ACI Research Paper #15-2022

Surveying the Singapore Urban Labour Market Using Online Job Postings

Terence HO

Jing Zhi LIM

Lucas SHEN

August 2022

Please cite this article as:

Ho, Terence, Jing Zhi Lim, Lucas Shen, "Surveying the Singapore Urban Labour Market Using Online Job Postings," Research Paper #15-2022, *Asia Competitiveness Institute Research Paper Series (August 2022)*

Surveying the Singapore Urban Labour Market Using Online Job Postings

Terence Ho* Jing Zhi Lim† Lucas Shen‡

August 2022

Contents

*Ho is Associate Professor in Practice at the Lee Kuan Yew School of Public Policy.
Email: sppwlth@nus.edu.sg
†Lim is Research Assistant at the Asia Competitiveness Institute, Lee Kuan Yew School of Public Policy.
Email: jingzhilim@nus.edu.sg
‡Shen is Research Fellow at the Asia Competitiveness Institute, Lee Kuan Yew School of Public Policy.
Email: lsys@nus.edu.sg

Public and official data sources on the local labour market (such as vacancy rates, recruitment rates, employment changes, median salaries, value-added per worker, etc.) are readily available at an aggregated level. However, the extent to which certain skills are in demand and how hiring managers value workers of different experiences and profiles remain opaque.

In this report, we use an early tranche of online job postings data from JobTech, Singapores leading AI talent intelligence startup. The data on 2018–2021 job ads from the online job boards constitutes close to the full universe of local online job ads for that period and includes the usual set of information on a job ad, such as the qualifications required and the full text data of the job descriptions. As a complement to traditional surveys, we use these data to monitor local labour market dynamics and understand what hiring managers demand of workers.

We start by summarising key findings in our report in seven points, followed by a brief reflection on policy takeaways. The main body of our report unpacks our descriptive findings into six broad themes that are collected into sections. The themes are: (1) how online job ads are able to track the labour market, (2) labour market rigidity in terms of hiring based on formal qualifications, (3) the demand and change in demand for skills, (4) the state and trends in AI-related jobs and skills, (5) trends in working from home during the Covid-19 pandemic, and (6) how the government-run MyCareersFuture job board is differentiated from the other job boards.

Each theme is self-contained and has a slightly different subset or merged dataset. Each section therefore starts by describing the data used for the charts in the section before summarising the results. The data and decisions we make regarding the data are described in full in a dedicated Materials and Data section at the end of the report.

Key Findings

1. **Online job ads track the market well.**
   The first-order concern with surveying the labour market using online job ads is that job ads may not correspond one-to-one with job openings. We use vacancy rates from the Labour Market Survey to find that job ads move well with job vacancies: more vacancies usually correspond to more job ads. This comovement suggests that the jobs ad data we have can reasonably capture vacancies in the labour market.

2. **Qualifications. It's *the* Gateway.**

   Despite concerted efforts to upskill workers and encourage continuing adult education [1], the market continues to put weight on formal education requirements, which most job advertisements still state. The share of job ads that do not require formal education rose in 2021, but these jobs are associated with lower salaries, suggesting that the rise in such jobs ads is due more to a compositional shift in the labour market than to hiring flexibility.

3. **Not all tertiary specialisations are valued equally.**

   By presenting jobs as collections of skills, we observe that not all skills are valued equally. Skills related to data and computer science are frequently demanded for bachelors jobs, masters jobs, and PhD jobs. While the bachelors and masters jobs also frequently demand other non-STEM skills such as project management or marketing, the demand for skills in PhD jobs is extremely skewed towards AI. Most jobs requiring a PhD are in fact seeking employees with highly specialised data or AI expertise.

4. **Seven of the fastest-growing skills are STEM**

   We observe a clear acceleration of demand for STEM-related skills, with seven of the ten fastest-growing skills being those related to technology and engineering. In particular, these skills are related to data analytics, general programming for product deployment, and cloud computing. In contrast, only one in ten of the fastest falling skills demand is STEM-related (SAP).

5. **Machine learning dominates AI.**

   Six sub-categories of AI-related skills display a steady trend in which machine learning strongly dominates other AI categories by share of AI skills for all periods in this sample. While we observe that the media coverage on artificial intelligence typically highlights deep learning algorithms and robotics, the market currently requires more skills related to machine learning. The dominance of machine learning over robotics, in particular, tells us that on the AI frontier, the software component trumps the hardware.

6. **Full Work from home is not the norm**

   As our analyses and sample period coincide with the Covid-19 pandemic, we also examine the proportion of job ads that indicate that the advertised job can be done remotely. According to our higher resolution tracking of work-from-home jobs ads, hirers respond promptly to policy changes. The surge in work-from-home job ads occurred during key periods when Covid-19 first gained global media attention and when the government made key announcements. While our estimate is mediated by the dependence of the recall rate of work-from-home jobs on our search list, the evidence suggests that the most optimistic estimate of (new) jobs that can work from home fully is only 2%.

7. **Government-run job board is differentiated**
   Comparing the government-run MyCareersFuture job board with the other job boards reveals tangible differences. First, except for rare cases, job postings on MyCareersFuture always disclose the salary range for the advertised job. Since salary is often a key concern for job-seekers, this helps to reduce labour market friction and gives MyCareersFuture an advantage over other job boards. Second, we see that MyCareersFuture has a larger share of the SGUnited jobs that support the local workforce, as it aims to mitigate the adverse impact of the Covid-19 pandemic.

Policy Takeaways

Online job postings are a useful complement to labour market data and official statistics as they help to shed light on job requirements such as skills and qualifications in a timelier manner. For instance, it would be useful to track the number or frequency of job postings that require skills qualifications other than formal degrees or diplomas, as this would indicate that recognition of skills is gaining traction. Also, it would be useful for policymakers to know if the proportion of skilled jobs requiring relevant experience is changing, as employers are encouraged to tap on the broader pool of workers who can be trained even if they do not already possess relevant experience.

The data suggest that STEM skills, in particular programming and machine learning, are in high and growing demand. Accordingly, efforts to train students and workers in these skills should be stepped up to prepare the workforce to thrive in the digital economy. While working from home over the past two years was necessitated by the COVID-19 outbreak and safe management measures, it is important to know to what extent WFH will remain a permanent fixture going forward. In particular, how will workers in different sectors divide their time between the office and home on a permanent basis? This has implications for office space demand and urban planning, including the transport system.

Theme 1: Tracking the Labour Market

Data for this section

The data for this section comes from combining JobTech data on online postings with the official labour market survey statistics. We use two sets of official statistics: 1) job vacancy rate and 2) resignation and recruitment rates from the Labour Market Survery, which covers establishments with at least 25 employees. Since the official statistics are only available to us as quarterly data, we similarly aggregate the real-time job posting data to the quarterly-year level. We use our curated source of industry classification for job postings and match it to the official statistics for the broad industry level analyses. Finally, the quarterly data we use is always indexed to the first quarter of 2020, so that all subsequent data points in the series are interpreted as a percentage change in the data since Q1 of 2020.

Job postings tracking official statistics

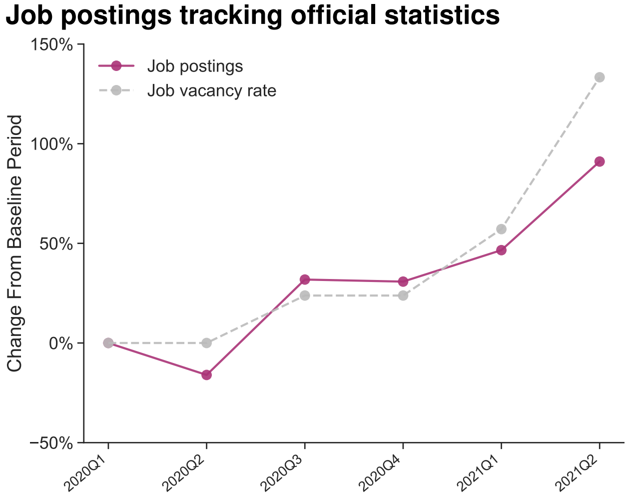

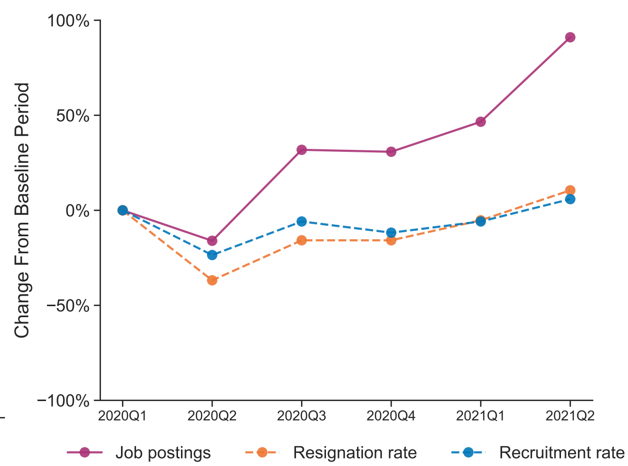

Real-time job postings track the market well

We find that job postings track the labour market well after taking the more recent 2020–2021 cut of the data. This is especially true for job vacancy rates. In particular, our job posting data tracks job vacancy rates extremely closely, from Q1 2020 to Q2 2021 (the last quarter in our sample with complete data).

