this is the Lilly Ledbetter case. Ledbetter, who was employed at Goodyear for more than 20 years, sued Goodyear after discovering that she made less than her male colleagues. Ledbetter argued her case all the way to the Supreme Court, but lost due to statutory limitations on the amount of time an employee is legally given to charge their employer with discriminatory wage practices. Fortunately, in 2009, the U.S. Congress enacted the Lilly Ledbetter Act - overturning the statutory limits that prevented Ledbetter from winning her case – and clearing the way for workers to seek redress against unscrupulous employers. See: The National Women's Law Center (January 29, 2013). Lilly Ledbetter Fair Pay Act: http://www.nwlc.org/resource/lilly-ledbetter-fair-pay-act-0

AMT Labor Market Frictions

Market frictions refer to different "transaction costs" that people and firms incur when participating in a given labor market (Coase, 1937; Williamson, 1979). Here, we focus on two types of transaction costs, namely, the cost of: (1) labor market information, and (2) market entry. The extent of these costs—and, by definition, the market frictions (The Royal Swedish Academy of Science, 2010) that they generate—influences how competitive a labor market is. In perfectly competitive labor markets, economists usually assume employers and employees have perfect information about critical factors that determine the costs and benefits of employment, such as the availability or location of jobs, as well as "the actual characteristics of the jobs available" at any given time (Cahuc and Zylberberg, 2004: 518). A frictionless market is one in which "the product or service traded is standardized, and all properties are known to the buyer as well as the seller by assumption" (Mortensen, 2011).

Imperfect Information

When labor market information is costly rather than costless, it is considered "imperfect," and a source of labor market frictions. This violates a primary assumption of perfect competition in that such competition, in a neoclassical model of markets, depends on all parties having access to the information they need to make rational decisions. The AMT's platform, by design, skews task and reputation information to favor requesters at the disadvantage of crowdworkers, and epitomizes imperfect information in a labor market for this reason.

As mentioned previously, Amazon integrates a worker reputation system into the platform in the form of worker approval ratings. On AMT, if a crowdworker's reputation receives a demerit, for example, because a requester rejects that worker's HIT, then the crowdworker will have fewer HITs made available to them in the future. "The uncertainty associated with HIT payment complicates" crowdworkers' "work and reduces their effective wage" (Silberman, Ross, Irani and Tomlinson, 2010).

Workers' reputations also seem to determine their qualification for a Master's Account.[20] This feasibly reduces a crowdworker's total potential earnings, where the extent of these losses are unknown. Moreover, even with the approval rating, it is difficult to observe worker quality on a specific task *a priori* as the rating does not necessarily indicate how a worker would perform on a new task. This forecloses the opportunity for a requester to use information about a worker's individual human capital to determine whether or not to hire them for a specific task.

More importantly, Amazon's reputation system is one-sided in that it only signals to requesters how well crowdworkers have performed in the past. And it does not indicate to crowdworkers how well requesters have behaved as employers at all. 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. John Horton (2010b) reports that many critics have scrutinized the number of unscrupulous requesters on crowdsourcing platforms. In fact, as analyses of online crowdsourcing forums indicate, unscrupulous requesters frequent the AMT market often, sometimes to commit cybercrime, or arbitrarily reject work and deny payment to crowdworkers (Silberman, Ross, Irani and Tomlinson, 2010). Crowdworkers have no mechanism to remedy the wrongdoings of requesters directly through a reputation system on the platform. So while requesters can punish crowdworkers they deem bad actors by withholding payments or reporting them to AMT, crowdworkers have no mechanisms to do the same. For crowdworkers, this fact makes the cost of finding good HITs on AMT higher than it would be if perfect information about requesters was publicly available.

[20] Again, as noted above, Amazon has not publicly released the specific guidelines they use to allocate Masters Accounts. But, based on discussions among self-identified Master Account holders on a worker-based discussion group examined in larger study, workers need to have a combination of high approval ratings and HIT count (number of tasks completed successfully) to receive this special Account status.

Tellingly, in response to this asymmetry in information, researchers Silberman and Irani (2013) developed TurkOpticon[21], a browser extension that allows AMT workers to rate requesters and view the ratings of requesters by fellow workers. Our surveys and ethnographic data indicate that workers in our study have widely adopted this tool. Several participants, during interviews, noted that TurkOpticon was one of the first tools that they read about in online worker forums and the one that they adopted early on to more efficiently identify the "good jobs." The widespread use and popularity of TurkOpticon speaks to the magnitude of market inefficiency that imperfect information about Amazon Mechanical Turk (AMT) creates.

In recent years, online forums frequented by crowdworkers have noticeably proliferated. Crowdworkers use online forums to share information about the quality of HITs available on AMT and swap recommendations for the requesters who post them. Notably, online forums provide labor market information that is external to the AMT market environment (platform). Economic theories of perfect competition assume perfect information about "exchange in a centralized market" is available from internal sources, and that "information about the goods and services traded as well as the price" is completely known to all those participating (Mortensen, 2011). AMT clearly does not provide much information to requesters or crowdworkers, and where Amazon does institute mechanisms to distribute information, it advantages requesters at the expense of crowdworkers.

