Economic Shocks and Skill Acquisition: Evidence from a National Online Learning Platform at the Onset of COVID-19*

Ina Ganguli† Jamal Ibrahim Haidar‡ Asim Ijaz Khwaja§ Sam Stemper¶

Basit Zafar||

March 25, 2022

Abstract

We study how large shocks impact individuals' skilling decisions using data from the largest online learning platform in Saudi Arabia. The onset of the COVID-19 pandemic brought about a massive increase in online skilling, and demand shifted towards courses that offered skills, such as telework, likely to be immediately valuable during the pandemic. Consistent with a model where individuals trade off reskilling costs with their expectations of future labor market conditions and their duration of work, we find that shifts into telework courses were largest for older workers. In contrast, younger workers increased enrollments in courses related to new skills, such as general, occupation-specific, and computer-related skills. Using national administrative employment data, we provide suggestive evidence that these investments in skills in early 2020 helped users maintain employment over the course of the pandemic.

*This paper was funded through support from the Ministry of Labor and Social Development (MLSD) and the Human Resources Development Fund (HRDF) of the Kingdom of Saudi Arabia. We would like to thank Dr. Jomah al-Hamid, Mr. Mohammed Al-Shewaier, Mr. Abdullah Qasim, Ms. Shahera Al-Zahib, Ms. Hana Al-Belihees, and Mr. Mohammed Al-Salman at HRDF for their continued support and valuable input, without which this project would not have been possible. We are thankful for the support of additional staff at HRDF and NLO who were instrumental in facilitating this project, including Abdullah Al-Sabani, Sameer Alamri, Anas Al-Kharashi, Hadeel Al-Abdulwahed, Alanoud Al-Dous, and Mohammed Al-Mughaiseb. Ghaida Alanazi, Luma Akkad, Yazen Kashlan, Soha Shami, Fay Terrett, and Gabriela Lecaro provided excellent research assistance.

†University of Massachusetts Amherst. Email: iganguli@umass.com

‡Harvard Kennedy School. Email: Jamal_Haidar@hks.harvard.edu

§Harvard Kennedy School. Email: khwaja@hks.harvard.edu

¶Harvard Kennedy School. Email: stemper@g.harvard.edu

||University of Michigan. Email: basitak@gmail.com

1 Introduction

COVID-19 is the most rapid and widespread labor market shock in recent history. In a matter of weeks in early 2020, work changed dramatically throughout the globe. In the short-run, many jobs went remote as governments initiated social distancing measures and individuals opted to work from home voluntarily.[1] It is yet unclear the degree to which the COVID shock will bring about long-term changes in the labor market. Both these short and long-term affects likely have important implications for skill acquisition.

In this paper, we study the impact of COVID-19 on skill acquisition by utilizing data from Doroob, the largest online learning platform in Saudi Arabia. This setting has several features that make it well-suited to examine the many potential avenues through which the COVID-19 shock could affect skill development. First, the high-frequency nature of data captured by Doroob, combined with the rapid spread of COVID-19 in Saudi Arabia and similarly rapid government response (Algaissi et al., 2020), enables us to identify precisely how activity on Doroob changed in response to this unprecedented health and economic shock. Second, most of the courses on Doroob are fairly short (most courses are estimated to take roughly three hours to complete), so our empirical work allows us to characterize a fine gradation of individual behavioral responses, in the form of weekly changes in course preferences over time. Third, Doroob offered courses on teleworking skills prior to the pandemic. These courses were motivated by a broader national initiative aimed at facilitating flexible work arrangements for job seekers, particularly women.[2] As work shifted online, courses on teleworking, arguably the skill with the most immediate need in response to the COVID shock, were available to interested Saudis. Finally, use of the platform was widespread—with over two million users (in a country with a population of 34 million) prior to the COVID shock—and users of the platform are diverse along a number of dimensions, such as age, gender, education, and employment

[1] Brynjolfsson et al. (2020), Dingel and Neiman (2020).

[2] Saudi to boost 'decent and proper' jobs for women and the disabled, Alarabiya News, March 15, 2017.

status. While most prior work has focused on the relationship between economic shocks and human capital among an age-specific group of individuals,[3] our work is unique in that it identifies how a large economic shock affects individuals at different points in their life-cycle–roughly half of the individuals in our analyses are age 25 or older.

Our empirical work compares user behavior in the months before the onset of COVID-19 in Saudi Arabia to the weeks immediately after. We compare this difference to differences in behavior over the same periods in 2018 and 2019. This setup is analogous to a difference-in-differences approach; the effect of the COVID-19 shock is identified based on differences in behavior over time in 2020 versus the same differences in prior years. Identification rests on a parallel trends assumption that, absent the COVID-19 shock, user behavior in March and April 2020 would have evolved similarly to March and April in the prior two years. Trends in enrollments in the weeks preceding the COVID-19 shock support this assumption.

We find that the COVID shock led to a large increase in activity on the platform, along every margin we examine. Aggregate daily data indicates that daily user registrations increased from roughly 200 per day to over 2,000 at their peak in April 2020. New, COVID-induced users were also more engaged on the platform than prior registrants, enrolling in approximately 0.4 more courses in their first week on the platform, relative to a pre-COVID mean of 2.1 courses per user. This cohort of new users was also more persistent on the platform, exhibiting relatively larger increases in their second and third weeks on Doroob. Pre-existing Doroob users also increased their usage of the platform substantially. Relative to identically-defined cohorts of users in prior years, COVID induced a 150 percent increase in weekly course enrollments.

Interestingly, this increased activity was also accompanied by systematic changes in which courses were popular post-pandemic and these changes differed further by user

[3] For example, middle school students in Adukia et al. (2020), high school students in Atkin (2016), or college students in Charles et al. (2018) and papers described in Patnaik et al. (2020).

age. We find that telework courses, a relatively less popular set of courses prior to COVID-19, exhibited larger increases in enrollments than all others: increases of over 1200 percent among both new and existing users. These increases were driven by relatively older Doorob users, particularly users above 40 years old. These age effects are large; compared to prior cohorts, new users over age 40 enrolled in 0.18 more telework courses during their first week on the platform, while for users age 18 to 24 it was 0.04 courses. The opposite is true of other courses; increases in general skills, occupation-specific, and computer courses were driven by younger Doroob users, particularly users between 18 and 25 years old.

Additionally, we provide evidence of heterogeneity in responses by gender and employment status. Among women, increase in enrollments were concentrated in computer and occupation-specific courses, whereas men increased enrollments most in telework courses. Moreover, users who were either employed or students exhibit relatively larger increases in course-taking than users who were jobseekers.

Finally, we provide suggestive evidence that these investments in skills during early 2020 helped Doroob users maintain employment over the course of the pandemic. To do so, we link user enrollments data with national administrative data on private sector employment, and compare employment rates between two sets of prior Doroob users: users who responded to the pandemic by taking additional courses and users who did not. We find that the prior Doroob users who returned to the platform in 2020 were more likely to be employed in the private sector in September 2020, even after controlling for a rich set of covariates, including prior course enrollments. However, we caution a causal interpretation of these estimates, as our variation in post-COVID course-taking likely reflects unobservable user characteristics that may covary with employability.

We motivate our findings through a life-cycle model of skill acquisition. In the model, agents invest in different skills over time based on the anticipated future value of these skills in the labor market. As in canonical models of human capital acquisition, younger agents invest more heavily in acquiring skills; they have a longer time horizon to reap the benefits of such investment. When skill prices vary over time—for example, telework skills may be particularly valuable in the short-run, and relatively less valuable in the long-run—older users invest relatively more in skills that are more immediately valuable. Younger users, on the other hand, invest relatively more in skills that are valuable over the long-run.

Our work relates to a long theoretical literature on skill acquisition over the life-cycle, dating back to Becker et al. (1964) and Ben-Porath (1967), both of whom emphasize the role of age in human capital investment. In these models, human capital investment declines with age because older agents have less time to earn returns on these investments. More recently, Cunha and Heckman (2008), Sanders and Taber (2012), and Cavounidis and Lang (2020) have offered theoretical models of multi-dimensional human capital acquisition, which allow for investments in different skills over time: cognitive versus non-cognitive skills or firm-specific versus general skills, for example. Our results highlight how age and multi-dimensional skill acquisition interact, an interaction highlighted specifically by Cavounidis and Lang (2020). In this setting, the COVID-19 shock induced both older and younger workers to acquire new skills, but older workers directed their skill acquisition efforts towards teleworking courses—courses that would allow them to apply their existing skill set rather than develop skills for a new occupation or career entirely.

Empirically, related literature has focused on educational responses to economic shocks. Many of these papers analyze how economic shocks, such as changes in the skill intensity of exports (Atkin, 2016; Blanchard and Olney, 2017) or shocks to local labor markets (Charles et al., 2018; Adukia et al., 2020) affect levels of educational achievement (e.g. high school dropout rates, college attendance). Generally, this literature finds that individuals increase their educational investment in response to increases in the skill intensity of local labor demand, and vice versa. However, these papers typically estimate changes to levels of human capital investment, rather than the skill content of investments. Our setting allows us to estimate how the COVID-19 shock affected both the amount of skill acquisition (as reflected in Doorob enrollments) as well the composition of skills acquired.

Patnaik et al. (2020) review the literature on college majors, the most well-studied form of heterogeneous human capital. Much of this literature focuses on estimating the elasticity of college major choice with respect to earnings. The subset of this literature that focuses on changes in demand-side labor market returns includes Beffy et al. (2012), Long et al. (2015), Blom et al. (2021), Bedard and Herman (2008), and Abramitzky et al. (2019). Our work differs from these studies in that we focus on a particular economic shock—COVID-19—and focus on a broader population of students, jobseekers, and workers, rather than only college students. This focus allows us to compare educational responses across age groups—comparisons that prove to be empirically relevant in our setting.

Finally, our work relates to the literature on gender, human capital, and occupational decisions. Many papers document large differences in occupational preferences between men and women (e.g. Mas and Pallais, 2017). For example, Wiswall and Zafar (2018) demonstrate that, among American college students, these preferences affect their choice of major (and ultimately their occupational choice and wages). In the Saudi context, these differences may be particularly noteworthy, given the changing role of Saudi women in the labor force, as documented in Miller et al. (2019), Aloud et al. (2020), and Bursztyn et al. (2020).

We proceed as follows. Section 2 describes the onset of COVID-19 in Saudi Arabia, the platform we study, and provides some theoretical motivation. Section 3 describes our data and methodology. Section 4 outlines our results, and Section 5 concludes.

2 Setting

2.1 COVID-19 in Saudi Arabia

Saudi Arabia's experience with COVID-19 mirrors that of many countries that experienced a large-scale outbreak starting in March 2020. The Saudi Ministry of Health reported the first case of COVID-19 in the country on March 2, 2020.[4] This first case originated abroad but community spread within Saudi Arabia accelerated in the weeks that followed. While it took 13 days for Saudi Arabia's cumulative case count to rise from 1 to 100, it took less than one month to reach 1,000 cumulative cases and less than two months to reach 20,000 cumulative cases. Panel A of Figure 1 compares the series of new COVID-19 cases in Saudi Arabia to the United Arab Emirates,[5] and the United States.

Over the same period, the Saudi government implemented a series of increasingly stringent measures aimed to stop community spread of the virus. On March 16, Saudi Arabia issued a stay-at-home order and closed government offices.[6] On March 20, Saudi Arabia suspended domestic flights and mass land transport.[7] Two days later, King Salman announced a nation-wide curfew restricting movement between 7 p.m. and 6 a.m.[8] Restrictions on travel and recreational activities continued on and off over the next year.

[4] MOH Reports First Case of Coronavirus Infection, Saudi Ministry of Health, March 2, 2020.

[5] The UAE neighbors Saudi Arabia and is also a member of the Gulf Cooperation Council. Among GCC members, Saudi Arabia is the most populous and the UAE is second.

[6] Kingdom's government decides to suspend attendance at workplaces in all government agencies for period of (16) days except for health, security, military and electronic security center, Saudi Press Agency, March 16, 2020.

[7] Saudi Arabia suspending domestic flights, mass land transport in fight against COVID-19, Arab News, March 20, 2020.

[8] Saudi Arabia reports 51 new cases, total now 562, Saudi Gazette, March 24, 2020.

Saudi Arabia lifted many restrictions in March 2021,[9] and as of June 2021, more than 450,000 people in Saudi Arabia had tested positive, more than 7,400 died, and nearly 15 million vaccine doses had been administered.[10]

In an attempt to mitigate the economic impact of the pandemic, the Saudi government announced a number of support packages. Specifically, the government devoted $61 billion to target the private sector, including a 60 percent wage subsidy for Saudi private sector salaries, postponement of government dues, and support for the banking and small and medium enterprises.[11] Despite these efforts, the Saudi economy experienced its largest contraction in more than three decades, shrinking by 4.1 percent in 2020.[12] At its peak in the second quarter of 2020, Saudi unemployment stood at 15.4%, 3.1 percentage points higher than in the same period of the prior year.[13] More recently, unemployment rates have fallen and economic growth followed large swings in the oil sector, falling during the first quarter of 2021 and rebounding in the two quarters thereafter.[14]

2.2 The Doroob E-Learning Platform

We study how the COVID-19 shock affected skill acquisition over this period of rapid change and uncertainty. Our setting is Doorob, the largest online learning platform in Saudi Arabia.

Doroob is a national e-learning and skills training platform designed to enable mil-

[9] With the virus relatively contained, Saudi Arabia lifts most pandemic restrictions., New York Times, March 7, 2021.

[10] Ritchie et al. (2020)

[11] Kingdom of Saudi Arabia; Government and institution measures in response to COVID-19., KPMG, November 18, 2020.

[12] Saudi Economy Grew 2.8% in Fourth Quarter As Covid Impact Eased, Bloomberg, February 10, 2021.

[13] Saudi unemployment spikes as virus-hit economy shrinks by 7% in second-quarter, Reuters, September 30, 2020. Labor market statistics Q2 2020, Saudi General Authority for Statistics

[14] Saudi Arabia: Unemployment falls to lowest in five years, AlJazeera, June 30, 2021. Saudi Economy Grew at The Fastest Rate in Nearly a Decade, Bloomberg, November 9, 2021.

lions of Saudis of varying backgrounds to upgrade their skills and become more employable.[15] The platform is available at https://doroob.sa/ and offers free online courses in Arabic and English through which users can build skills across a diverse set of domains, including sector-specific skills—such as telecommunication, retail, and hospitality—as well as occupation-specific courses—related to secretarial skills, photography and insurance representative, for example. Doroob also offers courses on general and interpersonal skills, including English language and leadership skills. In addition, the platform offers certificates to allow users to demonstrate their progress on Doroob and to help employers in assessing and verifying worker skills. These certificates are recognized by various employers in Saudi Arabia, and are meant to help users during their job search process. The Doroob platform is sponsored by the Saudi Human Resources Development Fund.

The platform has expanded significantly in the last few years. Since August 2014, nearly two million unique users have registered for Doroob—a non-trivial share of the Saudi population of 34 million.[16] Over the same period, these users account for over 30 million course enrollments. Panel B of Figure 1 displays cumulative users and enrollments on the platform, after implementing a set of sample restrictions described below. As of 2020, Doroob offered nearly 400 unique courses, spanning a wide range of skills. The number of courses offered on the platform has grown over time from 122 in 2016 to 258 in 2020. The series of unique courses over time is shown in Figure 2, alongside the number of enrollments in each course type. Changes in course composition and availability over time suggests that selection of users into the platform may change over time; our analysis below will take this into account.

Users on Doroob include students, jobseekers, and the currently employed. To better understand the population of Doorob users, we conducted a survey of roughly 1,000 new

[15]Saudi citizens over the age of 18 are eligible to enroll on the platform. Doroob Help: Common Questions.

[16]Population, total - Saudi Arabia, The World Bank.

Doroob users in February 2020, which provides useful background on the characteristics and motivations of new Doroob users who joined the platform prior to the COVID shock. We solicited survey responses via a pop-up on the Doroob website during the registration process and via direct emails to new Doroob registrants. Summary statistics for key variables are shown in Table 1 and discussed below.

Roughly half of Doroob users are women, and most users are young, with an average age of 27. As we show later, these characteristics are generally consistent with the broader Doroob population, based on demographic characteristics captured during the Doorob registration process.

A small share of our survey respondents, roughly 6 percent, were unemployment insurance (Hafiz) beneficiaries. The Hafiz unemployment insurance program requires beneficiaries to take Doroob courses as a way to show their commitment to upskilling. Because these users are prescribed a set course curriculum to receive unemployment benefits, we exclude these users in Table 1 and throughout our later analysis.

Doroob users are directed to the platform through numerous venues: either online or though colleagues, friends, or family. Doroob users sign up for the platform primarily to gain skills for the purpose of employment. In our survey data, over 60 percent of new users reported not being employed and seeking employment. The next most popular category, comprising 15 percent of users, was users who are “currently employed and plan to change jobs in the next year." Consistent with roughly 75 percent of Doroob users looking for work, when asked why they registered for Doroob, 82 percent of users responded that they were “interested in learning a new skill" and 53 percent responded that they “want to get a job."

In addition, we solicited self-reported rates of job search behavior. 32 percent of users report using a job search platform or engaging in job training over the past 6 months. 27 percent report sending resumes or job applications over the same period, and 19 percent report contacting an employer/agency or interviewing for a job.

Doroob users can voluntarily enroll in hundreds of courses across a range of subjects. Courses are free to take, and most courses take roughly three hours to complete. Upon registering for the platform, approximately half of Doroob users enroll in exactly one course in the next 30 days. Roughly 20 percent of Doroob users enroll in two courses over this period.[17] Throughout this paper, we separate courses into four large and distinct categories which individually account for a large share of Doorob enrollments: general skills courses, occupation-specific courses, computer courses, and telework courses. Table 2 provides a list of the most popular courses within each of these four categories during our during our 23-week analysis period. Survey data indicates that Doroob users primarily select courses based on personal preferences and anticipated labor market effects, i.e., whether they think the course will increase their employability or salary.

2.3 Conceptual Framework

We provide theoretical motivation for COVID-19's effect on skill acquisition in our context. To build intuition on how individuals make skilling decisions and how these may be affected by the pandemic, we offer a conceptual framework in which individuals invest in different skills over time. Details of the model can be found in Appendix A; this section sketches the model and notes main predictions.

We study the skill investment decisions of a representative agent in the context of a discrete time model similar to Sanders and Taber (2012). In this model, agents allocate time in each period between (a) investment in skills and (b) working for a wage. This setup is similar to the canonical life-cycle human capital model in Ben-Porath (1967), with one important difference: skills are multi-dimensional.

In each period, agents earn a wage based on their stock of skills and the prices these

[17]These figures are proportions of users who register for Doroob and enroll in at least one course. Our data excludes users who registered for the platform but did not enroll in any courses. The distribution of enrollments for different analysis populations is shown in Appendix Figure B1.

skills earn in the labor market. Skill prices can change over time–for example, teleworking skills may be more valuable in the short run and less valuable in the long-run.