Diving into subsets of the data, we find that the topline job postings track the official statistics

Job postings tracking official statistics by occupation

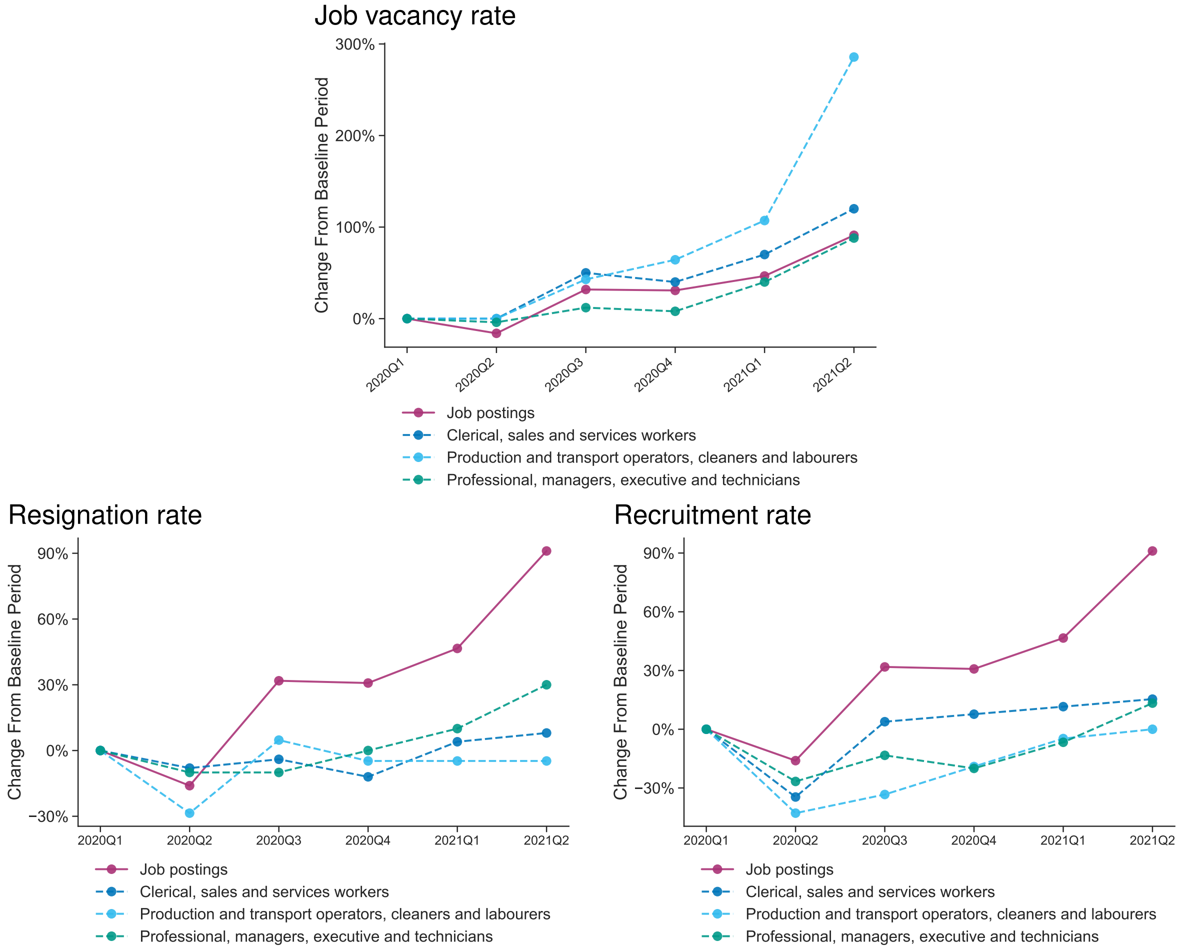

when split by broad occupation group. Since we have the industry classification of job postings, we also show that the job postings do best tracking manufacturing and "others". The fact that job postings do not track the construction sector well is consistent with how vacancies in the construction sector are not typically posted online (see Theme 6: Government vs. Private Job Boards). One conjecture for why services job postings systematically undertracks services job vacancy might be that one job post in the services sector represents many vacancies at once.

Overall, the job postings data tracks the small Singapore urban labour market well and allays concern that job postings are only capturing turnover, to the extent that the job posts track vacancies much better than resignation. The comovement of vacancies and job posts also allays concern that the median job post represents multiple vacancies. Nonetheless, we note that not all market subgroups will be precisely captured in (e.g., construction) with real-time job postings.

Job postings tracking official statistics, by industry

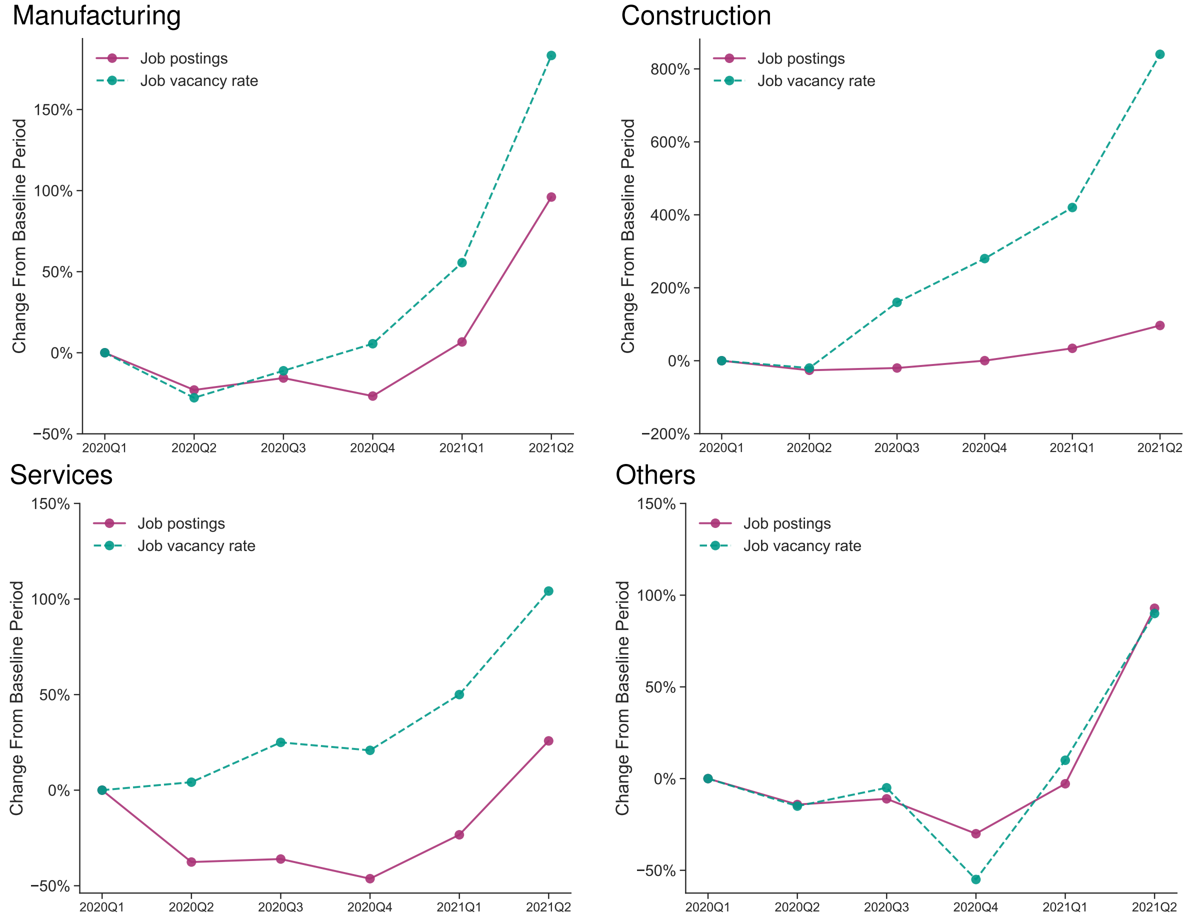

We further emphasise that while we observe that online job ads track the labour market well, caution should be exercised with projecting trends into the future [2, 3]. The composition of online job ads that can predict the movement of job vacancy rate, recruitment rate, and other measures from traditional sources in the current sample period may change or may no longer have utility in tracking future market trends.

Theme 2: Market Rigidity

Data for this section

For the data in this section, we aggregate the real-time job postings data to get the share of jobs by qualifications per year. As many jobs mention multiple qualifications, we use the lowest stated qualification as a minimum-stated requirement (i.e. we classify a job posting that mentions both Bachelor's and Masters's as a Bachelor's-level job). To compute salaries, we use the midpoint of the stated minimum and maximum monthly salary for each job posting and then take the median of these midpoints.