A critical point is that AMT does not post labor market information directly on the platform, and for this reason, crowdworkers spend substantial amounts of time searching for information on the Internet (Yuen, King and Leung, 2012).[22] Most economists will agree that time spent searching for work is a cost job seekers bear in order to secure employment. Normally wages and salaries are considered the return workers receive for investments made searching for jobs (Diamond, 2011;

[21] See *Turkopticon*: http://turkopticon.ucsd.edu/
[22] For a discussion of increased income from better labor market information, please see: Argrawal, Horton, Lacetera, et. al., 2013

Mortensen, 2011; Pissarides, 2011). However, when jobs are broken down into piecemeal tasks, which pay pennies a piece to workers who have to do hundreds of them an hour to make less than the U.S federal minimum wage, then the return wages are thought to provide workers becomes highly questionable.

In order to assess the value of online forums as a resource offsetting the lack of perfect information in the AMT Marketplace, we posted an experimental mapping HIT to the AMT platform. Specifically, we asked respondents to identify their current geographic location and tell us how they found our mapping HIT. As Figure 3 demonstrates, the majority of traffic to the HIT came from online forums opposed to general searches conducted on the AMT platform itself. If the AMT Marketplace provided sufficient information, we would expect the opposite of what we found through our mapping HIT. The fact that tens of thousands of registered users generate traffic on multiple forums dedicated to finding good AMT HITs indicates the scale of this market inefficiency. As such, online crowdsourcing forums offer further evidence of the imperfect nature of AMT labor market information. It is important to emphasize that this forum activity does not indicate that the transaction costs associated with finding HITs on AMT are necessarily reduced by forum use. This is because forums are another search mechanism. Data about forum traffic therefore can only illustrate some of the search costs that are incurred by crowdworkers. It is critical to understand that workers' search efforts are far from free. It is a cost or rent borne by those who are actively looking to find decent work online. These costs also indicate that information about the AMT labor market is imperfect, and the labor market is therefore non-competitive.

Figure 3. Online Forum Use by Crowdworkers to Find HITs on AMT

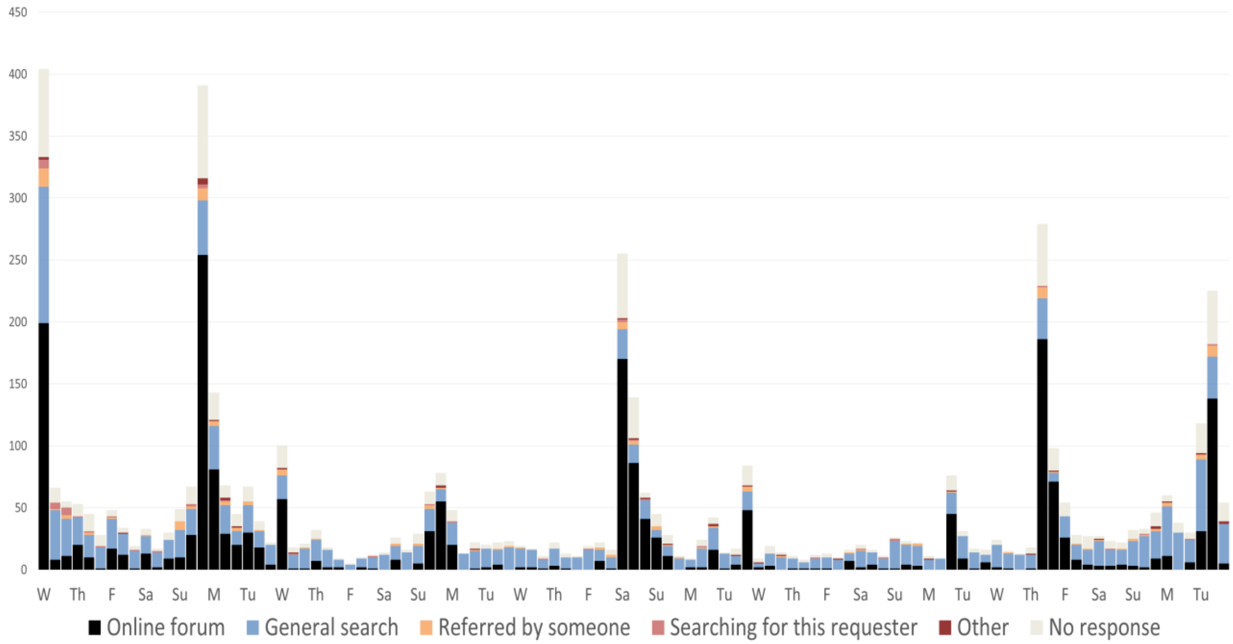

Based on 4856 responses to a task (HIT) posted on AMT, April 23 – May 28, 2014.
Graph credits: Gregory Minton.

Search activity within labor markets that are marred by imperfect information produce another telling outcome: poor employer-employee matches (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. When considering the model of perfect competition, this means that "how wages are set is problematic" and "'market clearing' [establishing the equilibrium between supply and demand] in the usual sense of the term, is impossible" (Mortensen, 2011). 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). This might explain why AMT's reported work quality is not responsive to changes in price (Mason and Watts, 2010).

Market Entry

As mentioned above, if labor market information is imperfect, then market entry cannot be considered unfettered. Searching for people to hire and jobs to accept becomes costly business in an information-poor market. Search activity aside, other transaction costs can increase market costs, working against a model of a perfectly competitive market. One of the most significant market costs and barriers to entry for crowdsourcing market participants is the price of broadband connectivity and distribution of the infrastructure and devices people need to access the Internet. Unfortunately, the crowdsourcing literature often ignores these expenses, which are largely, if not entirely, absorbed by crowdworkers. Figures 4 and 5 use self-reported locations of participants in our mapping HIT to illustrate the effects of entry costs. Specifically, Figures 4 and 5 show the location of AMT crowdworkers, living in the United States and India respectively, by degree of Internet access.[23]

[23] Internet access refers to general broadband coverage for a given county/district, and not the percentage of people who have broadband subscriptions or direct access to Internet services in their home.