In equilibrium, individuals select the optimal level of skilling to equalize the present return on time spent working to the discounted future return on time spent investing in skills. We list the four main predictions of the model, and how they relate to the COVID-19 shock studied here, below.

1. *Skill investment is higher when future skill prices are higher.* Intuitively, forward-looking, earnings-maximizing individuals will concentrate their skill investment in the skills that are most well-paid. In the context of COVID-19, long-term shifts in the value of different skills are uncertain, but many commentators argue that COVID-19 accelerated pre-existing trends towards remote work, e-commerce, automation, and other high-skill work.[18] Doroob was created as part of a broader effort to train Saudi nationals for jobs in many of these areas. Conceptually, we predict that individuals will increase investment in many of these Doroob courses in response to the COVID shock.

2. *Investment is higher when wages in the current period are lower.* Individuals face a tradeoff between earning wages in the current period and investing in skills for the future; higher current-period wages entails greater opportunity cost of skilling, lowering investment. In the context of the COVID-19 shock, many individuals saw their effective wage become zero, as their places of work were temporarily or permanently closed. More broadly, stay-at-home orders limited the set of non-work activities available to Saudis, and thus lowered relative value of free time spent not on Doroob. Our model predicts that skilling will increase in response to wage declines.

3. *In all skills, investment is weakly lower for older agents.* As in similar models, indi-

[18] McKinsey and Company, February 18, 2021. "The future of work after COVID-19."

viduals tend to concentrate their skill investments in early periods; younger agents devote more time to skill acquisition than older agents. This difference reflects the longer time horizon that younger workers have. Skills that can be deployed over a 50-year career have higher present value than skills deployed over a 10-year career. Given that young workers have a longer working horizon in the post-COVID economy, we predict that their skilling decisions will increase most in response to the COVID shock.

In our empirical work, we also consider differential responses between individuals who are currently working, currently jobseekers, or currently students (based on information provided by the user during registration). While these distinctions are outside the scope of our formal model, a similar prediction is that students may exhibit larger enrollment responses than workers or jobseekers, for reasons very specific to the pandemic: school closures likely lowered students' investment in skills via traditional educational institutions (secondary schools and universities), so students may have substituted lost schooling with skilling through other channels, like Doroob. In this sense, some of the increase in skilling may reflect substitution of skilling across domains (e.g. from in-person classes at a university to online courses on Doroob).

4. *Older individuals invest relatively more in skills that are valuable in the short-run.* The presence of multi-dimensional skills introduces interactions between age and skill content. Intuitively, a greater share of older agents' remaining working life is concentrated in periods where short-term skills will be valuable; concentrating investment in these skills is relatively more important in their wage path. One short- to medium-term effect of the COVID shock was it's effect on remote work. In the weeks and months following the COVID shock, teleworking skills became much more valuable as many jobs transitioned to remote work. While these short- to medium-term effects are well-established, the long-term value of teleworking skills

in the labor market is unclear. In this light, our conceptual framework predicts that older users will invest most substantially in telework skills, relative to younger users.

Alternatively, heterogeneity by age may arise in response to different initial levels of skill development. If older individuals have different initial (pre-COVID) skill levels, they may exhibit different responses to the shock. Namely, if older individuals may have relatively lower levels of teleworking skills prior to the COVID-19 shock, they may exhibit higher increases in telework investment in response.

3 Data and Methodology

3.1 Data Samples

To estimate how the COVID-19 shock affected skilling decisions on Doroob, we analyze user-level course enrollments data over the period between January 2018 and June 2020; thus, we are limited to studying the short-term impact of the pandemic. Throughout, we exclude users who were younger than 18 or older than 65 when they signed up for Doorob. We additionally exclude users who were directed to Doorob via the Saudi unemployment insurance program, Hafiz; these users are prescribed a set course curriculum to receive unemployment benefits, so their behavior on the platform largely reflects these course requirements, rather than independent course selection.

The first column of Table 3 shows summary statistics for active Doorob users in 2020 (after making the sample restrictions described above). Over 200,000 unique users enrolled in courses on Doroob over this period. This population is relatively young, with an average age of 27.4, and over half of users under 25 years old. Roughly half of these active Doroob users are women. 52 percent of users report having a bachelor's degree.

Our data contains enrollments over time for different users, enabling us to make numerous comparisons across users and over time. We focus on three 23-week periods in early 2018, 2019, and 2020, and estimate the effects of the COVID shock by comparing changes in 2020 enrollments to those in prior years. We refer to users whose 2020 enrollment behavior is in our sample as the “2020 Cohort" and the prior two cohorts as the “2019 Cohort" and “2018 Cohort."

Additionally, we distinguish between two margins through which skilling decisions may respond to the COVID shock. First, prior Doroob users may engage with the platform differently in response to COVID. For example, users may return to the platform to take additional courses in response to the COVID shock or may take different types of courses than they would have, absent the COVID shock. Broadly, these analyses relate to existing users. The second and third columns of Table 3 summarize characteristics of the 2020 Cohort and the 2018/2019 Cohorts of existing users, respectively. We describe the selection criteria we use to identify these cohorts below.

Second, the COVID shock may have induced individuals who had not used the platform before to register for Doroob. These analyses relate to new users. The fourth and fifth columns of of Table 3 summarize characteristics of the 2020 Cohort and the 2018/2019 Cohorts of new users, respectively. We describe the selection criteria we use to identify these cohorts below.

In the two sections immediately below, we describe the data and methodologies used in our analyses in more detail.

3.2 Existing Users

Our existing users analyses aim to identify how existing Doroob users changed their skilling decisions in response to the COVID-19 shock. To do so, we define three cohorts of users. The first—which we refer to as a 2020 Cohort—is comprised of users whose enrollment decisions in 2020 may have been affected by COVID-19. The other two—2018 and 2019 Cohorts—are users whose enrollment decisions in 2018 and in 2019 were not affected by COVID-19. Panel A in Figure 3 illustrates this strategy.

Our 2020 Cohort consists of users who had joined Doroob and were active on the platform prior to the spread of COVID-19. Specifically, we consider all users who joined Doroob and took at least one course on the platform between July 1, 2019 and December 31, 2019: before concerns about COVID-19 had become widely-known.[19] We analyze this sample's course-taking behavior over a 23-week period from January 5, 2020 to June 13, 2020. This period includes 10 weeks prior to the Saudi stay-at-home order on March 16, 2020, and the 13 weeks thereafter.

In addition, to implement our difference-in-differences strategy, we further define two cohorts using the same selection criteria and analysis period, but shifted one or two years earlier. Specifically, 2018 and 2019 Cohort users are those who joined Doroob and took at least one course on the platform between July 1, 2017 and December 31, 2017 or July 1, 2018 and December 31, 2018, respectively. We analyze these users' behavior analogously over the first 23 weeks in the next calendar year (i.e., from January to mid-June). As in the treatment cohort, this period includes 10 weeks prior to the week of March 15 and 13 weeks thereafter.[20]

Using 2018 and 2019 Cohorts as "controls" allows us to mitigate the effects of two potential sources of bias if we were to estimate effects using only the single-difference event study for our 2020 Cohort. First, with only one cohort, we would be unable to distinguish between variation over the calendar year (i.e., seasonality or calendar week effects) and the effects of the COVID-19 shock. The inclusion of control cohorts allows us to control for calendar week fixed effects directly. Second, with one cohort, we could not distinguish between the effects of maturity on the platform and the effects of COVID-19. With a control cohort, the effects of platform maturity should be reasonably similar, as both co-

[19] December 31, 2019 marked the date on which the WHO County Office in China was notified of a cluster of viral pneumonia cases in Wuhan (Carvalho et al., 2021).

[20] For our treatment cohort, the analysis period runs from January 5, 2020 to June 13, 2020. For the control cohorts, this period is January 7, 2018 to June 16, 2018 and January 6, 2019 to June 15, 2019.

horts joined the platform in the six month period prior to their analysis period. However, the selection of individuals across cohorts (in terms of observables and unobservables) may change. While this is not directly testable, a parallel trends assumption between the treatment and control cohorts suggests that this is not a serious concern.

The second and third columns of Table 3 provide summary statistics on the users in our panel. Our 2020 Cohort of existing Doroob users consists of roughly 27,000 users. Together, 2018 and 2019 Cohorts consists of nearly 125,000 users. The differences in the relative size of these cohorts is driven by the relatively large user growth during the second half of 2017 and 2018, relative to the user growth during the second half of 2019. These trends are reflected in Panel B of Figure 1; the slope of the cumulative registrations over time reflects the rate of user growth. This slope is steeper during the second half of 2017 and 2018, than in the second half of 2019. While there is differential growth across the years, our identification strategy does not rely on similar growth across years. The larger concern is differential selection across cohorts, which can be mitigated through evidence of parallel trends across cohorts.

Among both groups, approximately half are women. While the treatment and control samples differ on some observable characteristics, our identification strategy does not require balance on observables. Instead, we rely on a parallel trends assumption that, absent the COVID-19 shock, the pre-COVID differences in user behavior between these two groups would have remained constant. Our event study specifications provide evidence of parallel trends in the weeks preceding the COVID-19 shock. In addition, in Appendix C we provide results based on a coarsened exact matching (Iacus et al. (2012)) algorithm that matches users in the treatment group to users with similar characteristics in the control group. As described later, these results are qualitatively similar to our main estimates.

To analyze the behavior of these users, we construct a full user-by-week panel. For all users and all 23 weeks during the analysis period, we count the number of weekly enrollments per user, both overall and in subject-specific courses (e.g. telework courses or occupation-specific courses).

With these data, we estimate the effect of COVID-19 on user behavior via a two-period difference-in-differences specification that includes fixed effects for cohorts t and calendar weeks w.

$$y_{iwt} = \beta\big[PostMarch15_w \times \mathbb{1}\{t = 2020\}\big] + \mu_t + \lambda_w + \varepsilon_{iwt}. \tag{1}$$

$PostMarch15_t$ is equal to one for weeks March 15 and later. The binary variable represented by $\mathbb{1}\{t = 2020\}$ is equal to one for 2020 (i.e. for the 2020 Cohort), and zero otherwise. Our coefficient of interest, β, reflects the difference in post-March 15 enrollments per user y_{iwt} in 2020 compared to the identically-defined set of control users in the same week during prior years; recall that March 15 is the date when Saudi Arabia issued a stay-at-home order and closed government offices. We also estimate Equation 1 with treatment interacted with binary variables for different demographic characteristics, such as age groups, to estimate heterogeneity by age.

To provide a week-by-week analysis of the impact of COVID-19 we estimate a similar equation that estimates over-time effects flexibly via calendar week fixed effects:

$$y_{iwt} = \sum_{\substack{k=-10 \\ k \neq -1}}^{12} \beta_t \big[\mathbb{1}\{k = w\} \times \mathbb{1}\{t = 2020\}\big] + \mu_t + \lambda_w + \varepsilon_{iwt}. \tag{2}$$

Equations 1 and 2 are identical, with the exception of how they estimate the dynamic effects of the COVID-19 shock. Equation 1 summarizes the estimated effect of the COVID-19 shock with an individual coefficient, β, whereas Equation 2 estimates dynamic, week-specific effects of the COVID-19 shock. Here, our event study coefficients, β_t, can be interpreted as differences in user behavior y_{iwt} in 2020 compared to the identically-defined set of control users in the same week during prior years. Dropping the week in which

$k = -1$ implies that differences between control and treated cohorts are normalized such that the difference in the week prior to March 15 is equal to zero.

Recent advances in the difference-in-differences literature have raised concerns about bias in two-way fixed effects estimates (Goodman-Bacon, 2021). In our context, our estimate $\hat{\beta}$ in Equation 1 is an unbiased estimate of the average treatment effect on the treated because treatment timing is not staggered (Baker et al., 2022).

Table 3 shows that, on some dimensions, the 2020 Cohort and the 2018 and 2019 Cohorts exhibit differences in baseline characteristics. In Appendix C, we show that our results with respect to existing users are largely similar if we balance user characteristics across these two cohorts. To do so, we implement coarsened exact matching (Iacus et al., 2012) to match users in the treatment group to users with similar characteristics in the control group. This matching algorithm produces weights for each user that balance baseline characteristics across the two groups. Appendix C compares the unweighted event study results to results weighted by the coarsened exact matching algorithm.

3.3 New Users

In our analyses of new users, we consider the effects of the COVID-19 shock on the choices new users make during their first week after registering for the platform. For a specific week, we define a new user as any user who has joined Doroob in that week and taken at least one course during their first 7 days after joining. For example, new users for the week of March 15, 2020 are users who joined between March 15, 2020 and March 21, 2020 (inclusive). Our analyses compare enrollment patterns across cohorts of new users over time.

By construction, our sample only includes users who took at least one course after registering (we do not have information on users who enroll but do not take any course subsequently). In these analyses, we compare the choices of cohorts of new users who joined Doroob before the COVID shock to the choices of cohorts of new users who joined after the COVID shock in 2020 to the choices of users who joined Doorob in the same weeks in 2018 and 2019. Panel B of Figure 3 illustrates this strategy.

To estimate the effect of the COVID shock on choices new users make once they are on the platform, we analyze the enrollment behavior of users in their first week on the platform. As in our prior analyses, we compare these choices to the choices of new users from the year prior. For user i who joined Doroob in calendar week w in year t, denote the user's course enrollments in their first week with y^1_{iwt}. We estimate the equation:

$$y^1_{iwt} = \beta \left[PostMarch15_w \times \mathbb{1}\{t = 2020\} \right] + \mu_t + \lambda_w + \varepsilon_{iwt}. \tag{3}$$

As before, β measures the difference between changes in user behavior that occurred coincident with the COVID-19 shock, relative to the changes in user behavior over the same period in prior years. $PostMarch15_t$ is equal to one for users who joined Doroob after March 15th, and $\mathbb{1}\{t = 2020\}$ identifies the 2020 Cohort (versus the 2019 and 2018 Cohorts).

However, Equation 3 differs from our analysis of existing users in that it considers each user's enrollments in their first week; β therefore measures changes across cohorts of new users, as opposed to changes within the same cohort.

Finally, we estimate an event study specification of the form below.

$$y^1_{iwt} = \sum_{\substack{k=-10 \\ k \neq -1}}^{12} \beta_t \left[\mathbb{1}\{k = w\} \times \mathbb{1}\{t = 2020\} \right] + \mu_t + \lambda_w + \varepsilon_{iwt}. \tag{4}$$

Here, β_t measures dynamic effects of the COVID-19 shock, reflecting the week-specific difference between enrollments among the 2020 Cohort relative to the 2019 and 2018 Cohorts.

In addition to the specifications in Equations 3 and 4, we run versions of these regressions that measure differences in user persistence between COVID shock-induced cohorts and others. To do so, we replace y^1_{iwt} on the left-hand side of these equations with y^2_{iwt} or y^3_{iwt}, which capture user i's total enrollments in their 2nd and 3rd week since they joined Doroob, respectively; we keep all other parameters the same. If COVID-induced cohorts revert back to typical behavior of new users in their second or third week, we would expect these effects to be zero. Alternatively, if COVID-induced cohorts are more active on the platform in the weeks after they join, these effects will be positive.[21]

The fourth and fifth panels of Table 3 provide summary statistics on the users in new users analysis. Our treatment cohort is much larger—over 125,000 users—than the combined size of both control cohorts. This difference reflects partly the growth of the platform over time and partly the effect that the COVID-19 shock had on enrollments, as shown visually in Panel B Figure 1. Again, our analysis will depend on a parallel trends assumption, that is pre-COVID, differences in new users between the treatment cohort and control cohorts did not exhibit any trend.

4 Results

We present three sets of results. First, we present broad evidence that levels of activity on Doroob increased dramatically across all margins we analyze and nearly all course types. Second, we provide evidence on the relationship between course content and user demographic characteristics, specifically age and gender. Finally, we present suggestive evidence on the relationship between COVID-induced enrollments and labor market out-

[21]These estimates likely understate the effect of the COVID shock on user persistence, because the effects of the COVID shock may affect some users who joined Doroob immediately prior to the COVID shock. For example, consider users who joined the week prior to the onset of the COVID shock. While these users' first-week enrollments precede the March 15th stay-at-home order, their second-week enrollments fall on the week of March 15. Thus, their enrollment patterns likely reflect, to some degree, the effect of the COVID shock. However, because these users were not part of the COVID-induced cohort, their second and third week enrollments serve as controls in our analysis of COVID-induced changes in persistence.

comes.

4.1 Levels of Activity on Doroob

We first present results with respect to the volume of activity on Doroob, which exhibited a massive increase among both existing and new users. The aggregate patterns in Panel B of Figure 1 demonstrate that enrollments and registrations increased in response to the COVID-19 shock. Our analyses of existing and new users estimate how the COVID-19 shock affected average weekly enrollments per user, and how these changes varied across course types.

Our main results with respect to existing users are shown graphically in Figure 4. Panel A of Figure 4 displays the raw means of weekly enrollments per user, separately for all three cohorts of existing users: 2018, 2019, and 2020 Cohorts. Prior to mid-March, these series move mostly in tandem and fell between 0.02 and 0.06 enrollments per user per week.[22] All series have a slight downward slope, reflecting users engaging less with the platform over time; recall that all three cohorts joined Doroob in the second half of the year prior, months before the analysis period reflected in Figure 4. This trend breaks in the week of March 15, when the 2020 Cohort dramatically increased their enrollment activity, from roughly 0.02 enrollments per week to 0.06 enrollments per week, on average. Over the same period in 2018 and 2019, enrollment patterns were reasonably flat.

[22] The raw means in Figure 5 shows that enrollments generally exhibit a downward trend over time. This trend reflects user attrition over time. This attrition was particularly steep for 2019 Cohort, which generates some significant coefficients on pre-COVID week indicators. In Appendix B, we show that these significant pre-period event study coefficients are driven primarily by users who take exactly one course during their first 30 days on Doorob. Separating our analyses between users who took exactly one course (who we refer to as “low attachment users") and users who took at least 2 courses during this period (“high attachment users") produces event study estimates that exhibit less extreme differences in pre-COVID enrollments, with post-period estimates of roughly similar magnitude.

Panel B of Figure 4 displays corresponding event study estimates, netting out changes in the 2018 and 2019 Cohorts. Estimates in Figure 4 correspond to β_t in Equation 2 and show how the COVID-19 impact varies over time. Consistent with the raw averages in Panel A, these estimates suggest that the COVID-19 shock increased average enrollments per user by roughly 0.04 per week. A typical course on Doroob is 180 minutes long, so a 0.04 course increase is equivalent to roughly 7 minutes per user per week.

Panel A of Table 4 presents average treatment effect estimates: the corresponding estimates of the impacts on new users from Equation 1. First, consistent with Figure 2, Column 1 in Table 4 suggests that average enrollments on Doroob increased by approximately 0.04 enrollments per week. Relative to pre-COVID weekly enrollments among the 2020 Cohort, this effect is approximately 0.1 standard deviations. This estimate is based on user-by-week data; multiplying estimates by 13 (the number of post-COVID-19 shock weeks in the panel) suggests that the COVID-19 shock increased total enrollments among this group by roughly 0.5, equivalent to 90 minutes in total. Columns 2 to 5 estimate effects across different course types. These results suggest a broad increase in enrollments across all course types.

