Towards a Geographic Understanding of the Sharing Economy: Systemic Biases in UberX and TaskRabbit

JACOB THEBAULT-SPIEKER, University of Minnesota
LOREN TERVEEN, University of Minnesota
BRENT HECHT, Northwestern University

Despite the geographically-situated nature of most sharing economy tasks, little attention has been paid to the role that geography plays in the sharing economy. In this article, we help to address this gap in the literature by examining how four key principles from human geography – distance decay, structured variation in population density, mental maps, and "the Big Sort" (spatial homophily) – manifest in sharing economy platforms. We find that these principles interact with platform design decisions to create systemic biases in which the sharing economy is significantly more effective in dense, high socioeconomic status (SES) areas than in low-SES areas and the suburbs. We further show that these results are robust across two sharing economy platforms: UberX and TaskRabbit. In addition to highlighting systemic sharing economy biases, this article more fundamentally demonstrates the importance of considering well-known geographic principles when designing and studying sharing economy platforms.

Categories and Subject Descriptors: H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Computer-supported cooperative work*

Additional Key Words and Phrases: sharing economy; geography, residential segregation, Big Sort, mental maps, distance decay, population density, location-aware computing; mobile crowdsourcing

1. INTRODUCTION

In the past few years, consumers have flocked to sharing economy services. UberX connects "microentrepreneur" drivers to customers needing a ride over 5 million times a day [Tepper 2016], TaskRabbit's user base grew by 340% in 2015 [Yeung 2016], and more than 60 million people have used Airbnb as of mid-2016 [Airbnb 2016]. The popularity of sharing economy platforms has been attributed to increased convenience and lower prices for consumers [Silverstein 2014] and new sources of income for microentrepreneurs [O'Brien 2015]. These and other benefits have led many to speculate that the sharing economy will become the dominant consumer paradigm of the 21st century ([Milbourn 2015; Nunberg 2016; Hempel 2016]).

While much attention has been paid to the economics and labor conditions of UberX, Airbnb, TaskRabbit and similar services (e.g. [Musthag and Ganesan 2013; Ikkala and Lampinen 2015; Dillahunt and Malone 2015; Raval and Dourish 2016; Rosenblat and Stark 2015]), there has been much less focus on another factor that is critical in nearly all sharing economy platforms: *geography*. Regardless of whether we consider ride-hailing services (e.g. UberX, Lyft), peer-to-peer rental services (e.g. Airbnb), mobile crowdsourcing services (e.g. TaskRabbit), or even non-commercial sharing economy platforms (e.g. CouchSurfing), geography plays a key role. For instance, in the case of ride-hailing, a driver must travel from his or her current location to the location of the ride requester and drive the requester to a desired destination. For Airbnb, customers must decide where to stay, and prices are in part defined by the geographic context of each option. In TaskRabbit – a well-known platform that allows people to "outsource household errands and skilled tasks" [TaskRabbit 2016] – a microentrepreneur ("tasker") commutes to the task requester's location and/or to the locations involved with the specific errand.

The goal of this article is to better understand the role of geography in the sharing economy. We work towards this goal through two studies that provide evidence that the following four established principles from the field of human geography are key factors in the relative success of the sharing economy: (1) *residential clustering* (i.e. the

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

"Big Sort" [Bishop 2008]), (2) *structured patterns in the geographic variation of population density* across metropolitan areas [Brunn, Jack Francis Williams, et al. 2003], (3) *distance decay* (i.e. as the distance between two locations increases, interaction between them tends to decrease, e.g. [Stewart 1948; Reilly 1931; Bjelland et al. 2013; Flowerdew and Aitkin 1982]), and (4) *mental maps* (e.g. [Ladd 1967; Matei et al. 2001; Gould and White 1986]).

Through our consideration of these four principles, we demonstrate the critical importance of a geographic lens when examining the sharing economy, showing that this lens can reveal structural inequalities that would be otherwise invisible. Focusing on TaskRabbit and UberX in the Chicago region, we find that the four geographic principles lead to structural geographic biases in which the sharing economy is more effective in some types of areas than other types of areas. Namely, sharing economy platforms appear to succeed in areas with high socioeconomic status (SES) and population density and struggle in areas with low SES and low population density. Our evidence, for instance, shows that people in poor neighborhoods and outer-ring suburbs in the Chicago region wait longer for UberX cars and will have a harder time finding a TaskRabbit worker ("tasker") to complete a given errand.

Additionally, as in many parts of the world, population density and SES in our study region (the Chicago area) are correlated strongly with membership in certain protected classes, in particular those defined by race and ethnicity. This relationship results in an unfortunate corollary to our findings: in some cases, the sharing economy appears to be most effective in areas with fewer minorities and much less effective in black and Latino neighborhoods.

Our work triangulates our high-level findings across multiple methodological approaches. These approaches include controlled experiments, the analysis of qualitative survey responses, and (to our knowledge) the first use of an advanced geostatistical technique known as *spatial Durbin* modelling in the human-computer interaction literature. Spatial Durbin models are emerging as a best practice in the social and natural sciences, and are an approach that we believe can be broadly useful for studying the sharing economy and in the growing "geographic human-computer interaction literature" [Hecht et al. 2013] more generally.

While this work focuses on the descriptive analysis of the geography of sharing economy platforms, our studies also provide evidence for potential solutions to the challenges we identify. For instance, we observe that very few TaskRabbit microentrepreneurs live in low-SES neighborhoods and that large microentrepreneur-to-job distances contribute to the higher prices and decreased willingness to accept jobs in these neighborhoods (a manifestation of distance decay). Since in TaskRabbit, UberX, and most other sharing economy services, low-SES individuals have a harder time satisfying microentrepreneur enrollment requirements (e.g. microentrepreneurs must have a bank account in many cases), this likely reduces microentrepreneur participation in low-SES neighborhoods and thereby diminishes the effectiveness of the overall platforms in these neighborhoods. Below, we discuss how our results – in combination with a geographic perspective on the sharing economy – suggest that removing or relaxing some of these microentrepreneur requirements may address some of the problems raised by our findings.

In summary, this article makes the following contributions to the literature on the sharing economy:

(1) We present **a broad examination of the role geography plays in the sharing economy** and solidify the importance of a geographic perspective in

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

the sharing economy literature. In particular, **our results point to the influence of four key principles from human geography**: “Big Sort” residential clustering, geographic variation in population density, distance decay, and mental maps.

(2) We present evidence that the interaction between common sharing economy platform design decisions and these four geographic principles lead to **structural geographic biases in the sharing economy**, biases that reinforce existing advantages. Specifically, our results suggest that high-population density, high-income neighborhoods receive the largest benefits from the sharing economy and poor urban neighborhoods and outer-ring suburbs receive fewer benefits.

(3) We find evidence that, due to the pervasive correlation between poverty and race/ethnicity in the United States and many other parts of the world, **in many cases, black and Latino neighborhoods tend to be less well-served by the sharing economy.**

(4) We discuss the **design implications of our research**, including evidence from our studies that points to means by which the benefits of the sharing economy may be more widely distributed.

(5) Finally, this work makes a lower-level, methodological contribution: **this article introduces spatial Durbin models** to the human-computer interaction literature and discusses why spatial Durbin modeling is important for robustly understanding many sharing economy geospatial processes (and many geographic HCI processes more generally). To support the wider adoption of Durbin modeling, we have released our modeling code with this article.

Below, we first describe in detail the work that motivated this article, including providing brief overviews of our four geographic principles. We then present our methods and results for our TaskRabbit study. Next, we describe our geographic analysis of UberX, introduce spatial Durbin modeling, and show how Durbin modeling can robustly identify structural geographic biases in UberX wait times. We conclude with a summative discussion that cuts across both studies and provide an overview of design implications for the sharing economy.

2. RELATED WORK

2.1. Relevant Principles from Human Geography

This article examines the sharing economy with a lens informed by geography, and specifically the large branch of geography known as *human geography* (c.f. [Bjelland et al. 2013]). Four key principles from human geography define our geographic lens: (1) residential clustering (i.e. the “Big Sort”), (2) structured geographic variations in population density, (3) distance decay, and (4) mental maps. It is likely that other principles from human geography also play a role in the variable success of the sharing economy, a point to which we return in the discussion section. However, we focused on these four principles because (1) they have been found to figure into many similar geographic processes (e.g. transportation geography) and (2) they have been observed to play a role in other online social systems that have geographic footprints (see below). This led us to hypothesize that these four factors would also have an impact on key sharing economy processes given the sharing economy’s inherently geographic nature. Below, we describe each of these four principles in more detail.

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Residential clustering – in which people of similar characteristics reside close to one another – is a key property of the human geography of nearly all places around the world, and has been well-known in human geography and related fields for decades (e.g. [McKnight 1997; Bjelland et al. 2013]). For those familiar with the social networks literature, residential clustering can be understood as a type of "spatial homophily", and indeed this term has been used to describe similar phenomena (e.g. [Zhang and Pelechrinis 2014; McPherson et al. 2001]). Within North America, residential clustering occurs along racial, ethnic, and socioeconomic lines, among other dimensions [Brunn, Jack F. Williams, et al. 2003; McKnight 1997]. As we note below, clustering in some North American cities has become somewhat extreme, with tremendous socioeconomic (and other) gradients occurring across a metropolitan area. For instance, as of 2010, in the New York City metropolitan area, 78 percent of black residents would have to move to match the geographic distribution of white residents of the metropolitan area [Frey 2015; Frey and Myers 2005]. The same is true for 62 percent of those of Latino descent. Residential clustering in the United States (along with its concordant challenges) was the subject of the prominent book "The Big Sort" by Bill Bishop [2008], which has led to the widespread use of the term "Big Sort" to refer to residential clustering. As such, we adopt this terminology in this article.

A longstanding subject of interest in the economic and urban geography communities has been understanding and modeling variations in population density across urban (and rural) areas (see [Brunn, Jack Francis Williams, et al. 2003] for an introduction and overview). These variations occur in structured – but diverse – patterns in cities and regions around the world. In North America, due to the character of local transportation networks, work/life behaviors and other factors, areas with very high population density tend to occur in city centers, which can have high socioeconomic and low socioeconomic status regions (as per the "Big Sort" phenomena). Outside city centers, density nearly always decreases, and in the suburban areas around cities (prior to entering rural areas), one usually finds low density, high-SES regions. These patterns are very much present in our Chicago region study area and play a key role in explaining the structured geographic biases in the sharing economy we see below. It is important to note that outside of North America, population density patterns (and related SES patterns) can vary, resulting in different impacts on the sharing economy. While we briefly address implications for non-North American cities below, future work should seek to extend our research to the other metropolitan (and rural) structures that exist around the world. Brunn et al. [2003] provides an overview of different metropolitan area structures that may be useful for this investigation.

The third major geographic principle considered in this article is distance decay, or the tendency for interaction between two places to decrease as the distance between the places increases [Dennett 2012; Bjelland et al. 2013; Reilly 1931; Stewart 1948]. Distance decay plays a role in a tremendous variety of human geographic process (and many processes from physical geography as well), with trade patterns [Disdier and Head 2008] (trade declines with distance), transportation behavior [Reilly 1931] (destination choice is often largely defined by distance), and information dissemination [Takhteyev et al. 2012; Odlyzko 2015; Wheeler and Mitchelson 1989] being some of the most well-known processes in which distance decay is a primary factor. Closely related to distance decay is the modeling of distance as a cost function in economic geography, leading to location-allocation problems [ESRI 2016] (e.g. what's the optimal place to put my Coca-Cola bottling plant given transportation costs of water, syrup, etc.). Distance decay has also been observed in communication and collaboration patterns in online communities. Liben-Nowell, for instance, established that roughly two-thirds of

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

friends on Live Journal in 2004 could be attributed to a notion of geographic distance [Liben-Nowell et al. 2005], and similar phenomena have been observed in other social networks [García-Gavilanes et al. 2014; Takhteyev et al. 2012; Scellato et al. 2011]. Along the same lines, researchers have also shown that even Wikipedia contributions are subject to distance decay, with the likelihood of an editor contributing to an article about a place being a function of distance to the place [Hecht and Gergle 2010; Sen et al. 2015]. As we will show below, distance decay – when coupled with Big Sort processes along SES dimensions – makes distance an indirect agent of structural geographic bias in the sharing economy.

Finally, mental maps are, broadly speaking, the representations of the world that each individual has in their minds, both in terms of the geometry of the world and the attributes of those geometries. Work on mental maps dates back at least to Lynch's well-known 1960 book *The Image of the City* [Lynch 1960], and has continued for many decades, including prominent works by Gould and White [1986] and Matei et al. [2001]. Matei et al. focused on how communication infrastructures (e.g. mass media), coupled with "Big Sort" phenomena, has resulted in dramatically varying "comfort" levels across metropolitan areas. Namely, people from one type of area – e.g. high-SES areas that tend to be populated by people of certain races and ethnicities – feel unsafe and otherwise "uncomfortable" in other types of areas, and vice versa. Critically, the mental maps literature also points to a certain degree of ignorance associated with these comfort contours. That is, people tend to be less knowledgeable about places they feel less comfortable, both in terms of important attributes of these areas like crime rates [Matei et al. 2001], but also in terms of the geometries of these areas [Ladd 1967]. Mental maps and distance decay also have overlap, with knowledge about an area being in general inversely associated with the area's distance from one's home region [Gould and White 1986]. Our findings below point to mental maps – especially the associated knowledge and comfort factors – as playing a key role in sharing economy microentrepreneur decisions. This is particularly true of their decisions about whether or not to provide service in a certain area.

A key theme present in our human geographic principles is that socioeconomic status plays an important role: SES is one of the primary dimensions on which a "Big Sort" occurs in most cities around the world, SES and population density have important interplay (especially with respect to low-density suburbs), and people's mental maps and corresponding comfort and knowledge levels tend to vary across neighborhoods of differing SES [Ladd 1967; Matei et al. 2001]. As such, below, we adopt SES as a primary query mechanism with which to explore these geographic factors, using SES as an independent variable in both studies. As we will see, SES indeed sheds light on the impact of our human geographic properties in the sharing economy. We augment SES as an independent variable with other variables of interest, particularly targeting distance to capture a detailed picture of distance decay and population density to understand how its variation affects the sharing economy.

It is important to note that there are also other themes present in our geographic principles, most notably ethnicity and race, which are also important "Big Sort" dimensions and mental map determinants. Indeed, a number of factors in North America's history and present have led SES and race and ethnicity to be strongly correlated. In our studies below, in addition to using SES, we also discuss implications for race and ethnicity where we have sufficient data to support this analysis.

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

2.2 Measuring Geographic Disparity in Social Computing Platforms

Sharing economy platforms exist in a broader universe of geospatial sociotechnical systems, which also includes geowikis (e.g. OpenStreetMap and Cyclopath [Priedhorsky et al. 2010]), physically-situated citizen science projects (e.g. [Sheppard et al. 2014; Sullivan et al. 2009; Moran et al. 2014]), among others (e.g. location-based social networks, geotagged social media). Research has shown that these systems can have substantial variations in geographic coverage (e.g. [Haklay 2010; Mashhadi et al. 2012; Mooney et al. 2010; Quattrone et al. 2014; Zielstra and Zipf 2010]). For instance, Quattrone et al. show that more egalitarian (measured as a lower Power Distance) countries with higher incomes (GDP) have better geographic coverage in OpenStreetMap [Quattrone et al. 2014]. Similarly, Haklay et al. find that within Britain, the most disadvantaged areas (according to the Index of Deprivation, an aggregate metric of SES factors) tend to have worse coverage than those areas that are less disadvantaged [Haklay 2010].

