

Article

COVID-19 Impacts on the IT Job Market: A Massive Job Ads Analysis

Maura Cerioli [†], Maurizio Leotta ^{*,†} and Filippo Ricca [†]

Dipartimento di Informatica, Bioingegneria, Robotica e Ingegneria dei Sistemi (DIBRIS), Università di Genova, 16146 Genova, Italy

* Correspondence: maurizio.leotta@unige.it

[†] These authors contributed equally to this work.

Abstract: The COVID-19 pandemic has had a significant impact on many economic sectors. The most severe immediate financial effects of COVID-19 include job losses and decreased hiring, and we can expect generalized lower economic growth in the long term. While such phenomena are there for all to see, measuring their scope is complex. In this work, we performed a massive analysis of job postings (ads) taken from LinkUp, a popular job search web platform, to better understand the occupational trends in IT. We analyzed about nine million ads for computer and mathematical experts to measure the impact of the virus on the IT job market. We also extended our investigations to almost 109 million advertisements (about 300 GB of data) for all kinds of positions to overview the effects of COVID-19 on the job market at large. The results show that the COVID-19 crisis hit the job market during the first two quarters of 2020, causing the number of job advertisements to drop across all sectors (except one). Specifically, the IT sector lost between 15% and 48% of the ads, depending on the specific professional figure. Since the last quarter of 2020, the ad numbers have recovered for some sectors, and by the first 2021 quarter, all of them have more job ads than in the last five years. Finally, we used text analysis to understand the trends of interest in teleworking. We found that in the second quarter of 2020, the number of ads explicitly mentioning telework was almost three times the average of the previous quarters.

Keywords: job market; content analysis; job advertisements; IT professions; SARS-CoV-2; coronavirus



Citation: Cerioli, M.; Leotta, M.; Ricca, F. COVID-19 Impacts on the IT Job Market: A Massive Job Ads Analysis. *Electronics* **2023**, *12*, 3339. <https://doi.org/10.3390/electronics12153339>

Academic Editor: Domenico Ursino

Received: 26 June 2023

Revised: 26 July 2023

Accepted: 1 August 2023

Published: 4 August 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The effect of the COVID-19 pandemic on the world economy has been of incredible and unprecedented strength. For example, the Chinese economy shrunk by 6.8% (on an annual basis) in the first quarter of 2020. The economies of the Eurozone have shrunk by 14.8% (annualized rate). The US economy shrunk at a rate of 4.8% (annualized) in the first quarter of 2020 [1]. These values are impressive.

Job loss is the most severe short-term economic consequence of COVID-19. In contrast, lower economic growth will be the expected long-term effect (see the Indian Society of Labor Economics (ISLE) report [2]). For example, the United States had an unemployment rate of 14.7% in April 2020, and researchers estimate that 42% of pandemic-induced layoffs will result in permanent job loss [3].

This study aims to understand and experimentally measure the effects of the COVID-19 pandemic on the IT job market and compare it with other sectors. We believe that this work is valuable to better understand the impact of COVID-19 on the IT job market. Furthermore, it will also spread some light on which professional figures and types of work have experienced the most significant crisis, providing decision makers with the means to help them.

To achieve this goal, we use traditional *data analysis* techniques such as *content analysis* [4,5] and *text mining* [6] applied to a massive set of job advertisements taken from

a popular Web job-search platform, LinkUp [7]. LinkUp is a modern openwork platform. All jobs available on the platform are sourced and updated daily from company websites. In this way, the platform can provide job boards and media websites with real-time access to millions of job ads.

The job posting evaluation method has potential due to the huge data volume in the LinkUp DB (about 300 GB of textual data). In addition, the statistics cover geographically distributed advertising, which helps provide a reliable look into the job situation in general.

In this work, we consider almost 109 million ads to analyze the impact of COVID-19 on the entire job market. Moreover, we focus on more than 9 million ads to understand the specific effects on the IT market. To provide an answer to our research questions, we compare the number of job ads published in the previous years with those observed during the peak months of the pandemic up to the end of the critical period, in this way avoiding seasonal fluctuation interference.

Among the effects of COVID-19 considered in this paper, there is also the impact on remote work. Some companies and organizations were ready to move to fully remote working because, for example, they had adopted that modality for part of their workforce and had the needed infrastructure and policies in place. Thus, they smoothly transitioned to teleworking for all their employees. However, many others may never have experienced remote working and, though in possession of a flexible policy, experienced more troubles. Overall, the teleworking experience imposed by the crisis was satisfactory in most cases. Thus, some big market players, such as Microsoft, Amazon, Facebook, and Google, have allowed remote work. Others, such as Twitter, have allowed it indefinitely [8].

The increased focus on teleworking caused by the pandemic needs is reflected and can be measured by analyzing the number of occurrences of terms related to teleworking in ads. Using text analysis, we have found that the percentage of ads mentioning teleworking (or some of its synonyms) in the post-COVID-19 era is 2.5 times that before. The ratio was even higher for some specific job categories that can be effective from remote locations. For instance, for *software quality assurance engineers and testers*, the ratio is more than three times, and for *software developers for systems software*, it is five times.

The work presented here extends our previous preliminary poster paper [9] in several directions.

- We investigate a more extended period. In the original poster, we only had data up to June 2020. Thus, we had to limit our analysis to the peak of the crisis. In this paper, instead, we work on ads published up to March 2021. In this way, we can see the entire evolution of the market from before the pandemic crisis to the early stages of the subsequent economic recovery.
- We deepen our analysis of the impact of COVID-19 on the entire job market by providing different views to show its health status before and after the restrictions imposed to contrast the pandemic.
- We focus our attention on the IT job market, analyzing the impact of COVID-19 on the specific roles in this field (new research question **RQ2**). Thus, we improve our understanding of the effects of the crisis on this particular sector through a finer granularity analysis. Indeed, by discussing less broad fields, better insights can be obtained.
- We investigate the evolution of interest in working from home (telecommuting) from before the crisis to its end (new research question **RQ3**). We focus on the IT job market for this examination, analyzing this aspect role by role.

The paper is organized as follows. Section 2 describes our content analysis study on the LinkUp data, presenting the goal, the derived research questions, and the dataset. Section 3 shows the study results. Finally, related works are discussed in Section 4, and we draw our conclusions in Section 5.

2. Study Definition

Information about job market trends can provide valuable insights for research studies. Since several years ago, most job positions have been recruited through ads on specialized websites such as LinkedIn, Indeed, and LinkUp. Thus, we evaluated various web portals hosting job advertisements. In particular, we looked at the ones reported as the most relevant by many online guides. Among them, we have selected LinkUp [7] for the following reasons:

- It is a job search website collecting verified up-to-date job listings directly from employer websites (so the information it contains is reliable from both a content and temporal perspective);
- It provides the possibility of having full access to the raw data (a crucial aspect to enable automating the massive analyses carried out in our work);
- It also provides a complete history of the ads archive (fundamental for the temporal analysis performed in this study spanning over many years);
- It is a widely employed portal that contains a massive number of job ads.

For this reason, some years ago, we entered into an agreement with LinkUp to obtain access to their data for research purposes and have proposed a few studies based on these data (see, e.g., [9,10]). In particular, to carry out our current investigation about the impact of COVID-19, we analyzed almost 109 million (108,694,360) ads published on the LinkUp portal in the last five years.

2.1. Goal

The main goal of this study is to understand the effects of the COVID-19 pandemic on the IT job market. This study can be beneficial for IT professionals who want to achieve the following:

- (Benefit B1) Improve their understanding of the phenomenon and the effects of a historical event such as the COVID-19 pandemic on the IT job market compared to other sectors;
- (Benefit B2) Be ready for similar events to minimize their impact;
- (Benefit B3) Quantify the impact of teleworking in technical professions and the possible consequences in the world of work.

Understanding which job positions have been affected more and the reasons behind the experienced impact can help devise countermeasures to make such job positions more resilient to a possible future pandemic. An increase in telecommuting could lead to a redefinition of specific roles to require a smaller amount of physical presence in the office. However, the redefinition could lead to a restructuring of the work organization.

2.2. Research Questions

Considering the goal mentioned above, we defined the following *research questions* (RQs):

RQ1 *How has the number of job advertisements published for each job sector changed over time?*

The first RQ aims to understand the effects of the pandemic on the new positions in the entire job market. We compare the number of job ads published during the last years with those created during the first year of the pandemic, sector by sector (e.g., “Food Preparation and Serving-Related” or “Computer and Mathematical”; see Section 2.3). An alternative would have been to use the active ads at any given time. However, two opposite reasons may require advertisers to close an ad: the job is filled, or the position has been canceled. A short time between publishing and closing is a healthy and quick market symptom in the former case. In the latter, it reflects a loss of job positions in a critical situation. On the contrary, adopting the created, instead of active, job ads as a measure allowed us to avoid any possible ambiguity but take a snapshot of the advertisers’ expectations.

RQ2 *How has the number of IT job advertisements published for each IT sub-sector changed over time?*

The second RQ is similar to the first, but it applies to a more focused and restricted sample, wherein we consider only the IT sector. Thus, the job profiles' granularity increases. The greater level of detail of the *O*NET* classification (see Section 2.3) we use to capture the concept of occupation profile allows us to see if some areas are less affected than others by medical emergencies. Such information, in particular, can help professionals choose their specialization in fields in a way that is safer from the viewpoint of resilience to a pandemic.

RQ3 *How has the pandemic influenced the interest in telecommuting in technological professions?*

To answer that question, we automatically count the occurrences of synonyms for telecommuting, such as teleworking or working from home, and compare their numbers across the quarters to see how much more popular the COVID-19 crisis makes this topic.

Given the nature of this study, which is mainly descriptive (it describes some conditions or factors found in a population in terms of its frequency and impact), to answer all the RQs, we applied quite exclusively descriptive statistics and showed our findings using charts.

While **RQ1** aims to understand the effects of the pandemic on the new positions in the entire job market (including all sectors and therefore also the IT sector), **RQ2** applies to a more focused and restricted sample, i.e., the IT sector.

Concerning the relationship between RQs and benefits, the results of **RQ1** and **RQ2** could help to achieve **B1** and **B2**. Conversely, the results of **RQ3** could be helpful in achieving **B3**.

2.3. Dataset Details

The data source used in this study consists of the complete dataset downloaded from LinkUp. It contains all the job ads published in the last six years (April 2015 to March 2021), listing almost 109 million records (108,694,360). Each job advertisement includes a plurality of relevant data, such as company name, city, state, country, creation and deletion date, full-text description, and *O*NET* code. The Occupational Information Network [11] (*O*NET*) is an online and accessible database, defined under the sponsorship of the US Department of Labor/Employment and Training Administration (USDOL/ETA). It lists numerous occupational definitions to characterize jobs regarding the skills and knowledge required, performed work, and typical work settings. The *O*NET* classification is constantly updated. For instance, novel *O*NET* codes are added to the database to describe new job descriptions. On the contrary, *O*NET* can be dropped if necessary (even if a set of more detailed occupations more often replaces them). Overall, the *O*NET* classification is reasonably fine-grained since it comprises about a thousand codes (the exact number depends on the year considered).

