Isn't that NEET? Increasing Engagement with Online Learning Tools for Unemployed Youth

Elizabeth Dhuey[1], Jill Furzer, Michal Perlman and Linda White

Abstract

Online learning tools have had a dramatic increase in use recently. Nevertheless, causal evidence around the effectiveness of online tools is minimal, and user engagement remains a significant limitation. As one potential mechanism to increase user engagement, we deployed informational nudging to induce higher online learning use in a randomized controlled trial design. Working with local employment service providers, we provided online learning licences to youth not employed or in education and training (NEET), with some randomly assigned to receive informational nudges. We incorporate nudging into more prescriptive settings in two additional studies, including an interactive onboarding conference and a job skills training program. Overall, our results show higher online learning engagement from baseline and that use increased with the intensiveness of engagement supports. We find little effect of informational nudging on online engagement on its own. However, we identify substantial peer spillover effects in use due to nudging, suggesting that nudging increased the direct or perceived value of the online learning tool and motivation to engage.

Funding: This research is funded by a grant from the Social Sciences and Humanities Research Council of Canada (SSHRC) grant numbers 890-2018-0020 and 435190367-S. Accelerate internship for Jill Furzer was funded by MITACS grant number IT16170.

[1] elizabeth.dhuey@utoronto.ca. University of Toronto, 121 St. George Street, Toronto, ON, M4J5B7, Canada

Introduction

The onset of the COVID-19 pandemic moved most education online almost overnight in March 2020. However, causal evidence on the benefits of learning online for education and career outcomes is limited. The available causal evidence, focusing primarily on university-enrolled students, shows mixed results (Figlio et al., 2013; Bowen et al., 2014; Joyce et al., 2015; Alpert et al., 2016). Nevertheless, limited engagement with online learning hinders much of this evaluation (Bergdahl, Nouri, and Fors, 2020: Henderson, Selwyn, and Aston, 2017; Hew and Cheung, 2014; Hew, Qiao, and Tang, 2018; Kim et al., 2020; Terras and Ramsay, 2015). Furthermore, the focus of previous research on university-enrolled students may not be representative of the individuals most likely to benefit from online learning and tools, such as those not employed or in education and training who may need to rely more on online tools to acquire job training skills at a low cost. As job displacement is most likely to affect those without a university or college degree (Brotman, 2016), estimating online learning outcomes for this group is vitally important.

In this paper we use an experimental design to understand willingness to engage with online platforms and subsequent education or employment outcomes for youth not employed or in educational training (NEET). Youth aged 18 to 29 were enrolled in one of three programs after initiating contact with a local employment service provider (ESP). In the first program, participants received a licence directly from their local ESP. In the second program, youth received a licence during a one-day job skills conference in which one of the activities was to set up their online learning program account. The third and final program incorporated the online learning tool within an intensive six-week information technology (IT) skills training program. Half of the youth participants in each study were randomized to receive additional bi-weekly informational nudges to encourage online learning usage. In the final program, we implemented a two-stage randomization design to account for potential peer spillover effects (Duflo and Saez, 2002). Over eleven IT training cohorts, we first randomized what proportion of each cohort would receive informational nudges. In a second randomization stage, we randomized the individuals within each cohort, corresponding to a specified treatment proportion, to receive treatment.

Within these three programs, we tested the role of informational nudges and other engagement supports, namely onboarding and embedment, on online learning engagement in terms of activation, usage, and types of content assessed. The randomized encouragement design circumvents issues of selection into online learning use. As we provide all participants with the learning platform free of change, cost barriers to online learning use are also alleviated.

Across all three studies, we found higher engagement with online learning than typically seen for this particular platform. Activation rates ranged from 50 to nearly 100 percent, moving linearly with the program's intensity. Nevertheless, we find little evidence that simple additional informational nudges changed usage behaviour for participants between treatment and control groups. We fail to reject the null hypothesis of zero treatment effect for minutes used or the number of logins both overall and by individual months.

Nevertheless, we do find evidence suggesting peer spillover effects in engagement due to informational nudging in our final study involving integration of online learning into an intensive IT training program. When some but not all individuals in a given training cohort were randomized to receive information nudges, use increased in both the treatment and control groups. Control group engagement with the online learning platform increased as more peers within their given cohort received treatment. One theory of this spillover effect is that the nudge places true value on the commodity. Suppose the nudge points individuals in the treatment group towards practical online learning courses and improved job placement outcomes because of skills acquired from these additional videos. In that case, treatment group members may encourage control group members to watch these videos to improve their own outcomes. A second hypothesis is that the nudge increases only the perceived value of the online learning platform if control group members see that some members receive additional information that they did not receive. Through aggregated data, we are able to see program graduation and employment outcomes three months post-graduation from program for treatment and control group but do not find differences between treatment and control groups. This may again be because of peer spillover effects or because there are no real effects garnered from the additional online learning due to informational nudging. We find a strong correlation between outcomes between treatment and control groups, but smaller correlations with the proportion of individuals under treatment in a given cohort. Due to data limitations, we are unable to assess true treatment and peer spillover effects. Estimating the benefits of online learning tool use on job placement outcomes remains an area of future research.

As a final contribution, we assess how engagement with online learning changed with the COVID-19 pandemic and the subsequent social distancing policy responses. We find a large increase in use immediately following the move online, with treatment group participants more likely to increase usage in the post-pandemic period. After the initial wave of the pandemic, usage returned to near normal levels and no differential use between treatment and control groups were detected. These results suggest that informational nudging was useful for treatment group participants to highlight potential course content in the early pandemic when many were motivated to seek additional beneficial past-times while social distancing. When pandemic fatigue set in, or as social distancing rules loosened over the summer 2020 this motivation bias appeared to disappear.

