ACLA WORKING PAPER SERIES

ABOLISHING POLICY-INDUCED DROPOUT IN DUTCH HIGHER EDUCATION DUE TO COVID-19. EXPLORING TRENDS IN FUTURE ACADEMIC PERFORMANCE

MELVIN VOOREN
ILJA CORNELISZ
MARTIJN MEETER
CHRIS VAN KLAVEREN

Working Paper 20222

http://www.acla.amsterdam/workingpapers-wp20222

AMSTERDAM CENTER FOR LEARNING ANALYTICS
van der Boechorststraat 7
1081BT, Amsterdam

Abolishing Policy-Induced Dropout in Dutch Higher Education due to Covid-19. Exploring Trends in Future Academic Performance

Melvin Vooren[1] Ilja Cornelisz[2] Martijn Meeter[3] Chris van Klaveren[4]

25th October 2022

Abstract

This study uses the unique opportunity of the Covid-19 pandemic to examine the future academic performance of students who would otherwise be academically dismissed. We utilize a Machine Learning model based on solely pre-pandemic information to account for potential Covid-19 effects on academic performance. We show that students continue studying at the same pace in year two, regardless of receiving a waiver or being dismissed. Extrapolating the results, we show that students who would normally have been dismissed will continue progressing at the same pace towards graduation. We discuss the implications for the ongoing discussion on academic dismissal policies.

Keywords: academic dismissal; policy-induced dropout; machine learning

[1]Vrije Universiteit Amsterdam, Amsterdam Center for Learning Analytics, Faculty of Behavioral and Movement Sciences M.Vooren@vu.nl

[2]Vrije Universiteit Amsterdam, Amsterdam Center for Learning Analytics, Faculty of Behavioral and Movement Sciences I.Cornelisz@vu.nl

[3]Vrije Universiteit Amsterdam, Amsterdam Center for Learning Analytics, Faculty of Behavioral and Movement Sciences M.Meeter@vu.nl

[4]Vrije Universiteit Amsterdam, Amsterdam Center for Learning Analytics, Faculty of Behavioral and Movement Sciences C.P.B.J.van.Klaveren@vu.nl

1 Introduction

It is a well-established fact that too many students leave university programs before completion. Within the member states of the Organization for Economic Co-operation and Development (OECD), the average dropout rate in higher education amounts to 31 per cent. Countries like the United States and Italy have higher rates of around 50 percent, whereas Flanders and Denmark are at the lower end of the spectrum, at around 15 percent. The Netherlands sits around the OECD average (Source: Education at a glance: Educational finance indicators, 2015 = 100). There are several personal, institutional, and societal costs that are associated with college dropout. These include but are not limited to: forgone tuition fees, reduced earnings, reduced efficiency, and a decreased supply of skilled labor.

Because dropout is such a widespread and costly phenomenon, institutions have used many different programs to reduce dropout. According to the literature, most remediation programs are found to have no impact, or a negative impact, worsening dropout figures instead of improving them (Deelen and Kuijpers, 2016). Financial aid programs, aimed at providing extra funds to extend the study duration are found to have a positive impact when they are combined with other support programs (Angrist, Lang, and Oreopoulos (2009); Page and Scott-Clayton (2016)). However, it has been largely ignored that a sizable proportion of dropout can be policy-induced.

In countries where academic dismissal policies are enforced, students are required to leave the program in case of poor academic performance, or after a fixed number of exam failures. In the Netherlands, there is a nation-wide policy that students have to leave at the end of their freshman year after poor academic performance, generally defined as having accumulated less than 70 per cent of the credits in the first year, i.e. 42 out of 60 ECTS. As a result, the policy increases first-year dropout rates. The students who are forced to drop out as a result of academic dismissal are referred to as policy-induced dropouts (PIDs).

In an earlier study using data from the Netherlands, (Sneyers and Witte, 2017) have found that the policy increases graduation rates for the remaining students. However, other research by Cornelisz, van der Velden, de Wolf, and Van Klaveren (2020) have found that those students who are forced to leave the program due to academic dismissal actually switch to a similar program at the same university, or the same program at a different university, and that for students forced to switch because they are just below the threshold of dismissal, the likelihood of obtaining a degree is similar to that of students is similar to that just above the threshold for dismissal.

