Understanding the Skill Provision in Gig Economy from a Network Perspective: A Case Study of Fiverr

KEMAN HUANG, Sloan School of Management, Massachusetts Institute of Technology, USA
JINHUI YAO, School of Electrical and Information Engineering, University of Sydney, Australia
MING YIN, Department of Computer Science, Purdue University, USA

The recent emergence of gig economy facilitates the exchange of skilled labor by allowing workers to showcase and sell their skills to a global market. Despite the recent effort on thoroughly examining who workers in gig economy are and what their experience in gig economy are like, our knowledge on how exactly workers provide their skills in gig economy, and how worker's strategies of providing and expanding skills relate to their success in gig economy is still lacking. In this paper, we conduct a case study on a prominent gig economy platform, Fiverr.com, to better understand the provision of skills on it through large-scale, data-driven analysis. In particular, we propose the concept of "*skill space*" from a network perspective to characterize the relationship between different skills by measuring how frequently workers provide different skills together. Through our analysis, we reveal interesting patterns in worker's provision of skills on Fiverr. We then show how these patterns change over time and differ across subgroups of workers with different characteristics. In addition, we find that providing a set of skills that are highly-related with each other correlates with a better overall performance in gig economy, and when workers expand their skillsets, expanding to a new skill that is highly-related to the existing skills takes less time and is associated with better performance on the new skill. We conclude by discussing the implications of our findings for gig economy workers and platform in general.

CCS Concepts: • **Human-centered computing → Human computer interaction (HCI)**; **Collaborative and social computing**;

Additional Key Words and Phrases: Gig economy; Skill Space; Skill Provision; Skill Expansion; Skill Relatedness

ACM Reference Format:
Keman Huang, Jinhui Yao, and Ming Yin. 2019. Understanding the Skill Provision in Gig Economy from a Network Perspective: A Case Study of Fiverr. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 132 (November 2019), 23 pages. https://doi.org/10.1145/3359234

1 INTRODUCTION

The past decade has seen a rapid growth of the *gig economy*. During the early days of gig economy, "microtask"—short-term, small-sized work that has low requirements on specialized skills—is the primary form of work exchanged on gig economy platforms like Amazon's Mechanical Turk. More

Authors' addresses: Keman Huang, Sloan School of Management, Massachusetts Institute of Technology, 245 First Street, E94-15th Floor, Cambridge, MA, 02142, USA, keman@mit.edu; Jinhui Yao, School of Electrical and Information Engineering, University of Sydney, Sydney, Australia, jinhui.yao@gmail.com; Ming Yin, Department of Computer Science, Purdue University, West Lafayette, MA, 47907, USA, mingyin@purdue.edu.

2573-0142/2019/11-ART132 $15.00
https://doi.org/10.1145/3359234

recently, online labor markets and app-based platforms such as Upwork[1], Fiverr[2] and TaskRabbit[3] have emerged to directly connect employers and a global pool of workers for completing significantly more skilled work such as graphic design and software development. In addition to facilitating the matching of skills between supply and demand across global economies, the gig economy has also enabled an unconventional, temporary work arrangement that is embraced by an increasingly large portion of the workforce due to its relative flexibility, low barrier to entry, and potential to supplement primary income [18].

The rapid growth of gig economy workforce has sparked much research in better understanding workers in the gig economy. For example, a substantial amount of work has been done in analyzing the size and demographical decomposition of workers [11, 22, 23, 25]. Empirical studies are carefully designed to estimate worker's earning in the gig economy [6, 7]. More recently, much attention has been given to examining how the experience of workers in gig economy varies with their gender, race, and social-economic status [7, 13, 19], and perhaps surprisingly, it has been found that biases that are prevalent in traditional labor markets still present in new work environments that the gig economy offers [13].

Another angle that is critical for further advancing our knowledge of gig economy workers—especially those workers who complete skilled work in the gig economy—yet substantially under-explored is how workers provide skills in the gig economy and how the ways that workers provide skills affect their success. Indeed, compared to workers in traditional employment who repeatedly exercise the same set of skills that is often predefined in the job requirement and they are a master of, a common perception of workers in gig economy is that they can work on gigs that fit their skillsets, enjoy more freedom in expanding their skills based on interests and demand, and even practise the newly learned skills through the gig work. Empirical, data-driven studies on how exactly workers provide skills in the gig economy and whether worker's strategies on skill provision associate with how successful they are as gig workers, however, are largely lacking, potentially due to the limited amount of open data available. This work seeks to fill this gap to obtain a better understanding on the provision of skills in the gig economy. In particular, we aim to answer the following research questions:

- **RQ1**: From a macro point of view, what are the empirical patterns in the provision of skills in gig economy, e.g., which skills are frequently provided together by workers and which ones are rarely provided together?
- **RQ2**: How do these empirical patterns evolve over time?
- **RQ3**: Do workers with different characteristics (e.g., from different countries) show different patterns in providing skills?
- **RQ4**: From a micro point of view, how does the way that an individual worker provides skills relate to her overall performance, and how does the way an individual worker expands skills relate to her performance on the new skill?

These research questions are important for several reasons. First, the different nature of gig economy compared to the traditional employment implies the possibility of fundamentally different ways for workers to supply skills in their gig work. Workers with different characteristics may also exercise distinct, ever-changing strategies when participating in the gig economy. A close examination of the ways that workers provide skills in the gig economy could thus help us gain insights into how this new form of work is interpreted and utilized by the public over time. Second, as

[1]www.upwork.com
[2]www.fiverr.com
[3]www.taskrabbit.com

work in the gig economy is often atomic and can serve as building block for high-level complex jobs, understanding the empirical patterns in the provision of skills may inform us on the relationship between different skills and help to quantify the degree of relatedness among skills. Such knowledge can be very useful from the marketplace point of view for balancing the supply and demand of skills, such as recommending workers who have "related" skills to complete work which require skills that are in shortage. Third, understanding the association between how workers provide skills and their success in the gig economy can further lead to practical implications for effective skill recommendations that guide workers to efficiently advance their "careers" in the gig economy.

To answer these questions, we conduct a case study with Fiverr—one of the prominent online labor markets—using a dataset consisting of over 540,000 gigs and 250,000 workers. In particular, in this study, we approach our research questions from a network perspective. We propose the concept of "*skill space*" to characterize a *network of skills* established by their "relatedness," which is measured through the probabilities for different skills to be provided together by the same worker.

Constructing the skill space based on the dataset reveals a number of interesting patterns on how Fiverr workers provide skills in the gig economy and how these patterns change over time. For example, workers tend to provide similar skills within the same domain or complementary skills in different domains together, and more workers choose to provide skills which are closely-related to other skills or skills which belong to a tight-knit group of related skills. Workers have gradually adjusted the type of skills they provide over the past few years, while they also constantly expand the number of skills that they supply, effectively leading to the fact that more pairs of skills are frequently provided together over time. We further find that workers with different characteristics indeed show distinct patterns in providing skills. Based on the skill space that we construct, we quantify each worker's skill provision and expansion strategy by considering how related the skills a worker provides are and how much the new skill a worker is developing relates to the worker's current skillset. We find that workers who provide skills that have high degree of relatedness are associated with better overall performance in the gig economy. In addition, workers who expand into a new skill that is highly-related to their existing skills complete such expansion faster and also perform better on the new skill. These findings have practical implications for workers and market/platform providers of the gig economy, and we discuss these in the conclusion.

2 RELATED WORK

The gig economy today has facilitated the global exchange of both low-skill, low-pay, and short-term tasks (i.e., microtasks) and work that requires specialized skills, pays well, and costs much longer time. Microtasks are often exchanged on online crowdsourcing platforms like Amazon's Mechanical Turk and FigureEight [17, 28]; on these platforms, it is typically the demand-side of the labor (i.e., employer, or "requesters") who posts tasks on the platform and then workers can browse all tasks and decide which ones to work on. On the other hand, highly-skilled work is often exchanged on online labor markets that are driven by the supply-side—it is workers who advertise what they can do on the markets and employers will then decide which workers to hire. In this paper, we use "*gigs*" to refer to the skilled work advertised by workers themselves on online labor markets, to differentiate it from "tasks" on crowdsourcing platforms that are posted by employers. Intuitively, by deciding which gigs to offer on a market, workers effectively decide what skills to supply in the gig economy.

Early research on *skills* in the gig economy typically concerns skills exchanged through microtasks on crowdsourcing platforms, and they often come from a demand-side point of view by asking how to help employers find workers who have sufficient skills to satisfy the needs of their tasks. In this context, early studies often consider only a *single* skill and solve problems like how a worker's capability on this skill can be accurately estimated [20], and how can the employers identify the most qualified worker for a skill through social networks [5]. More recently, an increasing number of studies have been conducted to examine scenarios where each worker possesses *multiple* skills [2, 9, 15, 27, 31, 32]. Different algorithms have been proposed to match a task to the "right" worker who is most capable on the skill that the task requires, while the estimation of worker's capacity on each skill keeps updated in an online fashion. Notably, early work typically assumes that skills are independent from each other, with only a few exceptions trying to characterize relationships among skills. For example, a taxonomy or ontology was used to describe the *hierarchy* of different skills [21, 24], which enables employers to optimize task assignment when exact matching of skills is not possible.

