Avoiding the South Side and the Suburbs: The Geography of Mobile Crowdsourcing Markets

Jacob Thebault-Spieker
GroupLens Research
University of Minnesota
thebault@cs.umn.edu

Loren Terveen
GroupLens Research
University of Minnesota
terveen@cs.umn.edu

Brent Hecht
GroupLens Research
University of Minnesota
bhecht@cs.umn.edu

ABSTRACT

Mobile crowdsourcing markets (e.g., Gigwalk and TaskRabbit) offer crowdworkers tasks situated in the physical world (e.g., checking street signs, running household errands). The geographic nature of these tasks distinguishes these markets from online crowdsourcing markets and raises new, fundamental questions. We carried out a controlled study in the Chicago metropolitan area aimed at addressing two key questions: (1) What geographic factors influence whether a crowdworker will be willing to do a task? (2) What geographic factors influence how much compensation a crowdworker will demand in order to do a task? Quantitative modeling shows that travel distance to the location of the task and the socioeconomic status (SES) of the task area are important factors. Qualitative analysis enriches our modeling, with workers mentioning safety and difficulties getting to a location as key considerations. Our results suggest that low-SES areas are currently less able to take advantage of the benefits of mobile crowdsourcing markets. We discuss the implications of our study for these markets, as well as for "sharing economy" phenomena like UberX, which have many properties in common with mobile crowdsourcing markets.

Author Keywords

Mobile crowdsourcing; volunteered geographic information;

ACM Classification Keywords

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – Computer-supported cooperative work.

INTRODUCTION

Mobile crowdsourcing markets are a recent development in crowdsourcing that has attracted substantial attention from both academia and industry. These markets are similar to purely online (or virtual [35]) markets like Amazon Mechanical Turk, but offer tasks that are situated in the physical world. For instance, the mobile crowdsourcing market Gigwalk has worked with Microsoft Bing to collect 3D panoramas [39]. Another market, TaskRabbit, suggests that people "outsource household errands and skilled tasks" (e.g. minor home repairs, grocery shopping, cleaning one's house, assembling Ikea furniture) [40].

Despite the growing interest in mobile crowdsourcing markets (e.g. [25,35]), no work has focused on understanding these markets from a geographic perspective, a perspective that is fundamental to the physically-situated nature of mobile crowdsourcing tasks. As a result we have little understanding of the geography of mobile crowdsourcing markets. For instance, the relationship between a task's location and the price at which a crowdworker will perform the task is unknown. Even more fundamentally, the same is true for the relationship between a task's location and the likelihood of finding a mobile crowdworker willing to do the task at all.

This paper seeks to address this gap in the literature through a survey deployed on *TaskRabbit,* a popular mobile crowdsourcing market. Through quantitative and qualitative analysis of our survey results, we find that two types of areas are disadvantaged on TaskRabbit: areas with low socioeconomic status (SES) like the "South Side" of Chicago and suburban (and rural) areas. People in both types of areas have access to significantly fewer mobile crowdsourcing workers. Additionally, in many cases, they will have to pay more to get a task completed. With respect to low-SES areas, this result is especially problematic as it suggests that when it comes to mobile crowdsourcing, it is "expensive to be poor" [9] and people who live in low-SES areas are less able to leverage the efficiencies of mobile crowdsourcing (e.g. spending time on longer-term goals that may increase one's SES versus "crowdsourceable" household labor).

Our quantitative and qualitative analyses reveal several mechanisms behind this unequal geography of mobile crowdsourcing markets. Our statistical modeling shows that both distance to a task and the SES of a task area influence whether a worker is willing to accept a task, and that

CSCW 2015, March 14–18, 2015, Vancouver, BC, Canada.

http://dx.doi.org/10.1145/2675133.2675278

distance is the dominant predictor of task price. Qualitative responses to our survey suggest that perceptions of higher crime in lower SES areas reduces workers' willingness to go to these areas to complete a task. Additionally, very few survey respondents reported living in low-SES areas, leading to longer travel times – and increased prices – even when people in these areas are able to find a crowdworker. This is particularly the case where the socioeconomic residential segregation that is endemic to most major metropolitan areas leads to large low-SES districts (e.g. the South Side of Chicago, west Oakland, north Minneapolis).

This work also contributes to our understanding of the role of gender in crowdsourcing markets. In particular, our survey reveals that women are far more likely than men to avoid doing a task, primarily for safety and distance reasons. As such, although women made up the majority of our respondents, the number of available tasks is somewhat smaller for them than it is for male workers. A larger pool of women are likely competing for a smaller pool of tasks, which may have an effect on wages over time.

We next discuss related work, followed by a description of the study we deployed on the mobile crowdsourcing market TaskRabbit. We then present both the quantitative modeling and qualitative analysis results of our study. Finally, we discuss in more detail the implications of our findings.

RELATED WORK

Our work is informed by research on several topics: general properties of crowdsourcing markets, motivations of mobile crowdworkers, and geographic coverage biases in geowikis (e.g. OpenStreetMap) and social media VGI (e.g. geotagged tweets and photos).

Non-mobile crowdsourcing markets

Non-mobile (or "virtual" [35]) crowdsourcing markets like Amazon's Mechanical Turk (AMT) have been the subject of intense study. Work in this area ranges from novel applications like word processors backed by the crowd [3], to complex architectures exploring the best way to organize people [16,19], to understanding and supporting the labor dynamics of crowd workers [15]. One of the more common uses of AMT is data processing, ranging from translation of text [8] to annotation of images [29].

More directly relevant to this work is research on the demographics of workers on AMT (e.g. [14,26,30]). As AMT gained popularity and success, the demographics of workers started to shift, from predominantly US-based workers using AMT as a second income source, towards a very large proportion of workers based in India using AMT as a full-time job [14,30]. This shift in the demographics of AMT, and the role that crowdsourcing plays as an income source for these groups has led some to build technology to try to address the power imbalance between workers and Amazon [15].