Each *O*NET* code is composed of two parts: the first two digits represent the *family*, i.e., the more general job sector; then, the following six digits describe the specific job profile. For example, the family *Computer and Mathematical* is associated with code 15. Under this family, there are *O*NETs* such as *Business Intelligence Analyst* (code 15-1199.08) and *Web Developers* (code 15-1134.00). The *O*NET* information allowed us to cluster and count the advertisements by their offered job sector for **RQ1** using the *families* and by the specific occupation for **RQ2** and **RQ3** using the complete code.

2.4. Dataset Preprocessing

The job ads made available by LinkUp are primarily written in English, although the jobs were advertised in more than two hundred different countries.

Job data are shown in two separate tables: *job descriptions* and *job records*. The *job description* is the complete ad text, which has an average length of 3500 chars. The *job record* consists of the metadata used for the ad listing, including the title, *O*NET*, and creation date. The information in these job records is sufficient for most of our analyses. The MD5 hash of the job offer *URL* is the primary key for both tables and connects the two parts as an external key from the *job description* to the *job records*.

To help the reader understand the data we manipulate, we present some synthetic data in Table 1 and Figure 1. We had to resort to synthesizing our example, as we cannot use actual entries from the database because of the terms of our agreement with LinkUp. However, to obtain a good-as-real entry, we started with data from advertisements and elaborated on them as follows:

- We used the MD5 of a string not representing an URL as the primary key to avoid conflicts with entries in the database;
- We mixed up parts of different entries to synthesize our example, making the result more representative and less recognizable at the same time;
- We replaced the advertising company's and other organizations' names with fakes;
- We slightly toggled the numeric data for anonymity's sake;
- We used generative artificial intelligence [12] to rephrase the description text beyond recognition.

Table 1. An example of (synthesized) job record data.

CSV File Header	Generic CSV File Line
hash	935dc4041a27c42e68bcf651f5f04f95
title	Automation Testing Engineer
company_id	00000
company_name	TheBestOfAll
city	Marburg
state	MR
zip	35037
country	DEU
created	2017-02-18T19:11:00+00:00
last_checked	2017-04-15T02:11:30+00:00
last_updated	2017-03-18T09:47:00+00:00
delete_date	2017-05-01T00:00:00+00:00
unmapped_location	FALSE
onet_occupation_code	15-1253.00

Though generated by database tables, the data shared by LinkUp are distributed through standard text format files.

Each job record is a line in a CSV file, following the schema in Table 1, where the reader can easily recognize the data parts discussed above. As each entry contains 14 distinct pieces of information, we transposed the lines of the CSV file for readability's sake. Thus, the header is the first column, and the synthesized example is the second column.

The job descriptions, instead, are in XML format. Figure 1 shows our artificial example. The hash element is the same as in the job record part in Table 1, so the data in the two figures refer to one entry. The description element shows the most common parts we find in ads, including a short introduction to the job environment, a brief summary of the job itself, the list of the necessary qualifications to fill the position, and a few technical keywords, to convey the employer's expectations better. The final result is slightly shorter than the average entry, with 3378 characters.

To answer our research questions, we needed to analyze the data, and the file format used to store them was found to be inadequate. Thus, we processed them to extract the relevant pieces of information, storing them in aggregate form when possible and using them to populate an MS-SQL database for finer analyses when courser-grained data would be insufficient to answer our research questions.

```

<job>
  <hash>
    935dc4041a27c42e68bcf651f5f04f95
  </hash>
  <description>
    <![CDATA[
      TheBestOfAll is a financial technology (fintech) company that offers banking solutions to
      the general public. Last year, we experienced the highest growth rate among companies in
      Europe, as recognized by \textit{The Economist}. Our company has garnered support from leading
      technology investors and is currently valued at EUR 5.7 billion. \n \nWe have fostered a
      collaborative and inclusive culture where team members support one another, acquire
      knowledge, and innovate to propel the company and its customers forward. Consequently, we
      have been honored as a "Best Working Place in Northern Europe" for five consecutive
      years. We have also received annual accolades from BestEmployeeReviews for having one of
      the "Best SW Engineering Departments." \n \n Join our team of 1350 talented and dedicated
      professionals if you are interested in tackling significant challenges and making a
      meaningful impact on people's lives. \n \nAbout the Role: \n \nOur quality assurance (QA)
      team assumes end-to-end responsibility for building a high-quality product. They employ
      automation and functional and security testing to validate a wide range of products,
      ensuring a seamless banking experience for our customers. As a QA automation developer,
      you will collaborate directly with our product teams, participating in application
      specification reviews and architectural design sessions from a quality perspective. You
      will also devise testing strategies for various features and functionalities. Within your
      first year, you will gain an understanding of how to deliver high-performing software at
      a rapid pace of development. \n \nOur core technologies include: \n \nC#, C++, SQL,
      Selenium WebDriver, TestNG, Git, Maven, Jenkins, ArgoCD, REST Assured, SumoLogic,
      Kubernetes, Docker, and Gatling. \n \nResponsibilities: \n \nConduct functional, performance,
      and security tests by deploying disposable environments on demand, incorporating
      microservices with different versions and configurations and customizing database
      content. \nDevelop test automation in tandem with feature development, utilizing our
      functional and performance test frameworks. \nCustomize, configure, and patch services
      and deploy infrastructure and code changes up to the production stage. \nAnalyze log
      files and troubleshoot production issues by automating and reproducing them in a
      non-production environment. \nPlan, design, and implement strategies and solutions to
      validate features by test automation. \nQualifications: \n \nProficiency in developing
      test automation using C#/C++ and Selenium, with experience utilizing integrated
      development environments (IDEs) such as Visual Studio Code or Eclipse. \nExperience
      deploying services in diverse environments, such as Kubernetes, and employing deployment
      tools such as ArgoCD and Jenkins. \nCapability of enhancing the efficiency of existing test
      infrastructure using Docker, Selenium, and headless browser executions. \nSome experience
      with log aggregators such as Kibana or SumoLogic, as well as application monitoring tools
      such as Better Stack, Dynatrace, or Sentry.io. \nComfortable working in an Agile
      methodology, such as, for instance, Scrum. \nWe are an equal opportunity employer and value
      diversity within our company. We do not discriminate on the basis of gender, sexual
      orientation, marital status, age, religion, race, color, national origin, veteran status,
      or disability status.
    ]]>
  </description>
</job>

```

Figure 1. An example of (synthesized) job description data.

Figure 2 illustrates the details of our data processing methodology. We implemented it in C#, which provides libraries to interact with CSV and XML files and has an extremely efficient and natural way to work on relational databases in its *entity framework*. On the left-hand side, we see the first part of the process, which entailed the analyses of the *job records* from the CSV files. We worked sequentially on each CSV file by processing one of their lines at a time. For each record, using the *O*NET* to classify the ads by their job sector and the creation date, we counted the number of ads published each month for each job sector and stored the results in a temporary key-value table (a dictionary in C#) to be saved in an Excel file at the end of the process. Such data are sufficient to answer our first two research questions. Moreover, Excel provides appropriate analysis tools to further cluster the data by coarser-grained job sector classification or by quarters and graphically present the aggregation results.

To answer our last research question, however, we needed to investigate the textual contents of each ad. Thus, we saved the record to be analyzed in an MS-SQL database to support more sophisticated queries. Since the last research question is limited to the IT job sectors, we only saved those records with an *O*NET* in that category, and thus did not waste space.

When we completed our work on all CSV files, the process on the left-hand part of Figure 2, we had two artifacts: an Excel file containing the number of ads for each *O*NET*, each year, and each month, and a database with an entry for each IT-related job advertisement. Such entries only contained information from the *job record* files. We used

the process presented on the right-hand side of Figure 2 to complete them by adding their description.

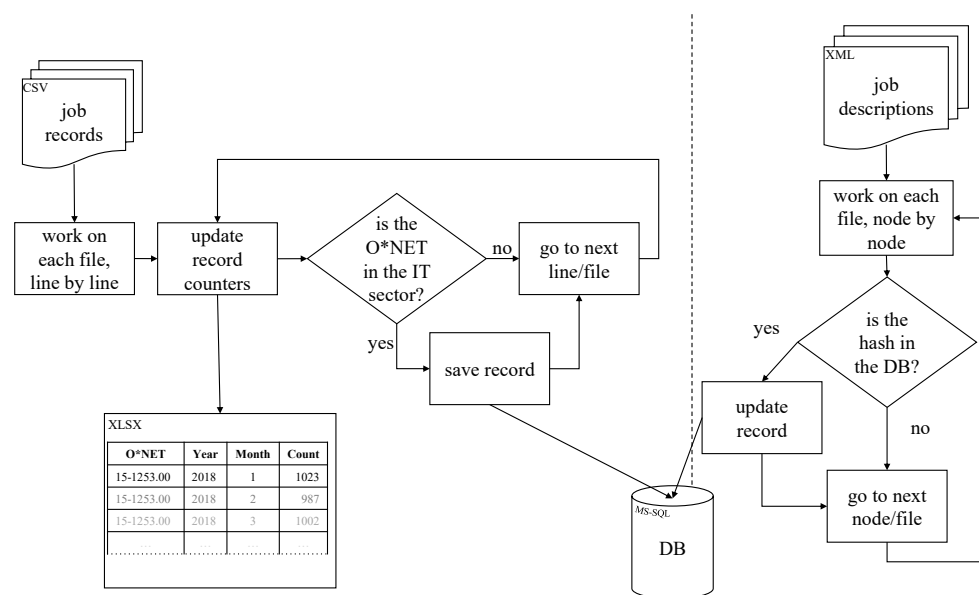


Figure 2. Data processing: DB population methodology.

In this phase, we started with the *job description* files in XML format, and for each node representing a job description, we verified whether it referred to an entry in the database. In that case, we used the node's description part to update the corresponding database entry; otherwise, we discarded the node. Since the hash was a foreign key from *job descriptions* to *job records*, and we initialized the database with all and only the *job records* in the IT sector, our process corresponded to saving the descriptions of job ads in the IT sector. At the end of this update, we dropped the very few *job records* that were without an associated *job description*.

Finally, we used a *full text indexing* method (FT-index in the following) of the columns containing ad titles and descriptions to obtain faster and more efficient queries. FT-index is a standard feature of MS-SQL servers that builds an index associating each word with the position(s) where it appears in each cell. Roughly speaking, creating a *full text index* means that, for each row in the table and each column to be indexed (in our case, Title and Description):

1. The content of that cell is tokenised in words, breaking the strings where spaces, tabs, ends of lines, punctuation symbols, numbers, and such occur;
2. The *stop words*, that is, those words too short or too common to be significant, such as articles or auxiliary verbs, are dropped;
3. The index is updated, adding references to the row, the column, and the positions of all occurrences of each remaining word.