The estimates described above incorporate only changes among existing Doroob users, and do not fully capture the magnitude of changes in skill acquisition induced by the COVID-19 shock. We next consider how the COVID-19 shock affected the enrollments of new users on Doroob. Our first set of results for new users, shown visually in Figure 5, document how enrollments on Doroob changed in response to the COVID-19 shock. Figure 5 displays estimates of β_t in Equation 4. The weekly series in Figure 5 are noisy, reflecting the variation in the composition of new users from week to week. However, the general trend suggests that enrollments among post-COVID shock cohorts were higher, indicated by the generally positive weekly estimates in Figure 5. These estimates suggest that the COVID shock increased the number of first-week enrollments among new Doroob users by roughly 0.5 enrollments.

Panel B of Table 4 displays corresponding estimates of Equation 3 across different course types. First, consistent with Figure 5, Column 1 in Table 4 suggests that average first-week enrollments on Doroob increased by approximately 0.44 following the COVID-19 shock. This increase is consistently positive across course types, with one exception: computer courses, which appear to fall in popularity among new users. An additional 0.44 courses corresponds approximately to an estimated additional 80 minutes spent on the platform during each user's first week.

Increases across course types can be difficult to compare, given the differences in the relative popularity of each course type. For example, in our existing users analysis, roughly half of the the 2020 Cohort's pre-COVID enrollments were in general skills courses, but only 12 percent were in computer-related courses. To enable consistent comparisons across course types, we repeat these analyses after dividing each cohort's weekly enrollments by their average pre-COVID enrollments (separately for each course type). This normalization enables us to interpret difference-in-difference estimates from Equation 1 as relative increases; a coefficient of 1 indicates a 100 percent increase.

These results are shown in Table 5. Panel A in Table 5 shows results with respect to existing users. Consistent with the results in described above, the COVID-19 shock induced existing users to increase their enrollments by nearly 150 percent. This increase was roughly equivalent (in relative terms) for courses related to general or occupation specific skills, slightly lower for computer skills, and substantially larger for telework courses. Among existing users, telework courses increased over 1300 percent.

As before, we present normalized outcomes in Table 5 shows normalized results with respect to existing users, enabling comparisons across course types. The COVID-19 shock increased first-week enrollments among new users by roughly 20 percent, an increase that appears reasonably similar across all courses overall, general skills courses, and occupation-specific courses. Computer courses exhibit a 9 percent decrease in popularity. As before, telework courses are a massive outlier, exhibiting an increase of over 1200 percent.

Finally, we consider whether COVID-induced users were more persistent on the platform, relative to users that joined in prior weeks. To do so, we estimate Equation 3, replacing first-week enrollments with enrollments in a users' second or third week. The results are shown in Table 6, where Panel A displays results for first-week enrollments (identical to those in Table 4) and Panels B and C reflect second- and third-week enrollments, respectively. We highlight three points from Table 6. First, the increased activity of new, COVID-induced users on Doroob extends into their second and third week on the platform, as indicated by consistently positive coefficients in Panels B and C. Second, while the magnitude of these reductions falls over time—from 0.44 enrollments in week 1 to 0.27 to 0.15 in weeks 2 and 3—the *relative* magnitude increases over time. Specifically, the effect size in week 1 is roughly 20 percent of the pre-COVID mean, whereas the effect size in weeks 2 and 3 is roughly 100 percent of the pre-COVID mean. Finally, the week-one reduction in enrollments in computer-related courses appears to be short-lived; consistent with all other courses, computer enrollments in weeks two and three exhibit large, positive increases.

Broadly, this evidence on the levels of investment is consistent with the model described in Section 2. For one, we see the largest increases in skill investment in the skill whose demand rose the most substantially—telework. In the context of our model, this increase reflects the increase in the future skill prices of telework. Second, investment in skills rose broadly across all course types. In the context of our model, the most likely explanation is the lower opportunity cost of skill investment. The combination of stay-at-home orders and COVID-19 related job losses lowered the cost of online skill investment, leading to broad increases in course-taking on Doroob. Of course, some of this response could be substitution away from in-person skill investments, such as college courses. Later, we show that increases in course-taking were largely similar—and universally positive—among students, jobseekers, and the employed alike.

While our main analyses isolate over-time variation in exposure to the COVID shock, we additionally explore variation across regions in Appendix D. To do so, we compare enrollment responses among users in provinces that experienced larger COVID-19 outbreaks in early 2020 to enrollment responses among users in other provinces. While all increases in enrollments are found in provinces with high- and low-COVID incidence, areas that experienced relatively larger outbreaks generally exhibit larger responses among both new and existing users.

4.2 Heterogeneous Enrollment Responses to the COVID-19 Shock

Next, we consider how responses to the COVID-19 shock differed across demographic groups.

4.2.1 Age and Enrollments

We first consider how course content and age interact. Figure 6 provides suggestive evidence of our main effects, which consider the relationship between age and course content. Figure 6 displays enrollments by course type over time, separately for users under 30 and users age 30 or older. Vertical lines in Figure 6 correspond to March 15, at the onset of the COVID shock. Across all course types, users of all ages increase daily enrollments substantially. Among three panels: computer courses, general courses, and occupation-specific courses, increases in enrollments are equivalent or slightly larger among users under 30 years old. The opposite is the case for telework courses, where older users exhibit a much larger increase.

Our formal analyses disaggregate these comparisons between existing and new users. Among existing users, we estimate Equation 1 with fixed-effects and interactions for binary variables indicating different age groups. Throughout, all age interactions are interpreted relative to the youngest group: 18 to 24 year old users.

Panel A of Table 7 shows estimates of these regressions for existing users. Column

1 displays estimates for all courses overall. OLS estimates of enrollments exhibit little systematic differences between age groups and responses to the COVID-19 shock. General skills and computer courses exhibit some evidence of a negative age gradient; the response to COVID-19 was lower among 25 to 29 year olds, relative to 18 to 24 year olds for these courses. However, these patterns do not appear to extend to older users.

In contrast, telework courses, exhibit a sharp, positive age gradient. Relative to younger existing users, older ones were much more likely to enroll in telework courses. In response to the COVID shock, users between the age of 18 and 24 increased their weekly enrollments in telework courses by 0.001 enrollments per user per week. Users age 30 to 39 and age 40 to 65 exhibit increases that are four times larger.

Among new users, stronger age patterns emerge. Panel B of Table 7 displays estimates of Equation 3 with treatment interacted with binary variables for age groups. In their first week, younger users who joined the platform post-COVID shock were much more likely to enroll in general skills, occupation-specific, or computer courses during their first week in the platform. Meanwhile, older users were much more likely to enroll in telework courses. These increases were large and economically significant. Relative to users age 18 to 24, new users between age 40 to 65 enrolled in 0.14 more telework courses during their first week on the platform.

In the context of our conceptual model, the different age gradients of different courses reflect different skill prices over time. Telework courses were highly valuable in the immediate post-COVID period, but may be relatively less valuable in the future. Our model predicts that, while skill investment generally is concentrated among the young, older users' skill investment is concentrated among skills with the most immediate value. An alternative explanation is also possible: younger users may already have acquired the necessary telework skills prior to the pandemic.

4.2.2 Gender and Enrollments

Next, we consider how course content and gender interact. Our analyses of existing users estimate Equation 1 with treatment interacted with a binary variable equal to one for women. Panel A of Table 8 shows estimates of these regressions. When we consider all courses, general courses or occupation specific courses (Columns 1 to 3), estimates of differential effects among women are imprecise; effects are statistically significant at 10 percent for occupation-specific courses and otherwise insignificant. Computer and telework courses exhibit heterogeneous impact by gender. Compared to men, women were more likely to respond to the COVID shock by enrolling in computer courses (in Column 4). The opposite is true of telework courses, shown in Column 5; relative to men, women were less likely to enroll in these courses in response to the COVID shock.

The results reveal greater heterogeneity of impact among new users. Panel B of Table 8 displays estimates of Equation 3 with treatment interacted with a binary variable equal to one for women. Relative to men, women who joined Doroob in response to COVID were much more engaged with the platform, particularly with respect to occupation-specific and computer courses. Men, on the other hand, were much more likely to enroll in telework courses upon joining Doroob. The interpretation of these patterns is not fully clear: it could be that females were already well-placed to deal with the changes induced by COVID-19 because of prior social restrictions that limited their mobility.

In Appendix E, we investigate the interaction of age and gender in our context. We do so for two reasons. First, age and gender are correlated in our data; men on Doroob are typically older than women on Doorob.[23] Second, the age and gender composition of our new users sample changes over time, as shown in Appendix Figure E1. Both of these facts raise the possibility that measured age effects may instead reflect the effects of gender, or vice versa. Appendix E demonstrates that these effects are indeed distinct.

[23] Among our existing users sample, the median man was 26 years old, whereas the median women was 24 years old. Among our new users sample, these figures were 28 and 23.

4.2.3 Labor Force Status and Enrollments

Table 9 considers responses by labor force status. We group users into four groups, according to information self-reported during Doroob registration: students, employed workers, jobseekers, and missing (users who do not report a current status during registration). All interaction terms reflect differences relative to the missing group.

Among existing users in Panel A, Column 1 of Table 9 indicates that students and employed workers exhibited the largest increases in enrollments in response to the COVID shock. Compared to jobseekers, students and employed workers exhibit a response roughly 30 percent larger. This relationship holds across all course types, with one exception: telework courses. For those courses, employed workers exhibit a larger response to the COVID-19 shock than students and jobseekers.

Among new users, Panel B of Table 9 shows that the increase in post-registration enrollments among newly-registered users is driven primarily by students and employed workers; these users enrolled in between 1 and 1.5 more courses in their first week on the platform. Similar to above, we see that telework responses were extremely largest among employed workers, increasing their enrollments by over 0.15 enrollments their first week.

Consistent with our conceptual model, these results may reflect a relaxation of time constraints for employed workers and students. In the absence of in-person work and in-person education, the opportunity costs of these users' time fell; these effects were likely weaker among jobseekers.

An alternative explanation with respect to increased enrollments among employed workers is more direct—employers or coworkers may have recommended that employees take courses on Doroob following the stay-at-home order in March 2020. In Appendix F, we find some evidence of increasing employer concentration in registrations, suggesting that employers or coworkers may have been encouraging employees to sign up for Doroob in response to the COVID shock. In the months following the COVID shock, the share of new Doroob users who signed up during the same week and worked at the same firm rose from less than one percent to roughly three percent. This increase constitutes a small share of the overall increase in registrations, but suggests that employer or coworker recommendations may have accelerated the sharp growth in Doroob users.

4.3 Employment and COVID-Induced Enrollments

To assess the labor market impact of COVID-induced enrollments, we ask whether users who returned to take courses on Doroob following the COVID shock were more likely to be employed in the months that followed. To do so, we analyze administrative data collected by the Saudi social insurance agency General Organization for Social Insurance ("GOSI"). This GOSI data captures all private-sector jobs in Saudi as of September 19, 2020, 6 months after the onset of the COVID shock and 3 months after the end of our course enrollments analysis.[24]

We link GOSI data to user-level course-taking data from Doroob for all users in the existing users 2020 Cohort: users who joined Doorob in the second half of 2019. With this data, we assess whether users who enrolled in courses following the COVID-19 shock were more likely to be employed in GOSI data. We run linear probability models of the form:

$$Employed_i = \beta_0 + \beta_1 PostCOVIDCourses_i + \gamma \mathbf{X} + \varepsilon_i, \tag{5}$$

where $Employed_i$ is a binary variable equal to one for users who are employed in GOSI data, and zero otherwise. $PostCOVIDCourses_i$ reflects the number of courses taken by user i over the post-COVID period: March 15, 2020 to June 13, 2020. $\mathbf{X}$ is a matrix that includes demographic characteristics (fixed effects for age, gender, employment status at registration, and how the user was directed to the platform) as well as controls for pre-

[24] A separate database covers public sector employment, to which we do not have access. It is worth noting that Doorob was created in part to encourage private sector employment among Saudis. Thus, studying impact on private sector employment is informative.

COVID-19 shock courses during (fixed effects for number of courses taken in each of the four course categories).

Baseline results are shown in Table 10. For ease of interpretation, we multiply the dependent variable by 100, so coefficients can be interpreted as percentage points. Of the 27,000 users in the 2020 Cohort of existing users, over 23,000 users have national identifiers that allow us to merge them to GOSI data. Within that group, approximately 7 percent were employed in the private sector as of September 2020.

Column 1 of Table 10 indicates that a one-course increase in post-COVID Doroob enrollments is associated with a 0.24 percent increase (e.g. from 7.08 percent to 7.33 percent) in the likelihood of employment. Column 2 adds controls for demographic characteristics, which reduces the estimated effect size to 0.13 percentage points. Additionally adding controls for pre-COVID enrollments in Column 3 does not appear to change the estimated effect size.

In Columns 4 through 6, we repeat the analysis, separating enrollment counts across the four course categories used previously. Occupation-specific enrollments exhibit positive and statistically significant effects on employment. Estimated effects for enrollments in other course types are imprecise and statistically insignificant.

These results warrant two caveats. First, this analysis is suggestive in nature and does not isolate quasi-random variation in enrollments. Instead, observed enrollments may correlate with unobserved characteristics, biasing our estimates. To the degree that unobserved characteristics covary positively with both enrollment patterns and employment, our estimates are biased upwards. Second, and more practically, nearly half of the employed Saudis work in the public sector.[25] Thus, these estimates are based on only a subset of all Saudi jobs and do not reflect effects on employment levels overall.

[25] For example, official statistics in Q1 2019 report that 1.4 million Saudi nationals work in the public sector, versus 1.7 Saudi nationals who work in the private sector. Labour Market First Quarter 2019, General Authority for Statistics, Table 3.

5 Conclusion

The COVID-19 pandemic brought about a massive shift in the nature of work across the globe. While long-term effects of the economic and health shock are yet unknown, workers, firms, and other institutions continue to anticipate further changes to labor market: continued remote work and sectoral reallocation, for example.

To what degree are workers adjusting their skilling decisions in response to these changes in the labor market? The evidence in this paper suggests that workers are responding quickly to changes in the economy; adjustments in skilling decisions were swift—within a week of stay-at-home orders—and persistent—continuing for weeks and months afterwards. Moreover, these adjustments appear to be consistent with theoretical predictions with respect to short- and long-term exposure to labor market changes. Our simple human capital model provides a theoretical basis for this intuition: when skill prices vary over time, older users invest relatively more in skills that are more immediately valuable, such as teleworking skills. More broadly, we find little evidence that there are subsets of workers whose skilling decisions were not affected by the COVID shock. COVID appears to have altered the choices of young and old, as well as men and women—though not identically.

Future research should assess the nature and magnitude of *long-term* changes to the labor market, as well as the elasticity of individual educational choices to these changes. Relative to traditional education institutions, online platforms, such as Doroob, appear to enable workers to respond much more quickly to short- and long-term fluctuations in the labor market. The degree to which the presence of online platforms affects educational choices and downstream labor market outcomes constitutes a useful venue for further study. In addition, whether the shift in course-taking is actually welfare improving for the individuals is not clear. A better understanding of that is required for policy prescriptions.

References

Abramitzky, R., V. Lavy, and M. Segev. "The Effect of Changes in the Skill Premium on College Degree Attainment and the Choice of Major." National Bureau of Economic Research Working Paper 26420.

Adukia, A., S. Asher, and P. Novosad. "Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction." *American Economic Journal: Applied Economics* 12, (2020) 348–76.

Algaissi, A. A., N. K. Alharbi, M. Hassanain, and A. M. Hashem. "Preparedness and response to COVID-19 in Saudi Arabia: Building on MERS experience." *Journal of Infection and Public Health* 13, (2020) 834–838.

Aloud, M. E., S. Al-Rashood, I. Ganguli, and B. Zafar. "Information and Social Norms: Experimental Evidence on the Labor Market Aspirations of Saudi Women." National Bureau of Economic Research Working Paper 26693.

Atkin, D. "Endogenous Skill Acquisition and Export Manufacturing in Mexico." *American Economic Review* 106, (2016) 2046–85.

Baker, A. C., D. F. Larcker, and C. C. Wang. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics* 144, (2022) 370–395.

Becker, G. S. et al. *Human Capital*. New York: Columbia University Press (1964).

Bedard, K. and D. A. Herman. "Who goes to graduate/professional school? The importance of economic fluctuations, undergraduate field, and ability." *Economics of Education Review* 27, (2008) 197–210.

Beffy, M., D. Fougere, and A. Maurel. "Choosing the Field of Study in Postsecondary Education: Do Expected Earnings Matter?" *Review of Economics and Statistics* 94, (2012) 334–347.

Ben-Porath, Y. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy* 75, (1967) 352–365.

Blanchard, E. J. and W. W. Olney. "Globalization and human capital investment: Export composition drives educational attainment." *Journal of International Economics* 106, (2017) 165–183.

Blom, E., B. Cadena, and B. J. Keys. "Investment over the Business Cycle: Insights from College Major Choice." *Journal of Labor Economics* .

Brynjolfsson, E., J. J. Horton, A. Ozimek, D. Rock, G. Sharma, and H.-Y. TuYe. "COVID-19 and Remote Work: An Early Look at US Data." Technical report (2020). National Bureau of Economic Research Working Paper 27344.

Bursztyn, L., A. L. González, and D. Yanagizawa-Drott. "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia." *American Economic Review* 110, (2020) 2997–3029.

Carvalho, T., F. Krammer, and A. Iwasaki. "The first 12 months of COVID-19: a timeline of immunological insights." *Nature Reviews Immunology* 21, (2021) 245–256.

Cavounidis, C. and K. Lang. "Ben-Porath Meets Lazear: Microfoundations for Dynamic Skill Formation." *Journal of Political Economy* 128, (2020) 1405–1435.

Charles, K. K., E. Hurst, and M. J. Notowidigdo. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review* 108, (2018) 2947–94.

Cunha, F. and J. J. Heckman. "Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Journal of Human Resources* 43, (2008) 738–782.

Dingel, J. I. and B. Neiman. "How many jobs can be done at home?" *Journal of Public Economics* 189, (2020) 104235.

Goodman-Bacon, A. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225, (2021) 254–277.

Iacus, S. M., G. King, and G. Porro. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis* 20, (2012) 1–24.

Long, M. C., D. Goldhaber, and N. Huntington-Klein. "Do completed college majors respond to changes in wages?" *Economics of Education Review* 49, (2015) 1–14.

Mas, A. and A. Pallais. "Valuing Alternative Work Arrangements." *American Economic Review* 107, (2017) 3722–59.

Miller, C., J. Peck, and M. Seflek. "Integration Costs and Missing Women in Firms." National Bureau of Economic Research Working Paper 26271.

Patnaik, A., M. J. Wiswall, and B. Zafar. "College Majors." National Bureau of Economic Research Working Paper 27645.

Ritchie, H., E. Ortiz-Ospina, E. M. Diana Beltekian, J. Hasell, B. Macdonald, C. A. Charlie Giattino, L. Rodés-Guirao, and M. Roser. "Coronavirus Pandemic (COVID-19)." *Our World in Data* Https://ourworldindata.org/coronavirus.