In this work, we focus on mobile crowdsourcing markets rather than virtual markets like AMT. The inherently geographic nature of mobile crowdsourcing markets results in substantially changed dynamics – e.g. distance and environmental effects – relative to AMT and other virtual crowdsourcing markets.

Mobile crowdsourcing marketplaces

Mobile crowdsourcing marketplaces are the geographic counterpart to virtual crowdsourcing markets like AMT [35]. A number of companies have sprung up to build mobile crowdsourcing marketplaces (e.g. GigWalk [41], Field Agent [42], and TaskRabbit [40]). While these marketplaces are still relatively new, they have quickly attracted interest from researchers. Initial work has focused on participant behavior. For example, Teodoro et al. [35] conducted a qualitative study to investigate the motivations of workers in TaskRabbit and Gigwalk. They found that monetary compensation and control of working conditions (time of day, rate of pay, the tasks they do) were primary factors for joining these systems.

Rather than studying mobile crowdworkers on existing platforms, Alt et al. [1] independently developed a mobile crowdsourcing system. They asked people to complete tasks using a smartphone and observed their behavior. They found that workers were more willing to do tasks before and after business hours, tasks that were near their home or office, and very simple tasks (e.g., taking photos).

Although neither Teodoro et al. nor Alt et al. focused on the geography of mobile crowdsourcing markets as we do here, both observed that how far people have to travel appears to influence their attitude toward a task, a finding that is a reflection of the important principle of distance decay in geography (or that as the distance between two locations increases, interaction between them tends to decrease, e.g. [10]). Similarly, Musthag and Ganesan [25] studied distance-to-task as a factor in a study of the behavior of "power users" in a mobile crowdsourcing system. They found that power users travel further per day, earn more per mile travelled, and account for 84% of total earnings and 80% of the completed tasks in the system. As detailed below, we enrich our understanding of the role of distance in mobile crowdsourcing markets by including it in a model of decisions related to pricing and willingness to complete a task. We learned additional context (e.g. the relationship of distance to other factors like SES) and new specifics (e.g. quantifying the relationship of distance to price).

Coverage problems in mobile crowdsourcing

Mobile crowdsourcing markets exist in a broader mobile crowdsourcing universe, which also includes phenomena like geowikis (e.g. OpenStreetMap and Cyclopath [27]), and physically-situated citizen science projects (e.g. [31]). Research has shown that these systems can have drastic geographic coverage problems (e.g. [12,21,24,28,37]). For instance, Quattrone et al. [28] show that countries with higher GDP and a lower Power Distance (more egalitarian) have better coverage in OpenStreetMap. Similarly, Haklay et al. [12] find that within Britain, the most deprived areas (according to the Index of Deprivation, an aggregate metric of SES factors) tend to have worse coverage than those areas that are less deprived.

Volunteered geographic information (VGI) [11] is a concept from the geography and geographic information science communities that is closely related to mobile crowdsourcing. According to Goodchild [11], VGI is "the widespread engagement of large numbers of private citizens, often with little in the way of formal qualifications, in the creation of geographic information". However, despite this "widespread engagement", VGI systems [20,23] also have been shown to have socioeconomic biases. Li et al. [18] studied Flickr and Twitter, demonstrating that low SES areas and rural areas both have worse coverage (less data) than higher SES and urban areas. Similarly, Hecht and Stephens [13] found that people from rural areas produce less social media VGI (e.g. Twitter, Flickr, Foursquare) per capita than their urban counterparts.

This research

Prior research has demonstrated coverage biases in mobile crowdsourcing systems that correlate with differences in socioeconomic status and presents qualitative evidence that distance matters to crowdworkers as they consider what tasks to take on. We extend prior research in several ways. (1) We formalize the decision-making of crowdworkers as consisting of two sequential steps: **willingness** to do a task, and how much compensation they require for the task **(**i.e. the **price** of the task**)**. (2) We do a **controlled study** of crowdworkers from one of the most popular marketplaces for mobile crowdwork, TaskRabbit. (3) We use **formal statistical modeling** to quantify the effect of distance-to-task and SES status of the task area on workers' decisions regarding willingness and compensation. (4) We enrich the results of our formal model through **qualitative analysis** that reveals the reasons behind workers' decisions.

RESEARCH QUESTIONS

As noted above, mobile crowdworkers must ask themselves two questions when confronted with a potential task: "Am I **willing** to do this task?", and second, "Is the **price** acceptable?" (or in the case of some markets, "What price would I ask?"). Our two research questions in this study are targeted at better understanding each step of this process from a geographic perspective. Our first question, aimed at the first step, is thus:

RQ-Willingness: *Where will participants in mobile crowdsourcing markets be willing to go to complete tasks?*

The related work outlined above led us to hypothesize that two geographic factors may influence workers' willingness to complete a task. First, the SES-related coverage biases observed in domains like geowikis and volunteered geographic information suggest that SES may play a role in this step of the process. Second, within geography and related fields, distance is often thought of as a cost function [10]. This leads to concepts like distance decay, which has been observed in mobile crowdsourcing systems. Therefore, we hypothesized that:

H-Willingness: *As distance to a task increases, willingness to complete the task will decrease (H-Willingness-Distance), and as the SES of a task's surroundings increases, willingness to complete a task will increase (H-Willingness-SES).*

Once crowdworkers have decided they are *willing* to complete a task, they must then decide on a *price* for the task. In TaskRabbit, this can either be accepting an offer, or making a bid. To address this second stage in the decision-making process, we ask our second research question:

RQ-Price: *How does geography affect how much participants in mobile crowdsourcing markets request in payment?*

The literature discussed in the previous section led us to hypothesize that SES and distance would also affect the price a worker charges for a task. More specifically, we hypothesized that:

H-Price: *Participants will demand a higher price for tasks farther away from where they work and live (H-Price-Distance) and in lower SES areas (H-Price-SES).*

STUDY DESIGN

To address the above research questions and evaluate the corresponding hypotheses, we developed a survey and deployed it on TaskRabbit. TaskRabbit is a well-known mobile crowdsourcing market used by task requesters for the completion of physically-situated tasks ranging from delivering flowers to helping build IKEA furniture to helping task posters move large items.