Though this process is expensive, it makes subsequent queries very fast, and we had to perform it just once.

Using an FT-index, we can look for the following:

- (a) A word or a phrase;
- (b) A group of words all appearing in the text within a given distance, preserving or not their order;
- (c) Any *inflection* of a word (singular and plural for nouns, or any verbal forms for verbs).

Finally, the clauses of such queries can be logically combined by the usual Boolean connectives *and*, *or*, and *not*, thus obtaining a powerful and expressive language for textual searches.

3. Results

For each RQ, this section reports the procedure followed to define the answer and the obtained results.

3.1. RQ1: Evaluating the Interest for O*NET Families over Time

To answer the first research question (RQ1), we grouped the job advertisements by the families of their O*NET (which represents the job sector) and their creation month and year.

To complete this first analysis, we considered the data from April 2015 to March 2021. The total number of advertisements created in those six years is 108,694,360 and is distributed over all the 23 families grouping the O*NET codes. The various families have a very different total number of ads, ranging from 42,693 to 15,719,970.

We decided to include only families with at least a million ads in our analysis (see Table 2); in this way, we minimized statistic fluctuations.

Table 2. Families with at least a million ads (above the line) and the remaining families not considered in this study (below the line).

Family	# of ads	% of ads
Sales and Related	15,719,970	14.46%
Office and Administrative Support	13,224,261	12.17%
Healthcare Practitioners and Technical Management	11,510,420	10.59%
Computer and Mathematical	10,501,101	9.66%
Food Preparation and Serving Related	10,224,150	9.41%
Business and Financial Operations	9,190,981	8.46%
Transportation and Material Moving	7,806,985	7.18%
Healthcare Support	4,548,740	4.18%
Architecture and Engineering	4,003,897	3.68%
Installation, Maintenance, and Repair	3,186,812	2.93%
Education, Training, and Library	3,007,151	2.77%
Production	2,369,824	2.18%
Community and Social Service	2,069,859	1.90%
Protective Service	1,915,928	1.76%
Arts, Design, Entertainment, Sports, and Media	1,903,146	1.75%
Building and Grounds Cleaning and Maintenance	1,872,074	1.72%
Personal Care and Service	1,719,363	1.58%
Life, Physical, and Social Science	1,432,409	1.32%
Construction and Extraction	1,412,351	1.30%
Construction and Extraction	679,832	0.63%
Legal	294,397	0.27%
Farming, Fishing, and Forestry	58,016	0.05%
Military Specific	42,693	0.04%

Figure 3 shows the number of ads created each month for each O*NET family. We report the families in decreasing order from bottom to top with respect to their number of ads. At a glance, we can see that all families were better off than in April 2015 from the second quarter of 2017 to the first quarter of 2020.

Analyzing the overall trends, it is also possible to recognize both yearly patterns, such as the local minimum in December, and market crises, such as, for example, the after-effects of the Chinese stock market crash, which is visible at the end of 2015 and the beginning of 2016. We decided to highlight the yearly patterns, adopting different colors for the dates on the horizontal axis, using black for December, gray for June, and red for April (representing the minimum on the COVID-19 crisis). However, it is still challenging to capture the big picture. Thus, in Figure 4, we provide a heat map of the number of advertisements. Each row shows the data of a family with a red–yellow–green background from the minimum to the maximum for that family.

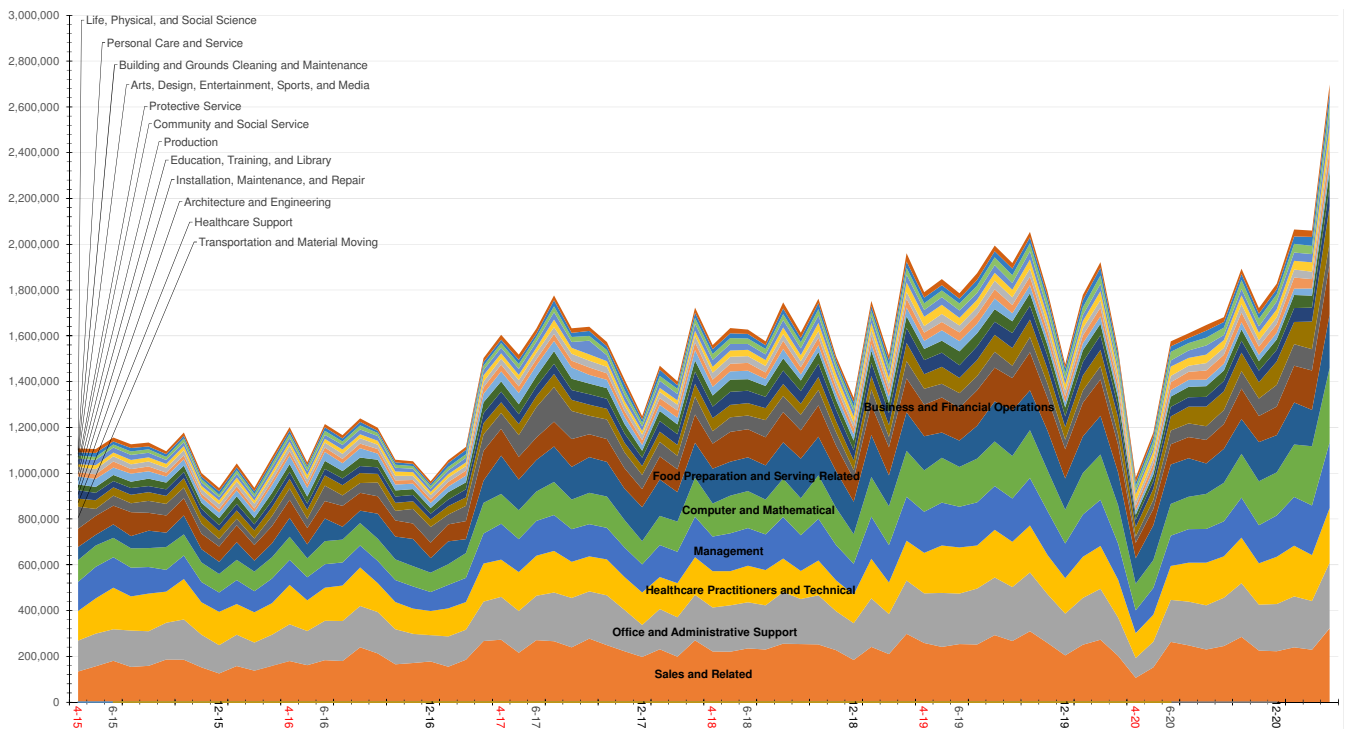


Figure 3. RQ1: number of ads by O*NET family in the years 2015–2021.

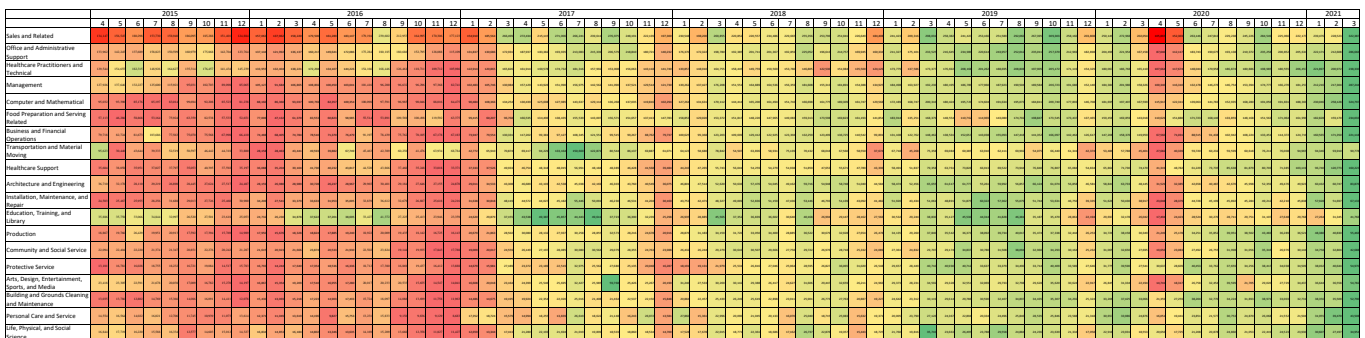


Figure 4. RQ1: heat map of ads number by O*NET family (from April 2015 to March 2021).

Though the values in the individual cells are hard to read, the colors provide a high-level view of the job offer trends and allow the reader to see that the overall trend in the last few years was positive for most families, changing from red to green from left to right until March 2020.

Moreover, we can spot seasonal trends by comparing the color changes for the same month columns in different years. Perhaps the most evident is the December columns, which are more red than the adjacent ones. This phenomenon is known as the *holiday hiring freeze*, and though not all sources agree, there is a large consensus about its existence. For instance, Przystanski in [13] presented data similar to ours, showing a monthly fluctuation with a significant minimum in December. The reasons for a smaller number of positions offered in December are debatable (see, e.g., [14] for an analysis of this phenomenon). The most widely accepted is a minor interest in hiring on the part of the prospective employers busy dealing with the end of the financial year and a smaller number of candidates, possibly caused by the scarcity of open positions, which interact in a feedback loop. An analysis of seasonal fluctuations in the job market is beyond the scope of our research. However, we must be aware that they exist and may hide global trends behind the cyclic changes caused by yearly patterns when comparing each month to the next.

Clustering the data year by year, a more straightforward pattern emerges. In Figure 5, we present a course-grained heat map of ad numbers by O*NET family before the COVID-19 crisis in April 2020. Each cell contains the number of ads in twelve months, from April to March of the following year, because our dataset covers a period starting in March 2015 and the COVID-19 crisis began in April 2020. Out of the 19 families, all but three have their maximum in the last column, and 13 (including the IT sector, *Computer and Mathematical*) present a positive trend in the previous (at least) three years. Thus, the COVID-19 crisis hits a healthy job market for all families but three (*Transportation and Material Moving, Education, Training, and Library, and Arts, Design, Entertainment, Sports, and Media*). Indeed, for *Food Preparation and Serving-Related*, we regard the inversion between the last two columns as a minor glitch, being on the order of 2.26%. Moreover, it could be due, at least in part, to bad performance in the first 2020 quarter when many people started to be aware of risks related to spending time in crowded places and deserted restaurants on a voluntary basis.

