



Talent Flow Analytics in Online Professional Network

Richard J. Oentaryo¹ · Ee-Peng Lim² · Xavier Jayaraj Siddarth Ashok² · Philips Kokoh Prasetyo² · Koon Han Ong³ · Zi Quan Lau³

Received: 16 January 2018 / Revised: 8 June 2018 / Accepted: 22 July 2018 / Published online: 12 August 2018
© The author(s) 2018

Abstract

Analyzing job hopping behavior is important for understanding job preference and career progression of working individuals. When analyzed at the workforce population level, job hop analysis helps to gain insights of talent flow among different jobs and organizations. Traditionally, surveys are conducted on job seekers and employers to study job hop behavior. Beyond surveys, job hop behavior can also be studied in a highly scalable and timely manner using a data-driven approach in response to fast-changing job landscape. Fortunately, the advent of online professional networks (OPNs) has made it possible to perform a large-scale analysis of talent flow. In this paper, we present a new data analytics framework to analyze the talent flow patterns of close to 1 million working professionals from three different countries/regions using their publicly accessible profiles in an established OPN. As OPN data are originally generated for professional networking applications, our proposed framework repurposes the same data for a different analytics task. Prior to performing job hop analysis, we devise a job title normalization procedure to mitigate the amount of noise in the OPN data. We then devise several metrics to measure the amount of work experience required to take up a job, to determine that the duration of a job's existence (also known as the job age), and the correlation between the above metric and propensity of hopping. We also study how job hop behavior is related to job promotion/demotion. Lastly, we perform connectivity analysis at job and organization levels to derive insights on talent flow as well as job and organizational competitiveness.

Keywords Talent flow · Job hop · Network analysis · Centrality

K. H. Ong and Z. Q. Lau: Work done during internship at Living Analytics Research Centre.

✉ Philips Kokoh Prasetyo
pprasetyo@smu.edu.sg

Richard J. Oentaryo
richard.oentaryo@mclaren.com

Ee-Peng Lim
eplim@smu.edu.sg

Xavier Jayaraj Siddarth Ashok
xaviera@smu.edu.sg

Koon Han Ong
koonhan.ong.2013@accountancy.smu.edu.sg

Zi Quan Lau
ziquan.lau.2015@sis.smu.edu.sg

¹ McLaren Applied Technologies, Singapore, Singapore

² Living Analytics Research Centre, Singapore Management University, Singapore, Singapore

³ Singapore Management University, Singapore, Singapore

1 Introduction

Job hop is a common behavior observed in any workforce. As a person hops from jobs to jobs, he or she acquires new skills and potentially gains higher income. Every job hop captures an important decision made by the person as well as an attempt of the hiring organization to acquire talent. When job hop behavior is analyzed at the workforce level, it will yield insights about the workforce, job pool and employers.

Such insights have been traditionally obtained using surveys on employers and job seekers. For example, the US Bureau of Labor Statistics (BLS) conducts annual surveys with approximately 146,000 businesses and government agencies to collect employment data.¹ The surveys yield useful information about job demand, job supply, income, working hours, etc. While surveys can be a powerful instrument to gather direct user input, they are usually not scalable. In the case of the BLS surveys, they cover less

¹ www.bls.gov/ces/.

than 1% of all US businesses. Moreover, as fast-changing technologies (such as sharing economy [8]) begin to impact job demand quickly, it is critical to explore new ways to obtain job-related insights.

Past studies [10, 13, 16] also tend to study jobs and organizations in isolation, without considering them as connected networks and how the networks capture talent flows from jobs to other jobs, and from organizations to other organizations. A lack of this network view prevents us from analyzing the ways people build their career, and competition among organizations for talent. For example, some job changes could be promotions, while others could just be lateral and even demotions. The network view is also crucial in studying how the competitions among jobs and organizations would eventually impact job creation and talent attraction.

In contrast, online professional networks (OPNs) are fast becoming a marketplace for resume posting, candidate hunting, and job searching. Representative examples of OPN are LinkedIn, Xing, and Viadeo.² Detailed job activity data at the individual user level are now publicly available in these OPNs, as soon as the users update their profiles. These data can be analyzed to derive interesting behavioral insights about jobs and organizations, as well as to build services that can benefit both employers and job seekers, e.g., a service that helps employers find suitable employees and another service that helps job seekers find suitable jobs.

Objectives In this work, we focus on using data from one of the world's largest OPNs to analyze job hops and talent flow. To support our analysis on hops within an organization and across organizations, we first develop several metrics that measure the amount of experience required for a job and how established/recent a job is, from the viewpoint of the people holding the job.

We also aim at studying how the job hop behavior of a workforce is related to job promotion/demotion. This is a topic often discussed based on anecdotal examples [1, 9]. A better approach is to conduct a large-scale data science study. This will give much broader insights on job hop patterns, particularly useful in human resource recruitment and career coaching.

Finally, our research aims at analyzing talent flow based on job hop behavior and measuring the capabilities of each job and organization in attracting, supplying, and competing for human capital. To this end, we create a weighted directed hop network among jobs and organizations, develop different centrality measures for the job and organization nodes, and evaluate them by manual

inspection or by comparing with other attributes such as organization size.

Contributions To accomplish the above objectives, we develop a new data analytics framework to clean, aggregate, and derive talent flow insights from OPN data. The proposed framework constitutes a generalization of our earlier work [18], featuring two major extensions:

- Prior to talent flow analytics, we introduce job title translation, parsing, and normalization steps as additional data cleaning/preprocessing steps in order to mitigate noise in our OPN datasets.
- To demonstrate the applicability of our proposed framework, we conduct an extended study using OPN datasets of working professionals from three countries/regions (i.e., Singapore, Switzerland, and Hong Kong) with diverse workforce profiles. Our study reveals a number of interesting findings and insights that unveil similarities and differences among countries/regions.

All in all, the main contributions of this work are as follows:

- We present a talent flow analytics framework to facilitate a data science approach to analyze talent flow among jobs and organizations. This outlines the essential steps to repurpose the OPN data for talent flow study that complements the traditional surveys.
- We devise several key metrics to analyze talent flow networks, with the aim to answer several research questions connecting talent flow with career progression, user attributes (e.g., working experience), and user career behavior (e.g., promotion and demotion). We conduct empirical studies using the key metrics to explain some patterns in our datasets.
- We justify the applicability of our approach through extensive empirical study using OPN datasets from three different countries/regions. The results reveal interesting insights on the similarities and differences of the talent flow patterns across countries/regions.

Paper outline The remainder of this paper is organized as follows. Section 2 first provides a survey of related works. Section 3 in turn describes the proposed talent flow analytics framework and the dataset used in our study. Details of the talent flow network construction and talent analytics approaches are described in Sects. 4 and 5, respectively. Section 6 presents the key insights and discussion on the results. Finally, Sect. 7 concludes this paper.

² LinkedIn—www.linkedin.com; Xing—www.xing.com; Viadeo—www.viadeo.com.

2 Related Work

Research on job and workforce movements has been around for decades [6, 10, 15, 17, 21]. Topel et al. [21] analyzed 15 years of job changing and wage growth of young men from longitudinal employee–employer data. Long et al. [15] studied the labor mobility in Europe and the USA. Moscarini et al. [17] measured worker mobility across occupations and jobs in the monthly Current Population Survey data from 1979 to 2006. More recent survey-based studies [6, 10, 19] have revealed that the younger employees are more likely to switch jobs and employers/companies than the older ones. Friedell et al. [5] showed that there is a discrepancy between younger generations' expectations in the workplace and older generations' perception of those expectations.

In general, these studies traditionally relied on surveys, census, and other data such as tax lists and population registers, which require extensive and time-consuming efforts to collect. Often times, surveys may focus on selected workforce segments or industries. Such an approach thus cannot be easily scaled up or replicated across many segments/industries.

With the wide adoption of OPNs, there is a rapidly growing interest to mine the online user data from the OPNs to understand job and workforce movements as well as career growth. For example, State et al. [20] analyzed the migration trends of professional workers into the USA. Xu et al. [22] combined work experiences from OPNs and check-in records from location-based social networks to predict job change occasions. Chaudhury et al. [3] analyzed the growth patterns of the ego network of new employees in companies.

An important aspect in OPNs is job hop. Job hop data capture a wide range of signals that can help understand the performances of organizations, talent sources, job market, professional profiles, as well as career advancement. Cheng et al. [4] modeled job hop activities to rank influential companies. Xu et al. [23] generated and analyzed job hop networks to identify talent circles. Kapur et al. [11] devised a talent flow graph to rank universities based on the career outcomes of their graduates. They applied their approach to two specific workforce segments: investment banker and software developer.

Users' career paths have also been utilized to model professional similarity for use in job recruitment process [24]. In this work, a sequence alignment method was used to quantify similarity between two career paths. Liu et al. [14] devised a multi-source learning framework that combines information from multiple social networks to predict the career path of a user. While the approach is interesting, their work focused only on four job categories, namely

software engineer, sales, consultant, and marketing. Recently, Li et al. [13] proposed a survival analysis approach to model career paths for turnover and career progression. However, their study was conducted on within-organization career paths (i.e., inside a company) for talent management.

Additionally, career trajectory similarity has been proposed to identify individuals who share similar career histories with some given user-provided ideal candidates so that the former can be returned as talent search results [7, 24]. Xu et al. [23] defined a job transition network with vertices and edges representing organization and talent flow between two organizations for a time period, respectively. From the network, talent circles each covering a set of organizations with similar talent exchange patterns are detected. It has been shown that talent circles can improve talent recruitment and job search.

On a related track, Xu et al. [22] analyzed job change patterns using OPN data and correlated these patterns with human activity data from a location-based social networking site. They also proposed a set of features to predict future job changes to be made by some employee.

Our research The work presented in this paper differs from the above-mentioned works in several unique ways. Firstly, we introduce quantitative metrics to measure how much work experience is required to take up a job and how recent/established a job is, and examine their relationships with the propensity of hopping. Secondly, we compute the level gain of job hops so as to analyze promotion/demotion of employees which, to our best knowledge, has been missing in the previous studies. Additionally, we perform an extensive study on talent flow and competition by analyzing both job-level and organization-level hop networks, without being restricted to specific workforce segments or industries. Last but not least, our study involves multiple countries or regions, thus providing more comprehensive insights on how talent flows compare among different countries or regions.

3 Dataset and Method

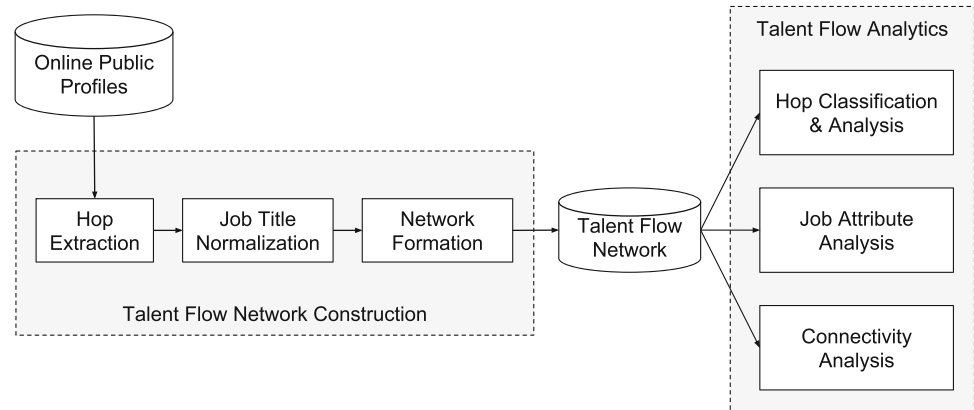
This section provides an overview of the OPN dataset considered in this work as well as of our proposed approach for talent flow analytics.

3.1 Dataset

To facilitate our empirical studies, we extract online public profiles from one of the world's largest OPNs. In particular, our data collection involves extracting the public profiles of all OPN users in three countries/regions, i.e., Singapore,

Table 1 Statistics of the OPN datasets used in this study

Statistics	Singapore	Switzerland	Hong Kong
Number of user profiles	1,674,432	1,419,129	432,525
Number of core user profiles	502,123	377,590	82,672
Number of organizations	151,638	124,462	45,652


Fig. 1 Overview of the proposed talent flow analytics framework

Switzerland, and Hong Kong. We collect a list of public profiles from public users directory associated with the target city or region. We also extract organization profiles mentioned in these public user profiles. As given in Table 1, the datasets consist of 1.6 and 1.4 M user profiles involving about 151 and 124 K organizations for Singapore and Switzerland, respectively. The Hong Kong dataset is smaller with 432 K user profiles and 45 K organizations.