Considering this, the academic dismissal policies are quite costly from the student perspective because switching to a similar program elsewhere usually implies moving to a different town and starting all over again.

Ost, Pan, and Webber (2018) estimate the returns to college persistence for students around the academic dismissal threshold using data from 13 public universities in the state of Ohio. The issue with academic dismissal thresholds is that the minimum GPA or amount of credits is publicly known. Therefore, students and their professors can manipulate their GPA or achieved credits by exerting more effort or by being more lenient on the student in terms of grading. Hence, using the number of credits by the end of the year as a running variable in a regression discontinuity setup is problematic as it does not properly identify the treatment.

This study uses an alternative strategy to identify the causal effect of the academic dismissal policy on academic outcomes. In the first year of the Covid-19 pandemic, the Dutch government abolished first-year academic dismissal in the beginning of the second semester, on March 19, 2020. As a result of this, students who would otherwise have been academically dismissed were allowed to continue their studies in year 2. This circumstance provides a unique opportunity to monitor the performance of PIDs in the second year of the program.

However, the Covid-19 crisis also comes with a methodological challenge, because semester 2 performance could have been affected by other factors induced by the Covid-19 pandemic, such as grading leniency, on-line learning, or mental health effects. In order to address these methodological concerns, this paper has the following three consecutive research objectives. The first objective is determining the number of students who drop out due to academic dismissal, i.e. the PIDs. Secondly, we examine second-year performance of the students who we identify as PIDs and were allowed to continue as a result of Covid-19. Finally, we use these results to calculate the average study duration for PIDs until obtaining a bachelor's degree.

The main challenge in identifying PIDs in the Covid-19 cohort is that achieved credits at the end of the program cannot be compared with achieved credits in earlier cohorts, due to the aforementioned Covid-19 effects on student performance, including mental health effects. To address this, we use a machine learning model to predict expected credits at the end of year 1 based on only pre-covid-19 characteristics, i.e. performance in semester 1 of year 1 and individual background characteristics. Because the first semester was not affected by the Covid-19 crisis, we can compare the expected credits of students in the Covid-19 cohort with the expected credits of those in earlier cohorts. Doing so, we can identify the PIDs and observe their performance in year 2. Within a large research university in the Netherlands, we identify 22.4 per cent of all dropouts in year 1 as PIDs.

As a next step, we examine second-year performance of PIDs. In the earlier cohorts, we can see that there is a substantial share of students below the academic dismissal threshold who are allowed to continue in year 2, despite the fact that they should have been dismissed based on their performance. In other words, these students received a waiver for the academic dismissal policy (which may be gven out on request by program examination boards if one can successfully plead "special circumstances"). We then compare the second-year performance of the students below the academic dismissal threshold in those earlier cohorts (i.e., those with a waiver) with the students below the threshold in the Covid-19 cohorts. We see that the performance of PIDs is not significantly different from that of students who got a waiver. Because the second-year performance of students with a waiver is thus the same as the performance of the PIDs, we assume that this remains the same in later years as well.

2 Background

2.1 Programs aimed at reducing dropout

2.1.1 Reducing information asymmetries

The higher educational dropout literature describes three categories of programs targeted at reducing program dropout. The first category of programs is aimed at addressing information asymmetries or mismatch between student and educational program. The idea behind this is that by improving the information that students have about the programs they want to enroll into, they are able to make a more well-informed choice about their study choice. In the Netherlands, all prospective bachelor's students are forced to follow an 'academic matching' program as part of the enrollment process. Prospective students are required to participate in the matching program before they can enroll into any bachelor's program. The program consists of a short course that is related to the academic content of the chosen bachelor's program. The content varies by program and may include assignments and reading course literature. The program is concluded with a test. Based on the test results and a questionnaire, students receive a personalized study advice whether he or she has made the right choice. Eventually, the goal of these programs is to reduce dropout by preventing enrollment of less-capable and less-motivated students.

Deelen and Kuijpers (2016) analyze the effect of this particular policy on the probability that applicants withdraw their application before the start of the program in a randomized controlled trial. The applicants were randomly assigned to the academic matching program. The results suggest that the academic matching program did not reduce program enrollment in the bachelor's program compared to the control group. If the program is not effective at reducing enrollment, we can conclude that it is also not effective at reducing dropout. Van Klaveren, Kooiman, Cornelisz, and Meeter (2018) examine the effect of the policy on first-year enrollment and dropout rates. The authors find that providing students with personalized information on their future academic performance in the program increases enrollment rates by 25%, and does not reduce first-year dropout. In addition, they find that students are generally overly positive about their future performance, which may explain the increase in enrollment rates but absence of an effect on first-year dropout rates.