Sanders, C. and C. Taber. "Life-Cycle Wage Growth and Heterogeneous Human Capital." *Annual Review of Economics* 4, (2012) 399–425.

Wiswall, M. and B. Zafar. "Preference for the Workplace, Investment in Human Capital, and Gender." *The Quarterly Journal of Economics* 133, (2018) 457–507.

Figure 1: Time Series Data: COVID-19 Case and Activity on Doroob

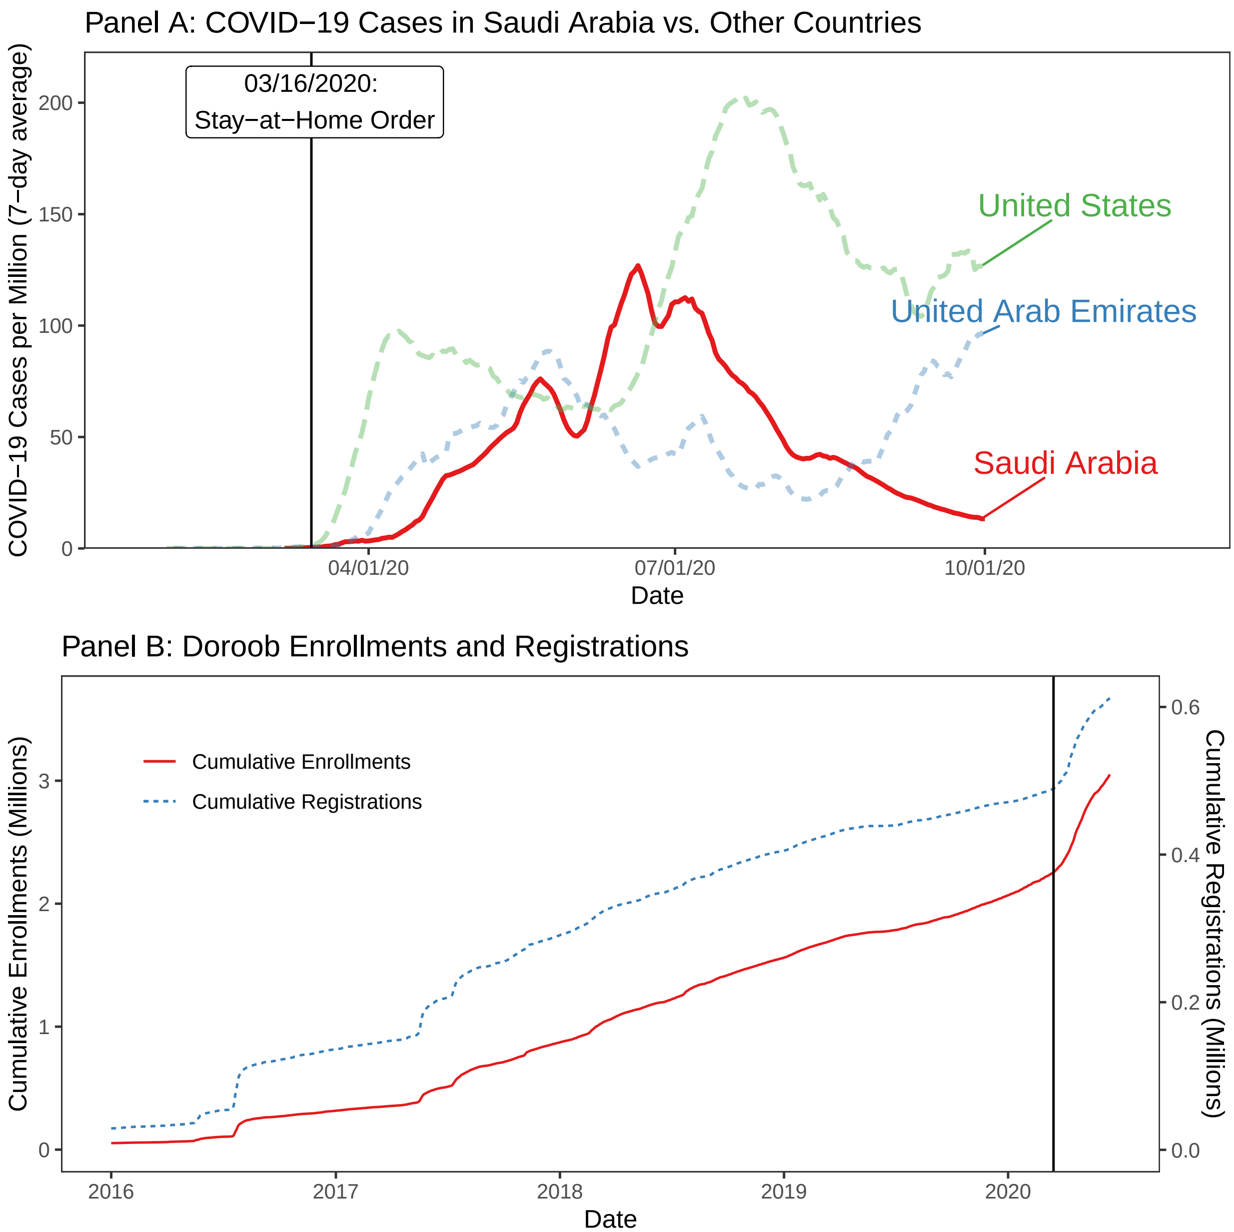

Notes: Panel A displays the 7-day moving average of new COVID-19 cases per million for Saudi Arabia, United Arab Emirates, and United States. COVID-19 cases data are from Ritchie et al. (2020). Panel B displays the cumulative users and enrollments on Doroob between January 1, 2016 and June 13, 2020. Totals exclude users who were younger than 18 or older than 65 when they signed up for Doorob and users who were directed to Doorob via the Saudi unemployment insurance program, Hafiz. In both panels, vertical line denotes March 16, 2020.

Figure 2: Unique Courses on Doroob over Time

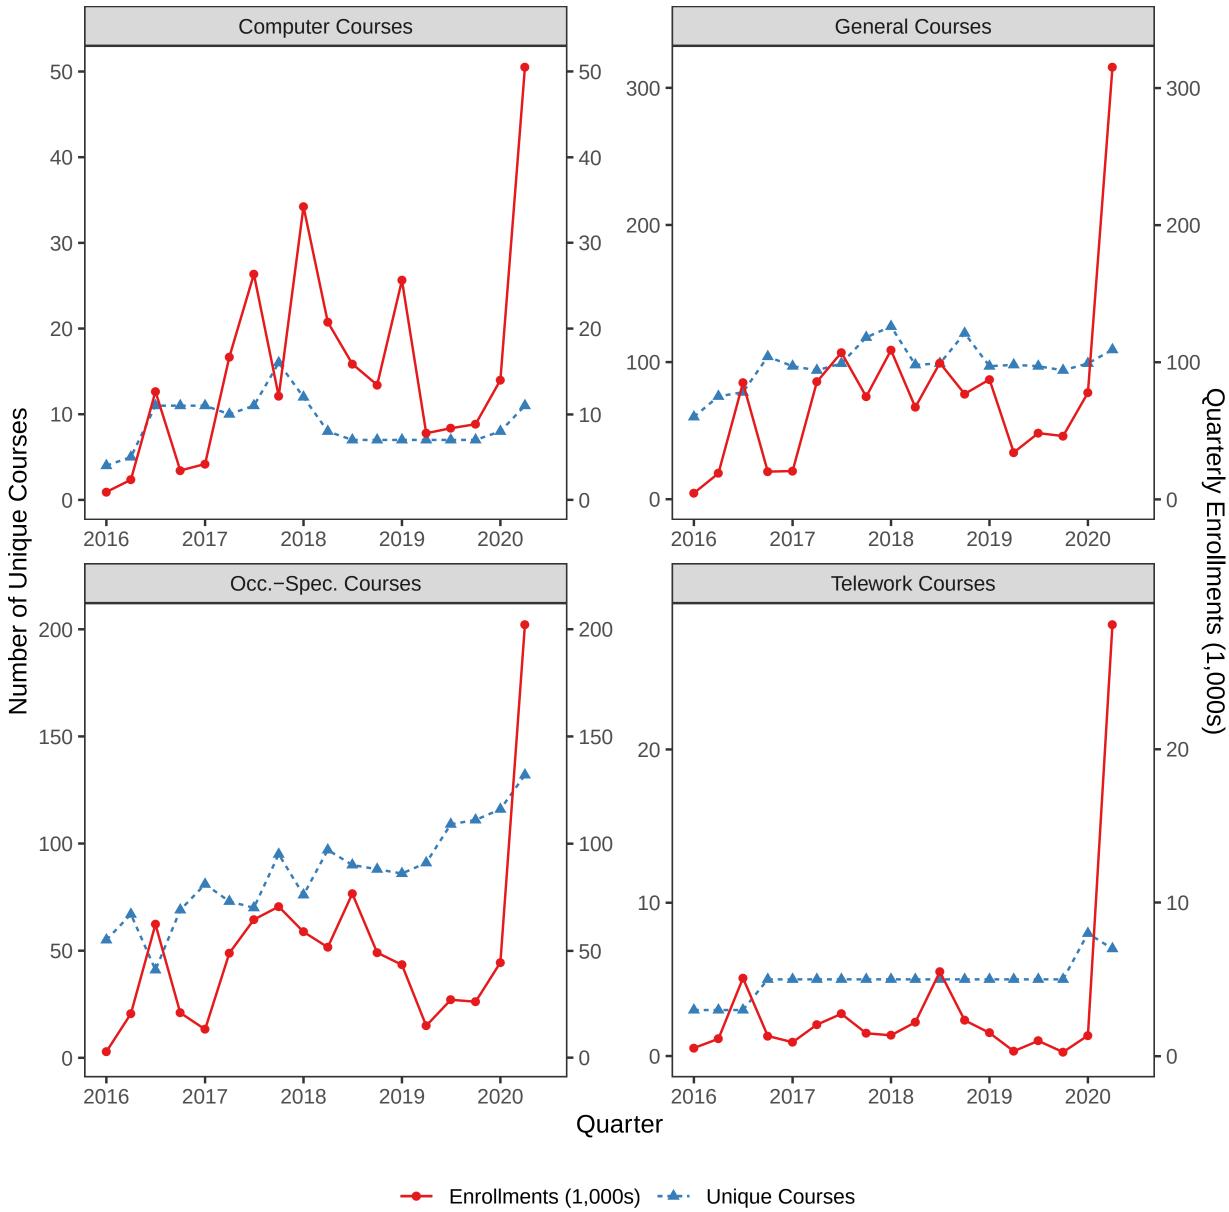

Notes: Figure displays the number of unique courses on Doroob each quarter since Q1 2016, separately for each course type.

Figure 3: Data Construction Illustration

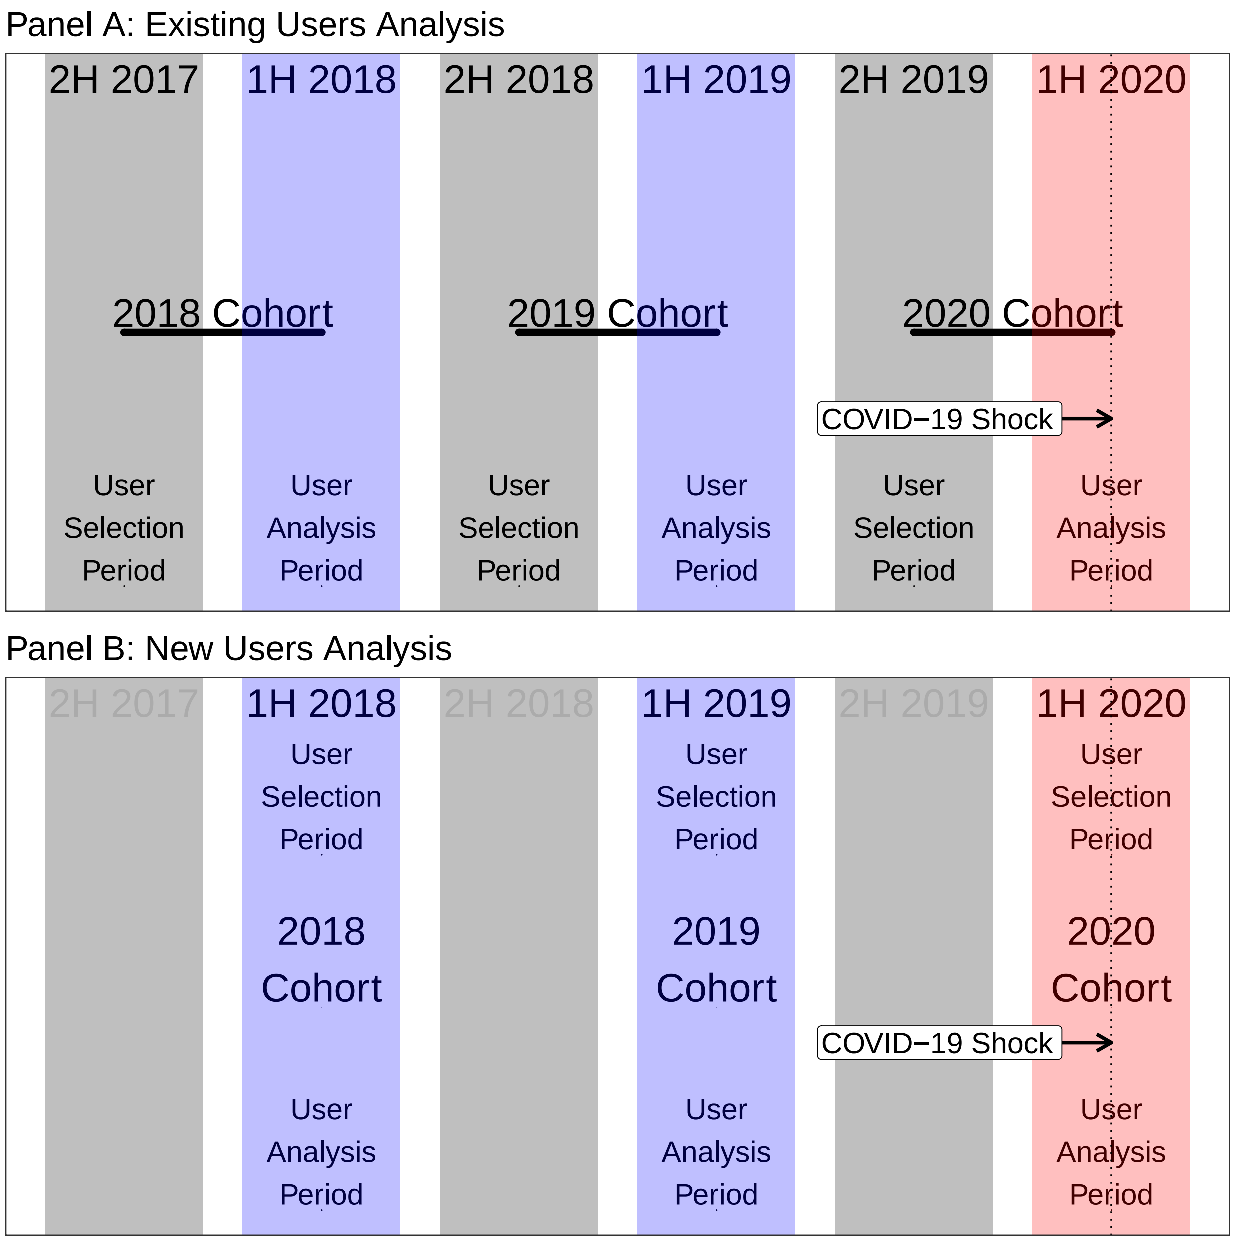

Notes: Figure illustrates the data used in our analysis of existing users and new users. To construct our existing users sample, we define two sets of users in our 2018 and 2019 Cohorts–users who joined Doroob between July 1, 2017 and December 31, 2017 or between July 1, 2018 and December 31, 2018. We define our 2020 Cohort as users who joined Doroob between July 1, 2020, and December 31, 2018. For all groups, we analyze subsequent course enrollment behavior in the subsequent months of January to June. To construct our sample of new users, we define two sets of users in our 2018 and 2019 cohorts, in the 23 weeks centered on March 15–users who joined Doroob between January 7, 2018 and June 16, 2018 or between January 6, 2019 and June 15, 2019. We define our 2020 Cohort as users who joined Doroob between January 5, 2020 and June 13, 2020.