This work highlights that engagement with online learning tools remains one of the main barriers to online learning benefits. We find that simple low-cost informational nudging does little to incentivize participants to increase engagement with the platform. We do, however, find that more intensive onboarding to online learning in person with knowledgeable technicians helps to improve use. Embedding online learning within an intensive in-person learning program also helps to encourage use. We additionally find strong peer effects in use, with control group participants seeking out additional engagement if more of their peers are receiving informational nudging. However, we cannot rule out whether these peer spillover effects are from true value-add of online learning on job outcomes or merely perceived value-add. As demonstrated with a large uptick in use during the COVID-19 pandemic's first wave, and subsequent drop in use, individual motivation to use online learning remains a major factor in use.

Background: Experimental Design

Starting in October 2019, we partnered with a local civic engagement organization to provide an online learning tool, LinkedIn Learning, at no cost to youth aged 18 to 29 in the Greater Toronto Area (GTA) who identified as unemployed and were working with a local employment service provider (ESP). The partner organization works to improve access to skills development and learning opportunities for youth, including racialized youth, youth living in poverty, and new immigrants.

This project was based on an initial pilot study implemented in 2018/2019 where NEET youth were provided a free licence to LinkedIn Learning (a platform that typically costs $1200). The youth advocacy organization invested significantly in curating a video curriculum and original content creation on foundational soft skills and digital literacy, as well as skills on learning, resiliency and financial literacy. Examples of foundational soft skills in the curated content included effective communication in the workplace, collaboration skills, prioritization and organization skills. Digital literacy skills included 101 tutorials on Microsoft Excel, collaboration tools, and troubleshooting common PC errors. As the NEET population enrolled in the study was ethnically and racially diverse, the youth advocacy organization also made a concerted effort to choose or produce videos with diverse representation. Workshops and webinars were also incorporated to encourage use. In the pilot study, through a combination of curated playlists and regular notifications, the activation rate for participants was around 54 percent for this particular online learning platform, the typical activation rate for the platform in other settings is 12 percent.

To assess whether this difference represented a meaningful change in online learning use, and to target the role of informational nudges in inducing this use, a randomized controlled trial (RCT) was implemented. In this RCT, individuals were randomized to be in either a control group, receiving only a free licence, or a treatment group, where they received a free license and regular information notifications to encourage use, specifically for the curated curriculum content. The organization distributed 1,900 individual licences total from October 2019 to September 2020. These licences were procured in September 2019. After a 6-month period, if an individual had not activated a licence it was reassigned to a new individual.

Recruitment and Randomization

There were three main recruitment mechanisms for this trial. First, youth were recruited through employment service providers (ESPs) from across the Toronto area (Study A). Secondly, youth participants were recruited through an all-day youth conference in the Toronto area (Study B).

Finally, youth were recruited through an intensive IT skills training program (Study C). We discuss the first two recruitment arms first.

For youth recruited via a youth conference or an ESP, a simple two-arm treatment and control group study was employed. Half of the participants within each recruitment cohort were assigned to a control arm where they had access to the full online learning platform, including curated curriculums for skills development and playlists of videos. The remaining youth were assigned to a treatment arm. Treatment arm participants received access to the online learning platform but additionally received bi-weekly email notifications about the content available and new learning pathways. Treatment arm participants were also invited to participate in a workshop on online learning. Initial licence distribution occurred at the youth conference. This youth conference is a flagship event of our partnered non-profit that provides targeted workshops on skills development for youth nominated by their youth support workers. Working with conference delegates, the organization distributed online learning licences to 90 youth. Staff from the online learning platform were on hand to help youth activate their licence keys and set up their accounts.[1]

Treatment randomization was based on initial sign-up for the youth conference. The organization relied on spot randomization to assign participants to be in the treatment or control arms, aiming for an equal distribution of treatment and control participants within the youth conference participants. Ninety youth registered for the conference, with 45 assigned to treatment and 45 to control. Given the sensitivity of collecting demographic information from a population of marginalized youth, no pre-study stratification was performed. The treatment and control groups

[1] There were approximately 250 youth participants at the conference. One hundred and fifty-seven of these were affiliated with the group IT training group, which comprise the pool from which Study B draws. We thus exclude these 157 from the Study B sample.

may have therefore been unbalanced across demographic factors such as age, gender or race. We list this as a limitation of the study.

The second recruitment strategy relied on licence distribution through local ESPs. We recruited 647 youth in total through this method. As with the conference recruitment, participants were spot randomized to either the treatment or control group as they were added to the study. This entailed equal assignment to treatment and control for those recruited via an ESP. For individual cases we randomly assigned on an individual basis. A project staff member was available to aid youth in setting up their LinkedIn profile and online learning account to ensure activation. Overall, 327 youth were randomly assigned to the treatment group, a treatment assignment rate of 50.5 percent.

In our third and final recruitment strategy, we worked with an information technology (IT) job skills training program. This job skills development program runs three to four cohorts at different locations, 4 times per year. There was a total of 11 cohorts over the course of the study. The smallest cohort had 33 members and the largest had 106 participants. The total recruited cohort across all cohorts was 736. Three cohorts were present at the conference to receive the first round of licences and activations. The remaining cohorts received license keys and activation assistance through their training program workers. The control group received licence access and regular support for their training program, including workshops and other engagement. The treatment group received push notifications to encourage use of the online platform.

As these job skills training programs have a high degree of interaction, spillover effects between peers within a cohort are plausible. Bias in peer effects may overstate estimates (Evans, Oates and Schwab, 1992). To build peer effects into the design, we used a two-step randomization process in line with Crépon et al (2013). Here, French youth who had recently graduated from college or university but who had yet to be able to find employment were randomly assigned to employment service providers. As spillover effects (on employment outcomes) were likely within the employment service areas, the authors first randomly assigned a proportion P for treatment assignment in an area. We replicate the assigned values of P in this study as 0, with no one assigned to treatment, 25, 50, 75 or 100 percent assigned to treatment.

In the second step, individuals within a given cohort were then assigned individually to either treatment or control for a given proportion P of eligible youth in the cohort. This design allows for testing of differences in the pooled results between simple treatment and control groups but can further allow for testing of spillover effects for those not assigned to treatment. Likely spillover effects would be in the form of positive peer spillover effects. If workshops and nudging successfully engage treated individuals more, control group peers may also engage in these activities more. Each P value is assigned twice, outside of P = 0.75, which was randomly assigned three times. No location was assigned the same P value twice, allowing for isolation of separate impacts of P and location-specific effects.