Figure 4. U.S. Crowdworker Location by Internet Access (AMT)

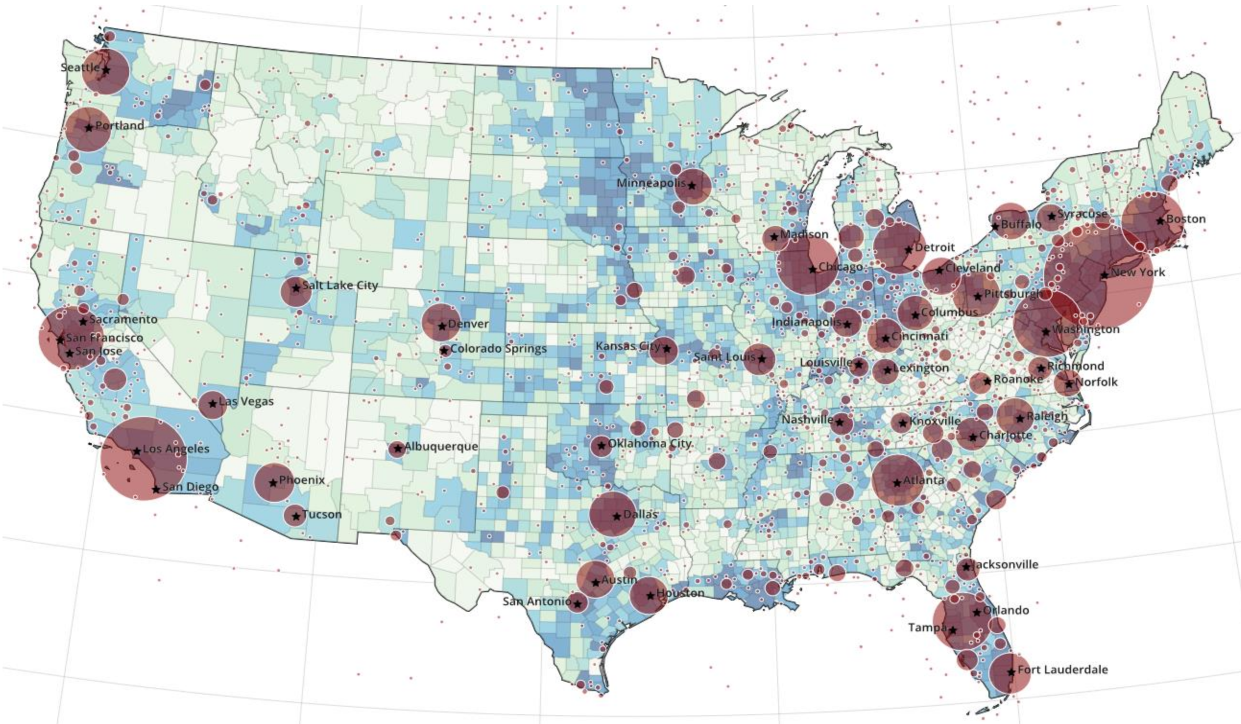

Self-reported locations for about 10,000 participants in a map task (HIT) on AMT.
Coloration of counties/districts is by an estimate of *Internet access*.[24]
Map credits: Gregory Minton.

Figure 4 and 5 demonstrate that crowdworkers tend to live in urban areas that have more Internet access. Intuitively this makes sense. But we concretely link and affirm much of what the literature on information and communication technologies for development tell us about the Digital Divide: the Internet offers economic rewards to people who can afford it (DiMaggio and Bonikowski, 2008).[25]

[24] USA data for estimated Internet access by county/district was scraped from the National Broadband Map website: http://www.broadbandmap.gov
[25] In both figures, self-reported crowdworker locations are demarked in red, and map coloration depicts Internet access at the county/district level.

Figure 5. India Crowdworker Location by Internet Access (AMT)