Figure 4: Response to the COVID-19 Shock among Existing Users

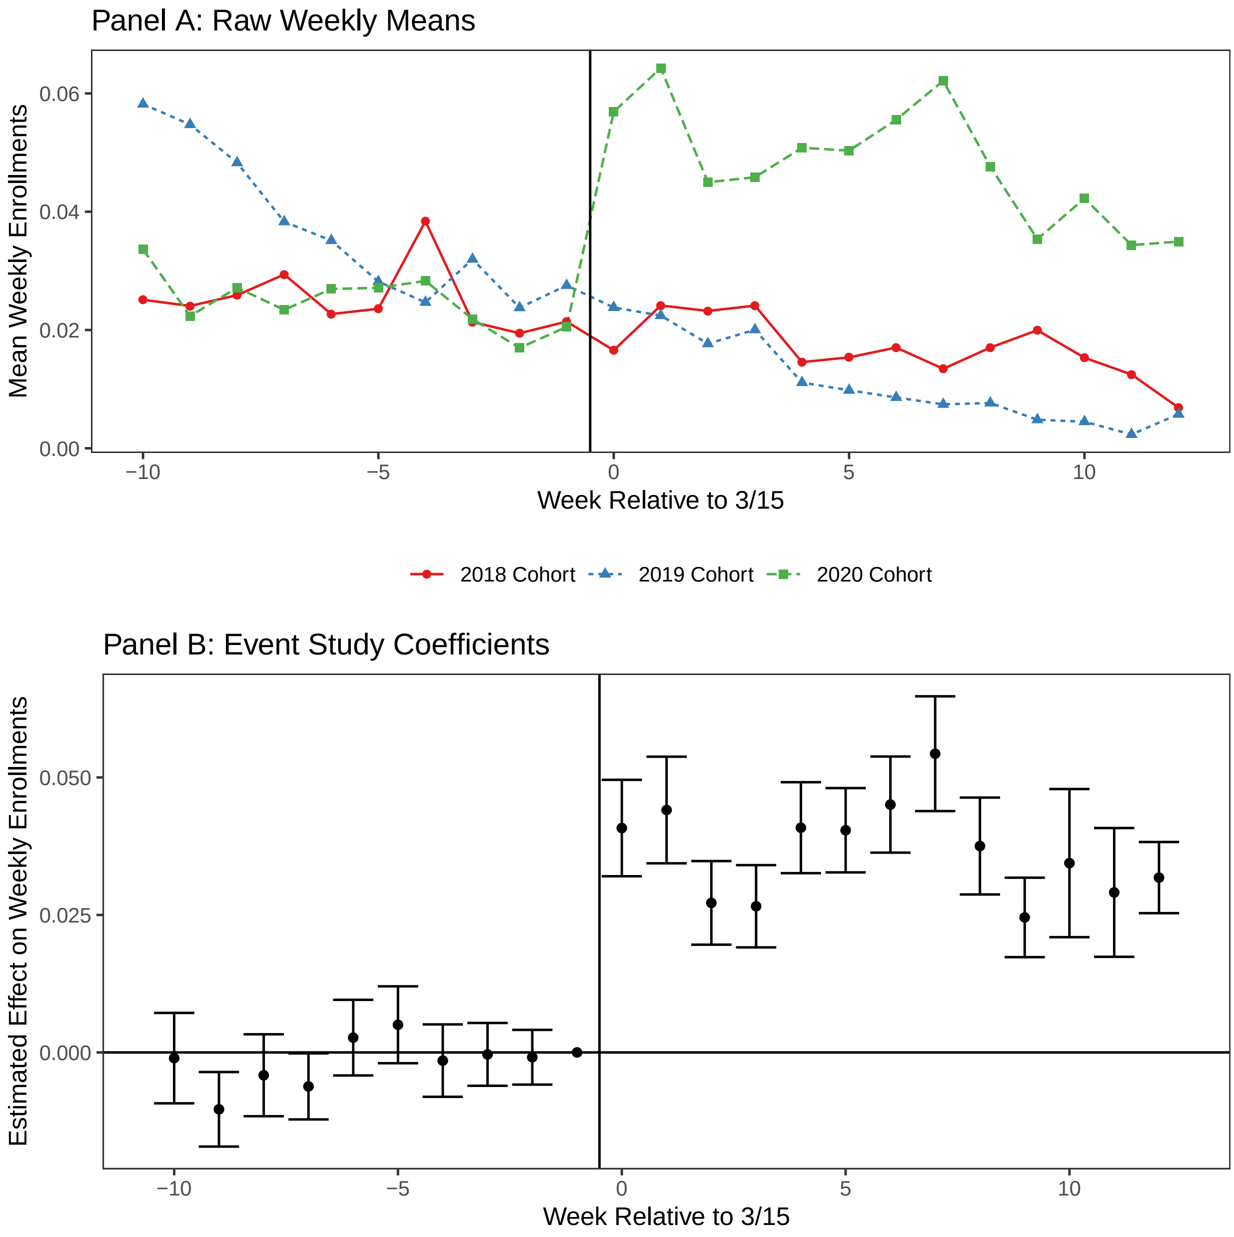

Notes: Figure displays how the COVID shock affected enrollments among existing users. The top panel displays the raw means of weekly enrollments per user, separately for 2018, 2019, and 2020 Cohorts. The bottom panel displays estimates of β_t in Equation 2. These estimates control for cohort and calendar week fixed-effects. Error bars represent 95% confidence intervals.

Figure 5: Response to the COVID-19 Shock among New Users

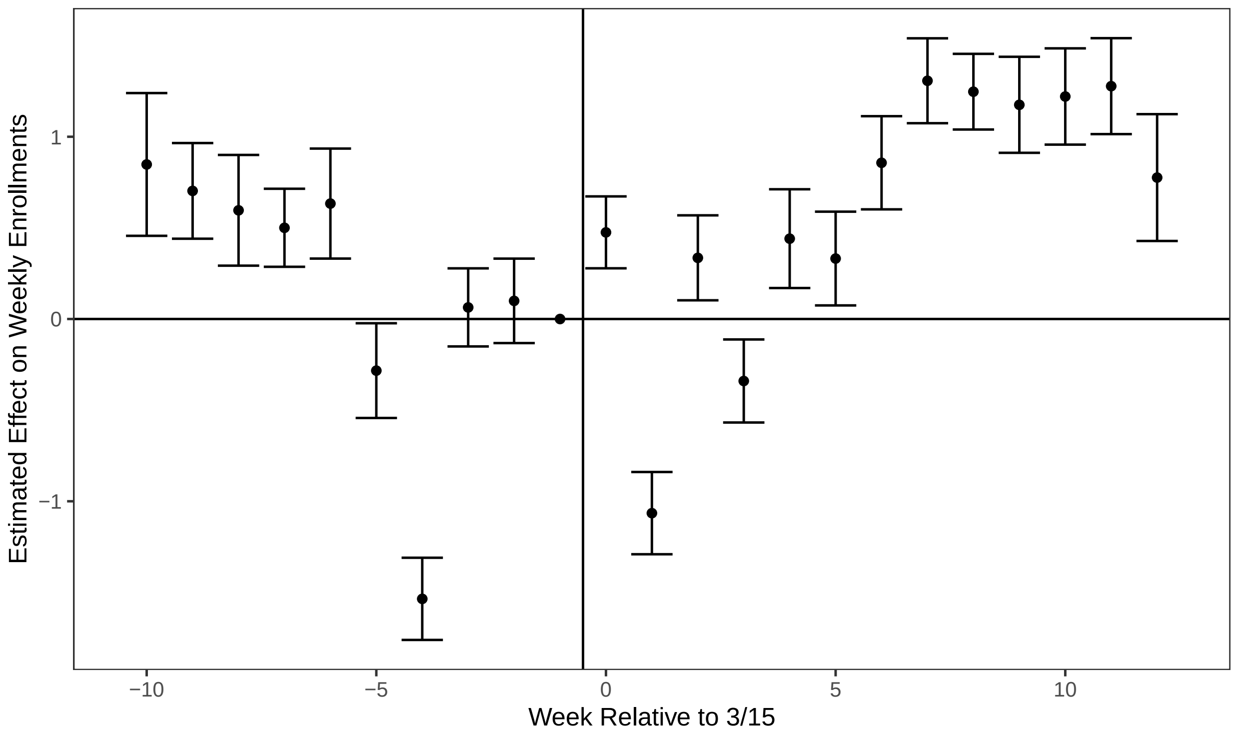

Notes: Figure displays how the COVID shock affected enrollments among existing users. Plotted coefficients are estimates of β_t in Equation 4. These estimates control for cohort and calendar week fixed-effects. Error bars represent 95% confidence intervals.

Figure 6: 2020 Doroob Enrollments by Age and Course Category

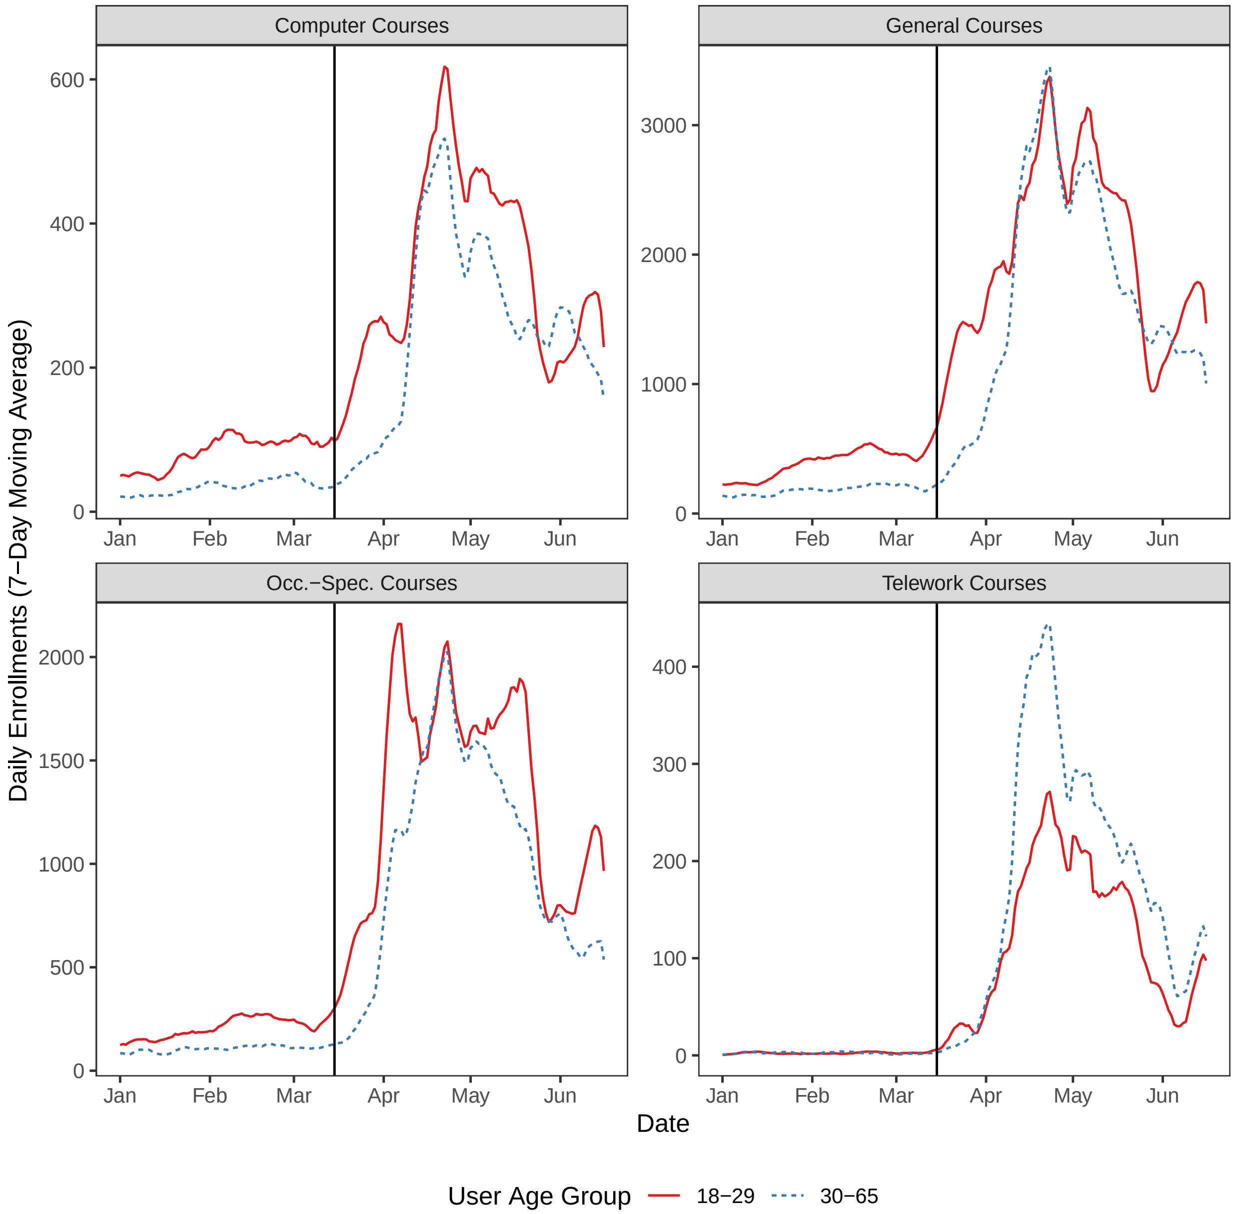

Notes: Figure displays the 7-day moving average of total course enrollments in 2020, separately for each course category and age groups: 18-29 years old and 30-65 years old. Age groups are determined by the age of each user upon registering for Doorob. Averages exclude users who were younger than 18 or older than 65 when they signed up for Doorob and users who were directed to Doorob via the Saudi unemployment insurance program, Hafiz.

Table 1: Survey Responses of Doroob Users in February 2020

Statistic	Mean	St. Dev.	N
Demographic Characteristics			
Female	0.48	0.50	962
Age	27.49	11.67	962
Employment Characteristics			
Employed; no plan to change jobs	0.11	0.31	962
Employed; plan to change jobs	0.15	0.36	962
Not employed; seeking employment	0.61	0.49	962
Not employed; not seeking employment	0.13	0.34	962
Job Search Behavior Over the Past 6 Months			
Used job training/job search platform	0.32	0.47	962
Sent out resumes/applications	0.27	0.44	962
Contacted employer/agency, or interviewed	0.19	0.39	962
How Did You Hear About Doroob?			
From my employer	0.03	0.18	962
From my university or school	0.15	0.36	962
From a colleague or friend	0.16	0.36	962
From online or other media	0.43	0.50	962
From my family	0.17	0.38	962
Other	0.06	0.23	962
Why Did You Register for Doroob?			
I am interested in new skill	0.82	0.38	962
I want to get a job	0.53	0.50	962
Boss directed me to register	0.06	0.23	962
It's a requirement (Hafiz program)	0.10	0.30	962
It's a requirement (other)	0.04	0.19	962
Other	0.06	0.23	962
When Choosing Courses, How Will You Decide Which to Take?			
Personal Preferences (1-10)	8.02	2.45	962
Whether course will get me a job (1-10)	8.11	2.58	962
Whether course will get me higher salary (1-10)	7.53	2.91	962
Prefs. of family/friends (1-10)	4.06	3.47	962

Notes: Table summarizes responses to a survey of new Doroob users conducted in February 2020. Sample excludes users directed to Doorob via the Saudi unemployment insurance program, Hafiz. We solicited survey responses via a pop-up on the Doroob website during the registration process and via direct emails to new Doroob registrants. Users could select multiple options when asked "Why did you register for Doroob," so the sum of mean responses exceeds one.

Table 2: Most Popular Courses by Course Category: 1/5/20 - 6/13/20

Course Name	Number of Enrollments
General Skills Courses	
Leadership Basics	29,569
Labor Culture according to the Saudi Labor System	18,375
Self-Management	16,684
Smart Work Ethic	15,902
Communicating in the Work Environment	15,592
Occupation-Specific Courses	
Introduction to Human Resource: Tasks	22,830
Basics of Management	14,556
Project Management	11,326
Principles of Financial Accounting Part I: General Framework of Accounting	8,394
The Basics of Quality and Safety Standards for Recreational Events	6,559
Computer Courses	
Information Security	25,345
Introduction to Microsoft Excel	11,086
Workplace IT: Master Microsoft Software	7,771
Introduction to Microsoft Word	7,160
Introduction to Microsoft PowerPoint	4,759
Telework Courses	
Culture of Teleworking Administrative Side	13,512
Culture of Teleworking Technical Aspect	8,656
Solve Problems While Working Remotely	2,000
Communication Skills in Remote Work	1,859
Professional Development for Remote Workers	1,627

Notes: Table displays the five most popular courses and corresponding enrollments during the analysis period in 2020: January 5, 2020 to June 13, 2020.

Table 3: User Summary Statistics

		Existing Users		New Users	
	Active Users	2020 Cohort	2018/19 Cohorts	2020 Cohort	2018/19 Cohorts
Female	0.47	0.53	0.54	0.42	0.45
Age	27.41	26.36	27.30	28.50	27.89
Age: 18-24	0.49	0.55	0.44	0.44	0.45
Age: 25-29	0.17	0.16	0.22	0.16	0.20
Age: 30-39	0.24	0.21	0.25	0.26	0.24
Age: 40-65	0.10	0.08	0.08	0.13	0.11
College Degree	0.52	0.55	0.23	0.57	0.24
Student	0.27	0.40	0.12	0.31	0.15
Jobseeker	0.17	0.22	0.11	0.17	0.10
Employed	0.36	0.29	0.12	0.44	0.14
N	213,671	27,099	124,353	126,667	78,087

Notes: Table displays means of numerous characteristics among different sets of Doroob users. Active users are users who took at least one course betwen January 5, 2020 and June 13, 2020. Age represents age at the time of registration. Student, Jobseeker, and Employed are binary variables identifying each user's reported "current status"; these variable do not sum to one because they are missing for many users.

Table 4: DiD Effects on Enrollments

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
Panel A: Existing Users					
Post x 2020	0.038^{***}	0.021^{***}	0.012^{***}	0.003^{***}	0.002^{***}
	(0.002)	(0.001)	(0.001)	(0.000)	(0.000)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	3483396	3483396	3483396	3483396	3483396
Panel B: New Users					
Post x 2020	0.443^{***}	0.261^{***}	0.122^{***}	-0.027^{***}	0.087^{***}
	(0.041)	(0.025)	(0.018)	(0.010)	(0.002)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users. Panel A displays coefficient estimates of β in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Table 5: DiD Effects on Enrollments (Normalized Outcomes)

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
Panel A: Existing Users					
Post x 2020	1.408***	1.396***	1.408***	0.812***	13.802***
	(0.078)	(0.076)	(0.111)	(0.088)	(0.731)
Pre-COVID 2020 Mean	1.000	1.000	1.000	1.000	1.000
Pre-COVID 2020 SD	15.844	16.253	24.098	26.390	96.085
Num. obs.	3483396	3483396	3483396	3483396	3483396
Panel B: New Users					
Post x 2020	0.236***	0.220***	0.228***	−0.098***	12.565***
	(0.017)	(0.018)	(0.028)	(0.026)	(0.224)
Pre-COVID 2020 Mean	1.000	1.000	1.000	1.000	1.000
Pre-COVID 2020 SD	1.469	1.498	2.497	2.322	13.942
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users. All outcomes are normalized to reflect changes relative to pre-March 15th means (i.e. a coefficient of 1 reflects a 100 percent increase). Panel A displays coefficient estimates of β in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Table 6: DiD Effects on User Persistence

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
Panel A: New Users - Week 1					
Post x 2020	0.443^{***}	0.261^{***}	0.122^{***}	-0.027^{***}	0.087^{***}
	(0.041)	(0.025)	(0.018)	(0.010)	(0.002)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754
Panel B: New Users - Week 2					
Post x 2020	0.265^{***}	0.149^{***}	0.071^{***}	0.030^{***}	0.014^{***}
	(0.019)	(0.011)	(0.008)	(0.003)	(0.001)
Pre-COVID 2020 Mean	0.280	0.152	0.095	0.030	0.002
Pre-COVID 2020 SD	1.646	0.903	0.776	0.240	0.057
Num. obs.	204754	204754	204754	204754	204754
Panel C: New Users - Week 3					
Post x 2020	0.146^{***}	0.087^{***}	0.035^{***}	0.016^{***}	0.008^{***}
	(0.012)	(0.007)	(0.005)	(0.003)	(0.001)
Pre-COVID 2020 Mean	0.142	0.074	0.049	0.018	0.001
Pre-COVID 2020 SD	0.944	0.529	0.447	0.194	0.035
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among new users. Panels A, B, and C display coefficient estimates of in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first, second, and third week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p$<$0.1; **p$<$0.05; ***p$<$0.01).