On the other hand, not much research studies skills in the gig economy from a supply-side perspective, and our knowledge about how workers provide and expand skills through gigs on online labor markets is still limited. So far, there are only a few observational evidence suggesting that workers tend to provide gigs that belong to the same category (and thus require the same skill) or categories that share similar skillsets [16, 21]. Besides, it was also found that while workers indeed expand the range of skills that they provide over time, the marginal utility for them to develop new skills also diminishes [16]. Given our limited knowledge of the supply of skills in the gig economy, a systematic, data-driven study on how exactly workers provide skills in the gig economy is of great needs and importance. Such study can not only deepen our understandings on the ways that the public participate in this emerging new economy, but can also be helpful for addressing the skill supply-demand imbalance in the long-term—appropriate recommendations or interventions guiding workers to provide skills that are in need can only be made with a solid knowledge on how skills are currently provided by workers. This paper thus fills this gap by conducting a large-scale, empirical analysis on the provision of skills in gig economy. Different from previous studies, we propose to characterize the relationships between skills through a *network* and following this idea, we build a "*skill space*" from the empirical data. This is inspired by the concept of "product space" [14] in development economics which formalizes the relatedness between products traded in the global economy.

3 BACKGROUND AND DATA

3.1 The Fiverr Platform

We conduct our case study on a popular online labor marketplace, Fiverr, to empirically understand the provision of skills in gig economy.

More specifically, Fiverr is an online platform where registered workers can create a profile on it to include a list of "gigs" that they can perform. For each gig that a worker lists, she specifies the price and estimated delivery time, and she also needs to categorize it using a predefined hierarchical structure—Fiverr defines 9 possible categories of work for the platform, such as Graphics and Design, Programming and Tech and Writing and Translation. We refer to each category as a *domain*. Within each domain, Fiverr further defines subcategories to characterize work that requires specific skills, leading to a total of 104 subcategories across 9 domains. For example, the subcategory "*Legal Writing*" in the Writing and Translation domain requires skills on "writing legal terms of service and contract drafts." Thus, a worker's assignment of subcategories to gigs indicates the range of skills that she possesses; so in this study, we refer to each subcategory as a "*skill*." Note that workers can only assign a single subcategory (i.e., skill) to a gig, and workers may add new gigs to their profiles over time where the type of skills needed in the new gigs can be either the same or different from the worker's existing gigs.

On the other side of the market, potential employers can browse all available gigs on the platform, contact the workers for gigs that they are interested in for further information, and purchase gigs

Domain	9
Skill	104
Gigs	548,475
From Accessible Workers	522,391 (95.24%)
Average # of Reviews per Gig	15.86
Average Star Rating Score for a Gig	4.78
Worker	259,135
Accessible Worker	248,541 (95.91%)
Average # of Reviews per Worker	33.56
Average Star Rating Score for a Worker	4.67

Table 1. Overview of the dataset.

that satisfy their needs. Once a purchase is made on a gig, the worker of that gig will deliver the work through Fiverr, and the employer has the option of writing a review and rating the work quality of the gig on a scale of 0 (lowest) to 5 (highest) stars after they receive the deliverables.

Finally, based on workers' activity, performance, and reputation, Fiverr assigns a platform-wide "level" to each worker (i.e., New, Level 1, Level 2, and Top Rated), and workers with higher level obtain benefits such as the capability to offer more gigs and price gigs higher. Fiverr also updates the assignment of worker's level periodically.

We decide to conduct our case study with Fiverr for two reasons: First, different from many other platforms like Upwork where skills are represented as worker-generated, free-form labels, Fiverr predefines a set of skills and organizes them in a 2-layer hierarchy. This helps us to limit the scope of our study without loss of generality. Second, the amount of publicly available data for Fiverr is abundant.

3.2 Data Collection and Overview

We performed a crawl of the Fiverr website in May 2017, and fetched all gigs that were listed on Fiverr before the time of crawl. As a result, we obtained a dataset consisting of 548,475 unique gigs that were listed on Fiverr between February 2010 and May 2017. For each gig, we obtained its detailed information, such as the worker who offers the gig, the registered date representing when the gig was listed on Fiverr, the total number of reviews, and the average star rating score of the gig. Doing so, we find that these 548,475 gigs are provided by 259,135 unique workers. On average, each worker provides 2.12 gigs that are associated with 1.60 different skills, though workers show substantial variance in the number of gigs and skills that they provide (see Figures 1a and 1b). For each worker whose profile is still accessible[4], we further fetched detailed information on the worker, including the worker's level, country, time of registration on Fiverr, existence of social media accounts as identity verification, education history (if provided), list of certifications received (if provided), the average response time to employer's contacts, and the most recent date of delivery. Table 1 summarizes the dataset we collected, and Figures 1c-1f show the distributions for the number of reviews and average star rating score of a gig (or a worker).

4 CONSTRUCTING THE SKILL SPACE

To understand how workers provide skills on the gig economy platform Fiverr, we first use the collected dataset to construct a network of skills, or a "*skill space*," to characterize the relatedness between skills, similar to the "product space" in [14]. Such a network is built through analyzing how frequently pairs of skills are provided together by the same worker.

[4]A worker's profile becomes inaccessible when she quits the market, or when her account is disabled by the platform operators.

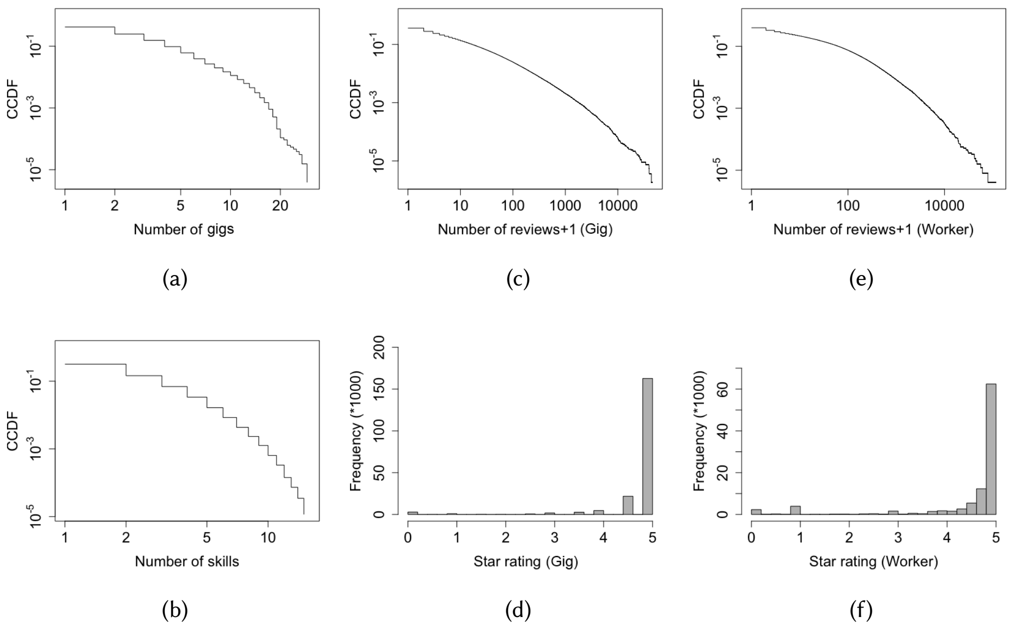

Fig. 1. 1a: Complementary cumulative distribution function of the number of gigs listed by a worker; 1b: Complementary cumulative distribution function of the number of skills provided by a worker; 1c: Complementary cumulative distribution function of the number of reviews given to a gig; 1d: Histogram of the average star rating score of a gig; 1e: Complementary cumulative distribution function of the number of reviews given to a worker (computed by aggregating all reviews given to all the gigs that the worker provides); 1f: Histogram of the average star rating score of a worker (computed by averaging all ratings given to all the gigs that the worker provides).

In particular, each node in the network represents a skill; a directed edge $E(i, j)$ exists if skills i and j have been provided together by some worker, and its weight $w(i, j)$ is defined as the fraction of workers who provide both skills i and j among workers who provide skill i. Doing so on the collected dataset, we get a network consisting of 104 nodes and 10,180 edges. Note that this is *not* a complete network as 266 (5.2%) pairs of skills have never been provided together by any worker in our dataset. We refer to this network as the *Skill Provision Co-occurrence Network (SPCN)*.