Sample Size and Minimal Detectable Effects

For the estimation of user engagement on the platform, we present the minimal detectable effect sizes across the final sample size numbers in each study (Appendix Table A.1). We had sufficient power in the sub-samples to detect one to four hours difference in total use, or 5 to 20 individual logs, depending on the study. The MDE sizes are each calculated based on the true sample size numbers and the realized mean and standard deviation of the control group samples in each study.

Data Overview and Descriptive Differences

To assess use of Linkedin Learning on employment outcomes we rely on user engagement with the online platform compiled from the partner organization via the online learning tool user portal. These user engagement reports contain information on specific courses viewed, total hours viewed, and number of completed courses by month. For Study C, we additionally have information on employment and education outcomes at the three- and six-month mark, aggregated to location and treatment assignment, as provided by the training organization. We review simple differences in outcomes by treatment and control group first. Given the randomization design, these differences should provide valid causal treatment effects.

Figure 1: Use since month of licence procurement, by treatment status

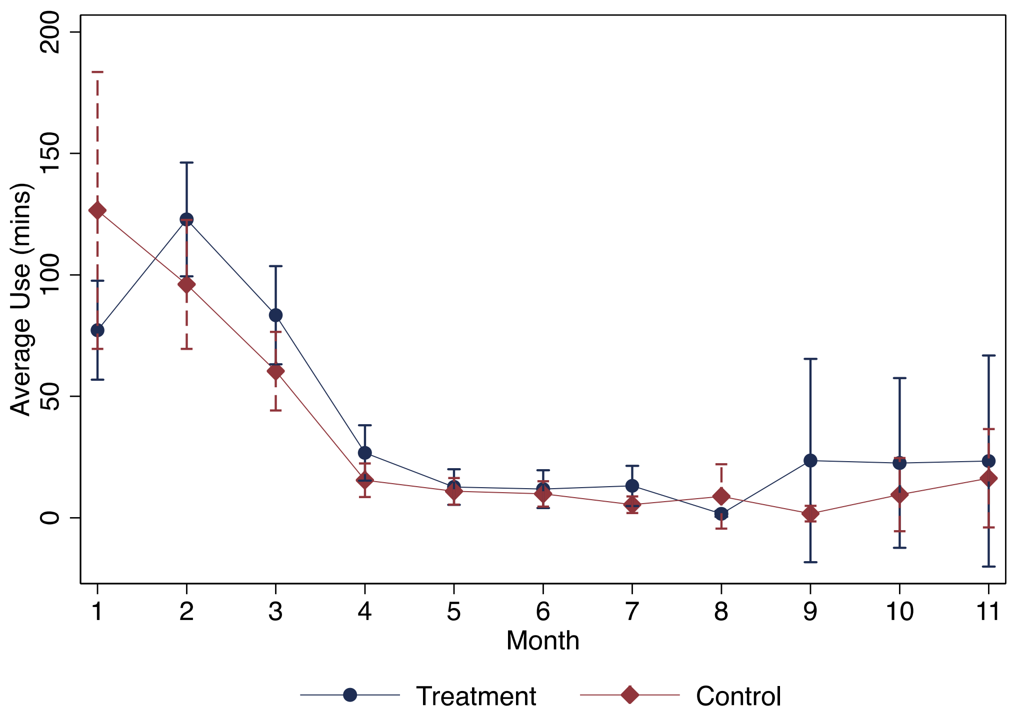

In Figure 1, we pool together all individuals across all three studies by treatment status. There are 747 individuals in treatment groups and 725 in control groups. Pooled together, we see high rates of average use in the first few months of use, but no discernable differences in use between treatment and control groups. Overall use drops precipitously after month three, likely corresponding to the end point for the IT training component of Study C. We see some increased use in use in months 9 to 11. As this study ran over one year and licences were given out throughout the year, only individuals from Study B have full data available to eleven months. These months therefore correspond to the months June to August 2020 for this group.

Examining individual studies, we start with the least prescriptive study (Study A) for participants who received online learning licences from general employment service providers (ESP). In Panel A of Table 1, participants who received licences from a general ESP experience a high rate of activation when compared to typical activation – estimated to be around 12 percent. We do not see evidence of a difference in activation for the treatment group versus the control group. There is some evidence of higher use for the treatment group in the first month, by 9.10 (19.96) additional minutes. This number is smaller than the minimal detectable effect of one hour or 60 minutes; thus, this difference is not statistically significant. We find no statistically significant differences in use across any monthly usage rate or the total usage in logs or hours. The lack of significance in these findings may be due to power issues or the lack of difference in program usage between the treatment and control groups. Looking only at activators in Appendix Table A.1, there are some differences in days to activation for those receiving informational nudges. There are also differences in terms of total hours. Treatment group participants who activated used the online learning platform for 3.58 (1.97) additional hours. This mostly stemmed from higher use in the first three months after activation.

In Study B, participants were given more intensive onboarding to the online platform at the time of licence activation via a two-day conference. With this additional onboarding support, we see even higher rates of activation than the general ESP group, up from 56 percent to 66 percent.

Panel B demonstrates evidence of an 8-percentage point increase in activation for the treatment group, but this is not statistically significant given the small sample size. There is also some evidence of higher use for the treatment group, especially in the first month, with 13.11 (10.02) additional minutes watched. However, we find no statistically significant differences, likely due to power issues relating to the sample size but also minimal differences in program usage between the treatment and control groups. Looking only at activators, there are no statistically significant differences in days to activation, use by month, or overall (Appendix Table A.1, Panel B).

