Objectives

In the past decade, people have increasingly been performing short-term jobs, or gigs, on temporary-work platforms like Uber, Upwork, and TaskRabbit. This trend has given rise to discussion, in academic and policy circles, about a “gig economy,” and whether these new forms of work are good or bad, on balance, for society. Yet this debate mostly takes place in an empirical void: how the typical worker fares in the gig economy is not well understood. Do gig workers manage to build good careers on these platforms, or does this sort of work just take job security away from others without providing it to them? Can someone feasibly do many different types of work, for many employers, for many years? These questions are unanswered in part because traditional survey instruments are not helpful for gathering representative information on a population, like gig workers, which is hard to define before surveying begins. We will use funding from the SSHRC Insight Development Grant to gather representative information using a new survey method, to better understand how gig workers distribute their efforts across different platforms and types of work. Our objective with this research is to speak to a larger debate in work and employment relations about the possible trade-offs between employment flexibility and employment security.

Context and Literature Review

My co-applicant’s and my interest in this phenomenon grows out of our prior work. My dissertation work focused on the effects of performance on organizational culture. Rather than focusing on survey or other data that was gathered by companies, I harvested self-reported data from several hundred thousand workers who were employed in a diverse set of organizations and industries. The type of questions you can ask with such data, such as whether a certain form of organizational culture causes certain types of workers to stay or quit, is very different from the type of questions you can ask with company data. My co-applicant has published several studies on working conditions and the turnover of jobs, and has also therefore noticed the limitations of relying on organizational datasets to track people across jobs and organizations. In each case, we have seen how a data-gathering approach that worked in one labor-market regime has become less useful as employment has become less tied to one firm.

Labor markets have changed profoundly in recent decades, shifting away from steady, permanent employment and towards non-standard work arrangements (Cappelli and Keller, 2013). Katz and Krueger (2016) estimate that 94% of the net employment growth in the United States from 2005 to 2015 is attributable to non-standard work. A growing percentage of non-standard work involves *gig work*, which can be defined as short-term tasks solicited on demand through technologically mediated platforms. Common examples include automobile rides requested via Uber’s platform and freelance work completed on the Upwork website, although today there are myriad platforms focusing on different types of tasks.

We do not know whether or how gig work and the gig economy have restructured the jobs performed through such platforms, or the work experience of those who work. Some accounts of the gig economy posit that many workers have abandoned traditional jobs and careers, instead constructing work as a series of gigs on various platforms. These arguments presume that workers can earn sufficient income and benefits from these gigs to sustain themselves, while also enjoying non-wage benefits like flexibility and autonomy. Others posit that gig work has had little impact on traditional employment, with most gig workers just using gigs to supplement standard employment. This perspective typically treats gig work as a form of contingent or precarious labor that is characterized by insecurity, hidden costs, and few or no health or retirement benefits (Spreitzer et al., 2017). For example, some Uber drivers incur significant costs around automobile maintenance and depreciation, which can significantly reduce their take-home pay.

Each of these views has important implications, both for researchers’ theories of labor and employment and policymakers’ approach to worker welfare and security. Scholars continue to debate whether the proliferation of nonstandard work arrangements has reduced worker welfare. Some suggest that nonstandard work is highly polarized, between high-skilled jobs that offer good pay and non-monetary benefits and low-skilled jobs that exploit workers.

These debates go on in part because gathering reliable data on gig workers has proven difficult. Existing empirical work on the gig economy has two severe limitations. First, it typically treats the *company* as the unit of analysis (Chen et al., 2017). Thus a study might analyze Uber's data on its own drivers, to answer questions about how those drivers interact with that platform (ex. Cook et al., 2018). These studies give us insight into worker earnings and satisfaction, but they cannot answer questions about whether and how workers interact *across* multiple platforms. As useful as such studies are, they can only answer questions about what workers do when and while they work for one company. This leaves basic questions, like whether and how much workers spread their time across different gig-work platforms, unanswerable. Second, the data sources that treat the *worker* as the unit of analysis, such as the U.S. Bureau of Labor Statistics' Contingent Worker Supplement survey, focus their data collection on respondents' primary job, potentially missing supplementary ones. Such data is therefore biased towards workers who do gig work as their main form of employment. Furthermore, these surveys rarely delineate gig work accurately from other types of non-standard work, like part-time or contingent wage work for a single employer. (The Bureau of Labor Statistics failed to collect reliable data on electronically mediated employment in its 2017 Contingent Worker Supplement survey (Bureau of Labor Statistics, 2018).)