2.1.2 Remediation programs

The objective of remediation programs is to get skills of at-risk students to the required levels. Because programs in this category are more extensive than those in the first category, they are also costlier. Calcagno and Long (2008) use a regression discontinuity design to estimate the causal effect of math and reading remediation programs on the outcomes of college students in Florida, United States. The results point at lower first-year dropout rates. However, there does not seem to be an effect of remedial programs on the probability of completing courses or eventual graduation. Therefore, the authors conclude that remedial programs are not effective in promoting long-term study success.

Using data from an urban community college system in the United States, Scott-Clayton and Rodriguez (2015) estimate the effects of remediation programs on credits and degree completion. Although half of the students in the community college system attend remedial courses, the authors find that remedial education does not improve students' skills. To the contrary, the results point at a crowding-out effect: instead of taking college-level courses, the students take remedial courses. This explains the negative effect of remedial education on long-term study success. Valentine et al. (2011) conducted a meta-analysis on the effectiveness of remediation programs. The results of this meta-analysis point at small effects on preventing first-year dropout.

Based on this, the authors do not recommend adopting remediation programs as a general policy to reduce dropout and increase long-term graduation rates. Scott-Clayton, Crosta, and Belfield (2014) recommend using screening tools such as standardized tests to target remediation programs at those students who require remediation without sacrificing study success of those who do not need it.

2.1.3 Financial aid programs

Financial aid programs are aimed at reducing dropout consists of financial aid programs. The objective is to provide extension of study duration, reducing financial constrains to complete the program. Compared to the aforementioned programs, financial aid programs are the costliest. Page and Scott-Clayton (2016) review the research on the relation between socioeconomic gaps, financial aid, and academic achievement. They argue that while high-income students are able to rely on their parents, low-income students face financial constraints, and therefore rely on financial aid programs. However, financial aid on its own does not promote academic achievement among susceptible students – other support programs are required as well for them to have an effect. Using data from a randomized controlled trial at a satellite campus of a large Canadian university, Angrist, Lang, and Oreopoulos (2009) estimate the effects of offering financial incentives for good grades among college freshmen. The authors report a positive effect of financial incentives on student achievement in the first year, when combined with other academic support services. This combined treatment had an effect on the grades and academic achievement for women, but not for men.

2.2 Policy-induced dropouts

Taken together, out of all policies aimed at reducing program dropout, financial aid programs are the most effective. This might relate to the fact that access to higher education is not the same for each student (see Scott-Clayton and Sacerdote (2016) and Page and Scott-Clayton (2016)). High tuition fees and different tax credits for higher education expenses might explain differences in program dropout among different groups of students. Dropping out might be the most economically sound decision for students with a low socio-economic status because the costs of staying in the program simply outweigh the costs of dropping out. This might explain why programs do not effectively reduce student dropout.

Although many different policies have been administered to address program dropout, it has been largely ignored that a sizable proportion of dropout can be policy-induced. In countries such as the United States, academic probation programs prevent further course enrollment of students with a low grade point average (GPA). In other countries, too, there are academic dismissal policies (ADP) in place. In the Netherlands, students are required to leave the program in the case of poor academic performance. This ADP is enforced at the end of their freshman year, increasing first-year dropout rates. This dropout is policy-induced.

Within a difference-in-differences framework, Sneyers and Witte (2017) compare Dutch university programs that introduced an ADP to programs that did not. The authors confirm that the ADP increases first-year dropout rates, but leads to higher graduation rates for the remaining students. They also find that the ADP decreases overall student satisfaction. It is no surprise that an ADP that forces poor-performing students to drop out in the first year increases first-year dropout rates. The question that arises however is what happens to those students who are forced to drop out. Cornelisz, van der Velden, de Wolf, and Van Klaveren (2020) investigate what happens to the students who were academically dismissed from the Economics and Business bachelor's program at the University of Amsterdam in the Netherlands. By linking the data from this program to administrative records, the authors estimate the long-run academic outcomes of academic dismissal. The authors identify the causal effect of academic dismissal in a regression discontinuity framework by comparing students around the academic dismissal threshold. They find that academic dismissal does not alter the propensity of graduation, nor does it change study delay. They also find that 43.4% of the dismissed students enroll into the same program at a different university, whereas 41.9% of them enroll into a program within the same academic domain. Taken together, the conclusion from Cornelisz et al. (2020) is the proportion of graduation students around the academic dismissal cutoff if comparable, and that the students who were academically dismissed generally switched to a similar program or the same program elsewhere.