	From 4-15 to 3-16	From 4-16 to 3-17	From 4-17 to 3-18	From 4-18 to 3-19	From 4-19 to 3-20
Sales and Related	1,884,573	2,275,082	2,907,615	2,826,480	3,060,455
Office and Administrative Support	1,721,906	1,849,545	2,287,490	2,374,976	2,650,073
Healthcare Practitioners and Technical	1,783,091	1,643,714	1,892,983	1,787,243	2,244,093
Management	1,314,893	1,222,249	1,735,334	1,986,115	2,208,194
Computer and Mathematical	1,068,029	1,173,795	1,581,293	1,973,200	2,203,317
Food Preparation and Serving Related	732,944	1,002,190	1,772,681	1,890,410	1,847,694
Business and Financial Operations	928,087	960,055	1,234,083	1,490,678	1,738,800
Transportation and Material Moving	536,778	659,082	1,069,334	742,751	706,387
Healthcare Support	443,070	478,054	571,382	652,571	833,295
Architecture and Engineering	345,308	362,005	533,159	656,940	696,860
Installation, Maintenance, and Repair	344,114	390,017	538,003	582,507	605,962
Education, Training, and Library	361,083	379,914	465,871	392,076	422,809
Production	223,624	236,369	345,863	383,551	424,500
Community and Social Service	258,299	250,958	330,215	335,278	378,452
Protective Service	197,679	212,773	276,068	353,228	417,412
Arts, Design, Entertainment, Sports, and Media	225,718	216,561	353,745	347,659	355,641
Building and Grounds Cleaning and Maintenance	170,141	183,302	261,501	297,301	390,049
Personal Care and Service	158,277	157,299	236,490	253,000	306,743
Life, Physical, and Social Science	187,095	163,849	233,903	257,437	280,057

Figure 5. RQ1: Heat map of ad numbers by O*NET family clustered in 12-month ranges (from April 2015 to March 2020).

We placed the beginning of the COVID-19 crisis for the job market in spring 2020 because, though the first documented case of COVID-19 happened in December 2019, the virus did not spread around the world until March 2020. Moreover, the most significant economic impact took place in the second and third 2020 quarters, with many countries applying restrictions to travel and gatherings in an uncoordinated way, following strategies dictated by local needs. By the fourth 2020 quarter, the economy started to recover, with traveling normalizing and people circulating freely.

Since the most evident economic impacts of the pandemic lasted about a year, we needed a finer granularity to understand what happened. Therefore, we compared the trends of the different families quarter by quarter starting from 2015. Figure 6 shows the comparison results. The data used for each family are the percentage of ads for that family in that quarter over the total ad number (for that family). In this way, it was possible to compare different families' trends, even if their absolute numbers of ads were on different scales. The percentages highlight the importance, for a given family, of a specific quarter. Whenever the percentage is over 4.2% (represented by the horizontal line in each diagram), the ads created in that particular quarter are above the average, as we have six years' data, representing 24 quarters in total.




Figure 6. RQ1: Percentage distribution of ads by O*NET family in the quarters of 2015 to 2021 (from Q2 2015 to Q1 2021).

At the top of Figure 6, we started with the second-quarter data (in different shades of red). The shorter 2020 columns for all families but *Food Preparation and Serving-Related* stand out and clearly indicate that the job market was hit hard. The situation is slightly better for the third quarters (in different shades of blue), with three families (*Arts, Design, Entertainment, Sports, and Media*, *Healthcare Support*, and *Protective Service*) performing better than the previous year. Moreover, many gaps are smaller, hinting at more minor problems. The last quarters (in different shades of green) show further improvements, with 12 out of 19 families improving their performances with respect to 2019. Finally, the bottom graphs show that the first 2021 quarter was better than the 2020 quarter for all families and outperformed all previous years for all families except *Education, Training, and Library*.

We considered two aspects to analyze how the job market coped with the COVID-19 challenge more in-depth. First, the length of the positive trends interrupted by the pandemic was interesting as a refinement of the rough-grained analysis in Figure 5. Indeed, it measures the health status of the job market for that family. The second aspect worthy of notice was the improvement/deterioration intensity because gap heights shed some light on the relevance of the phenomena.

Both aspects can be partially investigated using Figure 6. However, comparing the columns within the individual graphs is a challenging task. Hence, to make understanding the positive trend lengths easier, in Figure 7, we counted how many columns on the left were shorter than the one of a given quarter, starting from the first quarter of 2019, one year before COVID-19 hit. Thus, any cell contains the number of years in a positive trend up to that quarter. In other words, its value measures the length of the positive span in years. Zero corresponds to a quarter with a lower percentage than the previous year’s same quarter, where the growth stopped. We highlighted such negative points in red. Vice versa, positive numbers represent positive trends, and the higher the value, the longer the trend and the deeper the green highlighting them.

	2019				2020				2021
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1
Sales and Related	3	1	1	1	0	0	0	0	1
Office and Administrative Support	3	4	1	3	0	0	0	0	1
Healthcare Practitioners and Technical	3	1	1	1	4	0	0	2	5
Management	3	3	3	3	4	0	0	0	5
Computer and Mathematical	3	4	4	4	4	0	0	5	5
Food Preparation and Serving Related	3	0	4	4	0	1	0	5	1
Business and Financial Operations	3	3	3	3	4	0	0	0	5
Transportation and Material Moving	0	1	0	1	0	0	0	2	1
Healthcare Support	3	4	1	3	4	0	2	4	5
Architecture and Engineering	3	3	4	4	0	0	0	0	1
Installation, Maintenance, and Repair	3	4	1	3	0	0	0	0	1
Education, Training, and Library	0	1	1	1	0	0	0	2	1
Production	3	3	3	4	4	0	0	5	5
Community and Social Service	0	3	1	3	1	0	0	4	2
Protective Service	3	4	1	4	4	0	2	5	5
Arts, Design, Entertainment, Sports, and Media	3	3	1	3	0	0	2	4	1
Building and Grounds Cleaning and Maintenance	3	4	4	3	4	0	0	4	5
Personal Care and Service	3	3	1	3	4	0	0	0	5
Life, Physical, and Social Science	2	1	3	3	0	0	0	4	1

Figure 7. RQ1: The length in years of positive trends quarter by quarter by O*NET family.

We can see that the 2019 quarters are almost all in (some shade of) green, and, in particular, the data for the fourth quarter (column Q4 of 2019) show a positive trend for all families. Thus, we would expect the 2020 data to follow the current positive trend. However, in the first 2020 quarter, nine families saw their number of ads decrease. Two of them, *Transportation and Material Moving* and *Education, Training, and Library*, were already in a bumpy trend. However, the others (*Sales and Related*, *Office and Administrative Support*; *Food Preparation and Serving-Related*; *Architecture and Engineering*; *Installation, Maintenance, and Repair*; *Arts, Design, Entertainment, Sports, and Media*; and *Life, Physical, and Social Science*) were all steadily improving until that moment. The second quarter saw a relevant fall for all families except for *Food Preparation and Serving-Related*. This particular family gained from 4.07% to 4.72%, a value that is still below the 4.87% from the second quarter of 2018. Furthermore, in the third quarter, *Food Preparation and Serving-Related* fell back from 5.62% to 4.90%, confirming that the apparent improvement of the second quarter was only due to bad performance in 2019 versus 2018 and 2017. In the third quarter, *Healthcare Support* gained from 5.50% to 6.03%, as expected given the relevance of healthcare support during the pandemic. *Protective Service* and *Arts, Design, Entertainment, Sports, and Media* improved marginally (from 5.34 to 5.55 and from 4.92 to 5.01, respectively). Finally, the last 2020 quarter showed a more extensive improvement, with most families back in a positive trend. The first 2021 quarter confirmed the positive change, with all families showing a positive trend.

To precisely quantify the gains and losses of ads for different families, tackling the second aspect, that of the phenomenon intensity, Figure 8 presents the difference between the advertisements of a quarter, say N , and the same quarter of the previous year, say N_{old} , normalized by N_{old} to obtain data significantly comparable across rows and columns. Thus, the value in a cell of Figure 8 represents the percentage of gain/loss in ads number for a given quarter and family with respect to the same quarter of the previous year. Therefore, Figure 7 measures the trend lengths, while Figure 8 shows the intensity of the phenomena. We can think of the latter as the derivative of the former.

	2019				2020				2021
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1
Sales and Related	7.18%	11.4%	9.8%	16.3%	-3.2%	-30.8%	-10.9%	-5.1%	9.0%
Office and Administrative Support	13.24%	12.9%	19.0%	19.9%	-4.2%	-43.0%	-18.7%	-1.8%	21.5%
Healthcare Practitioners and Technical	6.88%	25.3%	40.7%	28.9%	9.7%	-36.3%	-9.5%	10.4%	24.3%
Management	19.38%	13.3%	15.1%	16.8%	0.7%	-36.0%	-21.4%	-5.1%	31.6%
Computer and Mathematical	26.16%	16.9%	19.0%	11.2%	1.0%	-31.5%	-19.0%	6.9%	48.7%
Food Preparation and Serving Related	10.49%	-16.5%	7.2%	2.4%	-3.1%	16.0%	-12.9%	0.0%	22.0%
Business and Financial Operations	24.49%	19.9%	17.1%	25.4%	5.8%	-47.7%	-31.8%	-14.4%	30.0%
Transportation and Material Moving	-6.60%	9.0%	-12.7%	3.9%	-16.8%	-33.9%	-3.1%	45.3%	78.1%
Healthcare Support	28.82%	24.3%	45.5%	34.0%	11.6%	-26.0%	9.5%	29.7%	78.4%
Architecture and Engineering	20.18%	10.2%	11.7%	6.6%	-3.3%	-40.3%	-29.4%	-7.8%	20.1%
Installation, Maintenance, and Repair	8.82%	5.3%	11.4%	2.9%	-3.7%	-40.3%	-16.9%	-1.8%	25.4%
Education, Training, and Library	-2.74%	12.8%	26.2%	1.9%	-10.8%	-43.6%	-28.1%	2.3%	14.2%
Production	7.45%	10.9%	19.9%	10.3%	2.0%	-26.7%	-0.8%	21.1%	40.1%
Community and Social Service	-0.07%	4.4%	21.4%	14.8%	11.6%	-25.9%	-12.4%	3.1%	20.6%
Protective Service	66.63%	45.2%	24.8%	8.6%	0.1%	-13.5%	4.0%	5.2%	35.1%
Arts, Design, Entertainment, Sports, and Media	3.54%	9.3%	7.6%	6.4%	-12.8%	-36.3%	1.7%	2.3%	59.2%
Building and Grounds Cleaning and Maintenance	13.51%	20.4%	32.5%	39.8%	32.8%	-3.5%	-2.5%	10.3%	22.1%
Life, Physical, and Social Science	7.10%	10.3%	31.8%	16.0%	26.2%	-12.2%	-2.1%	-2.3%	28.4%
Personal Care and Service	35.12%	23.6%	17.8%	15.9%	-15.4%	-19.4%	-15.1%	14.3%	41.8%

Figure 8. RQ1: The relative gaps quarter by quarter by O*NET family.