To overcome these empirical limitations, we plan to pilot a new survey of gig workers in Montreal and Toronto using the novel method of respondent-driven sampling (RDS). The RDS method resembles snowball sampling, where the researchers do not know the population from which to define their sampling frame in advance, and instead start with some "seed" respondents and gather recommendations for others. Unlike snowball sampling, though, RDS enlists the respondents themselves to recruit new respondents. Because the social structure of the referral process is then known, one can weight the data to produce representative results for the population.

Research Questions

We will use funding from the SSHRC Insight Development Grant to 1) design and pilot a respondent-driven sampling (RDS) survey of gig economy workers in two Canadian cities, Montreal and Toronto; 2) analyze the data to answer to answer several important research questions that are of interest to academics, policymakers, and the general public; and 3) write and disseminate the results to academics and the broader public.

The RDS method will allow us to answer several fundamental research questions, including whether and how much gig workers move across different employment platforms; whether and how they move between gig work and more formal employment; what types of different tasks they perform on various platforms; and what level of interaction they have with other gig-economy workers.

First, what kinds of workers participate in the gig economy? We will examine worker heterogeneity in demographics, education, earnings, and participation within and across gig-work platforms. Do the heaviest participants differ from more casual ones, for example?

Second, how does gig work intersect with other types of employment? It is unclear whether gig work substitutes for or complements other work arrangements. Do many workers structure their primary work as a series of gigs? Do most gig workers use gig work to replace income when unemployed? Is gig work more common when workers need supplemental money, as during the holidays or after incurring unforeseen expenses?

Third, what is the skill content of the tasks that gig workers perform? Gig work decouples specific tasks from general jobs, allowing gig workers to perform a more disparate set of tasks. One can do people's laundry through TaskRabbit, for example, without working in or for a laundry. The discreteness of such tasks removes the social stigma that might be associated with older jobs, and so gig workers might be more willing to do work that was previously performed by jobs they would be unwilling to take. How much do workers mix high- and low-skilled tasks? Does this vary by worker characteristics? And do workers use one platform to do multiple tasks, or spread their efforts across platforms?

Fourth, how do workers' social networks affect their gig work? Social-network contacts help people secure jobs and succeed at work. Our survey method requires us to collect data on respondents' networks, which can also help us understand how workers' network contacts influence their awareness and selection of gig opportunities, and the strategies they use succeed on the job. Do people with more contacts also secure more gigs, for example? Is a wider social network associated with more varied gig work, or with more gigs on a given platform?

Theoretical Framework

In traditional surveys, one develops a representative sample by first defining the population and then by randomly selecting people from the population. Sampling frames can be more or less complicated, but they almost always require first defining the population. Gig workers resemble a hidden population, though: one that researchers cannot define before they start gathering responses.

A classic way to study such populations is snowball sampling, where the researcher identifies a few initial respondents, asks those respondents for further recommendations, follows up with those recommendations, and repeats the process. The problem with snowball sampling is that we have no way to relate the various respondents to the larger population, and thus no way to judge whether and how the sample is biased.

Respondent-driven sampling was developed to deal with this problem (Heckathorn, 1997). The key difference between snowball sampling and RDS is that, in RDS, the respondents themselves recruit future respondents. The researcher can then map out the social structure of the referral network. Under a reasonable set of assumptions, this network of referrals can be treated as a first-order Markov process. We can skip the details of that process here and point out that this lets us weight the responses by the structure of the social network in order to recover unbiased estimates of the larger population.