Table 1: Treatment Differences in Online Learning Engagement

	Overall		Treatment		Control		Treatment Effect
	Mean	S.D.	Mean	S.D.	Mean	S.D.	(Std. Err.)
	Panel A: General ESP (Study A)						
Activation	0.56	0.50	0.56	0.50	0.56	0.50	0.001 (0.039)
M1 Use	59.50	251.29	64.03	269.65	54.93	231.61	9.102 (19.96)
M2 Use	11.09	65.76	13.09	80.33	9.08	46.79	4.017 (5.23)
M3 Use	3.33	25.82	3.31	28.86	3.34	22.41	-0.026 (2.05)
Total Use (hrs)	1.40	4.72	1.53	5.15	1.27	4.23	0.266 (0.37)
N	635		319		316		
	Panel B: Specialized Onboard (Study B)						
Activation	0.66	0.48	0.7	0.46	0.62	0.49	0.08 (0.10)
M1 Use	12.69	47.44	19.32	64.48	6.21	18.74	13.11 (10.02)
M2 Use	16.32	70.59	16.27	81.98	16.37	58.29	-0.10 (15.05)
M3 Use	8.96	57.22	5.33	16.43	12.5	79.11	-7.17 (12.18)
Total Use (hrs)	0.89	2.53	0.99	2.11	0.79	2.91	0.20 (0.54)
N	89		44		45		
	Panel C: IT Training (Study C)						
Activation	0.95	0.22	0.96	0.19	0.94	0.24	0.02 (0.02)
M1 Use	149.93	791.64	96.56	311.74	205.67	1084.18	-109.11* (58.28)
M2 Use	207.36	464.19	230.27	426.15	183.43	500.32	46.85 (34.21)
M3 Use	139.87	344.25	161.8	378.89	116.96	302.74	44.84* (25.35)

Total Use (hrs)	9.51	21.03	9.73	19.5	9.28	22.55	0.45 (1.55)
N	736		376		360		
Standard errors in parentheses: * p<0.05, ** p<0.01, *** p<0.001.							

In the final study (Study C), online learning was embedded with an IT training program. We see extremely high engagement overall, with 95 percent of individuals activating. There is a 2.4 percentage point difference in activation between treatment and control groups. Looking at use, we see that the control group had a much larger rate of use in the first month, with 109 additional minutes or 1.8 hours. In subsequent months, those assigned to treatment start to engage more with the content and have a sustained increase in use over the control group for almost all subsequent months. Overall, this use totals to 8.6 additional logs and 0.45 additional hours. Nevertheless, based on the sample size, we are not powered to detect an effect size. We next turn to a formal regression estimation strategy for studies B and C.

Empirical Strategy

Main Specification

To estimate precise effects of treatment group assignment on online learning there are additional confounding factors for which we may want to control. While we have limited information available on individual level characteristics of the NEET youth in this study, there are a number of factors related to timing of activation, location, and the employment service providers which may help with precision of estimates. We can use a formal regression estimation for Studies A and C. For Study B, all participants activated at the same time and no other information is available. Our baseline estimating equation across all these two studies takes the form:

$$Y_i = \alpha_1 + \beta_1 Z_i + X_i \partial_1 + \varepsilon_{ic} \qquad (1)$$

Where Y_{ic} represent online learning use and engagement for individual i. Treatment status is measured through Z_i and the estimated treatment effect is estimated through β_1. Any available controls are included in X_i. We run separate equations for each study. For Study A we notably include fixed effects for the employment service provider with whom individual i is affiliated. For Study C, we include cycle and location fixed effects to control for the four separate locations for the IT training as well as an indicator for the precise cycle a cohort started in over the 11 job training cohorts running throughout the year.

Peer Effects

Specific to the intensive job skills training program in Study C, we wish to measure the peer effects of treatment assignment. Given that this study consists of cohorts of students who are meeting in intensive learning environments with a high degree of interaction, it is likely that there might be spillover effects between treatment and control groups. This might be in the form of either treatment group participants telling control group participants about their additional informational treatments, making control group participants engage more. There may also be outcome spillover effects if treatment group participants learn specific skills that are passed on to peers, or if improved employment outcomes for treatment group peers create network effects leading to better employment outcomes for their control group peers. The potential for negative externalities is also present if online learning in the treatment group is robust enough to generate better employment outcomes that crowd out the employment outcomes of untreated individuals. All of these effects could serve to understate the treatment and control group differences (Miguel and Kremer, 2004). This is commonly referred to as contamination bias (Heckman, Lalonde and Smith, 1999). To measure these externalities, we estimate the following equation:

$$Y_{ic} = \alpha_2 + \beta_{25} Z_{ic}P_{25c} + \beta_{50} Z_{ic}P_{50c} + \beta_{75} Z_{ic}P_{75c} + \beta_{100} Z_{ic}P_{100c} +$$
$$\delta_{25} P_{25c} + \delta_{50} P_{50c} + \delta_{75} P_{75c} + X_{ic}\partial_2 + \omega_{ic} \quad (2)$$

Here assignment to treatment is measured through Z_{ic} and indicators for proportion of treated individuals within a cohort are assigned via P_{xc}. Thus β_{25} is the treatment effect for an individual assigned to treatment in a cohort with 25 percent of individuals treated. The spillover effect of being in the control group within a cohort with at least some but not all individuals treated is given through δ_x. To test the relationship of peer spillover effects we are interested in seeing if the δ_x coefficients are either jointly zero or either increasing or decreasing in P. If they are jointly zero, then there are no peer effects. If they are positive and increasing in P there are strong positive peer effects on engagement. Any negative δ_x would signal either disengagement when looking at use outcomes. Negative effects on employment would single crowding out effects. Control variables include the number of participants in the cohort and the number of treated participants in each cohort, as well as location fixed effects, number of months of data, and whether activation occurred in 2019 or 2020.

To generate more power to detect both treatment and externality effects we can also pool across all treatment assignment proportion groups:

$$Y_{ic} = \alpha_3 + \beta_3 Z_{ic}P_c + \delta_3 P_c + X_{ic}\partial_3 + \epsilon_{ic} \quad (3)$$

Here P_c is any non-zero treatment assignment cohort. The coefficient δ_3 is thus the effect of being untreated in a treated cohort and β_3 is the effect of being treated in a treated cohort. All other controls are analogous to equation (2).