We deployed our survey in an organic fashion by posting it as a task to TaskRabbit's Chicago metropolitan area site just as a typical task requester would post a task. Only TaskRabbit workers local to the Chicago area could fill out the survey, for which we paid respondents $5 in 15-minute intervals, capped at an hour (e.g. a person who took more than 15 minutes but less than 30 would receive $10).

To add context to our survey results, we first asked respondents a number of questions about themselves, such as their gender, their preferred mode of transportation, and their activity level on TaskRabbit. We also asked respondents to select their home census tract[1] on a map we provided, and to do the same for census tracts that they visited at least once a month.

---

[1] Census tracts are geographic areas defined by the U.S. census and "generally have a population size between 1,200 and 8,000 people, with an optimum size of 4,000 people" [36].

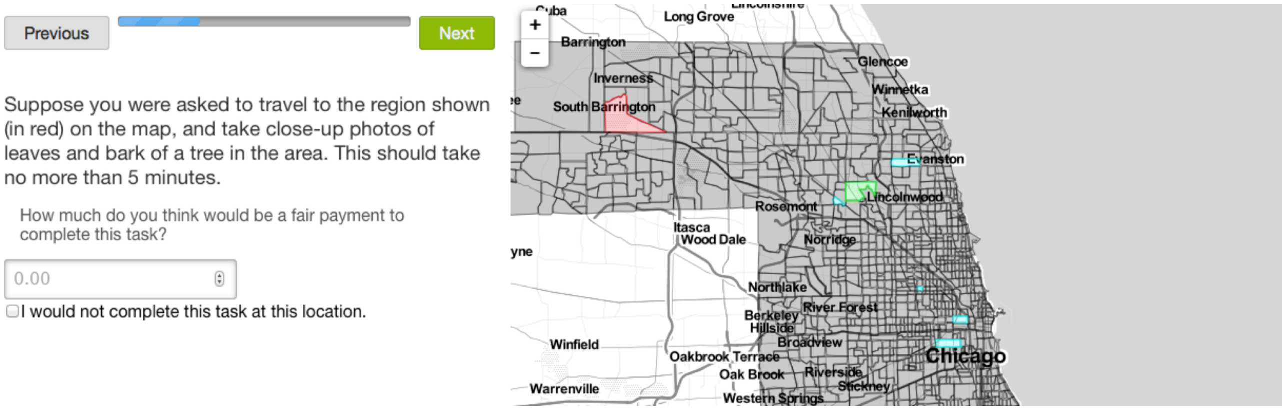

Figure 1: A screen grab from the survey taken by workers on TaskRabbit. The green census tract is the worker's self-reported home tract and the blue tracts are those that the worker reported visiting at least once a month. The red tract is the tract about which the worker is currently being questioned (Note: screen grab is cropped for space and, for privacy reasons, the figure does not depict an actual crowdworker's responses).

After respondents answered questions about themselves, they began the main portion of the survey, which was dedicated to investigating *RQ-Willingness* and *RQ-Price*. This portion of the survey involved prompting respondents with census tracts in Cook County, Illinois, which contains Chicago and many of its suburbs[2]. For each census tract, the respondent had the option to either check a box labeled "I would not do this task at this location" or to name what they felt would be a fair price to complete the task (see Figure 1) were they to be asked to do it by a task requester.

In order to control for task-specific attributes, each tract was randomly assigned one of three hypothetical tasks designed to vary the level of engagement with the local area (Table 1). These tasks ranged from one that could be done without leaving a vehicle to one that required interaction with a person in the area. Tasks were designed so that they would not take more than five minutes.

Task	Engagement Level
Task 0: Suppose you were asked to travel to an intersection in the region shown (in red) on the map, and photograph all of the signs at that intersection. This should take no more than 5 minutes.	Low
Task 1: Suppose you were asked to travel to the region shown (in red) on the map, and take close-up photos of leaves and bark of a tree in the area. This should take no more than 5 minutes	Medium
Task 2: Suppose you were asked to travel to the region shown (in red) on the map, visit someone's home, and ask the owners to respond to a single question about local politics. This should take no more than 5 minutes	High

Table 1. Tasks in our survey and their hypothesized engagement level.

Each participant received 20 census tracts (18 unique). Fourteen of the tracts were randomly selected (without replacement), and these tracts were the tracts that were considered in our quantitative analysis below. In order to enrich our qualitative understanding of the geography of mobile crowdsourcing markets, we also considered four special-case tracts: the highest-income and the lowest-income tracts and the highest-crime[3] and the lowest-crime tracts. The other two tracts were repeated from the randomly chosen set of 14 in order to verify intra-rater reliability. The repeated tract was presented no fewer than 5 tracts after the original.

Upon seeing and responding to all 20 of the tracts with either a price or by stating that they would not complete the task, we asked respondents several open-ended questions whose answers were entered into text boxes. Specifically, we asked respondents about how they made their pricing decisions and why they would not complete certain tasks (if they checked that box at least once).

RESULTS

Forty respondents completed the survey, which we ran during Spring 2014. 57.5% of respondents were women (42.5% men). The median respondent performed a task on TaskRabbit between once a week and once every two weeks. 30% of respondents indicated that they complete multiple tasks per week, while only 20% of respondents indicated that they complete a task once a month or less.

[3] As reported by the Chicago Police Department (Cook County-wide crime data is not available).