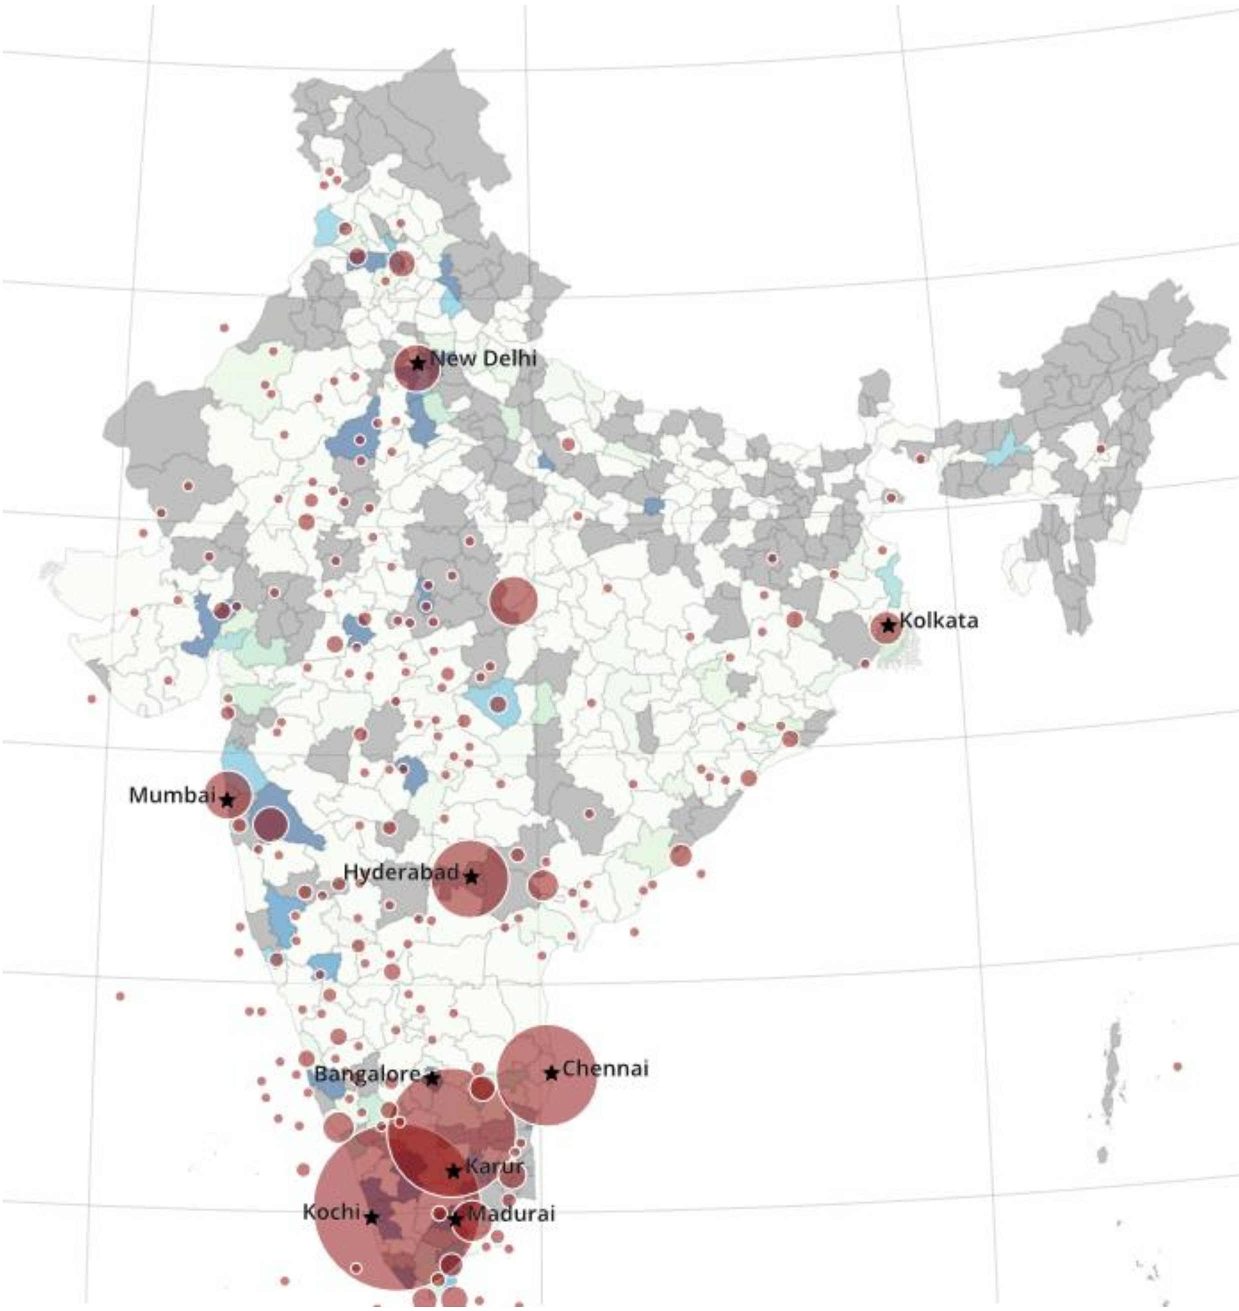

Self-reported locations for about 10,000 participants in a map task (HIT) on AMT.
Coloration of counties/districts is by an estimate of *Internet access*.[26]
Map credits: Gregory Minton.[27]

[26] India data for estimated Internet access by county/district was retrieved from the Indian Census Department and other sources, available here: http://geocommons.com/overlays/18608

Both maps also demonstrate that Internet access limits the mobility of AMT labor market entrants to those who can both afford to connect to it but also, particularly in the case of India, to individuals who live in locations where governments and NGOs have invested in infrastructure—electrical grids, broadband connections, and readily available computer literacy—that make the Internet feasible.[28] Our survey and ethnographic data also suggest that other technical factors might impede crowdworker mobility according to the type of HITs posted to the AMT market. In rural India, for instance, information obtained from participant interviews indicates that some crowdworkers avoid doing tasks that they believe their current Internet capacity (or, speed) will prevent them from doing in the amount of time allotted to complete the task. Delays when downloading or doing tasks online can negatively impact crowdworker job performance reputations, and hence, their income earning potential. It is readily understandable then, why crowdworkers are risk averse to doing tasks that require more bandwidth then they perceive they have. Costly barriers to market entry indicate market frictions exist, and possibly prohibit prices (wages) from effectively incentivizing higher quality work outcomes, because other factors like Internet latency, for example, prevent crowdworkers from performing optimally, as they would if such technical impediments did not exist. Since such market frictions are likely one cause of inefficient outcomes for requesters and crowdworkers, competition is more likely imperfect, and for this reason, requesters could conceivably enough market power to behave in a monopsonistic fashion.