Along the same lines, Li et al. studied geotagged social media platforms (Flickr and Twitter), and demonstrated that low-SES areas and rural areas both have worse coverage (less data) than higher SES and urban areas [Li, M. F. Goodchild, et al. 2013]. Similarly, researchers have found that found that people from rural areas produce less geotagged social media (e.g. posts to Twitter, Flickr, or Foursquare) per capita than their urban counterparts [Hecht and Stephens 2014] and that this information is less likely to be produced by locals [Johnson, Sengupta, et al. 2016], and Johnson et al. identified that peer production crowdsourcing is less effective at describing urban areas than it is rural areas [Johnson, Lin, et al. 2016]. Even location-based games with social components (e.g. Pokémon GO) have been found to have similar coverage issues [Colley et al. 2017]. As we will see below, it is likely that many of these findings can be attributed to the same geographic principles discussed in this article. Exploring this in more detail is an important direction of future work.

2.3 Sharing Economy Research

Sharing economy platforms (e.g. TaskRabbit, Uber, and Airbnb) have become a subject of intense public discussion, which has led to increased attention from researchers (e.g. [Edelman and Luca 2014; Edelman et al. 2015; Quattrone et al. 2016; Ikkala and Lampinen 2015; Dillahunt et al. 2016; Lee et al. 2015; Raval and Dourish 2016; Thebault-Spieker et al. 2015; Ge et al. 2016; Hughes and MacKenzie 2016; Ma et al. 2017]). Initial work has focused on addressing non-spatial issues, usually involving the adaptation of research questions from the virtual crowdwork literature (e.g. [Kittur et al. 2011; Bernstein et al. 2010; Ipeirotis et al. 2010; Benson et al. 2015]) to the sharing economy context. For example, Teodoro et al. [2014] conducted a qualitative study to investigate the motivations of workers in TaskRabbit and Gigwalk (a platform broadly similar to TaskRabbit but with different primary use cases). They found that monetary compensation and control of working conditions (time of day, rate of pay, the tasks they do) were primary factors in workers' motivation to participate in these platforms as micro-entrepreneurs. Alt et al. [2010] independently developed an experimental system similar to TaskRabbit. They asked people to complete tasks using a smartphone and observed their behavior. They found that workers were more willing to do tasks that were, for example, relatively straightforward (e.g., taking photos) and that could be done before and after business hours. Ikkala and Lampinen [2015] explored the social role that payment plays in an Airbnb study, and discuss how payment modifies the social relationship between hosts and guests.

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

One thread of sharing economy work has focused on non-commercial "peer-to-peer exchange" platforms (a term used instead of "sharing economy", which some have problematized [Schor 2014; Schor and Fitzmaurice 2015]). One prominent example in the HCI community is the body of work on timebanking, or time-based currency made possible through technological support (e.g. [Shih et al. 2015; Bellotti et al. 2014]). A thread of work related to these non-commercial systems focuses on the social dimension of the commercial sharing economy. This thread suggests that even in commercial systems, there is a social 'economy' between the worker and the person receiving service, and that this dimension is a critical attribute of what people like about and expect from these systems [Heyman and Ariely 2004]. While our work here focuses on large, commercial sharing economy platforms, examining the role of geography in non-commercial peer-to-peer-exchange platforms is an important direction of future work.

Another major focus of sharing economy research has involved examining the challenges associated with being a sharing economy "microentrepreneur". Lee et al. [2015] found that tensions arise between supervisory task assignment algorithms and microentrepreneurs in ride-hailing systems like UberX and Lyft. Rosenblat and Stark [2015] consider the power structures that arise from the reputation system in UberX, and what effect his has on drivers. In a similar vein, Raval and Dourish [2016] find that part of what sets "good" drivers apart in UberX is the emotional labor they carry out, and argue for the importance of recognizing this labor. Glöss et al. [2016] compare the differences in work and perspectives between taxi and Uber drivers. Ahmed et al [2016] consider a very similar juxtaposition to Glöss et al., but focus on a different, international context: the Ola ride-hailing platform in India. Ola connects passengers to rickshaw rides, similar to UberX. Ahmed et al. explore the differences between auto-rickshaw drivers who do not use a sharing economy platform, and those who do.

Recent work has examined the relationship between demographics and worker participation in sharing economy platforms, and this research played a major role in informing the research present in this article. For instance, Lee et al. [2015] found that UberX drivers often turned off 'driver mode' in the Uber app when they are near areas where they feel unsafe or avoided unsafe areas entirely, a finding that provides key context to some of our results below. Dillahunt and Malone [2015] identified that there are barriers to participation in the sharing economy for people who live in low-SES areas. Dillahunt and Malone's work, in particular, provides important scaffolding for our design implications as we discuss below. Along the same lines, Edelman and Luca [2014] found that black Airbnb hosts systemically earn less than non-black hosts, and that users with stereotypically African American names Airbnb are less likely to be accepted as guests compared to identical profiles with stereotypically white names [Edelman et al. 2015]. A similar preliminary set of results was recently identified in UberX in Seattle and Boston by Ge et al. [2016] with respect to wait times and cancellations.

Our research is most directly motivated by recent work on the sharing economy that has begun to identify geographic phenomena as potential factors of interest. For instance, Teodoro et al. and Alt et al. (as well as others) observed that how far people would need to travel to a task appears to influence their attitude toward the task. This is a finding that we both replicate and formalize in a controlled experiment on TaskRabbit, identifying this phenomenon as a manifestation of distance decay. Through modeling and qualitative analysis, we are also able to provide the first evidence that distance decay in the sharing economy – coupled with "Big Sort" residential self-selection – has substantial effects on the availability of sharing

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

economy services and the price of these services (and, subsequently on the bias in the geographic effectiveness of these services).

Similarly, Quattrone et al. [2016] explore the geographic and demographic factors that contribute to Airbnb growth and penetration in London. Through this study, they make policy recommendations based on their findings that would allow regulators to be more responsive to the changing attributes of Airbnb. Expanding on Diakopoulos's past work on UberX surge pricing [2015] (variable prices based on demand) in a recent blog post for the *The Washington Post,* Stark and Diakopoulos [2016] describe initial explorations of UberX availability (with a particular focus on surge pricing) with regard to socioeconomic and demographic attributes of the Washington, D.C. area. Hughes and MacKenzie [2016] performed a similar analysis in Seattle. Among other extensions of this work (see below), we are also able to replicate and formalize the findings from both of these studies using a robust statistical framework (spatial Durbin modeling) that can provide new insight given the spatial properties of relevant data.

In summary, this article builds on existing literature by demonstrating that four geographic principles – "Big Sort" phenomena, structured variations in population density, distance decay, and mental maps – can play a key role in defining the relative effectiveness of the sharing economy in a given region. Critically, we also show how these principles interact with design decisions in sharing economy platforms to create important structural geographic biases that disadvantage people living in low-density and poor areas. The robustness of these contributions is supported by this work being the first to examine multiple sharing economy platforms (and multiple *types* of platforms) with a geographic lens and by our adoption of a new statistical framework from the domain of spatial statistics (spatial Durbin modeling). We believe this framework will prove useful for other researchers in future examinations of the sharing economy. Overall, this article, supported by the previous work in this space, paints a clear two-part picture: (1) when it comes to the sharing economy, geography matters and (2) one way it matters is that key human geography principles interact with sharing economy design decisions to create structural geographic biases.

3. STUDY 1: TASKRABBIT (MOBILE CROWDSOURCING SHARING ECONOMY PLATFORM)

We begin the discussion of our empirical work with our analysis of TaskRabbit. TaskRabbit is a canonical example of the mobile crowdsourcing branch of the sharing economy. As noted above, it is used by task requesters for the completion of physically-situated tasks such as delivering flowers, building IKEA furniture, and helping task posters move large items [TaskRabbit Support 2016].

The goal of our TaskRabbit study [Thebault-Spieker et al. 2015] was to understand the effectiveness of the TaskRabbit platform in different regions with respect to our four geographic principles. To address this goal, we first had to define effectiveness within a TaskRabbit context. We did so by decomposing the notion of TaskRabbit effectiveness into two basic dimensions: (a) the ability to find a microentrepreneur to complete a task (i.e. the *willingness* of a worker to do a task) and, if a microentrepreneur is willing to do the task, (b) the *price* at which the microentrepreneur will complete the task. These dimensions led directly to our two research questions for this study, each of which we explicate in turn immediately below.

Research Question #1: RQ-Willingness: Where will participants in TaskRabbit be willing to go to complete tasks?

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

As noted above, SES plays a key role in three of our human geographic principles: the "Big Sort" (i.e. people of similar SES cluster together), population density variation (i.e. some parts of metropolitan areas like the suburbs tend to be wealthier than others), and mental maps (e.g. people who live in higher SES areas tend to know less about low-SES areas and have low comfort levels in these areas). As such, in this study and in the UberX study below, we used SES as a straightforward probe into the function of these principles in the sharing economy. This was a decision that turned out to be supported in our results (see below). In particular, the human geography literature on our principles suggests that workers would be less willing to complete tasks in low-SES areas, which amounted to our first hypothesis: **H-Willingness-SES**.

The one principle that is not directly addressed through an investigation of SES in this context is distance decay. As such, we also included distance to a task (from a worker's frequently visited areas) as an independent variable. Distance decay suggests that as this distance increases, willingness to complete a task should go down: **H-Willingness-Distance**.

Research Question #2: RQ-Price: How does geography affect how much participants in TaskRabbit request in payment?

At the time of data collection, TaskRabbit had a straightforward auction system for tasks in which workers would bid on tasks posted by users. As such, we believed that the amount workers would charge for a task would be subject to similar processes as their willingness to do the task. Specifically, we hypothesized that distance and task price would be positively correlated (indeed, cost can be a primary mechanism for distance decay, e.g. [Cochrane 1975]) (**H-Price-Distance**) and SES and price would be inversely correlated (**H-Price-SES**).

It is important to note that TaskRabbit's pricing model is subject to frequent iteration, as is the case with many aspects of most sharing economy platforms. As of this writing, TaskRabbit's model has changed to a more complex approach that involves several options for workers and requesters. However, the model still incorporates relatively significant user input in some cases, making it more liable to principles from human geography, something that is in theory not the case with UberX (although driver behavior with respect to surge pricing problematizes this notion [Diakopoulos 2015]). We highlight the role of pricing model design, the differences between UberX and TaskRabbit with respect to pricing, the relationship between willingness and pricing, and innovation in this area in our Discussion section below.

3.1 TaskRabbit study design

To address the above research questions and evaluate the corresponding hypotheses, we developed an experiment and recruited TaskRabbit workers as participants. This recruitment was done in an organic fashion by posting tasks 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 participate in our experiment, for which we paid participants $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 results, we first asked participants a number of questions about themselves, covering topics such as gender, preferred mode of transportation, and activity level on TaskRabbit. We also asked participants to select their home

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

census tract[1] on a map we provided, and to do the same for census tracts that they visited at least once a month.

After participants answered questions about themselves, they began the main portion of the experiment. This portion of the experiment involved prompting participants with census tracts in Cook County, Illinois, which contains Chicago and many of its suburbs[2]. For each census tract, the participant had the option to either check a box labeled "I would not do this task at this location" (RQ-Willingness) or to name what they felt would be a fair price to complete the task (RQ-Price). We did not ask TaskRabbit workers to complete the tasks, only to say if they would complete them and at what price. Figure 1 shows an example of the experiment interface.

Each tract was randomly assigned one of three hypothetical tasks designed to vary the level of engagement with the local area (Table I), an important variable considering the mental maps literature and its findings relating to geographically-variable comfort levels (e.g. [Matei et al. 2001]). These tasks ranged along an engagement spectrum from a task that could be done without leaving a vehicle to a task that required interaction with a person in the area. Tasks were designed so that they would not take more than five minutes.

Each participant received 20 census tracts. 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 in our study area and, in according with Matei et al.'s work on mental maps [2001], the highest-crime[3] and the lowest-crime tracts. The remaining two tracts were repeated from the randomly chosen set of 14 tracts to verify intra-rater reliability. Each 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 participants several open-ended

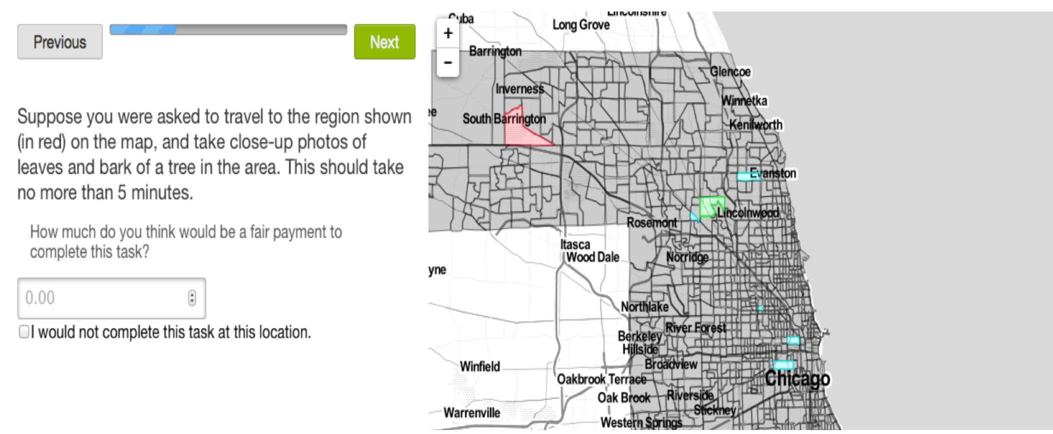

Fig. 1: An example of what participants saw in our TaskRabbit experiment. The green census tract is the worker's self-reported home tract and the blue tracts are those that the worker reported visiting at least once a month. The red tract is the tract about which the worker is currently being questioned (Note: the image is cropped for space and, for privacy reasons, the figure does not depict an actual worker's responses).

[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" [US Census Bureau Geography 2010].
[2] Cook County has a total of 1,317 census tracts.
[3] As reported by the Chicago Police Department (Cook County-wide crime data is not available).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

questions whose answers were entered into text boxes. Specifically, we asked participants about how they made their pricing decisions and why they would not complete certain tasks (if they checked that box at least once).

3.2 TaskRabbit Results

Forty participants completed the experiment (20% of active TaskRabbit workers in Chicago at the time of the study), which we ran during Spring 2014. 57.5% of participants identified as women (42.5% men), which aligns well with gender distribution in the platform overall [TaskRabbit 2016]. The median participant performed a task on TaskRabbit between once a week and once every two weeks. 30% of participants indicated that they complete multiple tasks per week, while only 20% of participants indicated that they complete a task once a month or less.