The RDS method relies on a key assumption: though the population is hard for outsiders to identify, members of the population are known to one another. This assumption has been used to motivate studies of other populations, like intravenous drug users and undocumented workers, where it holds. On the other hand, an RDS survey of sex workers was aborted when it became clear that most such workers were kept socially isolated from one another. Based on our own conversations with gig workers, we think that knowledge of others is the norm, and thus RDS is a reasonable approach to take.

The RDS method requires collecting three pieces of data. The first is the respondent's personal network size, or the number of contacts they have who are also gig workers. The second is the respondent's serial number, and the third is the serial number of the respondent's referrals. These are generated as follows. First, the researchers identify a set of initial respondents or "seeds" and incentivize them to take the survey, usually with a cash payment. These seeds are then given a certain number of coupons that they can give to other potential respondents they know. The seed gets an additional payment whenever one of the people they referred completes the survey. Those referrals in turn become seeds themselves, and the process repeats. The serial numbers on the coupons are generated in such a way that the referrer for each one can be recovered; this is what lets the researchers map out the referral network. This process also means that, while RDS is not snowball sampling, the process does snowball: the number of respondents quickly grows in later survey waves. It is important to note that such coupons are almost always physical coupons, usually collected from the respondent when they complete the survey. This is to help ensure that referrals actually happen through social interaction between peers, rather than through indirect means like posting the coupon on the internet.

RDS relies on a validated mathematical model of the recruiting process to produce sample weights to adjust for the non-random peer recruitment (RDS web). One result is that it is possible to calculate both population estimates and estimate the uncertainty associated with those estimates. When used correctly, RDS produces asymptotically unbiased estimates, i.e. any bias disappears with a sufficiently large sample size (Salganick and Heckathorn, 2004). We will use the free Respondent Driven Sampling Analysis Tool (RDSAT) to analyze the data.

Crucially, because the unit of observation is a worker, this type of survey can ask worker-centric questions, like whether and how respondents spread their effort across platforms. And because an academic survey is easier to develop and pilot than a government survey, this approach can experiment more quickly and flexibly to find questions that solicit the right type of information about gig work.

This indeed was the initial motivation for using RDS survey methods to study groups like drug users and undocumented workers: to gather information that could influence policy but was not available through other means.

Methodology

We first need to develop a survey instrument. This includes drafting the questions for the survey, designing referral coupons, putting a respondent-tracking database into place, and conducting sample interviews.

We will focus on Uber drivers for our initial seeds. We think that Uber drivers are a good pilot population, for several reasons. First, because driving for Uber is a gig job that requires direct contact with the public, drivers are easier to approach and recruit. Second, while we do not know the exact number of Uber drivers in the Montreal area, we know from using the app that it is substantial. While Uber drivers may be our initial seeds, they can recruit workers who use other platforms as well.

We want to develop interview procedures that ensure confidentiality and data security. Our plan is to ask respondents to show us a referral coupon and some form of identification during the interview (again, to prevent multiple responses), but we want only to record responses, serial numbers, and contact information; and furthermore to record contact information separately from the other data.

In this study, we plan to conduct interviews using video-conference software. Gig work is mediated through internet platforms, so we know the target population has a smartphone or computer and a basic level of digital literacy. Doing video interviews lets us verify subject identities and serial numbers without having to collect and store physical copies of sensitive information. It also greatly simplifies the logistics and reduces the cost of meeting with respondents to interview them. This is particularly important for a second wave of data collection in Toronto: once initial seeds have been recruited, the interviewers themselves can work from Montreal, where the project is based. It is useful in this regard to note the existence of one-use video-conferencing sites like appear.in, which allow people to create video-conference sessions in a browser and invite people without having to install additional software, create accounts, or share personal contact information.