Demand for education qualifications

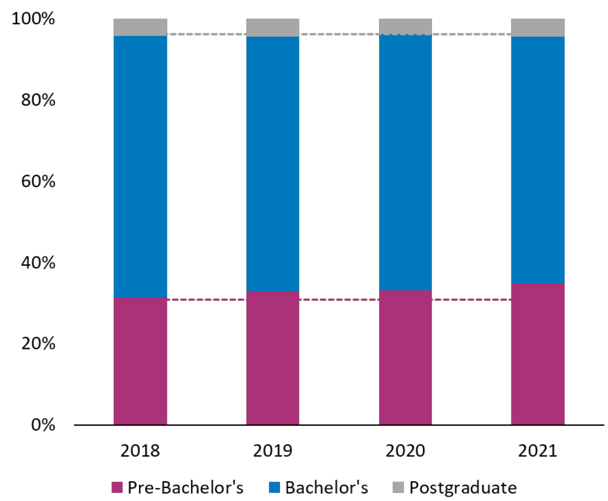

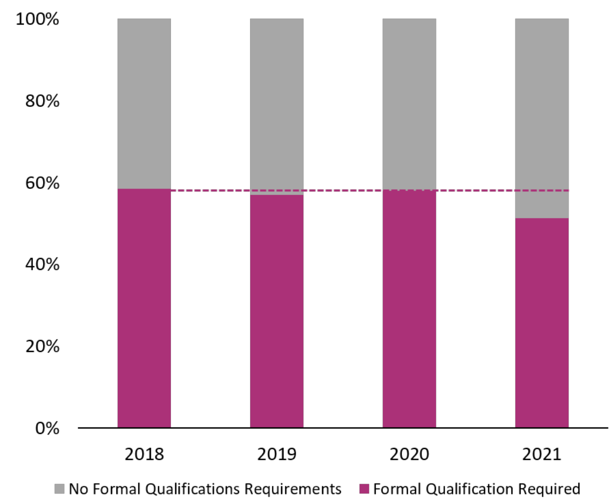

Share of informal and low educational requirement jobs on the rise

Based on the four years of data we have, we observe that the share pre-bachelor jobs (diplomas and NITEC) have increased gradually, from 31% to 35%, suggesting that the market is more flexible regarding the formal qualifications of job candidates. As a result of dividing job postings by whether they state a formal qualification, we find that the share of jobs that do not require formal diplomas and degrees rose by approximately seven percentage points in 2021 as compared to previous years.

Salary for jobs with and without formal requirements

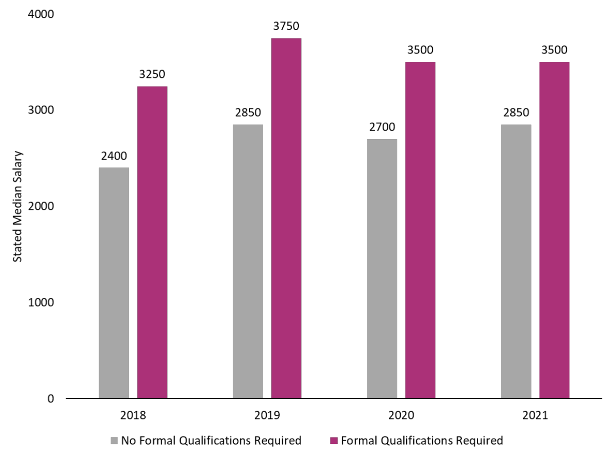

Salary for jobs with different education requirements

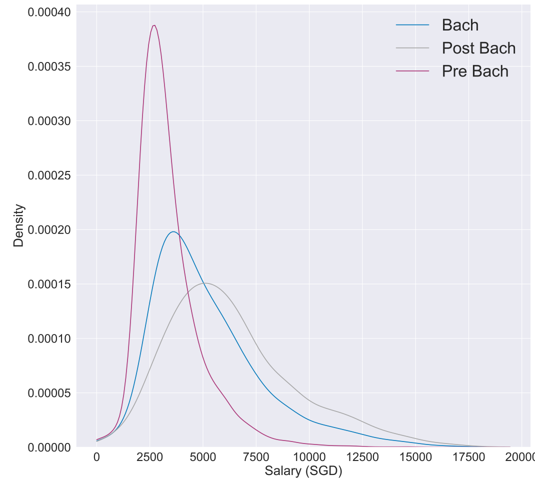

Informal and low educational requirement jobs have lower salaries

In light of the above observations, it appears the labour market is becoming more flexible when it comes to hiring. This is desirable if the market is more receptive to candidates with lower formal qualifications so long as they possess the relevant skills or years of experience. Additionally, it bodes well for initiatives to upskill workers through continuous adult education.

A closer look at salaries, however, offers a vastly different interpretation. No significant change has been observed in the salary for jobs that do not require formal education or training in the past three years. In fact, the salary for these jobs is systematically lower than those requiring formal qualifications. We also find modal salaries for pre-graduate are lowest compared to jobs which require higher educational attainment. In other words, the increasing share of jobs without formal requirements reflects a compositional change in the economy rather than a change in hiring practices.

Theme 3: What Skills Do the Market Demand?

Data for this section

This section uses list of skills per job posting curated by JobTech. The skills word clouds are constructed by up-weighting skills mentioned frequently in a qualification-level and by down-weighting skills mentioned frequently across all qualification levels. Qualifications for jobs are as defined in Theme 2: Market Rigidity. To examine how the demand for specific skills has changed over time, we bin job postings in our sample period's first and last twelve months. We then compute skill shares based on what proportion of job postings in the two periods mention the skills in the job description. The change in demand for a skill is then the percentage point difference for that skill for the last twelve months versus the first twelve months of the data. We keep only the top 200 skills by share of jobs listing the skills.

Differentiated skills for different workers

Here, we represent jobs as a bag of skills using the list of skills curated by JobTech. For the pre-bachelors qualifications, NITEC and diploma, we see that employers essentially want general skills such as maintenance management, liaison, and willing to take up ad-hoc duties. In comparison, Bachelor's and master's jobs involve more technical skills, such as general programming, computer science, and data analytics. As expected, the dominant skills for Bachelor's and masters also span a wide range, with more general business skills demanded, such as regulatory compliance and budget management.

On the other hand, skills demanded for PhD jobs are very different from all other jobs, where there is a clear emphasis on roles that require deep domain expertise, such as life science and chemistry. The skills often do not involve business functions, indicating that most PhDs are hired to be technical leaders, in contrast to general business managers. It is also clearly evident that the technology sector is a key employer of PhDs in the industry, demanding AI and computer science-related skills such as machine learning, algorithms, and programming.

Finally, the skills for jobs where no formal requirement is listed also look different from the other jobs and resemble jobs that are pre-bachelor jobs. To this extent, representing jobs as a bag of skills also provides evidence that the perceived flexibility in the labour market, where the share of jobs that do not require formal qualification is rising, is spurious and likely driven by the composition of job vacancies rather than a change in hiring flexibility.