Results

Baseline Regressions

Results estimating treatment effects for Study A and C via regression analysis are found in Table 2. For Study A, we estimate similar treatment effects as measures in Table 2, with 15.97 (12.36) additional minute of use with informational nudging compared to 0.27 hours (Panel A). Once controlling for timing of activation and incorporating fixed effects for the general ESP with which a participant was working, we find 17.22 (14.88) additional total minutes of use but are still not powered to detect effect sizes. We interestingly detect fewer logins to the online learning tool for those in treatment, with 2.70 (1.30) fewer logs. These results suggest that informational nudging increases concentrated use, with less logs but more time per log. For Study C, using simple treatment and control group differences, we find only 0.45 additional hours of use with treatment. Our baseline estimate in column (1) are largely equivalent to this number, with an estimated 27.01 (213.01) additional minutes of use with treatment. However, we find that when accounting for location and cycle timing fixed effects, the treatment effect turns negative. There were specific locations that appeared particularly good at encouraging use within the embedded IT training program. Further incorporating the proportion of individuals receiving treatment in a given cohort, we estimate that those in treatment has 241.65 (219.48) fewer minutes of use and 5.78 (11.25) fewer logins. Nevertheless, there were likely peer spillover effects in Study C that have yet to be accounted for. We turn to this next.[2]

[2] Note that since we have no other covariates for the Study B, regression and simple differences are the same and thus we do not present these results. Combining Studies A and B into a pooled design also does not give enough power to detect effect sizes.

Table 2: Baseline Regressions Study B and C

Panel A: Study A

	Total Minutes Used			Total Logs		
	(1)	(2)	(3)	(1)	(2)	(3)
Treatment Effect	15.9667	14.9012	17.2243	-1.9754	-2.0485	-2.6967*
	(15.3644)	(14.0350)	(14.8762)	(1.2277)	(1.1969)	(1.3004)
Timing Covariates		x	x		x	x
ESP Fixed Effects			x			x
N	628	628	628	628	628	628

Panel B: Study C

	Total Minutes Used			Total Logs		
	(1)	(2)	(3)	(1)	(2)	(3)
Treatment Effect	27.0089	-154.7720	-241.6547	8.5552	-7.4891	-5.7777
	(213.0160)	(181.6166)	(219.4772)	(15.8235)	(9.0798)	(11.2523)
Location FE		x	x		x	x
Cycle FE		x	x		x	x
Proportion Treated			x			x
N	736	736	736	736	736	736

Note: Standard errors in parentheses: * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Standard errors clustered at cohort*treatment level.

Peer Spillover Effects

In Table 3, we present regression estimates of peer spillover effects within Study C. To review, all cohorts were randomized to have either 0, 25, 50, 75 or 100 percent of cohort members receiving treatment nudges. Cohorts without any participants assigned to treatment can be thought of as a “super control group” baseline comparison group against which we will compare all treatment effects. For those in treatment, we see some variation in engagement by proportion treated. Engagement based on total use (measured in minutes or hours) increased for all treated individuals, except for those in the 50 percent randomization group. Online logins appeared to decrease across all treatment proportions, as did use in the first month.

More compelling is that these results suggest large spillover effects for use to the control groups with some but not all cohort peers in treatment. We were interested in whether the estimated coefficients for spillover effects by proportion in treatment were jointly zero and reject this null hypothesis. Use appeared to mostly increase as more members of a cohort were in treatment. We find this for total minutes and hours of use in the first month and overall. There is a less clear relationship between logins. Nevertheless, number of logins do not necessarily quantify meaningful use, as in the other studies we also see more use but less logins.

In a final specification to test peer spillover effects, we pool across all cohorts and assess the impact of treatment in a cohort with any treated individuals (Table 4). We find lower estimated engagement from treatment assignment when treatment is defined in this way. This engagement includes 5.3 (2.3) fewer logs in the first month and 295.5 (209.5) fewer minutes (nearly 5 hours) use overall. However, we also observe peer spillover effects in this pooling strategy, with 391.6 (129.6) additional minutes of use in month one for those in a control group within a treated cohort. These results seem to suggest that when a value is assigned to a commodity, here in the form of additional informational nudges indicating its value, people use it more. In one scenario, control group members may hear from other members that the videos on which information nudges were provided had some material value in job training and securing employment. Another scenario may involve control group members feeling as if they were accidentally not given pertinent information for the course and watching the videos independently.

Table 3: Spillover Effects in Online Learning Usage						
	Activation	Total Use (mins)	Total Logs	Total Use (hrs)	Logs in Month 1	Use in Month One (mins)
Mean (S.D.)	0.95 (0.22)	570.52 (1262)	53.77 (77.89)	9.51 (21.03)	10.19 (17.72)	149.93 (791.64)
	(1)	(2)	(3)	(4)	(5)	(6)
p(25)*Treated	-0.042	65.043	-10.747	1.084	-7.599	-47.721
	(0.039)	(316.467)	(43.258)	(5.274)	(7.377)	(63.534)
p(50)*Treated	-0.022	-712.069**	-22.659	-11.868**	-6.557*	-579.750+
	(0.013)	(193.603)	(13.039)	(3.227)	(2.333)	(317.058)
p(75)*Treated	0.062	41.942	9.970	0.699	-0.249	41.273
	(0.061)	(205.107)	(13.781)	(3.418)	(2.118)	(30.267)
p(100)*Treated	-0.348***	1490.550***	-92.311***	24.842***	-23.818***	1866.197***
	(0.033)	(234.158)	(18.513)	(3.903)	(4.148)	(50.666)
Control p(25)	-0.082***	1009.050***	51.060***	16.818***	9.722***	697.168***
	(0.011)	(78.714)	(10.977)	(1.312)	(1.877)	(16.481)
Control p(50)	-0.191***	985.291***	-4.756	16.422***	-4.979**	1191.532***
	(0.010)	(106.272)	(7.700)	(1.771)	(1.404)	(161.305)
Control p(75)	-0.343***	1777.405***	-20.588	29.623***	-6.825+	1801.488***
	(0.041)	(277.480)	(12.263)	(4.625)	(3.259)	(37.528)
N Cohort	-0.001***	-6.371*	1.313***	-0.106*	0.117*	-14.538***
	(0.000)	(2.531)	(0.192)	(0.042)	(0.042)	(0.501)
N Treated	0.007***	-31.781***	2.074**	-0.530***	0.532***	-41.390***
	(0.001)	(5.801)	(0.472)	(0.097)	(0.103)	(1.252)
Location FE	x	x	x	x	x	x
N	733	733	733	733	733	733

