CAGE

Job Search during the COVID-19 Crisis

CAGE working paper no. 473

May 2020

Lena Hensvik
Thomas Le Barbanchon
Roland Rathelot

WARWICK
THE UNIVERSITY OF WARWICK

UKRI
Economic and Social Research Council

Job Search during the COVID-19 Crisis

Lena Hensvik
Thomas Le Barbanchon
Roland Rathelot *

May 2020

Abstract

This paper measures 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. First, the labour demand shock in Sweden is as large as in the US, and affects industries and occupations heterogeneously. Second, the scope and direction of search change. Job seekers respond to the shock by searching less intensively and by redirecting their search towards less severely hit occupations, beyond what changes in labour demand would predict. The redirection of job search changes relative hiring costs, and has the potential to amplify labour demand shifts.

Keywords: coronavirus, search intensity, search direction, labour demand shock, job vacancies, online job board

JEL Codes: J22, J23, J21, J62, J63, J64, E24

*Hensvik: Uppsala University. Le Barbanchon: Bocconi University. Rathelot: University of Warwick. We would like to thank the Swedish Public Employment Service for making the data available. We gratefully acknowledge financial support from the ERC StG 758190 "ESEARCH", Forte and IFAU. Ana Teles provided excellent research assistance. All errors are our own. Hensvik is also affiliated to IFAU, CEPR, Le Barbanchon to CEPR, IGIER, IZA, J-PAL, Rathelot to CAGE, CEPR, J-PAL.

1 Introduction

In the aftermath of the COVID-19 pandemic, economic activity falls all over the world. Early April, IMF forecasts that global GDP will fall by 3% in 2020, and even by 5.9% in the U.S. and 7.5% in the euro area.[1] Still according to the IMF, unemployment is expect to rise by as much as 3 percentage points in the euro area. Unemployment in the US sore from below 4% to 14.7% early May.[2] One key question is how workers will react to this crisis. Depending on whether job search reacts to the shock, the supply side of the labour market may amplify, or compensate, the consequences of the shock on labour demand.

This paper provides the first empirical evidence about the impact of the COVID crisis on job search. We analyse real-time data on job search from the job board of the Swedish public employment service *Platsbanken*. Following a large decrease in the number of new vacancies posted, we document whether job seekers search more or less intensely and how they adjust the direction of their search.

First, we show that the number of new vacancies posted online decreases drastically in the aftermath of the sanitary crisis: since early March 2020, employers post around 40% less vacancies. The labour demand shock differs across industries. Some industries, such as hotels, and restaurants, reduce hiring by as much as 185%, probably because of a sharper demand shock on their product market triggered by social distancing recommendations. There is also a strong heterogeneity in the labour demand shock across occupations. In particular, we find that occupations characterised by less home-working suffer more than occupations with a higher share of hours worked from home.

The sharp reduction in vacancies is followed within two weeks by a 10% reduction in the number of clicks on vacancies by *Platsbanken* users. We further find that the average number of clicks per user decreases simultaneously, which nets out the direct aggregate effect on the number of vacancies available to click on. The reduction in job search intensity also holds in within-worker designs. This provides strong empirical support for procyclical job search intensity. From the employers side, we also find that the number of clicks per vacancy decreases. Even though employers face less competition to attract applicants as fewer vacancies are online, their vacancies receive less attention, so that the vacancy-level tightness overall increases.

Beyond search intensity, our data allow to document the impact of the crisis on the direction of search. Do workers re-direct their search towards markets where they are in higher de-

[1] IMF World Economic Report of April 2020, published on 6 April 2020.
[2] BLS' *The Employment Situation*, 8 May 2020

mand? First, *Platsbanken* users significantly shift their vacancy clicks towards occupations with high home-working index. The observed search redirection goes beyond how counterfactual job seekers clicking randomly on available vacancies would shift their vacancy clicks. Second, we split occupations in two groups – resilient vs non-resilient – according to their differential evolution of vacancy inflows between February and March 2020. We find that the share of clicks towards resilient markets increases (beyond their market share in vacancies). Overall, this provides empirical evidence that the direction of job search is dynamic and reacts to composition shocks on labour demand.

From the employers' point of view, the re-direction of job search following the COVID shock has heterogeneous impacts. Employers posting jobs in high home-working or resilient occupations receive more clicks per vacancy than employers posting in other occupations. This suggests that the endogenous response of job search to labour demand shifts is likely to amplify them. As vacancies in resilient occupations attract more attention, recruitment processes may speed up, decreasing recruitment costs, which would induce these employers to open up new vacancies.

We first contribute to the recent literature documenting the effects of the COVID-19 crisis on labour markets. Dingel and Neiman (2020) provide an ex-ante description of US jobs, which would be on hold because of severe social-distancing measures, while Mongey and Weinberg (2020) describe their workers. Brynjolfsson et al. (2020) and Bartik et al. (2020) use ad-hoc surveys to provide evidence for the massive switch to home-working during the first weeks of April. Kahn et al. (2020) documents the extent and heterogeneity in the drop of labour demand in the US using online vacancy data and new UI claims in March and April 2020. Our contribution is to combine real-time administrative data from both sides of the labour market to precisely document the effects of the COVID-19 crises on job search.

Our analysis relates to the empirical literature on job search that use data from online job boards (Marinescu, 2017; Belot et al., 2018; Marinescu and Rathelot, 2018; Banfi and Villena-Roldan, 2019; Marinescu and Wolthoff, 2020; Kudlyak et al., 2020; Brown and Matsa, 2020). The only paper using click data that focuses on the response of job search to labour market conditions is Faberman and Kudlyak (2019). They find that the number of applications per job seeker is higher in metropolitan areas where unemployment is higher, using cross-sectional variation. In a within-user design, we leverage the COVID shock to study the job search response, both its intensity and its selectivity.

We also contribute to the literature evaluating the cyclicality of job search intensity. The early literature, which leverages time-use data, finds that the time unemployed spent searching jobs correlates negatively with local unemployment rates, although the correla-

tion is not statistically significant (Krueger and Mueller, 2010; Krueger et al., 2011; Krueger and Mueller, 2012, 2016).[3] Merging time-use data with the CPS on longer time periods, Gomme and Lkhagvasuren (2015) confirm that search effort is procyclical. Using the same survey data but relying mostly on cross-state variations for identification, Mukoyama et al. (2018) find that search effort is countercyclical. In this paper, we find that job search intensity is procyclical, leveraging a large labour demand shock, and we document that search direction is also "procyclical", as search redirects towards resilient occupations in bad times.

The paper proceeds as follows. We describe the Swedish institutional background in Section 2 and the data in Section 3. We document the labour demand shock induced by the COVID crisis in Section 4. We estimate the response of job search intensity in Section 5, and of the direction of search in Section 6. We conclude in Section 7.

2 Background

The first Swedish case of COVID-19 was confirmed on January 31 in a traveller from China. As soon as the second week of March, community spread was confirmed and various measures were taken with the aim of slowing down the spread (or "flattening the curve"). These measures have been relatively mild compared to other countries, primarily relying on voluntary compliance with the Public Health Authority's social distancing guidelines. During the second week of March, the Public Health Agency made several formal announcements, ordering that all residents keep a distance from each other; that high schools and universities should be closed, and that workers should work remotely to the extent possible. Gatherings were also limited to 500 people; a restriction that was further tightened to 50 two weeks later.[4]

Figures A1a-A1f show the time evolution of time spent at different places according to data provided by Google's COVID -19 Mobility Reports. They suggest that the public announcements were followed by substantial drops in time spent in workplaces (-24 %), retail and recreation (-41 %) and transit stations (-36 %), while time spent at home and in parks increased (+84 % and +6% respectively). While the response in Sweden is weaker than in European countries with stricter social distancing measures such as Norway, Denmark and France, it is more similar to the drop in mobility in the US.[5]

[3]State-level unemployment rate in Krueger and Mueller (2010); national unemployment in the cross-country study Krueger and Mueller (2012); county-level unemployment rates in the New Jersey survey Krueger et al. (2011); Krueger and Mueller (2016)

[4]In contrast to many other countries, primary schools have remained open throughout.

[5]Figure A2a shows that the number of COVID-related deaths have risen dramatically since the beginning

In response to the crisis, the Swedish government has taken several measures to protect jobs and workers. On the firm side, this implies a payroll tax reduction and the short-time work system, which allows firms to reduce working time for their employees by 20, 40 or 60 percent (March to May also 80 percent).[6] On the worker side, several measures have been taken to extend unemployment insurance coverage and increase benefit levels during the crisis.[7]

3 Data

The primary data source for the analysis is online data consisting of all posted job ads and the search activity on Sweden's largest job board (Platsbanken), which is operated by the Swedish Public Employment Service (PES). On Platsbanken, firms can post vacancies and screen applicants. Users can search and view ads and apply to posted vacancies.

On the vacancy side, the data contain rich information about the posted job, such as the occupation, location, start- and end date, working hours, skill requirements, etc. The data inform about the first date of publication on the website, when users can start to view the job ad, and the deadline date for applications. There is also a firm identifier, which allows us to map each vacancy to firm-level industry codes according to the Swedish SNI classification. In Sweden, employers do not post wages. We thus assign to vacancies their occupation mean wage in 2018 obtained from official statistics.[8] In addition, we add information about the home-working prevalence of occupations. Our primary measure is derived from the American Time Use Survey (ATUS). For each occupation, we calculate the mean share of hours worked at home over hours worked at the workplace and home (Hensvik et al., 2020).[9] We also complement the analysis with the home-working index provided by Dingel and Neiman (2020) as well as the one in Mongey et al. (2020).[10]

On the job seeker side, our data allow us to follow users over time via an anonymised identifier. For each user, we have information about the vacancy id of the viewed ad and a time stamp. Importantly, our data contain the ad views of all users, both those searching from their computer and those using their phone.

of March, amounting to over 2,000 deaths by the end of April.

[6]By the beginning of May, 15 percent of Swedish firms have taken up this (temporary) short-term work subsidy).

[7]For a more detailed description of the Swedish policy-response, see Hensvik and Skans (2020).

[8]The wage data (*Strukturlonestatistiken*) are downloaded from Statistics Sweden's web page, see www.scb.se. 2018 is the latest year available.

[9]The calculations are based on more than 30,000 workers observed between 2011 and 2018.

[10]The alternative measures assess the tele-workability of occupations according to the description of their tasks (based on O*NET.)

The working dataset contains the search activity from January 2020 to April 2020, which amounts to more than 90 millions clicks on vacancies. We also add search activity from January 2019 to April 2019 as a control group when we compute difference-in-difference estimates of the COVID-19 effects. On the labour demand side, we observe vacancies available on Platsbanken over the same period in 2020 and in 2019. This amounts to just over 400,000 clickable job ads.

4 The impact of the crisis on labour demand

4.1 The aggregate impact

The aim of this section is to describe how labour demand changes on the Swedish labour market in the wake of the COVID 19 pandemic. Our primary measure of labour demand comes from vacancy postings on Platsbanken, the largest job board in Sweden. We measure the changes in labour demand by the average daily inflow of new vacancies per week. The left panel of Figure 1 shows the evolution of the daily inflow of vacancies and jobs from January to mid April in 2020, while the right panel compares the change to 2019. The inflow is stable until week 10 and experiences a sharp and persistent drop in week 11 (which starts on March 9th). Difference-in-difference estimates from models comparing the change before and after week 10 in 2019 and 2020 suggest a reduction by 40% in the inflow of vacancies, and 34% on the inflow of jobs (see Table B1 in Appendix). The magnitude of the decline in vacancy postings is very similar to what has been documented in the US (Kahn et al., 2020).

We corroborate this large labour demand shock by showing that the number of layoff notices increases sharply in March 2020, compared to previous months or to the same period in 2019 (Appendix Figure A2b), and that the number of registrations of job seekers at the Swedish Public Employment Service increased by 70% (Appendix Table B1, col. (4)). We also note that the timing of the drop in vacancy postings coincides with the reductions in mobility, as measured by Google mobility reports (Figures A1a-A1f).

4.2 Heterogeneity by industry, and occupation

The substantial drop in vacancy postings following the COVID outbreak could reflect several factors. Some local labour markets will be impacted directly by the social distancing guidelines through reductions in product demand and in derived labour demand. Other

local labour markets are likely impacted more indirectly, either as suppliers of industries directly impacted by social-distancing measures or as suppliers of foreign industries from countries more severely impacted by the pandemic. Furthermore, social distancing recommendations apply among co-workers in workplaces where there are no close contact with customers. This may further slow down activity in industries producing upstream in the input-output network.[11] Finally, overall increased uncertainty is expected to cause large output contraction (Baker et al., 2020)

To better understand the nature of the pandemic crisis, this section describes which industries and occupations have been most severely affected. Starting with a simple variance decomposition of the labour demand drop between occupations and industries we find that about two thirds of the drop is explained by industries while one third can be attributed to occupations. This suggests that the shock is not operating through uniform reductions in impacted industries, and that different occupations within industries are affected differentially. This fact would be consistent with firms updating their production functions during the crisis. Below, we describe in more detail the industries and occupations most severely affected by the shock.