The first four data columns refer to 2019, so before COVID-19 was identified. Thus, they provide a baseline for the analysis of what happened when COVID-19 hit. Though most values are positive, it is interesting to note their marked variability. The extensive range of values in the same row shows that many factors may significantly influence a family's performance. The same factors also have an impact in the following years, combining their effects with those of COVID-19. Thus, analyzing a single value might be imprecise and carry some bias. However, when similar changes apply to several cells in the same row or column, we can infer that they are caused by one new factor and are not due to the synergy of many random effects.

The heat map stresses the worse 2020 performance. Indeed, 67 out of 76 cells have a smaller value in 2020 than the corresponding cells for 2019. Seven of the remaining nine cases can be explained as families starting a positive trend at the end of 2018 or beginning of 2019 (this is the case for the first quarter and the *Healthcare Practitioners and Technical*; *Community and Social Service*; *Building and Grounds Cleaning and Maintenance*; and *Life, Physical, and Social Science*) sectors in crisis (third and fourth *Transportation and Material Moving* quarters), or, simply, minor, immaterial differences (fourth *Education, Training, and Library* quarter). The last two cases are more related to COVID-19 and show different situations. The fourth *production* quarter in 2020 marks the start of a *bouncing-back* trend. With the end of the economic crisis in sight, the productivity sector improved dramatically in the last 2020 quarter and even more in the first 2021 quarter, which outperformed the previous five years for all families, as shown in Figure 6. *Transportation and Material Moving* saw one of the more substantial recoveries in the last quarter of 2020 and the first quarter of 2021. The regained freedom of movement and the need to move products around partially justify the significant gains, but the effects are amplified by the problematic situation of the sector in the corresponding quarters one year before.

Another sector that did exceedingly well is *Healthcare Practitioners and Technical*. Given the relevance of the healthcare system during the pandemic and the steady progress shown in the first quarters since 2015, this result is not surprising.

Finally, *Food Preparation and Serving-Related* shows a 16% improvement and thus is the only family associated with a positive value in the second 2020 quarter. However, as we have discussed previously, such a result is probably due to its bad 2019 performance combined with the substantial increment of home delivery services [15].

Let us now focus on the losses in the second 2020 quarter, the most dramatic time of the pandemic, when half of the industrial world went through lock-downs that were more or less dire. The negative gap sizes show that all the less affected families are related to work positions that much needed in a pandemic. Indeed, they are *Building and Grounds Cleaning and Maintenance*, *Personal Care and Service*, and *Protective Services*, all of which are

needed to keep shared spaces hygienic, take care of infected and not auto-sufficient people, and regulate accesses to public places, workplaces, and shops. The majority of the families taking the worst blows are associated with segments of productivity clearly affected by a medical crisis, such as, for instance, *Education, Training, and Library*; *Business and Financial Operations*; and *Architecture and Engineering*. However, the relevant losses of *Healthcare Support* (−26%) and *Healthcare Practitioners and Technical* (−36.3%) are not expected. They can be only partially justified by the increase in those families seen in the first quarter. The bad performances of *Computer and Mathematical* are quite unexpected as well because working online became much more relevant at the time, and that should have increased the need for computer and technical support. We will discuss this aspect in greater detail in the next section.

Summary. In conclusion, the COVID-19 crisis affected the job market in 2020, especially in the second quarter. The effects in the first quarter are mostly limited to a loss of improvement. However, the *Sales and Related*; *Office and Administrative Support*; *Food Preparation and Serving Related*; *Architecture and Engineering*; *Installation, Maintenance, and Repair*; *Education, Training, and Library*; *Arts, Design, Entertainment, Sports, and Media*; and *Life, Physical, and Social Science* families registered a smaller number of ads than in the first quarter 2019. In the second quarter, all families but one lost jobs in significant percentages. The recovery started from the third quarter for a few families and steadily gained momentum, affecting more and more families. The first 2021 quarter shows encouraging data and better performance than in the last five years for almost all families.

3.2. RQ2: Evaluating the Interest for O*NET in the Computer and Mathematical Family over Time

Analogously to what we did for **RQ1**, to answer **RQ2**, we clustered the job ads by their O*NET and their creation month and year and considered the data for the last six years from April 2015 to March 2021. The overall number of ads created in that interval for the O*NETs in the *Computer and Mathematical* family is 9,486,182, distributed over 35 O*NETs, with a very different share of the ads, ranging from 164 for *Mathematical Technicians* to 2,552,673 for *Software Developers, Applications*. To minimize statistic fluctuations, we decided to analyze only O*NETs with at least 80,000 ads; see Table 3.

Table 3. O*NET from the *Computer and Mathematical* family with at least 80,000 ads (above the line) and the remaining O*NET families not considered in this study (below the line).

O*NET	# of ads
Software Developers, Applications	2,552,673
Computer User Support Specialists	995,617
Network and Computer Systems Administrators	949,961
Computer Systems Analysts	814,503
Information Technology Project Managers	784,418
Computer Systems Engineers/Architects	609,669
Information Security Analysts	577,922
Software Quality Assurance Engineers and Testers	439,038
Web Developers	397,660
Software Developers, Systems Software	352,916
Operations Research Analysts	237,479
Computer and Information Research Scientists	147,354
Computer Programmers	144,298
Database Administrators	110,879
Computer Network Architects	83,547

Table 3. Cont.

O*NET	# of ads
Informatics Nurse Specialists	51,468
Business Intelligence Analysts	48,201
Actuaries	43,744
Clinical Data Managers	40,235
Search Marketing Strategists	27,625
Statisticians	26,518
Computer Network Support Specialists	14,307
Geographic Information Systems Technicians	9906
Web Administrators	6189
Document Management Specialists	4357
Telecommunications Engineering Specialists	4179
Data Warehousing Specialists	2715
Database Architects	2058
Geospatial Information Scientists and Technologists	2034
Video Game Designers	1494
Biostatisticians	1194
Mathematicians	914
Computer Occupations, All Other	691
Mathematical Science Occupations, All Other	255
Mathematical Technicians	164

Figure 9 shows the number of ads created monthly for each O*NET in decreasing order of ad number from bottom to top.




Figure 9. RQ2: number of ads by O*NET for the Computer and Mathematical family in the years 2015–2021.

The Computer and Mathematical O*NET family behaves similarly to the other families; they all generally exhibit positive trends up to the first 2020 quarter and have yearly patterns; see, for instance, Figure 10, where we provide a heat map of the number of ads. Each row shows the data of an O*NET with backgrounds spanning from red to green from the minimum to the maximum of that O*NET. As for Figure 4, the goal of Figure 10 is to stress the yearly patterns and the positive general trend for the readers, not to provide the data.

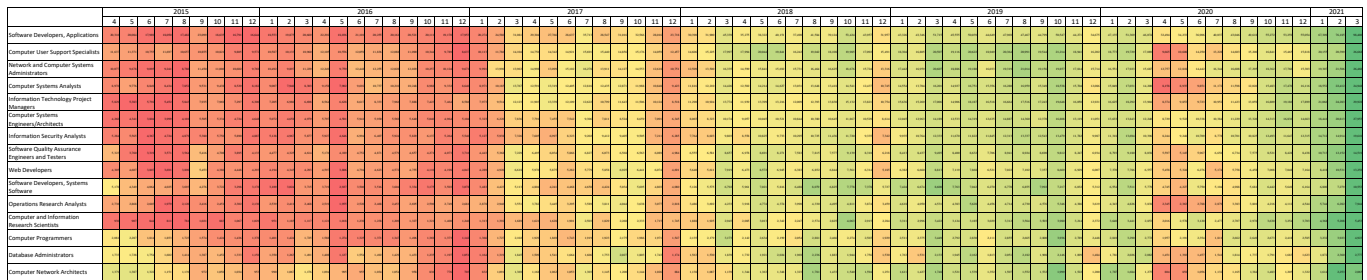


Figure 10. RQ2: Heat map of ad number by O*NET in the quarters of 2015–2021 (from Q2 2015 to Q1 2021).

Comparing the heat map for families in Figure 4 to that for each O*NET in Figure 10, we may notice that the color distribution in the latter is smoother. Indeed, clustering the data by year from April to next March, in Figure 11, we see that every O*NET except for Software Developers, Systems Software were steadily gaining in the last three 12-month periods before April 2020 (i.e., the start of the COVID-19 crisis). Figure 12, which presents the length of positive trends over a number of years O*NET by O*NET analogously to Figure 7 for various families, does not have ones and twos in the 2019 columns. That population shows that the ad number steadily increased for almost all O*NETs over the considered period.

	From 4-15 to 3-16	From 4-16 to 3-17	From 4-17 to 3-18	From 4-18 to 3-19	From 4-19 to 3-20
Software Developers, Applications	228,187	255,535	372,912	499,671	554,549
Computer User Support Specialists	128,850	134,507	179,526	219,455	231,534
Network and Computer Systems Administrators	121,316	138,764	167,781	197,676	217,134
Computer Systems Analysts	104,817	117,830	147,148	163,622	188,230
Information Technology Project Managers	78,683	101,079	138,551	164,404	198,474
Computer Systems Engineers/Architects	55,425	67,865	97,086	131,699	161,798
Information Security Analysts	62,704	71,059	99,535	123,001	144,186
Software Quality Assurance Engineers and Testers	51,105	57,324	77,292	96,927	102,297
Web Developers	51,706	55,329	72,071	82,984	86,307
Software Developers, Systems Software	45,621	43,381	59,692	83,771	81,097
Operations Research Analysts	29,023	32,011	42,910	51,224	54,845
Computer and Information Research Scientists	11,526	15,498	24,293	34,077	39,217
Computer Programmers	20,112	17,480	24,549	29,479	34,098
Database Administrators	17,663	16,254	19,752	22,422	23,479
Computer Network Architects	13,626	11,697	13,349	17,763	18,470

Figure 11. RQ2: Heat map of ad number by O*NET clustered in 12-month ranges (from April 2015 to March 2020).

	2019				2020				2021
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1
Software Developers, Applications	3	4	4	4	4	0	0	5	5
Computer User Support Specialists	3	4	4	3	0	0	0	0	1
Network and Computer Systems Administrators	3	4	4	3	0	0	0	0	1
Computer Systems Analysts	3	3	4	4	4	0	0	5	5
Information Technology Project Managers	3	4	4	4	4	0	0	5	5
Computer Systems Engineers/Architects	3	4	4	4	4	0	0	5	5
Information Security Analysts	3	4	4	4	4	0	0	5	5
Software Quality Assurance Engineers and Testers	3	4	4	3	0	0	0	4	1
Web Developers	3	0	4	3	0	0	0	0	1
Software Developers, Systems Software	3	3	3	0	0	0	0	0	1
Operations Research Analysts	3	3	4	4	0	0	0	0	1
Computer and Information Research Scientists	3	4	4	4	4	0	0	5	5
Computer Programmers	3	3	3	3	4	0	0	0	5
Database Administrators	3	3	3	3	4	0	0	0	5
Computer Network Architects	3	3	3	3	0	0	0	0	1

Figure 12. RQ2: The length in years of positive trends quarter by quarter and by O*NET.