2.3 Covid-19

In addition to knowing what happens to students who were academically dismissed, it is also important to know how they would have performed in the second year at the same university if they were not dismissed. This study exploits the unexpected Covid-19 pandemic to monitor the performance of PIDs in the second year of the program. We compare the students in the 2019-2020 academic year, starting in September 2019, to the students in earlier cohorts. In February 2020, the Covid-19 crisis came to Europe and started affecting the students in the Covid-19 cohort. Campuses in the Netherlands closed on the 12th of March. A week later, on the 19th of March 2020, it was announced that the first-year academic dismissal program had been cancelled nationwide. Thus, the first semester of year one of the Covid-19 cohort was similar to that of the earlier cohorts, whereas things were different in the second semester. This makes identifying PIDs a methodological challenge, because performance in the second semester in the Covid-19 cohort could have been affected by for instance on-line learning, mental health effects, or an increase in fraud due to on-line proctoring and grading leniency.

This study has three main research objectives. The first objective is to determine the number of students who drop out due to the ADP, or in other words, identifying the PIDs. The second objective is examining the second-year performance of PIDs who continued in the second year due to the cancellation the ADP as a result of Covid-19. The third and final objective of this study is to calculate the total study duration for PIDs until obtaining a bachelor's degree, if they are not academically dismissed.

3 Data

3.1 Institutional background

We use administrative data from the Vrije Universiteit Amsterdam (VU Amsterdam), which is one of two large publicly-funded research universities in Amsterdam (the other is the University of Amsterdam). It is a broad universities, offering a wide range of programs in all domains. Table 1 displays the student numbers at VU Amsterdam for the cohorts 2015 until 2019, for the first year and all years, respectively. There is a consistent upward trend in the number of first-year students. From 2015 to 2019, the number of first-year students has increased by 66%.

(.................................... Table 1 about here ....)

Table 2 shows first-year dropout statistics at VU Amsterdam. For the years 2015 to 2018, we observe an average first-year dropout rate between 25 to 30 percent. These first-year dropout rates at VU Amsterdam are lower than the total average dropout in the Netherlands mentioned in the Introduction section, but it is worth noting here that table 2 shows dropout after the first year only, in contrast to total dropout. The table also shows a lower first-year dropout rate for the 2019 cohort compared to the 2015-2018 cohorts. This is due to the cancellation of the ADP in the Covid-19 cohort.

(.................................... Table 2 about here ....)

When we plot the proportion of students continuing in the second year against the number of credits accumulated in the first year with respect to the academic dismissal threshold for both the Covid-19 cohort and the earlier cohorts, we see that there is a higher proportion of students continuing in the second year for the Covid-19 cohort. Figure 1 shows the proportion of students continuing in year two. The students have been binned based on their accumulated credits in the first year. The vertical line represents the academic dismissal threshold. Based on the ADP, students to the right of this threshold are allowed to continue in the second year, whereas the students on the left should be academically dismissed based on their academic performance in the first year. For most programs in our sample, the academic dismissal threshold is set at 42 out of 60 credits in year one. This is the case for 92 percent of the programs. The remaining 8 percent of the programs have set the academic dismissal threshold at 40 or 46 out of 60 credits.

(.................................... Figure 1 about here ....)

In Figure 1, the data of the 2015 to 2018 cohorts are represented in blue, whereas the 2019 or Covid-19 cohort is represented in red. We can see that also for the pre-Covid-19 cohorts, there is a significant share of students to the left of the academic dismissal threshold who continue in the second year. These students somehow got a waiver for the ADP. The area under the blue line represents the students who were allowed to continue in the second year in the earlier cohorts. These students can be labeled as *waivers*, because in the earlier cohorts, the ADP was in force. We can also see that in the Covid-19 cohort, a significant share of students to the left of the threshold did not continue despite the abolishment of the ADP. This group of students consists of two subgroups: the students who would not get a waiver but decided to drop out based on their own decision, and the *waivers* who decided to drop out anyway or switch to a different program.