Skills by qualification

NITEC

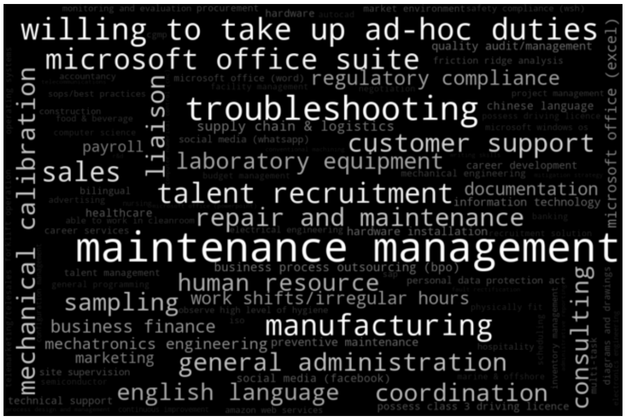

Diploma

Bachelors

Masters

PhD

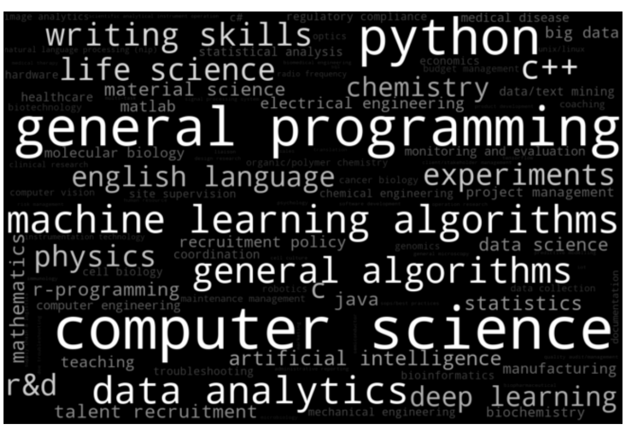

No requirement

Change in skills demanded

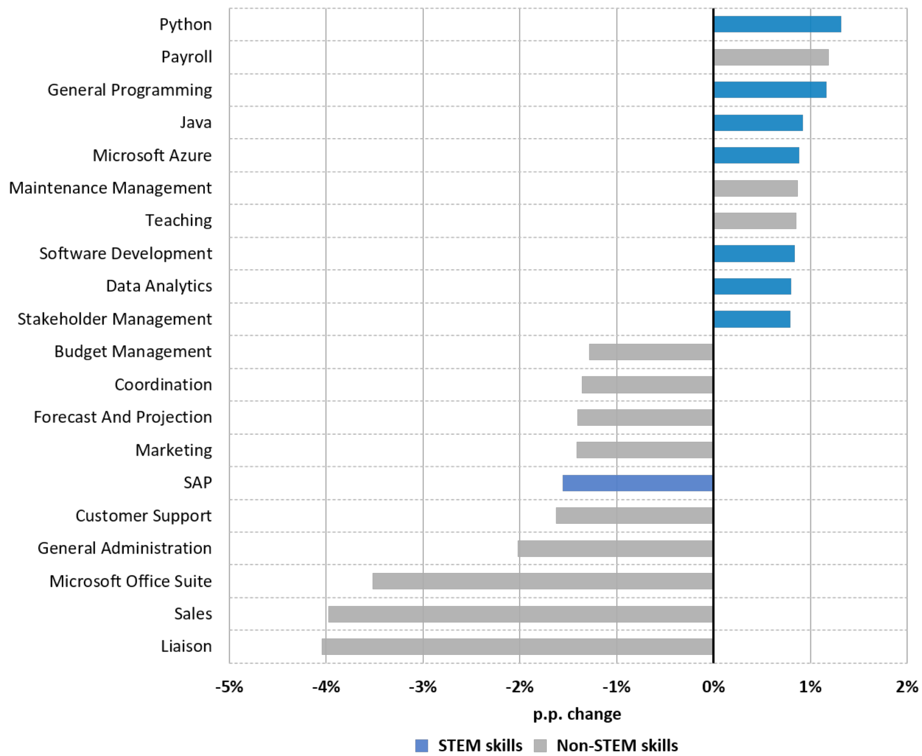

STEM skills gaining ground in the labour market

Focusing on the changes in skills demanded, even in the relatively short sample period, yields interesting observations about the type of skills that are changing the most. We focus on the top ten largest growth skills plus the top ten largest declining skills. Skills that saw the largest increases are related to STEM fields which relates to scripting languages (Python), programming languages (Java), a finding reflected in surveys on general trends [4, 5]. Other fast-growing STEM-related skills include relates to data analytics and cloud computing (e.g. Azure). Conversely, most of the largest decline in skills are non-STEM skills, except SAP (an enterprise resource planning software). The rest of the declining skills demanded are skills such as general administration.

Theme 4: The State of Artificial Intelligence in the Market

Data for this section

To identify job postings relating to “artificial intelligence” (AI) jobs, we use a simple case-insensitive string search using key phrases associated with i) artificial intelligence, ii) machine learning, iii) natural language processing, iv) neural networks, v) robotics, and vi) visual image recognition. Salary is the midpoint between the stated minimum and maximum salary in a job posting. Word clouds are constructed using the same methodology as described in Theme 3: What Skills Do the Market Demand?.

Trend and share of AI jobs

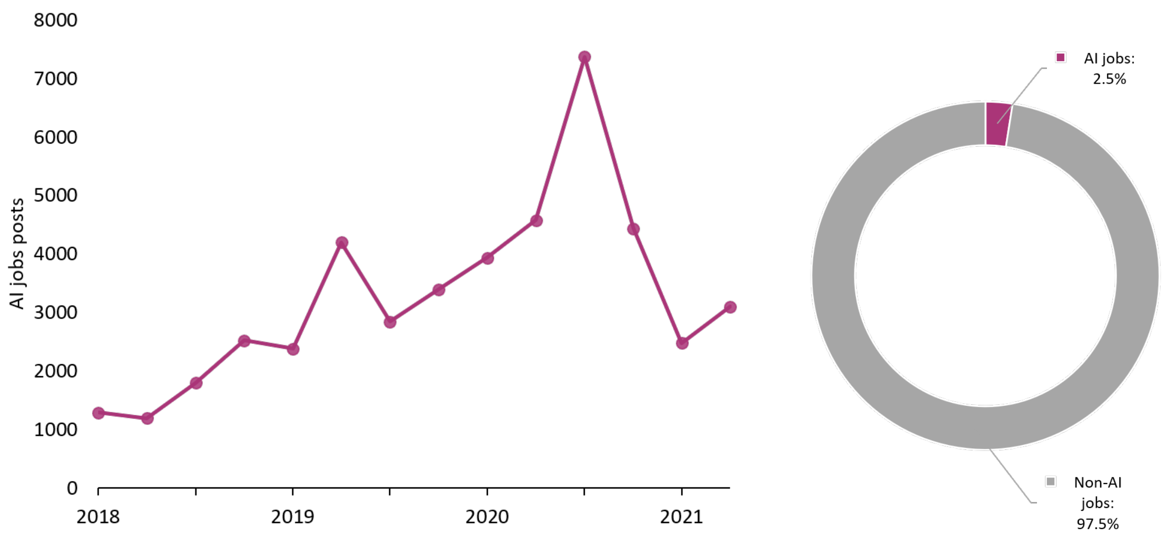

Strong growth and share of AI jobs

A trend analysis reveals growth in AI jobs from 2018 to 2020, before a dip occurs at the end of 2020, in the middle of the pandemic. The dip at the end of 2020 notwithstanding, AI jobs constitute a small but substantial 2.5% of shares in job postings. This finding is similar to the Stanford Institute for Human-Centered Artificial Intelligence’s report [6]. The report states the US, UK, Canada, Australia, and New Zealand all have a share of less than 1%, demonstrating the extent of AI job posting share in Singapore. The upward trend in AI jobs will likely be reinforced with initiatives to equip students (from middle school onwards) via free AI courses [7].

Six AI clusters

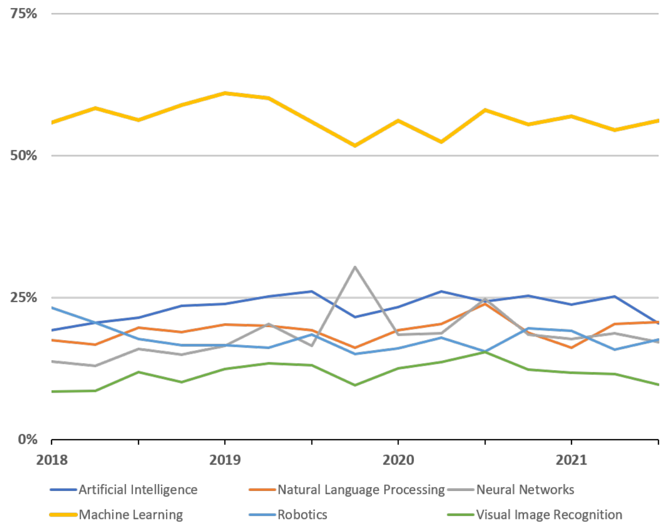

Salary of AI jobs

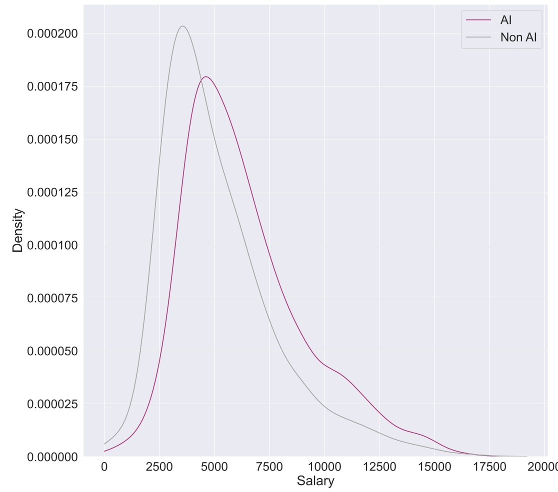

Machine Learning dominates AI jobs

Our analysis also shows that the stated salary in AI job postings is higher than those in the non-AI job postings. This is evident even when we use the median of the listed salaries to avoid extreme values on the right tail. Breaking down AI jobs in the six different clusters reveals that while there is no conspicuous trend in any of the six cluster, the machine learning cluster dominates with more than half the share of all AI job postings. Our Co-occurrence of AI clusters wheel further illustrates the dominance of the machine learning cluster. The (visual) image recognition and robotics clusters have low co-occurence with the other clusters. The remaining three clusters, artificial intelligence, neural networks, and natural language processing mostly only co-occur with the machine learning cluster. All these stand in contrast to the media coverage of AI which disproportionately emphasises robotics clusters.

The high share of machine learning is also corroborated by representing AI and non-AI job postings as a bag-of-skills. We see that the most represented skills for AI jobs relate to machine learning algorithms, general programming, and Python, a specific programming language with a mature ecosystem of off-the-shelf machine learning packages. We also observe some differences in soft-skills between AI and non-AI job postings. AI job postings, for example, are more likely to state a requirement for creative thinking, problem-solving, and analytical thinking. Other skills such as communication have continued relevance for all job postings.