RQ-Willingness

Because price is irrelevant if a worker will not complete a task, we first sought to understand the geography of workers' *willingness* to complete tasks. To do so, we built a logistic mixed effects model with three fixed effects:

- Distance to task from the closest census tract visited by the respondent at least once a month (as indicated in the survey)
- Median household income of the task tract, as an indicator of socioeconomic status. Many other socioeconomic variables are well known to be correlated with income (e.g. educational attainment, occupation). To reduce the effect of the long-tailed distribution of wealth, we log-transformed this variable. Median household income data was gathered from the United States Census' American Community Survey 2006-2010 dataset.
- Task ID, to make sure we understand the effect of distance and median income in the context of a given task.

The models' random effects were intercepts for respondent and by-respondent slopes for the effects of income and distance. The model's dependent variable was whether or not the survey respondent had checked the "I would not do this task at this location" box for a given tract. Overall, respondents indicated that they would not do 34% of the tasks. The few census tracts that had a reported median income of zero (e.g. the tract that consists of O'Hare International Airport and a few hotels[4].) were excluded

Fixed Effect	Estimate	*p*-value
Travel time (in hours)	-3.15 (0.99)	0.001
Log_2[Task tract income in $10k]	0.87 (0.36)	0.014
Task ID (baseline = Task 0)	1: 0.37 (0.40) 2: -0.92 (0.40)	0.003
Constant	1.81 (0.82)	0.028

Table 2: The results of our *willingness* model (p-value for Task ID was calculated with an ANOVA).

Fixed Effect	Estimate	*p*-value
Travel time (in hours)	10.10 (2.27)	<0.001
Log_2[Task tract income in $10k]	0.40 (0.52)	*n.s.*
Task ID (baseline = Task 0)	1: -1.73 (0.85) 2: 0.28 (0.87)	0.024
Constant	16.92 (2.90)	<0.001

Table 3: The results of our *price* model (p-value for Task ID was calculated with an ANOVA).

To operationalize distance, we used travel time rather than Euclidean distance to better match the lived experience of mobility in Chicago. We used the Google Distance API to calculate the off-peak travel time between the centroid of the task tract and the centroid of the nearest (to the task tract) tract that the respondent indicated visiting frequently (more than once a month). The API supports multiple transportation modes, and we calculated travel time with respondents' self-reported preferred mode of transportations.

Table 2 shows the results of our model. All fixed effects are significant. Returning to our hypothesis *H-Willingness*, we find that both *H-Willingness-Distance* and *H-Willingness-Price* are supported. Socioeconomic status of the task location and distance to the tract both have an effect on whether or not a worker on TaskRabbit is willing to accept a task, with SES having a significant positive relationship and distance having a negative one. According to the model, for every doubling of task area median income, there is a 2.38x increase in likelihood that a worker will accept a task. In other words, holding the other variables constant, our model suggests that the likelihood of a worker accepting a task will more than double if the task is in a tract with a median income of, for instance, $60K rather than a tract with a median income of $30K. As shown in Figure 2, $60K is a relatively normal median household income in northern Chicago and the Chicago suburbs, with $30K median household incomes common on the South Side.

With respect to travel time, our model indicates that for every hour of travel time there is a substantial decrease in willingness. Specifically, workers are about 4.3% as likely to be willing to do a task an hour away as they are one located in their immediate vicinity. We also saw that respondents were significantly less willing to do the most engaging task relative to the others ($p < .05$). Further investigating this phenomenon, especially with regard to its interaction with SES, is a direction of future work (see below).

Examining our willingness results in more detail, we found an interesting result with regard to gender. While 78% of women said they would not complete at least one task, the equivalent number for men was 53%. In addition, the grand mean willingness (mean of means) for women is 57.1%, for

[4] One respondent indicated living in this tract. While we did not consider samples where the proposed task was in this tract (and other zero-income tracts), we did include this user's responses about tasks in other tracts because there are reasonable residential options in this tract (though temporary ones).

men it is 77.7%. Our qualitative results below suggest that both distance and crime factors play a role in women's explicit willingness decisions, but these are the same factors also indicated by men. Although further research is needed, it is likely that women have a lower threshold for one or both of these factors.

RQ-Price

We next turned our attention to analysis of the price respondents indicated they would charge for a task, assuming they were willing to complete it. We began our analysis of the price data by ensuring that it had sufficiently high intra-rater reliability. We did so by calculating the Pearson's correlation coefficient between the first and second price judgments for repeated tracts. The coefficient was r = 0.96 across all respondents, indicating that respondents' pricing decisions were very consistent.

To understand the geography of prices on TaskRabbit and the effects of distance and SES on these prices, we built a linear mixed effects model with identical independent variables as our willingness model but with reported task price as dependent variable.

The results of this price model can be seen in Table 3. The table reveals that travel time was indeed positively related to price, supporting *H-Price-Distance*. Indeed, the model suggests that for every hour of travel time, price goes up at a rate of $9.97/hour.