Meanwhile, the core users in Table 1 refer to users who have at least one entry in the education, experience, and skill fields. As per Table 1, we have 502 K core users in Singapore, 377 K core users in Switzerland, and 82 K core users in Hong Kong. For some analyses that we perform in this work, such as on work experience and job level (see Sect. 5.2), these fields need to be available. In these cases, the relevant metrics are computed based on the core users only, and non-core users are excluded. For all other analyses such as hop extraction, job title normalization, and talent flow network construction (cf. Sect. 3.2), we include all users that are found in our data.

3.2 Talent Flow Analytics Framework

Our proposed talent flow analytics framework, as depicted in Fig. 1, consists of two key phases: *talent flow network construction* and *talent flow analytics*. In the first phase, we construct a talent flow network from the online public profiles. This comprises three steps: *hop extraction*, *job title normalization*, and *network formation*. In the hop extraction step, we collect all job hops from online public profiles. During job title normalization, we reduce the

**Fig. 2** Definition of job hop

duplicate job titles which occur due to variations in language, small typos, and non-standardized writing. For instance, “finance manager” is the same as “manager, finance,” “manager–finance,” “finance manager,” and “finance manger.” (Note that the last two finance manager variations are due to typos).

After job title normalization, we represent each job by its normalized job title and the company. We then craft the talent flow network based on transitions between jobs and perform talent flow analytics that encompasses three types of analysis, namely hop classification and analysis, job attribute analysis, and connectivity analysis. Each analysis studies a specific aspect of talent flow network. Firstly, hop classification and analysis focuses on analyzing types of job hop activity in the network whether the hop is an internal or external hop. Hop analysis will reveal the pattern of internal and external hops. Secondly, job attribute analysis aims to analyze talent flow with respect to job attributes and job hop attributes such as promotion and demotion. Finally, connectivity analysis strives to analyze talent flow behavior at the network level allowing us to

Table 2 Statistics of the extracted job hops

Statistics	Singapore	Switzerland	Hong Kong
Number of hops	4,561,881	4,027,083	917,617
Number of distinct job titles	795,249	1,094,061	195,659
Number of distinct job titles ($\text{min_sup} \geq 10$)	38,966	38,220	8982

determine important jobs and organizations. We describe the talent flow network construction and analytics phases in greater detail in Sects. 4 and 5, respectively.

4 Talent Flow Network Construction

In this section, we describe in greater detail the steps involved in the construction of a talent flow network (as per Fig. 1).

4.1 Hop Extraction

The first step to construct the talent flow network is to extract *job hop* from the online public profiles. Each *job* in talent flow network is defined as a triplet (t, c, i) , which represents job title t at organization c in industry i , respectively. Note that each organization c belongs to a unique industry i . We then define a *job hop* as a transition from one job to another job with *non-overlapping* time period. A job hop represents a talent flow from one job (or organization) to another job (or organization), and a collection of job hops forms a talent flow network (see Sect. 4.3).

Figure 2 shows an example of an OPN user who lists five jobs A, B, C, D , and E in his profile. In this case, the user is regarded as having only three hops, i.e., from job A to job B , from B to E , and from C to D . There is no hop from B to C , or from B to D , or from D to E , since they take place in an overlapping time period and are likely to be side activities of the user. To capture as many distinct job titles as possible for the next step (i.e., job title normalization), we include all users for this step, not only the core users.

Table 2 summarizes the statistics of the extracted job hops showing the number of job hops, and distinct job titles³ for Singapore, Switzerland, and Hong Kong datasets. From the extracted job hops, we collect more than 795,000 distinct job titles in Singapore, 1 million distinct job titles in Switzerland, and 195,000 distinct job titles in Hong

³ Distinct job titles refer to unique string of job titles. For example, “Java Developer” and “Java Software Engineer” are considered as different job titles, even though semantically they might correspond to the same job.

Kong. We discovered that job titles in Switzerland consist of many languages such as English, French, German, Italian, Spanish, and Portuguese. We also found that the majority of job titles in Hong Kong are in English and Mandarin.

Although we have a large numbers of distinct job titles, the job titles are typically very noisy, owing to two main reasons:

- **Data sparsity** This could be due to either poor/inaccurate naming of job titles or less popular jobs that have small number of occurrences. The first issue can be resolved by job title parsing and normalization, which we will describe in Sect. 4.2. In contrast, the latter issue cannot be resolved by job title normalization. To mitigate this effect, we define a threshold for each job title in the extracted hop collections to be at least 10 instances.
- **Job title variation** This variation is typically caused by language diversity and non-standardized job title naming. This can be addressed by job title parsing and normalization as well as job title translation for non-English job titles. This is further elaborated in Sect. 4.2.

4.2 Job Title Normalization

We develop a parser to normalize job title on the extracted hop collection. Job title normalization is important to reduce variations of same job titles such as “research director” and “director of research.” For a given job title, the parser normalizes the title into its constituent parts, allowing us to extract important *functional* components inside a given job title. We define the constituent parts of a given job title as follows:

1. *Primary function* indicates the main job role. Each job title must have at least one primary function. Thus, this part is compulsory.
2. *Domain* indicates the domain of a job role. This part is optional.
3. *Position* indicates the seniority level of a job role. This part is optional.
4. *Secondary function* indicates secondary job role. This part is optional.

Senior	Business	Development	Manager	(ASEAN)
Position	Domain		Primary Function	Additional Information
Assistant	Director	of	Research	
Position	Primary Function		Domain	

Fig. 3 Examples of parsed job titles

5. *Additional information* indicates extra information about the job title. It commonly appears inside a bracket. This part is optional.

To build such a parser, we devise grammar rules to cover valid job title syntax. The entire parsing process can be broken down into several steps:

1. **Lexical analysis** In this step, a lexical analyzer (lexer) tokenizes a job title input using defined regular expressions and matches them against dictionary files.
2. **Syntax tree generation** Using the tokens from lexer, the parser subsequently checks if these tokens adhere to the grammar rules and validates the syntax. If the syntax is valid, the concrete syntax tree will be generated from the selected rules.
3. **Extraction** In this step, the constituent parts of the job title are extracted from the concrete syntax tree.

We implement the parser using PLY⁴ parser tool. Few examples of the parsed job titles are shown in Fig. 3. Job titles that fail the lexical analysis and syntax tree generation steps are considered to have parsing errors.

Before parsing the job titles, some efforts are required to handle the non-English job titles in Switzerland and Hong Kong datasets. We use the Google Translate API⁵ to translate non-English job titles to English. The translation results are then validated by our job title parser. In this case, we expect the valid translated job titles to be parseable.

Figure 4 depicts the parsing error rates and the number of parsed job titles at different job title minimum support values on each dataset. In particular, Fig. 4a, c, e shows the number of distinct job titles and job parsing error rates at different minimum supports, while Fig. 4b, d, f shows the corresponding numbers of parsed job titles. We can see that the lower the minimum support, the higher the number of distinct job titles, but the parsing error rates increase. At the minimum support of 10, the parsing error reaches 15.75% for Singapore, 28.87% for Switzerland, and 17.44% for Hong Kong datasets. Following these processes, we obtain

⁴ <http://www.dabeaz.com/ply/>.

⁵ We use the googletans Python wrapper for Google Translate API—<https://pypi.python.org/pypi/googletans>.

32,828 parsed job titles in Singapore, 27,184 in Switzerland, and 7,415 in Hong Kong.

With job title parsing, it is expected that different job titles with the same constituent parts would map to the same parsed result. Among these job titles, we pick the most popular one as the normalized (i.e., canonical) job title and use it to substitute all the other job titles with the same constituent parts. Table 3 presents the statistics of the normalized job titles. As given in the table, the Singapore and Hong Kong datasets have 11.4 and 10.8% duplicate job titles, respectively. On the other hand, the Switzerland dataset has very high duplicate job titles, 34.10%. This may be attributed to the multilingual nature of the Switzerland data, containing many languages and different ways of naming job titles. After normalization, we finally have 29,084, 17,913, and 6614 normalized job titles in Singapore, Switzerland, and Hong Kong, respectively.

4.3 Network Construction

We use the extracted job hops and the normalized job titles to form our talent flow network. A talent flow network is a directed graph where each edge represents a job hop activity. Based on node type in the network, talent flow network can be classified into two types: (1) **job network**, where each node $v_{t,i}$ represents a canonical job title t in industry i , and (2) **organization network**, where each node v_c represents an organization c . Job network allows us to observe talent flow at job level, whereas organization network allows us to observe talent flow at organization level.

For the job network, a directed edge from node $v_{t,i}$ to node $v_{t',i'}$ represents a job hop activity from a node (t, i) to another node (t', i') . We also capture the number of user profiles moving from (t, i) to (t', i') as the **edge weight** $e_{(t,i) \rightarrow (t',i')}$. The same applies to the organization network, i.e., the edge weight $e_{c \rightarrow c'}$ represents the number of users moving from an organization c to another organization c' .

5 Talent Flow Analytics

This section elaborates the three types of talent flow analytics as shown in Fig. 1, namely: (a) hop classification and analysis, (b) job attribute analysis, and (c) connectivity analysis.

5.1 Hop Classification and Analysis

The hop classification and analysis essentially involve two types of job hop:

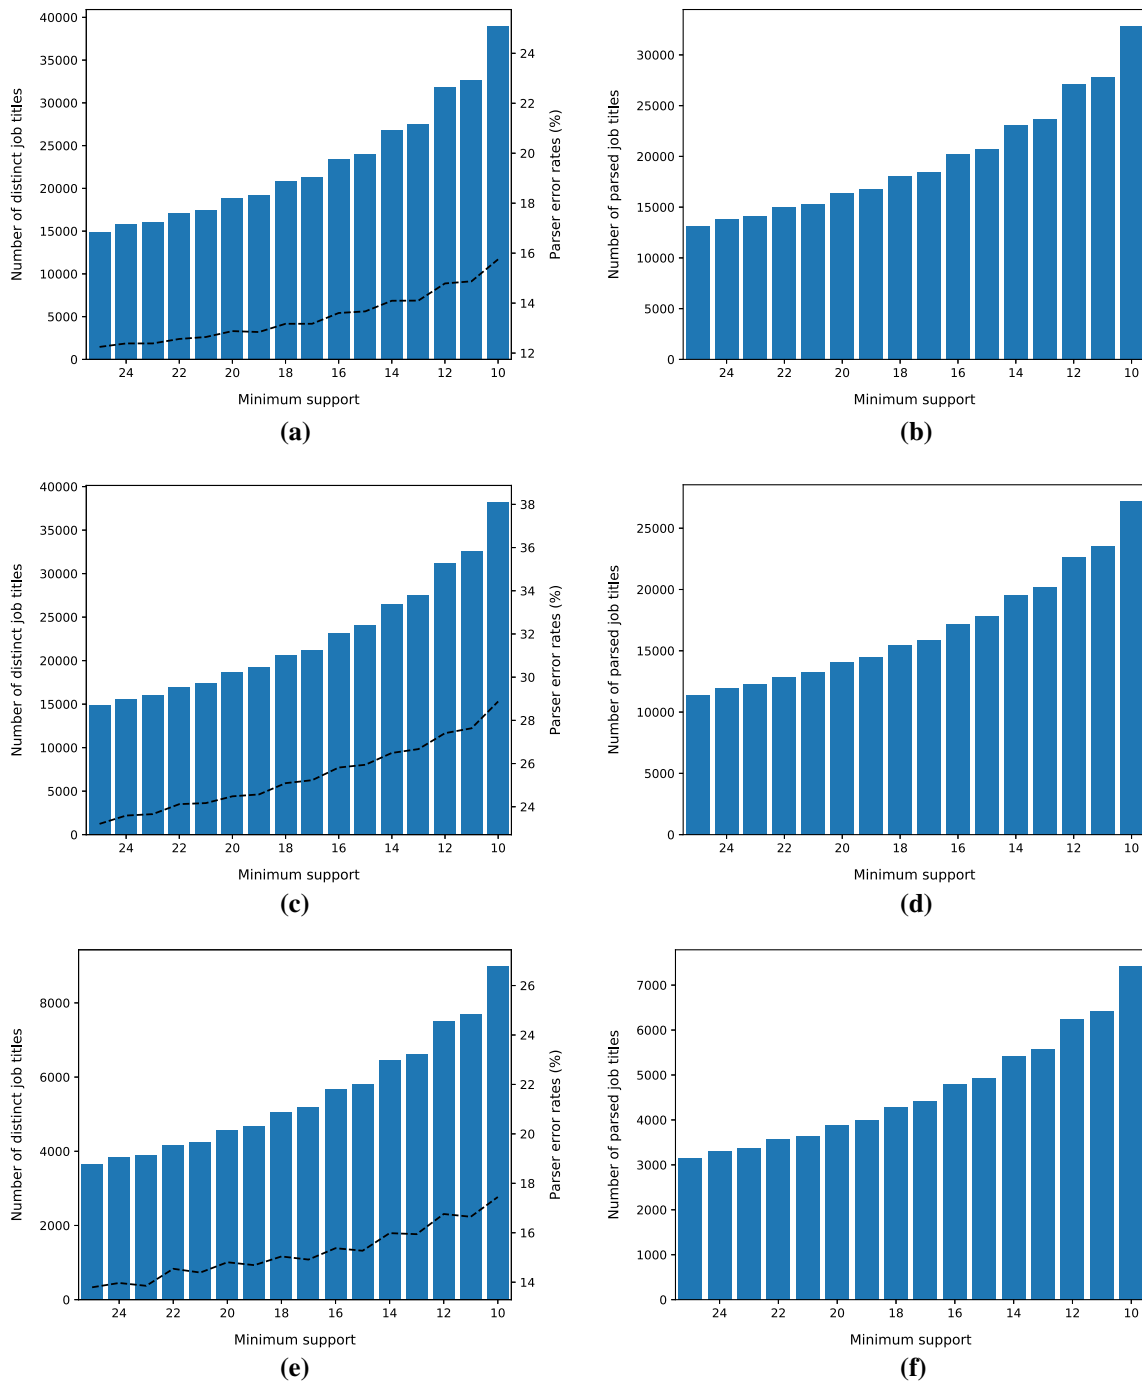


Fig. 4 Results of job title parsing. **a** Parsing error rates for Singapore data. **b** Parsed job titles for Singapore data. **c** Parsing error rates for Switzerland data. **d** Parsed job titles for Switzerland data. **e** Parsing error rates for Hong Kong data. **f** Parsed job titles for Hong Kong data

- External hop** This hop refers to transition from one job to another job, where the source and destination companies are *different*. That is, an external hop is a hop from job $j = (t, c, i)$ to job $j' = (t', c', i')$, where $c \neq c'$. By this definition, the origin job title t need not be the same as the destination title t' . Intuitively, two
 - jobs with the same title but at different companies should be treated as separate jobs.
- Internal hop** This is transition from one job to another, where the source and destination companies are *the same*, i.e., an internal hop is a hop from job $j = (t, c, i)$ to job $j' = (t', c', i')$, where $c = c'$ and $t \neq t'$. The latter

Table 3 Statistics of normalized job titles ($\text{min_sup} \geq 10$)

Statistics	Singapore	Switzerland	Hong Kong
Number of parsed job titles	32,828	27,184	7415
Number of duplicate parsed job titles	3744 (11.4%)	9271 (34.1%)	801 (10.8%)
Number of normalized job titles	29,084	17,913	6614

constraint ($t_j \neq t_{j'}$) is meant to avoid job title duplicates under the same company (e.g., a person may state three times that he/she is a civil engineer at company X, for he/she has worked on three construction projects under the same company). As such, we do *not* count a move from job j to j' where $t = t'$ and $c = c'$ as a (valid) internal hop.

5.2 Job Attribute Analysis

In our study, we want to tell how much career advancement people make in their jobs. We therefore first need to estimate the experience of a person holding a job. Secondly, to determine changes of job market over time, we need to estimate how long a job has existed. To fulfill the two goals, we introduce several key metrics that are applied to the core users, i.e., those with at least one entry in the education and skill fields. With the two fields, one can derive interesting attributes of jobs and the skill profiles of the three economies. We describe the key metrics in turn below.

- **Work experience** This refers to the duration since the graduation date of the most recent educational degree of a person till the time at which he/she finishes a particular job. For a person p with job title t at organization c , the work experience is:

$$wk_exp(p, t, c, i) = end_time(p, t, c, i) - grad_date(p) \quad (1)$$

where $grad_date(p)$ denotes the last graduation date as mentioned in his/her account profile. In our subsequent analyses, note that we consider only positive work experience, i.e., we exclude cases whereby $wk_exp(p, t, c, i) \leq 0$. This is due to the observation that most jobs taken prior to the last education in our data are typically of interim nature (such as internship), which may introduce bias in our analysis at higher (e.g., industry) level. Next, for a given job title t in industry i , the average (i.e., expected) work experience of the job title–industry pair (t, i) is given by:

Table 4 Hop classification statistics for Singapore dataset

	Promotion	Demotion	Total
External hop	2627	930	3557
Internal hop	2182	200	2382
Total	4809	1130	5939

Table 5 Hop classification statistics for Switzerland dataset

	Promotion	Demotion	Total
External hop	859	533	1392
Internal hop	1486	341	1827
Total	2345	874	3219

$$avg_wk_exp(t, i) = \frac{1}{|S_{t,i}|} \sum_{(p,t,c,i) \in S_{t,i}} wk_exp(p, t, c, i) \quad (2)$$

where $S_{t,i}$ is the set of (unique) person–job pairs having job title t in industry i . Examples of job title with high avg_wk_exp score across industries in our data are “Professor,” “Managing Director,” and “CEO,” whereas examples with low avg_wk_exp score are “Intern” and “Teaching Assistant.”

- **Job age** This is the duration from the start of a given job until the current date $curr_date$. It measures how recent or established a job is from the perspective of a person holding the job. For a person p with job title t at organization c , the job age is defined as:

$$job_age(p, t, c, i) = curr_date - start_date(p, t, c, i) \quad (3)$$

where $start_date(p, t, c, i)$ refers to the start date of the person p 's job title t at organization c of industry i . For a given job title t from industry i , the average (expected) age of the (job title, industry) pair (t, i) is therefore:

$$avg_job_age(t, i) = \frac{1}{|S_{t,i}|} \sum_{(p,t,c,i) \in S_{t,i}} job_age(p, t, c, i) \quad (4)$$

Table 6 Hop classification statistics for Hong Kong dataset

	Promotion	Demotion	Total
External hop	175	56	231
Internal hop	350	36	386
Total	525	92	617

Examples of job titles with high *avg_job_age* score across industries in our data are “Director,” “Systems Engineer,” and “Division Manager,” while examples with low score are “Data Scientist” and “Media Analyst.”

Based on the above metrics, we further derive several higher-level metrics, by aggregating over user profiles at either the job or organization level:

- **External hop fraction** The fraction of people who move out from an organization c to a different organization $c' \neq c$ over the (total) people hopping from organization c . Formally, for a given group of users g (e.g., work experience, job age, or skill count group), job title translation and parsing the external hop fraction is:

$$\%external_hop(g) = \frac{|\mathbf{P}_{c \rightarrow c'}^g|}{|\mathbf{P}_{c \rightarrow c'}^g| + |\mathbf{P}_{c \rightarrow c}^g|} \quad (5)$$

where $\mathbf{P}_{c \rightarrow c'}$ is the set of all user profiles belonging to group g who perform *external hops* from some arbitrary organizations c to *different* organizations $c' \neq c$. Conversely, $\mathbf{P}_{c \rightarrow c}$ is the set of user profiles belonging to group g who perform *internal hop* within the *same* organization c .

- **Job level** As different organizations offer jobs of different rewards and seniority levels (even for the

same job titles), we want to be able to measure them. Since our data do not carry any salary information, we estimate the seniority level of a job (t, c) by computing the *average work experience* over all users who mention job title t at organization c in their profiles:

$$job_level(t, c) = \frac{1}{|\mathbf{P}_{t,c}|} \sum_{p \in \mathbf{P}_{t,c}} wk_exp(p, t, c, i) \quad (6)$$

where $\mathbf{P}_{t,c}$ is the set of all people who include job (t, c) in their profiles. In the equation, i can be inferred from c . Intuitively, a job with longer average work experience implies that a longer time is required to achieve that position, and hence we can expect it to be a high-level job (e.g., CEO of a multi-national organization).

- **Level gain** This refers to the difference between the levels of two jobs within the same or different companies. A positive level gain can be loosely interpreted as a “promotion,” whereas a negative level gain loosely implies a “demotion.” Here, the “promotion” (“demotion”) does not necessarily mean a monetary increase (decrease), but more of an increase (decrease) in the level of work experience required. Formally, the level gain for hop from job (t, c) to job (t', c') is given by:

$$level_gain((t, c), (t', c')) = job_level(t', c') - job_level(t, c) \quad (7)$$

We note that, although there is no ground truth available in our OPN data, our manual inspections show that the level gain provides a reasonable proxy for a promotion or demotion. It is also worth mentioning that no zero level gain (i.e., neither “promotion” nor “demotion”) is found in our data.

Table 7 Network statistics of the Singapore dataset

Metric	Job graph	Company graph
<i>Basic</i>		
Number of nodes	30,531	16,112
Number of edges	45,412	43,499
Sparsity of adjacency matrix	0.005%	0.017%
<i>Strongly connected component (SCC)</i>		
Number of SCCs	25,340	11,698
Size of the largest SCC	5028 (16.47%)	4367 (27.10%)
Size of the second largest SCC	11 (0.036%)	3 (0.018%)
<i>Weakly connected component (WCC)</i>		
Number of WCCs	3864	2187
Size of the largest WCC	22,201 (72.72%)	12,705 (78.85%)
Size of the second largest WCC	22 (0.072%)	6 (0.037%)

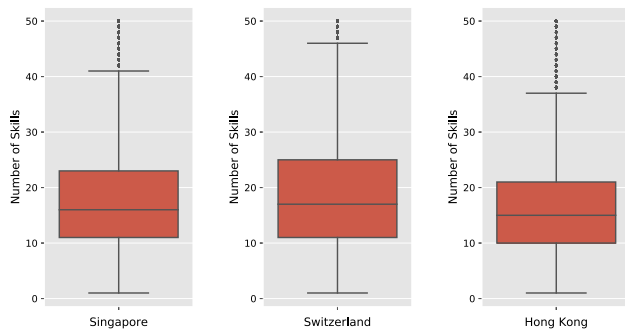


Fig. 5 Distribution of number of skills

5.3 Connectivity Analysis

To facilitate connectivity analysis, we utilize several network centrality metrics to measure node importance in both job and organization networks, as follows:

- In-degree centrality** This metric refers to the number of inbound (unweighted) edges for a node in the job or organization graph. The in-degree centrality can be interpreted as a measure of how *prominent* a job (or organization) is in a local sense—a high in-degree may imply that it attract talents from the immediate in-neighbors. For this metric, we do not take into account the edge weight information (i.e., the total number of incoming user profiles), as we want to minimize the support bias due to a large number of users for a given job (organization).
- Out-degree centrality** This is defined as the number of outbound (unweighted) edges for a node in the job or organization graph. We can use the out-degree centrality to measure how *influential* a job (or organization) is in a local sense—a high out-degree may be indicative of a talent supplier to the immediate out-neighbors. Again, we do not utilize the edge weight to compute this metric, so as to mitigate the support bias.
- PageRank centrality** This is a well known metric originally used to rank web pages [2]. PageRank views inbound edges as “votes,” and the key idea is that “votes” from important nodes should carry more weight than “votes” from less important nodes. In this work, we employ a *weighted* version of PageRank [12], whereby the transition probabilities for each (source) node is proportional to the (out-)edge weights divided by the weighted out-degree of the node. In the context of job and organization graphs, the weighted PageRank can be viewed as a measure of *global competitiveness*—a job or organization with high PageRank reflects a “desirable” destination point where the flow of talent is heading to. Here, we use edge weight, as the hop volume matters in determining where the flow goes