3.2.1 RQ-Willingness. Because price is irrelevant if a worker will not complete a task, we first sought to understand the geography of worker willingness. 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 participant at least once a month (as indicated in the experiment) [This helped us understand the role of distance decay].
- Median household income of the task tract, as an indicator of socioeconomic status (e.g. [Li, M. Goodchild, et al. 2013; Steward 2009]). As noted above, 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. [This helped us see the effect of the "Big Sort", population density effects, and mental maps]
- Task ID, to make sure we understand the effect of distance and median income in the context of a given task.

The model's random effects were intercepts for participant and by-participant slopes for the effects of income and distance. The model's dependent variable was whether the participant had checked the "I would not do this task at this location" box for a given tract. It is important to note that this model used standard mixed effects techniques rather than our more advanced spatial Durbin modeling approach, which is employed in our analysis of UberX data. We discuss the relationship between these two models and their appropriateness for each setting in Section 4.

Table I. Experiment Tasks and Their Hypothesized Engagement Level

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

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

To operationalize distance, we used travel time rather than Euclidean distance to better match actual 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 tract (to the task tract) that the participant indicated visiting frequently (more than once a month). The API supports multiple transportation modes, and we calculated travel time with participants' self-reported preferred mode of transportation.

Overall, participants 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 from further analysis.

Table II shows the results of our model. All fixed effects are significant and we find that both *H-Willingness-Distance* and *H-Willingness-SES* are supported. Socioeconomic status of the task location and distance to the tract both have an effect on whether a worker is willing to complete a task, with SES having a significant positive relationship and distance having a significant negative one.

The effect sizes are relatively large. 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 standard median household income in northern Chicago and the Chicago suburbs, with $30K median household incomes common on the "South Side" (as the southern part of Chicago is commonly known).

With respect to travel time, our model indicates that for every hour of travel time there is a substantial decrease in willingness to complete a task. In this case, the geographic interpretation is clear: this result directly validates a rather large presence of distance decay. Specifically, TaskRabbit workers are about 4.3% as likely to complete a task an hour away than they are tasks in their immediate vicinity.

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 the means for each participant) for women was 57.1% but for men it was 77.7%. Our qualitative results below suggest that both distance and crime factors play a role in willingness decisions by women (in part mediated by mental maps), but these

Table II: The results of our *Willingness* model

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

[4] One participant 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).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Table III: The results of our *price* model

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

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.

3.2.2. RQ-Price. We now turn our attention to our analysis of the price participants indicated that they would charge for a task (assuming they were willing to complete the task). We began this analysis 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 the repeated tracts. The coefficient was r = 0.96 across all participants, indicating that participants' pricing decisions were very consistent. To understand the effect geography has on task prices in TaskRabbit, we built a linear mixed effects model with identical independent variables as our willingness model but with reported task price as the dependent variable.

The results of this price model can be seen in Table III. This table reveals that travel time was positively associated with price, supporting *H-Price-Distance* and the distance-as-cost-function view of distance decay. Indeed, the model suggests that for every hour of travel time, the price goes up at a rate of $9.97/hour. 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.

This, however, is where the important role of "Big Sort" phenomena becomes clear: due to these phenomena, even though SES is not a significant predictor of price, we found that people who live in large low-income areas are indeed likely to be charged more for the same task. To understand how this works, consider Figure 2, which shows the self-reported home tracts of all 40 participants on top of a map of income by census tract in Cook County. Immediately visible in Figure 2 is that very few participants live in the heart of low-income areas. Indeed, most participants seem to live in middle-income areas next to the very high-income portions of northern Chicago (the "North Side"). Only a single participant lives well within the lower-income South Side of Chicago. As a result, low-income residents on the South Side are almost always a large distance away from any given TaskRabbit worker, making *distance an agent of higher prices* for these low-income neighborhoods. In other words, a low-income resident of the South Side would have to pay more to receive a given TaskRabbit service, for instance someone to take care of errands to make time for longer-term goals [Venkatraman 2010; Blow 2015]. Moreover, as per our findings above, South Side residents also likely have a harder time finding a TaskRabbit worker to accept a request for services in the first place.

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

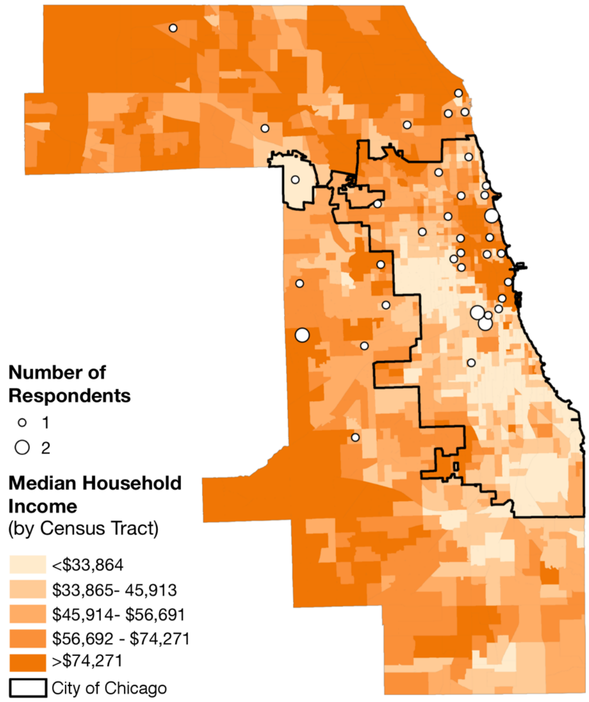

Fig. 2. Experiment participants' self-reported home census tracts and median income in Cook County, Illinois. Very few participants live in low-income tracts. Note that the low-SES "South Side" of Chicago (Chicago is outlined in black) has only one participant, and no participants live in the poorest parts of the South Side. Median income color classes are determined via the quantile method, meaning each class represents a quintile of the household income dataset. Participant are displayed at the centroid of their home census tract.

This result suggests a specific character for the effect of the Big Sort on the sharing economy. We found that most sharing economy workers live near (but not in) high-SES regions. If this result generalizes, people who live in large, low-income districts like the South Side will need workers to travel greater distances to get to their task locations, resulting in longer travel times, and, ultimately, higher prices. 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 more minimal. However, these smaller pockets may get rarer and rarer in a "Big Sort" world.

It is also important to point out that these "Big Sort" effects also play a role in willingness decisions. Since distance and willingness were found to be inversely associated, the fact that TaskRabbit workers live far away from large low-SES areas means that this inverse association will disproportionately affect people who live in these areas. Moreover, since we identified a separate effect in which willingness and SES are positively associated (as SES goes up, willingness goes up), the distance effects and SES effects likely compound each other to make the task willingness in large low-SES areas particularly low.

Lastly, although our consideration of population density largely lies in our UberX study, we do see an important effect for population density here. Figure 2 shows that workers are concentrated in the high-density city of Chicago rather than the low-density suburbs. As such, distance not only reduces availability of TaskRabbit and increases the price of TaskRabbit in lower SES areas, but does the same in suburbs. However, the situation in suburbs is generally quite different: as can be seen in Figure 2 (and is discussed in related work), suburban people have higher incomes than people

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

on the South Side of Chicago, and thus they can potentially afford the increased costs. In addition, in some cases, even people in somewhat remote suburbs are closer to one of our participants than a person in southern Chicago. That said, 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 [Brunn, Jack F. Williams, et al. 2003; Bjelland et al. 2013]). Where this is the case, services like TaskRabbit will likely be drastically more expensive and less available in these areas. As we note above, investigating these phenomena in cities with different socioeconomic segregation patterns is an important direction of future work.

3.3 TaskRabbit Qualitative Results

Thus far, our quantitative models have revealed evidence for the importance of distance decay, "Big Sort" phenomenon and, to a lesser extent, structured variations in population density, when considering the effectiveness of TaskRabbit. We have also seen these principles manifest in structural geographic biases in TaskRabbit, biases that lead TaskRabbit to be both more expensive and less available in low-SES regions in our study area. We now turn to our qualitative results to attempt to help understand *why* these dynamics exist. To do so, a single investigator looked for themes in the textual survey responses, focusing on ideas related to our four geographic principles. Below, we outline the results of this analysis, which identified qualitative data relevant to mental maps (and their interactions with "Big Sort" phenomenon) and distance decay.

3.3.1. Mental Maps. A theme that was very clear in our participants' answers to why they ticked the "will not do [a task]" box is the importance of mental maps, and specifically a large region of low comfort levels in their mental maps corresponding to Chicago's South Side (and to a certain degree the "West Side"). Participants reported that these low comfort levels were driven mostly by perceptions of high crime. Indeed, some participants' responses read as if they came directly out of Matei et al.'s study that examined the role of crime in neighborhood-level comfort assessments in mental maps. For example, consider this response from P27:

> "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." (P27)

P9 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 requesters upon completion), completing a large number of tasks, and not violating any of TaskRabbit's policies. P9 offered similar feedback to P27:

> "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." (P9)

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

P4, a relatively new resident of Chicago, wrote that the comfort level overlay in her mental map also led to similar decisions about whether or not to accept a task. In this case both poverty and crime are mentioned:

> "I only moved to Chicago last May, 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." (P4)

P39 specifically addressed 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." (P39)

P16 was very explicit about how he makes decisions between the contradictory signals from his comfort layer and the desire to increase his income:

> "Whether or not my assumptions of lack of safety were correct, I wouldn't put myself in danger for a few dollars" (P16)

The quotes above make it clear that a key sharing economy decision-making process – whether or not a microentrepreneur agrees to accept a task – is subject to classic mental map effects. These are effects that have been observed in geography and related fields for decades (e.g. [Gould and White 1986; Matei et al. 2001]). As has been observed by Gould and White [Gould and White 1986], Matei et al.[2001], and others, humans tend to ascribe large regions of their mental maps with positive and negative emotions, with Matei et al. [2001] specifically focusing on comfort levels assigned to neighborhoods in mental maps as a function of perceived crime in those areas. Our qualitative results suggest that this is a primary driver behind the results of our willingness model.

These findings also dovetail with recent findings by Lee et al., who, as noted above, did qualitative work with UberX drivers. Lee et al. identified that UberX drivers often manually disable their availability to the UberX platform when they are traveling through what they perceive to be unsafe neighborhoods, a finding that can be easily understood through a mental maps lens. This is also the direct UberX analogy to a TaskRabbit worker not accepting tasks in specific neighborhoods, and is a point to which we return in our UberX study.

Another theme in the above responses that is also present in Gould and White's and Matei et al.'s results is a lack of geographic nuance in mental maps. While the South and West sides do indeed experience much higher levels of crime than other parts of Chicago, there are pockets of these areas that are quite safe [Chicago Tribune 2016]. However, "Big Sort" processes have driven most people of higher SES out of the South Side and the West Side (as well as many people of White, non-Latino descent; see below), likely leading to large regions that are unknown to people who do not live there, both in terms of geometry and personal comfort levels. This is roughly analogous to a finding observed in mental maps of another major urban city, Boston [Ladd 1967]. The mental maps of our participants clearly are not sufficiently nuanced to support knowledge of the lower crime pockets on the South and West side.

Before moving on to our analysis of the presence of distance decay themes in our qualitative results, it is important to point out that prior work (e.g. [Matei et al. 2001]) suggests that, even though SES and crime are the only two attributes directly cited by

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

our participants in their willingness decisions, race and ethnicity may also be involved. This is a point we address in our discussion below.

3.3.2. Distance Decay. Participants' qualitative feedback supports the finding from our quantitative modeling exercise that proximity of the task location is a very important factor in task willingness and pricing decisions. Here, P4 explicitly discusses the role of distance decay 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." (P4)

Our qualitative data also shed some of light on the specific form of distance decay in TaskRabbit pricing and willingness decisions. Distance decay can take many and multiple forms (e.g. gravity models (e.g. [Reilly 1931; Flowerdew and Aitkin 1982; Stewart 1948]), thresholds (e.g. [Agarwal 2001]), and our qualitative results suggest that distance decay in this context contains a threshold component. Four participants explicitly or implicitly mentioned thresholds in explaining why they said they would not complete a specific task:

> "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." (P31)
>
> "Getting there would take me longer than actually completing the task." (P39)
>
> "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." (P16)
>
> "I didn't think any price would be worth the commute and risk while still offering even a marginally fair price." (P23)

More specifically, these participants suggested that when the cost of commute time (either in raw time or money) rises above a certain level (in two cases the financial or temporal cost of the task), they would no longer be willing to accept the task. This feedback should help guide future work involving modeling distance decay's role in willingness decisions in the sharing economy.

3.4 Summary of TaskRabbit study

Above, we have seen that three of our human geographic principles – the "Big Sort", distance decay, and mental maps – play a key role in the effectiveness of the sharing economy. We have also shown that these principles manifest in structural geographic biases in TaskRabbit, in which people who live in high-SES regions in the urban core gain most of the benefits of TaskRabbit's rendition of the sharing economy (at least in the Chicago area). These biases also mean that TaskRabbit is both more expensive and less accessible to people in low-SES areas. We have also observed a smaller role for our fourth geographic principle, structured variation of population density, observing that prices are also higher (and to a lesser extent, service is less available) in high-SES, low-density suburbs as well. Below, we explore whether these same trends persist in an entirely different rendition of the sharing economy: the well-known ride-hailing platform, UberX.

4. STUDY 2: UBERX (RIDE-HAILING SHARING ECONOMY PLATFORM)

As noted above, the goal of our UberX study is to identify whether the key findings from our TaskRabbit study – the importance of the four geographic principles and their

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

manifestation through specific geographic biases – generalize to UberX. Most (if not all) studies of the sharing economy thus far have focused on a single sharing economy platform. By taking this multi-platform approach, we aimed to gain a more general understanding of the role of geography in the sharing economy (rather than a platform-specific understanding).

We focus our attention in this section on an analysis of UberX *wait times*, a dimension of effectiveness related to our willingness variable in TaskRabbit. In UberX, drivers can show their willingness to pick up a passenger by accepting or rejecting a fare, by avoiding (or spending time in) certain areas or by selectively turning on and off their availability as they approach certain areas (as was mentioned above). All of these expressions of willingness manifest in the amount of time a potential customer has to wait before an UberX driver arrives at her/his location. Moreover, these wait times can be automatically obtained, affording a quantitative understanding of geographic effectiveness just as was the case with TaskRabbit. It is important to note that we do not analyze price in UberX, as this is determined automatically by a basic formula in all cases (a point to which we return below).

More specifically, we structure our investigation of UberX wait times around the following research question:

Research Question #3: RQ-Wait Times: How does human geography affect UberX wait times?

Motivated by the geographic principles considered in this paper and in analogy to our TaskRabbit study, we made two hypotheses with respect to this question. First, we hypothesized that, all other factors being equal, wait times would be higher in low-SES areas than in high-SES areas (**H-Wait Times-SES**). Secondly, we also hypothesized that structured variations in population density would be a significant factor in wait times, with denser areas having more convenient access to UberX (**H-Wait Times-Population Density**).

In addition to highlighting the similarities and differences in the geography of UberX versus that of TaskRabbit – as we will see, there are far more similarities than differences – this section also makes a methodological contribution to the sharing economy literature. Specifically, to test our hypotheses, we adapt *spatial Durbin modeling*, an advanced spatial statistical technique from the natural and social sciences, and show how it is often critical for conducting robust geographic sharing economy analyses. To our knowledge, this study represents the first use of spatial Durbin modeling in the human-computer interaction community[5], and we expect that this work can provide statistical assistance for other sharing economy researchers and researchers in other domains who encounter similar types of spatial data. As such, we dedicate a significant portion of the methods section below to explaining the character, intuition, and proper execution of spatial Durbin modeling.