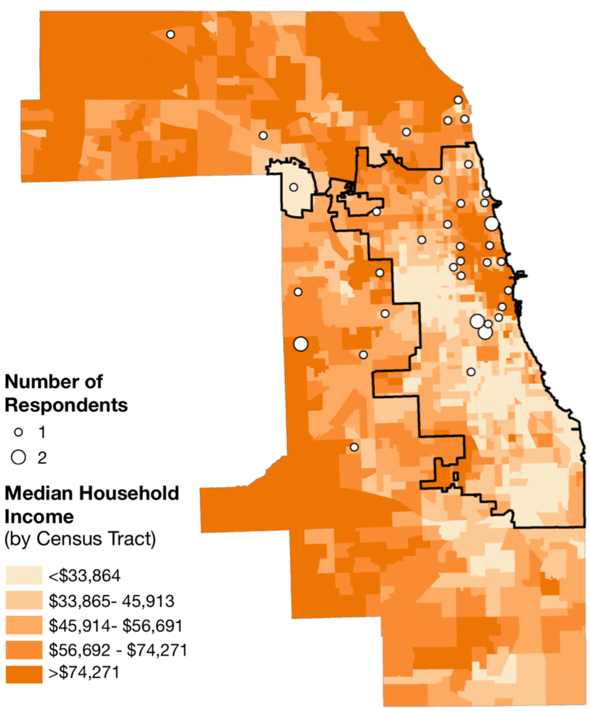

Figure 2. Survey respondents' self-reported home census tracts (centroids) and median income in Cook County, Illinois. Very few respondents live in low-income tracts. Note that the low-SES south side of Chicago (Chicago is outlined with a bold border) has only one respondent, and no respondents live in the poorest parts of south Chicago. Median income color classes are determined via the quantile method, meaning each class represents a quintile of the household income dataset.

Task tract income, on the other hand, was not significant; the median household income of the tract does not have a significant effect on price. In other words, *H-Price-SES* was not supported.

However, examining the geography of where our respondents live provides additional context for the role of distance in TaskRabbit. Figure 2 shows the self-reported home tracts of all 40 respondents on top of a map of income by census tract in Cook County. Immediately visible in Figure 2 is that very few respondents live in the heart of low-income areas. Indeed, most respondents seem to cluster around the very high-income portions of northern Chicago. Only a single respondent lives well within the lower-income South Side of Chicago. As a result, a low-income resident of the South Side would have to pay more for mobile crowdsourcing services (e.g. someone to take care of errands to make time for longer-term goals) than someone in wealthier areas of Chicago, and is likely to have a harder time finding someone willing to accept this request for services in the first place.

This suggests that the character of the socioeconomic residential segregation that is common to many metropolitan areas around the world – i.e. the "Big Sort" [4] – may play an important role in the price of mobile crowdsourcing market tasks. Where this segregation results in large-area low-income districts like the South Side, the people who live in these districts will live far from mobile crowdworkers, resulting in longer travel times, and higher prices for tasks. However, where low-SES pockets are much smaller (e.g. the lower income pockets in the suburbs just north of Chicago), the effect on travel time, and therefore price, will be minimal.

It is also important to note that distance not only disadvantages lower SES areas, but also disadvantages people who live in distant suburbs. Figure 2 shows that workers are concentrated in the dense city of Chicago rather than the suburbs. However, as can also be seen in Figure 2, suburban people have higher incomes than people on the South Side of Chicago, and thus they can potentially afford the increased costs. In addition, in many cases, even people in somewhat remote suburbs are closer to one of our respondents than a person in south Chicago.

Further, while United States suburbs tend to be relatively wealthy, the opposite is true in many cities around the world (e.g. France and Latin America [6,10]). Where this is the case, mobile crowdsourcing will likely be drastically more expensive and less accessible in these areas. As we note below, investigating these phenomena in cities with different socioeconomic segregation patterns is an important direction of future work.

QUALITATIVE RESULTS

Thus far, our quantitative models and our descriptive statistics have revealed a number of properties of the geography of mobile crowdsourcing, e.g. crowdworkers are less willing to do tasks in lower SES areas, and prices are higher in areas far from the areas frequented by crowdworkers. We now turn to our qualitative results to attempt to help understand *why* these dynamics are at play in mobile crowdsourcing markets. To do so, a single investigator looked for themes in the survey responses, focusing on ideas related to the geographic concepts *site* and *situation.* We use this framing to discuss the results of our qualitative data.

Site and *situation* are core concepts from geography used to describe the attributes of an area [10,17,22]. Broadly stated, *site* attributes are properties of an area that are not directly dependent on other areas (e.g. mineral deposits, quality of terrain, local weather). *Situation* attributes describe connectivity of one place to another (e.g. located on a major highway connected to a big city, has a seaport).

In the context of this work, *site* attributes refer to reputation of a census tract, the perception of safety in that tract, poverty in that tract, and so on. *Site* attributes do not change with the crowdworker being asked. *Situation* attributes, on the other hand, are those that vary by relative location with respect to the crowdworker. These include properties like easy access to the location, travel cost with respect to mode of travel, the time of day the task was requested for, distance to the location, etc. We will discuss each of these themes in turn.

Site attributes

Perceptions of high crime in the task's census tract was a site attribute that was very commonly mentioned by our respondents as a reason they ticked the "will not do" box. For instance, respondent 27 (R27) wrote:

I think the high incidence of gang-related crime makes many Chicagoans too nervous to visit some parts of the city. We always refer to Chicago as being a "city of neighborhoods" but the truth is that many Chicagoans feel uncomfortable visiting a huge portion of our city. The nature of the crimes that occur on the South and West Sides (gang-related) makes me particularly nervous because there's nothing you can do to prepare/protect yourself. I realize that I might have some biases but it's less about location for me and more about crime rate. I do wish Chicagoans (and visitors) could feel more comfortable exploring and enjoying more neighborhoods without worrying about crime."(R27)

In the above quote, R27 discusses the unease she perceives among others in Chicago and the (reported) reputation that the South Side and West Side of Chicago have for having large amounts of violent, gang-related crime.

R39 specifically addresses her gender as part of the reason she did not consider certain tasks, saying:

"I wouldn't feel safe in some areas as a female by herself." (R39)

R9 is a member of the TaskRabbit Elite. This is a designation one can earn within TaskRabbit after earning an average rating of 4.9 stars (ratings are given by task posters upon completion), completing a large number of tasks, and not violating any of TaskRabbit's policies. R9 offered similar feedback to R27:

I am an Elite member of TaskRabbit and I do a lot of tasks. I do not do tasks anything below the loop of Chicago [i.e. the South Side] *so it has to be on the north side for me to work. It is purely for safety concerns. (R9)*

R4, a relatively new resident of Chicago, similarly makes decisions about where to participate within the city based on the reputation of certain areas with respect to crime. In this case, she explicitly mentions poverty as well, providing qualitative support for our SES results.