[28] On the impact of the digital divide see, for example, Allen, Steven G. 2001. "Technology and the Wage Structure." Journal of Labor Economics 19:440-483.; Anderson, Ben. 2008. "The Social Impact of Broadband Household Internet Access." Information, Communication & Society 11:5-24; Attewell, Paul. 2001. "The First and Second Digital Divides."Sociology of Education.74:252-259; Autor, David H. 2001. "Wiring the Labor Market." The Journal of Economic Perspectives 15:25-40; Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" The Quarterly Journal of Economics 113:1169-1213; DiMaggio, Paul, Eszter Hargittai, Coral Celeste, and Steve Shafer. 2004. "Digital Inequality: From Unequal Access to Differentiated Use." Pp. 355-400 in Social Inequality, edited by Kathryn Neckerman. New York: Russell Sage; World Internet Project. 2010. "World Internet Project: International Report 2010." Los Angeles, CA: USC Annenberg School Center for the Digital Future.

Non-Causal Evidence of Monopsony Power

Emphatically, the labor market data required to test competing economic theories is often lacking (Robinson, 1959; Manning, 2003). Where labor economists test monopsony power in specific labor markets, they often find data to meet some conditions, while lacking data to meet other defining features of imperfect competition (Manning, 2003). For this reason, many labor economists recommend looking for factors that are a known result of monopsony power, rather than just trying to find the factors said to cause monopsonistic competition (Robinson, 1959; Manning, 2003).

The bellwether test for monopsony is to see if a firm faces an upward-slopping labor supply curve (Robinson, 1959; Manning, 2003). An upward-slopping supply curve signals that labor market supply does not readily change when prices (wage rates) do. If labor market supply behaves differently when supply curves slope upward, it is reasonable to question whether work quality might also respond differently to price under these circumstances (Robinson, 1959). In order to test the responsiveness of labor market supply to price (i.e., the price elasticity of labor market supply), we must be able to obtain information about the number of workers who are paid to do each HIT within a HIT Group. Economists calculate the labor supply to a particular firm through data found in standard metrics like the marginal product of labor, total revenue, marginal revenue of product, and marginal labor costs (Robinson, 1959). Economists use these measurements, in turn, to estimate the price elasticity of labor market supply (Robinson, 1959; Manning, 2003). For AMT, we must at least know how many crowdworkers do each HIT, how many products (total output) results from the HITs crowdworkers do, and the price (wage) crowdworkers receive for the HITs they complete.

Since this information is not available for AMT, we must look for other labor market features for indices of the labor market supply's elasticity. A consequence of inelastic labor market supply is that the number of employers in a labor market will be highly concentrated. For this reason, we looked at the literature and other data to see whether we could say this is true for the number of requesters actively participating in the AMT market.

Requester Concentration on AMT

Although the total number of requesters posting tasks on AMT seems large, seemingly countering our argument that this is a monopsonistic market, studies have found the range of requesters on AMT to be highly concentrated. Ipeirotis (2010) found that the "top requesters" generate "more than 30 percent of the overall activity in the market"; where top requesters comprise a sparse "0.1 percent of […] total requesters" on the AMT. This trend is consistent with the known consequences of monopsonistic competition, as prominently accounted for by the economic literature (Manning, 2003). Ipeirotis (2010) remarks that, "th[is] high concentration is not unusual for any online community." He suggests, further, "There is always a long tail of participants that has significantly lower activity than the top contributors" (Ipeirotis, 2010). This paper has sought to understand why this might be in order to offer an alternative narrative about work quality.

Our narrative about the quality of work on crowdsourcing platforms like Amazon Mechanical Turk (AMT) argues that work quality is not intrinsic to the traits of workers. We do not believe crowdworkers are inherently bad actors, so poorly skilled to be unemployable elsewhere, or are seeking to "game the system." Alternatively, we believe crowdsourcing platforms are, at present, poorly designed labor markets. These markets tend to select for low quality outcomes, not because crowdworkers are flawed, but because our understanding of what people require in order to do crowdsourcing work qualitatively well, is (Robinson, 1959).

Figure 6 shows that the top 1 percent of requesters on AMT post approximately 60 percent of the total rewards available on the AMT marketplace (Ipeirotis, 2010), and 10 percent of all requesters post 90 percent of total rewards (wages) available. Figure 7, "shows how this activity is distributed, according to the value of the HITs posted by each requester" (Ipeirotis, 2010).