Then, the area between the blue and red lines in Figure 1 represents the difference in drop-out rates of students between the Covid-19 cohort and the earlier cohorts. However, we cannot conclude that this difference stems from the ADP, because the performance in the second semester of the first year in the Covid-19 cohort might have been affected by other Covid-19 effects. As mentioned before, these effects may include the effects of online learning, grading leniency, and mental health effects. Because of this, it is likely that the group of students referred to as *policy-induced dropouts* (PIDs) is smaller than the area between the two lines to the left of the threshold.

3.2 Machine learning model

In order to eliminate the aforementioned Covid-19 effects on the proportion of students continuing in the second year, we only consider pre-Covid-19 data to predict the actual credits by the end of the first year. In the Covid-19 cohort, the academic year started in September 2019. In March 2020, the Covid-19 virus arrived in the Netherlands and the students were affected by the Covid-19 crisis. The first-year ADP in the Netherlands was officially cancelled on 19 March 2020. Therefore, all variables that were collected before the academic started, including survey and register data about student characteristics such as high school GPA, as well as achievement in the first semester of the first year of the program (ending on February 1), are not affected by these other Covid-19 effects.

We use these data in a machine learning model to predict achieved credits by the end of the first year. The objective of the machine learning model consists of two steps. First, we predict achieved credits by the end of the first the academic year based on semester 1 data in the absence COVID-19. Second, we use these predictions to identify students eligible to continue in the 2nd academic year, free of Covid-19 effects. We use a LASSO estimation model (Hastie et al., 2009). The Lagrangian equivalent of the LASSO model can be written as:

$$\beta^{lasso} = argmin_{\beta}\{\frac{1}{2}\sum_{i=1}^{N}(y_i - \beta_0 - \sum_{j=1}^{p} x_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{p} |\beta_j|\} \tag{1}$$

Here, the term $\frac{1}{2}\sum_{i=1}^{N}(y_i - \beta_0 - \sum_{j=1}^{p} x_{ij}\beta_j)^2$ represents the estimation model without penalty, and $\lambda\sum_{j=1}^{p}|\beta_j|$ represents the imposed penalty to prevent over-identification. We use cross-validation and out-of-bag-testing, where we randomly select 70% of the data as a training set with 10-fold cross-validation. For the cohorts 2015-2018 we randomly select 30% to use as an out-of-bag test set. Then, the prediction model parameters are used to predicts achieved credits at the end of year one for the 2019 cohort.

3.3 Conditional Mean Independence

We make our predictions Conditional Mean Independent (CMI) of student background differences observed across training cohorts. To do so, we define the set of variables containing student characteristics at the beginning of the educational program for which there are mean differences at the student cohort level as $\sum_{m=1}^{M} z_m$ and the set of other characteristic variables as $\sum_{j=1}^{J} x_j$. The Lagrangian equivalent for the CMI-LASSO can then be written as:

$$\beta^{lasso} = argmin_{\beta}\{\frac{1}{2}\sum_{i=1}^{N}(y_i - \beta_0 - \sum_{m=1}^{M} z_{im}\beta_m - \sum_{j=M+1}^{J} x_{ij}\beta_j)^2 + \lambda \sum_{j=M+1}^{J} |\beta_j|\} \qquad (2)$$

Table 3 shows descriptive statistics of our training set and the 2019 cohort. We have 11,492 observations in our training set, which corresponds to 70 percent of the students in the 2015 to 2018 cohorts. We observe statistically significant differences in the share of international students. This can be explained by the fact that the share of international students has been increasing over the years, which is shown in the comparison of the different cohorts in Table A1 in the Appendix. We also observe differences in the high school track, high school GPA, and high school math grade. These variables are non-penalized in the CMI-LASSO estimation model.

(.................................... Table 3 about here ....)

3.4 Machine learning results

In Table 4 we report the coefficients of the CMI-LASSO model. The upper part of the table gives the coefficients for the non-penalized variables. We include higher-order polynomials for the math grade to facilitate the CMI-LASSO estimation. We include the polynomials up to the seventh order, because including polynomials of an order higher than seven leads to multicollinearity.