Industry For each 1-digit industry, we compute the difference in the inflow of vacancies before and after week 10 in 2020, net of the inflow change in 2019 (reported in Appendix Figure A3). In other words, we define the date of the shock using the observed break in vacancy inflows and estimate the difference-in-difference in weekly vacancy inflows by industry around that date. While the shock has a negative impact on all industries, some industries are substantially more severely affected. In particular, we see larger drops in industries where social-distancing measures are likely to bind, such as hotels and restaurants (-185%), entertainment (-150%) and retail trade (-92%). The impact is much more moderate in the health and education sector, in real estate and in public administration and defence (reductions around -30%).

At the outset of the crisis, the Swedish government has declared some industries as essential.[12] We show that the decline in the number of vacancies posted is parallel in essential vs. non-essential industries (Appendix Figures A4), consistently with the finding of Kahn

[11] Barrot et al. (2020) analyze the impact of social-distancing rules on GDP in France, when taking sectoral interdependencies into account. They find large impacts of social distancing rules on upstream sectors.

[12] The Swedish government declared as essential industries: "Electricity, gas, steam and air conditioning supply", "Financial and insurance activities", "Wholesale and retail trade; repair of motor vehicles and motorcycles", "Manufacturing", "Human health and social work activities", "Information and communication", "Water supply; sewerage, waste management and remediation activities", "Public administration and defence; compulsory social security" and "Transportation and storage".

et al. (2020) in the U.S. This could be explained by all industries anticipating the slow-down of future aggregate demand, and thus reducing hirings.

Occupation Second, we analyse the differential labour demand shock by occupation. We first isolate the ten most shrinking and the ten most growing occupations according to the difference-in-difference estimates by 3-digit occupation (see Appendix Table B2). Among the ten occupations with the largest decrease in vacancy inflow, we find waiters and bartenders (-225%), dentists (-126%), and fast-food workers (-123%). Journalists and health care specialists are examples of occupations relatively resilient to the health crisis.

To further describe the nature of the shock, Figure 2 shows the evolution of posted vacancies in 2019 and 2020, when vacancies are characterised by the occupation wage, by whether or not the vacancy belongs to health care occupations, and by the occupation home-working prevalence (see Appendix Table B3 for corresponding difference-in-difference estimates). The drop in vacancy postings is larger in occupations at the bottom of the wage distribution (see Figure 2a). However, there is no clear break in week 10, which calls for some caution in attributing this drop to the COVID outbreak.[13] In contrast, Figure 2b shows a marked increase in the fraction of posted vacancies towards health care occupations. The increase by 5 percentage point clearly illustrates the rapidly increasing demand for health care personnel during the COVID-outbreak.

Figure 2c shows the home working shares of posted vacancies, derived from the American Time-Use Survey, as described in Hensvik et al. (2020). On average, the home-working share increases by 0.5 percentage points after the shock, suggesting that jobs where social-distancing measures bind less, as they are already done at home, seem to resist more to the health crisis shock. In Appendix Figure A6, we report results of the same analysis using alternative home-working measures, as computed by Dingel and Neiman (2020) and Mongey et al. (2020), who use O*NET descriptions of occupations. We find that these measures are not as predictive of the differences in the evolution of labour demand across occupations.

[13] Appendix Figure A5 shows the evolution of the median wage, suggesting a clearer increase at the time of the crisis.

5 The impact of the crisis on job search intensity

5.1 Aggregate trend

This section describes the aggregate evolution in job search. Apart from within-week daily seasonality, daily clicks on Platsbanken are fairly stable in February 2020. They decrease abruptly by around 10% mid-March 2020, and follow a downward trend in the following weeks (see appendix Figure A7).[14]

The evolution of aggregate vacancy views is driven by both the number of job ads available on the website and the search behaviour of workers. The previous section shows that firms post fewer vacancies since early March 2020, which contributes to lower aggregate clicks even when workers exert the same job search effort. To account for the evolution of available jobs, we now turn to econometric models in which the outcomes are the number of daily clicks per user and daily clicks per vacancy.

5.2 Econometric specification

We take the job seeker's perspective, and consider as outcome daily clicks per user. We estimate the following equation:

$$y_{id} = \sum_{w} \beta_w \mathbf{1}\left[d \in w\right] + \mu_i + t_{id}\delta + \varepsilon_{id} \tag{1}$$

where y_{id} is the log of the number of times user i clicked during day d, β_w are week fixed effects and the parameters of interest, μ_i are user fixed effects, and t_{id} is the number of days since the user first clicked on Platsbanken. This last covariate controls for duration dependence in job search behaviour (Faberman and Kudlyak, 2019). In the main specification, the estimation sample includes user-days (i, d) when there is at least one click. We thus focus on the intensive margin of job search behaviour. In additional specifications (reported in the Appendix), we also consider the extensive margin of job search, whether users click at least once in a given day. We then run the estimation considering the inverse hyperbolic sine of the number of clicks, which allows to keep in the estimation sample the user-days (i, d) when there is no click.

[14] We cannot directly compare the March 2020 evolution to the trend in March 2019 because of missing data in late February and early March 2019. However, appendix Figure A7 further illustrates that standard week-by-week variations in daily clicks are lower in magnitude than what we witness in March 2020.

Similarly, from the vacancy perspective, we estimate the following equation:

$$y_{jd} = \sum_{w} \beta_w \mathbf{1}\left[d \in w\right] + \lambda_j + X_{jd}\gamma + \varepsilon_{jd} \tag{2}$$

where y_{jd} is the log of the number of clicks on vacancy j on day d, β_w are week fixed effects, λ_j are vacancy fixed effects, X_{jd} are time-varying characteristics of the vacancy, including dummies for the first day and the last day it was published, as well as a flexible functional transformation of number of days since publication, that aims to capture duration dependence in clicks. Our main estimation sample includes for each vacancy days with at least one click. We check the robustness of the results in total margin specifications with the inverse hyperbolic sine of clicks on vacancy as an alternative outcome.

In all models, we take week 6 as the reference, and we cluster inference at the user (for regression 1), or the vacancy level (for regression 2).

5.3 Impact on individual search intensity

Figure 3 plots the week effects from Regressions (1) and (2). The top left panel in Figure 3 shows the estimates of the coefficients on the week dummies in the regression where the log number of clicks per job seeker is the outcome, and there is no user fixed effects. Conditional on clicking at least once in the day, the average number of clicks is fairly stable from one week to the other until week 11 in mid-March 2020: it hovers between -5% and 2.5% (in deviation from week 6). After week 12 (which starts on March 16), the number of clicks experiences a sharp decrease of around 15%. Given the large sample size, this drop is highly statistically significant. The drop is not due to seasonal effects, as the same estimates for year 2019 do not show any decline in the number of clicks per user.

The decline of clicks per user could be partly due to the decrease in the number of vacancies available on the job board. In the bottom-left panel of Figure 3, we run regression (2) of the log number of daily clicks per vacancy in a specification without vacancy fixed effects. The choice of the outcome restricts the sample to vacancy-days with a least one click, and focuses the estimation on the intensive margin of vacancy-level tightness. We find a sharp and persistent drop in daily clicks per vacancy of around 25% after week 12. The shift is very clear and significant, compared to both the situation before week 12, and the same weeks in 2019. From the perspective of individual vacancies, the situation is paradoxical. Even though vacancies face less competition, as the total number of vacancies decreases, and a larger pool of potential applicants, as the number of incoming job seekers increases sharply, our analysis shows that each vacancy receives less attention in the aftermath of the

crisis. The effective number of job seekers per vacancy appears to be smaller after week 12. This result could be explained by the impact of the crisis on individual behaviour, or could be due to composition effects. Both the pool of potential applicants and available vacancies may become more negatively selected since March 2020, with lower-effort job seekers and lower-quality jobs.

In the right-side panels, we control for composition effects by introducing user fixed effects (top panel) and vacancy fixed effects (bottom panel). We find a significant drop in the clicks per user in week 12 of around 15%. The magnitude of the effect is similar to the estimate from the regression without fixed effects. In the bottom right panel, we find a different pattern for our estimates. While the cross-sectional estimates are stable over time (except for the week-11 shock), the within-vacancy week effects show negative trends both in 2019 and 2020. This result resists our controls for duration dependence. However, we still find a drop in clicks per vacancy by 5% between week 10 and week 11. Overall, this provides stronger supportive evidence for a reduction in job search intensity.

In Appendix Figure A8, we show that job search also decreases when considering both margins, intensive and extensive. The total magnitude of the decrease is smaller though, as it is the sum of both a decrease in the share of users clicking at least once, and of the intensive margin effect (scaled by the extensive margin). The decrease amounts to around 0.02-0.03 in the clicks per user and to 0.1 in the clicks per vacancy. This result is robust whether we consider cross-sectional estimates or within-user/vacancy estimates.

6 The impact of the crisis on the direction of job search

While job search becomes less intense, one key question is whether it redirects towards specific jobs. Does the direction of search change following the labour demand shock of the COVID crisis? Does any re-direction of job search lead to differential impact on recruitment across employers?

6.1 How workers change the nature of their job search

In Section 4.2, we have documented that labour demand responses to the shock are heterogeneous by occupations. Do job seekers react to the change in the occupational composition of the vacancy pool?

We construct four outcomes y_{id} for user i in day d, corresponding to the occupations of the vacancies that the job seeker has clicked on:

- their average wage corresponding to the occupation of the vacancies user i click on during day d (targeted wage),
- the share of healthcare occupations (targeted healthcare),
- the average home-working index (targeted home-working) as computed in Hensvik et al. (2020),
- the average *resilience* index (targeted resilience): we attribute to each occupation the difference-in-difference estimate corresponding to the impact of the shock on the inflow of vacancies in that occupation (see e.g., Table B2).

Figure 4 plots the estimates of the week fixed effects (red dots) from within-user Regression (1), using the search-direction measures as outcomes. By construction, only user-days with at least one click are included in the estimation sample. We find that the same user has a slightly higher propensity to click on vacancies with higher wages over time: from January to April 2020, targeted wages increase by around 1% (top left panel). While the increase in targeted wages is gradual over time, the evolution of the other search-direction outcomes is kinked: flat until week 10, then sharply increasing. The share of clicks on healthcare occupation vacancies increases by 2 percentage points over the month of March 2020 (top right panel). The targeted home-working index increases by almost 0.008 by the end of April 2020 (bottom left panel). This represents 5% of the cross-occupation standard deviation of the home-working index (see Hensvik et al., 2020). Finally, the targeted resilience sharply increases by 0.08 during March 2020 (bottom right). The average vacancy clicked on belongs to an occupation whose change in aggregate number of vacancy creations is 8% above the trend.

As already explained above, the change in search direction could be driven by the evolution of available vacancies. We plot in blue dots the week fixed effects estimated in counterfactual regressions under the assumption of random search. For every click in the data, we impute a counterfactual clicked vacancy, that is drawn randomly from the pool of vacancies available on that day. In the top panels of Figure 4, we find that random searchers closely follow the actual increase in targeted wages and healthcare content. This leaves little room for behavioural response on top of composition changes of available vacancies. In the bottom panels, we also find that random searchers would click more frequently after March 2020 on occupations that are intensive in home-working, and that are resilient in terms of labour demand. However the magnitude of the increase for random searchers is lower than for actual searchers: one half for the home-working index, and one fourth for the resilience index.

Importantly, the change in search direction shown in Figure 4 is not driven by composition changes in the pool of job seekers, as we control for user fixed effects. Appendix Figure A9 shows the robustness of the results of regressions that do not include fixed effects. In the Appendix, we also report the changes in the search directions using home-working indices based on the task-content of occupations. We find that clicks after March 2020 are also re-directed towards occupations with high ONET-based indices of home-working (Appendix Figure A10). This contrasts with the downward trend of random-search that the two ONET-based measures deliver.

In a nutshell, job search is not sluggish. The direction of job search seems to react quickly, targeting resilient occupations to a greater extent than their share in available jobs.

6.2 The differential impact of job search changes on hiring

This section documents the effect of changes in search direction on employers posting vacancies of different types. To do so, we split the sample of vacancies according to their occupations, and we run a variation of Regression (2) of the number of clicks per vacancy, in which we include an occupation-group interaction term. We use the following specification.

$$y_{jd} = \sum_{w} \delta_w \mathbf{1}\left[j \in \mathcal{O}\right] \times \mathbf{1}\left[d \in w\right] + \sum_{w} \beta_w \mathbf{1}\left[d \in w\right] + \lambda_j + X_{jd}\gamma + \varepsilon_{jd} \tag{3}$$

where y_{jd} is the log of the number of clicks on vacancy j on day d, $\mathcal{O}$ is a vacancy subsample of interest, either resilient occupations, or occupations that are more often worked from home. All other notations have been previously defined in Regression (2). Our parameters of interest are the weekly fixed effects δ_w. They identify the weekly deviation of the number of clicks on vacancies belonging to subsample $\mathcal{O}$ compared to clicks on the complement set of vacancies.