The SPCN may capture both real empirical skill provision pattern and some noise. To extract only the real pattern, we next generate a statistically validated network based on the SPCN that we have got. Specifically, for each of the 548,475 gigs included in our dataset, by fetching the worker who provides the gig and the skill that the gig is associated with, we obtain 548,475 worker-skill profiles. We then keep the worker list in these profiles unchanged, while randomly shuffle the skill list using the Fisher-Yates shuffle algorithm [10] to get a random set of worker-skill profiles. As a toy illustrating example, imagine the real worker-skill profiles are {(Worker A, Skill 1), (Worker A, Skill 2), (Worker A; Skill 1), (Worker B, Skill 3), (Worker B, Skill 2), (Worker C, Skill 2)}. After the random shuffle, one random set of worker-skill profiles may be {(Worker A, Skill 2), (Worker A, Skill 1), (Worker A; Skill 2), (Worker B, Skill 1), (Worker B, Skill 2), (Worker C, Skill 3)} (note that the order of workers does not change while the order of skills is shuffled). In the random worker-skill profiles, each worker "provides" the same number of skills as that in the real worker-skill profiles, and collectively, each skill is "provided" for the same number of times as that in the real worker-skill profiles, too. Thus, the random worker-skill profiles generated through this approach offer a realistic

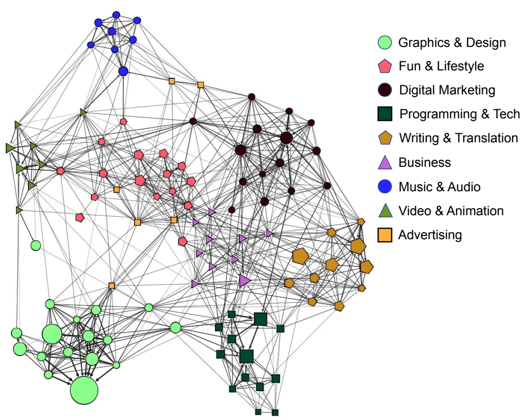

(a) The skill space

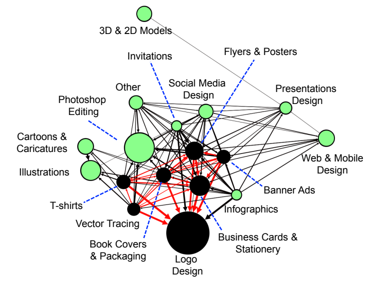

(b) Subnetwork for the Graphics and Design domain

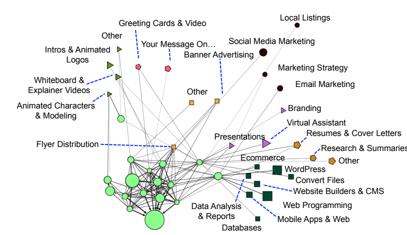

(c) Skills connected to the Graphics and Design domain

Fig. 2. 2a: The space of skills supplied in online labor market, Fiverr. The size of a node is proportional to the number of workers who provide the corresponding skill, and the thickness of an edge is proportional to the edge weight. 2b: The subnetwork for skills in the Graphics and Design domain (17 skills and 160 edges in total). Black nodes and red edges highlight the largest clique in this domain. 2c: Connections between skills in the Graphics and Design domain and skills in other domains.

simulation of how workers would provide skills if they randomly pick which skills to provide together.

Next, given one set of random worker-skill profiles, we again generate the skill provision co-occurrence network corresponding to it, which we refer to as the *random SPCN*. This random SPCN then reflects how frequently different skills would be provided together if workers decided to provide skills according to the chosen set of random profiles. We repeat this process for 1,000 times, that is, we generate 1,000 sets of random worker-skill profiles and correspondingly produce 1,000 random SPCNs.

Intuitively, if workers who provide skill i indeed tend to provide skill j as well, we would expect that in the original SPCN produced by the real worker-skill profiles, $w(i, j)$—which, recall, is the fraction of workers who provide both skills i and j among workers who provide skill i—is relatively high, while for the majority of the 1,000 random SPCNs, the weight of the directed edge from i to j is relatively low. To capture this intuition, we define the empirical p-value, $p(i, j)$, for the edge $E(i, j)$ in the original SPCN as the fraction of random SPCNs for which the weight of the directed edge from i to j is at least $w(i, j)$—the larger the p-value is, the less likely $E(i, j)$ is a real skill provision pattern. We keep an edge $E(i, j)$ in the statistically validated network only if $p(i, j) < 0.05$, which results in a network of 104 nodes and 1,283 edges. We refer to this statistically validated network as the "skill space" and analyses hereinafter are conducted based on this network, except for the analyses in Section 5.2, which are based on a series of additional skill spaces that we construct for different years to study how the empirical patterns in skill provision evolve over time.

Figure 2a gives a visual representation of the skill space. Intuitively, if an edge exists between skills i and j in the skill space, it means these two skills are frequently provided together by workers and thus there is likely a degree of relatedness between them. Interestingly, there are 546 pairs of skills for which directed edges for the two directions both exist in the skill space, leading to a total of 1,092 edges, while the rest of 191 edges only exist in a single direction between skill pairs. A closer look at the data suggests that among these 191 single direction edges, the majority of them (i.e., 77.5% of them, 148 edges) are from a skill with a smaller number of providers to another skill with a larger number of providers, which may indicate that workers with "niche skills" tend to provide both their niche skills and the high-level, more general skills at the same time.

5 RESULTS

In this section, we answer the 4 research questions on the provision of skills in gig economy with respect to the platform that we choose to study, i.e., Fiverr, based on the skill space we obtain for it.

5.1 RQ1: Empirical Patterns in Skill Provision

We start by analyzing the empirical patterns emerged in the provision of skills on Fiverr.

First, we conduct an in-depth case study on how skills are supplied by workers in one specific domain—Graphics and Design, which is the domain that has the largest number of providers on Fiverr according to the dataset that we collect. Figure 2b shows the *subnetwork* corresponding to skills in the Graphics and Design domain (i.e., 17 nodes and 160 edges). Visually, it is clear that within this domain, different sets of skills have significantly different degree of relatedness. For example, the largest clique (highlight in black nodes and red edges in Figure 2b) we find for this subnetwork includes 7 skills: *Banner Ads, Book Covers & Packaging, Flyers & Posters, Business Cards & Stationery, Logo Design, T-shirts,* and *Vector Tracing*. The average weight for edges in this clique is 0.14, implying that on average, given any two skills in this clique, 14% of workers who provide one skill also provide the other skill. On the other hand, the subnetwork also contains some sets of skills where not a single edge exists between any pair of skills in the set, and the largest size for such set is 5—for example, *Logo Design, Illustration, Photoshop Editing, Web & Mobile Design* and *3D & 2D Models* are 5 skills that are disconnected from one another. These observations suggest that "*clusters*" exist within skills of the same domain. A plausible explanation is that some skills are more similar (e.g., require similar knowledge, share similar tools, etc.) hence workers possess one skill likely possess the other as well, while some other skills are more inherently distinct and thus workers rarely provide them together. To further see the varying degree of relatedness between skills, we run a hierarchical clustering algorithm on all skills in the Graphics and Design domain[5] and results are shown in Figure 3. Clusters shown in the figure appear to suggest important distinctions on the nature of different skills—the red cluster contains skills that mostly require knowledge on design techniques, tools and principles, while the green cluster is for skills that are mostly related to drawing and the orange cluster represents the skill that needs understanding on front-end web development.

Moreover, skills in one domain can also relate to skills in other domains. Figure 2c shows all cross-domain connections for skills in the Graphics and Design domain. In total, there are 71 cross-domain edges, connecting Graphics and Design skills with skills in all other domains except for Music and Audio. Interestingly, many of these edges connect skills in different domains that

[5]We set the number of clusters as 6, which is the minimum number that gives us a structure where all disconnected pairs of skills are in different clusters.

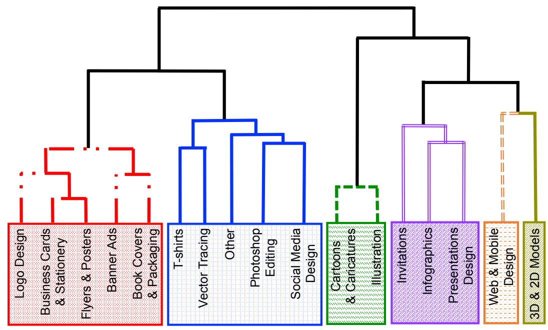

Fig. 3. Hierarchical clustering results for skills in the Graphics and Design domain.

are all needed in the workflow for completing some common complex job. For example, the skill *Web & Mobile Design* is closely related to many skills in the Programming and Tech domain such as *Web Programming* and *WordPress*, and these skills can all be components in workflows for building websites or mobile applications. That is to say, besides providing *similar* skills in the same domain, workers also tend to provide *complementary* skills from different domains that serve the same high-level purpose. Thus, "skill relatedness" defined through the skill space may capture either similarity or complementarity of different skills.

The insights that we obtain from our case study on skills in the Graphics and Design domain also apply to other domains, but the degree to which skills within a domain are related to each other or relate to skills in other domains vary across different domains. To quantify such differences, we use *edge density* and *weighted edge density* to measure how closely skills within one domain are related to each other. Specifically, given a domain d and its corresponding subnetwork $S(d)$ with n nodes and m edges, the edge density is defined as $\frac{m}{n(n-1)}$, i.e., the fraction of edges that actually exist in the skill space among all edges that could possibly exist in a network with n nodes [29]. The weighted edge density further considers the edge weight and is defined as $\frac{\sum_{E(i,j)\in S(d)} w(i,j)}{n(n-1)}$. Intuitively, higher values in edge density and weighted edge density imply that skills within a domain are more tightly connected and thus more related to each other. Results for the comparisons on these metrics are shown in Figures 4a-4b, which suggest that from the worker skill provision point of view, Music and Audio and Video and Animation are the two domains for which skills in them are most closely related, while the Fun and Lifestyle domain contains skills that are most loosely related.