The basic logistics for a respondent would be as follows. First, the respondent would be given a recruitment coupon, either by the study team or by a previous respondent. That card will be made up as a postage-paid postcard, one that has a unique serial number, a brief description of the study, eligibility requirements, and a contact email address for the researchers. The respondent can contact the researchers and set up an interview. During the interview, the interviewer will note the serial number and retire it from the list of open requests, gather contact information for the respondent, and conduct the survey. The respondent can then drop the postcard in the mail to the study team. Once they receive the card, the team will mail the respondent a cheque with compensation for their participation and three more recruitment coupons for them to distribute to other potential respondents. When *those* recruits take the survey, the study team will also mail the relevant recruitment compensation to the referrer. The team will monitor the number of outstanding requests and average response rate to ensure that subject recruitment is halted within budget.

We plan to conduct structured interviews. The resulting responses will be entered into a database which we can then analyze. Quantitative data analysis is the most straightforward part of this project, and one where we have considerable experience. Virtually all analysis can be done with existing software packages, like Stata or R. RDS survey responses have to be weighted depending on the respondents included and the questions asked. Douglas Heckathorn, who developed the RDS method, also released a free Respondent Driven Sampling Analysis Tool to do this weighting.

Anticipated Outcomes

This funding will let us draft the RDS survey instrument, develop relationships with partner organizations, administer the survey, pay respondents, analyze and interpret the resulting data, and co-author the academic and popular articles that will result from the project.

This pilot study will show us the feasibility of scaling up the survey to other cities in Canada and abroad. If the pilot study is successful, we plan to apply for a SSHRC Insight Grant to launch a full survey in multiple cities in Canada and the US, such as Vancouver, Boston, and San Francisco. We think that the biggest obstacle to feasibility here is actually the size and type of social networks that gig-economy workers have with one another. If most gig workers interact with employment platforms but not with one another, then the RDS assumption will be violated and we will be unable to obtain unbiased population estimates. Fortunately, we can gather sequential information about workers' social networks in the course of conducting the survey itself; we do not have to run the entire study to learn whether the assumption holds. What is more, the characteristics of gig workers' contact networks are *themselves* an interesting research subject, so we think that this pilot study could yield findings that would result in at least one publishable article even we conclude that RDS methods are not suitable for a larger study. Thus we think that the risk and reward involved in this project are balanced nicely.

Originality, Significance, and Expected Contribution to Knowledge

To our knowledge, there is no prior survey that collects detailed data from a representative sample of gig-economy workers. The results have implications for sociologists interested in understanding the gig economy amidst the rise of non-standard work arrangements, economists interested in understanding why workers may select into gig work, and management scholars grappling with the increased market power of gig platforms and the consequences for worker welfare. This kind of survey could help us address several open questions in work and employment research, such as whether gig work lowers the barriers to entry faced by women and visible minorities in many occupations, and whether gig work can be leveraged into more secure jobs.

We also anticipate making a substantial methodological contribution, by learning how RDS can be used to study workers in the labor market. Successful implementation would allow the method to be reproduced and scaled in other labor markets in Canada and abroad to study gig work, as well other forms of non-standard employment that are also hard to capture with existing survey methods.

Our results would have significant implications for policymakers who can regulate the gig economy and assess whether gig work is on the whole beneficial for society. The wider public has also displayed increased anxiety about new employment arrangements' replacing standard jobs for some workers. This anxiety hinges on whether gig work is worth the initial investment, and whether gig workers have outside employment options.

Impact beyond the Social Sciences and Humanities

We anticipate that our results will provide useful information for people considering gig work. Are other people managing to build careers out of gig work, or are such jobs better viewed as temporary work during transitions between standard employment? People considering gig work face considerable uncertainty in assessing whether the benefits of these jobs outweigh the time and resources required. Such information will also help organizations that provide assistance to gig workers. Should these organizations encourage workers to transition to other forms of employment? Or should they help workers construct viable careers within the gig economy?

We have kept this potential impact in mind when drafting our knowledge-mobilization plan. We think that, in addition to academics, there are three groups that would be interested in this project's results. The first are current and gig workers, as mentioned above, and we want to be sure to produce materials that are readable by and useful to them. The second group is policymakers, for whom concrete information about life in the gig economy can help shape future regulation. The third group is the general public. Like many people inside and outside academia, we have frequent conversations with people about what the gig economy is, where it is going, and what it implies for our own work in the future.