The remainder of this section will proceed as follows: we first introduce the datasets we utilize in our analysis of UberX wait times. Next, we describe spatial Durbin modeling and explain why it is essential for understanding the geography of the

[5] A search for "spatial Durbin" in the ACM Digital Library returned no results, although we also employed the same Durbin modeling approach in work that will be published at ACM SIGCHI 2017 [Colley et al. 2017]

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

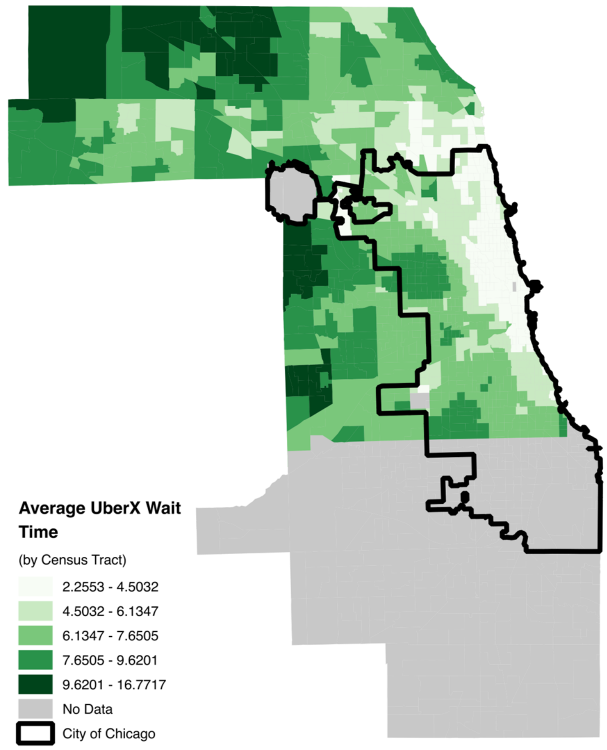

Fig. 3. Average UberX wait times in Cook County, Illinois. In grey are tracts excluded from our analysis because (a) they were not within UberX's operating area (the large block to the south) as of time of data collection or (b) the Uber API did not provide a single ETA value for these tracts.

sharing economy in many cases. Following our discussion of methods, we then present the results of our models, highlighting similarities and differences with our TaskRabbit results and discussing connections to our four geographic principles.

4.1 Datasets

In every metropolitan area where UberX provides service, there is a defined region where the service is available. At the time this analysis was performed (late 2014), this region did not encompass all of Cook County. Specifically, there are 275 census tracts (shown in Figure 3) on the southern end of Cook County that were outside of UberX's operating area. Thus, we excluded these tracts from our study. It is important to note that UberX's service area has since expanded. We return to this issue in the Discussion section, in which we highlight potential "early access" benefits provided to certain types of areas over others in geographic social computing systems.

The tracts in which UberX did not provide service at the time of analysis are systematically poorer than the tracts in which it did offer service. This provides our first evidence that some of the factors associated with effectiveness in TaskRabbit – particularly the "Big Sort" along the SES dimension – play the same role in UberX. More specifically, tracts excluded from UberX's service area had an average median household income of $53,122 ($sd$ = $19,247), whereas the tracts served by UberX have an average median income of $56,369 ($sd$ = $29,761).

We used Uber's Time Estimate API (available through their developer website) to measure wait times in the 1,041 Cook County census tracts within the UberX operating area. We sampled the centroid of each census tract every hour for a period

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

of 7 days, leaving us with 168 samples per tract in the ideal case. Uber's API never provided wait times for the three census tracts containing the Chicago area airports, which is likely due to Chicago city ordinances that prohibited Uber from providing service at airports at the time [Hilkevitch 2014][6]. Thus, we excluded these tracts from our data, leaving us with 1,038 census tracts for our analysis.

In some cases, Uber's API did not provide a wait time every time a tract was sampled, the reasons for which are unclear. However, ninety-eight percent of tracts in our sample received a wait time more than 80% of the time, and the census tract with the fewest samples received valid responses approximately 60% of the time (100 of 168). We compute the mean of all wait times for each census tract, and use this *average wait time* as the core dependent variable in our analysis of UberX.

To examine H-Wait Times-SES, we used the same SES data as we did for our TaskRabbit study: median household income (MHI) (in $10K increments) from the United States Census' American Community Survey 2006-2010 dataset. Again, we log-scaled this variable to reduce the effect of a long-tailed wealth distribution. To examine H-Wait Times-Population Density, we utilized United States census data on the number of people per square kilometer in each tract. We also log-scaled this variable, again to reduce the skew of a long-tailed population density distribution.

Finally, we note that while the methods we used to study TaskRabbit enabled us to compute a "distance from home region" variable, this was not the case for our UberX work: UberX's API does not give a regular starting point for a given UberX driver. As such, in this study, our ability to speak to distance decay directly is limited (although, as is discussed below, our results below do suggest several indirect findings).

4.2 Spatial Modeling

If the datasets we considered in this study had not been not geospatial, our modeling task would have been straightforward. Specifically, using standard regression modeling techniques like ordinary least squares (OLS), we could have assessed whether there were significant relationships between MHI and population density (independent variables) and UberX wait times (dependent variable), as well as determined the effect sizes of these relationships.

However, the geospatial nature of our datasets demand that our methods be considerably more sophisticated. For instance, let us consider a wealthy census tract that is surrounded by less wealthy census tracts (e.g. as is the case near the University of Chicago on the "South Side"). One might imagine that this tract's UberX wait times may be affected by the fact that its neighbors are less wealthy. After all, UberX drivers might not want to drive through poorer areas to get to this tract, which they may perceive to be less safe (as we saw in the case of TaskRabbit), or they may prefer to stay in an area that has consistent and widespread high incomes. Conversely, a poorer tract near richer tracts may see opposite effects. However, traditional regression modeling assumes that all samples (tracts) are independent and cannot incorporate potentially critical information about a tract's neighbors' income in its estimates of wait times for the tract. In other words, one can think of traditional regression modeling as failing to account for Big Sort effects when applied in many types of sharing economy analyses (and other types of analyses in geographic HCI).

More generally, when a dataset consists of specific geographic locations associated with attributes – which is the case for UberX wait times – it is common for individual data points to be affected by other nearby data points, or to be *spatially autocorrelated*

[6] These ordinances have since changed [Coale 2016].

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

[Anselin 1988]. The Big Sort is an instance of spatial autocorrelation in which the variables of interest are demographic in nature.

The presence of spatial autocorrelation, including in the Big Sort case, means that the data from one location often is not independent of data from neighboring locations, violating a core assumption of traditional regression modeling techniques like OLS. The output of these techniques – significance, effect sizes, etc. – are all conditioned on the assumption of independence of observations in the independent and dependent variables, something that often will not be true with geographic sharing economy data. As such, specialized techniques are needed, not just to acquire more statistical power and understanding, but also simply to gain reliable insight on the associations in question and the role of spatial relationships in these associations.

When engaging in spatial statistics, a frequent first step is to model the spatial structure of the study area. A common approach to generating a representation of this structure – and the approach we use in this work – is called a *Queen's weights matrix* (there are also other distance-based schemes, e.g. *k-nearest-neighbors*). With a Queen's matrix, the neighbors of a given census tract (or polygon more generally) are assumed to be all the tracts (or polygons) that are directly adjacent through either an edge or a vertex (similar to the moves available to a queen in chess).

Once the spatial structure has been encoded, we can turn our attention to modeling spatial effects between neighbors. These effects can be of three distinct types [Manski 1993]:

1. ***A correlated spatial relationship***: unknown factors lead to similar outcomes (e.g. wait times) between two neighboring locations.
2. ***An endogenous spatial relationship***: an *outcome* (e.g. wait time for a given census tract) for one location is dependent on the *outcomes* (e.g. wait times) of neighboring locations.
3. ***An exogenous spatial relationship***: an *outcome (*e.g. wait time for a given census tract) for one location is associated with the *predictors* of neighboring locations (e.g. the median income of its neighboring tracts).

In the context of our modeling exercise, this means that a given census tract's UberX wait time may be (1) similar to its neighbors because of some unmeasured factors, and/or (2) dependent on the wait times of its neighbors, and/or (3) dependent on the population density and median income of its neighbors (our predictors / independent variables).

Until recently, the majority of the focus in geostatistical modeling has fallen into two camps: modeling the correlated relationships (#1 above) by accounting for any spatial relationships using the error term (known as a *spatial error model*), and modeling the endogenous spatial relationship (#2 above) using the weights matrix to *lag* (or spatially weight) the dependent variable (known as a *spatial lag model*). The third type of spatial relationship – exogenous spatial relationships – had until recently largely been ignored in the natural and social sciences, let alone in the human-computer interaction literature. While most treatments of spatial data in the human-computer interaction literature do not consider spatial autocorrelation at all and instead utilize standard regression techniques for spatial data, to the extent that spatial models have been used, they have been the more traditional spatial lag and spatial error models (e.g. [Johnson, Sengupta, et al. 2016; Johnson, Lin, et al. 2016; Malik et al. 2015]).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Spatial error models may be sufficient when model interpretation is unimportant and addressing independence assumptions while maximizing predictive power is the only consideration [Elhorst 2010]. However, recent work in the spatial statistics literature has argued that both endogenous and exogenous relationships need to be examined when using spatial modeling to shed light on the underlying spatial processes [Elhorst 2010; Yang et al. 2013], as we are doing here (and as is common in the HCI community's consideration of spatial data more broadly). In this vein, a more generalized modeling approach – the *spatial Durbin model* – that accounts for both endogenous and exogenous relationships has begun to be recommended as best practice [Elhorst 2010; LeSage and Pace 2009].

A spatial Durbin model can be understood as having multiple versions of each variable corresponding to the endogenous and exogenous spatial relationships discussed, supporting better interpretation of the spatial relationships in the data. In our case, a spatial Durbin approach models the average UberX wait time for a given tract as a linear combination of (1) the values of the independent variables in that tract (as is typical in OLS regression), (2) the average of the values of each independent variable (MHI and population density) in neighboring tracts according to the Queen's matrix (spatially exogenous relationships), (3) the average of the wait times in neighboring tracts calculated in the same fashion (spatially endogenous relationships, i.e. spatial lag term), and (4) an error term, which functions similarly to the error term in a traditional OLS. So, in other words, whereas a traditional OLS regression for our experiment would involve two independent variables and an error term, a spatial Durbin approach applied to our problem would have 2 + 2 (spatially exogenous) + 1 (spatially endogenous) = 5 independent variables and an error term.

In light of the tremendous interpretative advantages of spatial Durbin modeling – none of the key distinctions between the "direct" and "indirect" effects below would be possible without Durbin modeling – we employed spatial Durbin modeling as our primary analytical tool for understanding our UberX data. We were not able to apply spatial Durbin modeling to our TaskRabbit analysis for one critical reason: our TaskRabbit experiment required the employment of mixed effects models and, to our knowledge, mixed effects models *have not yet been integrated with spatial Durbin approaches*. Indeed, spatial Durbin approaches have only recently become feasible; up until several years ago, they were too computationally demanding for common practice [Elhorst 2010]. In the TaskRabbit case, mixed effects models were required to control for the lack of independence between observations gathered from the same individual. In concert with our university's statistical consulting center, we determined that the violations of independence that are due to observations coming from the same participant (which are handled by mixed effects modeling) were more serious than those due to spatial autocorrelation (thanks in large part to the fact that we were not considering many immediate neighbors in our TaskRabbit observations). Because no existing modeling approach (to our knowledge) allows us to account for both types of independence assumption violations, we used the mixed effect models discussed above. As we will see below, we found nearly identical high-level results in our TaskRabbit and UberX analyses, adding credence to both the high-level results and the modeling choices in each.

The majority of future geographic sharing economy research will likely not face the challenges we did with TaskRabbit and will be able to gain the advantages of spatial Durbin modeling as we do here with UberX. Indeed, the two studies that most directly resemble our research – Quattrone et al. [2016] and Stark and Diakopoulos [2016] – did not face the statistical challenges associated with our TaskRabbit study. To

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

facilitate easier adoption of spatial Durbin modeling, we are releasing our modeling code along with this article (see below).

4.2.1 Interpreting Spatial Durbin Models. In general, when examining the results of Durbin models, traditional outcome metrics like the value and significance of coefficients (betas) are far less useful for interpretation than a series of specialized metrics, in particular the *Rho* term, *indirect effect* values, and *direct effect* values.

The Rho term captures the effect of spatial diffusion in the dependent variable (wait times) as one neighbor's value affects another neighbor's value, which then affects another neighbor's value, and so on. More specifically, the Rho term encapsulates endogenous spatial relationships and in the context of our work, describes how a given tract's wait time should be affected by the wait times of its neighbors (endogenous spatial relationships). This term is not interpreted like a traditional model coefficient, but instead is multiplied by the spatially-weighted average of neighbors' measured wait times.

The need for direct and indirect effect values arises out of the fact that spatial dependence invalidates the interpretive benefits of model coefficients in traditional regression (e.g. OLS). Traditional approaches compute regression coefficients through partial derivatives of the regression formula with respect to each independent variable. Because of spatial dependence, however, when these partial derivatives are computed in a spatial model that incorporates data from the neighboring area, any given partial derivative will in turn be dependent on the values of the neighboring tracts' partial derivatives. These feedback loops, caused by the spatial structure, get built into the models. This means that the variable coefficients in a spatial model cannot be interpreted directly because the partial derivatives are not orthogonal [Elhorst 2010].

On the surface, these feedback loops seem like a challenge to the interpretative power of spatial Durbin modeling. However, LeSage and Pace [2009] introduced direct effects and indirect effects to explicitly address this issue. Indirect and direct effects are calculated by averaging across all of the relevant partial derivative values at every location (calculated based on the lagged value of the variable in question). We follow Yang et al. [2013] and use a Markov Chain Monte Carlo (MCMC) approach to randomly permute input data in order to estimate the *average effects* and generate an average over the permuted output[7].

Direct effects describe average relationships between the dependent and independent variables that are analogous to the relationships modeled in traditional regression approaches. Specifically, a direct effect for a given independent variable describes the average impact the value of that variable *at a specific location* has on the value of the dependent variable *at that location*. In the context of our work, this means, for instance, the effect the median household income of a given tract has on the wait times of that specific tract. Each independent variable has its own direct effect.

Conversely, indirect effects model the relationship between the value of a dependent variable at a given location and the values of independent variables at *neighboring locations* (spatially exogenous relationships). In our analysis, indirect effects capture, for instance, the effect of the average median household income of neighboring census tracts on a given census tract's wait time. Like is the case with direct effects, each independent variable has its own indirect effect value (so each

[7] Computed using the `impacts` command in the `spdep` R package [Bivand and Piras 2015; Bivand et al. 2013]

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

independent variable in our model has both a direct effect value and an indirect effect value).

When describing the results of a spatial Durbin model, it is considered best practice [LeSage and Pace 2009] to present the Rho term (endogenous effect), the modeled coefficients, and the direct and indirect effects (exogenous effects), but interpret only the endogenous and exogenous spatial relationships. This is due to the unclear meaning of the standard modeled coefficients. We follow this best practice below.