Co-occurrence of AI clusters

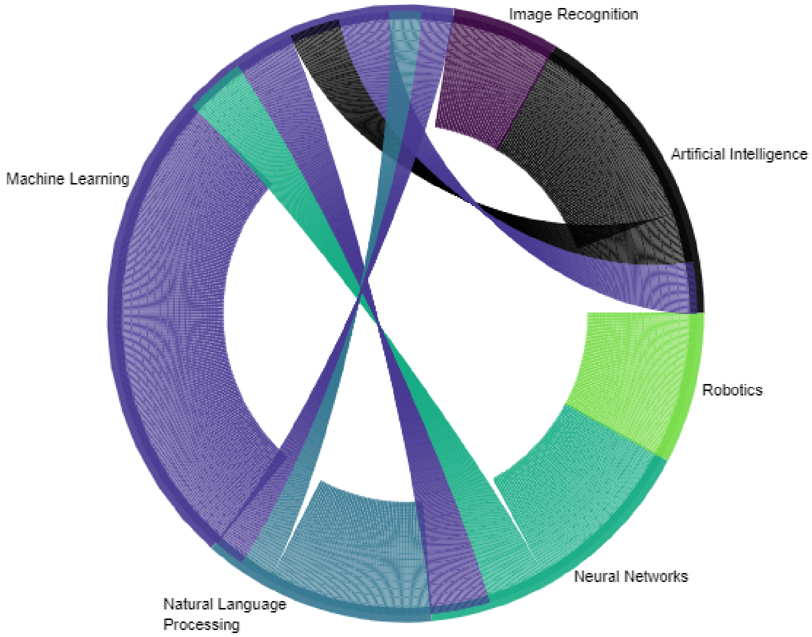

Skills for AI vs non-AI jobs

AI skills

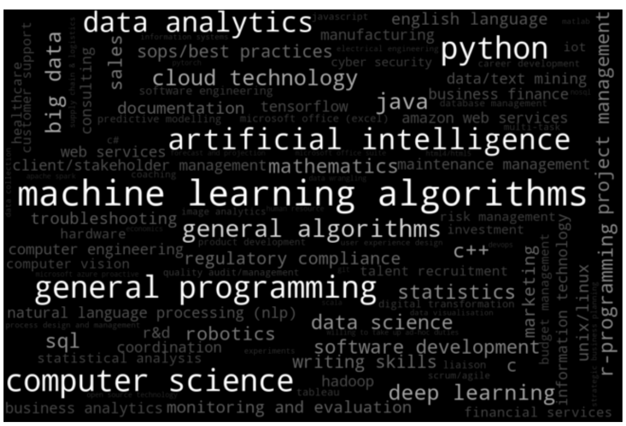

Non-AI skills

AI soft skills

Non-AI soft skills

Theme 5: Working from Home during the Covid-19 Pandemic

Data for this section

To identify job postings relating to "work from home" (WFH) jobs, we use a simple case-insensitive string search using key phrases associated with WFH (e.g., remote work). We then collapse the real-time job postings data to the week-of-year level to obtain the number of WFH jobs and the total number of jobs per week. Otherwise, the data for job postings by industry is similar to the one used in Theme 1: Tracking the Labour Market.

WFH jobs dominated by services sector

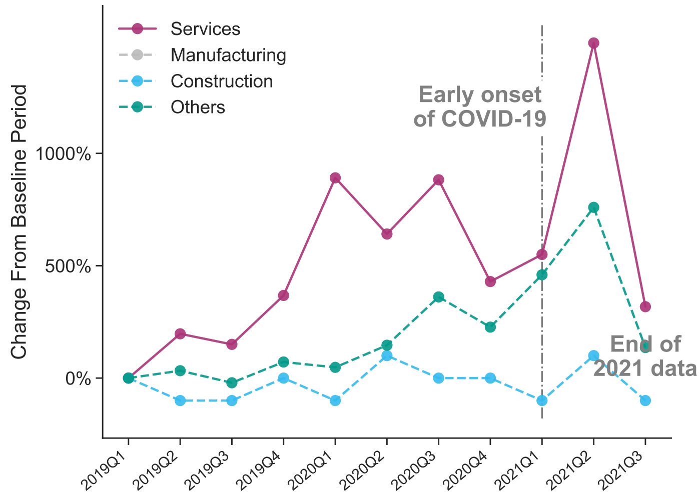

Who gets to WFH?

By splitting job postings into the four broad industry levels, we observe that the rise in WFH jobs is dominated by job postings linked to companies in the services sector. WFH jobs in the "others" sector also rise, but to a smaller extent. The construction sector has no discernible change in WFH jobs, which makes sense, while the data for the manufacturing sector is not shown because of the small number of WFH jobs in the base period.

Splitting job descriptions into WFH jobs and non-WFH jobs also reveals stark differences. WFH jobs tend to involve ICT employees in that these jobs require ICT skills (e.g., cloud technology, amazon web services) and general administration. Non-WFH jobs, on the other hand, require more ad-hoc duties, customer support, and maintenance management. These differences lend credence to differences in the working population that can and cannot WFH.

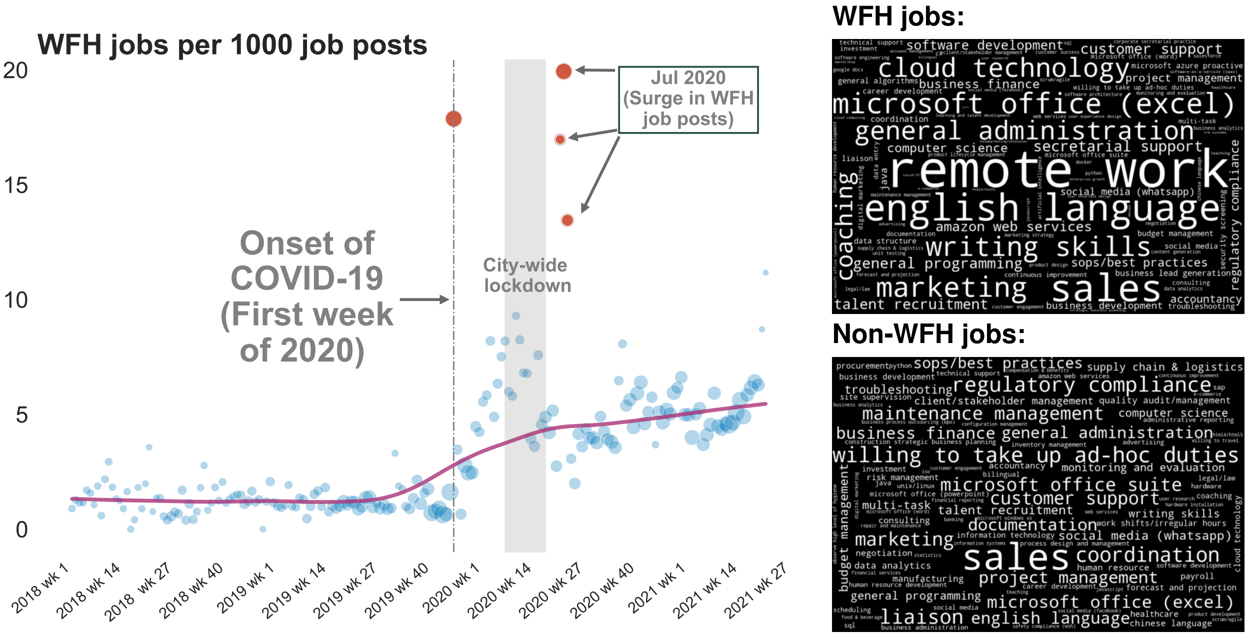

WFH jobs are sensitive to real-world changes and announcements

From the higher-frequency representation of the share of jobs that can WFH, we also observe that hirers are highly responsive to changes in the wider geopolitical landscape. Hirers start posting WFH jobs as early as the last week of December 2020, when the world is only starting to find out about the severity of Covid-19. We also observe three outliers for three weeks in July 2020, just when the city exits the lockdown, but where policymakers warned that working from home is to remain the default [8].

From surveys, the option to work from home is a key factor when employees assess job opportunities, with many preferring the option to do so at least partially [5, 9]. On the other hand, productivity changes remain a fundamental concern in the resistance to letting employees WFH [10]. Our data on online job posts suggest that the proportion of WFH job posts constitutes less than 2 per cent of total job posts, which says something about whether WFH will be the new norm. Nonetheless, we emphasise that a hybrid model, to various degrees, will likely constitute a larger share than what we find here.