Concentration of AMT Requesters by Rewards Posted

Figure 6

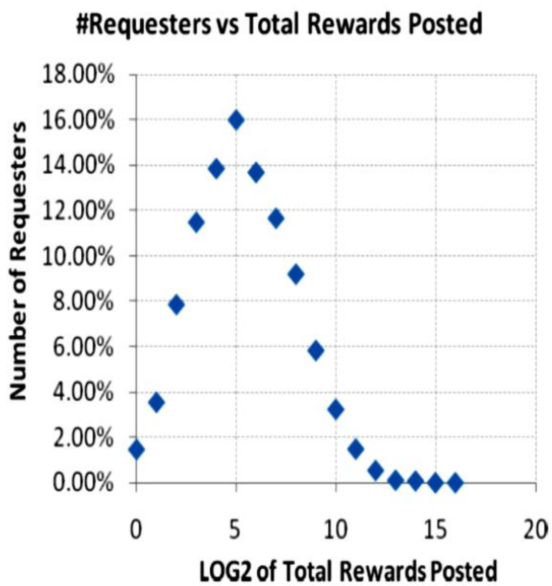

Figure 7

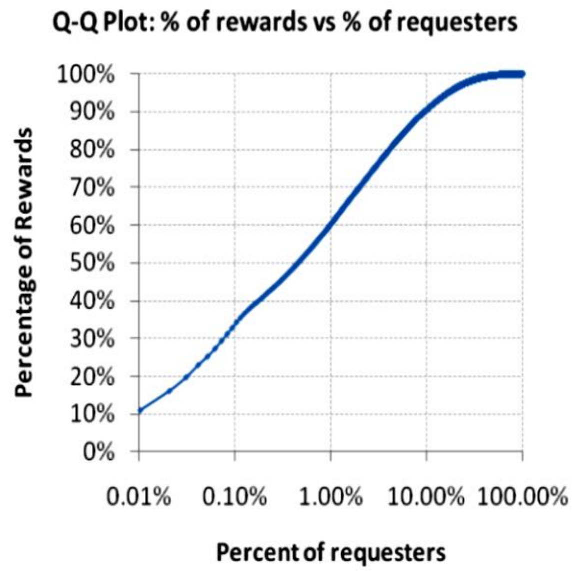

For original graphs, please see: Ipeirotis, P. (2010).

We take the high concentration of both requesters and the available HITs posted, to signal that income-earning opportunities are limited on AMT by the number of top, regularly active requesters, and the rewards for those tasks. In other words, despite the volume of HITs, there are few employers generating the bulk of work. This conclusion is supported by further studies that found that "the same level of activity" or distribution, was "demonstrated by workers" on AMT (Ipeirotis, 2010).

Section V: Consequences of Monopsony for Crowdsourcing Labor Markets

Imperfect competition – monopsony – in crowdsourcing markets like Amazon Mechanical Turk (AMT) has substantive consequences for requesters, crowdworkers, and the engineers and technologists who design these platforms. For requesters and platform designers, prices will not function as expected. We see this when pricing mechanisms continually fail to impact work quality (Mason and Watts, 2010). For crowdworkers, finding and accepting crowdsourcing jobs is both costly and risk-laden. There is never precise information about prices (Silberman, Ross, Irani and Tomlinson, 2010) and crowdworkers cannot accurately predict how requesters will behave. Equally, crowdworkers do not know which jobs are actually legitimate before incurring the cost of testing them out. Combined, these conditions mean that requesters both, ultimately, overpay for the total cost of labor, covering the cost of bad matches in time and poor work quality, while underpaying the crowdworkers who actively participate in and contribute quality work to the labor market. Alternatively, if prices could be accurately predicted, then opportunities to pay qualified crowdworkers better wages would be a more viable option for more requesters. For this reason, we next consider how to remedy monopsony power in online, commercial crowdsourcing markets like Amazon Mechanical Turk (AMT).

Section VI: Remedies for Monopsony Online

A few different technical remedies for monopsony online are possible for crowdsourcing labor markets. The solutions presented are given according to the most probable causes of monopsony power in the AMT market: non-competitive wage structures and labor market frictions, specifically imperfect information.

Solutions for Non-competitive 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. A preliminary example of how this could work is Dynamo.[29] Dynamo is a platform built by researchers for crowdworkers to share information, collaborate, and determine guidelines for academic requesters for setting wages and task design.[30] These guidelines speak specifically to the issue of fair payment for crowdwork on Dynamo. Workers are given a central role (voice) in deciding what constitutes fair pay.

That said, the Dynamo platform does not enable face-to-face or real-time wage bargaining. 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. For example, crowdsourcing platforms could incorporate online chat services directly into the platform, permitting requesters to talk directly to crowdworkers in real-time. Scalability, however, is one mentionable issue with 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, for example, could rapidly circulate opinions among participants. Some researchers have started to explore the role that systems-level visualizations can play in this regard. 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 and reduce market frictions.[31]

[29] See: http://www.wearedynamo.org/
[30] 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.
[31] See: http://www.wearedynamo.org/

Solutions for Imperfect Labor Market Information

While on the surface it might seem like 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, the reputation system is skewed in such a way as to create more information asymmetries than clarity on the platform. 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 some 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, crowdsourcing forums 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.[32] As helpful as these additional tools are, we argue that the dynamics of a healthy market demand that information is readily available to participants, embedded in the infrastructure of the marketplace itself. Online 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.

Vital pieces of information, from requester reputations to a real-time list of jobs and workers in the system, remain scattered across the Internet, requiring crowdwork labor market participants to absorb the costly scavenger hunt to make informed decisions. To correct this problem, communication and reputation systems on platforms like AMT need to be made fairer and more transparent. 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,

[32] See: http://wiki.wearedynamo.org/index.php/Guidelines_for_Academic_Requesters

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 accountably 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 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.

Section VII: Policy Considerations

Technical remedies for monopsony online 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. The legal implications around the exchange of labor on crowdsourcing platforms influence the scope of technical and policy remedies available to practitioners and crowdworkers alike. 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 an employer under the law. Most platform providers will not integrate technical fixes to their APIs that support workers through training, collaboration, and information-sharing, as such enhancements may suggest that the platform curates a workforce. As the class action lawsuit[33] pending against the editorial crowdsourcing site, *Crowdflower*, suggests, we have yet to legally decide what kind of employment crowdwork, technically, is (see NewScientist, February 2013). Additionally, most platforms do not directly set wage rates, and instead leave wage setting to people posting tasks. John Horton (2010a) 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." That the AMT platform is the market creator, it is likewise analogous to a government, the topography of a traditional labor market, or the factory shop of a company. 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 wellbeing and safety of both employers and their employees.