Table 7: DiD Effects on Enrollments by Age

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
			Panel A: Existing Users		
Post x 2020	0.038***	0.022***	0.011***	0.003***	0.001***
	(0.002)	(0.001)	(0.001)	(0.000)	(0.000)
Post x 2020 x 25-29	−0.008	−0.008***	0.002	−0.004***	0.001***
	(0.006)	(0.003)	(0.003)	(0.001)	(0.000)
Post x 2020 x 30-39	0.004	−0.002	0.003	0.001	0.003***
	(0.006)	(0.003)	(0.003)	(0.001)	(0.000)
Post x 2020 x 40-65	0.005	0.000	−0.001	0.002**	0.003***
	(0.012)	(0.006)	(0.006)	(0.001)	(0.001)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	3483396	3483396	3483396	3483396	3483396
			Panel B: New Users		
Post x 2020	0.447***	0.243***	0.131***	0.035***	0.038***
	(0.044)	(0.028)	(0.021)	(0.011)	(0.002)
Post x 2020 x 25-29	−0.009	0.092	−0.074	−0.060**	0.033***
	(0.103)	(0.064)	(0.049)	(0.026)	(0.005)
Post x 2020 x 30-39	−0.071	0.049	−0.154***	−0.057**	0.092***
	(0.118)	(0.069)	(0.056)	(0.025)	(0.005)
Post x 2020 x 40-65	−0.957***	−0.670***	0.097*	−0.527***	0.143***
	(0.150)	(0.104)	(0.058)	(0.045)	(0.008)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users, separately for users of different ages. Panel A displays coefficient estimates of β (and age group interactions) in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β (and age group interactions) in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Table 8: DiD Effects on Enrollments by Gender

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
			Panel A: Existing Users		
Post x 2020	0.040^{***}	0.021^{***}	0.014^{***}	0.002^{***}	0.003^{***}
	(0.003)	(0.002)	(0.002)	(0.000)	(0.000)
Post x 2020 x Female	-0.003	0.000	-0.004^{*}	0.002^{***}	-0.001^{***}
	(0.004)	(0.002)	(0.002)	(0.000)	(0.000)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	3483396	3483396	3483396	3483396	3483396
			Panel B: New Users		
Post x 2020	0.318^{***}	0.257^{***}	0.020	-0.063^{***}	0.104^{***}
	(0.064)	(0.039)	(0.029)	(0.016)	(0.003)
Post x 2020 x Female	0.222^{***}	-0.001	0.194^{***}	0.067^{***}	-0.038^{***}
	(0.078)	(0.049)	(0.036)	(0.019)	(0.004)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users, separately for men and women. Panel A displays coefficient estimates of β (and gender interactions) in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β (and gender interactions) in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Table 9: DiD Effects on Enrollments by Status

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
Panel A: Existing Users					
Post x 2020	0.008***	0.005***	0.004***	-0.001	0.001***
	(0.003)	(0.001)	(0.001)	(0.001)	(0.000)
Post x 2020 x Employed	0.056***	0.028***	0.017***	0.007***	0.004***
	(0.006)	(0.003)	(0.003)	(0.001)	(0.000)
Post x 2020 x Student	0.047***	0.029***	0.011***	0.006***	0.001***
	(0.004)	(0.002)	(0.002)	(0.001)	(0.000)
Post x 2020 x Jobseeker	0.036***	0.018***	0.011***	0.005***	0.002***
	(0.005)	(0.003)	(0.002)	(0.001)	(0.000)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	3483396	3483396	3483396	3483396	3483396
Panel B: New Users					
Post x 2020	-0.288	-0.123	-0.090	-0.139^{***}	0.064***
	(0.197)	(0.091)	(0.103)	(0.023)	(0.005)
Post x 2020 x Employed	1.543***	0.808***	0.406***	0.240***	0.089***
	(0.223)	(0.113)	(0.113)	(0.033)	(0.007)
Post x 2020 x Student	1.180***	0.599***	0.335***	0.278***	-0.032^{***}
	(0.207)	(0.101)	(0.107)	(0.028)	(0.006)
Post x 2020 x Jobseeker	0.843***	0.551***	0.166	0.125***	0.001
	(0.234)	(0.116)	(0.119)	(0.035)	(0.008)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users, separately for users with different labor force status upon Doroob registration. Panel A displays coefficient estimates of β (and status interactions) in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β (and status interactions) in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Table 10: Effect of Post-COVID Course Enrollments on Employment in September 2020

	Dependent variable:					
	P(Employed in Sept. 2020) x 100					
	(1)	(2)	(3)	(4)	(5)	(6)
All Enrollments	0.240***	0.125**	0.143***			
	(0.073)	(0.054)	(0.054)			
Computer Enrollments				−0.853	0.114	0.084
				(0.525)	(0.408)	(0.413)
General Skills Enrollments				−0.193	−0.073	−0.078
				(0.166)	(0.128)	(0.130)
Occ.-Spec. Enrollments				0.832***	0.390***	0.435***
				(0.198)	(0.150)	(0.149)
Telework Enrollments				1.327	−0.252	−0.106
				(0.906)	(0.687)	(0.691)
Dep. Var Mean	7.08	7.08	7.08	7.08	7.08	7.08
Demographic Controls	N	Y	Y	N	Y	Y
2019 Enrollments Controls	N	N	Y	N	N	Y
Observations	23,739	23,739	23,739	23,739	23,739	23,739
R^2	0.001	0.263	0.269	0.002	0.263	0.269
Adjusted R^2	0.001	0.261	0.263	0.002	0.261	0.263

Notes: Table displays relationship between post-COVID enrollments in Doroob courses and private sector employment status. Coefficients correspond to β_1 in Equation 5. Enrollment counts reflect total enrollments between March 15, 2020 and June 13, 2020. Outcome is a variable equal to 100 for users who were employed in the private sector as of September 19, 2020 and zero otherwise. Analysis is restricted to the 2020 Cohort of existing users: users who joined Doroob between July 1, 2019 and December 31, 2019. Users without individual identifiers in Doroob data, which are necessary to link to GOSI administrative data, are excluded from the analysis. Demographic characteristics include fixed effects for age, gender, employment status at registration, and how the user was directed to the platform. 2019 enrollment controls include fixed effects for number of courses taken in each of the four course categories prior to March 15, 2020. Standard errors clustered at the user level (*p$<$0.1; **p$<$0.05; ***p$<$0.01).

Appendix A A Model of Dynamic Investment in Multi-Dimensional Skills

We study the skill investment decisions of a representative agent in the context of a discrete time model similar to Sanders and Taber (2012). In this model, agents allocate one unit of time in each period between (a) investment in skills and (b) working for a wage. This setup is similar to the canonical life-cycle human capital model in Ben-Porath (1967), with one important difference: skills are multi-dimensional.

The model setup is as follows. Human capital H_t is an M-dimensional vector of skills, where $H_t^{(m)}$ denotes a workers stock of the m-th skill in period t. Workers produce m-specific human capital through a production function $\mathcal{H}^{(m)}$, which takes two inputs: previous human capital H_t and time s_t. Following Sanders and Taber (2012), we assume $\mathcal{H}^{(m)}$ takes the exponential form below.

$$H_{t+1}^{(m)} = \mathcal{H}^{(m)}\left(H_t, s_t^{(m)}\right) = H_t^{(m)} + A_m\left(s_t^{(m)}\right)^{\alpha} \tag{6}$$

where $s_t^{(m)}$ is time devoted to skill m in period t and $\alpha \in (0,1)$.

Wages are based on exogenous skill prices, denoted by an m-dimensional vector π_t. During period t, a worker with skills H_t earns

$$\underbrace{\pi_t' H_t}_{\text{wage per efficiency unit}} \times \underbrace{\left[1 - \sum_{m=1}^{M} s_t^{(m)}\right]}_{\text{efficiency units spent working}} . \tag{7}$$

Finally, users discount future earnings by $1/R$.

With this setup, we consider a three-period ($N = 3$) two-skill ($M = 2$) model and solve it by working backwards. Investing in skills in period 3 provides no return, so a worker's period 3 wage is given by

$$w_3(H_3, \pi_3) = \pi_3' H_3. \tag{8}$$

It follows that in period 2 the worker's value function is given by

$$V_2(H_2, \pi_2) = \left(1 - \sum_{m=1}^{2} s_2^{(m)}\right) \pi_2' H_2 + \frac{1}{R} \pi_3' H_3. \tag{9}$$

And in period 1 the worker's value function is given by

$$V_1(H_1, \pi_1) = \left(1 - \sum_{m=1}^{2} s_1^{(m)}\right) \pi_1' H_1 + \frac{1}{R} \left(1 - \sum_{m=1}^{2} s_2^{(m)}\right) \pi_2' H_2 + \frac{1}{R^2} \pi_3' H_3. \tag{10}$$

For each period t and skill m, the agent selects $s_t^{(m)}$ to equalize the present return on time spent working to the discounted return on time spent investing in skills. This tradeoff is captured by the agent's period 1 first-order condition below.

$$\underbrace{\pi_1 H_1}_{\text{pd. 1 wage}} = \underbrace{\frac{1}{R} \frac{\partial \mathcal{H}^{(m)}\left(s_1^{(m)}\right)}{\partial s_1^{(m)}} \pi_2'}_{\text{return to pd. 1 skilling in pd. 2}} + \underbrace{\frac{1}{R^2} \frac{\partial \mathcal{H}^{(m)}\left(s_1^{(m)}\right)}{\partial s_1^{(m)}} \pi_3'}_{\text{return to pd. 1 skilling in pd. 3}} \tag{11}$$

Plugging in the functional form for $\mathcal{H}^{(m)}$ given above, taking derivatives, and solving for $s_1^{(m)}$ yields the following equation for period 1 investment in skill m.

$$\pi_1 H_1 = \frac{1}{R} A_1 \alpha s_1^{(m)\alpha - 1} \pi_2^{(m)} + \frac{1}{R^2} A_1 \alpha s_1^{(m)\alpha - 1} \pi_3^{(m)} \tag{12}$$

$$= A_1 \alpha s_1^{(m)\alpha - 1} \times \left[\frac{1}{R} \pi_2^{(m)} + \frac{1}{R^2} \pi_3^{(m)}\right] \tag{13}$$

$$\rightarrow \quad s_1^{(m)} = \left[\frac{\alpha A_m \left(\frac{1}{R} \pi_2^{(m)} + \frac{1}{R^2} \pi_3^{(m)}\right)}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}} \tag{14}$$

This equilibrium equation above provides the basis for the propositions below.

Proposition 1. *Investment is higher when future skill prices are higher.*

$$\frac{\partial s_1^{(m)}}{\partial \pi_2^{(m)}} > 0 \tag{15}$$

$$\frac{\partial s_1^{(m)}}{\partial \pi_3^{(m)}} > 0 \tag{16}$$

Proposition 2. *Investment is higher when wages in the current period are lower.*

To see this, denote an agent's period 1 wage per efficiency unit as $w_1 = \pi_1 H_1$. Because skill prices are exogenous, variation across agents in w_1 reflects initial skill endowments. Agents with higher-paid skills will have higher wages. Below, we show that agents with higher-paid skills will invest less in skilling in equilibrium.

$$\frac{\partial s_1^{(m)}}{\partial w_1} = \underbrace{\left[\alpha A_m \left(\frac{1}{R}\pi_2^{(m)} + \frac{1}{R^2}\pi_3^{(m)}\right)\right]^{\frac{1}{1-\alpha}}}_{>0} \times \underbrace{\left[\frac{-1}{1-\alpha}\right]}_{<0} \times \underbrace{\left[\frac{1}{w_1}\right]^{\frac{1}{1-\alpha}+1}}_{>0} < 0 \tag{17}$$

Proposition 3. *In all skills, investment is weakly lower for older agents.*

We show this by considering an agent with 2 working periods rather than 3. We denote this agent's investment by $\tilde{s}$. Following the steps above yields the following equilibrium level of period-1 investment in skill m.

$$\tilde{s}_1^{(m)} = \left[\frac{\alpha A_m (\frac{1}{R}\pi_2^{(m)})}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}} \tag{18}$$

Comparing $\tilde{s}_1^{(m)}$ to $s_1^{(m)}$ shows that levels of skill investment are weakly lower among older agents (agents with fewer remaining working periods).

$$\left[\frac{\alpha A_m(\frac{1}{R}\pi_2^{(m)})}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}} \leq \left[\frac{\alpha A_m(\frac{1}{R}\pi_2^{(m)} + \frac{1}{R^2}\pi_3^{(m)})}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}} \tag{19}$$

$$\tilde{s}_1^{(m)} \leq s_1^{(m)} \tag{20}$$

Proposition 4. *Older individuals invest relatively more in skills that are valuable in the short-run.*

To see this, note that period 2 and period 3 skill prices take the following form.

$$\pi_2 = \begin{bmatrix} \pi_2^{(1)} \\ \pi_2^{(2)} \end{bmatrix} \quad \pi_3 = \begin{bmatrix} \pi_3^{(1)} \\ \pi_3^{(2)} \end{bmatrix}$$

Suppose skill 1, whose period 2 and 3 prices are represented by $\pi_2^{(1)}$ and $\pi_3^{(1)}$, respectively, is more valuable in period 1. $\pi_2^{(1)} > \pi_3^{(1)}$. Oppositely, suppose that skill 2 is more valuable in period 2. $\pi_2^{(2)} < \pi_3^{(2)}$.

Define θ as the ratio of an agent's skill investment in skill 1 versus skill 2 in period 1. (Mathematically, $\theta = s_1^1/s_1^2$.) Following from above, θ takes the following form for agents with 3 working periods remaining in their life-cycle.

$$\theta = \frac{s_1^1}{s_1^2} \tag{21}$$

$$= \frac{\left[\frac{\alpha A_1(\frac{1}{R}\pi_2^{(1)}+\frac{1}{R^2}\pi_3^{(1)})}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}}}{\left[\frac{\alpha A_2(\frac{1}{R}\pi_2^{(2)}+\frac{1}{R^2}\pi_3^{(2)})}{\pi_1 H_1}\right]^{\frac{1}{1-\alpha}}} \tag{22}$$

$$= \left[\frac{\alpha A_1(\frac{1}{R}\pi_2^{(1)} + \frac{1}{R^2}\pi_3^{(1)})}{\alpha A_2(\frac{1}{R}\pi_2^{(2)} + \frac{1}{R^2}\pi_3^{(2)})}\right]^{\frac{1}{1-\alpha}} \tag{23}$$

$$= \left[\frac{A_1(\pi_2^{(1)} + \frac{1}{R}\pi_3^{(1)})}{A_2(\pi_2^{(2)} + \frac{1}{R}\pi_3^{(2)})}\right]^{\frac{1}{1-\alpha}} \tag{24}$$

As before, we can denote θ for agents with 2 remaining working periods as $\tilde{\theta}$.

$$\tilde{\theta} = \frac{\tilde{s}_1^1}{\tilde{s}_1^2} \tag{25}$$

$$= \left[\frac{A_1(\frac{1}{R}\pi_2^{(1)})}{A_2(\frac{1}{R}\pi_2^{(2)})}\right]^{\frac{1}{1-\alpha}} \tag{26}$$

$$= \left[\frac{A_1(\pi_2^{(1)})}{A_2(\pi_2^{(2)})}\right]^{\frac{1}{1-\alpha}} \tag{27}$$

Below, we show that $\theta < \tilde{\theta}$.

$$\pi_2^{(1)} > \pi_3^{(1)} \geq 0, \quad 0 \leq \pi_2^{(2)} < \pi_3^{(2)} \quad \rightarrow \quad \pi_2^{(2)} * \pi_3^{(1)} < \pi_2^{(1)} * \pi_3^{(2)} \tag{28}$$

$$\pi_2^{(2)} * \frac{1}{R}\pi_3^{(1)} < \pi_2^{(1)} * \frac{1}{R}\pi_3^{(2)} \tag{29}$$

$$\pi_2^{(2)} * \frac{1}{R}\pi_3^{(1)} + \pi_2^{(2)} * \pi_2^{(1)} < \pi_2^{(1)} * \frac{1}{R}\pi_3^{(2)} + \pi_2^{(2)} * \pi_2^{(1)} \tag{30}$$

$$\pi_2^{(2)} * (\pi_2^{(1)} + \frac{1}{R}\pi_3^{(1)}) < \pi_2^{(1)} * (\pi_2^{(2)} + \frac{1}{R}\pi_3^{(2)}) \tag{31}$$

$$\left[\frac{(\pi_2^{(1)} + \frac{1}{R}\pi_3^{(1)})}{(\pi_2^{(2)} + \frac{1}{R}\pi_3^{(2)})}\right] < \frac{\pi_2^{(1)}}{\pi_2^{(2)}} \tag{32}$$

$$\left[\frac{A_1(\pi_2^{(1)} + \frac{1}{R}\pi_3^{(1)})}{A_2(\pi_2^{(2)} + \frac{1}{R}\pi_3^{(2)})}\right]^{\frac{1}{1-\alpha}} < \left[\frac{A_1(\pi_2^{(1)})}{A_2(\pi_2^{(2)})}\right]^{\frac{1}{1-\alpha}} \tag{33}$$

$$\rightarrow \quad \theta < \tilde{\theta} \tag{34}$$

Appendix B Existing Users Analysis: Low- vs. High-Attachment Users

Our analyses of existing users rely on a parallel trends assumption: absent the COVID shock, enrollments among 2020 Cohort users would have followed the same path as 2018 and 2019 Cohorts. In Figure 4, control cohort users who joined Doroob in 2H 2018 appear to exhibit negative pre-trends, relative to the treatment cohort, which consists of users who joined Doroob in 2H 2019. Negative pre-trends in this cohort generate significant coefficients on some pre-COVID relative week indicators in our event study.

In this appendix, we show that these pre-trends are driven primarily by users who took exactly one course during their first 30 days on Doroob: users we refer to as "low-attachment users." Removing these users from our existing users estimates produces estimates of slightly larger magnitude as those in the body of the paper, with little evidence of pre-trends in enrollment behaviors.

We first illustrate the distribution of "attachment" to Doroob among all three existing user cohorts. Figure B1 shows the distribution of enrollments for all three cohorts during the first 30 days on Doroob. Across all cohorts, the largest single group is users who enroll in exactly one course. This group comprises between 35 and 55 percent of each cohort. (Recall that our existing user analysis includes only users who enrolled in at least one course during their first 30 days on the Doorob platform, so there are no users in Figure B1 with zero enrollments.)

Figure B2 shows cohort-specific raw weekly mean enrollments for two groups: low-attachment users: users who took exactly one course during their first 30 days on Doroob, and high-attachment users: users who took more than one course over the same period. Visually, differences in pre-trends between the treatment cohort and the control cohorts appear to be much larger among low-attachment users. More specifically, negative pre-trends are concentrated among users who joined in the second half of 2018 (shown in

Figure B2 in blue).

More formally, Figure B3 displays event study coefficients separately for these two groups. The blue series in Figure B3 displays event study coefficients for low-attachment users. This panel shows some evidence of pre-trends—negative and statistically significant coefficients on relative week indicators in periods -9 and -8. The red series performs the same analysis, restricting the sample to high-attachment users. This analysis shows little evidence of pre-trends. The event study coefficients exhibit no evidence of systematic, differential trends between treatment and control cohorts prior to March 15.

In the body of the paper, Figure 4 finds coefficient estimates of between 0.025 and 0.05 enrollments per week in the post-COVID period. Both series in Figure B3 are roughly in line with these estimates, though high-attachment users (users in the red series of Figure B3) exhibit slightly larger coefficient estimates in the weeks immediately following March 15: estimates between 0.05 and 0.07 enrollments per user per week.