Moreover, to measure how closely skills in one domain relate to skills in other domains, we define two additional metrics—*cross-domain edge ratio (CER)*, which is the ratio between the number of cross-domain connections for one domain and the number of connections within it, and *cross-domain edge weight ratio (CEWR)*, which is the ratio between the total weight of cross-domain connections and within-domain connections for a given domain. Figure 4c reports the comparisons on these two metrics, which indicate that skills in the Advertising domain relate to skills in other domains most heavily while skills in the Graphics and design domain relate to out-of-domain skills the least.

Lastly, characterizing the relationships between skills via the skill space also implies different properties of skills in the context of this network, such as their degree of connectedness to the neighbors and locations in the network. To further examine how workers provide various skills differently, we consider the following three network metrics for a skill:

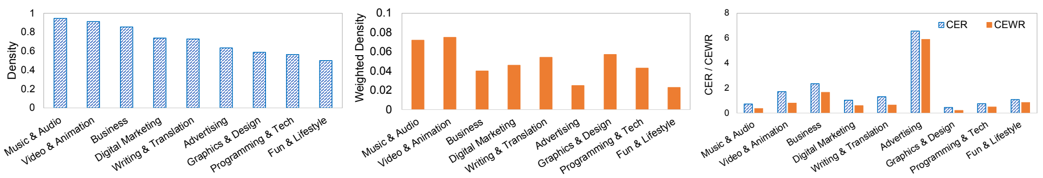

(a) (b) (c)

Fig. 4. Comparisons of skill relationships in different domains. 4a: Comparing the degree of skill relatedness within different domains through the density of edges within each domain. 4b: Comparing the degree of skill relatedness within different domains through the weighted density of edges within each domain. 4c: Comparing the degree of cross-domain skill relatedness through cross-domain edge ratio (CER) and cross-domain edge weight ratio (CEWR).

	# of Providers	**# of Gigs**	**Average Price**
Weighted degree	0.44^{***}	0.45^{***}	-0.07
Clustering coefficient	0.34^{***}	0.36^{***}	-0.03
Eigencentrality	$-0.18^{\dagger}$	-0.19^{*}	-0.02

Table 2. Pearson correlation coefficients between network metrics of a skill in the skill space and metrics characterizing the ways workers provide that skill. The statistical significance is marked as a superscript, with †, * and *** representing significance levels of 0.1, 0.05 and 0.001 respectively.

- *Weighted degree*: The sum of weights for all edges in the skill space that start from or end at the current skill. The higher the value, the more related the current skill is to other skills.
- *Clustering coefficient*: A measure introduced by [30] to quantify how likely the neighbors of a node (i.e., the current skill) are to form a clique. The higher the value, the current skill is related to a set of skills that are more related to each other.
- *Eigencentrality*: A measure characterizing how central a node (i.e., the current skill) is in a network [4]. The higher the value, the current skill relates to a larger number of more "central" skills[6].

Table 2 reports the correlation between these network metrics of a skill with a few measures quantifying how workers provide that skill, including the number of workers provide that skill, the number of gigs that are associated with that skill, and the average price that workers ask in a gig supplying that skill. It is observed that a larger number of workers provide skills that are more closely-related to other skills (i.e., skill with higher weighted degree), leading to more gigs being provided on those skills. Similarly, more providers and gigs are found for those skills that belong to a tight-knit group of related skills (i.e., skill with higher clustering coefficient).

On the contrary, a significantly *negative* correlation is detected between a skill's eigencentrality in the skill space and the number of workers/gigs associated with that skill. A closer look at the data suggest that skills that have higher eigencentrality values tend to be skills that are somewhat related to a substantially large number of other skills from different domains (i.e., the degree of the skill node is large yet the weight for edges associated with the node is low), such as "*Branding Service*" in the Busniess domain and "*Marketing Strategy*" in the Digital Marketing domain. Our results here then imply that, in general, workers prefer to provide skills that are closely-related to a set of mutually relevant skills rather than providing skills that are loosely-related to a large number

[6]Note that in the computation of both clustering coefficient and eigencentrality, we treat each edge in the network as binary—exist or not—and discard the weight information.

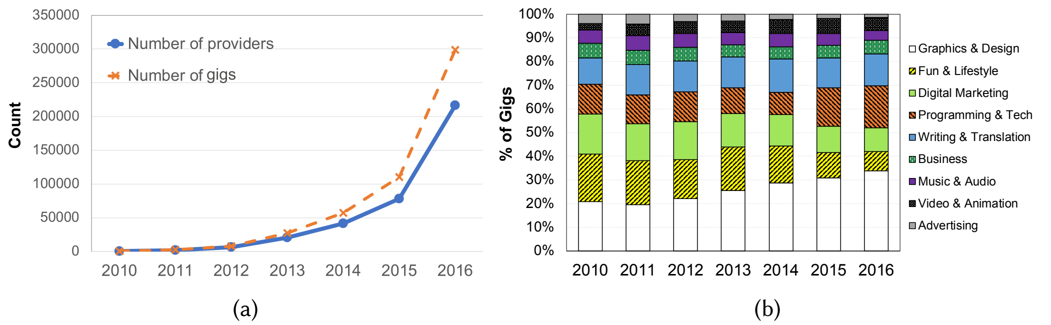

Fig. 5. 5a: Increases in the total number of Fiverr workers and gigs over time; 5b: the fraction of gigs supplying skills in each domain by the end of each year.

of skills. With respect to average price, however, we do not find it to be significantly correlated to any of the network metrics, suggesting that there is no clear relationship between a skill's property in the skill space and its price.

In sum, our examination on the empirical patterns in the provision of skills on Fiverr reveals that on one hand, skills in the same domain exhibit varying degree of relatedness with each other, leading to clusters of skills within the same domain which is not captured by domain-level characterizations of skills like that used in [21]; on the other hand, taxonomy or ontology of skills as used in [8, 24] may fall short in capturing the relationship between skills from different domains. This implies a highly sophisticated relationship between various skills in gig economy that can be best characterized through a network. Furthermore, it is also empirically observed that more workers tend to provide more gigs on skills that are closely-related to some other skills, potentially even forming a tight-knit set of coherent skills, rather than those skills that are loosely-related to a large number of other skills.

5.2 RQ2: Dynamics in Skill Provision Patterns

We now move on to our second research question to understand how the empirical patterns in skill provision among Fiverr workers evolve over time.

Figure 5a first shows that, unsurprisingly, between 2010 and 2016, there is an exponential growth in terms of both the number of workers providing skills on Fiverr and the number of gigs registered on Fiverr. To see what skills do workers provide on the market over the years, we compute the fraction of gigs that are associated with skills in each domain by the end of each year and results are presented in Figure 5b. It is clear from the figure that workers are providing significantly more gigs requiring skills in the Graphics and Design and Programming and Tech domains in recent years, while comparatively, the growth of skill provision in the Fun and Lifestyle and Digital Marketing domains is slower. Table 3 further lists the top 10 skills with the largest number of gigs provided by workers by the end of 2012, 2014, and 2016. Again, we can see that skills in the Graphics and Design domain become increasingly heavily supplied by workers on Fiverr over the years.

A more interesting aspect to explore with respect to the evolvement of skill provision is whether and how the patterns in the ways that workers provide multiple skills together change over time. To this end, we reconstruct the skill space following the methods we describe in Section 4, but restrict ourselves only to the gigs that are registered on Fiverr before a certain point in time. Figures 6a-6c depict the skill space by the end of 2012, 2014, and 2016, respectively. Clearly, from these figures, we find workers not only gradually adjust the type of skills that they provide (as shown by the relative size of nodes in the skill space across different years), but also tend to, over time, provide

By Dec. 2012	By Dec. 2014	By Dec 2016
Digital Marketing: SEO Services	Graphics & Design: Logo Design	Graphics & Design: Logo Design
Fun & Lifestyle: Other	Graphics & Design: Photoshop Editing	Graphics & Design: Photoshop Editing
Graphics & Design: Logo Design	Programming & Tech: Web Programming	Programming & Tech: Web Programming
Graphics & Design: Photoshop Editing	Digital Marketing: SEO Services	Programming & Tech: Wordpress
Programming & Tech: Web Programming	Programming & Tech: Wordpress	Digital Marketing: SEO Services
Writing & Translation: Article & Blogposts	Fun & Lifestyle: Arts & Crafts	Writing & Translation: Article & Blogposts
Digital Marketing: Social Media Marketing	Graphics & Design: Illustrations	Graphics & Design: Illustrations
Digital Marketing: Other	Fun & Lifestyle: Other	Writing & Translation: Translation
Programming & Tech: Wordpress	Digital Marketing: Other	Graphics & Design: Flyers & Posters
Writing & Translation: Other	Digital Marketing: Social Media Marketing	Video & Animation: Intros & Animated Logos

Table 3. Top 10 skills with the largest number of gigs provided on Fiverr by the end of 2012, 2014, and 2016 (lists are ordered).