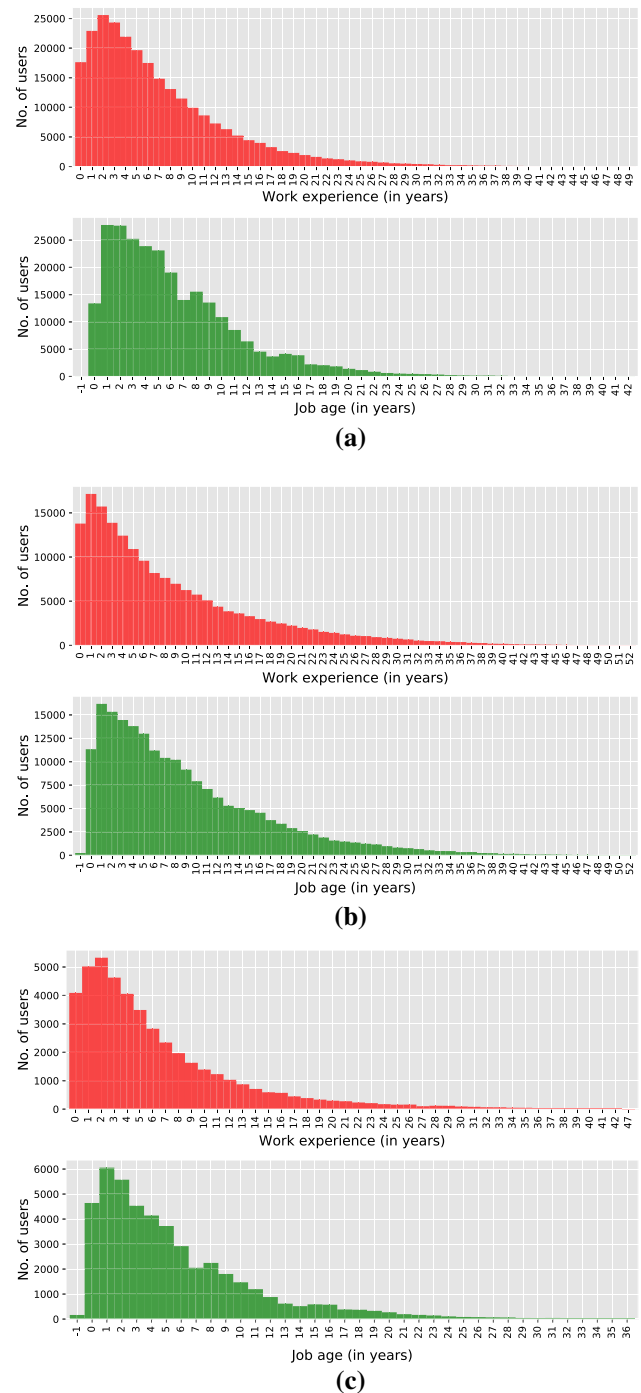


Fig. 6 Distributions of work experience and job age. **a** Singapore. **b** Switzerland. **c** Hong Kong

to. To avoid dead ends (i.e., nodes with zero out-degree), we allow our PageRank to perform random jump with the default “teleportation” probability of 0.15.

Table 8 Network statistics of the Switzerland dataset

Metric	Job graph	Company graph
<i>Basic</i>		
Number of nodes	5867	5704
Number of edges	8993	10,447
Sparsity of adjacency matrix	0.026%	0.032%
<i>Strongly connected component (SCC)</i>		
Number of SCCs	5139	4921
Size of the largest SCC	645 (10.99%)	765 (13.41%)
Size of the second largest SCC	16 (0.27%)	2 (0.035%)
<i>Weakly connected component (WCC)</i>		
Number of WCCs	1734	2296
Size of the largest WCC	2905 (49.51%)	2837 (49.74%)
Size of the second largest WCC	19 (0.32%)	5 (0.088%)

6 Insights and Discussion

Using the methodology and metrics described previously in Sects. 3–5, we present our empirical findings and analysis in this section. We begin our discussion with basic distribution analysis, followed by findings in each type of analysis within the talent flow analytics phase.

6.1 Distribution Analysis

We first analyze the distributions of several basic metrics, including skill count, work experience, and job age. We found that the *core user profiles* in our talent flow networks typically have about 10–25 skills. Figure 5 shows the box plots of the skill distribution. It is shown that Switzerland workforce tends to have (or rather declare) slightly more skills compared to Singapore and Hong Kong workforce. It also indicates that Hong Kong workforce has (or declares) the least number of skills. The maximum of 50 skills is due to the fact that our OPN imposes a maximum limit of 50 skills per user profile.

We also notice that most jobs in Singapore, Switzerland, and Hong Kong consist of young workforce, which has work experience of 5 years or less. This pattern holds for all the three datasets. Most users in our OPN data are relatively young in terms of work experience. This could be due to the younger users showing more interest in using OPN to conduct professional networking. On the other hand, there are only very few people who have worked for over 20 years. The most common work experience (i.e., the mode) is 2 years for Singapore and Hong Kong and 1 year for Switzerland.

Figure 6 presents the distributions of the work experience and job age across the three countries. The results




Fig. 7 Distribution of job level. **a** Singapore. **b** Switzerland. **c** Hong Kong

suggest that most jobs have been established for 1 year or more. On the other hand, only very few jobs have been established for more than 20 years. As with work experience, the most common job age is 1 year. The relatively young job age can be explained partly by the young user base, and partly by the sparsity of old but senior-level jobs. From the labor economics perspective, this suggests that attention need to be given to identifying and creating more senior jobs to support an aging workforce.

The job-level distribution shown in Fig. 7a reveals that most jobs in Singapore possess 4–6 years of experiences. Compared to jobs in Singapore, most jobs in Switzerland seem to have more senior jobs having 7–8 years of experiences (cf. Fig. 7b), and most jobs in Hong Kong seem to have more junior jobs with 3–4 years of experiences (cf. Fig. 7c). We also notice that the distribution of job level in Hong Kong is much shorter than that of Singapore and Switzerland and the maximum job level is 13 years. This is understandable nevertheless, considering the fact that Hong Kong began as a new special administration of China only after 1997.

We conduct further investigation by examining the job-level discrepancy of same job titles in Singapore vs. Switzerland as well as Singapore vs. Hong Kong. Figure 8 summarizes the results, providing an interesting insight that

Table 9 Network statistics of the Hong Kong dataset

Metric	Job graph	Company graph
<i>Basic</i>		
Number of nodes	1895	1935
Number of edges	2492	2982
Sparsity of adjacency matrix	0.069%	0.08%
<i>Strongly connected component (SCC)</i>		
Number of SCCs	1710	1762
Size of the largest SCC	151 (7.97%)	139 (7.18%)
Size of the second largest SCC	7 (0.37%)	14 (0.72%)
<i>Weakly connected component (WCC)</i>		
Number of WCCs	484	892
Size of the largest WCC	785 (41.42%)	763 (39.43%)
Size of the second largest WCC	32 (1.69%)	5 (0.26%)

explains the difference of job-level distribution in Fig. 7. We found that for the same job titles, jobs in Switzerland tend to have 2–3 longer years of experiences than jobs in Singapore, whereas jobs in Hong Kong tend to have equally long (or slightly shorter) years of experiences than jobs in Singapore.

Comparing Figs. 6 and 7, it can be observed that the distributions of job age and job level are not identical. This is expected, based on the definitions of the two metrics in Eqs. (4) and (6), respectively. In particular, job-level looks *forward* in time at how long an individual accumulates experience since his/her (last) graduation date, while job age looks *backward* in time at how long a job has been established until a current reference time point.

Last but not least, comparison of the job age discrepancy for same job titles in Singapore vs. Switzerland and Singapore vs. Hong Kong tells us that most jobs in Switzerland tend to have longer job ages than those in Singapore. This implies that the Switzerland workforce is ahead of the Singapore counterparts in terms of job establishment for the same job title. On the other hand, Hong Kong workforce tends to have lower job age than Singapore workforce, implying that the jobs in Hong Kong are less established than in Singapore. Figure 9 presents a summary of the job age comparison.

6.2 Hop Classification and Analysis

For this analysis, we started with an initial hypothesis that the propensity of external hop is potentially associated (correlated) with the work experience, job age, and number of skills. To test this, we conduct an investigation on how the external hop fraction [cf. Eq. (5)] varies with different

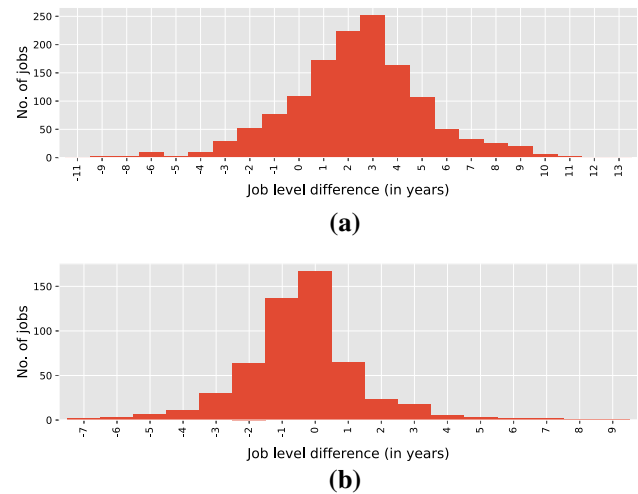


Fig. 8 Discrepancy of job level. **a** Singapore versus Switzerland. **b** Singapore versus Hong Kong

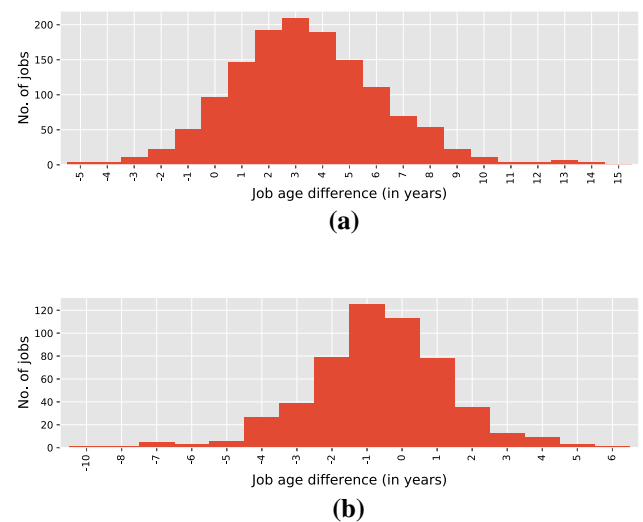


Fig. 9 Discrepancy of job age. **a** Singapore versus Switzerland. **b** Singapore versus Hong Kong

combinations of work experience, job age, and skill count groups.

Figure 10 shows the distribution of external hop fraction varying work experience, job age, and skill count group, whereby the minimum support was set to 100 for each bar in the plots. The figure reveals a number of key insights:

- External hops are generally very common in Singapore, Switzerland, and Hong Kong, regardless of work experiences, job age, and number of skills. The external hop fraction in Singapore, Switzerland, and Hong Kong