Note: Table presents regression results for Study C, accounting for potential peer spillover effects. Standard errors in parentheses are clustered at the training cohort level. Additional controls include the number of months in observations and year of activation. * p<0.05, ** p<0.01, *** p<0.001

Table 4: Spillover Effects in Online Learning Usage (Pooled Specification)						
	Activation	Total Use (mins)	Total Logs	Total Use (hrs)	Logs in Month 1	Use in Month One (mins)
	(1)	(2)	(3)	(4)	(5)	(6)
p(Any)*Treated	0.006	-295.541	-12.244	-4.926	-5.279*	-218.128
	(0.033)	(209.469)	(12.619)	(3.491)	(2.336)	(193.241)
p(Any)	-0.110*	252.671	-5.378	4.211	-3.952	391.634*
	(0.038)	(246.018)	(30.866)	(4.100)	(6.990)	(129.582)
N Cohort	-0.004***	2.923	0.101	0.049	-0.132	-0.206
	(0.001)	(5.767)	(0.707)	(0.096)	(0.157)	(2.602)
N Treated	0.001	4.868	0.275	0.081	0.183+	-1.119
	(0.001)	(4.776)	(0.438)	(0.080)	(0.086)	(2.800)
Location FE	x	x	x	x	x	x
N	733	733	733	733	733	733

Note: Table presents regression results for Study C, accounting for potential peer spillover effects. Standard errors in parentheses are clustered at the training cohort level. Additional controls include the number of months in observations and year of activation. * p<0.05, ** p<0.01, *** p<0.001

Graduation and Employment Outcomes

Specific to Study C, we have available aggregate graduation and job outcomes three months after the training period, by cohort and treatment status for all cohorts in the first two cycles. There are seven cohorts with this graduation and employment information available. In Table 5, we show mean rates of graduation from the IT training program and mean employment rate following the training program at three months. Those in the control group have an average graduation rate of 88 percent, while those in the treatment group graduated at a rate of 89 percent. We find a one percentage point difference in the mean employment rate following training. However, those in the control group have a slightly higher employment rate, at 47 percent versus 46 percent in the treatment group. Nevertheless, this time period did correspond to the beginning months of the COVID-19 pandemic and subsequent shutdown in the GTA. Therefore, these employment rates are much lower than usual and difference in employment between treatment and control are not statistically different.

Table 5: Employment Outcomes for Study C by Treatment Group					
	Control		Treatment		Difference
	Mean	s.d.	Mean	s.d.	
Graduated	0.88	0.01	0.89	0.10	0.01 (0.01)
Employed	0.47	0.04	0.46	0.02	0.01 (0.01)
Correlation with P	-0.39		-0.15		
Correlation with T			0.45		
N	67		90		157

As we only have aggregate, cohort and treatment level information on graduation and employment outcomes, we cannot estimate any meaningful treatment effects at the individual level. We can, however, assess how the treatment and control group outcomes correlate to the proportion of individuals treated and with each other. Both treatment and control group individuals appear to negatively correlate with the proportion of individuals treated but positively correlated with each other. As location appears to matter significantly for these outcomes, not much can meaningfully be said about these outcomes when we are unable to control for the specific location factors contributing to better or worse graduation and employment outcomes.

Changes in Use with COVID-19

In one final analysis, we assess how the use of the online learning tool changed with the onset of COVID-19 and the ensuing social distancing policies in Canada (Figure 2). Looking at weekly users by treatment group status, we can see a sharp increase in use after the new year and again around Ontario's social distancing orders (this was in the second week of March 2020, referenced via the dotted line). We see that more treatment group members were using the platform during this first wave. This increase in the number of weekly users remained through April but began to fall off substantially in May. After a sharp dip in June, use rose again. Notably, this later use appears slightly greater for control group members, potentially indicating that individual motivation spurs more use than informational nudges.

In terms of hours viewed overall, we again see a substantial increase in use directly following the lockdown (Panel B). However, the average number of hours viewed falls with time and does not appear to recover, even though the number of users rose over the same period. This effect potentially suggests that people were signing up for the online learning tool in later summer months but not actively engaging with the platform. In the final Panel (C), we measure the average number of hours per unique user on the site. We see that use increased with the lockdown by this measure but did not drop off after the initial first wave. Together, these figures suggest that use of the online learning tool increased in the pandemic period, but only for a select group of users who maintained consistent use. Informational nudging appeared to help with activation but did little to motivate actual engagement over this period.