As there are yearly patterns for each *O*NET* within the *Computer and Mathematical* family analogous to those for the families, we compared the trends of the different *O*NETs* quarter by quarter. Figure 13 presents the comparisons for each quarter starting in 2015. The data used in that figure are, for each *O*NET*, the percentage of ads for that *O*NET* in that quarter over the total number of ads for that *O*NET*. Thus, we can compare the trends of different *O*NET* classifications, even if their absolute numbers of ads are on different scales. The percentages capture the relevance of a specific quarter for a given *O*NET*. Whenever it is over 4.2% (the horizontal line in each diagram), the number of ads created in that quarter is above the average, as we have six years' data (24 quarters).




Figure 13. RQ2: Percentage distribution of ads by *O*NET* in the quarters of 2015–2021 (from Q2 2015 to Q1 2021).

The *uphill stair* pattern of almost all series in Figure 13 charts for 2019 data makes the positive trend of the pre-COVID-19 period patent. Figure 13 adds a quantitative perspective to the trend lengths from Figure 12. Comparing adjacent columns, we can gauge the steady improvements in the years before 2020.

In 2020, the positive trends came to a halt and reversed. In the first quarter of 2020 (bottom line, in shades of yellow), 7 out of the 15 *O*NETs* decreased their number of ads, and the others had minimal increments, showing they were grinding to a halt. All *O*NETs* took a drastic fall in the second and third quarters, as clearly visible in the topmost two charts in Figure 13, which all have shorter rightmost columns. The trend started to improve only in the last quarter of 2020, with half of the *O*NETs* showing better performance than in the previous year. Finally, the first quarter of 2021 is green for all sectors in Figure 7, and the bottom chart in Figure 13 shows distinctly taller rightmost columns.

Each column in Figure 13 represents the contribution of the corresponding quarter to the overall ads for an *O*NET*. Thus, the difference between two such columns compares the relevance of two quarters. However, the same gap between two columns may be more or less significant, depending on the column height. A 0.2% distance is impressive between 0.1% and 0.3% because it corresponds to tripling the performance. Vice versa, the same 0.2% distance between 3.7% and 3.9% is much less relevant. Therefore, to understand the significance of the losses and gains of ads for the different *O*NETs*, Figure 14 presents the percentage of gains/losses in ad numbers for a given quarter and *O*NETs* with respect to the same quarter of the previous year.

	2019				2020				2021
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1
Software Developers, Applications	28.6%	23.3%	18.6%	0.1%	4.5%	-26.1%	-5.4%	26.3%	67.3%
Computer User Support Specialists	16.3%	1.4%	14.1%	7.2%	0.0%	-41.2%	-29.3%	-14.6%	20.9%
Network and Computer Systems Administrators	29.9%	17.3%	21.9%	11.9%	-8.8%	-26.0%	-15.5%	-3.7%	34.1%
Computer Systems Analysts	15.6%	13.8%	21.2%	22.8%	4.1%	-42.7%	-28.4%	5.9%	42.4%
Information Technology Project Managers	33.1%	28.7%	34.5%	22.4%	1.9%	-43.7%	-31.1%	10.8%	54.2%
Computer Systems Engineers/Architects	47.6%	33.3%	30.6%	32.1%	1.3%	-27.6%	-16.0%	11.8%	69.2%
Information Security Analysts	29.1%	21.4%	17.7%	24.1%	7.2%	-20.2%	-20.7%	3.4%	39.1%
Software Quality Assurance Engineers and Testers	20.9%	3.3%	15.6%	5.6%	-1.3%	-32.5%	-21.1%	1.6%	43.7%
Web Developers	14.0%	-2.2%	7.0%	13.4%	-0.6%	-20.7%	-19.0%	-1.8%	48.3%
Software Developers, Systems Software	29.8%	12.1%	1.1%	-11.1%	-12.6%	-30.1%	-27.0%	-1.6%	24.1%
Operations Research Analysts	15.7%	10.7%	9.4%	9.9%	-1.5%	-48.2%	-28.2%	-3.6%	53.5%
Computer and Information Research Scientists	44.0%	25.5%	35.7%	3.3%	2.1%	-17.4%	-21.4%	10.7%	51.3%
Computer Programmers	11.3%	10.9%	12.4%	35.5%	6.4%	-15.9%	-14.4%	-16.3%	36.1%
Database Administrators	10.0%	4.3%	0.9%	5.7%	8.5%	-27.4%	-17.1%	-7.8%	19.0%
Computer Network Architects	42.3%	8.8%	1.2%	7.7%	-0.8%	-38.6%	-19.8%	-15.1%	25.0%

Figure 14. RQ2: The relative gaps quarter by quarter by O*NET.

Many factors affect the values in Figure 14, so that the results are difficult to explain without analyzing the text of a statistically significant percentage of the ads in detail. For instance, 77.5% of ads for the O*NET *Computer and Information Research Scientists* matched an FT-index search for (any inflectional form of) any among the following (full-text searches use quotes to denote phrases): “data science”, “data scientist”, bigdata, “big data”, “machine learning”, “artificial intelligence”, AI, “deep learning”, deeplearning, datamining, and “data mining”. Thus, we can explain the comparatively small loss for that O*NET (−17.4%) in the second semester of 2020 as a response to the increased need to analyze clinical, market, and user data to respond to the COVID-19 crisis. We plan to conduct similar investigations for the other O*NETs as well. Unfortunately, such further inquiries are incredibly time-consuming, as human intervention and manual fine-tuning are needed in most cases.

Fortunately, thanks to their job descriptions, we can reason other O*NET results out with less effort. Indeed, we can conjecture that *Computer Programmers* had the best performance (−15.9%) at the crisis peak, as their job description includes scripting and adapting applications. Thus, when a large part of work and customer interactions had to move online, requiring adaptations of significant parts of the existing systems, their work was much in need. The same reason can explain the third quarter’s performance (−14.4%), showing a loss of the same magnitude and the second-best result after *Software Developers, Applications* (−5.4%), though other O*NET improvements make this performance less impressive. In the fourth quarter, normality started to reestablish itself, with other O*NETs gaining back or having only slightly negative performances. Thus, the need for *Computer Programmers* was less felt, and the corresponding result (−16.3%) was the worst.

Analogously, we can conjecture that *Information Security Analysts* (−20.2%) and *Network and Computer Systems Administrators* (−26.1%) had their hands full, with many employees accessing data networks from home. Thus, their O*NETs were not as severely affected as others were in the second quarter.

Vice versa, the professions central to large project development, especially in the early phases of design, such as *Computer Network Architects* (Q2: −38.6%, Q3: −19.8%, Q4: −15.1%), *Computer Systems Analysts* (Q2: −42.7%, Q3: −28.4%), and *Information Technology Project Managers* (Q2: −43.7%, Q3: −31.1%), took the bluntest hits because, in such uncertain times, several large projects were postponed.

The positions needed by both large and small projects, such as *Computer Systems Engineers/Architects* (−27.6%), *Software Developers, Systems Software* (−30.1%), and *Software Quality Assurance Engineers and Testers* (−32.5%), fell in the intermediate area. *Web Developers* (−20.7%), who are often involved in small projects and given a positive trend for web usage, held their own in the second quarter.

The *Computer User Support Specialists* bad performances (Q2: −41.2%, Q3: −29.3%, Q4: −14.6%) are harder to explain. Indeed, we can expect an increased need for people providing technical help to non-IT computer users. However, we must consider that the crisis somehow lightened their workload, which includes assistance to new employees to learn the use of company facilities, company LAN maintenance, printing services, and the

like. Indeed, the newly hired were in a smaller number, and some headquarters shut down. Moreover, the workforce's contraction diminished the need for IT support, and it is possible that IT-savvy employees, who were previously involved in more gratifying projects, were reassigned to such tasks instead of letting go. Finally, many employees were forced to learn IT skills to maintain their productivity while working from home. Hence, companies possibly realized that they could cut their IT support costs, diminishing the professional help available to workers who were no longer completely ignorant.

Operations Research Analysts had a very bad 2020 (Q1: −1.5%, Q2: −48.2%, Q3: −28.2%, Q4: −3.6%), with the worst second-quarter performance of all *O*NETs* in the family and their negative trend starting in the first and ending in the fourth quarter, while most *O*NETs* were gaining in those quarters. We can attribute such a bad performance to the strategic role of operations research analysts, who were not highly in demand when most companies were focused on survival, putting on hold, or altogether canceling visionary projects. Indeed, the first 2021 quarter shows the relevance of that *O*NET* bouncing back (53.5%) on the wave of the growing optimism of the worldwide economy.

Database Administrators is another difficult-to-interpret category. Their loss of job openings started in the second quarter, and they had the best performance in the first 2020 quarter. However, they were later than most in recovering, having the first positive value only in the first 2021 quarter and with the most negligible improvement. Database administrators are needed to manage large data numbers. Thus, they are usually employed by medium to big companies whose core business is rarely in the ICT realms. We can then conjecture that, when facing the economic pandemic crisis, such companies focused their resources mostly on their core business and delayed investing in data management when possible. The limited percentage loss at the crisis peak suggests that database administrators were replaced when needed, but new data management projects were put on hold. Another factor to consider is that in 2019, the sector exhibited notable variance (and the same is true also for the previous years, though there is no clear seasonal pattern). Thus, it might be a fluctuating job market with quirks superimposed on the effects of COVID-19.

A final *caveat* is that the COVID-19 crisis was not the only factor influencing the job market. Indeed, we can see considerable variance in the 2019 data, which the forthcoming pandemic cannot have caused. Thus, we might overlook other factor contributions amplifying or obscuring the COVID-19 impacts by singling out specific values for individual analysis. Therefore, we restricted ourselves to discussing the data where contiguous values corroborated the reasoning.

Summary. In conclusion, the COVID-19 crisis affected the IT job market especially in the second and third quarters of 2020. The effects in the first quarter are mostly limited to a loss of improvement, though *Computer User Support Specialists*, *Network and Computer Systems Administrators*, *Software Quality Assurance Engineers and Testers*, *Web Developers*, *Software Developers*, *Systems Software*, *Operations Research Analysts*, and *Computer Network Architects* registered a smaller number of ads with respect to the first quarter of 2019. In the second and third quarters, all *O*NETs* lost jobs in significant percentages. The first 2021 quarter showed promising increases in job ads for all *O*NETs*.

3.3. RQ3: Analyzing Changes in Popularity of Telecommuting

To understand the impact of COVID-19 on the interest in *working from home* in the IT sector, we analyzed the text of the job ads for the ten *O*NETs* having the highest number of ads, listed at the top of Table 3. Using full-text searches, we counted the ads containing at least one synonym for *working from home*.