4.3. UberX SES Results

Table IV shows the model coefficients for our UberX wait times model. In our model, the *endogenous* relationship between a tract and its neighboring wait times is quite strong: 92% of the average wait time of immediate neighbors is contributed to the wait time of a given tract. This is intuitive: a tract should not have drastically different UberX wait times than its neighboring tracts due to the nature of wait times. For instance, if a tract has five neighbors and the sum of their average wait times is 1,500 seconds (mean = 300 seconds), the neighboring tracts would contribute 259 seconds (92% of the 300 second mean) to the predicted wait time of the tract in question.

Table V shows the direct and indirect effects of our independent variables. The table reveals that, when examining the entirety of UberX's service area across Cook County, population density is significant in both its direct and indirect effects (supporting H-Wait Times-Population Density). The indirect effects for population density suggest a strong and inverse relationship between population density and wait times (note that the indirect effect for population density in Table V is both large and negative). Specifically, if the average population density across all of a tract's neighbors were to double, we would expect a decrease in average wait time of approximately 70 seconds for that tract. This is not an extreme scenario: the mean population density in our study region is 6,183 people/km^2 and the standard deviation is 7,616 people/km^2. This means that tracts in a very dense area should expect much lower wait times than areas where there are fewer people per km^2.

While significant, the direct effects of population density – i.e. the role played by the population density of the tract in consideration itself – had a much smaller effect size. A doubling of a specific tract's density would only lead to an average decrease in wait time of approximately 3 seconds, for that tract. This is a trend that we will see repeated below: the characteristics of a tract's region appears to matter more than the characteristics of the tract itself.

More generally, the results in Table V substantiate the importance of the principle of structured variation of population density that we observed in our TaskRabbit study. Specifically, the sharing economy seems to be significantly less effective in the suburbs relative to the central city. For UberX, this finding is quite visible in Figure 3, where we see wait times of over 10 minutes in the very-low-density distant suburbs and under 3 minutes in dense urban cores.

While Table V strongly suggests that structured variation in population density is a prominent factor in UberX wait times, it displays less clarity about income's role. Table V shows no significant direct effects for income and only marginally significant indirect effects (providing little support for H-Wait Times-SES at this stage). With regard to these indirect effects, it appears that if a tract's neighboring region became

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Table IV: The Results of our Cook County UberX *Spatial Durbin* Model

Fixed Effects	**Estimate**	**p-value**
Rho (Weighted effect of neighbors' wait times)	0.92	< 0.001
$\log_2$[people / km^2] (population density)	-0.37	n.s.
lagged $\log_2$[people / km^2] (lagged population density)	-5.63	0.002
$\log_2$[Tract income in $10,000] (median income)	6.09	0.09
lagged $\log_2$[Tract income in $10,000] (lagged median income)	-10.21	0.02
Intercept	114.88	< 0.001

Note: Coefficients in grey are commonly reported, but not interpreted.

Table V: The *Direct* and *Indirect* Effects of our Cook County Model

Direct Effects	**Estimate**	**p-value**
$\log_2$[people / km^2] (population density)	-3.01	0.03
$\log_2$[Task tract income in $10,000] (median income)	4.02	n.s.
Indirect Effects	**Estimate**	**p-value**
$\log_2$[people / km^2] (population density)	-69.09	< 0.001
$\log_2$[Task tract income in $10,000] (median income)	-53.68	0.08

more wealthy, UberX wait times would decrease in that tract. However, at least at this point, we only have marginal confidence in this relationship.

Table V shows results for a model that considered the entirety of Cook County. However, many sharing economy decisions and debates occur at the municipal (city) level (e.g. [D'Onofrio and Thomas 2016; Gaines 2016; Byrne 2016]), where there tends to be less variation in population density. To understand the relationships between income, population density, and wait times within a central city itself – rather than an entire metropolitan area that includes suburbs and exurbs – we re-ran our model focusing only on census tracts that fall within the borders of the city of Chicago. We present the coefficients of this model in Table VI and the average direct and indirect effects in Table VII. Table VI shows that, as expected, in our Chicago model, the endogenous interaction between a tract's wait time and its neighboring wait times is quite high, just as it was for our Cook County model. A tract's immediate neighbors contribute 95% (*Rho* = 0.95) of the average of the neighboring wait times to the predicted wait time (as opposed to 92% in the Cook County model).

Table VII shows that we identified no statistically significant direct effects in our Chicago-only model. That is, the income and density values of a specific tract in Chicago do not seem to play a role in that specific tract's wait time. However, we do see a *significant and substantial indirect effect for median income*: if the income of an area goes up, wait times go down by a large margin (supporting H-Wait Times-SES). To be more specific, our model suggests that if a Chicago census tract's neighbors experienced a doubling in their average median household incomes, we would expect to see that tract's UberX wait time decrease by over 3 minutes and 10 seconds (190.6 seconds) on average. This suggests that while individual poor census tracts surrounded by higher-income tracts should benefit from the wealth of their neighbors, UberX is

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Table VI: The Results of our Chicago UberX *Spatial Durbin* Model

Fixed Effects	Estimate	p-value
Rho (Weighted effect of neighbors' wait times)	0.95	< 0.001
$\log_2$[people / km^2] (population density)	-0.07	n.s.
lagged $\log_2$[people / km^2] (lagged population density)	0.52	n.s.
$\log_2$[Task tract income in $10,000] (median income)	5.29	0.04
lagged $\log_2$[Task tract income in $10,000] (lagged median income)	-14.55	< 0.001
Intercept	33.29	0.05

Note: Coefficients in grey are commonly reported, but not interpreted.

Table VII: The *Direct and Indirect* Effects of our Chicago UberX Model

Direct Effects	Estimate	p-value
$\log_2$[people / km^2] (population density)	0.19	n.s.
$\log_2$[Task tract income in $10,000] (median income)	-0.41	n.s.
Indirect Effects	**Estimate**	**p-value**
$\log_2$[people / km^2] (population density)	8.95	n.s.
$\log_2$[Task tract income in $10,000] (median income)	-190.61	< 0.001

less effective in regions of Chicago where there is more widespread poverty. This can be seen in Figure 3, in which the low-income neighborhoods in the southern and western areas of Chicago have much higher wait times. We discuss how this may be related to mental maps in Section 5 below.

The results in Table VII can be read as strong support for the Big Sort's influence on the sharing economy within regions of similar population density (e.g. within the city limits of Chicago). While the SES of a specific tract does not seem to have an impact on sharing economy effectiveness in that tract (as instantiated by UberX wait times), the SES of a tract's neighborhood has a substantial impact the tract's wait times. This leads to high wait times in large low-SES areas (e.g. the South Side) and low wait times in large high-SES regions (e.g. northern Chicago). This also means that even if a neighborhood in a low-SES area begins to improve its SES, there will likely continue to be a damper on sharing economy effectiveness in this neighborhood. If those who argue that the sharing economy will become a dominant economic paradigm are correct, this is a troubling implication.

In the section that follows immediately below, we provide further discussion of these results in the context of our TaskRabbit results and four geographic principles.

5. DISCUSSION

In our investigation of the geography of the sharing economy, we examined two different sharing economy platforms using diverse methodologies that ranged from quantitative and qualitative analyses of survey results to the application of spatial Durbin autoregressive models on data gleaned from APIs. In both cases, however, we found very similar high-level findings regarding the geography of the sharing economy: the sharing economy is more effective in dense, wealthy neighborhoods and

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

significantly less effective in suburbs and low-income urban neighborhoods. Moreover, our results pointed to the underlying geographic principles responsible for these structural geographic biases: the Big Sort, structured variations in population density, distance decay, and mental maps.

In this section, we discuss the implications of these findings along several key dimensions: the role of race/ethnicity, suggested improvements to the design of sharing economy platforms, and directions for future work.

5.1 Examining our findings with a lens informed by race and ethnicity

As noted above, due to the tremendous economic inequalities that occur across racial and ethnic lines in the United States (and elsewhere), SES and race and ethnicity tend to be closely linked. Indeed, we observed that the percent of the population that self-identifies as *white (non-Latino)* has a strong correlation with income in our study area (r = 0.67). *White (non-Latino)* is a demographic variable provided by the U.S. census whose inverse (the percent of the population that does not identify as white or Latino) is often interpreted as the percent of the population that identifies as a racial or ethnic minority [US Census Bureau 1999].

Because of this correlation, we hypothesized that many of the patterns we saw with SES and the Big Sort would also occur with race and ethnicity and the Big Sort. To test this hypothesis, we replaced SES as an independent variable in both the TaskRabbit and UberX models. For TaskRabbit, we found that the percentage of the population that is white (non-Latino) was a marginally significant predictor of willingness (p = 0.06). As was the case with SES, we did not observe an explicit effect for white (non-Latino) with respect to price. However, due to the correlation above, poor neighborhoods tend to be minority neighborhoods in Chicago (and in many other places in the world), so we also observed the same distance decay effects with respect to white (non-Latino) as we did with SES.

We identified a similar finding for UberX wait times as we did for TaskRabbit willingness: in a version of our Chicago-only spatial Durbin model with SES replaced by the percent of the population that is white (non-Latino), we saw a significant indirect effect for percent white (non-Latino). If the area around a given tract changed from 0% white (non-Latino) to 100% white (non-Latino), we would expect to see that tract's UberX wait time decrease, on average, by *317 seconds (5 minutes and 17 seconds)*. In other words, if an entirely non-white area in the city of Chicago were to see a complete demographic shift towards being entirely white, our results suggest that tracts in that area may see UberX wait times decrease drastically, making the UberX service much more effective. This suggests that while individual non-white census tracts surrounded by white neighbors should benefit from better UberX service, UberX is *less effective in* regions of Chicago where there are large minority populations. Interestingly, these results dovetail with very recent findings by Ge et al. [2016] that suggest that people with African American-sounding names wait longer for UberX service in Seattle.

It is important to note that in the above results, the core geographic principles at work are no different than is the case with SES, they are simply *manifest* in race and ethnicity rather than SES. For instance, just as was the case for low-SES neighborhoods, many minority neighborhoods also see reduced TaskRabbit effectiveness due to the Big Sort and due to distance decay interacting with the Big Sort (with respect to the location of TaskRabbit workers' residences, which tend to be farther away from large minority districts than from large white (non-Latino) districts). Similarly, the Big Sort has the same effect on large minority neighborhoods

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

in UberX as it does on large low-SES neighborhoods. Indeed, as per the correlation mentioned above, many minority neighborhoods and low-SES neighborhoods are one and the same and thereby suffer from identical lower sharing economy effectiveness. More generally, a sharing economy platform that does not serve low-income people will, in general, also fail to serve non-white populations, at least in North America (a key concept in the sociological theory of *intersectionality* [Crenshaw 1989]). Similarly, it is likely that if we investigated sharing economy effectiveness across other demographic properties affected by the Big Sort (e.g. educational attainment, religion, age, political affiliations), we would identify similar findings.

However, regardless of the geographic mechanisms at work, just as was the case for SES, the structural racial and ethnic biases in the sharing economy identified in this section are quite important in their own right. Groups defined by race and/or ethnicity are protected classes in the United States [United States Congress 1964]. If these results are found to generalize across other cities – with sharing economy systems working better in white (non-Latino) neighborhoods than minority neighborhoods – this could become an important data point in the ongoing debate about the sharing economy occurring across the U.S. and around the world.

5.2 Additional Relevant Geographic Principles

In this article, we have discussed how four geographic principles play a key role in the sharing economy and result in structured geographic biases along SES lines (and those defined by race and ethnicity). However, these four principles are almost certainly not the only aspects of human geography that are important to consider when examining the sharing economy. For instance, two human geography principles worthy of exploration are *border effects* and *edge cities*.

Border effects are a well-known human geography principle that describe what occurs when two neighboring places that are on the opposite sides of an administrative boundary have tremendously different circumstances with respect to a variable of interest. Our results suggest that as some municipalities like Austin, Texas begin to place restrictions on sharing economy services (especially Uber and Lyft) [Burger 2015; Olsson 2016], border effects will become increasingly important in the sharing economy. Specifically, we (unsurprisingly) found that a census tract's wait times were highly dependent on neighboring tracts' wait times. This means that adjoined municipalities will likely suffer reduced sharing economy effectiveness when one municipality places restrictions on sharing economy services. For instance, our results suggest that Austin's suburbs are going to be severely affected by Austin's sharing economy-related decisions, even if they have no say in these decisions. Examining the effect of differing sharing economy regulations within the framework of border effects is an important direction of future work.

Our results also suggest that if and when the sharing economy becomes prominent in suburbs, "edge cities" [Garreau 1992] will lead the way and become secondary sharing economy hubs. An "edge city" is a concentration of work and leisure resources in a suburb that has good vehicle accessibility with respect to the rest of the metropolitan area, usually due to a nearby intersection of multiple freeways (e.g. a "ring road" and an intersecting highway) [Garreau 1992]. Common examples include Tyson's Corner, VA, Bloomington, MN, and the Rosslyn-Ballston Corridor in the Washington, D.C. metro area. Edge cities emerged in the second half of the twentieth century, becoming competitors with central business districts for shopping, employment, and entertainment services. The vehicle accessibility advantages that led to the agglomeration of traditional services in edge cities should also apply to sharing

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

economy services. Namely, relative to other suburbs, edge cities will be significantly less affected by the negatives associated with distant decay due to these accessibility advantages. Moreover, in the ride-hailing space, there is reason to believe that the limited availability of mass transit options in edge cities makes them even more suited to the sharing economy. Another more uncomfortable potential advantage for edge cities in the sharing economy is that they and their surrounding residential areas tend to be relatively high SES and populated by non-minority racial and ethnic groups. Overall, given the numerous properties of edge cities that interact with properties of the sharing economy, many branches of sharing economy research – e.g. work on improving sharing economy effectiveness in new regions, work seeking to understand its long-term impact on urban areas – should likely consider edge cities as in important near-term direction of inquiry.

5.3 Implications for "Geosociotechnical" Design

As noted above, in almost every case, *the structural geographic biases that are the result of the four geographic principles are not destiny for the sharing economy*: they are the outcome of interactions between these principles and the "geosociotechnical" design of existing sharing economy platforms. In other words, given the design choices in these platforms and the inherently geographic nature of the sharing economy, it is not a surprise that these principles manifest in the biases we observed. Indeed, awareness of these principles led us towards our SES-related hypotheses.

Fortunately, the important role of geosociotechnical design in the structural biases identified in this article means that *there is an opportunity to address these biases with design changes*. In the remainder of this sub-section, we outline a number of geosociotechnical improvements that could lead to the reduction of the SES, racial, and ethnic biases we identified in the sharing economy.