As a final step, we estimate changes in use due to the pandemic at the three-month post-lockdown and six weeks post-lockdown (Table 6). We frame this in a traditional differences-in-differences estimation to understand the impact of being in the post-lockdown period as a treatment group participant. First, we see that being in the post-shutdown period increased the number of unique users, total hours, average hours, and the number of custom courses accessed regardless of treatment status (Panel A). There are 15.91 (6.05) additional active users and 26.38 (10.77) additional hours of use in the post-period for those in the treatment group and receiving informational nudges. However, we see no difference between treatment and control groups in terms of average hours of use or custom content usage. Looking to cumulative data six months after the first COVID-19 social distancing orders were put in place, we see a maintained increase in use overall in the post-period, though

Figure 3: Changes in Online Learning Use Around COVID-19

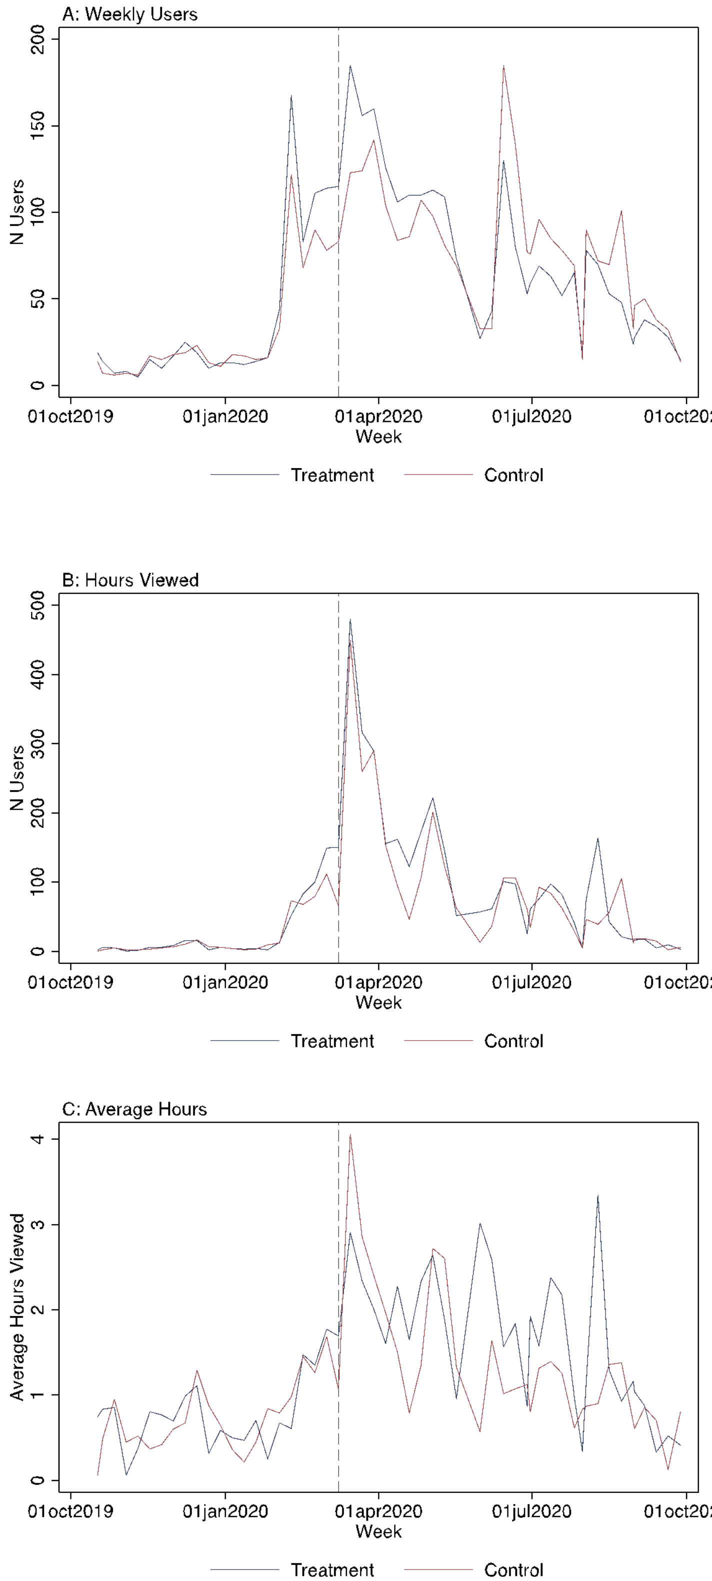

compared to the first three months, the number of unique users, total hours and average hours were all lower (Panel B). We also see any estimated increased use over this longer time frame due to informational nudging has moved towards zero.

Table 6: Change in Weekly Use Following COVID-19 Closures

	Panel A: Use Post-3 Months			
	(1)	(2)	(3)	(4)
	Users	Total Hours	Avg. Hours	Custom Content
Treatment	7.091*	6.401	0.054	1.591
	(3.117)	(4.438)	(0.075)	(1.586)
Post	70.164***	156.350***	1.410***	3.655*
	(10.125)	(39.665)	(0.314)	(1.785)
*Treatment*Post*	15.909*	26.384*	-0.152	-2.791
	(6.050)	(10.769)	(0.239)	(2.173)
N	64	64	64	64
R2	0.511	0.538	0.562	0.068

	Panel B: Use Post-6 months			
	(1)	(2)	(3)	(4)
	Users	Total Hours	Avg. Hours	Custom Content
Treatment	7.091*	6.401	0.054	1.591
	(3.067)	(4.366)	(0.074)	(1.561)
Post SD	47.396***	65.821**	0.591**	5.406
	(10.039)	(19.046)	(0.175)	(4.901)
*Treatment*Post*	-11.188+	8.280	0.248	-5.526
	(5.638)	(8.079)	(0.167)	(3.964)
N	106	106	106	106
R2	0.206	0.156	0.229	0.206

Note: Standard errors in parentheses and clustered at weekly level: * p<0.05, ** p<0.01, *** p<0.001 Post measured at and after the second week of March.

Conclusion

In this paper, we present the results of a randomized controlled trial to assess the use of informational nudging as a means to improve engagement with online learning platforms. We find that informational nudging alone was not enough to spur meaningful use of the platform. We observe that informational nudging helped to improve activations, especially during the COVID-19 pandemic. Yet, actual content watched did not appear to increase beyond the control group's use. With this we note power issues driven by sample sizes and minimal difference in activations between treatment and control groups as limitations of this paper.