Figure 5 plots the estimated δ_w week fixed effects for the subsample of high home-working occupations in the left-hand side panel, and of resilient occupation in the right-hand side panel (without vacancy fixed effects). For the year 2020 (blue dots), We find that vacancies in high home-working occupations attract 5 to 7% more clicks after week 11 than vacancies in non-resilient occupations. This is in sharp contrast with the periods before week 10 in 2020, and over the same period in 2019 (red dots), when resilient occupations do not attract more clicks. We can thus attribute the change in relative attractiveness to the COVID shock. Similarly, we find that vacancies in resilient occupations attract 10% more clicks after mid-March 2020 than vacancies in non-resilient occupations.

In Appendix Figures A11 and A12, we check the robustness of our results at the total margin response and when we include fixed effects respectively. Overall, we find that vacancies from high home-working or resilient occupations receive more clicks than other vacancies. As expected, the order of magnitude of the coefficient on week dummies is lower when we consider both the intensive and extensive margins of clicks per vacancy. Similarly, within-vacancy analysis yields lower COVID effects, but the patterns are qualitatively similar.

Heterogeneous effects on clicks per vacancy further confirm that COVID effects on search direction are not only driven by quantitative changes in vacancy composition. In the previous section, we use random-search as a counterfactual to assess the contribution of vacancy composition changes. We can also use effects on clicks per vacancy across occupation groups. If the increase in clicks per *user* to resilient occupations were only due to the fact that there are more vacancies from resilient occupations after mid-March 2020, we should not observe any increase in clicks per vacancy from resilient occupations, relative to non-resilient occupations. As we do just above, this section provides further evidence that job seekers redirect their search disproportionately to composition changes.

Importantly, these results across vacancy occupation groups suggest employers from resilient occupations actually benefit from the change in the nature of search of workers, relative to non-resilient occupations.

7 Conclusions

This paper leverages real-time data from the largest online job board in Sweden to study the reaction of job seekers to the COVID-19 sanitary crisis. We use clicks by users of the online platform to measure search intensity and search direction.

Labour demand, measured by the number of vacancies posted in the job board, experiences a brutal and persistent decrease since early March 2020. At the same time, the number of layoffs notifications is 10 times its 2019 level. The number of registration to the public employment service is roughly twice as large as in 2019. With a contracting labour demand and an increasing labour supply, one could expect, if job search individual behaviors remain as pre-crisis, that job-filling rates could increase, initiating a virtuous circle by lowering the cost of hiring for employers.

Our results suggest that job search reacted strongly to the crisis. The overall of clicks, instead of going up with the number of unemployed workers, is going down. User-level analysis confirms that each job seeker clicks less: the intensive-margin reduction in the

number of clicks is around 15%. From the point of view of employers, each vacancy receives not more but less attention from the supply side: the magnitude of the decline in the number of clicks received by each vacancy is around 30%. Including user and vacancy fixed effects does not change much the magnitude of the effects, suggesting that the decrease is due to the crisis impact on users, rather than to a change in the composition of users or vacancies on the platform.

Finally, we also study how Swedish job seekers change the targeting of the vacancies they click on. We document that, as a result of the crisis, job seekers become more likely to click on vacancies that belong to healthcare occupations, to occupations characterised by a high share of hours worked from home, and to occupations for which labour demand is more resilient to the crisis. The fact that job seekers update the direction of their search has consequences from the point of view of employers. We document that high-home-working vacancies receive between 5 and 7% more clicks due to the crisis, and resilient-occupation vacancies receive between 8 and 10% more clicks.

This last result is important to note for our understanding of job search. If the direction of job search was static, i.e., job seekers were not changing the occupations to which they applied, we would expect resilient vacancies to receive relatively less attention than those from occupations where the number of vacancies becomes very rare. If job search was overall random, and all job seekers were just clicking on vacancies randomly, all vacancies would receive the same number of clicks. Our results reject these models, but are compatible with some families of directed job search, where job seekers strategically revise the value of employment attached to the different occupations as a result of the crisis, or where information about vacancies is skewed towards vacancies posted in the more dynamic occupations.

References

Baker, S. R., N. Bloom, S. J. Davis, and S. J. Terry (2020): "Covid-induced economic uncertainty," Tech. rep., National Bureau of Economic Research.

Banfi, S. and B. Villena-Roldan (2019): "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 37, 715–746.

Barrot, J.-N., B. Grassi, and J. Sauvagnat (2020): "Sectoral Effects of Social Distancing," Working paper, SSRN.

Bartik, A. W., M. Bertrand, Z. B. Cullen, E. L. Glaeser, M. Luca, and C. T. Stanton (2020): "How Are Small Businesses Adjusting to COVID-19? Early Evidence from a Survey," Working Paper 26989, National Bureau of Economic Research.

Belot, M., P. Kircher, and P. Muller (2018): "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice," *The Review of Economic Studies*, 86, 1411–1447.

Brown, J. and D. A. Matsa (2020): "Locked in by Leverage: Job Search During the Housing Crisis," *Journal of Financial Economics*, 136, 623 – 648.

Brynjolfsson, E., J. J. Horton, A. Ozimek, D. Rock, G. Sharma, and H. Y. T. Ye (2020): "COVID-19 and Remote Work: An Early Look at US Data," Working paper.

Dingel, J. I. and B. Neiman (2020): "How Many Jobs Can Be Done at Home?" Working Paper 26948, National Bureau of Economic Research.

Faberman, R. J. and M. Kudlyak (2019): "The Intensity of Job Search and Search Duration," *American Economic Journal: Macroeconomics*, 11, 327–57.

Gomme, P. and D. Lkhagvasuren (2015): "Worker Search Effort as an Amplification Mechanism," *Journal of Monetary Economics*, 75, 106 – 122.

Hensvik, L., T. Le Barbanchon, and R. Rathelot (2020): "Which Jobs are Done from Home? Evidence from the American Time Use Survey," Working Paper DP14611, CEPR.

Hensvik, L. and O. Skans (2020): "COVID-19 Crisis Response Monitoring: Sweden," .

Kahn, L. B., F. Lange, and D. G. Wiczer (2020): "Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims," Working Paper 27061, National Bureau of Economic Research.

Krueger, A. B. and A. I. Mueller (2010): "Job Search and Unemployment Insurance: New Evidence from Time Use Data," *Journal of Public Economics*, 94, 298 – 307.

——— (2012): "The Lot of the Unemployed: A Time-Use Perspective," *Journal of the European Economic Association*, 10, 765–794.

——— (2016): "A Contribution to the Empirics of Reservation Wages," *American Economic Journal: Economic Policy*, 8, 142–79.

Krueger, A. B., A. I. Mueller, S. J. Davis, and A. Şahin (2011): "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High Frequency Longitudinal Data," *Brookings Papers on Economic Activity*, 1–81.

Kudlyak, M., D. Lkhagvasuren, and R. Sysuyev (2020): "Systematic Job Search: New Evidence from Individual Job Application Data," *Journal of Labor Economics*.

Marinescu, I. (2017): "The General Equilibrium Impacts of Unemployment Insurance: Evidence from a Large Online Job Board," *Journal of Public Economics*, 150, 14 – 29.

Marinescu, I. and R. Rathelot (2018): "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal: Macroeconomics*, 10, 42–70.

Marinescu, I. and R. Wolthoff (2020): "Opening the Black Box of the Matching Function: The Power of Words," *Journal of Labor Economics*, 38, 535–568.

Mongey, S., L. Pilossoph, and A. Weinberg (2020): "Which Workers Bear the Burden of Social Distancing Policies?" .

Mongey, S. and A. Weinberg (2020): "Characteristics of Workers in Low Work-from-Home and High Personal-Proximity Occupations," Becker-Friedman Institute White Paper.

Mukoyama, T., C. Patterson, and A. Şahin (2018): "Job Search Behavior over the Business Cycle," *American Economic Journal: Macroeconomics*, 10, 190–215.

Figures

Figure 1: Average daily inflow of vacancies/jobs per week

(a) Inflow of vacancies and jobs

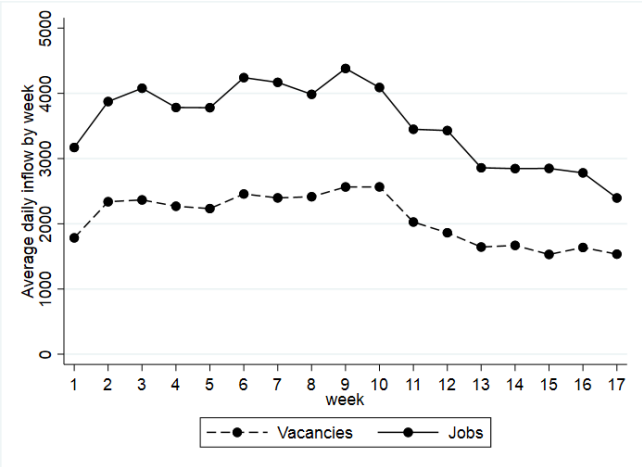

(b) 2019 vs. 2020

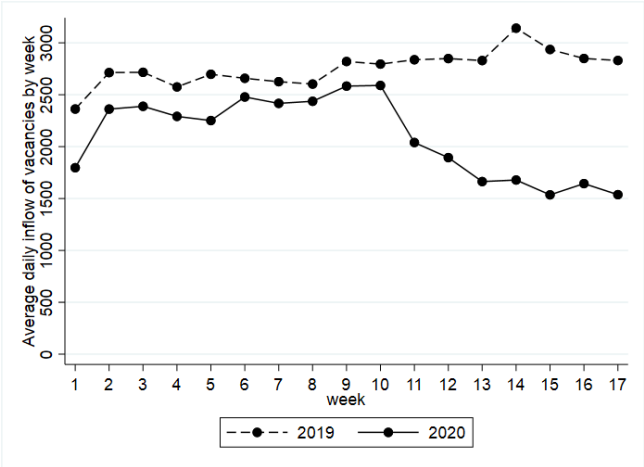

Note: This figure plots the weekly time series of the daily inflow of vacancies and jobs in 2020 (left panel), and of vacancies in 2019 and 2020 (right panel). On the x-axis, we have the rank of the week within the calendar year. In 2020, week 5 begins January 27 and ends February 2; week 10 is from March 2 to March 8; and week 13 from March 23 to March 29.

Figure 2: Trends in types of vacancies by week

(a) Mean log-wage

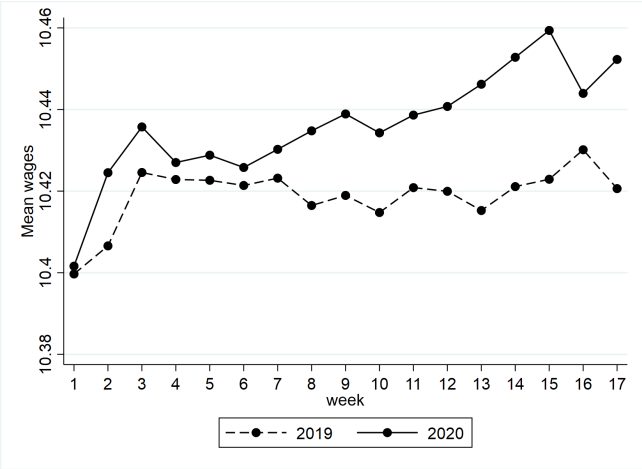

(b) Fraction in health care

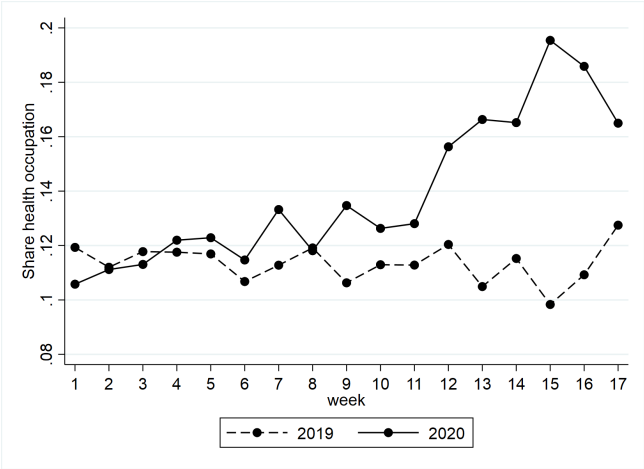

(c) Home working share

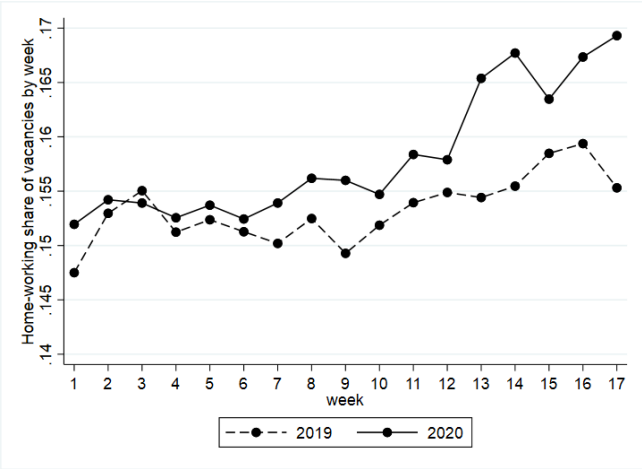

Note: The figure shows the evolution of different vacancy attributes inferred from the occupation of the vacancy. Mean wage is the average occupation wage in 2018; Share health care is the fraction of vacancies with the following occupation titles (SSYK codes): Medical doctors (221), Nursing professionals (222 & 223), Personal care workers in health services (532) and Health care assistants (533). Home working share is the mean share of hours worked at home over hours worked at the workplace and home derived from the American Time Use Survey (ATUS).