5.3.1. Using Design to Improve the Geographic Distribution of Microentrepreneurs. Our results suggest that the Big Sort residential geography of sharing economy workers – coupled with distance decay – is a primary causal factor behind the geographic variation in the effectiveness of sharing economy systems that we observed. For instance, in our TaskRabbit study, we found that workers were not willing to travel long distances for a task. This would not result in geographic disparities in the effectiveness of TaskRabbit if TaskRabbit workers were evenly distributed across the Chicago area. However, as noted above, TaskRabbit workers are heavily clustered in relatively wealthy and dense parts of the region (as per Big Sort processes), leading to higher prices and fewer jobs accepted in other types of areas. One might call these '(commercial) sharing economy deserts' as an analogy to 'food deserts' [Breneman and Ver Ploeg 2016], which tend to occur in similar types of areas. Along the same lines, with respect to UberX, Dillahunt and Malone [2015] found that few participants in a workshop on the sharing economy for job-seekers were even aware of sharing economy services.

These results suggest that recruiting new sharing economy workers in areas that suffer from the wrong end of the structural biases we identified would go a long way to eliminating these biases. For instance, it is likely that a relatively small number of TaskRabbit microentrepreneurs who live on the South Side of Chicago could have substantially diminished any price and willingness disadvantages in these areas. With respect to Dillahunt and Malone's findings, the same may be true for UberX wait times. Moreover, service quality improvements in disadvantaged areas could lead to a larger

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

(and more diverse) customer base, increasing the incentive to recruit more workers from disadvantaged areas.

The question then becomes: how can sharing economy platforms do such recruitment? Making different design choices can likely contribute to the answer. For instance, TaskRabbit requires that workers have a bank account to participate on their platform, with low-income and minority neighborhoods having a much greater percentage of the population that is "unbanked" [Sullivan 2013], inherently reducing the potential working population in these places. UberX additionally requires workers to own a car and have active insurance, among other requirements[8], which likely has a similar effect (Dillahunt and Malone [2015] found that two-thirds of people in a sharing economy workshop for disadvantaged communities did not have a car that met Lyft's requirements). Other worker restrictions may also play a role: both TaskRabbit and UberX require that workers pass a background check, and it is unclear if minor prior offenses would result in rejection (e.g. a minor drug possession arrest). Low-income neighborhoods have a higher rate of these minor offenses [Harris and Kearney 2014].

Some sharing economy platforms have begun to make design changes in this direction. For instance, recent efforts by Uber to provide banking services to its drivers [Leberstein 2016] and to make obtaining car leases easier for potential drivers [Business Wire 2016] could potentially address some of the problems that we identified in this article. Of course, these initiatives could also lead to new problems, including the serious risk of exacerbating debt-related challenges in disadvantaged areas, and these leases have been accused of being predatory [Newcomer and Zaleski 2016]. Our results suggest that research into these and other mechanisms for increasing worker participation rates in 'sharing economy deserts' should be a top priority for future work.

5.3.2. Addressing Workers' Mental Maps. Our TaskRabbit results showed that comfort levels in workers' mental maps played a key role in their willingness to accept tasks and specifically made them less likely to accept tasks in wide swaths of southern and western Chicago. A similar finding was identified by Lee et al. [2015], who found that UberX drivers turned off their availability when they were in neighborhoods they perceived as undesirable. In both cases, workers cited perceptions of crime as the reasons for their discomfort in certain areas. The mental maps literature, however, suggests that in many cases perception does not match reality. For instance, Matei et al. showed that comfort levels associated with regions in the Los Angeles metropolitan area were effectively uncorrelated with crime levels. Additionally, as noted above, our mental maps generally struggle to incorporate sufficient detail to be able distinguish pockets of low-crime areas in unfamiliar high-crime districts.

The gap between perception and reality in mental maps presents a potentially powerful opportunity for geosociotechnical design. One straightforward design improvement would be to provide workers with geographically-linked crime statistics in an easily-digestible format that would allow for design-making on-the-fly. In most areas, crime statistics are public information and could be surfaced via a map in a microentrepreneur app quite easily. A more interesting and likely more useful approach would be to provide this information in the context of a given task, e.g. when a TaskRabbit worker is deciding whether to accept a task or when an UberX driver is driving through a specific neighborhood. This information could take the form of basic

[8] https://www.uber.com/driver-jobs

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

crime statistics or, matching the norms of the sharing economy, could be reported as a "geographic reputation score". This score could take into account both public crime information as well as geographically-linked incident reports privately held by a sharing economy platform. Based on the work of Matei et al., it is likely that many areas that currently are associated with high discomfort would have high geographic reputation scores (and perhaps vice versa). If this information were made available to workers, it could address some of the TaskRabbit willingness and UberX wait time bias that we observed in this study[9].

5.3.3. A Role for "Sociotechnical Auditing". Our results add to evidence that auditing has an important role to play in protecting the sharing economy from bias, just as has been argued in more explicitly algorithmic domains (e.g. [Kay et al. 2015]) and for other sociotechnical platforms (e.g. [Sandvig et al. 2014; Annany et al. 2015]). Fortunately, the geographic techniques we developed and adopted here – especially our spatial Durbin modeling approach – can provide a useful lens in this process. The relative geographic effectiveness of sharing economy platforms in a given administrative district likely would be a valuable data point for the many sharing economy debates that are occurring around the world. Our experiments outlined above should be replicable in most (if not all) areas in which TaskRabbit and UberX are active, and our techniques should relatively easily generalize to similar platforms (e.g. Lyft).

To make repeating our work in other areas as straightforward as possible, we are releasing our UberX data collection and spatial Durbin statistical framework under an MIT license[10]. This package should allow someone with technical training to quickly repeat our UberX experiment in their area with relatively little effort. Moreover, it should be relatively straightforward to adapt our code to other outcome metrics besides wait time and to other similar sharing economy platforms.

Additionally, our approach here has been to examine system-wide effects, but the auditing of the sharing economy should also likely occur at the worker-specific level. By examining the geographic history of a given worker, it should be possible to determine if that person is exerting implicit or explicit bias in their pricing and willingness decisions. If the worker has a long enough history with a platform, techniques similar to those we described above can be employed (e.g. adapting wait times to job history-specific attributes). In more traditional workplaces, "substantive oversight of decision making" is one facet of minimizing workplace gender and racial bias [Bielby 2000], and it is intuitive that the sharing economy could learn from this body of literature. Correcting for implicit bias may sometimes be as simple as making decisions more legible, e.g. "98% of your completed jobs have been in areas that are at least 95% white (non-Latino)".

Lastly, while the term "algorithmic auditing" has been adopted for doing this type of auditing work in a technological context, this term is not ideal for the sharing economy. Algorithms play a role – especially in the case of UberX – but sharing economy platforms are *socio*technical, not just technical. As we have seen, it is human biases – in the form of the Big Sort, mental maps, etc. – that are the drivers of many of the structural geographic disadvantages that we observed in this study. As such, auditing in the sharing economy is "sociotechnical" auditing (and even perhaps

[9] It is important to note, however, that all attempts to create a geographic reputation score should be fully transparent, lest other forms of bias get implicitly or explicitly incorporated.

[10] https://github.com/jtsmn/uber_data_spatial_durbin_model

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

“geosociotechnical auditing”) and needs to adopt approaches from both the algorithmic auditing literature and the large literature on detecting bias in human decision-making, e.g. the implicit bias reduction technique mentioned immediately above.

5.3.4. Task-specific vs. Global Pricing. Our results above suggest that UberX’s decision to fix prices globally as a function of distance may have reduced pricing-related bias in its platform relative to TaskRabbit, which at the time of our analysis allowed for per-task pricing. Namely, whereas in TaskRabbit both price and willingness were entirely dependent on human decision-making processes that are subject to bias, in UberX, this is only true of willingness (manifest in wait times). TaskRabbit’s pricing model has changed since our study and now more closely aligns with that of UberX. Specifically, in most cases, workers now define hourly wages for categories of tasks, and are algorithmically presented to task requesters. As such, it is likely that the price-related biases we identified above are either reduced or manifest differently in the design of the TaskRabbit platform that is current as of this writing (like many sharing economy platforms, TaskRabbit is frequently changing its pricing structure).

However, basic economics tells us that price and willingness are not independent, and the relationship between the two was specifically addressed by a number of our TaskRabbit participants above. Specifically, when price controls are employed, shortages can emerge [Taylor 2006]. As such, if a sharing economy platform uses fixed pricing, and these prices are set too low for tasks in a specific area for whatever reason (e.g. distance, mental maps), willingness will likely drop in this area. This could, indeed, be a factor behind some of the relatively large wait time effect sizes observed in our UberX models. Better understanding the relationship between price and willingness in the context of the geosociotechnical design of pricing models is an important area of future work. Reimagining our work under a variety of different pricing models would be a good place to start.

5.4 Other Areas of Future Work

5.4.1 Gender and the Sharing Economy. One critical area of future work highlighted by this research is further examination of the relationship between gender and the sharing economy. We found in our TaskRabbit research that women were significantly less likely to be willing to do a task than men (willingness rates were about 20% lower). While we hypothesized that the effects associated with discomfort and mental maps may be exacerbated for women, resulting in the 20% difference, future studies will be necessary to (1) confirm this difference in other sharing economy contexts and (2) isolate its cause. The import of investigating these two points cannot be understated: if willingness is lower for women, it could have important effects on women’s ability to earn comparable amounts as men in the sharing economy: with less competition in areas perceived to be unsafe, men could charge higher rates. There have also recently been high-profile developments associated with the relationship between gender and the sharing economy that are worth studying and that may provide key sources of qualitative and quantitative data with regard to these issues, e.g. a ride-sharing service designed explicitly to serve the safety needs for women passengers by specifically hiring women drivers [Farivar 2016].

5.4.2. Temporal Bias in Access to Sharing Economy Services. In our UberX analyses, we observed that UberX was launched in a higher SES portion of the Chicago region before it became available to the metropolitan area more widely. This is a pattern we see at

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

a more global geographic scale as well, with many sharing economy platforms launching first in relatively high-SES, high-density metropolitan areas (e.g. San Francisco, New York) before their developers open them up to other metropolitan areas.

Relative to some of the other challenges associated with the sharing economy identified in this article, this 'temporal bias' is likely less significant, assuming wide launches eventually occur. However, one concern we have is that this temporal pattern may lead to higher SES individuals gaining first mover advantages both as consumers and microentrepreneurs, e.g. with regard to reputation socres. Better understanding the launch patterns of sharing economy platforms (and other geographic technologies, more generally) and their possible follow-on effects could be a valuable direction of future work.

5.4.3. Putting Sharing Economy Bias into Context. This article is interested in understanding the relative effectiveness of the sharing economy in different areas and the geographic mechanisms behind this variation, not in comparing the effectiveness of sharing economy platforms with their traditional economy equivalents. It may be, for instance, that UberX has significantly lower wait times than traditional cab companies in all areas of Chicago, regardless of the demographic makeup of a neighborhood. Similarly, TaskRabbit may open up new opportunities for acquiring low-cost paid help in small low-SES areas near high-SES regions.

Given that it is widely believed that sharing economy services will substantially displace their traditional economy equivalents in the near- and mid-term future [Milbourn 2015; Bernstein 2015], understanding the geographic variation in the effectiveness of these services is critical. However, many sharing economy-related debates have been framed as a comparison with traditional economy competitors. As such, it is important that future work provide much-needed robust data points with respect to this comparison. The methodological frameworks we developed above can be used for studies of this type. For instance, one could adopt our TaskRabbit experiment to collect data from taxi drivers. Similarly, our spatial Durbin modeling approach could be used with large-scale trip data collected by taxi companies and obtained by municipal governments [New York City Taxi & Limousine Commision 2016]. Indeed, we have completed early work comparing UberX to New York City's green and yellow cabs. This research suggests that UberX provides better service to areas with large minority populations compared yellow cabs. However, green cabs, which serve outer boroughs, significantly outperform UberX in this respect.

Of the comparisons between sharing and traditional economy services that have been made [Rivoli 2016; Avila 2016; Love 2016], one important factor has tended to be excluded: the informal economy services that often arise to address limitations in traditional economy services (e.g. [LeBlanc 1999; Resnick 2004; Suzuki 1985]). For instance, 'vernacular cabs' [Suzuki 1985] – ride-hailing services that are informally organized and have "fares based on negotiations or 'gentlemen's agreements'" – have existed in many low-income areas in the United States for years [Suzuki 1985]. In many ways, vernacular cabs (and related systems like sluglines [McDonald and Shubert 2016]) can be considered "peer-to-peer UberX" and their relationship to (digital) sharing economy services and traditional economy services should be a consideration in any comparative analysis of the sharing and traditional economies.

Vernacular cabs also present several intriguing possibilities for sharing economy researchers and practitioners. Can we develop technologies to support these networks in addition to (or instead of) attempting to adapt centrally-run commercial sharing

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

economy platforms to be more effective in low-SES areas? What would be the effect of having separate platforms for low-SES and high-SES areas? Would a peer-to-peer model work well in high-SES areas? Given the limited amount of information about vernacular cabs, likely the first step in this research direction is formative qualitative work on vernacular cab networks with an eye towards implications for design.

5.5. Limitations

The sharing economy is currently the subject of much political debate (e.g. [Foroohar 2016; O'Connor 2016; Craver 2016; Mihalopoulos 2016; Cornfield 2016]) and there are many sensitive issues associated with the socioeconomic and demographic factors examined in this work. As such, it is critical that readers consider the limitations of our research, especially when utilizing our research to inform their opinions about the politics surrounding the sharing economy. We outline the most important of these limitations below.

First and foremost, as is noted above, this work focuses exclusively on comparing the relative effectiveness of sharing economy platforms across various geographies; it does not examine the effectiveness of these platforms with respect to competing, non-sharing economy industries. This consideration is most important for our UberX results: although we showed that, for instance, within the city of Chicago (and possibly Cook County-wide) wait times are higher in poor neighborhoods, additional research needs to examine whether these wait times are higher or lower than that of taxis. However, regardless of the current geographic effectiveness of the sharing economy relative to traditional economy competitors, this paper has identified structural geographic biases in the sociotechnical design of multiple sharing economy platforms. This suggests that unless these biases are addressed directly, they will persist in a world in which the traditional economy competitors to sharing economy platforms no longer exist.

Secondly, in this paper, we have limited our focus to two specific commercial sharing economy systems. Researchers have also studied non-commercial sharing economy systems [Bellotti et al. 2014; Shih et al. 2015], and it is unclear whether the removal of explicit financial incentives would alter geographic dynamics in the sharing economy. Similarly, while this article furthers efforts in the HCI community to consider more than one platform when examining a new sociotechnical phenomenon (e.g. [Antin et al. 2011]), it was not possible to study all major commercial sharing economy platforms (let alone ones that may become prominent in the near future). Moreover, both TaskRabbit and UberX are frequently changing their geosociotechnical designs, creating variation even within these two platforms. Fortunately, the methodological frameworks outlined in this article provide guidance in engaging in further geographic sharing economy studies, whether these involve comparing commercial and non-commercial platforms, examining the effects of a design change, or investigating other related research questions. Researchers can also use our open-source geostatistical modeling infrastructure from our UberX analysis to examine a variety of sharing economy platforms in a variety of urban areas around the world (assuming spatially-referenced demographic information is available).

Third, although the sharing economy has spread to thousands of cities in dozens of countries, this study focused on one metropolitan area in a single country. Future work should seek to explore our research questions in other U.S. metropolitan areas and in metropolitan areas around world, especially those with significant different urban structures than exist in the United States [Brunn, Jack Francis Williams, et al. 2003].