In consideration of these points, we argue that the effort platform providers make to avoid costly legal responsibilities, contributes to the monopsonistic behaviors we observe in crowdsourcing labor markets. Therefore, the following policy fixes are recommended. First, treat crowdsourcing

[33] See Otey v. Crowdlower, Inc. et al: http://law.justia.com/cases/federal/district-courts/california/candce/3:2012cv05524/260287/124

labor markets according to their needs, and not those of traditional, offline markets. Doing so requires policy-makers 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 IRBs. 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.

Section VII: Conclusion

Numerous crowdsourcing studies, particularly on Amazon Mechanical Turk (AMT), offer preliminary evidence that the AMT labor market features monopsonistic behavior among employers. Specifically, observations from past research support key parameters given by models of monopsony, including: (1) employer-based wage setting, (2) ex-ante wage posting, (3) price discrimination, (4) substantive barriers to market entry, and other costly market frictions like asymmetric information problems. For this reason, we evaluated the robustness of models for competitive labor market supply against our proposed alternatives, namely those models describing monopsony power or monopsonistic competition. From our analysis, we conclude that the monopsony framework (see Manning, 2003) offers a more comprehensive, and holistic explanation for the labor market phenomenon, which models of perfect competition have not been able to provide (Robinson, 1959). Approaching the labor market from the lens of imperfect competition affords more opportunities to explore human factors and social forces pertinent to online, crowd labor exchanges that models of perfect competition, by definition, do not permit.

That said, the monopsony framework is not a perfect fit for a number of important reasons. In order to thoroughly, and empirically test for the prevalence of monopsony power in a given labor market, one must be able to measure the (in)elasticity of labor market supply to an individual firm (Manning, 2003). Since we lack data on how many crowdworkers do, and are paid for, each HIT posted to AMT in a HIT Group, we cannot accurately gauge the supply of labor on AMT, and its response to changes in price (wage rates). This speaks to one primary point this paper tries to make: more, and better, labor market data is required by economists in order for robust conclusions to be made about competing economic theories. We strongly concur with Alan Manning, who says, in this regard: "The problem here is the need for good experiments" (Manning, 2003).

Despite noted weaknesses, we still believe the monopsony framework provides a good starting point to question how alternative models for labor market supply might better explain crowdsourcing dynamics than the presumed model of "free market" competition, and the invisible hand that set its wages, a questionable story still pervading the economic literature today (Blundell and MaCurdy, 1998; Manning, 2003; Cahuc and Zylberberg, 2004). We emphasize that as new online markets emerge, the discipline of economics, as a whole, will need to analyze and rethink the standard models it employs and develop new frameworks to understand the new contexts by which people exchange labor online for pay.

References

1) Agrawal A, Horton J, Lacetera N and Lyons E (2013) Digitization and the Contract Labor Market: A Research Agenda. *National Bureau of Economic Research (NBER)*.
2) Ashenfelter O, Farber H and Ransom M (2008) Labor Market Monopsony. *Journal of Labor Economics*, 28(2): 203-210.
3) Ashenfelter O, Farber H and Ransom M (2010) Modern Models of Monopsony in Labor Markets: A Brief Survey. *Institute for the Study of Labor (IZA)*, Discussion Paper No. 4915: Bonn, Germany.
4) Ashenfelter O and Card D (2010) Lab Labor: what can labor economists learn from the lab? *Handbook of Labor Economics*, 4(A), North Holland/Elsevier: London.
5) Autor D (February 2013) The 'task approach' to labor markets: an overview. *Journal of Labor Market Research*.
6) Brabham, D (2013) *Crowdsourcing*. Cambridge, MA: MIT Press.
7) Blundell R and MaCurdy R (1998) Labor Supply: a review of alternative approaches. *Institute for Fiscal Studies (IFS)*, IFS Working Papers, No. W98/18: London.
8) Buhrmester M, Kwang T and Gosling S (2011) Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science* 6(3): 1-5.
9) Cahuc P and Zylberberg A (2004) *Labor Economics*. Cambridge, MA: The MIT Press.
10) Camerer C, Babcock L, Loewenstein G and Thaler R (1997) Labor Supply of New York City Cab Drivers: One Day at A Time. *Quarterly Journal of Economics*: 407-441.
11) Chandler D and Kapelner A (2013) Breaking monotony with meaning: Motivation in crowdsourcing markets. *Journal of Economic Behavior & Organization*, 90: 123-133.
12) Coase R.H (1937) The Nature of the Firm. *Economica*, 4(16): 386-405.

13) Diamond P (June 2011) Unemployment, Vacancies, and Wages. *American Economic Review*, 101: 1045-1072.

14) DiMaggio P and Bonikowski B (2008) Make Money Surfing the Web? The Impact of Internet Use on the Earnings of U.S. Workers. *American Sociological Review*, 73(2).

15) Dixon H (1987) A Simple Model of Imperfect Competition with Walrasian Features. *Oxford Economic Papers* 39: 134-160.

16) Ehrenberg R and Smith R (2012) Frictions on the Employee Side of the Market. In *Modern Labor Economics: theory and public policy* (11th edition), Prentice Hall: 131-145.

17) Felstiner A (2011) Working the Crowd: Employment and Labor Law in the Crowdsourcing Industry. *Berkeley Journal of Employment and Labor Law*, 32(143): 143-204.

18) Golub B and McAfee P.R (2009) Firms, queues, and coffee breaks: a flow model of corporate activity with delays. *Review of Economic Design*.

19) Horton J.J (2013) The Effects of Subsidizing Employer Search. *Social Science Research Network*, available at: http://ssrn.com/abstract=2346486

20) Horton J.J (2010a) Online Labor Markets. A. Saberi (Ed.), WINE 2010, LNCS 6484, pp. 515-522.