In addition to these non-penalized variables, the CMI-LASSO model has selected semester one credits, a dummy for female students, a dummy for a missing math grade, 18 individual course grades and 17 program dummies.

(. . . . Table 4 about here . . . .)

Table 5a shows the R-squared for the training set, the test set, and the 2019 cohorts, respectively. The R-squared for the training sets amounts to 80.9 percent. The out-of-bag R-squared for the test set is comparable: 82.1 percent. The out-of-bag R-squared for the 2019 cohort amounts to 80.1 percent. These relatively high values of R-squared show that the CMI-LASSO model is not only performing well in predicting end-of-year credits in the training set, but that the model also performs well in predicting end-of-year credits for the Covid-19 cohort. Table 5b shows the model accuracy for the 2015-2018 and 2019 cohorts. We see that accuracy based on both observed and predicted credits decreases due to cancellation of the ADP. The ratios indicate that the model prediction error is about similar for the Covid-19 cohort and the earlier cohorts.

(. . . . Table 5 about here . . . .)

4 Results

4.1 Identifying the policy-induced dropouts

In Figure 2, we plot the proportion of students continuing in the second year against the predicted credit bins at the end of the first year, based on only pre-Covid-19 information. The vertical line represents the academic dismissal threshold, and the values on the horizontal axis are relative to this threshold. This figure is essentially a repetition of Figure 1, but in Figure 2, the other Covid-19 effects are canceled out due to the fact that here we consider *predicted* credits, which are not influenced by Covid-19 effects. From this, it follows that the area between the red and blue lines now represents the PIDs.

(....................................Figure 2 about here....................................)

The lower panel of Figure 2 shows the kernel density estimates and histograms for the 2019 (Covid-19) and previous cohorts. The blue bars represent the students in previous cohorts, and the red bars represent those in the 2019 cohort. The students have been binned based on their predicted credits in year 1. The blue and red lines represent the fitted Gaussian kernel density functions for the previous and 2019 cohorts, respectively. The kernel density functions of the previous and 2019 cohort overlap. This means that by using the predicted credits, we capture the horizontal shift in credits due to the Covid-19 crisis, which represents the change in first-year performance and therefore achieved credits due to other Covid-19 effects. We are then left with the vertical shift, which represents the increase in the proportion of students continuing due to the abolishment of the ADP.

Given that the relative densities of the Covid-19 cohort and earlier cohorts are similar to the left of the academic dismissal threshold, we can determine the number of waivers, PIDs, and dropouts for the Covid-19 cohort (N=6,001). The number of waivers is represented by the area under the blue line. Taking into account the number of students in each bin to the left of the threshold in the Covid-19 cohort, the number of waivers amounts to 1,607 students. The area above the red line to the left of the cutoff gives an representation of the dropouts. Weighted for the number of students in the Covid-19 cohort, the amounts to 1,285 students. Then, the area between the red and blue lines to the left of the threshold weighted for the number of students in the bins in the Covid-19 cohort makes up the PIDs. In total, there are 371 PIDs. Both the waivers and the PIDs are affected by the ADP, which totals 1,978 students, and 18.8 percent of these are PIDs. Given a counterfactual where the ADP was in force in 2020, both dropouts and PIDs would drop out from the program in the first year: 1,656 students. Out of all first-year dropouts, 22.4 percent are PIDs.

4.2 Second-year performance

Figure 3 plots the accumulated credits in second year against the predicted credits in year one. The red line represents a polynomial fit for the Covid-19 cohort, left and right of the ADP threshold separately. The blue line shows the same for the earlier cohorts. On both the left and right-hand sides of the ADP threshold, the blue and red lines overlap. This shows that second-year performance in terms of accumulated credits is not significantly different in the Covid-19 cohort when compared to the earlier cohorts.

(. . . . Figure 3 about here . . . .)

(. . . . Figure 4 about here . . . .)

On the left-hand side of the ADP threshold, there is a difference in the composition of the students in the Covid-19 cohort in comparison to the earlier cohorts. In the earlier cohorts, we do not observe the PIDs, because these would have been dismissed due to the ADP. In other words, the blue polynomial to the left of the ADP threshold represents the waivers. In the Covid-19 cohort on the contrary, we do observe the PIDs in year two because they were not dismissed due to Covid-19. Therefore, we can say that the red polynomial to the left of the ADP threshold represents both the waivers and the PIDs. Since both polynomials on the left-hand side of the ADP threshold overlap, we can conclude that the second-year performance of PIDs and waivers is not statistically significantly different.