Figure 3: Regression of the number of clicks on week dummies, without fixed effects (left panels) and with fixed effects (right panels), clicks per user (top panels) and clicks per vacancy (bottom panels)

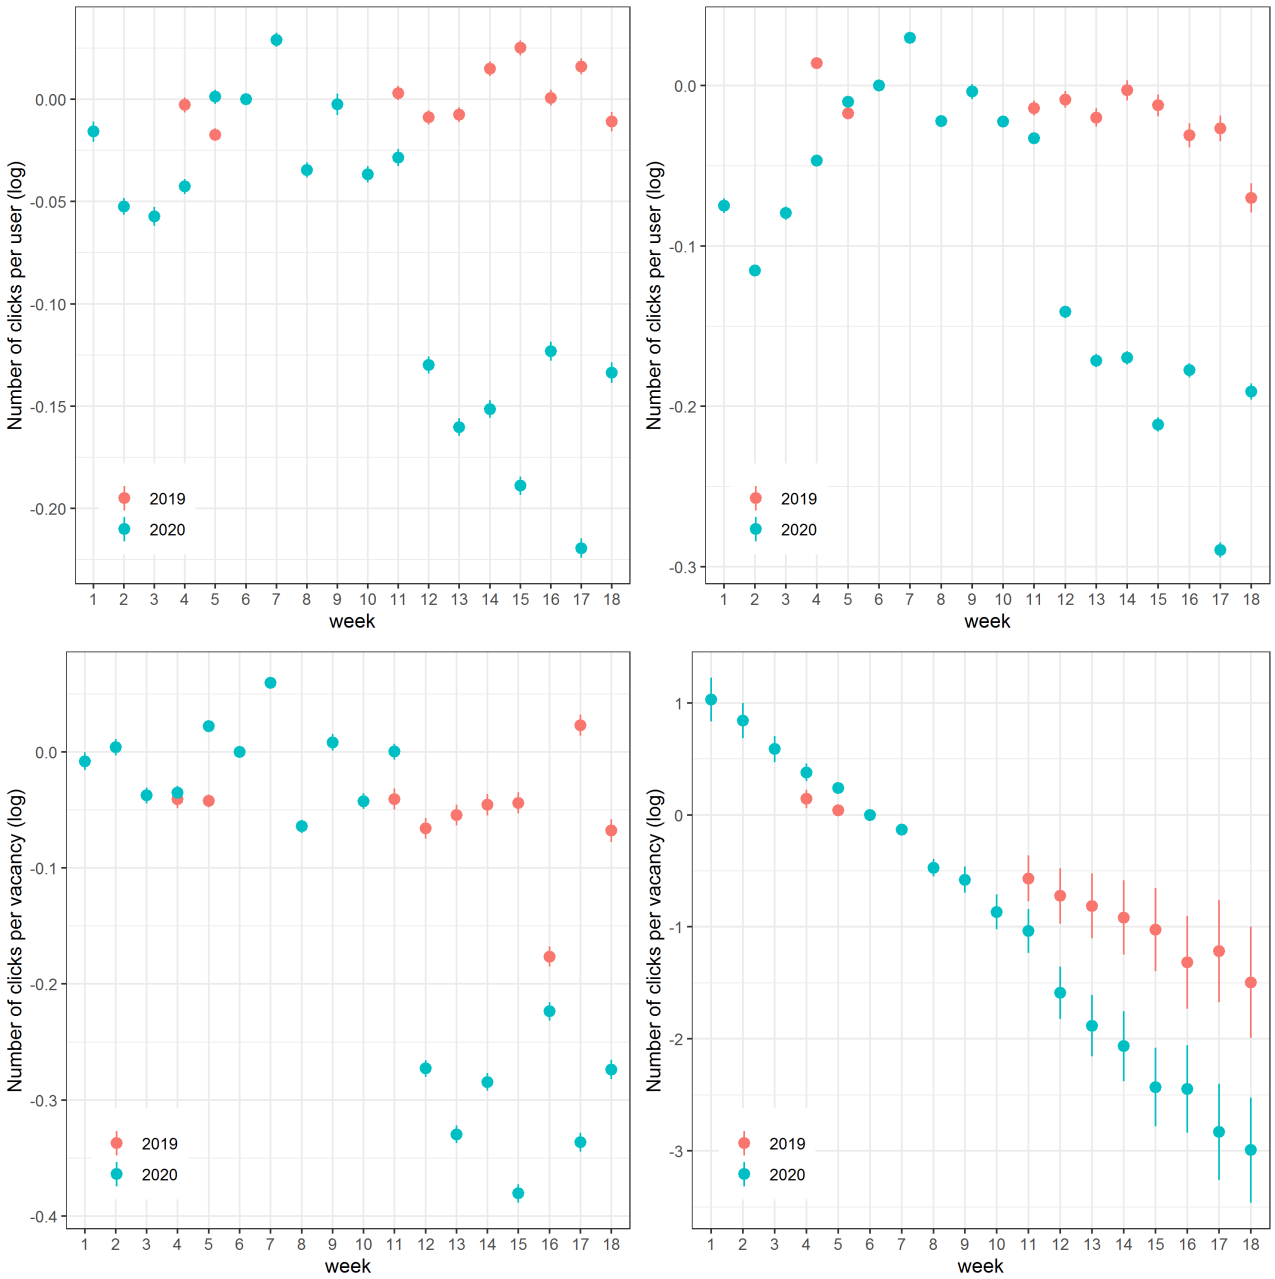

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows the weekly effects on the log of the number of clicks per user (top panels) and of clicks per vacancy (bottom panels). The plotted estimates correspond to the coefficients β_w in Regressions (1) and Regression (2). Blue dots are for 2020 estimates, and red dots for 2019 estimates. The 6th week of the year is chosen as the reference. Standard errors are clustered at the user level in per user regressions, and at vacancy level in per vacancy regressions. Vertical lines show 95% confidence intervals.

Figure 4: Regression of the targeted wage (top left-hand panel), of the healthcare occupation dummy (top right), of the home-working index (bottom left) and of the targeted resilience index (bottom right) on week dummies

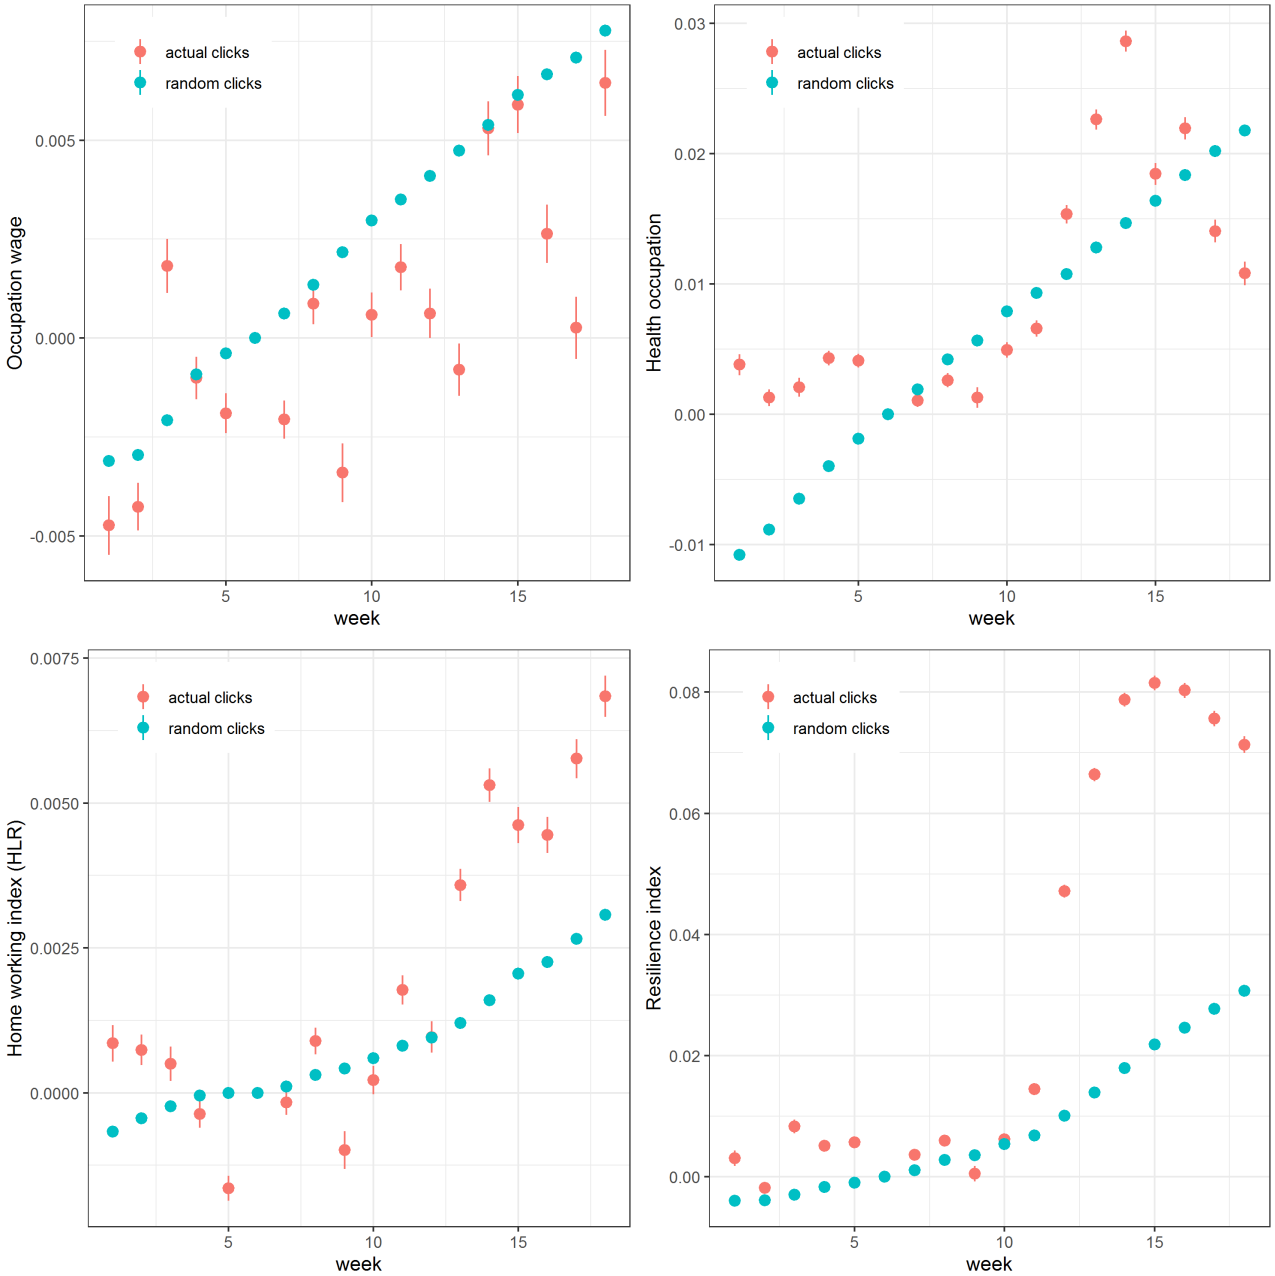

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows the weekly effects on the characteristics of the vacancies that users click on: average wage of the posted occupation (targeted wages), healthcare occupation, home-working index and resilience index. The plotted estimates correspond to the coefficients β_w in Regressions (1) with user fixed effects. Blue dots are for 2020 estimates on actual clicks, and red dots for clicks by counterfactual random searchers on available vacancies. The 6th week of the year is chosen as the reference. Standard errors are clustered at the user level. Vertical lines show 95% confidence intervals.