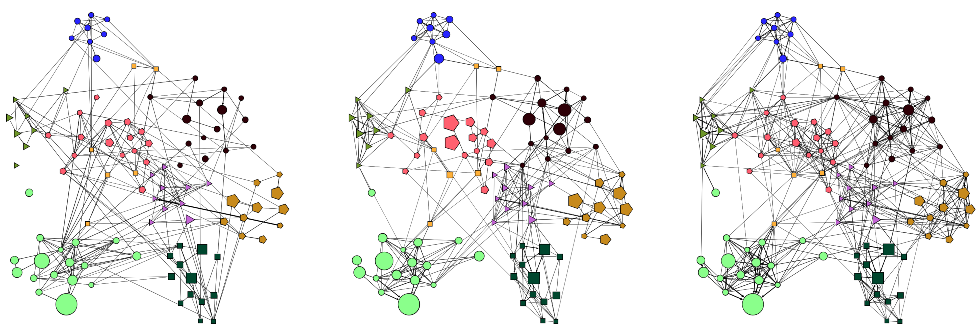

(a) By Dec 31, 2012 (476 edges) (b) By Dec 31, 2014 (632 edges) (c) By Dec 31, 2016 (1094 edges)

Fig. 6. The change of skill space over time. Within each plot, the size of a node is proportional to the number of workers who provide the corresponding skill, and the thickness of an edge is proportional to the edge weight.

many more pairs of skills that they previously seldom provide together—Indeed, while the skill space by the end of 2012 only contains 476 edges between skills and there are even some "isolated" skills (e.g., "*3D & 2D Models*" in the Graphics and Design domain) for which the providers don't seem to supply them together with any other skills, by the end of 2016, the number of edges in the skill space has reached 1094, double that in 2012. This indicates that over time, workers in the gig economy are changing the skill provision strategies, likely from concentrating on providing one particular skill to supplying a set of skills together.

A closer examination on the change of skill space suggest that by the end of 2012, 2014, and 2016, the fraction of edges in the skill space that connect two skills within the same domain is 63.0%, 70.1%, and 67.6%, respectively, and by the time we crawl the data (i.e., May 2017), this fraction further decreases to 62.6%. This implies that workers in the gig economy have first experienced a period when they increasingly provide more skills within the same domain together, and more recently they start to increasingly explore the potential of supplying skills from multiple different domains together. In addition, as shown in Table 4, comparing the skill space by the end of 2016 with that by the end of 2012, we also find that workers develop new "relationships" for skills in the Digital Markting domain the most—the pairs of skills within the Digital Markting domain that are frequently provided together by workers by the end of 2016 is 6 times as many as that by the end of 2012, and the pairs of skills between the Digital Markting domain and another domain that are frequently provided together by workers by the end of 2016 is 4.33 times as many as that by the end of 2012.

Domain	Ratio for Within-Domain Edges	Ratio for Cross-Domain Edges
Music and Audio	1.86	1.9
Video and Animation	2.3	2.67
Business	1.95	2.71
Digital Marketing	**6**	**4.33**
Writing and Translation	4.88	2.13
Advertising	**1.5**	2.15
Graphics and Design	1.79	**0.92**
Programming and Tech	2.2	1.09
Fun and Lifestyle	2.26	1.78

Table 4. The ratio between the number of within-domain (or cross-domain) edges existing in the skill space by the end of 2016 and the number of within-domain (or cross-domain) edges existing in the skill space by the end of 2012. The largest and smallest values in each column are highlight in red and blue, respectively.

Region	Country
Industrialized countries	Australia, Canada, France, Germany, Israel, Italy, Netherlands, Spain, United Kingdom, United States
South Asia	Bangladesh, India, Pakistan, Sri Lanka
Africa	Egypt, Morocco, Nigeria, Kenya, South Africa
Southeast Asia	Indonesia, Malaysia, Philippines, Singapore, Vietnam
Latin America	Brazil, Jamaica, Mexico, Venezuela
East Europe	Romania, Russia, Serbia, Ukraine

Table 5. The list of 32 countries that have more than 1,000 workers on Fiverr and their regional division.

5.3 RQ3: Differences between Workers in Providing Skills

Next, using the skill space, we examine how workers with different characteristics provide skills differently. In this study, we focus on understanding whether and how workers from different regions of the world show distinct patterns in providing their skills in the gig economy. We consider workers from 32 countries which are all the countries that have at least 1,000 Fiverr workers in our dataset, and we group these countries into 6 geographical regions as shown in Table 5.

First, we look into the "skill speciality" for workers in each geographical region. That is, given a particular region, we identify those skills for which within that region, the fraction of workers who provide the skill is higher than the global fraction, and we visualize this in Figure 7 through the color of nodes—the darker a node's color is, the larger fraction of workers in that region provide the corresponding skill compared to the global fraction. Visually, it is clear that there are substantial regional variations in the specialization of skills. For example, a significantly higher fraction of South Asian workers provide skills in the Graphics and Design and Programming and Tech domains, while workers from Africa provide skills in the Writing and Translation domain significantly more often than the global average. Even for skills in one domain like Writing and Translation, workers from different regions also have varying focuses—East European workers focus on *Translation* exclusively, Latin American workers mostly specialize on both *Translation* and *Transcription*, while African workers intensively provide a wide range of other skills in this domain like *Articles & Blogposts* and *Proofreading & Editing*.

Moreover, to see whether there is any regional difference in terms of the set of skills that workers tend to provide together, for each of the six regions, we construct the validated SPCN corresponding to that region using only the gigs provided by workers from that region. Edges in these regional validated SPCNs are highlight in red in Figure 7[7]. In other words, given a specific region, an edge

[7] We leave out edges in regional validated SPCNs that don't exist in the skill space, i.e., the global validated SPCN.

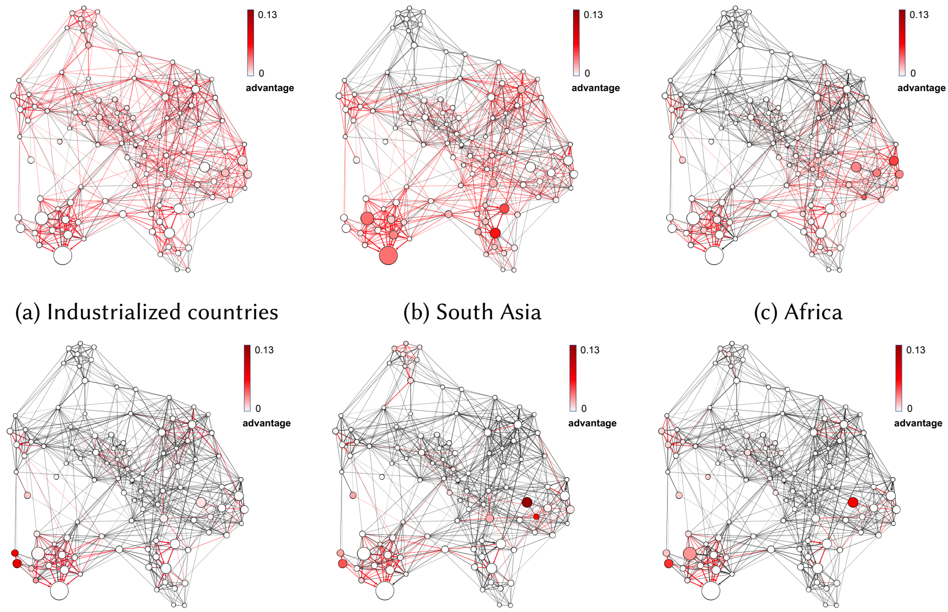

(a) Industrialized countries (b) South Asia (c) Africa

(d) Southeast Asia (e) Latin America (f) East Europe

Fig. 7. Differences in provision of skills for workers from different regions of the world. Within the subplot of a given region, a node is colored if workers of this region specialize in the corresponding skill and a darker color means a larger difference between the regional and global fraction of providers for the corresponding skill (i.e., "advantage"); an edge is highlight in red if that edge exists in the validated network for workers of the given region.

between skills i and j exists in its regional validated SPCN (and thus highlight in red in Figure 7) if these two skills are often provided together by workers *from that region*. Figure 7 shows that workers from different regions indeed have *distinct* skill provision patterns. For example, workers from industrialized countries (e.g., United States) and Latin America often provide skills in the Music and Audio domain together, while workers from South Asia, Africa and Southeast Asia seldom do so, which implies the high degree of relatedness between Music and Audio skills are largely due to skill provision patterns from workers in the former group. On the macro-level, we also observe that different from workers in industrialized countries who frequently provide various different pairs of skills within a domain or across different domains together, workers in Southeast Asia, Latin America and East Europe tend to frequently provide together only a few pairs of skills, mostly within certain domains.

5.4 RQ4: Skill Provision and Worker Success

Finally, we aim to answer our last research question: Do the ways that workers provide skills affect their success in the gig economy?

5.4.1 Skill Provision Strategy and Overall Performance. We first approach our research question from a static point of view. That is, we start off by investigating the associations between a worker's strategy on providing *a set of skills* and her *overall* performance on all the gigs/skills that she provides.

Dependent variables. We use three dependent variables to capture a worker's overall performance on the platform:

- *Number of reviews*: The total number of reviews a worker receives across all the gigs that she provides.
- *Fraction of rated gigs*: The fraction of gigs that a worker provides for which at least one review is received.
- *Average star rating score*: The average star rating a worker receives across all the gigs that she provides.

Intuitively, the more reviews a worker obtains, the more popular a worker's gigs are among employers; the larger a worker's fraction of rated gigs is, the more employers are willing to purchase different kind of gigs provided by the worker; and the higher a worker's average star rating score is, the more satisfied employers are with the quality of work produced by the worker. All of these, therefore, are used as indicators of a better overall performance of the worker in this labor market.

Independent variables. We then use two measures to quantify a worker's strategy in providing skills, mainly with respect to how "related" the skills provided by a worker is.