Figure B1: Distribution of Pre-Analysis Period Enrollments

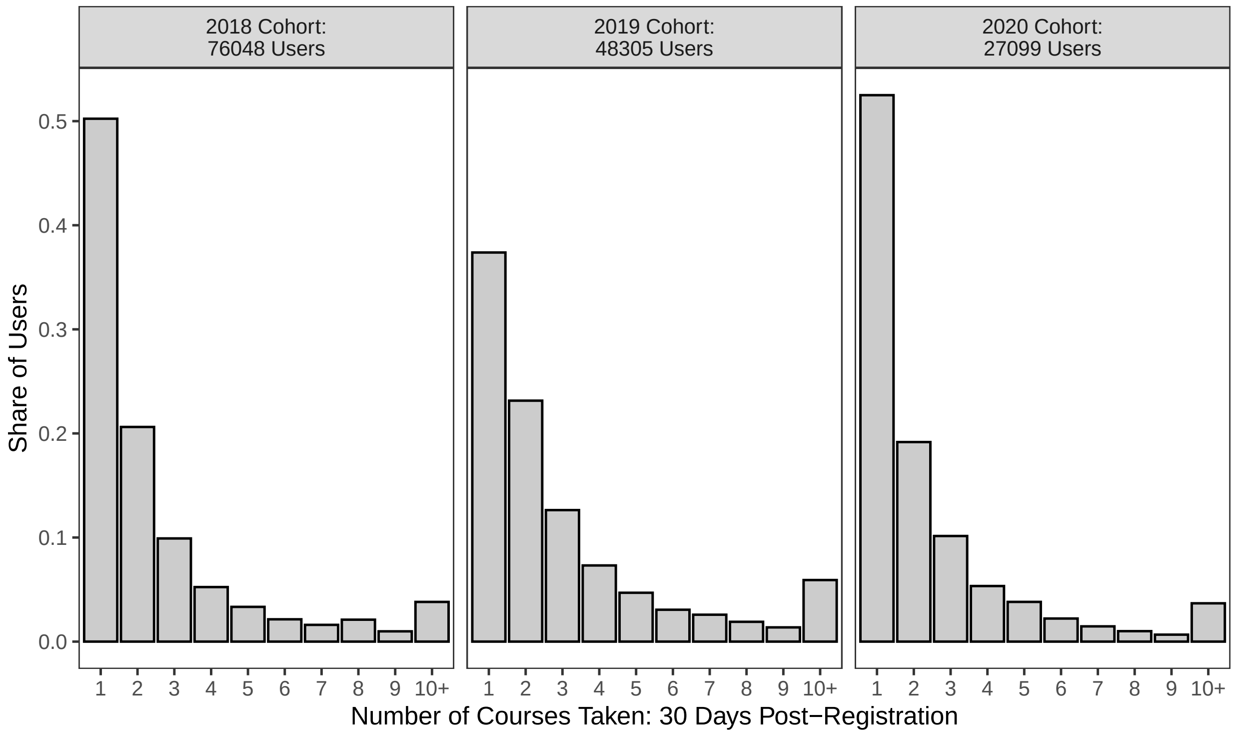

Notes: Figure displays the distribution of total enrollments during each user's first 30 days on Doroob, separately for all intensive margin cohorts. Users who registered for Doroob but did not take any courses in the following 30 days are excluded from the analysis.

Figure B2: Enrollments by Pre-Analysis Period Enrollments

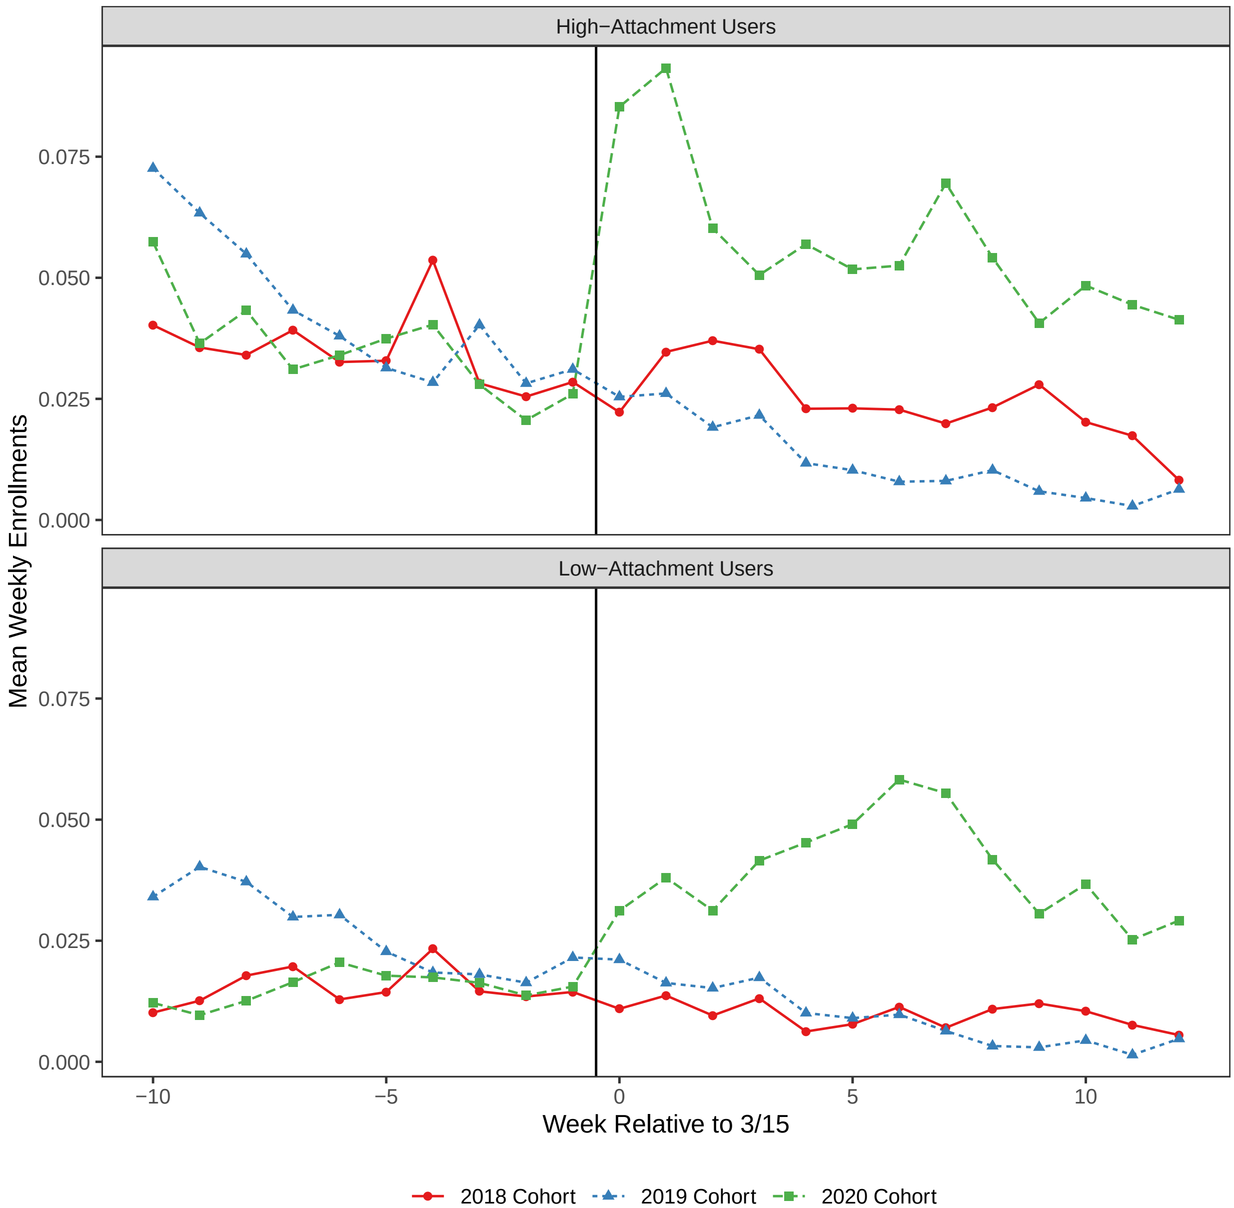

Notes: Figure displays the raw means of weekly minutes enrolled per user, separately for treated and control cohorts. The top panel restricts the analysis to users with exactly one enrollment during their first 30 days on Doroob. The bottom panel restricts the analysis to users with more than one enrollment over the same period.

Figure B3: Intensive Margin Response: Minutes Enrolled by Pre-Analysis Period Enrollments

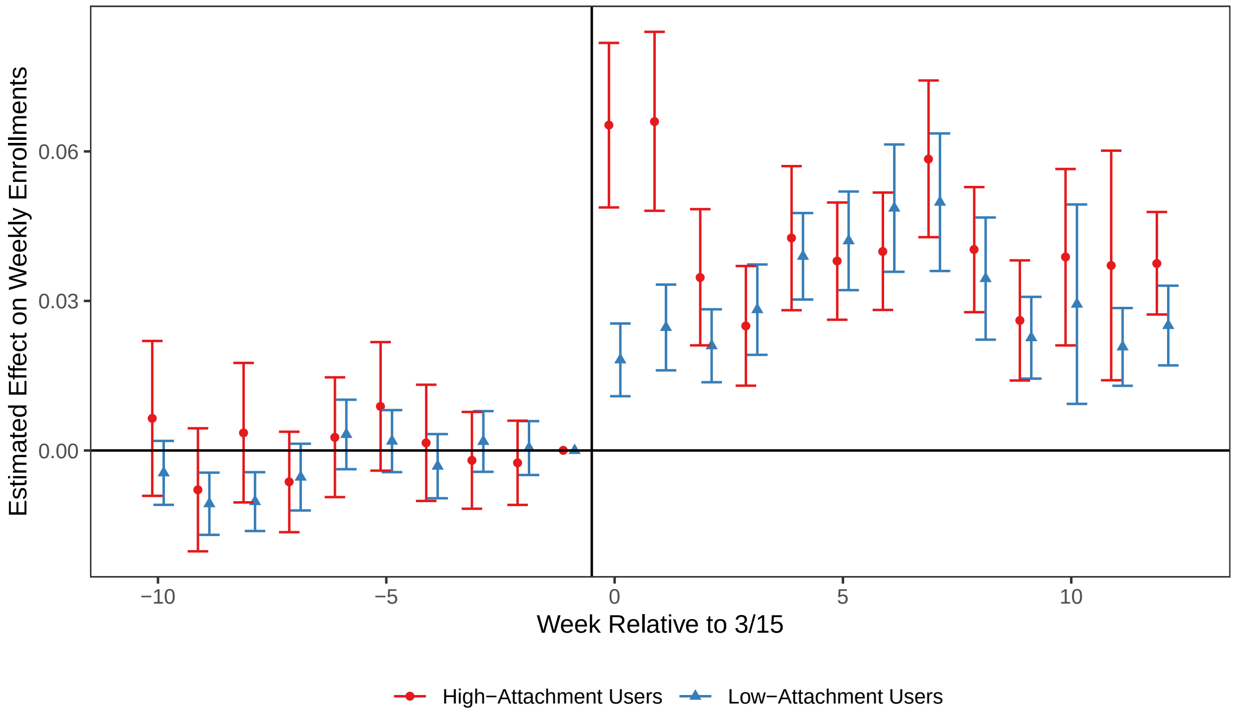

Notes: Figure displays estimates of β_t in Equation 2. These estimates control for user and calendar week fixed-effects. Error bands represent 95% confidence intervals. The top panel restricts the analysis to users with exactly one enrollment during their first 30 days on Doroob. The bottom panel restricts the analysis to users with more than one enrollment over the same period.

Appendix C Intensive Margin Analysis: Coarsened Exact Matching

As shown in Table 3, our intensive margin treatment and control cohorts differ in some observable characteristics. For example, users in our treatment cohort are roughly one year younger than those in the control cohort. In this appendix, we assess whether matching on user characteristics changes our main results.

Specifically, we use coarsened exact matching (Iacus et al. (2012)) to balance user demographics in the treatment and control cohort. We balance on all variables displayed in Table 3.[26] This procedure provides weights for all users in our sample based on the CEM algorithm; users in the control cohort that are more similar to those in the treatment cohort are assigned higher weights, and vice versa.

With these weights, we run our OLS event study models, weighted by CEM-weights. Figure C1 shows our main unweighted event study estimates alongside our CEM-weighted estimates for all courses. CEM-weighted results are nearly identical to the unweighted OLS results reported in the body of the paper.

[26] We conduct this matching procedure at the user level, and link the weights generated by this procedure to our user-by-week panel data.

Figure C1: Response of Existing Users: CEM-Weighted Results

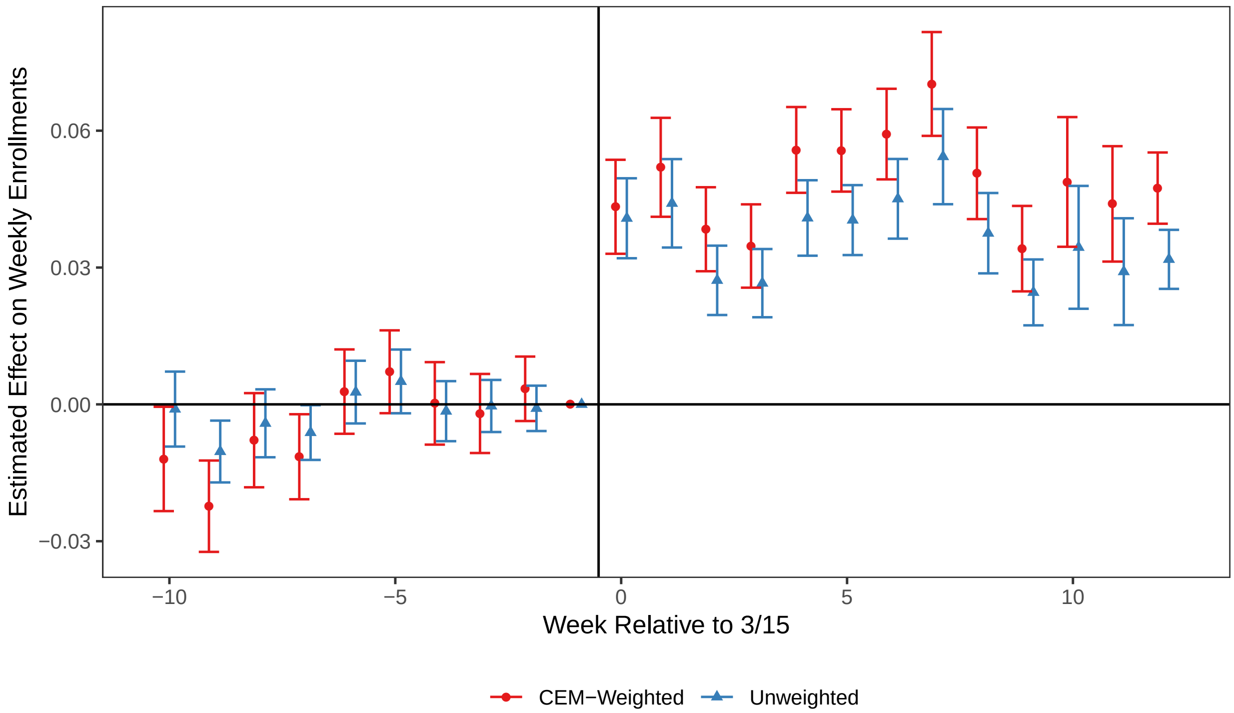

Notes: Figure displays how the COVID shock affected enrollments among existing users. Plotted coefficients are estimates of β_t in Equation 2. These estimates control for cohort and calendar week fixed-effects. Error bars represent 95% confidence intervals. CEM-weighted regressions use weights produced by a coarsened exact matching algorithm, as described in the text.

Appendix D Heterogeneity in Regional COVID-Rates

In this appendix we examine whether heterogeneity in responses across regions reflects patterns in the incidence of COVID-19 cases. First, we demonstrate that initial rates of COVID-19 cases in Saudi Arabia were concentrated in 4 of the 13 Saudi provinces. Second, we estimate difference-in-differences models with interactions with indicators for these 4 provinces. The results suggest that COVID-induced increases in course taking were slightly higher in areas with more COVID-19 cases.

Saudi Arabia is divided into 13 administrative regions, or provinces, defined geographically. Figure D1 displays a map of these provinces alongside the cumulative COVID-19 case rates per 1 million residents as of June 30, 2020. As shown in Figure D1, while all regions had some incidence of COVID-19 by June 2020, 4 regions have distinctly higher values than others. Specifically, Ar Riyad, Eastern Region, Mecca, and Medina all had rates above 6,000 cases per 1 million residents. No other region had rates above 3,500 cases per 1 million.

We assess whether these 4 regions—which we refer to as "High COVID Provinces"—had larger enrollment responses to the COVID shock. To do so, we estimate difference-in-differences models with interactions for users who are located in one of the 4 High COVID Provinces.

Table D1 displays results. Column 1 indicates that, among existing users, users in High COVID Provinces exhibited a 30 percent larger enrollment response to the COVID-shock (versus users in other provinces). This pattern is roughly similar among new users; among these users, those in High COVID Provinces exhibit a 45 percent larger enrollment response. (Not all users indicate a location upon registering for Doorob, which is why sample sizes in Table D1 are slightly smaller than those in the body of the paper.) These patterns persist across most course types and user populations, with the exception of new users' enrollments in telework courses, where responses from users in High COVID

Provinces were slightly smaller.