Figure 5: Regression of the log number of clicks per vacancy on week dummies interacted with high home-working occupations (left) and with resilient occupations (right)

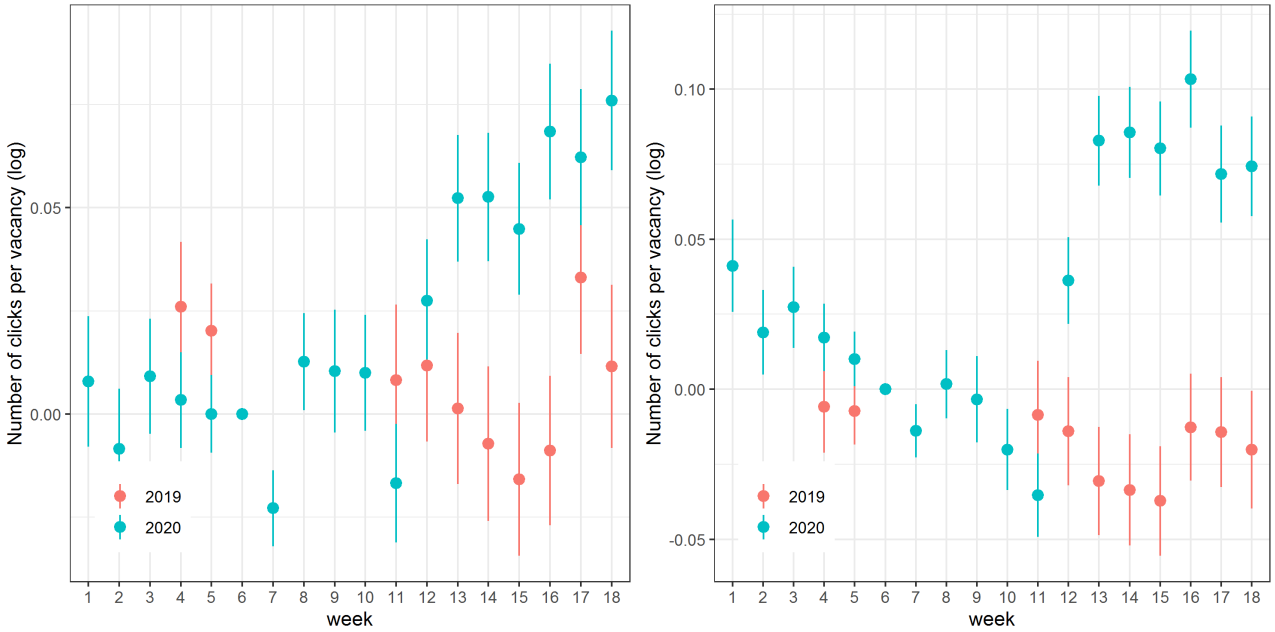

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows coefficient estimates from Regression (3) of the log of the number of clicks per vacancy. We plot the δ_w coefficients of the week effects interacted with high home-working occupations (left) and with resilient occupations (right). Blue dots are for 2020 estimates, and red dots for 2019 estimates. The 6th week of the year is chosen as the reference. Standard errors are clustered at vacancy level. Vertical lines show 95% confidence intervals.

Online Appendix

A Extra Figures

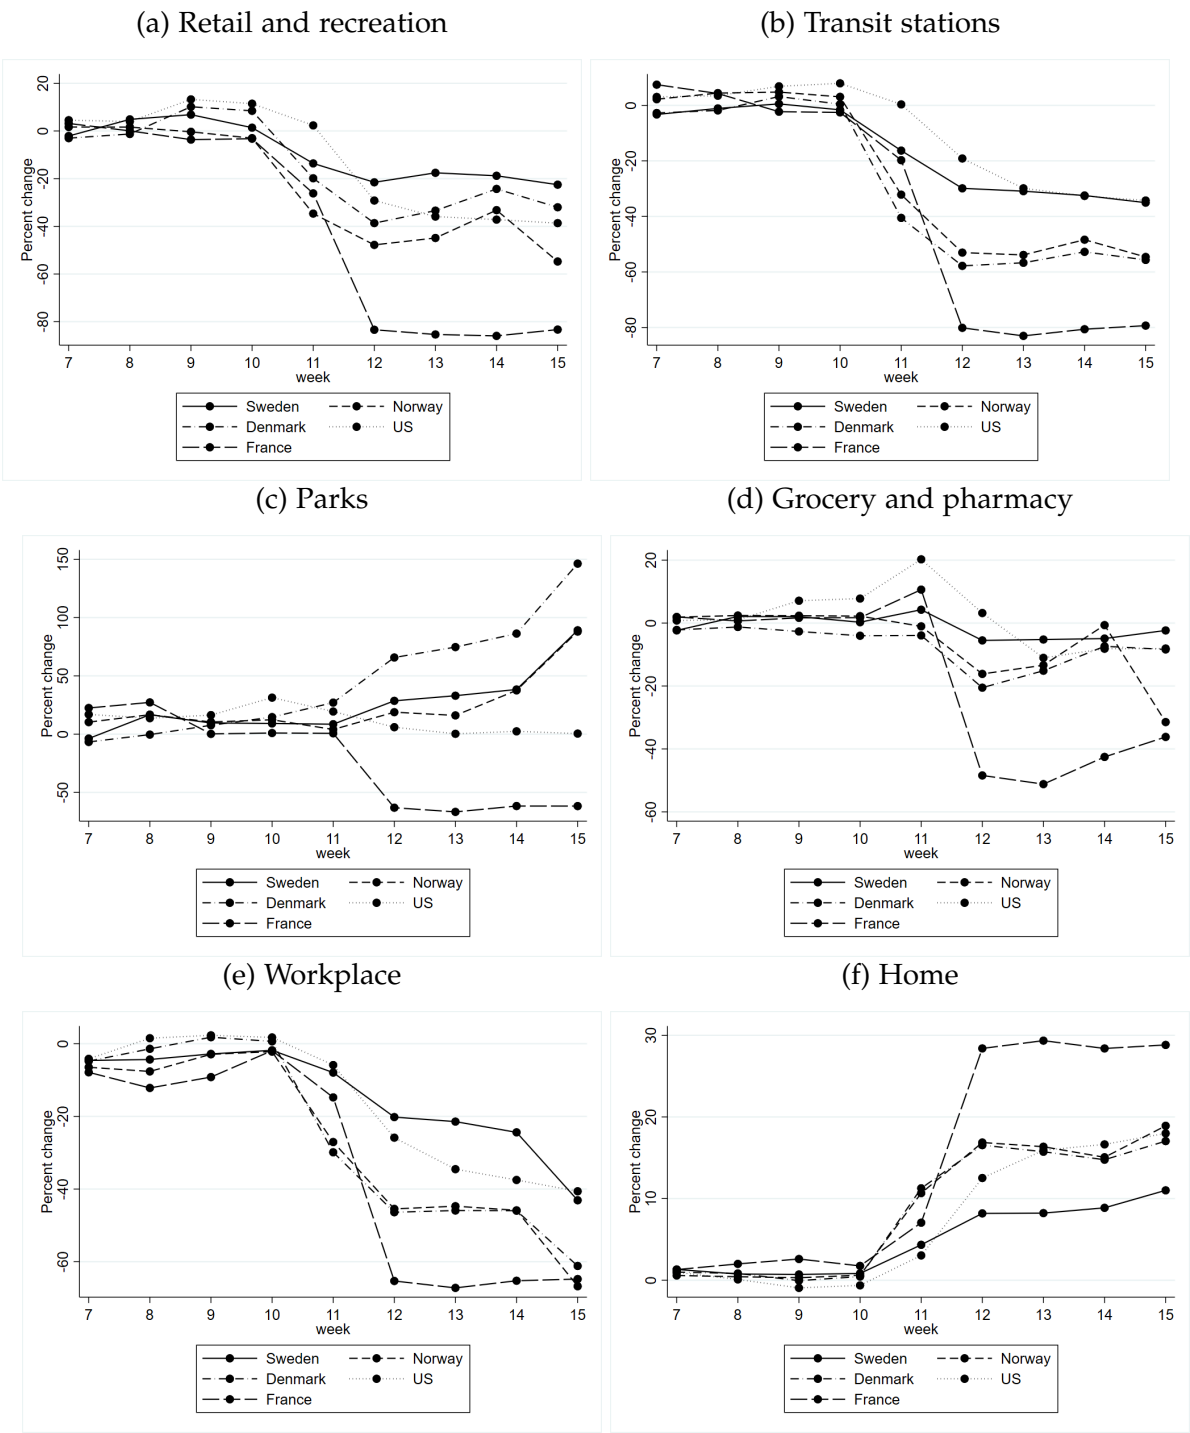

Figure A1: Percent change in time spent at different places in Sweden, Norway, Denmark, US and France

Note: The figures show the change in the time spent at different places provided by Google's COVID-19 Community Mobility Report. The data is drawn from users who have opted-in to Location History for their Google Account and the baseline is the median value, for the corresponding day of the week, during the 5-week period Jan 3–Feb 6, 2020. The figures show mobility trends for (a) places like restaurants, cafes, shopping centers, theme parks, museums, libraries, and movie theaters; (b) places like public transport hubs such as subway, bus, and train stations; (c) places like local parks, national parks, public beaches, marinas, dog parks, plazas, and public gardens; (d) places like grocery markets, food warehouses, farmers markets, specialty food shops, drug stores, and pharmacies; (e) places of work; (f) places of residence. The data and more information can be found at https://www.google.com/covid19/mobility/. Because the location accuracy and the understanding of categorized places varies from region to region, some cautions is warranted when interpreting the cross-country differences.

(a) Number of deaths

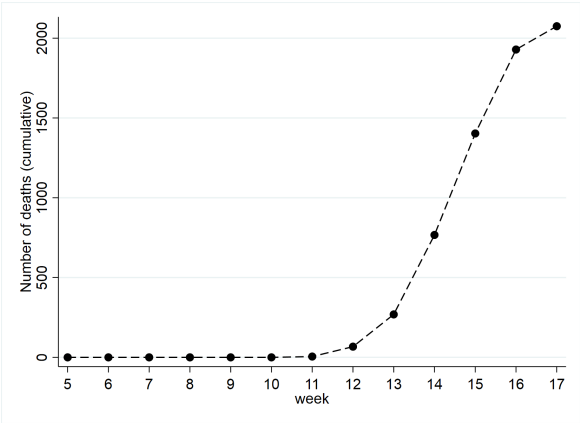

(b) Layoff notifications

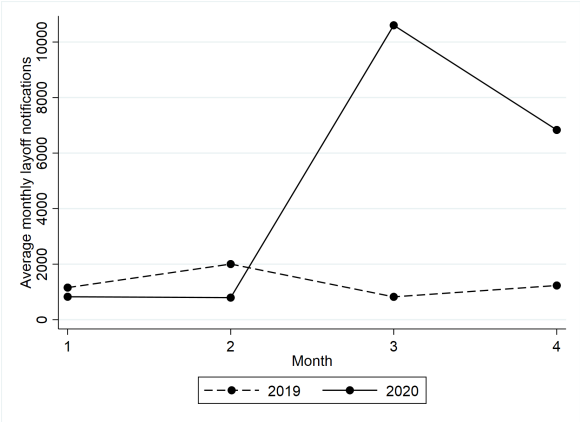

(c) Inflow into unemployment

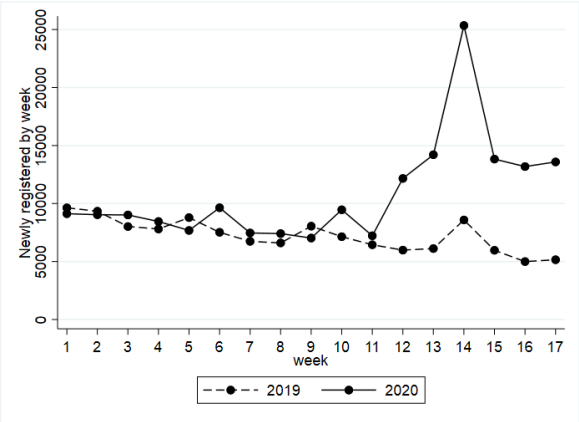

Figure A2: Trends in COVID-related deaths, layoff notifications and inflow of unemployed
Note: In this figure we show (a) the cumulative number of COVID-related deaths according to the Swedish National Board of Health and Welfare; (b) the average weekly number of people who was notified by month (the information is only available at a monthly basis, and the April observation reflects the first three weeks in April) and (c) the number of newly registered job seekers at the Public Employment Service by week in 2019 and 2020.

Figure A3: Percent change in vacancy inflow by 1-digit industry before and after week 10 relative to the percent change during the same period 2019.