Nevertheless, the usage rate from this study was sustainably higher than previously measured engagement with this specific online learning platform, with activation rates ranging from 50 to nearly 100 percent (compared to a previous average of 12 percent). Across the three separate RCTs employed in this paper, usage with the platform increased as additional supports were set up alongside free provision of the online learning licence. Our baseline study utilized employment service providers as touchpoints for participants to obtain licences. We also incorporated voluntary workshops to help participants sign-up and use the online learning tool. Furthermore, as participants had already sought out ESPs to seek employment opportunities, their motivation to engage in job skills training is likely higher than average. Together these factors likely helped to increase use. When we added an in-person conference led by representatives of the online learning tool, we saw an even further increase in activation and use. These results suggest that hands-on training and some combination of online and in-person use helps to improve engagement.

In a final level of online learning support, when online learning was embedded within an intensive IT job skills training program, we saw the highest activation rates and higher rates of continued use post-training. Again, these results suggest that providing free online learning tools on their own will not have a meaningful impact on engagement. Likely, coaching on online learning interface and the types of content available will help students to use online learning tools more.

Peer effects also appear to drive use with online learning. Use for control group members of a given cohort in the IT training program--those who did not receive additional informational nudging-- increased as more of their peers received nudges. Peer effects likely operate through word or month. This word of mouth may place a perceived additional value on the nudged content, leading to higher use for treatment and control groups. There may also be a tangible value driven by informational nudging if treatment group individuals had improved job prospects after watching the nudged content. In either scenario, nudging may increase participants' motivation to engage with content and maybe one reason why simple comparisons of engagement between treatment and control groups show minimal differences. Testing these specific channels remains an area for future research.

References

Alpert, W. T., Couch, K. A., & Harmon, O. R. (2016). A randomized assessment of online learning. American Economic Review, 106(5), 378-82.

Bowen, W. G., Chingos, M. M., Lack, K. A., & Nygren, T. I. (2014). Interactive learning online at public universities: Evidence from a six-campus randomized trial. Journal of Policy Analysis and Management, 33(1), 94-111.

Brotman, S. N. (2016, October 27). How digital readiness affects job retraining for labor market growth. Brookings Institute Blog. Retrieved from https://www.brookings.edu/blog/techtank/2016/10/27/how- digital-readiness-affects-job-retraining-for-labor-market-growth/

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., & Zamora, P. (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. The quarterly journal of economics, 128(2), 531-580.

Duflo, E., & Saez, E. (2002). Participation and investment decisions in a retirement plan: The influence of colleagues' choices. Journal of public Economics, 85(1), 121-148.

Evans, W. N., Oates, W. E., & Schwab, R. M. (1992). Measuring peer group effects: A study of teenage behavior. Journal of Political Economy, 100(5), 966-991.

Figlio, D., Rush, M., & Yin, L. (2013). Is it live or is it internet? Experimental estimates of the effects of online instruction on student learning. Journal of Labor Economics, 31(4), 763-784.

Heckman, J. J., LaLonde, R. J., & Smith, J. A. (1999). The economics and econometrics of active labor market programs. In Handbook of labor economics (Vol. 3, pp. 1865-2097). Elsevier.

Joyce, T., Crockett, S., Jaeger, D. A., Altindag, O., & O'Connell, S. D. (2015). Does classroom time matter?. Economics of Education Review, 46, 64-77.

Miguel, E., & Kremer, M. (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. Econometrica, 72(1), 159-217.

Appendix

Table A.1: Estimated Sample Size and Minimal Detectable Effect

	Estimated Sample Size			MDE	
Recruitment Mechanism	Control: Licence Only	Treatment: Notifications	Total	Total Use (hrs)	Total Logs
Study A: General ESPs	299	327	646	1.0	5.96
Study B: Guided Onboarding	45	45	90	1.5	20.3
Study C: Embedded Learning	360	376	736	4.3	16.1
Total	704	748	1,472		

Note: MDE=Minimal Detectable Effect. Assumes alpha = 0.05 and power = 0.80. Minimal detectable effect calculated based on control group means for each study.

Table A2: Treatment Effect Estimates, Activators Only							
	Panel A: Study A						
Days to Activation	7.31	18.23	9.11	23.51	5.53	10.54	-3.581+ (1.973)
M1 Use	105.54	327.52	113.48	351.48	97.51	302.15	-15.963 (34.658)
M2 Use	19.64	86.60	23.13	105.80	16.11	61.50	-7.019 (9.159)
M3 Use	5.89	34.17	5.85	38.21	5.93	29.63	0.075 (3.617)
M4 Use	10.10	67.01	16.08	91.63	4.12	23.35	-11.969 (8.199)
M5 Use	13.70	64.99	11.70	44.30	15.65	80.47	3.956 (9.712)
M6 Use	8.83	50.01	7.43	43.52	10.15	55.73	2.727 (8.364)
M7 Use	0.19	1.07	0.04	0.23	0.34	1.49	0.295 (0.209)
M8 Use	0.05	0.18	0.02	0.06	0.07	0.24	0.055 (0.056)
Total Logs	19.77	35.08	18.08	30.51	21.48	39.17	3.400 (3.709)
Total Use (hrs)	2.48	6.07	2.72	6.63	2.25	5.45	-0.468 (0.642)
N	358		180		178		358
	Panel B: Activators Only						
M1 Use	24.62	85.75	23.10	97.32	26.31	72.57	-3.21 (22.55)
M2 Use	13.51	70.04	7.57	19.22	20.09	100.21	-12.52 (18.34)
M3 Use	13.09	40.57	16.60	39.51	9.20	42.08	7.4 (10.62)
M4 Use	1.27	6.84	0.37	1.59	2.26	9.78	-1.88 (1.78)
M5 Use	8.66	29.01	9.00	30.27	8.29	28.09	0.71 (7.63)
M6 Use	10.00	40.37	14.06	54.47	5.51	12.80	8.55 (10.56)
M7 Use	19.95	32.34	17.61	21.51	22.54	41.47	-4.92 (8.48)
Total Logs	1.34	3.02	1.40	2.41	1.27	3.62	0.13 (0.79)
Total Use (hrs)	2.65	8.10	2.67	9.96	2.63	5.52	0.04 (2.17)
N	89		44		45		
Standard errors in parentheses: * p<0.05, ** p<0.01, *** p<0.001							