21) Horton J.J (2010b) The condition of the Turking class: Are online employers fair and honest? *Economics Letters* 111: 10-12.

22) Horton J.J and Chilton L.B (2010) The Labor Economics of Paid Crowdsourcing. *Proceedings of the 11th ACM Conference on Electronic Commerce (ACM-EC).*

23) Howe J.J (2006). The Rise of Crowdsourcing. *Wired Magazine*, *14,* 1-4.

24) Irani L and Silberman S.M. (2013). Turkopticon: interrupting worker invisibility in amazon mechanical turk. *CHI'13 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems:* 611-620.

25) Isard P (1977) How Far Can We Push the "Law of One Price"? *American Economic Review*, 67(5): 942-948.

26) Khanna S, Ratan A, Davis J and Thies W (2010) Evaluating and Improving the Usability of Mechanical Turk for Low-Income Workers in India. *ACM DEV'10*: London, United Kingdom.

27) Manning A (2003) Monopsony in Motion: Imperfect Competition in Labor Markets. Princeton, NJ: Princeton University Press.

28) Mason W and Suri S (2012) Conducting behavioral research on Amazon's Mechanical Turk. *Behav Res*, 44.

29) Mason W and Watts D (2010) Financial incentives and the performance of crowds. *ACM SigKDD Explorations Newsletter* 11(2): 100-108.

30) McAfee P.R, Mialon H.M and Mialon, S.H (2006) Does large price discrimination imply great market power? *Economic Letters* (92): 360-367.

31) McAfee P.R and Wiseman T (2008) Capacity Choice Counters the Coase Conjecture. *Review of Economic Studies* (75): 317-332.

32) Moore H.L (1907) Paradoxes of Competition. *Quarterly Journal of Economics*, 20(2): 211-230.

33) Mortensen D (June 2011) Markets with Search Friction and the DMP Model. *American Economic Review*, 101: 1073-1091.

34) Moscarini G (2005) Job Matching and the Wage Distribution. *Econometrica*, 73(2): 481-516.

35) Nabarun D, Freifeld C, Brownstein J, et. al. (2013) Crowdsourcing Black Market Prices for Prescription Opioids. *Journal of Medical Internet Research*, 15(8).

36) NewScientist (February 8, 2013) *Time to Focus on the welfare of online workers*. Available online: http://www.newscientist.com/article/mg21729033.400-time-to-focus-on-the-welfare-of-online-workers.html#.VATftnknJaQ

37) Piketty T and Saez E (2013) Top Incomes and the Great Recession: Recent Evolutions and Policy Implications. *International Monetary Fund Economic Review*, 61(3). Available online: http://eml.berkeley.edu/~saez/piketty-saezIMF13topincomes.pdf

38) Piketty T (1999a) The information-aggregation approach to political institutions. European *Economic Review*, 43: 791-800.

39) Piketty T (1999b) Can fiscal redistribution undo skills-biased technical change? Evidence from the French experience. *European Economic Review*, 43: 839-851.

40) Pissarides C (June 2011) Equilibrium in Labor Markets with Search Frictions. *American Economic Review*, 101: 1092-1105.

41) 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): 447-472.

42) Raff D and Summers L.H (1987) Did Henry Ford Pay Efficiency Wages*? Journal of Labor Economics*, 5(4): S57-S86.

43) Ranard B, Ha Y, Meisel Z, and et. al. (2013) Crowdsourcing – Harnessing the Masses to Advance Health and Medicine, a Systemic Review. *Journal of General Internal Medicine*, 29(1): 187-203.

44) Robinson J (1959) *The Economics of Imperfect Competition*. London: St. Martin's Press.

45) Rogstadius J, Kostakos V, Kittur A et al. (2011) An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets. *Association for the Advancement of Artificial Intelligence*: 321-328.

46) Schmitt J (2013) Why Does the Minimum Wage Have No Discernable Effect on Employment? *Center for Economic and Policy Research (CEPR)*: Washington, DC.

47) Shaw A, Horton J and Chen D (2011) Designing Incentives for Inexpert Human Raters. CSCW 2011, Hangzhou, China.

48) Silberman S, Ross J, Irani L and Tomlinson B (2010). Sellers' problems in human computation markets. *KDD-HCOMP'10, ACM*: Washington, DC.

49) Staiger D, Spetz J and Phibbs C (2008) Is There Monopsony in the Labor Market? Evidence from a Natural Experiment. *Journal of Labor Economics*, 28(2): 211-236.'

50) Stigler G.J and Sherwin R. A (1985) The Extent of the Market. *Journal of Law & Economics*, 28(3): 555-585.

51) The Royal Swedish Academy of Sciences (2010) Markets with Search Frictions. *The Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences*: 1-34. Available online: http://www.kva.se/globalassets/priser/ekonomi/2010/sciback_ek_10.pdf

52) Varian H (1987) Price Discrimination. *Center for Research on Economic and Social Theory (CREST)*, Working Paper No.87-26.

53) Weiss A (1980) Job Queues and Layoffs in Labor Markets with Flexible Wages. *Journal of Political Economy*, 88(3): 526-538.

54) Williamson O.E (1979) Transaction-Cost Economics: The Governance of Contractual Relations. *Journal of Law and Economics*, 22(2): 233-261.

55) Yuen M-C, King I and Leung K-S (2012) Task Recommendation in Crowdsourcing Systems. *CrowdKDD'12*: Beijing, China.