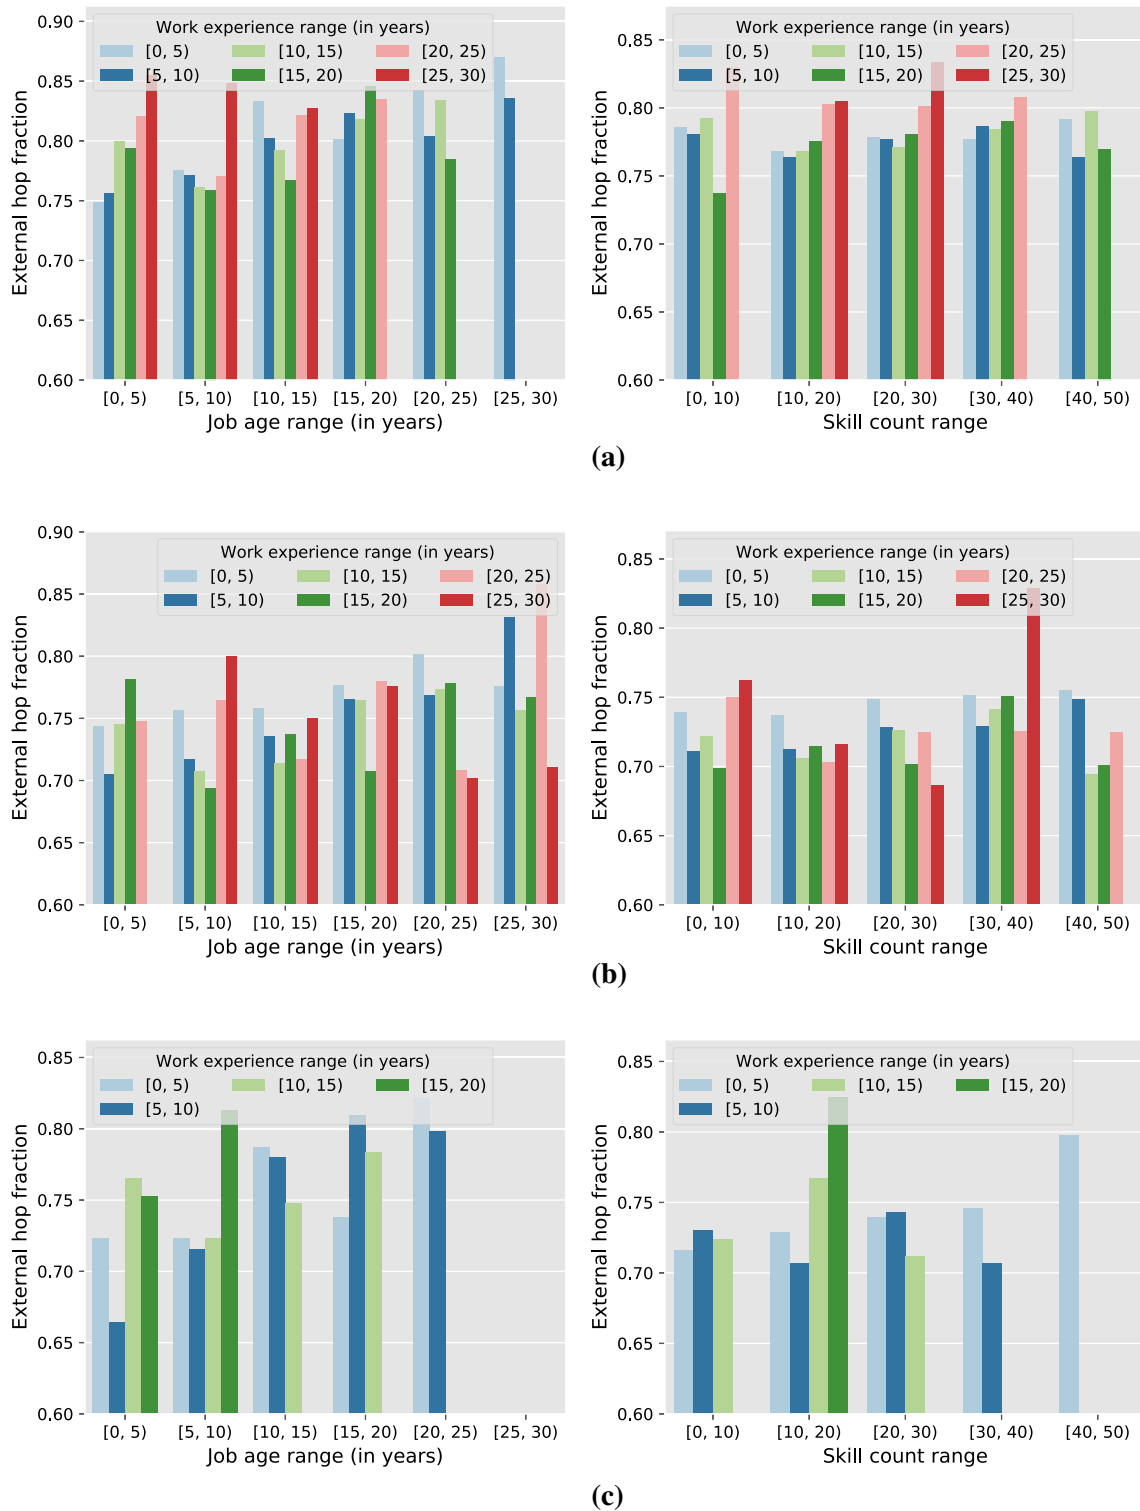


Fig. 10 Distributions of external hop fraction. **a** Singapore. **b** Switzerland. **c** Hong Kong

are generally greater than 75%, 70%, and 65%, respectively.

- We found no strong evidence that external hops are influenced by work experience, job age, and number of

skills, as no apparent patterns emerged on the charts. As such, we cannot accept our initial hypothesis. A plausible explanation is that other incentives that are unobservable from our data—such as monetary, work

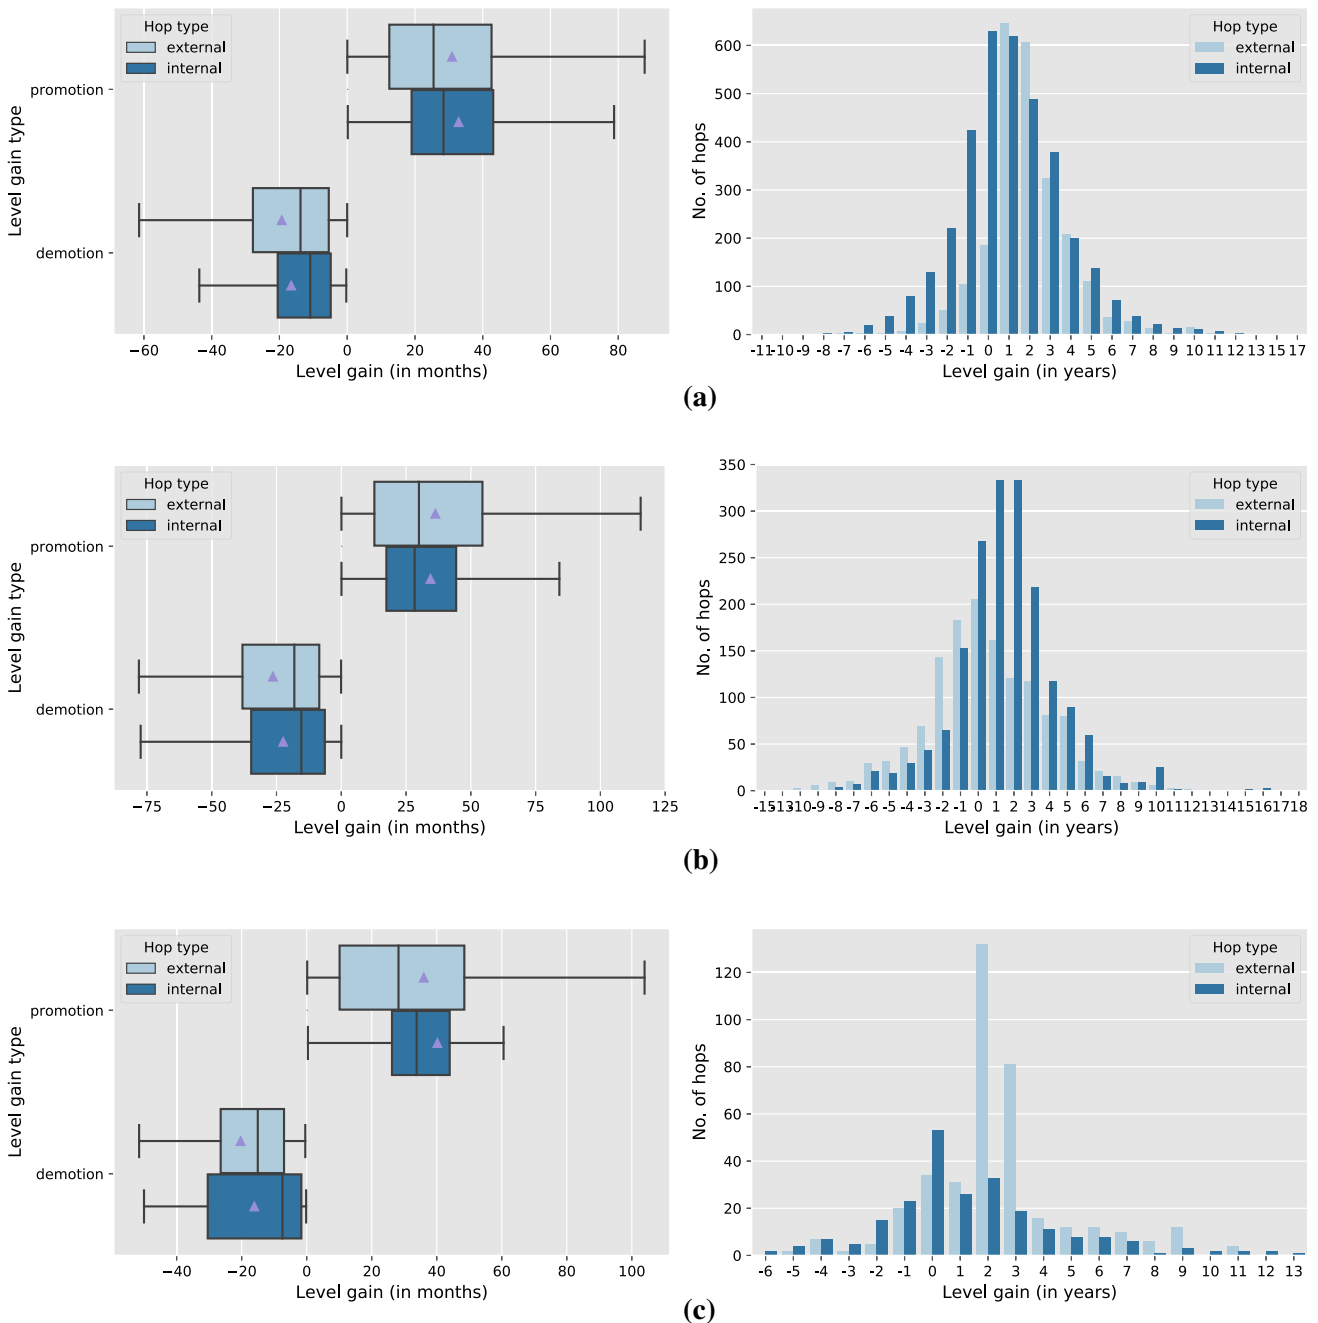


Fig. 11 Comparison of level gains for different hop types. **a** Singapore. **b** Switzerland. **c** Hong Kong

packages, and perks—might play more important roles in incentivizing external hops. Further investigations can be done by augmenting auxiliary information from other data sources from the relevant authorities, which is beyond the scope of our current work.

- For the Hong Kong data, we do not observe any users with work experience ≥ 20 years. This conforms with our earlier finding in Fig. 7, whereby the maximum job level found in the Hong Kong data is about 13 years. We can attribute this to the fact that Hong Kong is a

relatively young special territory of China, which started in 1997.

Naturally, one can extend the above analysis at more granular levels, such as company and industry levels within a country. It is worth highlighting, however, that the data at these levels are typically sparse and/or noisy. Indeed, we have observed in our OPN data that many small companies and industries have low support (i.e., low number of people and/or jobs). This would lead to unreliable metrics/statistics computation that would prevent us from deriving