The first step of that procedure was choosing the synonyms. We wanted to select the terms that were currently widely accepted as an alternative way to denote *working from home*. Thus, we started from those used in the Wikipedia article on *working from home* because the collaborative nature of that website promotes the usage of modern standard terminology. Then, for each candidate synonym, we verified its appropriateness, checking its definition in well-renowned dictionaries and in which contexts it was used online.

Finally, we collected the synonyms of such terms in dictionaries and performed a few online searches to keep only those that were broadly accepted. Indeed, many synonyms had a more specific meaning, were accepted only in some contexts, or had many other usages and would have been a source of false positives. At the end of our selection process, the synonym set consisted of *work from home*, *flexible workplace*, *remote work*, *telecommuting*, and *telework*.

Because of the ability of full-text searches to match inflectional forms, such synonyms were sufficient to find ads including plural forms or other verbal tenses and modes, such as *working from home*, *flexible workplaces*, and *telecommute*.

We randomly selected fifty ads that did not match the query and did not find false negatives. Therefore, we stopped looking for further synonyms.

We then manually inspected a random sample of the matching ads looking for false positives and found a few. Analyzing them, we understood that two classes of problems may cause them.

- The match correctly identified the usage of a synonym. However, it was in a negative context, so we could not count it as proof of interest in the subject. For instance, the sentence *Potential for Teleworking: No* is a correct match for *telework*, but the containing ad is not about a position offering some form of work from home.
- The full-text search captured sentences that should not be a match. For instance, one of the main challenges was *work from home* because *from* is a *stop word* (being a preposition), so it is not indexed. Thus, *work from home* matches *work at home* (which is fine), but also *work and home*, as in *balance work and home life* or in *across all platforms, from work and home to car and mobile*.

By manually inspecting the hits, we devised a heuristic to detect and exclude such false positives from the search. For the first kind of problem, we took advantage of another feature of full-text searches by considering the vicinity of the looked-for terms to other phrasal words. This way, we could exclude instances found in incorrect contexts, such as in the example above. For the second kind of problem, we harnessed full-text search and standard *LIKE* clauses together. Those are much less efficient and are sensible to spacing and punctuation but also consider stop words.

The ads selected by our query were not *all* those pertinent to job positions allowing some form of teleworking. For instance, companies with a well-established workplace flexibility policy may not explicitly refer to it in each ad. Moreover, some ads may convey the idea that teleworking is possible implicitly through the job description. Thus, the number of ads explicitly referring to teleworking indicates the topic's interest but does not directly measure the number of positions allowing one to work from home.

In Figure 15, we can see the percentages of ads explicitly referring to teleworking over the number of ads in the same quarter for the same *O*NET*. The first ten bar diagrams starting from the top left correspond to the ten most populous IT *O*NETs*, while the last on the bottom-right represents the total of the ten considered *O*NETs*. For each *O*NET*, the solid gray bars are the values of the 24 quarters (6 years) in temporal order from left to right, while the last two series on the right are the average over the quarters up to 2019 (checkered in yellow and black) and from 2020 on (striped in green and black), respectively. The two last series are the average of the *pre-COVID-19* and the *post-COVID-19* quarters, respectively, and their comparison lets us gauge the effects of COVID-19 on the issue.

Table 4 compares the average percentage of ads explicitly mentioning teleworking in the years before and after COVID-19 and their ratio. Considering the total data of all the ten listed *O*NETs*, we moved from 1.3% to 3.2%, thus seeing change on the scale of a factor of two-and-a-half. To evaluate if there is a statistically significant difference between these percentages (i.e., between the pre- and post-COVID-19 distributions reported in the two columns of Table 4), we analyzed them using a statistical test. Since the Shapiro–Wilk test provided normality (i.e., p -value > 0.05) for both distributions, we adopted a parametric test (this choice follows the suggestions given by ([16], Chapter 37)). In particular, we used the paired T-test to evaluate if there was a statistically significant difference between the

two distributions. Since the computed p -value was <0.01 , we concluded that the difference between the two distributions is statistically significant.




Figure 15. RQ3: Percentage distribution of ads concerning work from home.

Table 4. RQ3: Percentages of ads mentioning teleworking pre- and post-COVID-19.

<i>O*NET</i>	Average		Ratio
	Pre-COVID-19	Post-COVID-19	
Software Developers, Applications	1.0%	3.3%	3.2
Computer User Support Specialists	1.7%	3.1%	1.8
Network and Computer Systems Administrators	1.4%	3.9%	2.7
Computer Systems Analysts	1.1%	2.1%	1.9
Information Technology Project Managers	1.1%	2.5%	2.3
Computer Systems Engineers/Architects	1.2%	3.1%	2.7
Information Security Analysts	2.3%	5.0%	2.2
Software Quality Assurance Engineers and Testers	0.9%	2.9%	3.2
Web Developers	1.5%	3.1%	2.1
Software Developers, Systems Software	0.4%	2.3%	5.1
<i>Total</i>	1.3%	3.2%	2.5

Let us consider now the individual *O*NETs*. The increments range from about double for *Computer User Support Specialists* (1.8), *Computer Systems Analysts* (1.9), *Information Security Analysts* (2.2), and *Web Developers* (2.1), to about triple for *Software Developers, Applications* (3.2), *Network and Computer Systems Administrators* (2.7), *Computer Systems Engineers/Architects* (2.7), and *Software Quality Assurance Engineers and Testers* (3.2), and to the outstanding five-time value of *Software Developers, Systems Software* (5.1).

The series for the second 2020 quarter, the peak of the COVID-19 crisis, are labeled by their values, and we can see that they are taller than the previous columns. However, the cause of the astonishing results we have for the average is not an isolated spike. Indeed, we see tall columns for all quarters from the second quarter of 2020 on, towering over those of the pre-COVID-19 era for most *O*NETs*. Therefore, the increased interest in teleworking is not a short-time fad. It is becoming a well-established feature in job ads. This aligns with McKinsey’s report about the future of work after COVID-19 [17] “Remote work and virtual meetings are likely to continue, albeit less intensely than at the pandemic’s

peak". Additionally, other researchers found similar results. For instance, Kong et al. [18] highlighted that employees who found telecommuting during the pandemic to be a helpful option are often interested in continuing to work from home after the pandemic. Similarly, Da Silva et al. [19] found that the COVID-19 pandemic has altered work-from-home patterns, and both businesses and employees are still adjusting to these changes. Remote work will likely continue to be in significantly higher demand than it was before the start of the pandemic.

Summary. In conclusion, during the COVID-19 crisis, the sensibility of the need to work from home increased, and the percentage of ads explicitly mentioning it became about two and a half times higher after COVID-19 than it was before, on average. The highest impact is for *software developers of systems software*, with a ratio greater than five. Moreover, the increased interest in teleworking appears to not be a short-time fad.

4. Related Works

The COVID-19 pandemic represents a historical event and, for researchers from all over the world, a massive (and unplanned) experiment as well. For this reason, many are working to understand its effects from multiple viewpoints. Several papers and technical reports from relevant universities, institutions, or companies have focused on the impact of COVID-19 on the job market.

Some researchers adopted data sources similar to ours, that is, related to job ads [20–22]. For example, Hensvik et al. [20] measured the job-search responses to the COVID-19 pandemic using real-time data on vacancy postings and ad views on Sweden's largest online job board. Similarly to our results, they found that the number of new vacancies published online drastically decreased in the aftermath of the pandemic, with their results showing that, since March 2020, employers have posted around 40% fewer vacancies. The same holds for Holgersen et al. [21], who used the vacancy posting data from the Norwegian Labor and Welfare Administration job database to study the impact of the COVID-19 crisis on labor demand in Norway. They found that, in response to the COVID-19 pandemic, the vacancy postings from late February to the end of June 2020 declined by around 27% relative to the same period in 2019. Forsythe et al. [22] proposed a deeper understanding of how the labor market evolved over the COVID-19 crisis, using the data provided by Burning Glass Technologies, a company that scrapes, cleans, and codes job vacancies posted on the internet at a daily frequency. Analogously to our results, they found that US job vacancies collapsed by over 40%. Similarly to those mentioned above, our work provides an overview of the pandemic's impact on the labor market in general (RQ1). In addition, our work also provides a detailed view of the effects on job roles in the IT sector (RQ2) and an analysis of telecommuting (RQ3). That was possible thanks to the massive number of ads available in the LinkUp database, the granularity of their classification (i.e., the *O*NET* system), and the possibility of accessing their full-text description.

A set of interesting analyses and charts has been made available by SmartMarketData on their website [23]. This material is based, as in our case, on the LinkUp job ad database. Moreover, LinkUp also provides (for a fee) reports that give insights into the virus's impact on hiring and the economy [24]. Starting from the same raw dataset of those reports, various authors have produced independent analyses. Thus, the result presented here are unrelated to those provided by LinkUp and by SmartMarketData. The work in [10] also analyzes data from a large job ad repository but with different goals; specifically, they aimed to answer separate research questions about the distribution of the workforce between testers and developers, the relevance of the different testing techniques and tools, and the comparative interest in automatic and manual testing in the industry.

Following a different approach, several studies were based on surveys to gather relevant information directly from the population (aiming to generalize the results to the entire population). For example, Bick et al. [25] contacted around 5000 working-age adults selected as a representative sample of the US population to understand how many US workers shifted to remote work in the months after the pandemic outbreak. They found that

of all those employed in May, 35.2% worked entirely from home, whereas 8.2% reported doing so in February; in particular, highly educated, white, and high-income individuals were much more likely to shift to remote work and to maintain employment following the virus outbreak. These results align with our findings for **RQ3**, where we observed a peak in the job ads explicitly offering the possibility of working remotely. Other researchers analyzed specific aspects instead, such as Collins et al. [26], who examined the changes in mothers' and fathers' work hours from February to April 2020 using the panel data from the US Current Population Survey (approximately 60,000 households are surveyed each month). They found that mothers with young children reduced their work hours four to five times more than fathers.

Basile et al. [27] analyzed the effect of COVID-19 on the usage of the Reddit social news platform. In particular, they found that the Reddit users' posting behavior, in terms of submissions, comments, and scores, was strongly affected and modifications were uniform across countries. Additionally, the expression of emotion seen from users was amplified. The paper also provides an interesting overview of the literature concerning the effect of COVID-19 on social media; for example, Cinelli et al. [28] analyzed the diffusion of information about COVID-19 by conducting a massive data analysis on Twitter, Instagram, YouTube, Reddit and Gab. Gozzi et al. [29] analyzed the impact of media coverage and epidemic progression on collective attention in four countries: Italy, the United Kingdom, the United States, and Canada. Beyond news articles, videos, and Wikipedia page views, they analyzed Reddit posts and found that collective attention was mainly driven by media coverage rather than epidemic progression. Other authors analyzed the effect of the COVID-19 pandemic on other social media, such as Li et al. [30], who analyzed Weibo, or Nemes et al. [31], who analyzed Twitter.