I only moved to Chicago last May (2013) so I don't know much about the city except that there are large pockets of poverty, inequality and high crime. In terms of general areas of the city I understand that large swaths of the south side and west side include these pockets of poverty and high crime. Without specifics about which neighborhoods/blocks/streets are safe I essentially ruled out anything on the south or west side of the city. For the most part, I think the western suburbs are safe but I know nothing about the southern suburbs so I erred on the side of safety and avoided those areas as well. (R4)

Other participants raised this idea of generalizations about, and reputations of, unsafe areas as well:

Whether or not my assumptions of lack of safety were correct, I wouldn't put myself in danger for a few dollars (R16)

Situation attributes

As can be expected from our quantitative models, the respondents' qualitative feedback suggests that proximity or convenience of the task location is a very important factor in their pricing and willingness decisions. Here, R4 explicitly discusses the role of convenience in her pricing decisions.

Mostly how much of a pain it was going to be to get there. If it was a place I could stop by on my way to or from work or the gym= cheap. If it required getting in my car=more. If it required an extensive drive to a far flung suburb=more. (R4)

Others, like R16, discuss the role of transportation mechanisms and time it would take to travel to areas to complete tasks.

Other areas were too far from the Metra [the commuter rail system in Chicago] *to make it worth my while. Others were still close to the Metra but far enough away where the ticket round trip would be a bit pricy. (R16)*

Recognition of diminishing returns was also raised by other participants, like R23, R31, and R39, who discussed the tension between traveling long distances and being able to complete a task at a "fair" price:

"I didn't think any price would be worth the commute and risk while still offering even a marginally fair price." (R23)

"The distance was too far to justify any fair price for completing task. The price would have to be higher/greater than 25 dollars to justify it." (R31)

"getting there would take me longer than actually completing the task" (R39)

More specifically, these crowdworkers discuss needing to ensure they could afford to complete the task, and were balancing a trade-off between distance, the time it would take, and the cost of travel to these locations. The diminishing returns were both internal (e.g. it wouldn't make sense for them to travel), and external (e.g. it would be too expensive for the task poster to hire me to do this).

DISCUSSION

Above, we examined the geography of mobile crowdsourcing markets through quantitative and qualitative lenses, finding that the SES of a task area and distance to a task are key factors in access to the benefits of these markets. Below we discuss additional themes emergent from our findings and highlight our findings' implications for mobile crowdsourcing markets and related platforms like car-sharing services (e.g. UberX).

Relationship between Income and Crime

One important theme that emerges when comparing our quantitative and qualitative results is that perceived danger from crime may be an important mechanism through which low SES areas are disadvantaged in mobile crowdsourcing markets. In our qualitative feedback, crime was far more frequently discussed as a reason for declining a task than SES, but crime and SES are known to be highly correlated (e.g. [2,5,7]). Indeed, we obtained violent crime data from the Chicago Police Department (equivalent data for the entire study area was not available, since it extended outside the Chicago city limits) and found that the (inverse) Pearson's correlation with median household income was quite high (r = -0.70). In other words, it may be fear of crime that is at least partially driving the SES effects we saw in our model.

It is important to note that there is likely some disagreement between perception and reality when it comes to crime. For instance, while the South Side does have some very dangerous areas, it also has significantly safer areas. The safer areas – low SES or otherwise – may be grouped together with the more dangerous by crowdworkers, leading to less access to crowdworkers' services even in these safer South Side areas.

Similarly, it is also important to note that due to the tremendous economic inequalities that occur across racial and ethnic lines in the United States, our results suggest that specific racial and ethnic groups have differential access to mobile crowdsourcing market resources. Specifically, we observed that the percent of the population that self-identifies as *white (non-Latino)* (a designation used by the U.S. census) has a strong correlation with income in our study area (r = 0.67). Therefore, *because of these correlations,* many of the patterns we saw for SES factors may be true for race/ethnicity as well. To illustrate this point, we found that in a version of our willingness model that used percent white (non-Latino) as a fixed effect instead of income, white (non-Latino) was marginally significant (p = 0.06).

These relationships reflect a set of real and *perceived* correlations in our study area (and in the United States as a whole) between SES factors, neighborhood safety, and race/ethnicity. Our inquiry into the factors that influence mobile crowdworkers' willingness to travel to a particular area to do a task made it inevitable that these relationships and correlations would surface. We think it is intellectually appropriate to acknowledge them, situate them in proper context, and exercise caution in interpreting results based on them. We hope that other researchers looking at similar problems will do the same.

Implications for Ride-Sharing Platforms

Another interesting point of discussion ensues from the rise of car-sharing platforms like UberX, which have strong similarities to mobile crowdsourcing markets. In UberX, drivers (i.e. workers) have agency in terms of which regions they visit to try to find fares, just as TaskRabbit crowdworkers can choose which tasks to accept. As such, UberX may be subject to similar phenomena as we observed for TaskRabbit. Namely, our results suggest that people who live in low SES areas (or high crime areas) may find it significantly more difficult to take advantage of services like UberX, especially in terms of available drivers and, in certain areas, in terms of higher prices. Indeed, there is growing anecdotal evidence that this is the case, with Uber being increasingly accused of ride-sharing "redlining", or providing less service and higher prices in lower SES areas (e.g. [32,33]). Uber is currently the subject of widespread discussions with respect to transportation policy, and we hope this work can provide a related data point in these discussions. Repeating this work with UberX drivers is a valuable direction of future work.