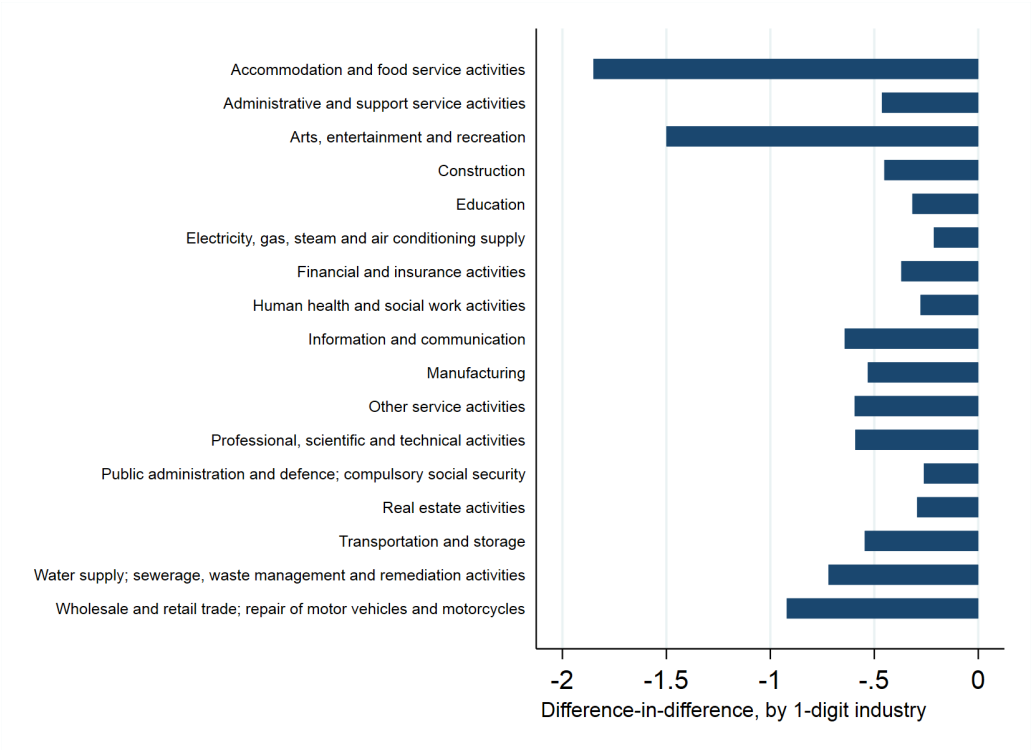

Figure A4: Average daily inflow of vacancies per week in essential and not essential industries

(a) Frequency

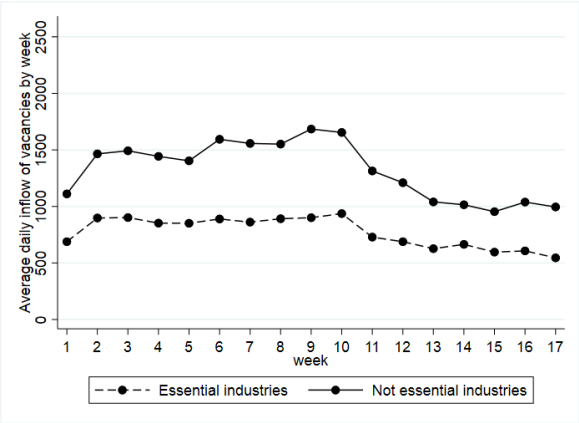

(b) Rates

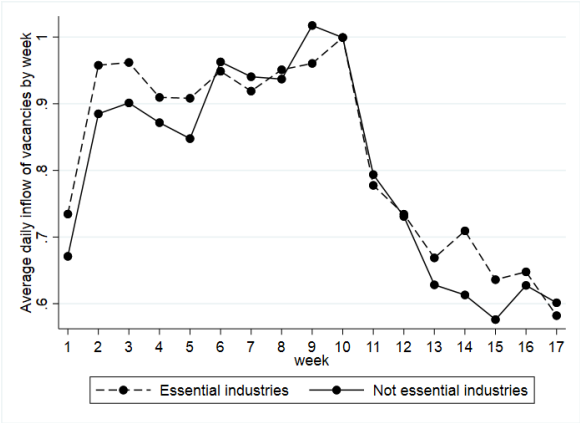

Note: Vacancies have been categorized using the Swedish industry classification SNI. Essential industries are the ones declared essential by the Swedish government: ”Electricity, gas, steam and air conditioning supply”, ”Financial and insurance activities”, ”Wholesale and retail trade; repair of motor vehicles and motorcycles”, ”Manufacturing”, ”Human health and social work activities”, ”Information and communication”, ”Water supply; sewerage, waste management and remediation activities”, ”Public administration and defence; compulsory social security” and ”Transportation and storage”.

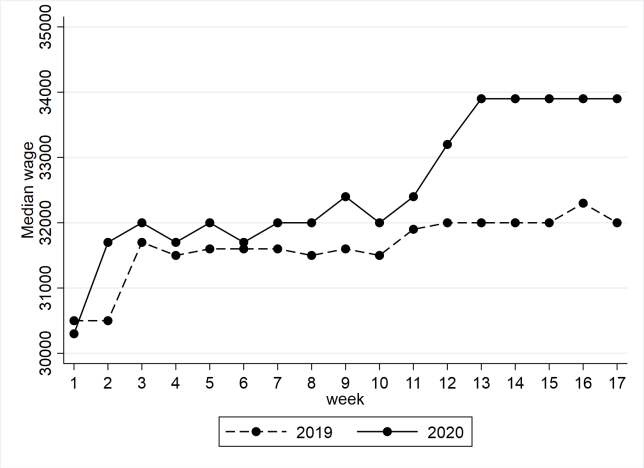

Figure A5: Median wages in occupation wages

Note: We start by computing the average wage for each occupation prior to the COVID crisis, what we name the occupation wage. The figure shows the median of the occupation wage in vacancy posted during a given week, in 2019 and 2020.

(a) Mean home working (D&N)

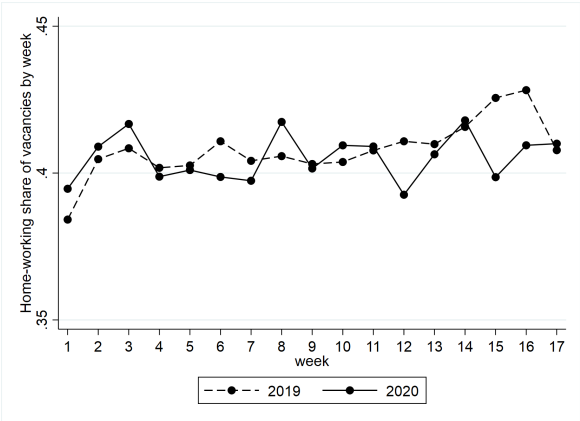

(b) Mean home working (MPW)

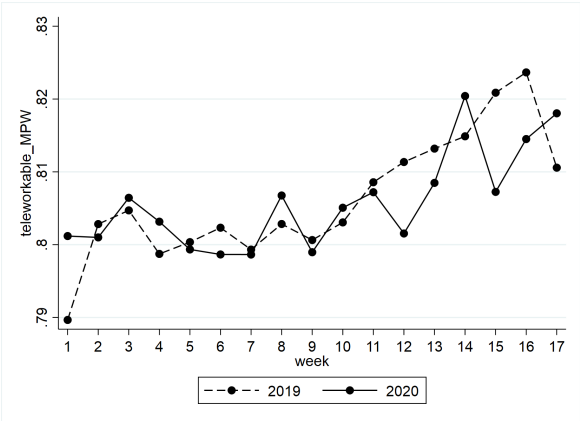

Figure A6: Alternative home working shares

Note: The figure shows how the home working shares of vacancies evolves on the website. In (a) we use the home working index from in Dingel and Neiman (2020) and in (b) the home working index from Mongey et al. (2020). In both cases, the index is based on the description of tasks in O*NET.

Figure A7: Number of daily clicks in 2020 (right-hand panel) and in 2019 (left-hand panel)

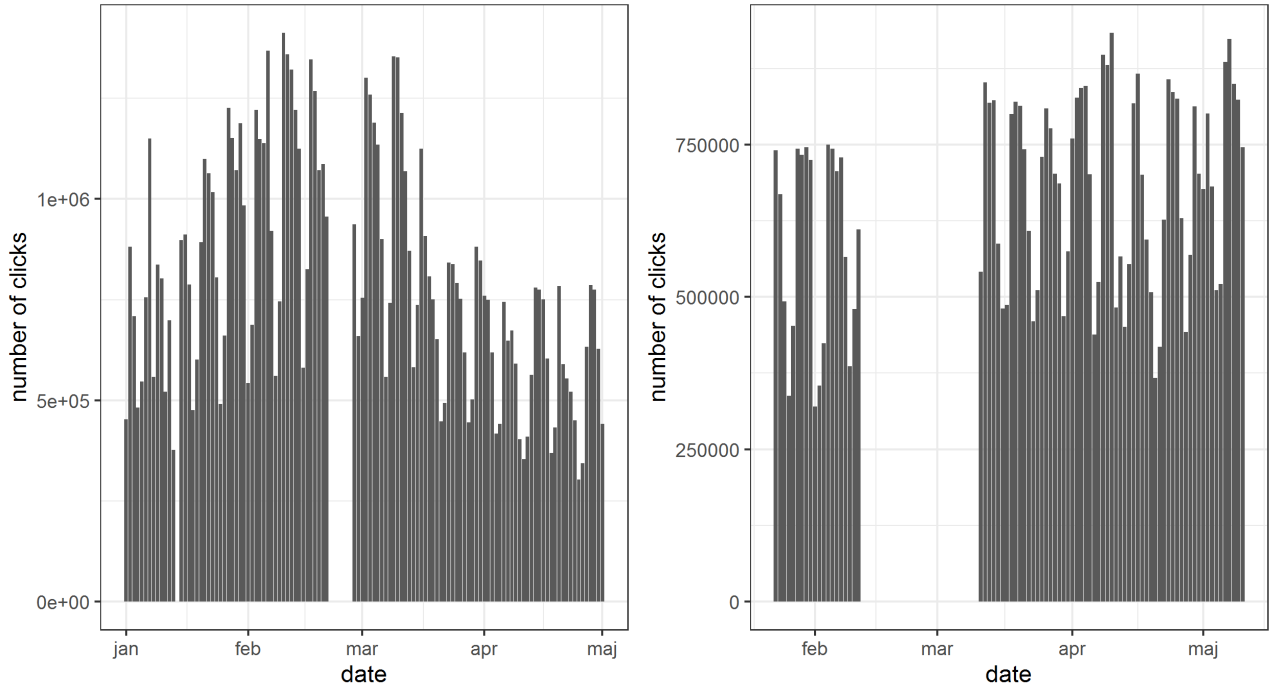

Note: These figure shows the number of clicks per day. Samples: Left hand side: Clicks between Jan 1st and April 11th 2020. Right hand side: Clicks between Jan 1st and May 31st 2019.

Figure A8: Regression of the number of clicks on week dummies **(total margin)**, without user fixed effects (left panel) and with fixed effects (right panel), clicks per user (top panel) and clicks per vacancy (bottom panel)