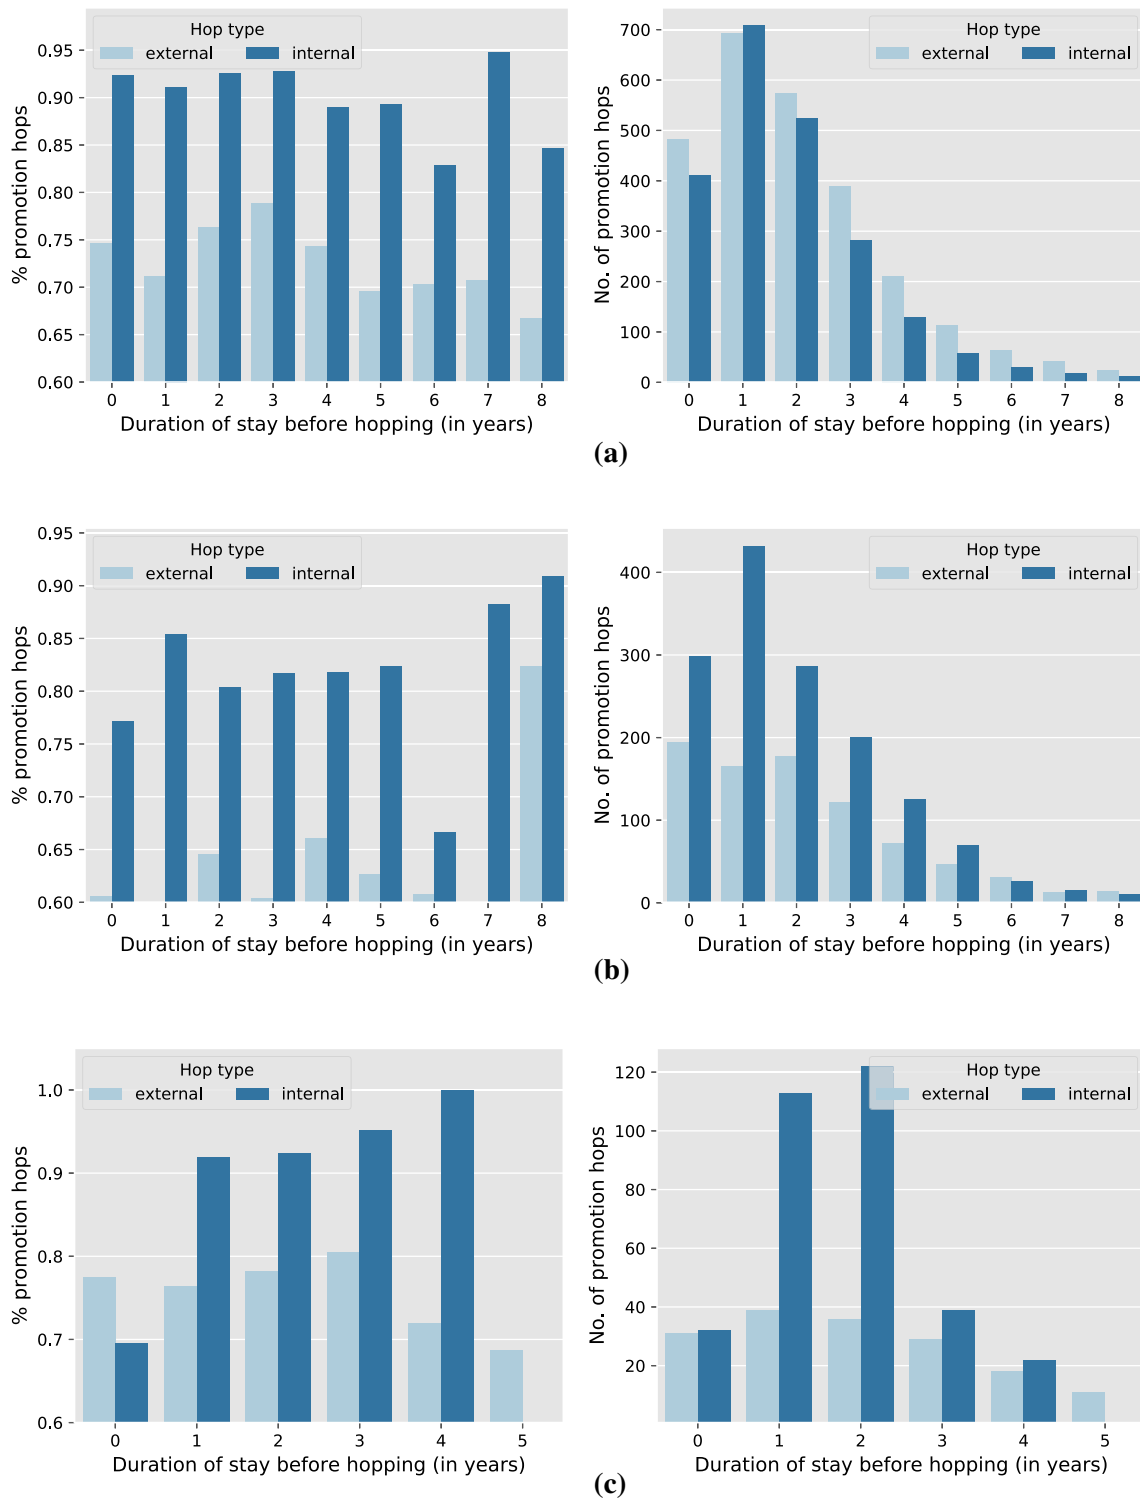


Fig. 12 Promotion hop fraction and counts for different durations of stay. **a** Singapore. **b** Switzerland. **c** Hong Kong

meaningful insights. Handling data sparsity using a more sophisticated analytics approach is beyond the scope of this paper, but is certainly an avenue worthy further investigation in the future.

6.3 Job Attribute Analysis

As promotion is often a cited reason for people leaving one job for another, we now conduct a promotion and demotion analysis by dividing the hops into external and internal

hops based on level gain (i.e., promotion vs. demotion). Tables 4, 5, and 6 give the statistics of promotion and demotion in Singapore, Switzerland, and Hong Kong, respectively, and Fig. 11 provides more detailed results. To get a reliable estimate of level gain—and in turn reliable promotion or demotion labels, we require both source and target jobs for each hop must fulfill the (default) minimum support of 10. As such, we do not include in the tables and figures hops that fail to meet the minimum support criterion.

From the results in Tables 4, 5, and 6, we can derive several conclusions:

- The Singapore workforce does substantially more external hops than internal hops. In contrast, Switzerland and Hong Kong workforces perform less external hops than internal hops. This suggests that the Switzerland and Hong Kong workforces are generally more “loyal” than that of Singapore.
- The probability of promotion is generally greater than that of demotion (for both external and internal hops across the three countries). Specifically, the Singapore dataset shows 81% promotion as compared to 19% demotion, Switzerland dataset shows 73% promotion as compared to 27% demotion, and Hong Kong dataset shows 85% promotion and 15% demotion. This matches the common intuition that a hop is more likely motivated by a job promotion than demotion.
- The Singapore workforce generally tends to seek promotion via external hops, whereas Switzerland and Hong Kong people prefer to seek promotion via internal hops. This relates back to our first point, about the Switzerland and Hong Kong workforces being more loyal than that of Singapore.

Figure 11 shows a more fine-grained detail in terms of the level gain distribution. It is evident that the majority of the level gain values are positive, again suggesting that hopping most likely involves promotion rather than demotion [i.e., $p(\text{promotion}) > p(\text{demotion})$]. We also found in all datasets that promotions generally give people a level gain of about 2 years, and demotions generally give people a job-level loss of about 1 year. The distribution curve at the right charts of Fig. 11 shows that the curves for external hops in Singapore and Hong Kong are at the right of that for internal hops. This pattern indicates that external hops in Singapore and Hong Kong tend to give people higher-level gain than internal hops. The pattern is quite the opposite for Switzerland. Internal hops in Switzerland results give higher-level gain compared to external hops. This finding is consistent with the previous finding in Table 5.

In addition, we investigate whether promotion hops vary with the duration of stay (at some job) before hopping. Figure 12 shows the promotion hop fractions (i.e.,

$p(\text{promotion}|\text{external hop})$ and $p(\text{promotion}|\text{internal hop})$) as well as promotion hop counts as a function of duration of stay prior to hopping. For these plots, we also set the minimum support threshold to filter out unreliable statistics. The right chart of Fig. 12 suggests that promotion hops most commonly happen after a person works for 1–2 years. However, the left chart of Fig. 12 indicates no obvious relationship between the duration of stay and promotion hop fraction. Regardless, it is again evident that the probability of promotion is higher for internal hops than for external hops.

As with the hop analysis in Sect. 6.2, it is possible to further extend the above-mentioned job attribute analysis to company and industry levels. However, we again observe data sparsity issue whereby many small companies and industries in our data have low support, potentially yielding unreliable metrics/statistics computation and inaccurate conclusion. Devising a better analytics approach to deal with data sparsity issue will be left for future work.

6.4 Connectivity Analysis

Network structure analysis In this section, we analyze the job hop behavior at the network level, which includes job and organization graphs. We set edge minimum support equals to two for this analysis. The basic statistics of the job and organization graphs are summarized in Tables 7, 8, and 9. We can conclude that all talent flow network graphs are sparse in general, having small number of edges relative to the squared number of nodes. We also examine the connectedness of the graphs by looking at the strongly connected component (SCC) and weakly connected component (WCC) metrics. The former checks for connectedness by following the directionality of the graph edges, whereas the latter ignores the directionality.

Overall, the results in Tables 7, 8, and 9 indicate that there exists a giant component for both job and organization graphs, and its size is significantly bigger than the second largest component. As such, we can conclude that our job and organization graphs are fairly well connected, in the sense that there exists a path between any two nodes within the giant components.

With the connectedness trait validated, we now examine the centrality properties of the nodes in our hop graphs. Figure 13 presents the complementary cumulative distribution functions (CDFs) of the in-degree, out-degree, and PageRank centralities for the job graph. It is shown that all three metrics exhibit heavy-tail, skewed distribution. We performed power-law fitting and obtained exponent terms of greater than 2 for all graphs, thereby indicating a scale-free phenomenon. Similar result was obtained for the organization graph, although the results are not shown here due to space constraint.

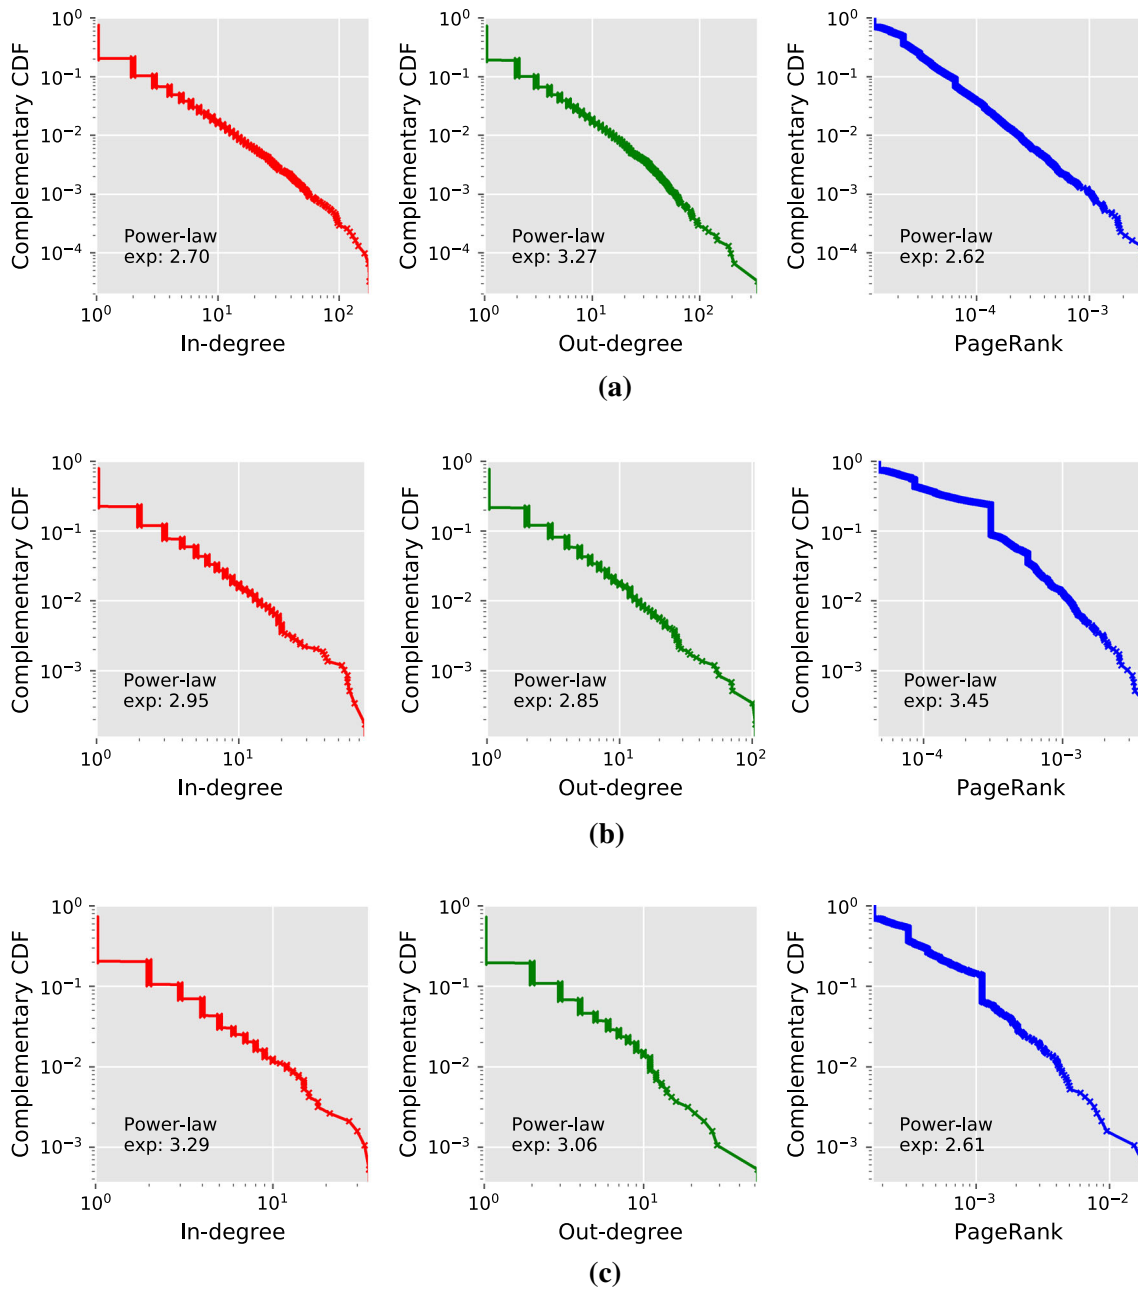


Fig. 13 Centrality distribution of job hop graph. **a** Singapore. **b** Switzerland. **c** Hong Kong

Job centrality analysis Next, we evaluate the top nodes having the highest centrality values in the job-level and organization-level graphs for Singapore, Switzerland, and Hong Kong, as shown in Figs. 14, 15, and 16, respectively. The results provide several interesting insights. For the job graph, we find that the top in-degree, out-degree,

and PageRank jobs are overall dominated by major industries.⁶

From the left charts of Figs. 14a, 15a, and 16a, we can see that the top in-degree nodes refer to those popular jobs in major industries that attract talents. Meanwhile, the middle charts of Figs. 14a, 15a, and 16a suggest that the top out-degree jobs are those that involve versatile skills (e.g., software engineer, consultant) or interim roles (e.g., intern). People having these jobs may thus be able to move

⁶ Major industry codes: 1 = Information Technology and Services, 3 = Banking, 4 = Financial Services, 7 = Accounting, 10 = Computer Software, 24 = Higher Education, 26 = Management Consulting.