5. Conclusions, Implications, Limitations, and Future Work

We decided to conduct this empirical study to better comprehend the impact of the COVID-19 pandemic on the IT job market and compare it with other sectors. The research method we adopted is content analysis. In particular, we analyzed slightly less than 109 million ads for the entire job market and more than 9 million ads for the IT market, taken from LinkUp, a relevant Web job-search platform. The answers to the research questions of this work derive from several comparisons between the numbers of job ads published during the crisis to those of the same period of previous years.

We can summarize the main findings from the study as follows. First, as expected, the COVID-19 pandemic has affected the entire job market, with the number of ads being significantly reduced at the time of the peak of the crisis in all sectors except the sector "Food Preparation and Serving-Related" (+16.0 percent). Second, the ad number in the IT sector dropped between 15% and 48% in the second quarter of 2020, depending on the specific professional figure. In particular, the data showed substantial job losses and a decrease in hiring for professionals associated with the development of large projects (e.g., *Computer Systems Analysts* and *Information Technology Project Managers*). On the contrary, minor losses occurred for professionals closer to the implementation aspects, such as *Computer Programmers*, *Computer and Information Research Scientists*, and *Information Security Analysts*. Third, during the COVID-19 crisis, companies' sensitivity towards remote work for IT professionals increased, and the number of ads explicitly mentioning that issue increased three times.

5.1. Implications of the Study

We adopted a perspective-based approach to judge the practical implications of our study, considering the manager, IT professional, and researcher perspectives. Understanding how the pandemic has changed the IT job market is useful for IT managers and professionals both in the immediate future, for example, to reposition professionals who are no longer needed or to find a new job, but also in the long term if a similar event were to repeat itself. For example, we have learned that in times of crisis, implementation jobs

seem to be much more stable and less risky than those related to requirement analyses and software design. The study mainly focuses on one aspect of the COVID-19 crisis in the world of work, i.e., teleworking, and it does it very simply using a single data source. This is relevant for the researcher, who could be interested in analyzing other aspects and adopting different data sources and more refined analyses. Another interesting aspect for researchers is to estimate whether the effect of the increase in teleworking that we have observed is transitory or will remain definitive.

5.2. Limitations and Future Research

The main limitation of this study is that it relies on only one data source. Thus, an interesting future work would be to integrate it with other sources. For this reason, in our future work, we intend to conduct interviews with employers and human resources managers and include ads from platforms other than LinkUp to extend the LinkUp data and deepen our understanding of the observed trends. Another limitation is that the data, though worldwide, are not uniform in their geographical distribution, so they could reflect untrue trends everywhere. More sophisticated web data mining and analysis techniques, such as those described in [32], might overcome this limitation, and we plan to adopt them in future work.

Author Contributions: All authors have contributed equally. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data sharing not applicable—no new data generated.

Acknowledgments: We want to show our gratitude to LinkUp [7] for providing us the possibility to access to the database containing the job advertisements (i.e., the raw data).

Conflicts of Interest: The authors declare no conflict of interest.

Disclaimer: Responsibility for any information and result reported in this paper lies entirely with the authors. LinkUp was not involved in any form of data analysis.

References

1. Barrero, J.M.; Bloom, N.; Davis, S.J. *COVID-19 Is Also a Reallocation Shock*; Working Paper 27137; National Bureau of Economic Research: Cambridge, MA, USA, 2020. [CrossRef]
2. The Economic Times: Job Loss Most Severe Immediate Impact of COVID-19: Survey. Available online: <https://economictimes.indiatimes.com/jobs/job-loss-most-severe-immediate-impact-of-covid-19-survey/articleshow/76323649.cms?from=mdr> (accessed on 25 June 2023).
3. University of Chicago: How 2 Out of Every 5 Jobs Lost during COVID-19 May Not Come Back. Available online: <https://www.weforum.org/agenda/2020/05/42-of-jobs-lost-during-covid-19-may-not-come-back> (accessed on 25 June 2023).
4. Gurcan, F.; Kose, C. Analysis of software engineering industry needs and trends: Implications for education. *Int. J. Eng. Educ.* **2017**, *33*, 1361–1368.
5. Krippendorff, K. *Content Analysis: An Introduction to Its Methodology*; SAGE Publications: Thousand Oaks, CA, USA, 2018.
6. Feldman, R.; Sanger, J. *The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*; Cambridge University Press: Cambridge, UK; New York, NY, USA, 2007.
7. LinkUp: Insightful Job Market Data. Available online: <https://www.linkup.com/> (accessed on 25 June 2023).
8. Sandler, R. Forbes: Here's When Major Companies Plan to Go Back to the Office. Available online: <https://www.forbes.com/sites/rachelsandler/2020/08/27/heres-when-major-companies-plan-to-go-back-to-the-office/> (accessed on 25 June 2023).
9. Cerioli, M.; Leotta, M.; Ricca, F. COVID-19 Hits the Job Market: An 88 Million Job Ads Analysis. In Proceedings of the 36th ACM/SIGAPP Symposium on Applied Computing (SAC 2021), Virtual, 22–26 March 2021; pp. 1721–1726. [CrossRef]
10. Cerioli, M.; Leotta, M.; Ricca, F. What 5 million job advertisements tell us about testing: A preliminary empirical investigation. In Proceedings of the SAC '20: The 35th ACM/SIGAPP Symposium on Applied Computing, Online Event, Brno, Czech Republic, 30 March–3 April 2020; Hung, C., Cerný, T., Shin, D., Bechini, A., Eds.; ACM: New York, NY, USA, 2020; pp. 1586–1594. [CrossRef]
11. Occupational Information Network (O*NET): Resource Center. Available online: <https://www.onetcenter.org/overview.html> (accessed on 25 June 2023).
12. OpenAI. GPT-3.5. Online. 2023. Available online: <https://openai.com> (accessed on 25 June 2023).
13. Przystanski, A. When Are Companies Hiring the Most? Available online: <https://blog.namely.com/what-time-of-year-do-most-companies-hire> (accessed on 25 June 2023).

14. Smith, M. The Holiday Hiring Freeze: Fact or Fiction? Available online: <https://www.topresume.com/career-advice/holiday-hiring-freeze-fact-or-fiction> (accessed on 25 June 2023).
15. Keane, J. Forbes: Glovo CEO Sees Acceleration and Consolidation in Food Delivery Post-COVID-19. Available online: <https://www.forbes.com/sites/jonathankeane/2020/06/15/glovo-ceo-sees-acceleration-and-consolidation-in-food-delivery-post-covid-19/> (accessed on 25 June 2023).
16. Motulsky, H. *Intuitive Biostatistics: A Non-Mathematical Guide to Statistical Thinking*; Oxford University Press: Oxford, UK, 2010.
17. Lund, S.; Madgavkar, A.; Manyika, J.; Smit, S.; Ellingrud, K.; Robinson, O. The Future of Work after COVID-19. McKinsey Global Institute. Online. 2021. Available online: <https://www.mckinsey.com/featured-insights/future-of-work/the-future-of-work-after-covid-19> (accessed on 25 June 2023).
18. Kong, X.; Zhang, A.; Xiao, X.; Das, S.; Zhang, Y. Work from home in the post-COVID world. *Case Stud. Transp. Policy* **2022**, *10*, 1118–1131. [[CrossRef](#)] [[PubMed](#)]
19. Da Silva, A.D.; Georganakos, D.; Weißler, M. How People Want to Work—Preferences for Remote Work after the Pandemic. ECB Economic Bulletin. 2023. Available online: https://www.ecb.europa.eu/pub/economic-bulletin/focus/2023/html/ecb.ebbox202301_04~1b73ef4872.en.html (accessed on 25 June 2023).
20. Hensvik, L.; Le Barbanchon, T.; Rathelot, R. Job Search during the COVID-19 Crisis. *J. Public Econ.* **2021**, *194*, 104349. [[CrossRef](#)]
21. Holgersen, H.; Jia, Z.; Svenkerud, S. Labor Demand during the COVID-19 Crisis in Norway: Evidence from Vacancy Posting Data. 2020. Available online: <https://ssrn.com/abstract=3663479> (accessed on 25 June 2023).
22. Forsythe, E.; Kahn, L.B.; Lange, F.; Wiczer, D. Labor Demand in the time of COVID-19: Evidence from vacancy postings and UI claims. *J. Public Econ.* **2020**, *189*, 104238. [[CrossRef](#)] [[PubMed](#)]
23. SmartMarketData: COVID-19. Available online: <https://www.smartmarketdata.com/covid-19> (accessed on 25 June 2023).
24. LinkUp: COVID-19 Data Packages. Available online: http://pages.linkup.com/DataPackages_COVID-19DataPackages.html (accessed on 1 May 2021).
25. Bick, A.; Blandin, A.; Mertens, K. *Work from Home after the COVID-19 Outbreak*; Federal Reserve Bank of Dallas: Dallas, TX, USA, 2020. [[CrossRef](#)]
26. Collins, C.; Landivar, L.C.; Ruppner, L.; Scarborough, W.J. COVID-19 and the gender gap in work hours. *Gend. Work Organ.* **2021**, *28*, 101–112. [[CrossRef](#)] [[PubMed](#)]
27. Basile, V.; Cauteruccio, F.; Terracina, G. How Dramatic Events Can Affect Emotionality in Social Posting: The Impact of COVID-19 on Reddit. *Future Internet* **2021**, *13*, 29. [[CrossRef](#)]
28. Cinelli, M.; Quattrocioni, W.; Galeazzi, A.; Valensise, C.M.; Brugnoli, E.; Schmidt, A.L.; Zola, P.; Zollo, F.; Scala, A. The COVID-19 social media infodemic. *Sci. Rep.* **2020**, *10*, 16598.
29. Gozzi, N.; Tizzani, M.; Starnini, M.; Ciulla, F.; Paolotti, D.; Panisson, A.; Perra, N. Collective Response to Media Coverage of the COVID-19 Pandemic on Reddit and Wikipedia: Mixed-Methods Analysis. *J. Med. Internet Res.* **2020**, *22*, e21597. [[CrossRef](#)] [[PubMed](#)]
30. Li, L.; Zhang, Q.; Wang, X.; Zhang, J.; Wang, T.; Gao, T.L.; Duan, W.; Tsoi, K.K.f.; Wang, F.Y. Characterizing the Propagation of Situational Information in Social Media During COVID-19 Epidemic: A Case Study on Weibo. *IEEE Trans. Comput. Soc. Syst.* **2020**, *7*, 556–562. [[CrossRef](#)]
31. Nemes, L.; Kiss, A. Social media sentiment analysis based on COVID-19. *J. Inf. Telecommun.* **2021**, *5*, 1–15. [[CrossRef](#)]
32. Smith, D.; Ali, A. Analyzing Computer Programming Job Trend Using Web Data Mining. *Issues Informing Sci. Inf. Technol.* **2014**, *11*, 203–214. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.