Theme 6: Government vs. Private Job Boards

Data for this section

This section focuses on the MyCareersFuture job board, the government-run job board with the second-largest share in job postings, and four other job boards with the most significant job postings share (JobStreet, jobsDB, Jobscentral, and eFinancialCareers). A key part of the findings in this section relates to the "SGUnited jobs". These are jobs under the March 2020 SGUnited Careers programme. This programme is part of a Covid-19 government response initiative that aims to support local workers and fresh graduates affected by the Covid-19 situation to access immediate short-term employment options that potentially lead to longer-term employment opportunities [11, 12]. Our search list to identify SGUnited jobs includes four terms: i) SGUnited, ii) trainee, iii) traineeship, and iv) mid-career.

MyCareersFuture and JobStreet tracks official statistics

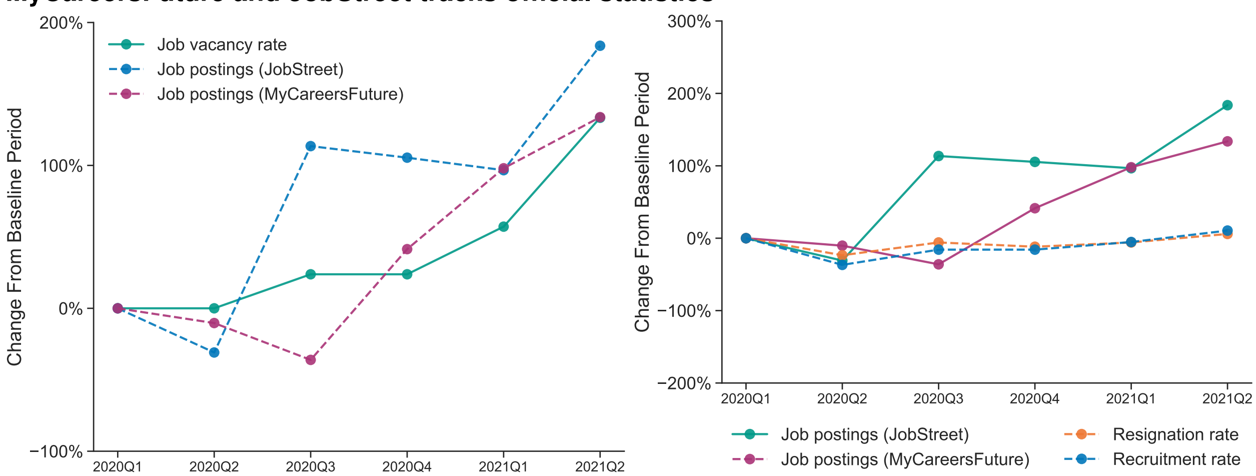

Among the differences between MyCareersFuture and other job boards is the fact that MyCareersFuture is a government-run job board run by the Singapore government, which went online in 2019. All other job boards are privately run. We focus on MyCareersFuture and the privately-run JobStreet, the two largest job boards which combine to hold about 55% of all job posts in the local labour market. MyCareersFuture aims to provide a platform for "Singapore Citizens and Permanent Residents with a fast and smart job search service to match them with relevant jobs, based on their skills and competencies" [13]. JobStreet, on the other hand, is a privately-run job

Share of job postings with salary stated

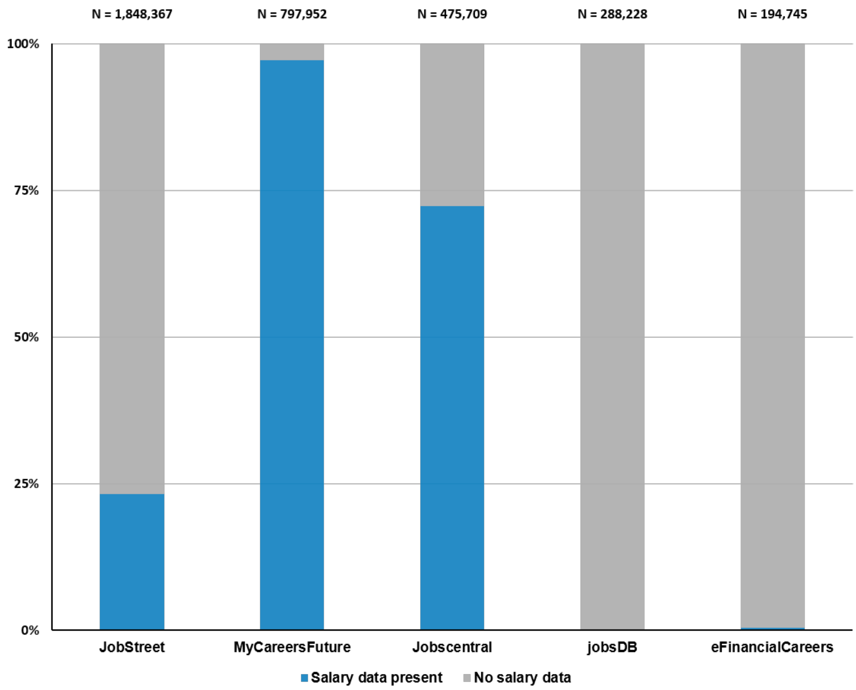

board that posts the largest share of job ads in the labour market. Like in Theme 1: Tracking the Labour Market, we examine how well job postings from MyCareersFuture and JobStreet track the official survey statistics. Our results show that even when we narrow down our data to these two job boards, job postings can still track vacancy rate, recruitment rate, and resignation rate.

MyCareersFuture is differentiated

Using the data from job postings, we find further differences between MyCareersFuture and the other job boards. First, MyCareersFuture is much more likely than any other job board to state the (minimum and maximum) salary in the job post. Almost all MyCareersFuture job posts state the salary. Comparatively, JobsCentral has the second-highest proportion of job postings with salaries stated at 73%. The number of jobs stating salaries drops dramatically in the next three largest job boards, with only 0–25% of job posts mentioning salaries. Salaries are often the most salient concern for people when looking for jobs, or at least the one that is most objectively verifiable. MyCareersFuture therefore has an advantage over other job boards. Additionally, differences in salary disclosure across job boards are of interest since salary disclosure is an essential part of the debate on friction in job search and workers' bargaining power [14]. The fact that MyCareersFuture job posts state salary indicates a move towards reducing friction in the

MyCareersFuture and JobStreet job mix

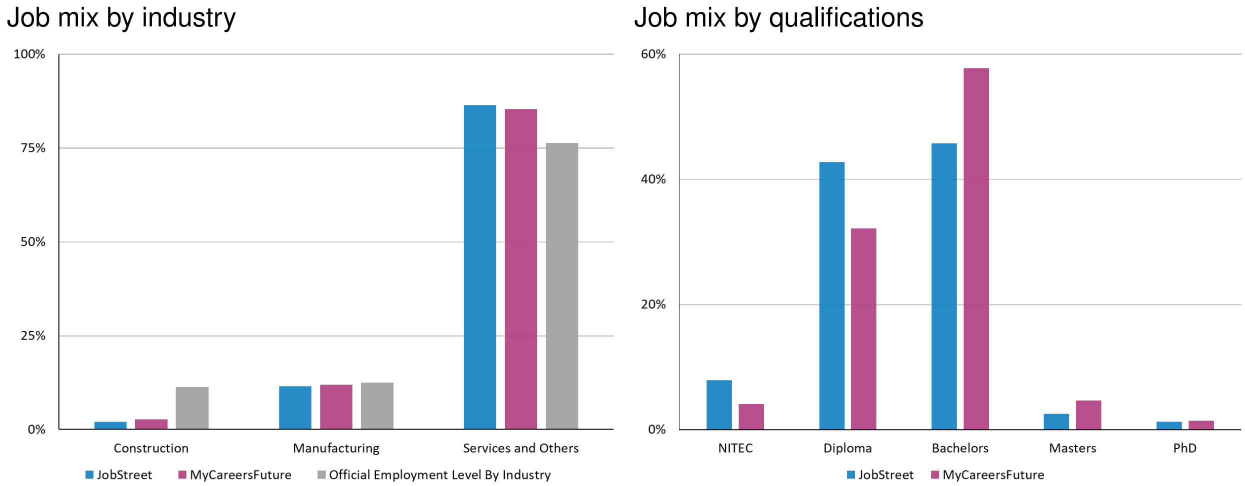

labour market.

Breaking down the mix of job postings by industry in the two job boards confirms that job postings are generally representative of the labour force stock, that is, the Official Employment Level by Industry. We observe that the Services and Others industry is slightly overrepresented online, which further confirms our finding in Theme 1: Tracking the Labour Market. While jobs posted on both boards frequently ask for formal qualifications, we observe that the requirements for education attainment levels on MyCareersFuture tend to be higher. This could be due to differing expectations from employers when advertising for jobs on different online portals and the self-selection of job seekers as well.

Supporting the local workforce through SGUnited jobs

A second observation we pick up is how the government-run MyCareersFuture holds a larger share of SGUnited jobs, which falls under a programme to provide around 18,000 traineeships and jobs where salary is co-funded by both the employer and the government [11, 12]. This programme was announced during the 2020 (Resilience Budget) to support the local workforce (Singaporeans and permanent residents). These traineeships include STEM in research and development labs and tech startups. We see that MyCareersFuture posts the largest share, 80 SGUnited jobs per 1000 job posts at the peak. These are for both SGUnited traineeships and jobs. The largest job board in our sample, JobStreet, only had around 15 SGUnited jobs per 1000 job posts for comparison. Summing SGUnited jobs for all the rest of the job boards confirms that MyCareersFuture is outstanding in its support for the SGUnited jobs initiative.