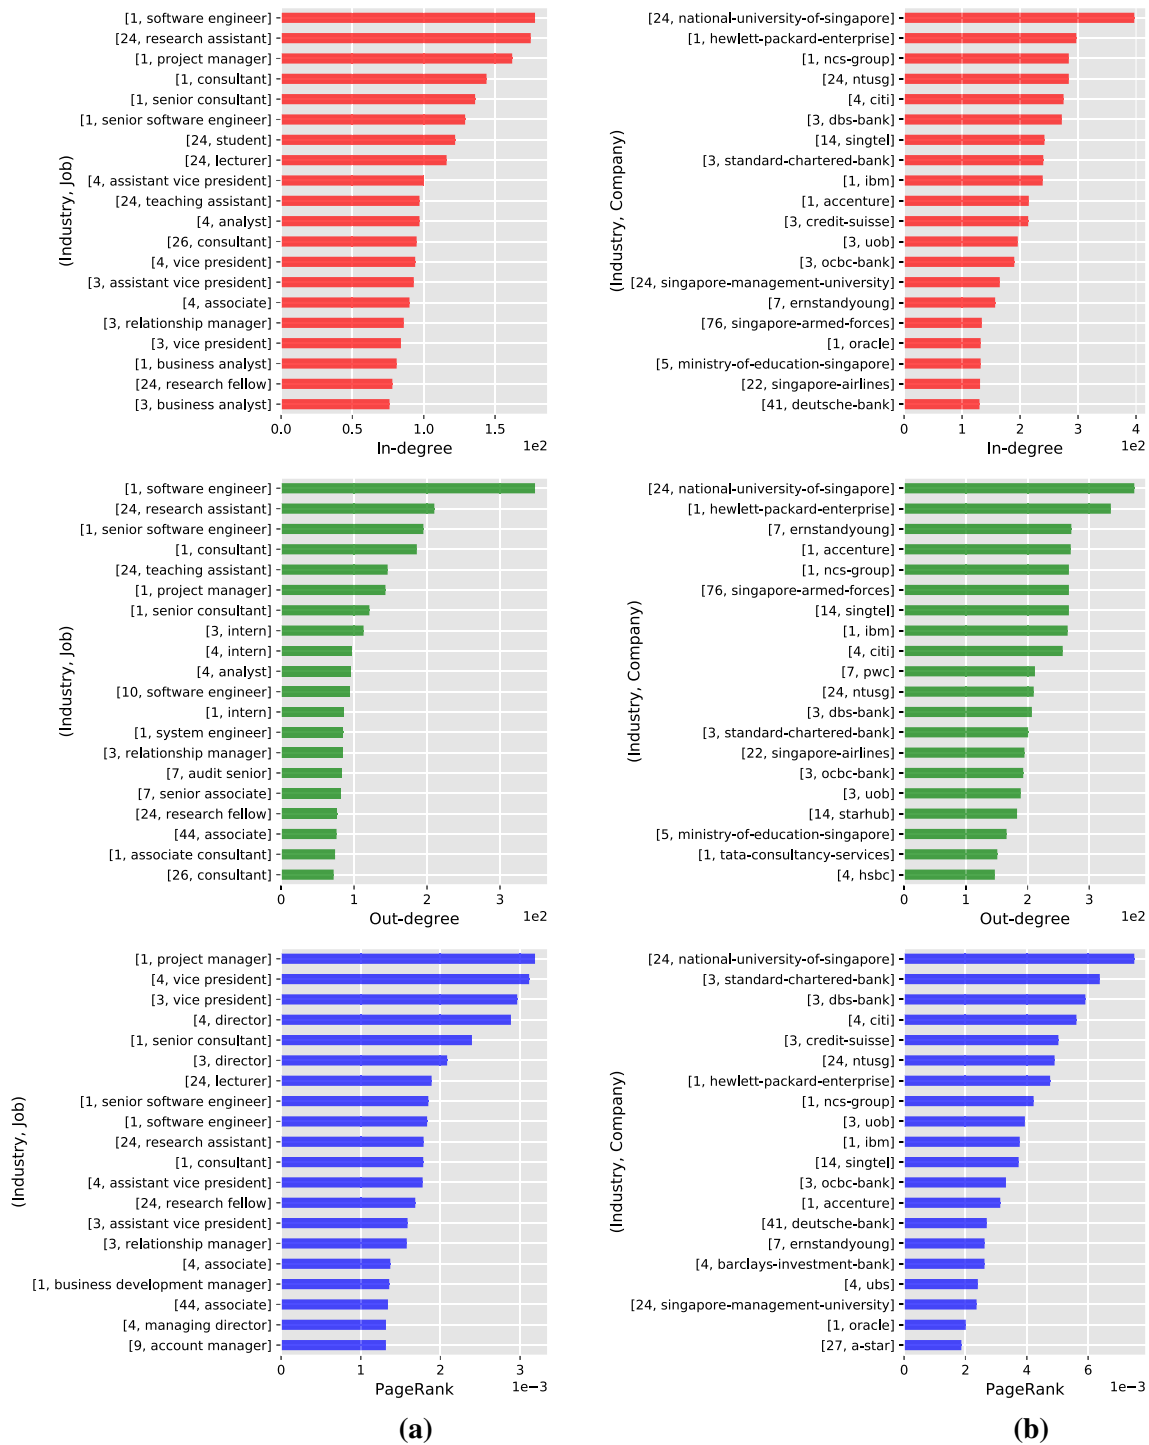


Fig. 14 Centrality of top jobs and companies in Singapore. **a** Top jobs in Singapore. **b** Top companies in Singapore

to more diverse range of jobs/organizations (i.e., talent supplier). Finally, the right charts of Figs. 14a, 15a, and 16a show that the top PageRank nodes correspond to high-level, managerial jobs (e.g., Director, Manager, Vice President). This conforms with our intuition on PageRank as a measure of job desirability (cf. Sect. 5.3).

Organization centrality analysis Figures 14b, 15b, and 16b show the top companies in the three countries/regions based on in-degree, out-degree, and PageRank. The top companies returned by these measures are large corporations. Given that different set of companies operate in these countries/regions, it is not feasible to compare top

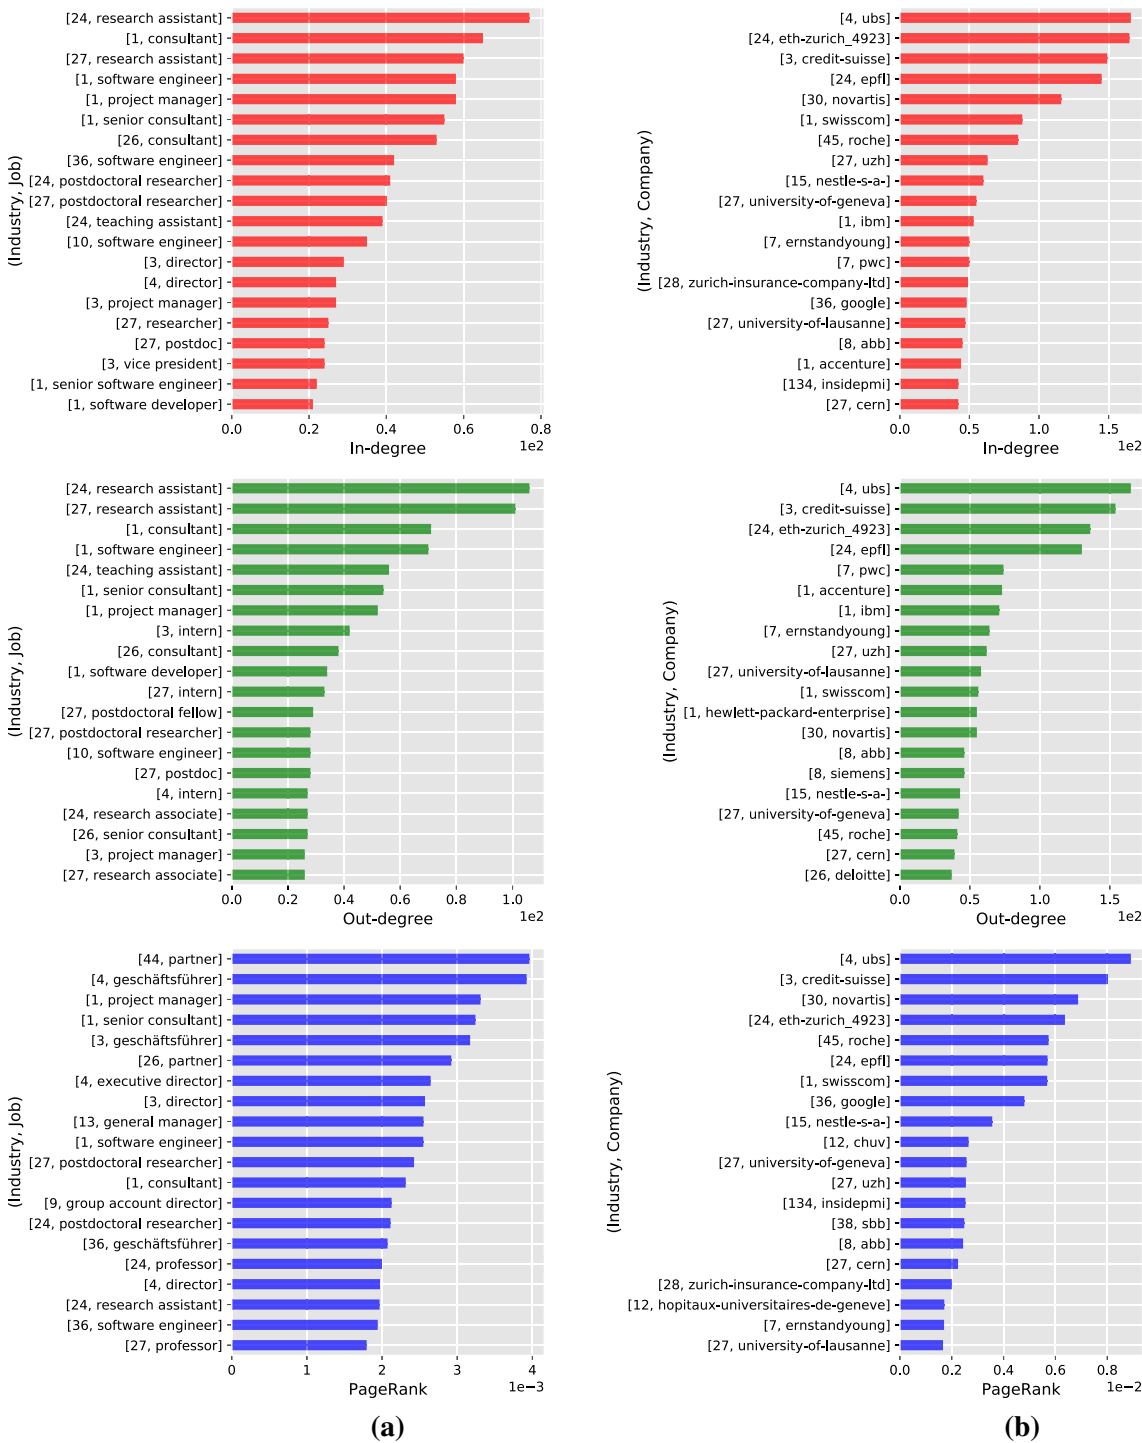


Fig. 15 Centrality of top jobs and companies in Switzerland. **a** Top jobs in Switzerland. **b** Top companies in Switzerland

companies across countries/regions. It is, however, noted that the top few companies of each country can be quite different when applying the different measures. Among the three countries/regions, Switzerland seems to have more top company overlaps between the measures.