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Fourth, now that we have established a foundation for the role of the four geographic principles considered here, researchers should likely consider these principles on an individual basis in more detail than we have done here. For instance, repeating some of the early work done on mental maps and the factors behind them (e.g. crime rates, perceptions of crime rates, media exposure) with sharing economy microentrepreneurs would shed a more detailed light on our findings.

Fifth, our TaskRabbit results relied on self-reported data rather than behavioral data. That is, we asked workers to indicate how much they *would* charge and whether they *would* do a task, rather than observing outcomes. Future work could seek to replicate our experiment with tasks that are fully completed. Similarly, our TaskRabbit study did not differentiate between 'preferred mode of transportation' and 'most common mode of transportation' (i.e. a person could 'prefer' a mode of transportation that they did not often utilize). Although this was unlikely to be an issue, future studies could improve upon ours by measuring the mode of transportation that participants actually use to complete tasks.

Finally, we believe that longitudinal analyses that expand the static snapshots in this paper are important directions of future work. The sharing economy is an incredibly fast-moving space: adoption rates are growing both on the consumer and the microentrepreneur side, policy is shifting, and (as noted above) geosociotechnical designs are constantly changing. It is unlikely that geography's role in sharing economy effectiveness will decline. However, the character of geography's role may change as the sharing economy develops.

6. CONCLUSION

In this article, we have demonstrated how four geographic principles – the "Big Sort", variation in population density, distance decay, and mental maps – result in structural geographic biases in the effectiveness of the sharing economy. These biases lead sharing economy services to be both more expensive and less available in low-SES areas and suburban areas than in high-SES and high-density urban areas. Moreover, SES and race/ethnicity are often strongly correlated in many parts of the world, and we observed that, at least in the city of Chicago itself, areas in which the population is more white (non-Latino) have better access to sharing economy services.

Overall, this article provides evidence that (1) in the sharing economy, geography matters, and geographic principles should be strongly considered in examinations of the sharing economy and (2) one way in which the importance of geography manifests is that key geographic principles interact with common design decisions in sharing economy platforms to create important biases in the effectiveness of the sharing economy. As discussed above, engaging with both of these takeaways can lead to 'geosociotechnical' design improvements in sharing economy platforms that reduce these biases, among other benefits.

7. ACKNOWLEDGEMENTS

The researchers would like to thank Shilad Sen, Aaron Rendahl, and Darren Gergle. This research was supported by NSF grants IIS-1707296, IIS-1218826, IIS-0808692 and a Microsoft FUSE Sharing Economy Research Award.

REFERENCES

Pragya Agarwal. 2001. Walter Christaller, Hierarchical Patterns of Urbanization. *eScholarship* (January 2001).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Airbnb. 2016. About Us - Airbnb. (June 2016). Retrieved June 28, 2016 from https://www.airbnb.com/about/about-us

Florian Alt, Alireza Sahami Shirazi, Albrecht Schmidt, Urs Kramer, and Zahid Nawaz. 2010. Location-based Crowdsourcing : Extending Crowdsourcing to the Real World. *October* (2010).

Mike Annany, Karrie Karahalios, Christian Sandvig, and Christo Wilson. 2015. Auditing Algorithms From the Outside: (2015). Retrieved November 16, 2016 from https://auditingalgorithms.wordpress.com/

L. Anselin. 1988. *Spatial Econometrics: Methods and Models*, Springer Science & Business Media.

Judd Antin, Ed H. Chi, James Howison, Sharoda Paul, Aaron Shaw, and Jude Yew. 2011. Apples to Oranges?: Comparing Across Studies of Open Collaboration/Peer Production. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. WikiSym '11. New York, NY, USA: ACM, 227–228. DOI:https://doi.org/10.1145/2038558.2038610

Joseph De Avila. 2016. Airbnb Will Collect Hotel Tax for Hosts in Connecticut. *Wall Str. J.* (June 2016).

Victoria M.E. Bellotti et al. 2014. Towards community-centered support for peer-to-peer service exchange: rethinking the timebanking metaphor. In ACM Press, 2975–2984. DOI:https://doi.org/10.1145/2556288.2557061

Alan Benson, Aaron J. Sojourner, and Akhmed Umyarov. 2015. Can Reputation Discipline the Gig Economy? Experimental Evidence from an Online Labor Market. (2015).

Michael Bernstein. 2015. The Future of Work: Working for the Machine. (August 2015). Retrieved August 10, 2015 from http://www.psmag.com/business-economics/the-future-of-work-working-for-the-machine

Michael S. Bernstein et al. 2010. Soylent: A Word Processor with a Crowd Inside. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology*. UIST '10. New York, NY, USA: ACM, 313–322. DOI:https://doi.org/10.1145/1866029.1866078

William T. Bielby. 2000. Minimizing Workplace Gender and Racial Bias. *Contemp. Sociol.* 29, 1 (2000), 120–129. DOI:https://doi.org/10.2307/2654937

Bill Bishop. 2008. *The Big Sort: Why the Clustering of Like-minded America is Tearing Us Apart*, Houghton Mifflin Harcourt.

Roger Bivand, Jan Hauke, and Tomasz Kossowski. 2013. Computing the Jacobian in Gaussian spatial autoregressive models: An illustrated comparison of available methods. *Geogr. Anal.* 45, 2 (2013), 150–179.

Roger Bivand and Gianfranco Piras. 2015. Comparing Implementations of Estimation Methods for Spatial Econometrics. *J. Stat. Softw.* 63, 18 (2015), 1–36.

Mark Bjelland, Daniel Montello, Jerome Fellmann, Arthur Getis, and Judith Getis. 2013. *Human Geography: Landscapes of Human Activities* 12 edition., New York: McGraw-Hill Education.

Charles M. Blow. 2015. How Expensive It Is to Be Poor. *N. Y. Times* (January 2015).

Vince Breneman and Michele Ver Ploeg. 2016. USDA Economic Research Service - Food Access Research Atlas. (June 2016). Retrieved June 30, 2016 from http://www.ers.usda.gov/data/fooddesert/

Stanley D. Brunn, Jack Francis Williams, and Donald J. Zeigler. 2003. *Cities of the World: World Regional Urban Development*, Rowman & Littlefield.

Stanley D. Brunn, Jack F. Williams, and Donald J. Zeigler. 2003. *Cities of the World: World Regional Urban Development* Third Edition., Rowman & Littlefield Publishers.

Ludwig Burger. 2015. Uber Germany retreats to Berlin, Munich. *Reuters* (October 2015).

Business Wire. 2016. Green Dot and Uber Announce "Uber Checking by GoBank®" | Business Wire. (March 2016). Retrieved June 20, 2016 from http://www.businesswire.com/news/home/20160317005466/en/Green-Dot-Uber-Announce-Uber-Checking-GoBank%C2%AE

John Byrne. 2016. Airbnb rules easily pass Chicago City Council despite vocal opposition. (June 2016). Retrieved June 30, 2016 from http://www.chicagotribune.com/news/local/politics/ct-chicago-city-council-airbnb-rules-met-20160622-story.html

Chicago Tribune. 2016. Chicago shooting victims. (June 2016). Retrieved June 15, 2016 from http://crime.chicagotribune.com/chicago/shootings

Rianne Coale. 2016. Traveling to O'Hare by train, cab and Uber: a comparison. (March 2016). Retrieved June 15, 2016 from http://www.redeyechicago.com/news/redeye-getting-to-o-hare-a-price-and-time-breakdown-20160309-story.html

R.A. Cochrane. 1975. A Possible Economic Basis for the Gravity Model. *J. Transp. Econ. Policy* 9, 1 (1975), 34–49.

Ashley Colley et al. 2017. The Geography of Pokémon GO: Beneficial and Problematic Effects on Places and Movement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Jsoh Cornfield. 2016. Eric Holder weighs in for Uber against fingerprint checks. (June 2016). Retrieved June 16, 2016 from http://bigstory.ap.org/article/01685274cc944196af9016ad0a875b72/eric-holder-weighs-uber-against-fingerprint-checks

Jack Craver. 2016. Council scrambling to find replacements for Uber and Lyft. (May 2016). Retrieved June 16, 2016 from http://www.austinmonitor.com/stories/2016/05/council-scrambling-find-replacements-uber-lyft/

Kimberle Crenshaw. 1989. Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics. *Univ. Chic. Leg. Forum* 1989 (1989), 139.

Adam Dennett. 2012. *Estimating flows between geographical locations:"get me started in"spatial interaction modelling*, Citeseer.

Nicholas Diakopoulos. 2015. How Uber surge pricing really works. *Wash. Post* (April 2015).

Tawanna R. Dillahunt, Airi Lampinen, Jacki O'Neill, Loren Terveen, and Cory Kendrick. 2016. Does the Sharing Economy do any Good? In *19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. ACM Press, 197–200. DOI:https://doi.org/10.1145/2818052.2893362

Tawanna R. Dillahunt and Amelia R. Malone. 2015. The Promise of the Sharing Economy among Disadvantaged Communities. In ACM Press, 2285–2294. DOI:https://doi.org/10.1145/2702123.2702189

Anne-Célia Disdier and Keith Head. 2008. The puzzling persistence of the distance effect on bilateral trade. *Rev. Econ. Stat.* 90, 1 (2008), 37–48.

Jessica D'Onofrio and Charles Thomas. 2016. Chicago approves new rules on Airbnb, ride-hailing services. (June 2016). Retrieved June 30, 2016 from http://abc7chicago.com/1396244/

Benjamin G. Edelman and Michael Luca. 2014. Digital Discrimination: The Case of Airbnb. com. *Harv. Bus. Sch. NOM Unit Work. Pap.* , 14-054 (2014).

Benjamin Edelman, Michael Luca, and Dan Svirsky. 2015. Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *Harv. Bus. Sch. NOM Unit Work. Pap.* , 16-069 (2015).

J. Paul Elhorst. 2010. Applied Spatial Econometrics: Raising the Bar. *Spat. Econ. Anal.* 5, 1 (March 2010), 9–28. DOI:https://doi.org/10.1080/17421770903541772

ESRI. 2016. Location-allocation analysis. (June 2016). Retrieved May 27, 2016 from http://desktop.arcgis.com/en/arcmap/latest/extensions/network-analyst/location-allocation.htm#ESRI_SECTION1_F8182D9F421E4EA4AEE11E7B360E1340

Cyrus Farivar. 2016. Chariot for Women, an Uber-like service, set to launch this month. (April 2016). Retrieved June 15, 2016 from http://arstechnica.com/business/2016/04/chariot-for-women-an-uber-like-service-for-females-set-to-launch-this-month/

Robin Flowerdew and Murray Aitkin. 1982. A Method of Fitting the Gravity Model Based on the Poisson Distribution*. *J. Reg. Sci.* 22, 2 (May 1982), 191–202. DOI:https://doi.org/10.1111/j.1467-9787.1982.tb00744.x

Rana Foroohar. 2016. How the Gig Economy Could Save Capitalism. *Time* (June 2016).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

William H. Frey. 2015. Census shows modest declines in black-white segregation. (December 2015). Retrieved June 15, 2016 from http://www.brookings.edu/blogs/the-avenue/posts/2015/12/08-census-black-white-segregation-frey

William H. Frey and Dowell Myers. 2005. Racial segregation in US metropolitan areas and cities, 1990–2000: Patterns, trends, and explanations. (2005).

Lee Gaines. 2016. Prompted by luxury treehouse, Schaumburg imposes tax on Airbnb rentals. (June 2016). Retrieved June 30, 2016 from http://www.chicagotribune.com/suburbs/schaumburg-hoffman-estates/news/ct-treehouse-airbnb-schaumburg-met-20160629-story.html

Ruth García-Gavilanes, Yelena Mejova, and Daniele Quercia. 2014. Twitter ain't without frontiers. In *the 17th ACM conference*. New York, New York, USA: ACM Press, 1511–1522. DOI:https://doi.org/10.1145/2531602.2531725

Joel Garreau. 1992. *Edge City: Life on the New Frontier* Reprint edition., New York: Anchor.

Yanbo Ge, Christopher R. Knittel, Don MacKenzie, and Stephen Zoepf. 2016. *Racial and Gender Discrimination in Transportation Network Companies*, National Bureau of Economic Research.

Mareike Glöss, Moira McGregor, and Barry Brown. 2016. Designing for Labour: Uber and the On-Demand Mobile Workforce. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1632–1643.

Peter Gould and Rodney White. 1986. *Mental Maps* 2 edition., London; New York: Routledge.

Mordechai Haklay. 2010. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environ. Plan. B Plan. Des.* 37, 4 (2010), 682–703. DOI:https://doi.org/10.1068/b35097

Benjamin Harris and Melissa Kearney. 2014. The Unequal Burden of Crime and Incarceration on America's Poor. (April 2014). Retrieved June 30, 2016 from http://www.brookings.edu/blogs/up-front/posts/2014/04/28-unequal-burden-crime-americas-poor-kearneym-harrisb

Brent Hecht et al. 2013. Geographic human-computer interaction. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*. New York, New York, USA: ACM Press, 3163. DOI:https://doi.org/10.1145/2468356.2479637

Brent Hecht and Darren Gergle. 2010. On the "localness" of user-generated content. In ACM Press, 229. DOI:https://doi.org/10.1145/1718918.1718962

Brent Hecht and Monica Stephens. 2014. A Tale of Cities: Urban Biases in Volunteered Geographic Information. (2014).

Jessi Hempel. 2016. Gig Economy Workers Need Benefits and Job Protections. Now. (January 2016). Retrieved January 29, 2016 from http://www.wired.com/2016/01/gig-economy-workers/

James Heyman and Dan Ariely. 2004. Effort for Payment A Tale of Two Markets. *Psychol. Sci.* 15, 11 (November 2004), 787–793. DOI:https://doi.org/10.1111/j.0956-7976.2004.00757.x

Jon Hilkevitch. 2014. Chicago halts Uber try at airport pickups. (May 2014). Retrieved July 24, 2015 from http://articles.chicagotribune.com/2014-05-06/news/chi-chicago-halts-uber-try-at-airport-pickups-20140506_1_ride-share-chicago-airports-uberx

Ryan Hughes and Don MacKenzie. 2016. Transportation network company wait times in Greater Seattle, and relationship to socioeconomic indicators. *J. Transp. Geogr.* 56 (October 2016), 36–44. DOI:https://doi.org/10.1016/j.jtrangeo.2016.08.014

Tapio Ikkala and Airi Lampinen. 2015. Monetizing Network Hospitality: Hospitality and Sociability in the Context of Airbnb. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW '15. New York, NY, USA: ACM, 1033–1044. DOI:https://doi.org/10.1145/2675133.2675274

Panagiotis G. Ipeirotis, Foster Provost, and Jing Wang. 2010. Quality management on Amazon Mechanical Turk. In *the ACM SIGKDD Workshop*. New York, New York, USA: ACM Press, 64. DOI:https://doi.org/10.1145/1837885.1837906

Isaac L. Johnson, Yilun Lin, et al. 2016. Not at Home on the Range: Peer Production and the Urban/Rural Divide. In ACM Press, 13–25. DOI:https://doi.org/10.1145/2858036.2858123

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Isaac L. Johnson, Subhasree Sengupta, Johannes Schöning, and Brent Hecht. 2016. The Geography and Importance of Localness in Geotagged Social Media. In ACM Press, 515–526. DOI:https://doi.org/10.1145/2858036.2858122

Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. New York, NY, USA: ACM, 3819–3828. DOI:https://doi.org/10.1145/2702123.2702520

Aniket Kittur, Boris Smus, Robert E. Kraut, and Robert Kraut. 2011. CrowdForge. In *the 2011 annual conference extended abstracts*. New York, New York, USA: ACM Press, 1801. DOI:https://doi.org/10.1145/1979742.1979902

Florence Ladd. 1967. A note on "the world across the street." *Harv. Grad. Sch. Educ. Assoc. Bull.* 12, 1967 (1967), 47–49.