SGUnited Job Posts by job boards

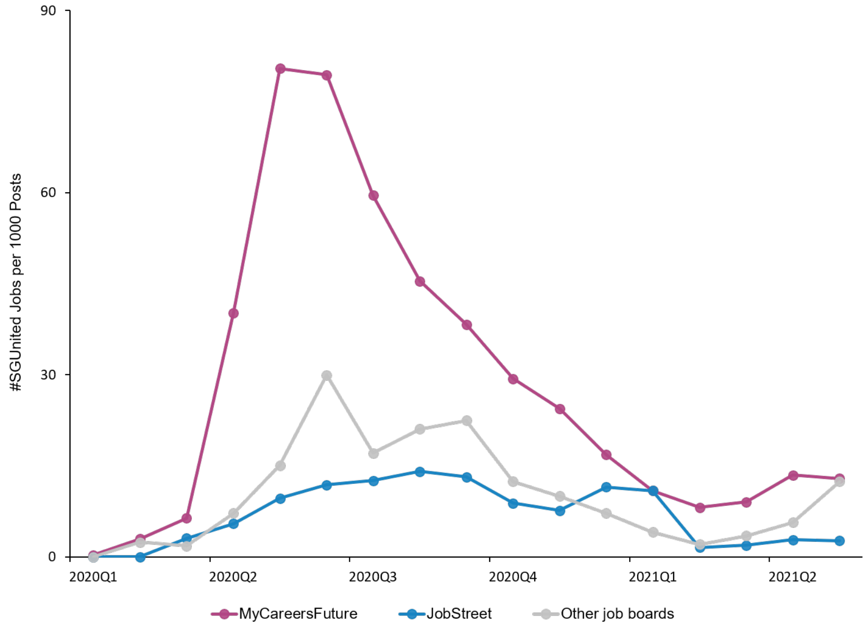

Breaking down the SGUnited jobs traineeships and mid-careers shows that while the initiative laid out 8,000 SGUnited traineeships and 10,000 SGUnited jobs, most of the SGUnited-related job postings we observe are the SGUnited jobs traineeships. Even at its peak in quarter 1 of 2021, only about 5 percent of the SGUnited-related postings are for SGUnited jobs, and most are still for the traineeships. One potential explanation is that since the public sector is taking the lead in creating the 10,000 jobs [11], many of the SGUnited jobs may end up on the job board for government positions which is not in our sample.

SGUnited jobs traineeships and mid-careers

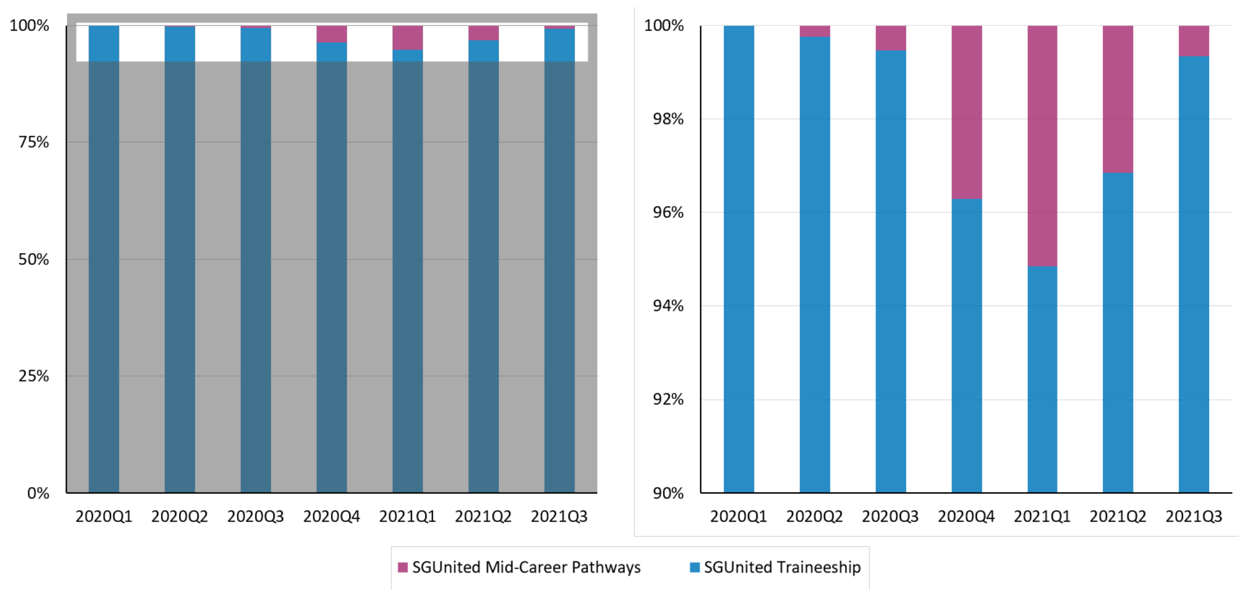

Materials and Data

Online job postings from JobTech

Our primary dataset of job postings is collected by JobTech, a Singapore-based startup. The full dataset contains 4.8 million job postings from over 200 sources online, dating from January 2018 to July 2021. In addition to the expected job description, the postings would also typically contain the name & a short description of the company, and the expected salary range where available. From the job description, JobTech extracts the particular skills required of the employee, which will be elaborated on below under **“weighted word clouds”**.

JobTech also performed deduplication of repeated job postings across time and sources to obtain a more accurate count of actual job vacancies in Singapore. Postings are considered duplicates and removed if posts have a high degree of textual similarity with an earlier posting within a 6 month period. After this process, the deduplicated data set contains 1.8 million unique job posts.

Official statistics

We obtain these two sets of official statistics from the Labour Market Survey under the Manpower Research and Statistics Department at the Ministry of Manpower:

- Job Vacancy Rate by Industry and Occupational Group, Quarterly

  (https://storage.data.gov.sg/job-vacancy-rates-by-industry-and-occupational-group-quarterly/resources/job-vacancy-rate-by-industry-level1-2021-12-16T03-40-12Z.csv)

- Average Monthly Recruitment/Resignation Rates by Industry and Occupational Group, Quarterly

  (https://storage.data.gov.sg/average-monthly-recruitment-resignation-rates-by-industry-and-occupational-group-quarterly/resources/average-monthly-recruitment-resignation-rate-topline-2021-12-16T03-34-30Z.csv)

The data documentation defines job vacancy rate as the total number of job vacancies divided by the total demand for labour at the end of the quarter. The total demand for labour is defined as the sum of the number of employees and job vacancies at the end of the quarter. The average monthly recruitment/resignation rate during a quarter is defined as the average number of persons recruited/resigned in a month during the quarter divided by the average number of employees in the establishment. The official statistics we use cover both public and private sector establishments with at least 25 employees.

These official statistics are available at the "Industry (Level1)" and at the "Broad Occupational Group" level, which we use in our analyses. We detail how we map job postings to the official statistics by industry level 1 below.

Orbis and company industry

To retrieve the 2-digit SSIC (Singapore Standard Industrial Classification) code of job postings in our dataset, we start by querying the *Orbis* database using the company name listed in the job posting. 86,421 (77.0%) of company names and 2,717,044 job postings (62.3%) have returned matches by Orbis. We manually look through the returned results and confirmed that the matched results are reasonable; most matched pairs with very different text occur because the matched company names are current names and while the queried company name are previous names or aliases.

Crosswalking company industry using different classification standards

The metadata returned by Orbis includes the NACE (Statistical Classification of Economic Activities in the European Community) industry classification. We crosswalk from the NACE industry codes to the ISIC (International Standard Industrial Classification of All Economic Activities) code and then from the ISIC to SSIC using official sources.

There are no one-to-many mappings from the NACE to the ISIC. However, one-to-many mappings from the ISIC to the 2-digit SSIC exist. These job postings, where the ISIC-SSIC mapping is one-to-many, are dropped in pipelines involving SSIC. In the end, for the online job postsings data before deduplication, we have 64,968 (74.7%) company names and 1,618,075 (37.1%) job postings successfully mapped to a 2-digit SSIC code.

The broadest industry clustering available under the official national statistics is the industry level 1 ("Industry (Level1)"). Only four industry categories are available: 1) manufacturing, 2) construction, 3) services, 4) others. The first three capture the largest employment size in the economy.

To go from the "Industry Level 1" to the 2-digit SSIC codes, we use the official SSIC 2020 classification report [15] and the Singapore Services Sector report [16].