- *Validated Edge Fraction (VEF)*: Suppose a worker provides n skills. VEF captures that among the $n(n-1)$ possible directed edges among the n skills, the fraction of directed edges that actually exist in the skill space.
- *Average Validated Edge Weight (Avg. VEW)*: The average weight for all directed edges in the skill space connecting skills that are provided by the worker.

As discussed before, a directed edge existing between two skills in the skill space indicates that those two skills are frequently provided together and are likely related. Thus, a larger value for either VEF or $Avg.VEW$ suggests that the set of skills provided by the worker are more related. A nuanced difference between these two metrics is that VEF captures *how many* of the skills a worker possesses are related, while $Avg.VEW$ characterizes among the worker's related skills, *how closely-related* they are to each other. Thus, a worker who provides a set of skills that are somewhat related to each other would have a high VEF and a low $Avg.VEW$, yet another worker who possesses a few pairs of skills that are very frequently provided together as well as some additional entirely unrelated skills will have a low VEF yet a high $Avg.VEW$.

Providing more related skills correlates with better worker performance. Table 6 reports the results of three groups of regression models where dependent variables are the number of reviews, fraction of rated gigs and the average star rating score, respectively, and quantifications of the relatedness between the skills that a worker provides (i.e., VEF and *Avg. VEW*) are used as independent variables[8]. We also control for other covariates like the number of skills the worker has and a variety of features characterizing the worker's usage of Fiverr[9].

In particular, as the data on the number of reviews are over-dispersed, we adopt Poisson regressions to properly analyze these data. Models 1a and 1b in Table 6 report the result, which reveal that the degree of a worker's skill relatedness, as measured by both the validated edge

[8]In this analysis, we restrict ourselves to workers who provide at least 2 different skills (hence skill relatedness can be measured) and have at least one review (hence worker performance is well-defined).

[9]"Social account," "Education history" and "Certifications" are three binary covariates indicating whether or not the worker provides the corresponding information. "Average response time" represents the average number of hours elapsed between an employer's contact and the worker's response; hence the smaller the number is, the faster a worker responds. "Most recent delivery" is encoded with 5 levels: within 1 day (Level 1), 1 day to 1 month (Level 2), 1 month to 1 year (Level 3); more than 1 year (Level 4); no record (Level 5).

	Number of Reviews		**Frac. of Rated Gigs**		**Avg. Star Rating**	
	Model 1a	Model 1b	Model 2a	Model 2b	Model 3a	Model 3b
Intercept	0.988***	1.004***	0.432***	0.451***	5.146***	5.139***
VEF	0.423***		0.033***		0.013	
Avg. VEW		2.219***		0.014		0.206***
Number of skills	0.019***	0.002***	-0.029***	-0.030***	-0.003	-0.004*
Level	1.515***	1.528***	0.151***	0.152***	0.121***	0.120***
Number of months registered	0.032***	0.033***	0.002***	0.002***	-0.007***	-0.006***
Social account	0.059***	0.050***	-0.002	-0.003	0.119***	0.119***
Education history	-0.122***	-0.134***	-0.007**	-0.007**	0.051***	0.050***
Certifications	0.063***	0.077***	-0.003	-0.002	-0.006	-0.007
Avg. response time	-0.001***	-0.001***	-0.000	-0.000*	0.001***	0.001***
Most recent delivery	-0.860***	-0.862***	-0.028***	-0.029***	-0.257***	-0.256***
R^2	0.680	0.680	0.320	0.318	0.191	0.191

Table 6. Regression models examining how the relatedness between skills a worker provides correlates with her overall performance. The statistical significance is marked as a superscript, with *, ** and *** representing significance levels of 0.05, 0.01 and 0.001 respectively.

fraction and the average validated edge weight, has a statistically significant *positive* correlation with the total number of reviews the worker receives. This indicates that, all else equal, choosing to provide a set of skills that are more related to each other is associated with obtaining more reviews. This observation is consistent with the conjecture that a higher degree of skill relatedness for a worker is correlated with a better overall performance of the worker on the platform. In addition, all covariates we include in the models are found to be significantly correlated with the number of reviews a worker can get. For example, the significantly positive coefficient estimated for the number of skills suggest that the more skills a worker possesses, the more reviews she accumulates. Surprisingly, the correlation between the existence of education history on a worker's profile and the number of reviews the worker gets is significantly negative, implying that workers who choose to provide their education history get fewer reviews in general. A closer look into the data suggests that compared to workers who do not provide their education history, on average, workers who choose to reveal their education history provide slightly smaller number of gigs and become registered Fiverr users for significantly shorter amount of time, which both may contribute to the smaller number of reviews received by them.

Similarly, Models 2a and 2b look into the relationships between the degree of a worker's skill relatedness and the fraction of rated gigs, and Models 3a and 3b explore the correlations between skill relatedness and the average star rating score a worker obtains. Here, we observe that for a worker, possessing *more* skills that are related to each other significantly correlates with having a *larger* fraction of gigs being rated, while possessing some *closely-related* skills associates with a *higher* average rating score. Again, this supports the conjecture that workers who have higher degree of skill relatedness are also the ones who perform well on the platform.

Statistically significant relationships are also detected for most of the covariates that we include in both sets of models, among which a particularly interesting one is the *negative* association between the number of skills a worker provides and her fraction of rated gigs (or average rating score). Combing together with our previous observations, these results imply that while workers who provide a large set of skills are less likely to achieve high performance (with respect to the worker's fraction of rated gigs and average star rating), workers who can provide a large set of skills *with high degree of relatedness* may have a higher chance to perform better. One possible explanation for this phenomenon is that providing related skills is less mentally-taxing than providing unrelated skills, hence the performance decrease caused by the allocation of a worker's limited resource (e.g., attention, time and energy) to different skills is smaller.

5.4.2 Skill Expansion Strategy and Performance on New Skills. Now, we approach our last research question from a dynamic point of view. Specifically, workers in the gig economy can constantly add new skills into, and thus expand, their skillsets. We are therefore interested in understanding how worker's strategy in expanding their skillsets correlates with their success on the new skills.

Formally, let's consider a particular worker w who provides k gigs in total. We first sort all these gigs according to the increasing order of the gig registration time, which results in a sorted gig list $G_w = \{g_w^1, g_w^2, \cdots, g_w^k\}$ and a sorted gig registration time list $D_w = \{d_w^1, d_w^2, \cdots, d_w^k\}$. By obtaining the skills associated with each gig, we further get a sorted skill list for the worker $S_w = \{s_w^1, s_w^2, \cdots, s_w^k\}$. Note it is possible that $s_w^{t_1} = s_w^{t_2}$ for some $t_1 \neq t_2$, as the worker may provide multiple gigs that are associated with the same skill at different time points.

We define a worker's "expansion" of skills as follows: Given two different skills i and j that the worker w provides, suppose the *earliest* gig in which the worker provides skill i is x (i.e., $s_w^x = i$ and for $x' \in \{1, 2, \cdots, x-1\}$, $s_w^{x'} \neq i$) and the *earliest* gig in which the worker provides skill j is y. When $x < y$ and the time elapsed between d_w^x and d_w^y is *at least* one day, we say worker w expands her skillset from skill i to skill j; the directed edge $e(i, j)$ is called an *expanded edge* for worker w, and the amount of time taken for this expansion to happen is $DG_w(i, j) = d_w^y - d_w^x$. We thus ask how the strategy of skill expansion—that is, the property of the expended edge $e(i, j)$—relates to the worker's performance on the new skill j.

Dependent variables. Given a specific expansion of skill from i to j (i.e., the skill expansion captured by $e(i, j)$), we use three dependent variables to describe worker w's performance on the newly developed skill j:

- *Expansion time*: The amount of time it takes to complete this expansion, that is, $DG_w(i, j)$.
- *Average number of reviews on the new skill*: The average number of reviews across *all* gigs associated with skill j that the worker provides after she expands to skill j.
- *Average star rating score on the new skill*: The average star rating score across *all* gigs associated with skill j that the worker provides after she expands to skill j.

Intuitively, a shorter expansion time, a higher number of reviews on the new skill, or a higher average rating score on the new skill all indicates a higher performance on the new skill.

Independent variables. Given a specific expansion of skill from i to j, we then use two measures to characterize the property of this skill expansion, with respect to how much the newly developed skill j relates to the existing skill i.

- *Existence in Skill Space*: A binary variable indicating whether the skill space includes the directed edge from i to j.
- *Validated Edge Weight (VEW)*: The weight of the directed edge from i to j in the skill space.

Similar as before, while both independent variables capture the degree of relatedness between the new skill j and the existing skill i, they have some nuanced difference: the existence in skill space focus on examining *whether* the two skills are related, while VEW quantifies how *closely-related* the two skills are.

Extending to a skill that is related to the existing skillset correlates with better performance on the new skill. As a first practice, we first look into how worker's performance on the new skill is affected by whether the skill expansion follows an edge in the skill space or not.