Sarah Leberstein. 2016. Uber's car leasing program turns its drivers into modern-day sharecroppers. *Quartz* (June 2016).

David LeBlanc. 1999. About Slugging. (1999). Retrieved June 30, 2016 from http://www.slug-lines.com/Slugging/About_slugging.asp

Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers. In ACM Press, 1603–1612. DOI:https://doi.org/10.1145/2702123.2702548

James LeSage and R. Kelley Pace. 2009. *Introduction to Spatial Econometrics*, Chapman and Hall/CRC.

David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, Andrew Tomkins, and Ronald L. Graham. 2005. Geographic Routing in Social Networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 33 (2005), 11623–11628.

Linna Li, Michael F. Goodchild, and Bo Xu. 2013. Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartogr. Geogr. Inf. Sci.* 40, 2 (March 2013), 61–77. DOI:https://doi.org/10.1080/15230406.2013.777139

Linna Li, Michael Goodchild, and Bo Xu. 2013. Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartogr. Geogr. Inf. Sci.* 40, 2 (March 2013), 61–77. DOI:https://doi.org/10.1080/15230406.2013.777139

Tessa Love. 2016. Study shows Airbnb's growth overcomes hotels. (May 2016). Retrieved June 17, 2016 from http://www.bizjournals.com/sanfrancisco/blog/2016/05/airbnbs-growth-overcomes-hotels-rooms.html

Kevin Lynch. 1960. *The Image of the City*, MIT Press.

Momin M. Malik, Hemank Lamba, Constantine Nakos, and Jürgen Pfeffer. 2015. Population Bias in Geotagged Tweets. In *Ninth International AAAI Conference on Web and Social Media*.

Charles F. Manski. 1993. Identification of Endogenous Social Effects: The Reflection Problem. *Rev. Econ. Stud.* 60, 3 (July 1993), 531–542. DOI:https://doi.org/10.2307/2298123

Afra Mashhadi, Giovanni Quattrone, and Peter Mooney. 2012. On the Sustainability of Urban Crowd-Sourcing for Maintaining Large-Scale Geospatial Databases Categories and Subject Descriptors. (2012).

Sorin Matei, Sandra J. Ball-Rokeach, and Jack Linchuan Qiu. 2001. Fear and Misperception of Los Angeles Urban Space A Spatial-Statistical Study of Communication-Shaped Mental Maps. *Commun. Res.* 28, 4 (August 2001), 429–463. DOI:https://doi.org/10.1177/009365001028004004

Xiao Ma, Jeffery T. Hancock, Kenneth Lim Mingjie, and Mor Naaman. 2017. Self-Disclosure and Perceived Trustworthiness of Airbnb Host Profiles. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. CSCW '17. New York, NY, USA: ACM, 2397–2409. DOI:https://doi.org/10.1145/2998181.2998269

Shannon McDonald and Lucas Shubert. 2016. Transportation Conditions and Solutions in Carbondale, Illinois. *J. Appl. Sci. Arts* 1, 1 (March 2016).

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

Tom L. McKnight. 1997. *Regional geography of the United States and Canada*, Pearson US Imports & PHIPEs.

Miller McPherson, Lynn Smith-Lovin, and James M. Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Annu. Rev. Sociol.* 27, 1 (2001), 415–444. DOI:https://doi.org/10.1146/annurev.soc.27.1.415

Dan Mihalopoulos. 2016. Mihalopoulos: Uber claims its mission goes beyond money. *Chic. Sun-Times* (May 2016).

Tad Milbourn. 2015. In the Future, Employees Won't Exist. *TechCrunch* (June 2015).

Peter Mooney, Padraig Corcoran, and Adam C. Winstanley. 2010. Towards quality metrics for OpenStreetMap. *Proc. 18th SIGSPATIAL Int. Conf. Adv. Geogr. Inf. Syst. - GIS 10* (2010), 514. DOI:https://doi.org/10.1145/1869790.1869875

Stuart Moran et al. 2014. Listening to the Forest and Its Curators: Lessons Learnt from a Bioacoustic Smartphone Application Deployment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. New York, NY, USA: ACM, 2387–2396. DOI:https://doi.org/10.1145/2556288.2557022

Mohamed Musthag and Deepak Ganesan. 2013. Labor Dynamics in a Mobile Micro-task Market. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. New York, NY, USA: ACM, 641–650. DOI:https://doi.org/10.1145/2470654.2470745

Eric Newcomer and Olivia Zaleski. 2016. Inside Uber's Auto-Lease Machine, Where Almost Anyone Can Get a Car. (May 2016). Retrieved June 30, 2016 from http://www.bloomberg.com/news/articles/2016-05-31/inside-uber-s-auto-lease-machine-where-almost-anyone-can-get-a-car

New York City Taxi & Limousine Commision. 2016. New York City Taxi & Limousine Commision Trip Data. (2016). Retrieved June 16, 2016 from http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml

Geoff Nunberg. 2016. Goodbye Jobs, Hello "Gigs": How One Word Sums Up A New Economic Reality. (January 2016). Retrieved January 29, 2016 from http://www.npr.org/2016/01/11/460698077/goodbye-jobs-hello-gigs-nunbergs-word-of-the-year-sums-up-a-new-economic-reality

Patrick Gillespie and Sara Ashley O'Brien. 2015. The gig economy: More people might have jobs than you think. (October 2015). Retrieved January 29, 2016 from http://money.cnn.com/2015/10/02/news/economy/jobs-report-gig-on-demand-economy/index.html

Sarah O'Connor. 2016. The gig economy is neither "sharing" nor "collaborative." *Financ. Times* (June 2016).

Andrew Odlyzko. 2015. The Forgotten Discovery of Gravity Models and the Inefficiency of Early Railway Networks. *Œcon. Hist. Methodol. Philos.* , 5-2 (June 2015). DOI:https://doi.org/10.4000/oeconomia.1684

Karen Olsson. 2016. The Short-Lived Career of an Austin Uber Driver. (May 2016). Retrieved June 15, 2016 from http://www.newyorker.com/news/news-desk/the-short-lived-career-of-an-austin-uber-driver

Reid Priedhorsky, Mikhil Masli, and Loren Terveen. 2010. Eliciting and Focusing Geographic Volunteer Work. In *CSCW '10: 2010 ACM Conference on Computer Supported Cooperative Work*. Savannah, GA.

Giovanni Quattrone, Davide Proserpio, Daniele Quercia, Licia Capra, and Mirco Musolesi. 2016. Who Benefits from the "Sharing" Economy of Airbnb? *ArXiv160202238 Phys.* (February 2016).

G. Quattrone, Afra Mashhadi, and Licia Capra. 2014. Mind the Map: The Impact of Culture and Economic Affluence on Crowd- Mapping Behaviours. In *CSCW '14*.

Noopur Raval and Paul Dourish. 2016. Standing Out from the Crowd: Emotional Labor, Body Labor, and Temporal Labor in Ridesharing. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. CSCW '16. New York, NY, USA: ACM, 97–107. DOI:https://doi.org/10.1145/2818048.2820026

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

William John Reilly. 1931. *The law of retail gravitation*, WJ Reilly.

Paul Resnick. 2004. Dynamic ride sharing in the Bay Area. *Paul Resnicks Occas. Musings* (August 2004).

Dan Rivoli. 2016. Report shows more Uber cars involved in fatal crashes than cabs. (June 2016). Retrieved June 17, 2016 from http://www.nydailynews.com/new-york/uber-argues-nyc-daily-crash-data-inaccurate-article-1.2675556

Alex Rosenblat and Luke Stark. 2015. *Uber's Drivers: Information Asymmetries and Control in Dynamic Work*, Rochester, NY: Social Science Research Network.

Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data Discrim. Convert. Crit. Concerns Product. Inq.* (2014).

Salvatore Scellato, Anastasios Noulas, Renaud Lambiotte, and Cecilia Mascolo. 2011. Socio-Spatial Properties of Online Location-Based Social Networks. In *Fifth International AAAI Conference on Weblogs and Social Media*.

Juliet Schor. 2014. Debating the sharing economy. *Gt. Transit. Initiat.* (2014).

Juliet B. Schor and Connor J. Fitzmaurice. 2015. 26. Collaborating and connecting: the emergence of the sharing economy. *Handb. Res. Sustain. Consum.* (2015), 410.

Shilad W. Sen, Heather Ford, David R. Musicant, Mark Graham, Oliver S.B. Keyes, and Brent Hecht. 2015. Barriers to the Localness of Volunteered Geographic Information. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. New York, NY, USA: ACM, 197–206. DOI:https://doi.org/10.1145/2702123.2702170

S. Andrew Sheppard, Andrea Wiggins, and Loren Terveen. 2014. Capturing Quality: Retaining Provenance for Curated Volunteer Monitoring Data. In *CSCW '14*. CSCW '14. New York, NY, USA: ACM, 1234–1245. DOI:https://doi.org/10.1145/2531602.2531689

Patrick C. Shih, Victoria Bellotti, Kyungsik Han, and John M. Carroll. 2015. Unequal Time for Unequal Value: Implications of Differing Motivations for Participation in Timebanking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. New York, NY, USA: ACM, 1075–1084. DOI:https://doi.org/10.1145/2702123.2702560

Sara Silverstein. 2014. Uber Vs. Taxi Pricing By City. (October 2014). Retrieved July 29, 2015 from http://www.businessinsider.com/uber-vs-taxi-pricing-by-city-2014-10?op=1

Jennifer Stark and Nicholas Diakopoulos. 2016. Uber seems to offer better service in areas with more white people. That raises some tough questions. *Wash. Post* (March 2016).

Judith Steward. 2009. Economic Status. (August 2009). Retrieved June 30, 2016 from http://www.macses.ucsf.edu/research/socialenviron/economic.php

John Q. Stewart. 1948. Demographic Gravitation: Evidence and Applications. *Sociometry* 11, 1/2 (1948), 31–58. DOI:https://doi.org/10.2307/2785468

Bob Sullivan. 2013. For many households, banking is a costly luxury. (April 2013). Retrieved June 15, 2016 from http://www.cnbc.com/id/100661088

Brian L. Sullivan, Christopher L. Wood, Marshall J. Iliff, Rick E. Bonney, Daniel Fink, and Steve Kelling. 2009. eBird: A citizen-based bird observation network in the biological sciences. *Biol. Conserv.* 142, 10 (October 2009), 2282–2292. DOI:https://doi.org/10.1016/j.biocon.2009.05.006

Peter T. Suzuki. 1985. Vernacular cabs: Jitneys and gypsies in five cities. *Transp. Res. Part Gen.* 19, 4 (July 1985), 337–347. DOI:https://doi.org/10.1016/0191-2607(85)90069-X

Jacki O'Neill Nicola J. Bidwell, Himanshu Zade, Srihari H. Muralidhar, Anupama Dhareshwar, Baneen Karachiwala, Tandong Neba Cedrick Syed Ishtiaque Ahmed. 2016. Peer-to-peer in the workplace: A view from the road. In ACM – Association for Computing Machinery.

Yuri Takhteyev, Anatoliy Gruzd, and Barry Wellman. 2012. Geography of Twitter networks. *Soc. Netw.* 34, 1 (January 2012), 73–81. DOI:https://doi.org/10.1016/j.socnet.2011.05.006

Accepted Feb. 2017 – This is a pre-print – Final version available shortly

TaskRabbit. 2016. Changing the Face of Diversity in Tech: Our CBC Tech 2020 African American Inclusion Plan. *Taskrabbit Blog* (April 2016).

TaskRabbit. 2016. TaskRabbit. (2016). https://www.taskrabbit.com

TaskRabbit Support. 2016. Which Category Should I Choose? (June 2016). Retrieved June 28, 2016 from http://support.taskrabbit.com/hc/en-us/articles/205309290-Which-Category-Should-I-Choose-

B. Taylor. 2006. Price Ceilings. (2006). Retrieved June 30, 2016 from http://economics.fundamentalfinance.com/price-ceiling.php

Rannie Teodoro, Pinar Ozturk, Mor Naaman, Winter Mason, and Janne Lindqvist. 2014. The motivations and experiences of the on-demand mobile workforce. In ACM Press, 236–247. DOI:https://doi.org/10.1145/2531602.2531680

Fitz Tepper. 2016. Uber has completed 2 billion rides. *TechCrunch* (July 2016).

Jacob Thebault-Spieker, Loren G. Terveen, and Brent Hecht. 2015. Avoiding the South Side and the Suburbs: The Geography of Mobile Crowdsourcing Markets. In *CSCW 2015*. ACM Press, 265–275. DOI:https://doi.org/10.1145/2675133.2675278

United States Congress. 1964. *Civil Rights Act of 1964*,

US Census Bureau. 1999. Explanation of Race and Hispanic Origin Categories. (September 1999). Retrieved June 30, 2016 from http://www.census.gov/population/estimates/rho.txt

US Census Bureau Geography. 2010. 2010 Geographic Terms and Concepts - Census Tract. (2010). Retrieved June 28, 2016 from https://www.census.gov/geo/reference/gtc/gtc_ct.html

Vijaysree Venkatraman. 2010. Time to Hire a Housekeeper? *Science* (June 2010). DOI:https://doi.org/10.1126/science.caredit.a1000056

James O. Wheeler and Ronald L. Mitchelson. 1989. Information Flows among Major Metropolitan Areas in the United States. *Ann. Assoc. Am. Geogr.* 79, 4 (December 1989), 523–543. DOI:https://doi.org/10.1111/j.1467-8306.1989.tb00275.x

Tse-Chuan Yang, Aggie J. Noah, and Carla Shoff. 2013. Exploring Geographic Variation in US Mortality Rates Using a Spatial Durbin Approach: Mortality and Spatial Durbin. *Popul. Space Place* 21, 1 (October 2013), 18–37. DOI:https://doi.org/10.1002/psp.1809

Ken Yeung. 2016. TaskRabbit's app update focuses on getting tasks done in under 90 minutes. (March 2016). Retrieved June 28, 2016 from http://venturebeat.com/2016/03/01/taskrabbits-app-update-focuses-on-getting-tasks-done-in-under-90-minutes/

Ke Zhang and Konstantinos Pelechrinis. 2014. Understanding Spatial Homophily: The Case of Peer Influence and Social Selection. In *Proceedings of the 23rd International Conference on World Wide Web*. WWW '14. New York, NY, USA: ACM, 271–282. DOI:https://doi.org/10.1145/2566486.2567990

Dennis Zielstra and Alexander Zipf. 2010. A comparative study of proprietary geodata and volunteered geographic information for Germany. *Conf. Geogr. Inf.* 1 (2010), 1–15.

Accepted Feb. 2017 – This is a pre-print – Final version available shortly