“Manufacturing” corresponds directly with industry section C for manufacturing which corresponds to 2-digit SSIC codes (industry divisions) 10-32. “Construction” corresponds directly with construction section F, which corresponds to 2-digit SSIC codes 41-43. “Services” is spread across different industry sections, and according to the services sector report, corresponds to the 2-digit SSIC codes 46, 47, 49–53, 55, 56, 58–63, 68, 90–93. We then map all remaining 2-digit SSIC codes to “Others.” We note that the Ministry of Manpower uses an alternative classification of “Others”, which includes “Agriculture, Fishing, Quarrying, Utilities and Sewage and Waste Management.”

Salary data

JobTech extracts the minimum and maximum expected salary from the job posting where available. Some of these numbers do not make sense based on our manual inspection. For example, to circumvent the disclosure of the expected salary range, some companies will enter placeholder values (e.g., $1 as salary). We remove these. Some companies will also advertise openings for jobs based overseas on local job portals, with salaries stated in foreign currencies. For instance, we traced one job posting with a stated salary in 15 million a month to a vacancy for a business manager in Jakarta handling a Singapore account. We believe that the number captures the Indonesian Rupiah rather than Singapore dollars. To deal with the foreign currency inflating salary values and other outliers, we drop job postings for salary-related analyses when the stated salary falls outside the 1st and 99th percentile.

Weighted word clouds

To construct the word clouds, we use the list of skills per job post provided by JobTech. JobTech curates four group of skills:

- Hard skills (e.g., screenwriting, drums, rhino 3d)
- Soft skills (e.g., business savvy, problem solving, mindfulness)

- Requirement skills (e.g., english, class 5 licence, good eyesight)
- Domain skills (e.g., sea trials, machine learning, zoology)

We drop soft skills in constructing the word clouds.

Effectively, skills are phrases of words (e.g., business savvy) and so are amenable to simple text processing. We start by concatenating the lists of non-soft skills by type of job (e.g., a bachelor degree job), and then construct a weighted "phrase" count. The count for each phrase and each group is upweighted by how frequent the phrase appears for a given group of job postings and downweighted by how frequent the phrase appears in all job postings so that phrases that appear frequently and commonly across all job postings are given smaller weights.

In the visualisation of the skills word clouds, both the brightness and size of the skills are coded such that they directly corresponds to its skill frequency.

Change in skills

The skills "friction ridge analysis" and "talent recruitment" have been removed from the list of the top ten largest growth skills due to various reasons arising from data quality. Microsoft Office (Excel) and Microsoft Office (Powerpoint) are skills that appeared separately in the top ten declining skills and have been aggregated to Microsoft Office Suite.

Identifying AI jobs

To identify job postings where the job scope relates to "artificial intelligence" (AI), we perform a simple upper and lower case-insensitive string search for phrases within the text of the job description. As artificial intelligence encompasses a broad spectrum of technologies, they are further grouped into 6 main categories [6]. Search phrases for each category are:

- **Artificial Intelligence:** Expert System, IBM Watson, IPSoft Amelia, Ithink, Virtual Agents, Autonomous Systems, Lidar, OpenCV, Path Planning, Remote Sensing
- **Natural Language Processing (NLP):** ANTLR, Automatic Speech Recognition, Chatbot, Computational Linguistics, Distinguo, Latent Dirichlet Allocation, Latent Semantic Analysis, Lexalytics, Lexical Acquisition, Lexical Semantics, Machine Translation, Modular Audio Recognition Framework (MARF), MoSes, Natural Language Toolkit (NLTK), Nearest Neighbor Algorithm, OpenNLP, Sentiment Analysis/Opinion Mining, Speech Recognition, Text Mining, Text to Speech (TTS), Tokenization, Word2Vec

- **Neural Networks:** Caffe Deep Learning Framework, Convolutional Neural Network (CNN), Deep Learning, Deeplearning4j, Keras, Long Short-Term Memory (LSTM), MXNet, Neural Networks, Pybrain, Recurrent Neural Network (RNN), TensorFlow
- **Machine Learning:** AdaBoost algorithm, Boosting, Chi Square Automatic Interaction Detection (CHAID), Classification Algorithms, Clustering Algorithms, Decision Trees, Dimensionality Reduction, Google Cloud Machine Learning Platform, Gradient boosting, Libsvm, Machine Learning, Madlib, Mahout, Microsoft Cognitive Toolkit, MLPACK, Mlpy, Random Forests, Recommender Systems, Scikit-learn, Semi-Supervised Learning, Supervised Learning, Support Vector Machines (SVM), Semantic Driven Subtractive Clustering Method (SDSCM), Torch, Unsupervised Learning, Vowpal, Xgboost
- **Robotics:** Blue Prism, Electromechanical Systems, Motion Planning, Motoman, Robot Framework, Robotic Systems, Robot Operating System, Robot Programming, Servo Drives, Servo Motors, Simultaneous Localization and Mapping, SLAM
- **Visual Image Recognition:** Computer Vision, Image Processing, Image Recognition, Machine Vision, Object Recognition

Co-occurence of AI clusters in job postings

We use the indicators for presence of keywords related to AI jobs described above to compute a matrix of co-occurrence in the six AI clusters. Values in the matrix cells indicate the number of job ads containing search terms relevant to both the column cluster and the row cluster. To increase visual clarity in the co-occurrence plot, we retain only the top 12 co-occurrence pairs.

Media coverage of AI

We focus on the two largest national news outlet in Singapore. Using the Google advanced search operators, we find that out of 2,000 search results involving AI (e.g. `site:www.channelnewsasia.com ("artificial intelligence" OR "AI")`), 307 of them involves robots or robotics (e.g. `site:www.channelnewsasia.com robot* AND ("artificial intelligence" OR "AI")`) but only 222 involves machine learning. For *The Straits Times* outlet, our findings are similar. Out of 5,880 search results, those involving AI is around twice that of machine learning (838 vs 467).

Identifying work from home (WFH) jobs

To identify job postings relating to WFH ("work from home") jobs, we perform a simple upper and lower case-insensitive string search for the key phrases within the text of the job description. Search phrases are: `work from home; working from home; work from anywhere; working from anywhere; remote work; flexible arrange; hybrid work; telecommut*; and work at home`. A manual validation of a small randomly drawn sample for job postings flagged as WFH jobs suggests that this simple approach has a 100% precision.

Identifying SGUnited jobs

To identify job postings that are introduced under March 2020 Resilience Budget's SGUnited Careers programme, we perform a simple upper and lower case-insensitive string search for phrases within the text of the job description. Search phrases are: `SGUnited; trainee; traineeship; mid-career; mid career`. A manual validation of a small randomly drawn sample for job postings flagged as SGUnited jobs suggests that this simple approach has a 100% precision.

Acknowledgements

We thank JobTech (Tan Jun Huang and Sarah Lim), a leading AI talent intelligence startup and data partner of ACI, who provided the underlying data and made the study possible. We thank our colleagues for helpful comments. We also thank Alyssa Nah Xiao Ting and Jingyuan Zhang for excellent research assistance.

References

[1] Budget Speech 2015 - Section C: Developing Our People.

[2] The Economist. A Real-Time Revolution Will Up-End the Practice of Macroeconomics, 2021.

[3] David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. The Parable of Google Flu: Traps in Big Data Analysis. *Science*, 343(6176):1203–1205, 2014.

[4] Pierre Carbonnelle. PYPL PopularitY of Programming Language, 2022.

[5] Stack Exchange. Stack Overflow Developer Survey Results, 2018.

[6] D Zhang, S Mishra, E Brynjolfsson, J Etchemendy, D Ganguli, B Grosz, T Lyons, J Manyika, JC Niebles, M Sellitto, et al. The AI Index 2021 Annual Report. *Human-Centered AI Institute*, 2021.

[7] Ng Wei Kai. Students Can Now Sign Up for Free Course in AI, Choose Pace of Study and Level of Difficulty. *The Straits Times*, 2022.

[8] Cindy Co. Work-From-Home to Remain as Default; Jobs Support Scheme Extended to July. *Channel News Asia*, Jun 2021.

[9] Firdaus Hamzah. 2 in 5 Workers in Singapore Will Not Accept a Job If Unable to Work Remotely or Have Flexible Work Hours: Survey, 2022.

[10] Lucas Shen. Does Working From Home Work? A Natural Experiment From Lockdowns. *Research Paper #2-2021, Asia Competitiveness Institute Research Paper Series*, 2021.

[11] Ministry of Finance Singapore. Ministerial Statement on Additional Support Measures in Response to COVID-19 Pandemic, 2020.

[12] See Kit Tang. COVID-19 Resilience Budget: Landmark S$48 Billion Package to Tide Singapore Through Unprecedented Crisis. *Channel News Asia*, Mar 2020.

[13] MyCareersFuture. Frequently Asked Questions for MyCareersFuture, 2021.

[14] Zoe B Cullen and Bobak Pakzad-Hurson. Equilibrium Effects of Pay Transparency. Working Paper 28903, National Bureau of Economic Research, June 2021.

[15] Ministry of Trade Department of Statistics and Republic of Singapore Industry. Singapore Standard Industrial Classification 2020, March 2020.

[16] Ministry of Trade Department of Statistics and Republic of Singapore Industry. Singapore Services Sector 2020, May 2022.