7 Conclusion

In this paper, we put forward a data analytics approach to study job hops at a large scale using the OPN data from multiple countries/regions. In conclusion, our study provides a few key takeaways:

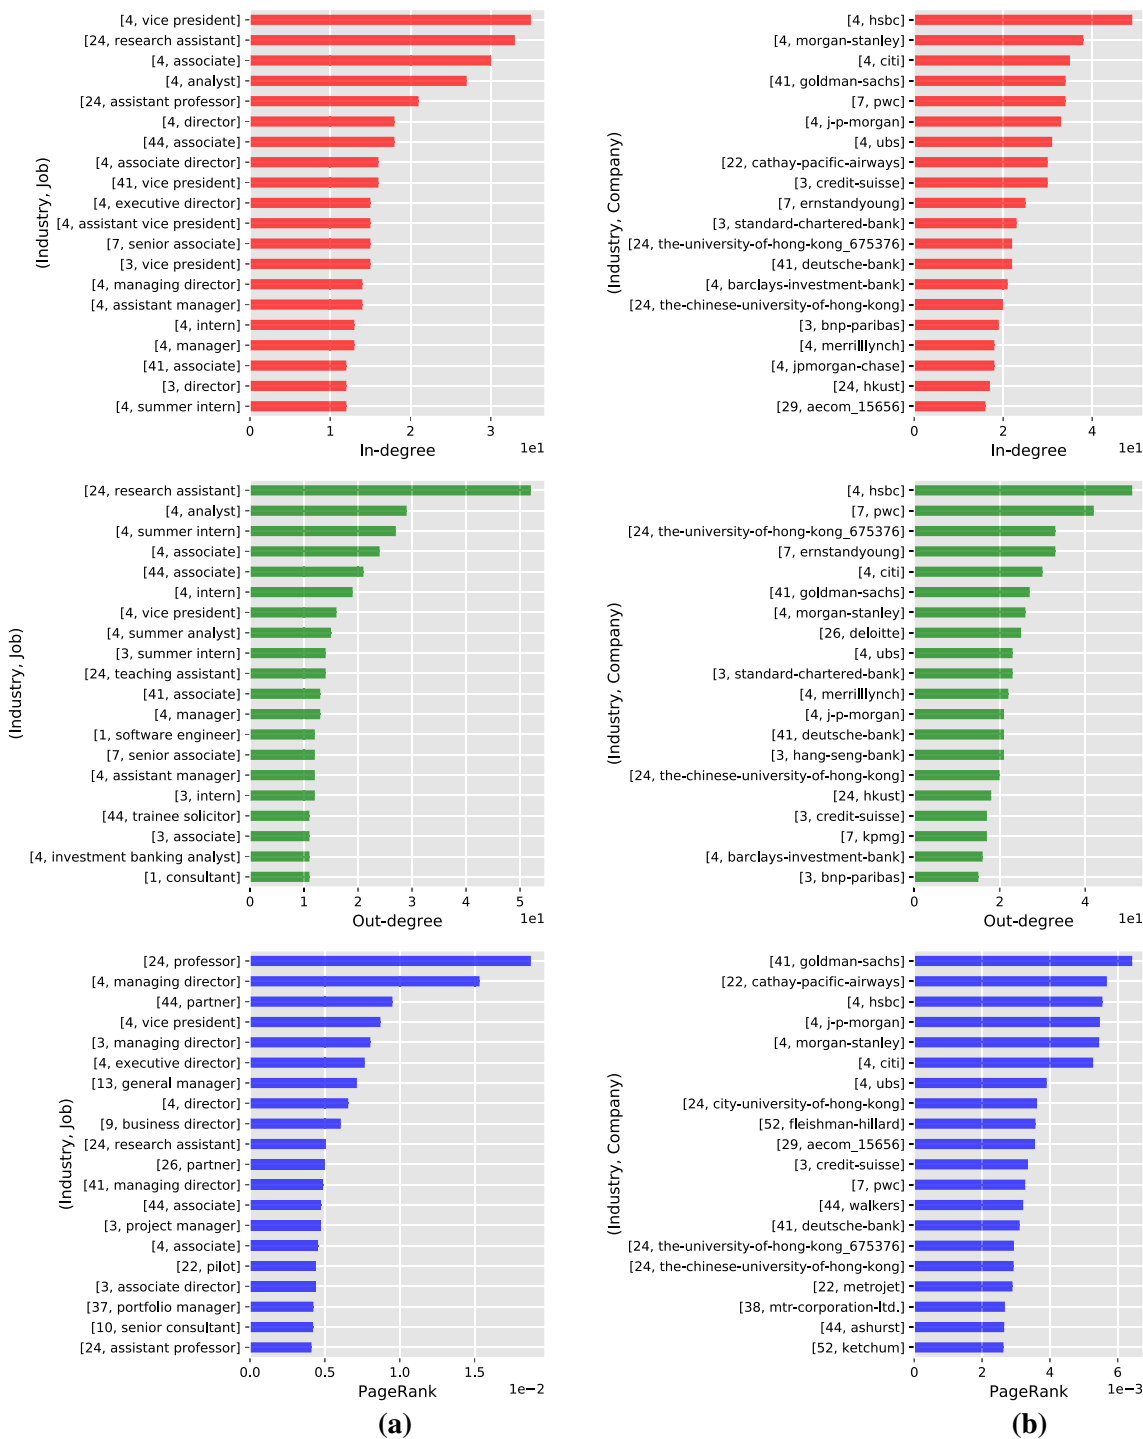


Fig. 16 Centrality of top jobs and companies in Hong Kong. a Top jobs in Hong Kong. b Top companies in Hong Kong

- We discover that external hops are not necessarily influenced by work experience, job age, and number of skills. We also observe that external hops are very common, and the Singapore workforce exhibits the highest external hop fraction among all the three countries/regions studied in this work.
- Our analysis on hop classification and job attribute demonstrate that: (1) external hops are very common; (2) job hopping involves promotions more likely than demotions, and people are more likely to get promoted due to internal hops than getting promoted due to

external hops; (3) promotion hops most commonly happen after a person works for 1–2 years.

- From our network connectivity analyses, we find that: (1) top in-degree job (organization) nodes are prominent jobs (companies) that attract talents, whereas top out-degree job (organization) nodes are influential jobs (organizations) that supply talents; and (2) job (organization) nodes with high PageRank refer to desirable, major jobs (organizations) that are well known for providing good career offering.

Our comparative study on the OPN data from Singapore, Switzerland, and Hong Kong has also enabled us to gain additional insights on the unique characteristics of the workforces in different countries/regions, such as:

- For the same job title, jobs in Switzerland tend to have 2–3 longer years of experiences than jobs in Singapore, whereas jobs in Hong Kong tend to have more or less comparable years of experience to jobs in Singapore.
- Most jobs in Switzerland tend to have longer job ages than those in Singapore users, suggesting that, for the same job title, the Switzerland workforce is ahead of the Singapore counterparts in terms of job establishment. In contrast, the Hong Kong users tend to have lower job age than the Singapore users, implying that jobs in Hong Kong are generally less established than those in Singapore.
- The resulting statistics of external and internal hops suggest that the Switzerland and Hong Kong workforces are generally more “loyal” than the Singapore workforce. This is evident from the significantly higher proportion of external hops (relative to internal hops) in the Singapore data.

The findings from this paper lead to a few possibilities. Firstly, we demonstrate that it is possible to repurpose the career histories of OPN profiles to study the job hop patterns of workforce within a country/region, and to compare across countries/regions. This vastly improves the scale and granularity of job hop study, which was traditionally done using surveys. Through our analysis, we show that the propensity to perform job hops is relatively higher among the young workforce than the older one. This could lead to two main concerns, namely: (i) the limited time to acquire adequate skills on the job among the young employees; and (ii) the unwillingness of companies to provide them skill training. These concerns may cost the workforce long-term’s skill development and productivity. To overcome these, more incentives may be introduced to encourage young employees to stay longer on their jobs. One could also increase the chance of job promotions among the younger employees.

Finally, our analysis also shows that job and organization graphs are well connected. We further define job centrality measures to determine attractive jobs and companies. Such measures allow jobs and companies to be ranked for applicants’ reference during job search. These measures can also be further refined to find attractive jobs and companies in specific industry domains.

Acknowledgements This research work is supported by the National Research Foundation, Prime Minister’s Office, Singapore, under its International Research Centres in Singapore Funding Initiative.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Alper J (1994) Job-hopping to greater career heights. *Science* 265(5180):1924–1925
2. Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Comput Netw ISDN Syst* 30(1):107–117
3. Chaudhury A, Kim M, Tiwari M (2016) Importance of first steps in a community for growth, diversity, and engagement. In: WWW, pp 19–20
4. Cheng Y, Xie Y, Chen Z, Agrawal A, Choudhary A, Guo S (2013) Jobminer: a real-time system for mining job-related patterns from social media. In: KDD, pp 1450–1453
5. Friedell K, Puskala K, Smith M, Villa N (2011) Hiring, promotion, and progress: Millennials’ expectations in the workplace. Working paper, St. Olaf College
6. Fuller S (2008) Job mobility and wage trajectories for men and women in the united states. *Am Sociol Rev* 73(1):158–183
7. Ha-Thuc V, Xu Y, Kanduri SP, Wu X, Dialani V, Yan Y, Gupta A, Sinha S (2016) Search by ideal candidates: next generation of talent search at linkedin. In: WWW, pp 195–198
8. Hamari J, Sjöklint M, Ukkonen A (2016) The sharing economy: why people participate in collaborative consumption. *J Assoc Inf Sci Technol* 67(9):2047–2059.
9. Hamori M (2010) Job-hopping to the top and other career fallacies. *Harv Bus Rev* 88(7–8):154–157
10. Joseph D, Boh WF, Ang S, Slaughter S (2012) The career paths less (or more) traveled: a sequence analysis of IT career histories, mobility patterns, and career success. *MIS Q* 36(2):427–452
11. Kapur N, Lytkin N, Chen BC, Agarwal D, Perisic I (2016) Ranking universities based on career outcomes of graduates. In: KDD
12. Langville AN, Meyer CD (2005) A survey of eigenvector methods for web information retrieval. *SIAM Rev* 47(1):135–161
13. Li H, Ge Y, Zhu H, Xiong H, Zhao H (2017) Prospecting the career development of talents: a survival analysis perspective. In: KDD, pp 917–925
14. Liu Y, Zhang L, Nie L, Yan Y, Rosenblum DS (2016) Fortune teller: predicting your career path. In: AAAI, pp 201–207
15. Long J, Ferrie J (2003) Labour mobility. In: Mokyr J (ed) *The Oxford encyclopedia of economic history*. Oxford University Press, pp 248–250

16. Maier C, Laumer S, Eckhardt A, Weitzel T (2015) Who really quits?: A longitudinal analysis of voluntary turnover among it personnel. *SIGMIS Database* 46(4):26–47
17. Moscarini G, Thomsson K (2007) Occupational and job mobility in the US. *Scand J Econ* 109(4):807–836
18. Oentaryo RJ, Ashok XJS, Lim EP, Prasetyo PK (2017) On analyzing job hop behavior and talent flow networks. In: *ICDMW*, pp 207–214
19. Schawbel D (2013) The cost of millennial retention study. Technical Report, Millennial Branding
20. State B, Rodriguez M, Helbing D, Zagheni E (2014) Migration of professionals to the US: evidence from LinkedIn data. In: *SocInfo*, pp 531–543
21. Topel RH, Ward MP (1988) Job mobility and the careers of young men. Working Paper 2649, National Bureau of Economic Research
22. Xu H, Yu Z, Xiong H, Guo B, Zhu H (2015) Learning career mobility and human activity patterns for job change analysis. In: *ICDM*, pp 1057–1062
23. Xu H, Yu Z, Yang J, Xiong H, Zhu H (2016) Talent circle detection in job transition networks. In: *KDD*
24. Xu Y, Li Z, Gupta A, Bugdayci A, Bhasin A (2014) Modeling professional similarity by mining professional career trajectories. In: *KDD*, pp 1945–1954