Crowdworker Recruitment and Background Checks

Although we have focused on the benefits of mobile crowdsourcing markets for task requesters, another benefit of these markets comes from the ability to find work as a crowdworker. However, our results also indicate that people in low SES areas again are disadvantaged in this respect. The vast majority of respondents reported living in census tracts whose median household income is well above the United States federal poverty line for a household of four (Figure 2). Indeed, the median of the median household incomes of respondents' census tracts was \$51,216 (over 200% of the poverty line). While it is possible that some respondents are much poorer than their tract's median, it is unlikely that this is true of a substantial number. In addition, only one respondent lives in the very poor western and southern areas of Chicago (and s/he lives in a relatively wealthier tract), meaning that that these populations are not yet leveraging TaskRabbit for work.

One reason for the underrepresentation of low SES crowdworkers is likely TaskRabbit's set of requirements for new crowdworkers. For instance, TaskRabbit mandates that crowdworkers pass a background check and have a smartphone and a bank account, both of which are found less often in low-income households than higher-income households (at least in the United States [34,38]). Interestingly, UberX also has requirements that would restrict lower income people from becoming drivers, e.g. owning a car, having insurance.

Our results suggest that reducing the barriers to entry for potential workers would be an effective way of increasing the accessibility of the benefits of mobile crowdsourcing markets. Not only would this bring in more low-income mobile crowdworkers, but many of these low-income workers would likely live in lower income areas, decreasing the distance to the nearest crowdworker for people in these areas. Our model indicates that this would both increase the number of workers willing to complete tasks and decrease prices for these people. Straightforward means of reducing barriers to entry would include providing mobile crowdworkers with a specialized smartphone (Uber does this) and paying workers through means that do not require a bank account.

It's "Expensive to be Poor" in Mobile Crowdsourcing

One of the major benefits of mobile crowdsourcing markets is that they allow task requesters to use their time more efficiently in an economic sense. For instance, consider a Silicon Valley worker with a high salary and many demands at her job. If she can find a TaskRabbit to do her laundry, pick up her packages, go to the grocery store for her, etc. at say, \$20/hour, she can focus on succeeding at work, which has much higher potential payoffs over the long term. Our results suggest that these efficiencies are harder to obtain through mobile crowdsourcing in low SES areas. This is a classic "expensive to be poor" [9] scenario in a new domain.

Other Directions of Future Work

An important direction of future work for us will involve looking at the implications of this papers' findings for rural areas. The fact that distance was significant both in willingness and in price does not bode well for the success of mobile crowdsourcing markets in rural contexts. Rural areas are some of the least-covered spaces in other types of mobile crowdsourcing (e.g. [13,37]) and it appears it is cost-prohibitive to use existing mobile crowdsourcing markets there as well.

Because of residential segregation, we might anticipate that distance and accessibility discrepancies are bi-directional, but this is an interesting direction for future work. Would locals to low-SES areas travel to higher-SES areas to perform crowdwork? What about equally-distant low-SES areas? Addressing each of these questions would add to our understanding of price and willingness in mobile crowdsourcing markets.

Finally, we also are looking further into the gender dynamics in mobile crowdsourcing markets. We hope to uncover additional details with regard to the differences in willingness between men and women that we saw in our survey, as well as better understand the potential economic implications (e.g. on overall wages).

Limitations

While the above work sheds new light on the role of SES and geography more generally in mobile crowdsourcing markets, our study also has several limitations that are important to discuss.

First and foremost, some participants in their qualitative feedback indicated that the time we allocated for a task had an effect on their decision making. For instance, TR5 wrote:

"If it was too far from my home or my job I would not travel to that location to take a picture for 5 minutes." (TR5)

This dynamic may have affected the participation and price dynamics of some of the locations. Therefore, our future work will involve investigating the role task duration plays (even in the context of flat-rate vs. hourly pricing).

Second, our study was performed on a single county. While we believe our results are generalizable to most cities, it will take additional research to be sure.

CONCLUSION

We presented evidence from a quantitative and qualitative survey of TaskRabbit indicating that mobile crowdsourcing markets advantage the advantaged and disadvantage the disadvantaged. Specifically, we show that the SES of a task area is associated with the willingness of crowdworkers to complete a task, with lower SES leading to fewer willing crowdworkers. We also discuss how distance to the task – which affects both the willingness of a worker to complete a task and the price at which the worker is willing to complete it – can act as an agent of SES in certain contexts due to endemic residential segregation in metropolitan areas. Qualitatively, we find that concerns about safety are a primary mechanism for our SES-related finding. By showing the uneven geography of mobile crowdsourcing markets, it is our hope that this work can help bring the benefits of mobile crowdsourcing (and related concepts like ride-sharing) to a wider audience.

ACKNOWLEDGEMENTS

This work was supported in part by NSF grants IIS-1218826 and IIS-0808692 and a Yahoo! ACE Award. We would like to thank Darren Gergle, Shilad Sen, Aaron Rendahl, and our GroupLens colleagues for providing valuable feedback during the course of this work.