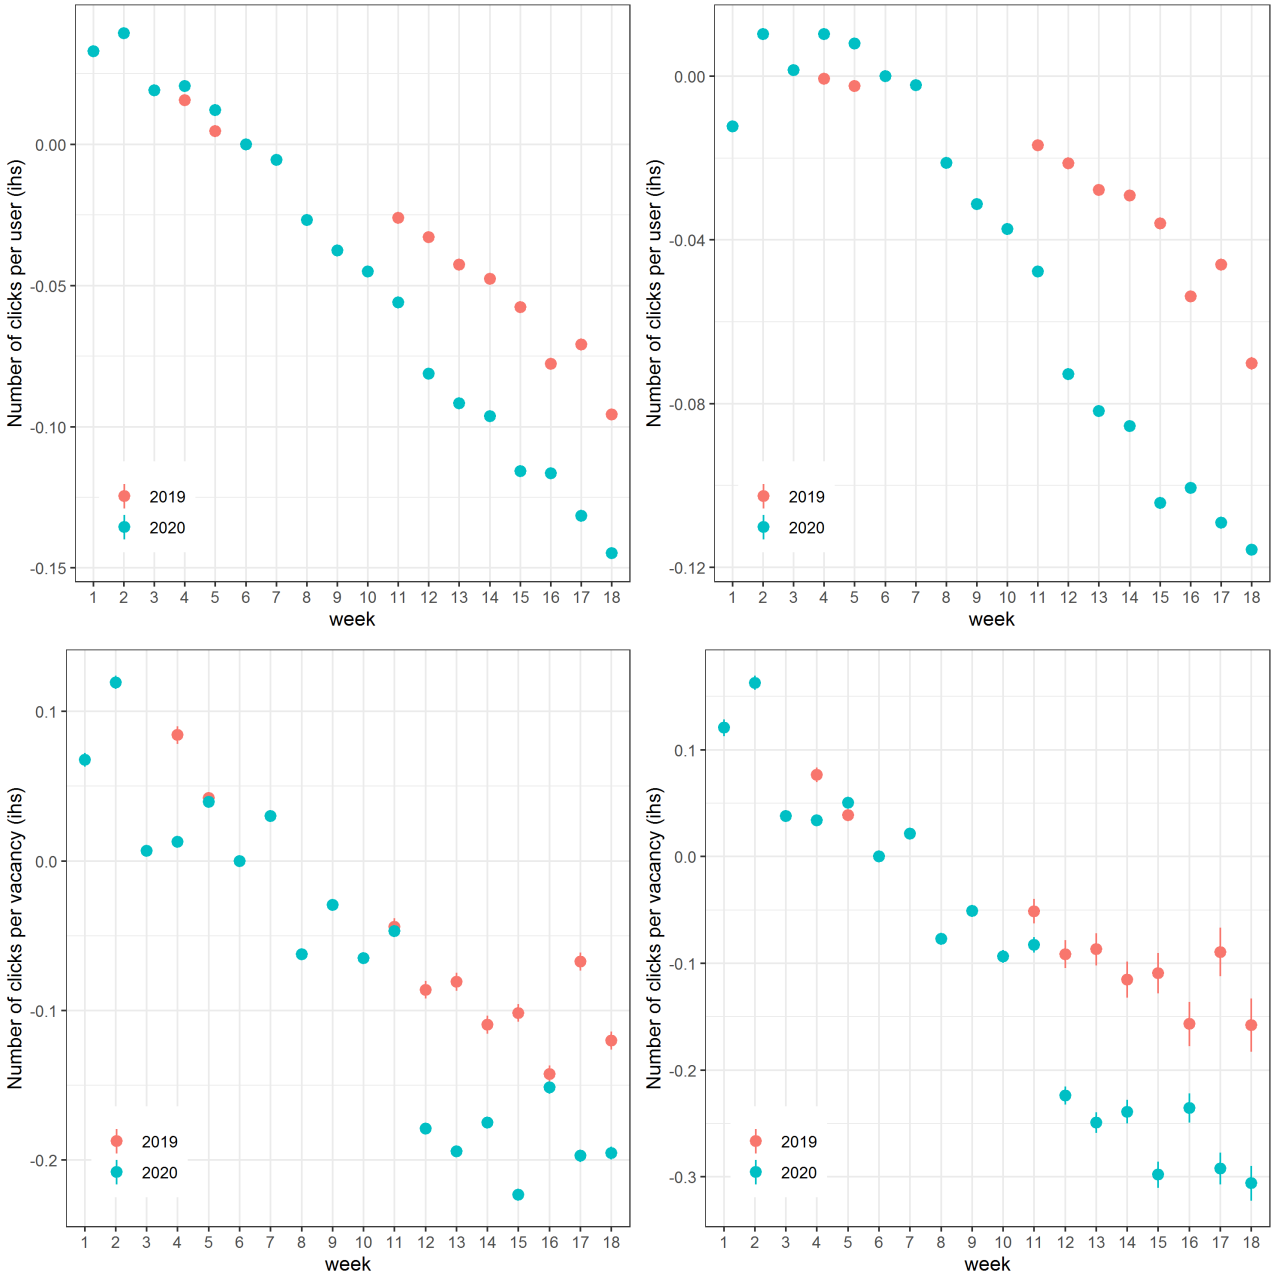

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows the weekly effects on the number of clicks per user (top panels) and of clicks per vacancy (bottom panels). We take the inverse hyperbolic sine of the number of clicks, thus including user x days and vacancy x days with zero clicks (total margin). The plotted estimates correspond to the coefficients β_w in Regressions (1) and Regression (2). Blue dots are for 2020 estimates, and red dots for 2019 estimates. The 6th week of the year is chosen as the reference. Standard errors are clustered at the user level in per user regressions, and at vacancy level in per vacancy regressions. Vertical lines show 95% confidence intervals.

Figure A9: Regression of the targeted wage (top left-hand panel), of the healthcare occupation dummy (top right), of the home-working index (bottom left) and of the targeted expansion index (bottom right) on week dummies, **without user fixed effects**

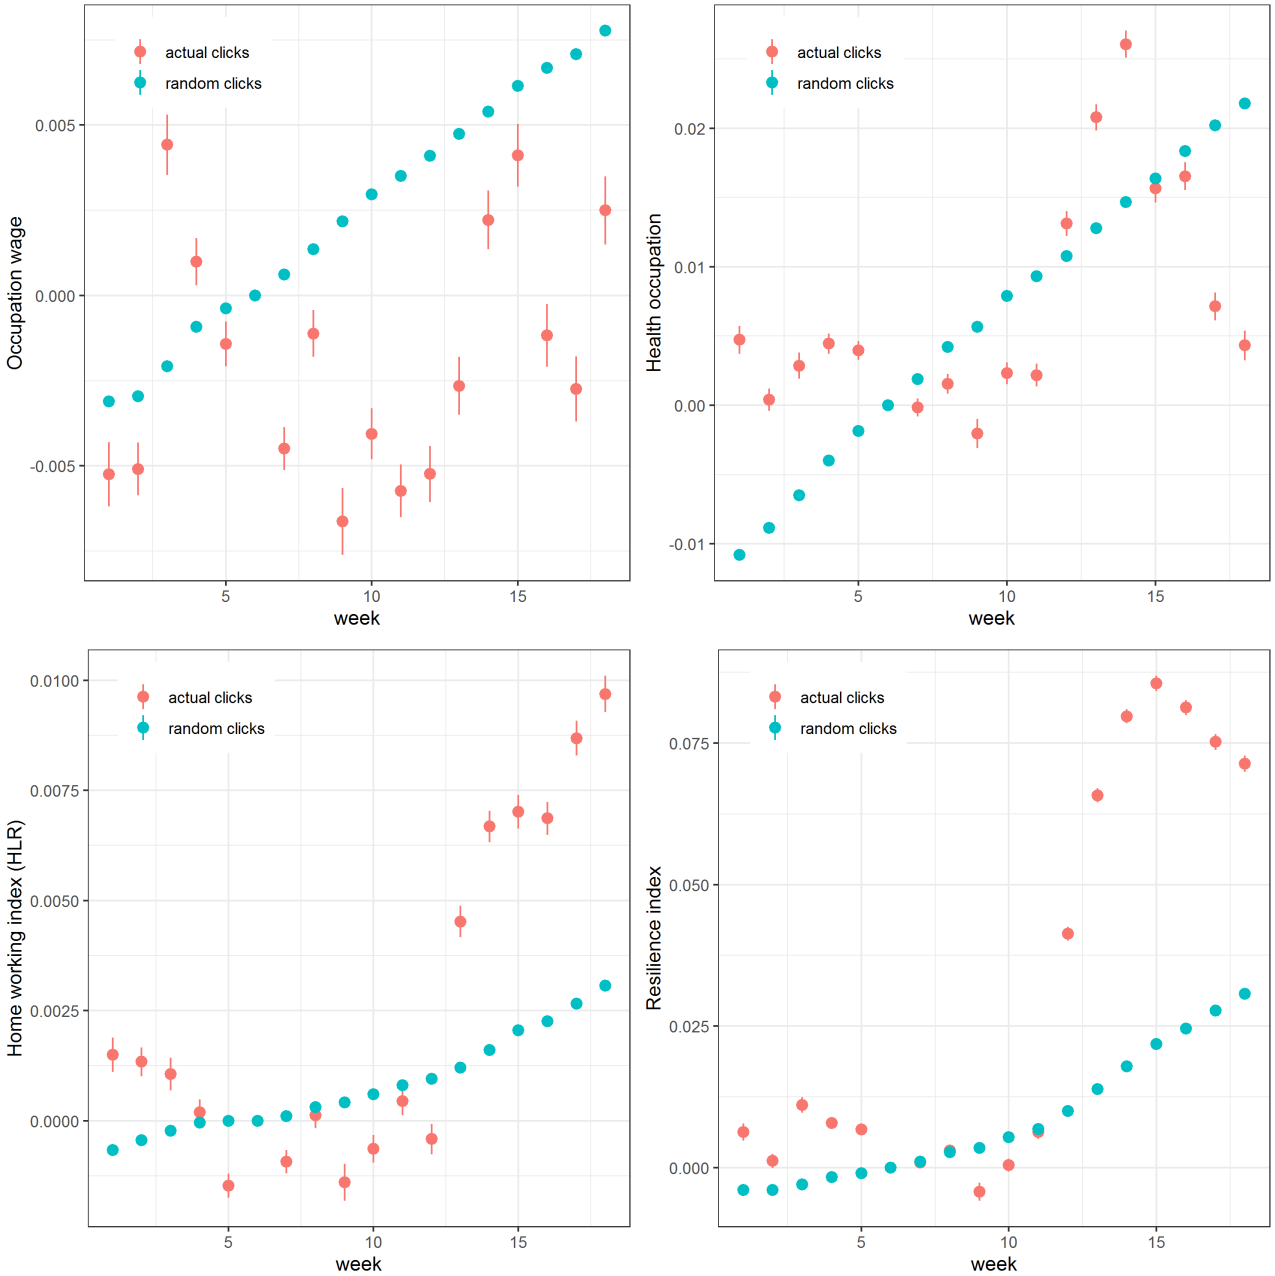

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows the weekly effects on the characteristics of the vacancies that users click on: average wage of the posted occupation (targeted wages), healthcare occupation, home-working index and resilience index. The plotted estimates correspond to the coefficients β_w in Regressions (1) without user fixed effects. Blue dots are for 2020 estimates on actual clicks, and red dots for clicks by counterfactual random searchers on available vacancies. The 6th week of the year is chosen as the reference. Standard errors are clustered at the user level. Vertical lines show 95% confidence intervals.

Figure A10: Regression of the MPW home-working index (top panel) and the DN home-working index (bottom panel) on week dummies: without fixed effects (left panel) and with fixed effects (right panel)

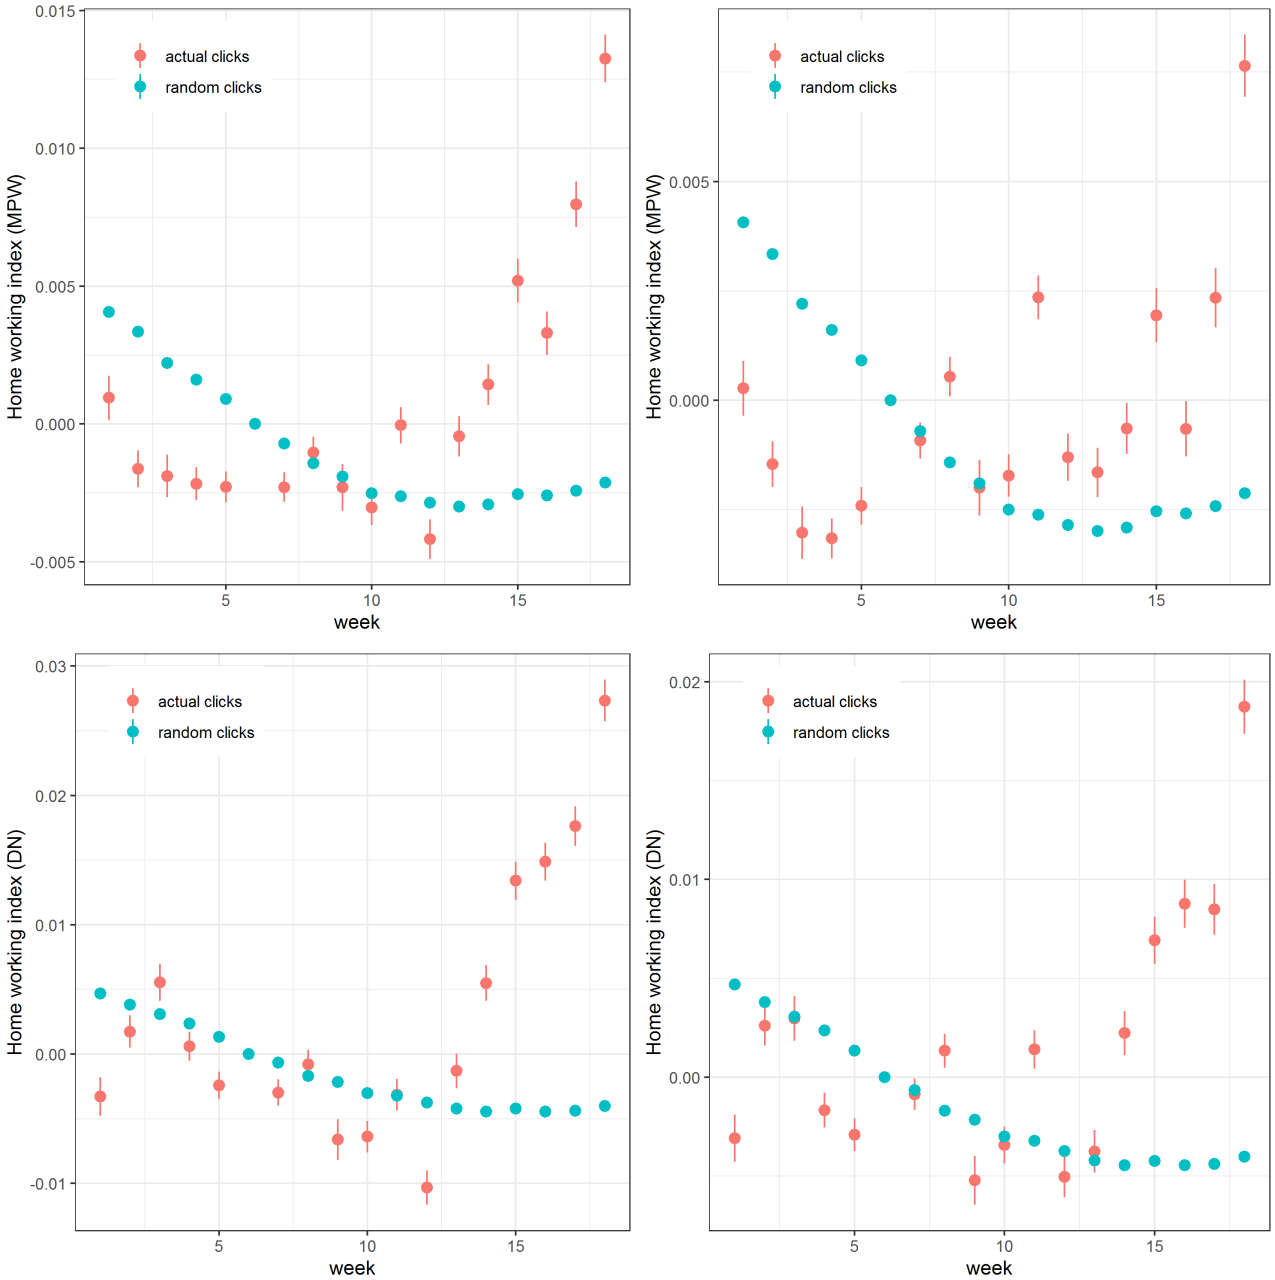

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows the weekly effects on the home-working index of the vacancies that users click on. The top panels use the home-working index in Mongey et al. (2020), the bottom panel the index in Dingel and Neiman (2020). The plotted estimates correspond to the coefficients β_w in Regressions (1) with or without user fixed effects. Blue dots are for 2020 estimates on actual clicks, and red dots for clicks by counterfactual random searchers on available vacancies. The 6th week of the year is chosen as the reference. Standard errors are clustered at the user level. Vertical lines show 95% confidence intervals.