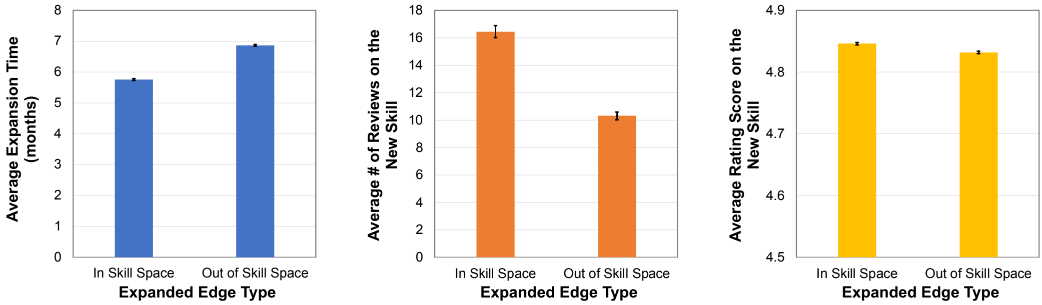

(a) Average expansion time (b) Avg. # of reviews on new skill (c) Avg. rating score on new skill

Fig. 8. Comparisons on speed and performance of skill expansions when the expended edge is in or out of the skill space. Error bars represent standard errors of the means.

Figures 8a-8c show the comparisons on the expansion time, the average number of reviews on gigs supplying the new skill, and the average star rating score on gigs supplying the new skill, respectively, when the expended edge is in or out of the skill space, averaged across all expanded edges for all workers. We clearly see that when workers follow an edge in the skill space to expand their skills, it takes workers *less time* to complete the expansion, and workers also get significantly *more* reviews and *higher* ratings in gigs where the new skill is used. Two-sided t-tests further confirm that all the differences are statistically significant ($p < 0.001$). In other words, we find that workers who choose to expand to skills that are related to their existing skills tend to perform better on the new skill.

To further see how the degree of relatedness between the new skill and existing skill correlates with the worker's performance on the new skill, we build three regression models where the dependent variables are the expansion time, the average number of reviews on the new skill, and the average star rating score on the new skill, respectively, while the independent variable is the validated edge weight (i.e., VEW) for the expended edge. We restrict this set of analyses for only those skill expansions that follow an edge in the skill space. Table 7 reports the results.

In particular, we use Poisson regression model for the expansion time data given it is over-dispersed, and we control a variety of covariates (e.g., the number of skills the worker already provides when the expansion happens, the worker's level, etc.). Result shown in Table 7 (the column for expansion time) suggests a statistically significant *negative* correlation between VEW and the expansion time, indicating that it takes workers *less* time to expand into a new skill that is *more closely-related* to the existing skill.

In addition, Table 7 (the right two columns) also reports results of regressions which examine the associations between the level of relatedness between the pair of skills in an expansion and the average number of reviews a worker receives on the new skill (or the average rating score a worker receives on the new skill). In these regressions, we include some additional covariates, that is, the average number of reviews a worker receives on her existing skill (or the average rating score a worker receives on her existing skill). It is shown that VEW *positively* correlates to both metrics of worker performance on the new skill. The positive association with the average rating score a worker can obtain on the new skill is also confirmed to be statistically significant. In other words, workers who expand to skills that are *more closely-related* to their existing skill are able to obtain significantly *higher* rating scores on gigs associated with the new skill.

	Expansion Time	Avg. # of Review (New)	Avg. Rating Score (New)
Intercept	0.209***	18.248***	4.735***
VEW	-0.337***	5.869	0.196***
Avg. # of reviews (Old)		0.104***	
Avg. rating score (Old)			0.020***
# of existing skills	0.089***	-3.786***	-0.007***
Level	0.154***	4.790***	0.084***
# of months registered	0.027***	0.620***	-0.002***
Social account	0.048***	0.118	0.050***
Education history	-0.016***	-1.326	0.023***
Certifications	-0.025***	0.463	-0.010
Avg. response time	0.000 ***	-0.001	0.000
Most recent delivery	-0.005*	-10.657***	-0.090***
R^2	0.299	0.106	0.060

Table 7. Regression models examining how the degree of relatedness between pairs of skills in an expansion correlates with worker performance on the new skill. The statistical significance is marked as a superscript, with * and *** representing significance levels of 0.05 and 0.001 respectively.

6 CONCLUSION AND DISCUSSION

This paper presents a large-scale, data-driven empirical case study on Fiverr to understand the provision of skills in gig economy. We adopt a network perspective to represent relationships between various skills provided in the gig economy and thus construct a “skill space” to characterize the relatedness between these skills. The skill space suggests interesting empirical patterns in how workers provide skills—on the one hand, clusters exist among skills of the same domain and workers tend to frequently provide similar skills within one cluster together; on the other hand, workers also often provide complementary skills in different domains that serve the same high-level purposes together. Workers also gradually adjust the type of skills they provide over time, and an increasingly large number of pairs of skills are frequently provided together by workers in recent years. In addition, workers from different geographical regions display significantly different patterns in how they provide skills. Importantly, we also find that how workers provide or expand their skills in online labor markets has important associations with how well they perform in these markets. Specifically, workers who provide a set of skills which have a higher degree of relatedness between each other tend to exhibit better overall performance, and workers who expand to new skills that are closely related to their existing skills also tend to complete such expansion faster and obtain better performance on the new skill.

Practical Implications of Our Findings. Our findings have many practical implications for workers and platform providers of the gig economy. For workers, as the high relatedness of skills for both worker’s provision and expansion of skills is shown to be associated with better worker performance, a straight-forward lesson here is that instead of simply following the demand and attempt to expand to skills that are in demand, it might be a better idea for workers to expand to skills that are highly-related to their existing skills which are also relatively in demand. On the other hand, workers who plan to develop a new skill that is distinct from their current skillsets may want to consider to first develop an intermediate skill that can bridge the gap between the new skill and existing skills (i.e., expand to a skill that is related to both the new skill and existing skills first) to ensure a high performance on the new skill. Yet, as we will discuss later in the limitations of our work, our findings can only suggest correlations between worker’s skill provision and expansion pattern and their performance. Rigorous causal analyses are needed for better assisting workers in their skill development on gig economy platforms.

From the platform provider's point of view, the "skill space" presented in this paper opens doors to opportunities for redesigning the operation of the gig economy platforms. First, the bottom-up emergent skill relationships as illustrated by the skill space can supplement the platform's existing pre-defined skill taxonomy, and even allow the platform to re-classify their skill taxonomy. For example, as we find that on Fiverr, skills in the Advertising domain barely relate to each other but highly relate to skills in other domains, Fiverr may want to consider moving skills under the Advertising domain to other more relevant domains[10]. Our findings can also be useful for guiding the design of advanced interventions in balancing the supply and demand of skills on gig economy platforms. For instance, inherent relationships between skills as revealed in our study can potentially be used for further optimizing the work assignment to match the right workers to right jobs and employers. In addition, platforms can also adopt algorithmic approaches to provide strategic guidance to workers (e.g., recommend skills to workers), which both help workers to develop successful "career paths" in online labor markets and nudge workers into areas that are in shortage of labor.

Importantly, we note that while our study indeed shows that workers of different characteristics, such as workers coming from different countries, provide skills in different ways, platforms should avoid stereotyping workers or even take advantage of these differences to provide differential treatment to workers. Platforms should offer equal opportunities and resources to gig workers with various backgrounds, and potentially make recommendations and suggestions to workers if needed, but leave the final right of decisions to workers themselves in terms of how they would like to participate in the gig economy. As an increasingly large population of workers choose the temporary work arrangements provided by gig economy platforms, platforms essentially function as a virtual workplace and thus it is critical for them to create a "culture" that embraces diversity and fairness. Such a culture would be beneficial for both workers' personal development and the platform's growth in the long term.

Limitations and Future Work. Our current study suffers from a number of limitations. First of all, in this work, we are only able to conduct a *case study* with one particular platform provider, Fiverr, to understand the provision of skills in gig economy, due to the limited amount of data accessible from other platforms. As a result, some subsets of results reported in this study may be attributed to the specific ways that Fiverr operates and may not generalize to other platforms.

However, we would like to highlight on a few methods and results that we believe are applicable for the broader contexts of gig economy: First, the key idea of this study, that is, characterizing the relationship between various skills through a network, is generally applicable to different platforms in the gig economy. We expect applying this methodology to a larger range of empirical data from different platforms would enable us to observe how skill provisions on different gig economy platforms are similar to or different from each other, and thus provide us with a more comprehensive picture on the provision of skills in gig economy today. Second, we conjecture that our findings of that the relatedness of the skills provided by a worker positively correlates with the worker's overall performance on the gig economy platform, and the degree of relatedness between an expanded skill and the existing skillset positively correlates with the worker's performance on the new skill, can be generalized to other platforms. Additional empirical studies need to be done with data from other gig economy platforms to confirm this conjecture. Third, many of the implications of this study as discussed above, such as leveraging the skill patterns emerged from

[10]In fact, as of June 2019, the Advertising domain has indeed been removed from Fiverr.

workers' provision of skills to supplement current skill taxonomy on gig economy platforms and to facilitate demand and supply matching or skill recommendation, can also be relevant for online labor markets other than Fiverr.

Another limitation comes from the kind of data we can obtain from Fiverr. As we collect our dataset by crawling the Fiverr website, our analyses are constrained by the range of data that are shown on the Fiverr webpages. For example, as shown in Tables 6 and 7, when we fit regression models to understand how the degree of relatedness among a worker's skills associates with the worker's overall performance as well as her performance on newly-developed skills, we find the fitted models have relatively large intercepts and low R^2. While the large intercepts estimated are consistent with recent findings which suggest that reviews on online platforms are substantially positively biased [1, 26], the relatively poor model fits could also be partly caused by the possibility that additional factors that are important for explaining worker performance are not captured by the models. Examples of such factors include the accuracy of the worker's gig descriptions and the communication messages between workers and employers, which are all beyond the access of a public web crawl.