REFERENCES

1. Alt, F., Shirazi, A.S., Schmidt, A., Kramer, U., and Nawaz, Z. Location-based Crowdsourcing: Extending Crowdsourcing to the Real World. *October*, (2010).
2. American Psychological Association (APA). *Violence and Socioeconomic Status*. 2010.
3. Bernstein, M.S., Little, G., Miller, R.C., et al. Soylent: A Word Processor with a Crowd Inside. *Artificial Intelligence*, (2010), 313–322.
4. Bishop, B. *The Big Sort: Why the Clustering of Like-minded America is Tearing Us Apart*. Houghton Mifflin Harcourt, 2008.
5. Browne, A., Salomon, A., and Bassuk, S.S. The Impact of Recent Partner Violence on Poor Women's Capacity to Maintain Work. *Violence Against Women 5*, 4 (1999), 393–426.
6. Brunn, S.D., Williams, J.F., and Zeigler, D.J. *Cities of the World: World Regional Urban Development*. Rowman & Littlefield Publishers, 2003.
7. Buka, S.L., Stichick, T.L., Birdthistle, I., and Earls, F.J. Youth Exposure to Violence: Prevalence, Risks, and Consequences. *American Journal of Orthopsychiatry 71*, 3 (2001), 298–310.
8. Callison-Burch, C. Fast, Cheap, and Creative: Evaluating Translation Quality Using Amazon's Mechanical Turk. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1*, Association for Computational Linguistics (2009), 286–295.
9. Ehrenreich, B. It Is Expensive to Be Poor. *The Atlantic*, 2014.
10. Fellmann, J.D., Getis, A., and Getis, J. *Human Geography: Landscapes of Human Activity*. McGraw-Hill, 2007.
11. Goodchild, M.F. Citizens as sensors: the world of volunteered geography. *GeoJournal 69*, 4 (2007), 211–221.
12. Haklay, M. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B: Planning and Design 37*, 4 (2010), 682–703.
13. Hecht, B. and Stephens, M. A Tale of Cities: Urban Biases in Volunteered Geographic Information. (2014).
14. Ipeirotis, P.G. *Demographics of Mechanical Turk*. 2010.
15. Irani, L.C. and Silberman, M.S. Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 611–620.
16. Kittur, A., Smus, B., Khamkar, S., and Kraut, R.E. CrowdForge: Crowdsourcing Complex Work. *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, ACM (2011), 43–52.
17. Lanegran, D.A. and Natoli, S. *Guidelines for Geographic Education in the Elementary and Secondary Schools*. Association of American Geographers, 1984.
18. Li, L., Goodchild, M.F., and Xu, B. Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartography and Geographic Information Science 40*, 2 (2013), 61–77.
19. Little, G., Chilton, L.B., Goldman, M., and Miller, R.C. TurKit: Tools for Iterative Tasks on Mechanical Turk. *Proceedings of the ACM SIGKDD Workshop on Human Computation*, (2009), 29–30.
20. Mashhadi, A., Quattrone, G., and Capra, L. Putting ubiquitous crowd-sourcing into context. *Proceedings of the 2013 conference on Computer supported cooperative work - CSCW '13*, (2013), 611.
21. Mashhadi, A., Quattrone, G., and Mooney, P. On the Sustainability of Urban Crowd-Sourcing for Maintaining Large-Scale Geospatial Databases Categories and Subject Descriptors. (2012).
22. McKnight, T.L. *Regional Geography of the United States and Canada*. Prentice Hall, 2004.
23. Mislove, A., Lehmann, S., Ahn, Y.-Y., Onnela, J.-P., and Rosenquist, J.N. Understanding the Demographics of Twitter Users. *ICWSM 11*, (2011), 5th.
24. Mooney, P., Corcoran, P., and Winstanley, A.C. Towards quality metrics for OpenStreetMap. *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '10*, (2010), 514.
25. Musthag, M. and Ganesan, D. Labor dynamics in a mobile micro-task market. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, (2013), 641.

26. Pavlick, E., Post, M., Irvine, A., Kachaev, D., and Callison-Burch, C. The Language Demographics of Amazon Mechanical Turk. *ACL '14*, (2014).

27. Priedhorsky, R., Masli, M., and Terveen, L. Eliciting and Focusing Geographic Volunteer Work. *CSCW '10: 2010 ACM Conference on Computer Supported Cooperative Work*, (2010).

28. Quattrone, G., Mashhadi, A., and Capra, L. Mind the Map: The Impact of Culture and Economic Affluence on Crowd- Mapping Behaviours. *CSCW '14*, (2014).

29. Rashtchian, C., Young, P., Hodosh, M., and Hockenmaier, J. Collecting Image Annotations Using Amazon's Mechanical Turk. *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, Association for Computational Linguistics (2010), 139–147.

30. Ross, J., Irani, L., Silberman, M.S., Zaldivar, A., and Tomlinson, B. Who Are the Crowdworkers?: Shifting Demographics in Mechanical Turk. *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, ACM (2010), 2863–2872.

31. Sheppard, S.A., Wiggins, A., and Terveen, L. Capturing Quality: Retaining Provenance for Curated Volunteer Monitoring Data. *CSCW '14*, ACM (2014), 1234–1245.

32. Silverstein, A. Uber and Lyft Are Being Accused of Redlining Again, But Is That Actually Happening? *Dallas Observer*, 2014.

33. Singer, C. Ridesharing and Redlining: Uber, Lyft, Race and Class. *Daily Kos*, 2014. http://www.dailykos.com/story/2014/05/27/1302417/-Ridesharing-and-Redlining-Uber-Lyft-Race-and-Class#.

34. Smith, A. *Smartphone Ownership 2013*. Pew Internet & American Life Project, 2013.

35. Teodoro, R., Ozturk, P., Naaman, M., Mason, W., and Lindqvist, J. The motivations and experiences of the on-demand mobile workforce. ACM Press (2014), 236–247.

36. United States Census. 2010 Geographic Terms and Concepts - Census Tract. http://www.census.gov/geo/reference/gtc/gtc_ct.html.

37. Zielstra, D. and Zipf, A. A comparative study of proprietary geodata and volunteered geographic information for Germany. *Conference on Geographic Information 1*, (2010), 1–15.

38. *2011 FDIC National Survey of Unbanked and Underbanked Households*. Federal Deposit Insurance Corporation, 2011.

39. Gigwalk Announces $6 Million In Funding, Expands Relationship With Microsoft, And Names New CEO. http://gigwalk.com/press/gigwalk-announces-6-million-in-funding-expands-relationship-with-microsoft-and-names-new-ceo.php.

40. TaskRabbit. https://www.taskrabbit.com.

41. GigWalk. http://gigwalk.com.

42. Field Agent. http://www.fieldagent.net.