Figure D1: Cumulative COVID-19 Case Rates per 1 Million as of 6/30/2020

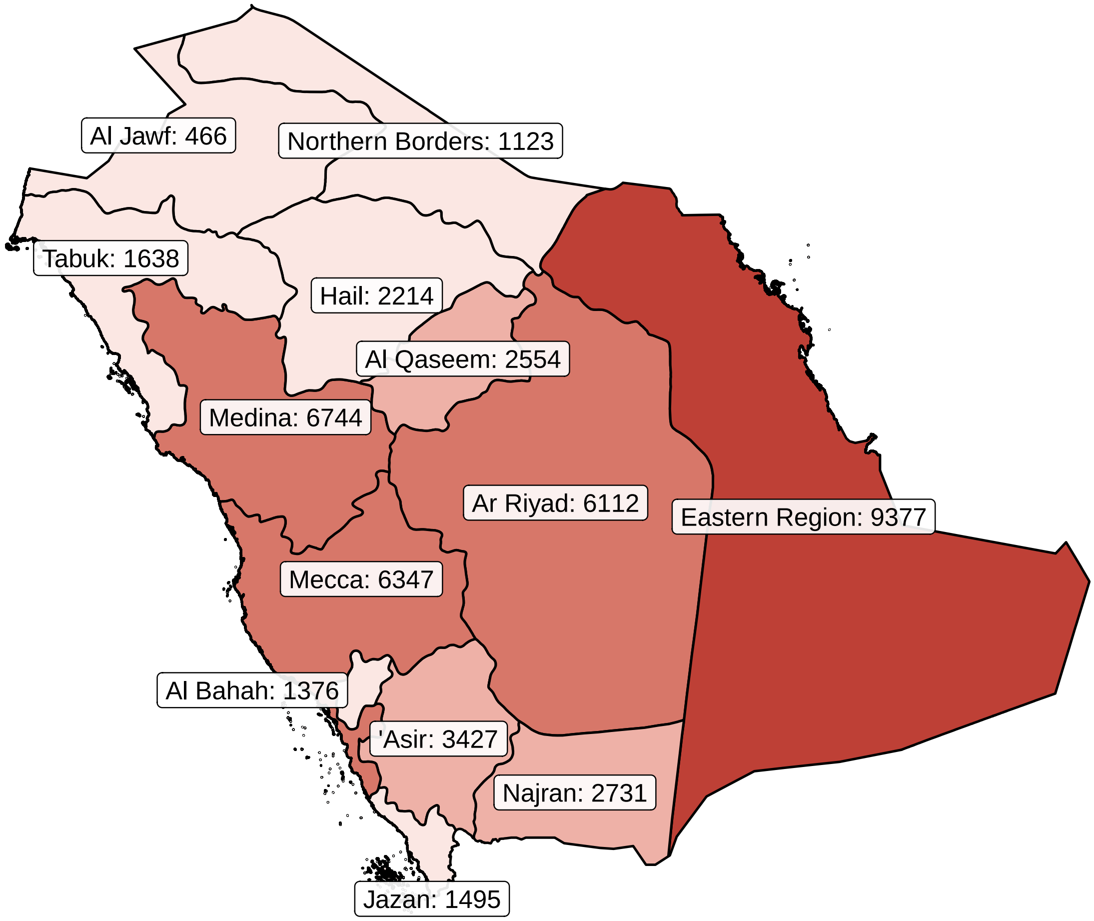

Notes: Figure displays the cumulative COVID-19 case rates per 1 million across Saudi provinces as of June 30, 2020. Case rates by provide data are from the Saudi KAPSARC COVID Data Portal. Colors reflect the displayed values: darker regions are those with higher cumulative COVID-19 case rates per 1 million as of June 30, 2020

Table D1: DiD Effects on Enrollments by Regional COVID Incidence

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
			Panel A: Existing Users		
Post x 2020	0.044***	0.025***	0.012***	0.004***	0.002***
	(0.004)	(0.002)	(0.002)	(0.001)	(0.000)
Post x 2020 x High Covid	0.013***	0.006**	0.005**	0.002**	0.000
	(0.005)	(0.003)	(0.002)	(0.001)	(0.000)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	1695859	1695859	1695859	1695859	1695859
			Panel B: New Users		
Post x 2020	0.726***	0.442***	0.123**	0.062**	0.099***
	(0.124)	(0.072)	(0.058)	(0.025)	(0.006)
Post x 2020 x High Covid	0.328**	0.114	0.217***	0.010	−0.013**
	(0.135)	(0.080)	(0.063)	(0.028)	(0.006)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	161198	161198	161198	161198	161198

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users, separately for users of in High- versus Low-COVID Provinces. Panel A displays coefficient estimates of β (and High-COVID Province interactions) in Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β (and High-COVID Province interactions) in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level (*p<0.1; **p<0.05; ***p<0.01).

Appendix E Age-by-Gender Effects

In this appendix we examine heterogeneity in responses to the COVID-19 shock with respect to age and gender. We present results with respect to age and gender heterogeneity (individually) in the body of the paper. In this appendix, we present results with respect to the interaction of these two characteristics.

To do so, we define four distinct groups of users: younger women, older women, younger men, and older men. Throughout, we define "younger" groups as users between age 18 and 29, and "older" groups as users between age 30 and 65.

Before showing results with respect enrollment patterns, we first show evidence that these groups exhibited different overall registration responses in the weeks following the COVID-19 shock. Figure E1 displays the weekly number of new users surrounding March 15 for 2018, 2019, and 2020, separately for the four age-by-gender groups defined above. All values in Figure E1 reflect percentage increases above their pre-March 15 values. Across all groups, Figure E1 illustrates that the COVID-19 shock brought about a massive increase in registrations; in the weeks following March 15, all four groups exhibit increases in registrations above 400%. These increases are not identical across groups; enrollments of older men increase by nearly 3,000% at their peak, higher than all other groups. Younger women exhibit the smallest relative increase; at their peak registrations for this group increased by nearly 500%.

We explore differential responses with respect to enrollment patterns in Table E1. We interpret these estimates in the context of our main results with respect to age and gender.

First, the results with respect to existing users suggest that, broadly, older women's enrollments were less responsive to the COVID shock than younger women or older men. Specifically, older men exhibit larger enrollment responses than younger men overall, and for 3 of 4 individual course types. Similarly, women exhibit larger enrollment responses overall and for the 2 largest individual course types, general and occupation-specific courses. However, older women exhibit relatively smaller responses across all course types, as indicated by the negative coefficient on the interaction between older users and female users. Consistent with our main results, older users' responses to the COVID shock were particularly large for telework courses. Among existing users, these responses were concentrated among men, as indicated by the negative coefficient on the interaction between older users and female users.

Second, among new users, age and gender effects tend to be larger than interaction terms between the two. Specifically, older users exhibit large reductions in enrollments in all non-telework courses, but relatively larger increases in telework courses. For telework courses, there is no evidence for differential age effects between men and women; this interaction term is small and not statistically distinguishable from zero. Enrollment responses among women are larger overall, an effect driven by occupation-specific courses. Consistent with the results in the body of the paper, new women on Doroob exhibit lower telework enrollment responses relative to men.

Figure E1: Weekly Registrations by Age and Gender Group

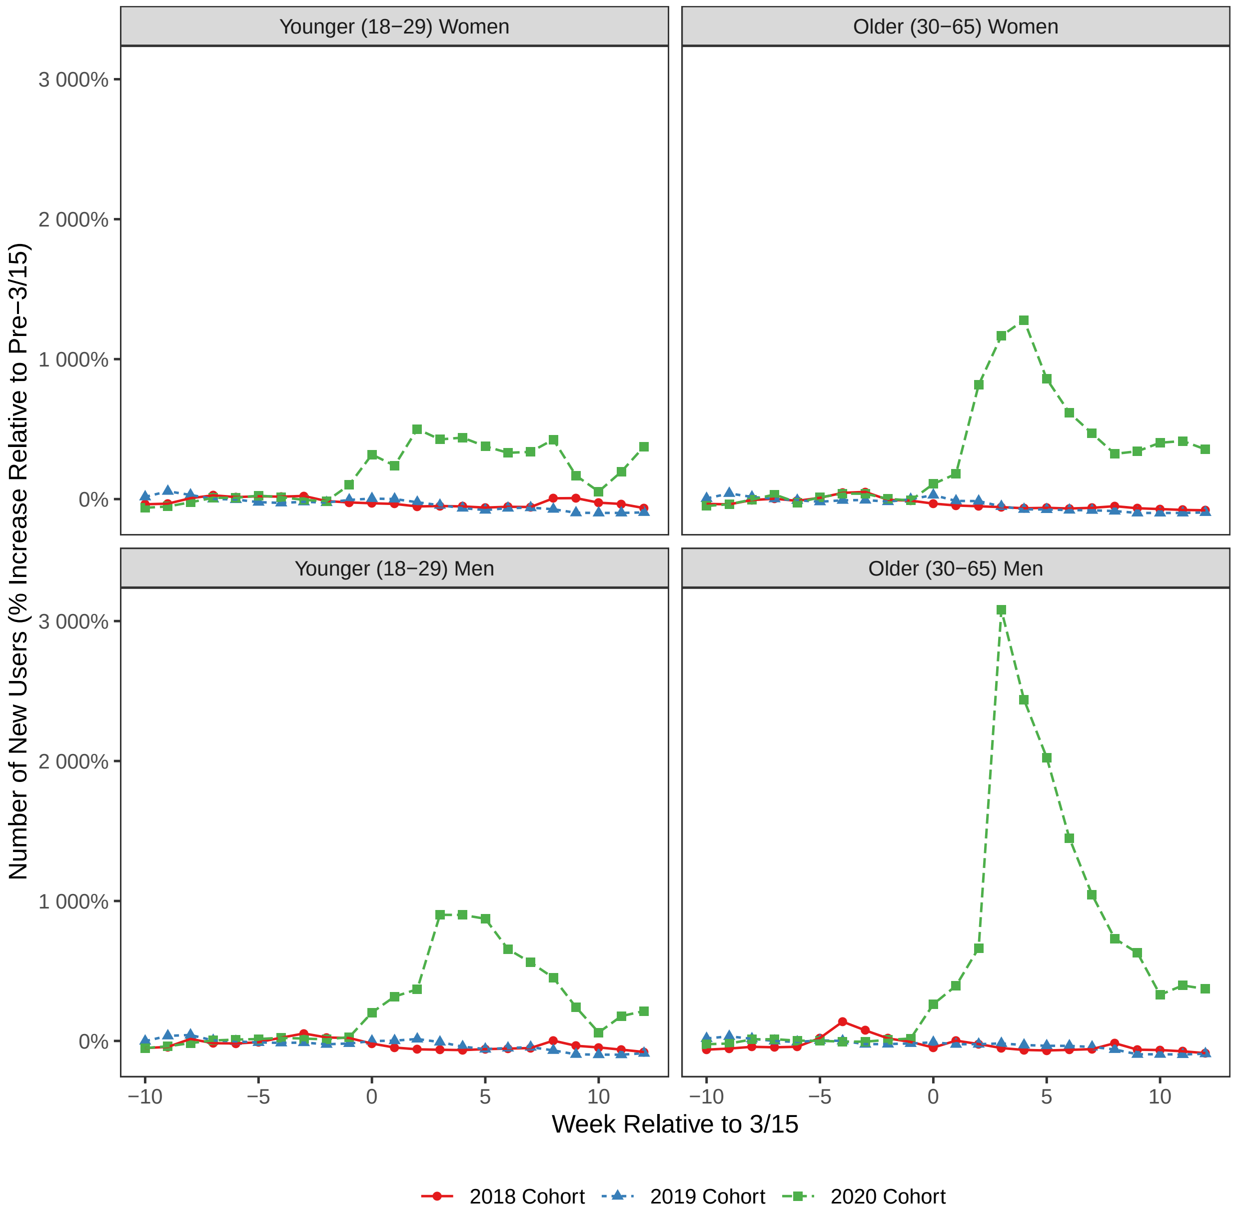

Notes: Figure displays how the COVID shock affected the number of new Doroob users, separately by age and gender groups. Displayed values represent percentage increases above each group's pre-March 15 mean (separately for each cohort).

Table E1: DiD Effects on Enrollments by Age and Gender Group

	All	General	Occ.-Spec.	Computer	Telework
	(1)	(2)	(3)	(4)	(5)
Panel A: Existing Users					
Post x 2020	0.030***	0.017***	0.012***	−0.000	0.002***
	(0.002)	(0.001)	(0.001)	(0.000)	(0.000)
Post x 2020 x Older (30-65)	0.028***	0.012***	0.007	0.006***	0.003***
	(0.009)	(0.004)	(0.004)	(0.001)	(0.001)
Post x 2020 x Female	0.011***	0.007***	0.000	0.005***	−0.001***
	(0.004)	(0.002)	(0.002)	(0.001)	(0.000)
Post x 2020 x Older (30-65) x Female	−0.044***	−0.022***	−0.013**	−0.008***	−0.002**
	(0.010)	(0.005)	(0.005)	(0.001)	(0.001)
Pre-COVID 2020 Mean	0.025	0.013	0.009	0.003	0.000
Pre-COVID 2020 SD	0.393	0.218	0.208	0.069	0.015
Num. obs.	3483396	3483396	3483396	3483396	3483396
Panel B: New Users					
Post x 2020	0.390***	0.324***	−0.004	0.007	0.062***
	(0.070)	(0.042)	(0.033)	(0.017)	(0.003)
Post x 2020 x Older (30-65)	−0.428***	−0.322***	−0.008	−0.190***	0.093***
	(0.145)	(0.088)	(0.066)	(0.034)	(0.006)
Post x 2020 x Female	0.182**	−0.074	0.253***	0.033*	−0.030***
	(0.082)	(0.052)	(0.039)	(0.020)	(0.004)
Post x 2020 x Older (30-65) x Female	0.086	0.209*	−0.129	−0.003	0.008
	(0.197)	(0.119)	(0.091)	(0.044)	(0.010)
Pre-COVID 2020 Mean	2.090	1.187	0.623	0.273	0.007
Pre-COVID 2020 SD	3.071	1.777	1.556	0.633	0.100
Num. obs.	204754	204754	204754	204754	204754

Notes: Table displays the estimated effect of the COVID-19 shock on enrollments among existing users and new users, separately for men and women in different age groups. Panel A displays coefficient estimates of β (and age, gender, and age-by-gender interactions) in

Equation 1 and estimate the effect of the COVID-19 shock on weekly enrollments of existing users. Panel B displays coefficient estimates of β (and age, gender, and age-by-gender interactions) in Equation 3 and estimate the effect of the COVID-19 shock on enrollments of new users during their first week on the platform. Columns reflect different course types. Examples of each course type are given in Table 2. Standard errors clustered at the user level ($^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$).

Appendix F New Users Analysis: Employer Registration Concentration

In this appendix we examine whether Doroob registration patterns suggest that employers encouraged their employees to join Doorob in the months following the COVID-shock.

To do so, we use the linked GOSI-enrollments data. We assess the presence of clustering behavior by examining the co-occurrence of registrations from users in the same week working at the same firm. Specifically, for each newly registered user, we count the number of users who registered from the same firm in the same week. (Note that our GOSI employment data is from September 2020, so our information on employment is from *after* these users registered for Doroob.) We then compare the distribution of same-firm-same-week registrations among newly-registered users pre- and post-COVID shock. The figure below compares these two distributions.

Most newly-registered Doroob users cannot be linked to our GOSI data; over 90 percent of new users fall into this category. However, among those who can, employer concentration appeared to rise following the COVID-shock. However, the scale of these differences is small.

Figure F1: Employer Registration Concentration

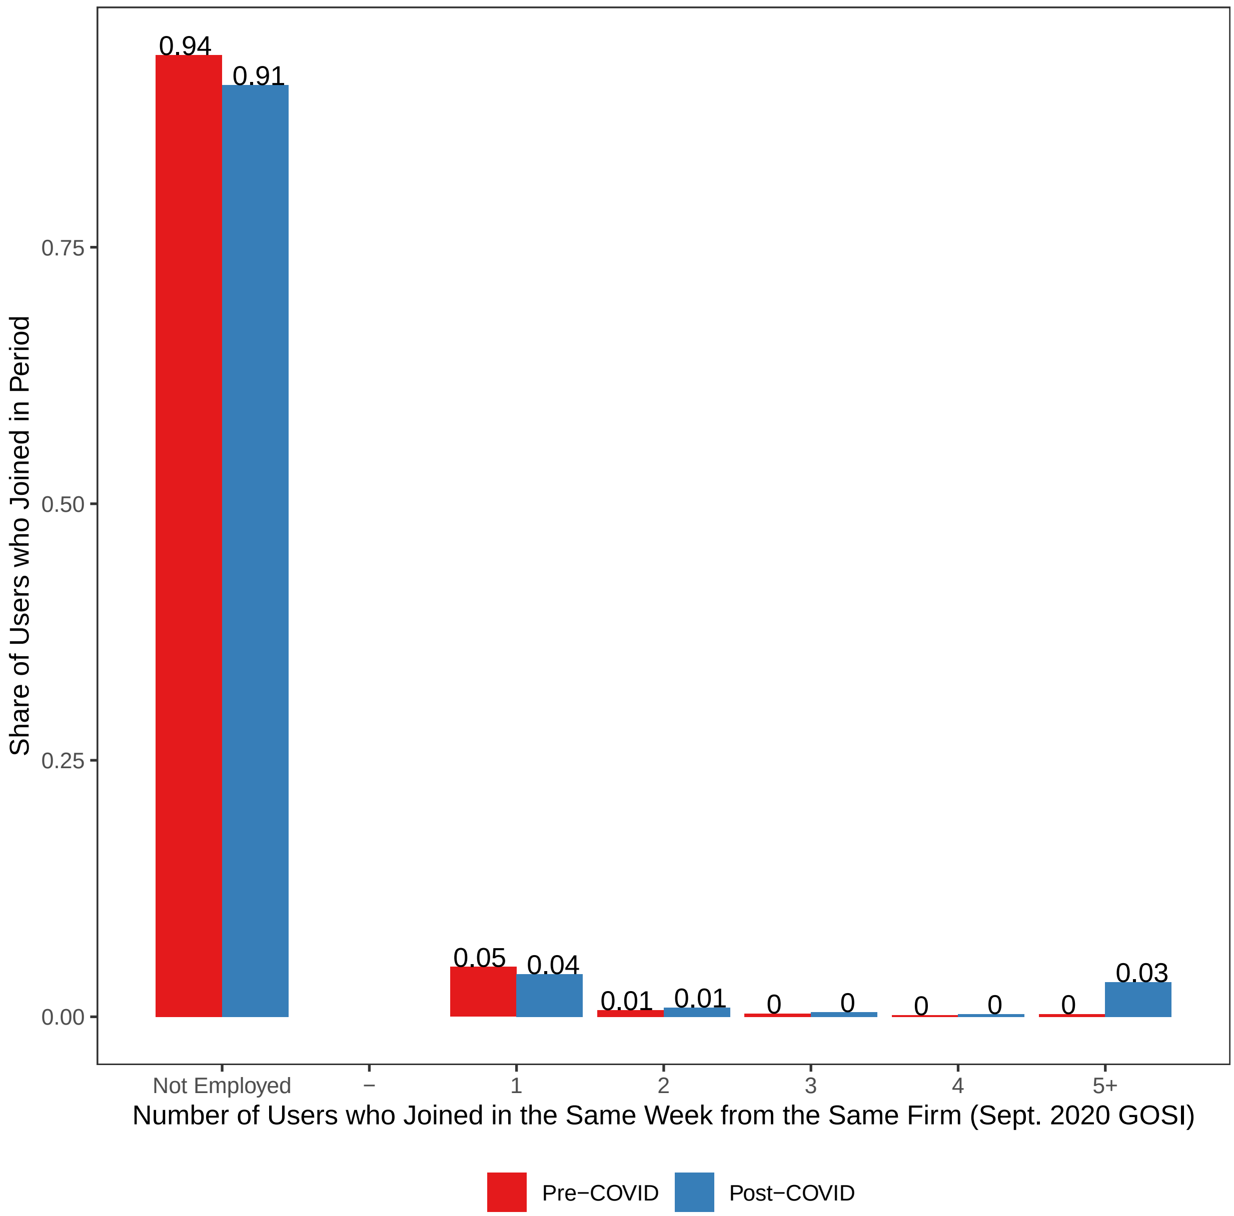

Notes: Figure displays the distribution of employer co-registration among new Doroob users. For each new user, we count the number of users working for the same firm who registered in the same week. Users who were the only individuals from their firms to sign up for Doorob in that week have value one. Users who were one of two individuals from their firms to sign up for Doroob in that week have value two, and so forth. The pre-COVID period is the three months preceding March 15, 2020, and the post-COVID period is the three months thereafter.