Figure A11: Regression of the log number of clicks per vacancy on week dummies interacted with high home-working occupations (left) and with resilient occupations (right): **total margin**

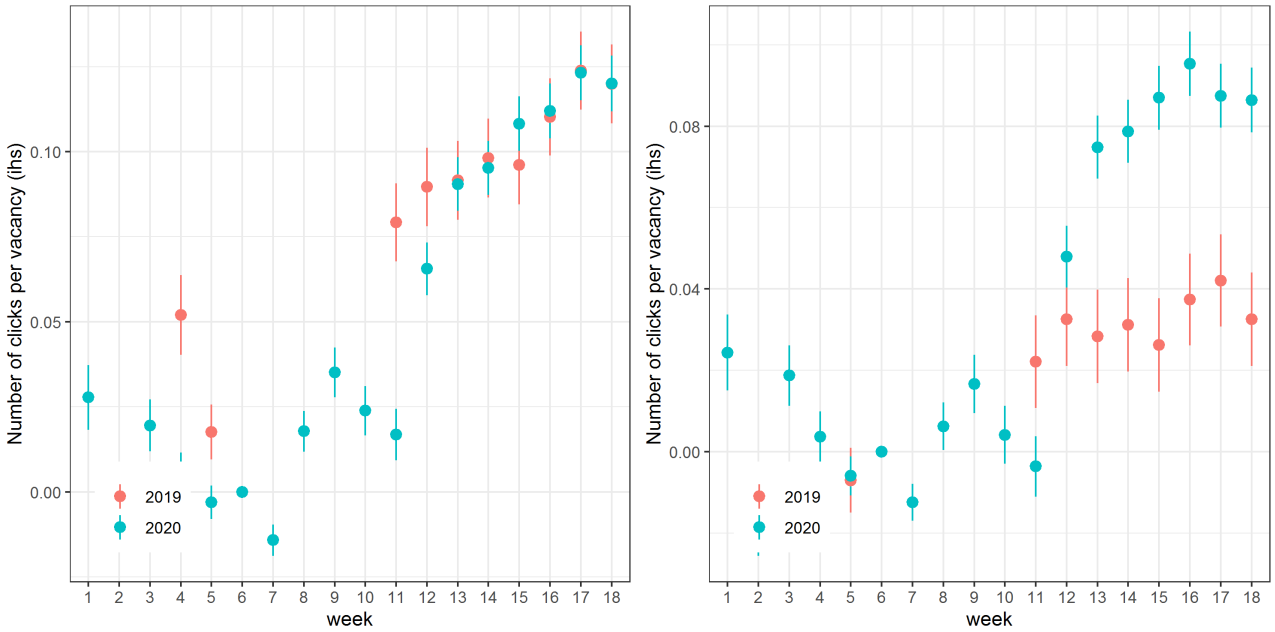

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows coefficient estimates from Regression (3) of the inverse hyperbolic sine of the number of clicks per vacancy (total margin including days without any clicks). We plot the δ_w coefficients of the week effects interacted with high home-working occupations (left) and with resilient occupations (right). Blue dots are for 2020 estimates, and red dots for 2019 estimates. The 6th week of the year is chosen as the reference. Standard errors are clustered at vacancy level. Vertical lines show 95% confidence intervals.

Figure A12: Regression of the log number of clicks per vacancy on week dummies interacted with high home-working occupations (left) and with resilient occupations (right): **with vacancy fixed effects**

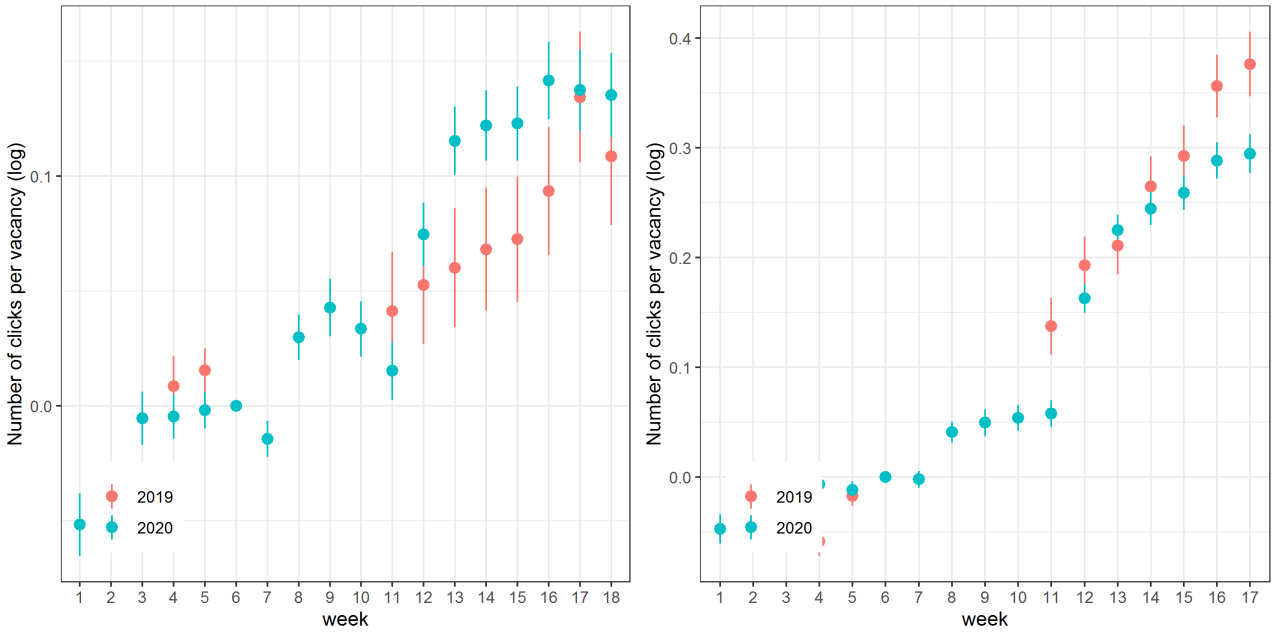

Sample: clicks on *Platsbanken* between between Jan 1st and May 3rd in 2020 and in 2019.
Note: This figure shows coefficient estimates from Regression (3) of the log of the number of clicks per vacancy, with vacancy fixed effects. We plot the δ_w coefficients of the week effects interacted with high home-working occupations (left) and with resilient occupations (right). Blue dots are for 2020 estimates, and red dots for 2019 estimates. The 6th week of the year is chosen as the reference. Standard errors are clustered at vacancy level. Vertical lines show 95% confidence intervals.

B Extra Tables

Table B1: Difference-in-difference estimates

Outcome:	(1) ln(Vacancies)	(2) ln(Jobs)	(3) By essential status	(4) ln(unempl.)
Post	-0.4070***	-0.3441***	-0.3955***	0.7308***
	(0.0489)	(0.0563)	(0.0423)	(0.1284)
Essential industry			-0.5243***	
			(0.0144)	
Essential industry × Post			-0.0027	
			(0.0408)	
Observations	34	34	68	34
R-squared	0.9551	0.9426	0.9792	0.8859

Note: The table presents difference-in-difference estimates from the following model: $ln(Inflow_{w,y}) = \theta(Treat_y \times Post_w) + \lambda_w + \lambda_y + \epsilon_{wy}$, where $Inflow$ is the inflow of vacancy ads (col. 1 and 3), jobs (col. 2) and unemployed (col. 4); *post* is a dummy taking the value one if the week is between week 11-17 (the comparison is week 1-10); *Treat* is a dummy taking the value one if $year = 2020$ and λ_w and λ_y are week and year fixed effects. Essential industries are the ones declared essential by the Swedish government: "Electricity, gas, steam and air conditioning supply", "Financial and insurance activities", "Wholesale and retail trade; repair of motor vehicles and motorcycles", "Manufacturing", "Human health and social work activities", "Information and communication", "Water supply; sewerage, waste management and remediation activities", "Public administration and defence; compulsory social security" and "Transportation and storage". Robust standard errors are reported in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Table B2: Ten most resilient and least resilient occupations: difference-in-difference estimation

Occupation label	**Estimated change in vacancy inflow Pre- vs. Post week 10, 2019 vs. 2020**	**std. err**
Ten least resilient		
Waiters and bartenders	-2.252	(0.215)
Dentists	-1.258	(0.243)
Fast-food workers, food preparation assistants	-1.230	(0.152)
Shop staff	-1.229	(0.194)
Culinary associate professionals	-1.133	(0.199)
Hairdressers, beauty and body therapists	-1.128	(0.268)
Postmen and postal facility workers	-1.105	(0.293)
Athletes, fitness instructors and recreational workers	-1.035	(0.435)
Cooks and cold-buffet managers	-1.007	(0.153)
Dental nurses	-0.996	(0.244)
Ten most resilient		
Authors, journalists and linguists	0.075	(0.278)
Specialists in health care not elsewhere classified	-0.058	(0.237)
Nursing professionals	-0.074	(0.167)
Personal care workers in health services	-0.105	(0.159)
University and higher education teachers	-0.132	(0.174)
Primary- and pre-school teachers	-0.144	(0.120)
Medical and pharmaceutical technicians	-0.145	(0.198)
Social work and counselling professionals	-0.180	(0.122)
Electrical equipment installers and repairers	-0.182	(0.148)
Car, van and motorcycle drivers	-0.191	(0.336)

Note: This table reports the 10 most/least shrinking occupations when we rank occupations according to their coefficients from the following model:
$ln(Inflow_{w,y}) = \theta(Treat_y \times Post_w \times Occ_o) + \lambda_{yo} + \lambda_{wo} + \lambda_{wy} + \epsilon_{wyo}$, where $Inflow$ is the inflow of vacancy ads; *Post* is a dummy taking the value one if the week is between week 11-17 (the comparison is week 1-10); *Treat* is a dummy taking the value one if $year = 2020$; Occ_o are occupation dummies; and λ_{yo} are year$\times$occupation fixed effects; λ_{wo} are week$\times$occupation fixed effects and λ_{wy} are week$\times$ year fixed effects. To make the table informative, we require that the pre-COVID vacancy share is at least 0.2 percent.

Table B3: Difference-in-difference estimates: vacancy types

	(1) log(Wage)	(2) Health occupations	(3) Home working HLR	(4) Home working D&N	(5) Home working MPW
Post week 10×2020	0.0151***	0.0473***	0.0055***	-0.0096*	-0.0053
	(0.0039)	(0.011)	(0.0018)	(0.0045)	(0.0033)
Observations	34	34	34	34	34
R-squared	0.9271	0.8411	0.9089	0.7964	0.8522

Note: The table presents difference-in-difference estimates from the following model: $y_{w,y} = \theta(Treat_y \times Post_w) + \lambda_w + \lambda_y + \epsilon_{wy}$, where y is given by each column; *post* is a dummy taking the value one if the week is between week 11-17 (the comparison is week 1-10); *Treat* is a dummy taking the value one if $year = 2020$ and λ_w and λ_y are week and year fixed effects. Column (3) shows the association with the share of working time spent at home calculated from ATUS in Hensvik et al. (2020) and in columns (4) and (5) we use the alternative measures relying on the description of tasks in O*NET used in Dingel and Neiman (2020) and Mongey et al. (2020). Robust standard errors are reported in parentheses *** p<0.01, ** p<0.05, * p<0.1.