In addition, this study only quantitatively examines the *correlation* between skill relatedness in worker's provision / expansion of skills and worker performance. Some of our observations align with conventional learning theory—for example, the "construction of knowledge" theory suggests that when students learn new skills by building new knowledge and skill on their prior knowledge, their learning performance can be enhanced [3, 12]. However, we note that our findings should *not* be interpreted as causal results, and further studies are needed to thoroughly understand the causal relationship between the worker skill development process and worker performance. For example, semi-structured interviews can be conducted with gig economy workers to qualitatively examine how workers decide which skills to expand and how they perceive the skill provision and expansion affects their success on gig economy platforms.

Another critical future direction of this work is to move beyond the question of "how do workers provide skills in the gig economy" to the deeper question of "why do workers provide skills in this way," especially in terms of what determines the evolvement in the workers' skill provision patterns. To this end, a qualitative, interview-based study will again help. It will also be interesting to supplement the current study with another examination on the relationship between different skills from the demand-side point of view to investigate how the empirical patterns for skill consumption and provision impact and co-evolve with each other. In addition, further studies can be carried out to understand whether and how the changes of policies or regulations on the gig economy platform and the changes of the general economic environment contribute to the evolvement of workers' skill provision patterns in the gig economy over time.

Finally, we note that as the gig economy has moved from being primarily composed of microtasks to containing more skilled work, there is a rich amount of research opportunities emerging for better understanding skilled-work-based gig economy from perspectives that are unique to it. Our study, in some sense, is an example of such kind of research—the skilled-work-based gig economy is often driven by the supply side and thus there is a need for obtaining deeper insights on the provision of skills from a supply point of view. We hope more research will be conducted in the future to thoroughly examine the unique aspects of skilled work in the gig economy and gain deeper knowledge of how skilled work and microtask based gig economies are similar to or different from each other.

ACKNOWLEDGMENTS

This work is partially supported by Cybersecurity at MIT Sloan, which is funded by a consortium of organizations, and Purdue University. The authors are grateful to the anonymous reviewers of

CSCW 2019 for their helpful comments to improve this study. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors alone.

REFERENCES

[1] Susan Athey, Juan Camilo Castillo, and Dan Knoepfle. 2018. Service Quality in the Gig Economy. (2018).

[2] Senjuti Basu Roy, Ioanna Lykourentzou, Saravanan Thirumuruganathan, Sihem Amer-Yahia, and Gautam Das. 2015. Task assignment optimization in knowledge-intensive crowdsourcing. *VLDB Journal* 24, 4 (2015), 467–491. https://doi.org/10.1007/s00778-015-0385-2 arXiv:1401.1302

[3] Victor A Benassi, Catherine Overson, and Christopher M Hakala. 2014. Applying science of learning in education: Infusing psychological science into the curriculum. (2014).

[4] Phillip Bonacich. 1987. Power and centrality: A family of measures. *American journal of sociology* 92, 5 (1987), 1170–1182.

[5] Alessandro Bozzon, Marco Brambilla, Stefano Ceri, Matteo Silvestri, and Giuliano Vesci. 2013. Choosing the Right Crowd: Expert Finding in Social Networks. In *the 16th International Conference on Extending Database Technology*. 637–648. https://doi.org/10.1145/2452376.2452451

[6] M Keith Chen, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen. 2017. *The value of flexible work: Evidence from uber drivers*. Technical Report. National Bureau of Economic Research.

[7] Cody Cook, Rebecca Diamond, Jonathan Hall, John A List, and Paul Oyer. 2018. *The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers*. Technical Report. National Bureau of Economic Research.

[8] Kinda El Maarry, Wolf-Tilo Balke, Hyunsouk Cho, Seung-won Hwang, and Yukino Baba. 2014. Skill ontology-based model for quality assurance in crowdsourcing. In *International Conference on Database Systems for Advanced Applications*. 376–387. https://doi.org/10.1007/978-3-319-05810-8 arXiv:9780201398298

[9] Ju Fan, Guoliang Li, Beng Chin Ooi, Kian-lee Tan, and Jianhua Feng. 2015. iCrowd: An Adaptive Crowdsourcing Framework. In *the 2015 ACM SIGMOD International Conference on Management of Data*. 1015–1030. https://doi.org/10.1145/2723372.2750550

[10] Ronald Aylmer Fisher and Frank Yates. 1938. *Statistical tables for biological, agricultural and medical research*. Edinburgh: Oliver and Boyd. 146 pages.

[11] Adi Gaskell. 2018. The Demographics Of The Gig Economy. (2018).

[12] Diane F Halpern and Milton D Hakel. 2003. Applying the science of learning to the university and beyond: Teaching for long-term retention and transfer. *Change: The Magazine of Higher Learning* 35, 4 (2003), 36–41.

[13] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1914–1933.

[14] C. A. Hidalgo, B. Klinger, A.-L. Barabási, and R. Hausmann. 2007. The Product Space Conditions the Development of Nations. *Science* 317, 5837 (2007), 482–487. https://doi.org/10.1002/em.20028

[15] Chien-Ju Ho and Jennifer Wortman Vaughan. 2012. Online Task Assignment in Crowdsourcing Markets. In *the Twenty-Sixth AAAI Conference on Artificial Intelligence*. 45–51. arXiv:1508.03593

[16] Keman Huang, Jinhui Yao, Jia Zhang, and Zhiyong Feng. 2016. Human-as-a-Service : An Empirical Study of Growth in Human Service Ecosystem. In *IEEE SCC*. 90–97. https://doi.org/10.1109/SCC.2016.19

[17] Ayush Jain, Akash Das Sarma, Aditya Parameswaran, and Jennifer Widom. 2017. Understanding workers, developing effective tasks, and enhancing marketplace dynamics: a study of a large crowdsourcing marketplace. *Proceedings of the VLDB Endowment* 10, 7 (2017), 829–840.

[18] Lawrence F Katz and Alan B Krueger. 2016. *The rise and nature of alternative work arrangements in the United States, 1995-2015*. Technical Report. National Bureau of Economic Research.

[19] Qing Ke. 2017. Service providers of the sharing economy: who joins and who benefits? *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 57.

[20] Marios Kokkodis and Panos Ipeirotis. 2014. The Utility of Skills in Online Labor Markets. In *Thirty Fifth International Conference on Information Systems*. 1–18.

[21] Marios Kokkodis and Panagiotis G. Ipeirotis. 2015. Reputation Transferability in Online Labor Markets. *Management Science* 1909, October 2017 (2015), 1–20. https://doi.org/10.1287/mnsc.2015.2217

[22] Farshad Kooti, Mihajlo Grbovic, Luca Maria Aiello, Nemanja Djuric, Vladan Radosavljevic, and Kristina Lerman. 2017. Analyzing Uber's ride-sharing economy. In *Proceedings of the 26th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, 574–582.

[23] James Manyika, Susan Lund, Jacques Bughin, Kelsey Robinson, Jan Mischke, and Deepa Mahajan. 2016. Independent work: Choice, necessity, and the gig economy. *McKinsey Global Institute* 2016 (2016), 1–16.

[24] Panagiotis Mavridis, David Gross-Amblard, and Zoltán Miklós. 2016. Using Hierarchical Skills for Optimized Task Assignment in Knowledge-Intensive Crowdsourcing. In *WWW'2016*. 843–853. https://doi.org/10.1145/2872427.2883070

[25] TJ McCue. 2018. 57 Million U.S. Workers Are Part Of The Gig Economy. (2018).

[26] Chris Nosko and Steven Tadelis. 2015. *The limits of reputation in platform markets: An empirical analysis and field experiment*. Technical Report. National Bureau of Economic Research.

[27] Habibur Rahman, Saravanan Thirumuruganathan, and Senjuti Basu Roy. 2015. Worker Skill Estimation in Team-Based Tasks. In *Proceedings of the VLDB Endowment*. 1142–1153. https://doi.org/10.14778/2809974.2809977

[28] Donna Vakharia and Matthew Lease. 2015. Beyond Mechanical Turk: An analysis of paid crowd work platforms. *Proceedings of the iConference* (2015), 1–17.

[29] Stanley Wasserman and Katherine Faust. 1994. *Social network analysis: Methods and applications*. Vol. 8. Cambridge university press.

[30] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world' networks. *nature* 393, 6684 (1998), 440.

[31] Yudian Zheng, Guoliang Li, and Reynold Cheng. 2016. DOCS : A Domain-Aware Crowdsourcing System Using Knowledge Bases. In *Proceedings of the VLDB Endowment*, Vol. 10. 361–372.

[32] Yudian Zheng, Jiannan Wang, Guoliang Li, Reynold Cheng, and Jianhua Feng. 2015. QASCA: A Quality-Aware task assignment system for crowdsourcing applications. In *the ACM SIGMOD International Conference on Management of Data*. 1031–1046. https://doi.org/10.1145/2723372.2749430

Received April 2019; revised July 2019; accepted August 2019