4.3 Time-to-degree

Now that we have shown that the second-year performance of the PIDs is not statistically significantly different from the waivers, we can explore the long-term performance of PIDs. Bin-wise, the waivers in earlier cohorts do not perform better than waivers and PIDs in the Covid-19 cohort. Therefore, the waiver-outcomes of earlier cohorts represent the counterfactual outcome for the PIDs. We can use this information to calculate the study duration for PIDs for obtaining a bachelor's degree.

Figure 4 shows the long-term development of total accrued credits for nine predicted-credit bins. In these nine graphs, the slope of the plotted line represents the average number of credits accrued per period. In other words, each panel gives us the average number of credits accrued per period for a given predicted-credit bin. For all cohorts, the development of the average total number of credits over time overlaps in each of the nine predicted-credit bins displayed in Figure 4. However, the slope of the plotted line decreases bin-by-bin when we go from high to low on the year-one predicted credits ladder. In other words, the study duration for obtaining a bachelor's degree for PIDs increases when first-year performance decreases, and this holds for both PIDs and waivers.

Regarding their study pace, students can be grouped based on their first-year performance. First-year performance predicts students' study pace and therefore their time-to-degree. However, there does not seem to be a clear cutoff in terms of first-year performance, study pace, and resulting study duration. In other words: lower first-year performance is associated with longer study duration, but not with higher dropout rates in the second year. Since the PIDs and waivers perform similarly in the second year, we can say that students below the ADP threshold would not perform worse in terms of dropout and study success - just that the study duration would increase.

5 Conclusion

From our results it seems that the waivers represent the performance of the PIDs. Based on this, we can conclude that after correcting for COVID-19 effects in year one, the pandemic did not lead to decreased performance in the second year of those students who were affected in year one. One way or the other, this could be either due to the fact the special circumstance leading to a waiver of the policy was not that special, or that this continued to affect the waiver students in such a way that they did not change their pace later on.

We can conclude that 22.4 per cent of first-year dropouts at a large research university in the Netherlands are policy-induced (PIDs). These PIDs do not perform significantly different than students who got a waiver for the academic dismissal program. It seems that students have a constant individual study pace, and from a student perspective, there does not seem to be a strict 'unsuitability' threshold in terms of year-one performance. In terms of educational inequality, this brings up an important issue as well. Why do we dismiss PIDs based on an arbitrary academic dismissal threshold if they perform equally well as students who got a waiver?

In short, we find that students have their own study pace, which does not seem to be affected by the academic dismissal policy. Students can be categorized in terms of their study pace based on their achieved credits in year one. From a normative policy perspective, any strictly enforced academic dismissal threshold is questionable. Left or right of any threshold, students have a similar study pace, and there does not seem to be an argument to dismiss part of these students on the basis of 'unsuitability' or labeling them as not being 'college material'.

References

Angrist, J., D. Lang, and P. Oreopoulos (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics 1*(1), 136–163.

Calcagno, J. C. and B. T. Long (2008). The impact and of postsecondary and remediation using and a regression and discontinuity approach and addressing endogenous and sorting and and noncompliance. NBER Working Paper Series 14194, National Bureau of Economic Research.

Cornelisz, I., R. van der Velden, I. de Wolf, and C. Van Klaveren (2020). The consequences of academic dismissal for academic success. *Studies in Higher Education 45*(11), 2175–2189.

Deelen, A. and S. Kuijpers (2016). The effects of additional study choice activities: Evidence from a randomized controlled trial. CPB Discussion Paper 338, CPB Netherlands Bureau for Economic Policy Analysis.

Ost, B., W. Pan, and D. Webber (2018). The returns to college persistence for marginal students: Regression discontinuity evidence from university dismissal policies. *Journal of Labor Economics 36*(3), 779–805.

Page, L. C. and J. Scott-Clayton (2016). Improving college access in the United States: Barriers and policy responses. *Economics of Education Review 51*, 4–22.

Scott-Clayton, J., P. M. Crosta, and C. R. Belfield (2014). Improving the targeting of treatment: Evidence from college remediation. *Educational Evaluation and Policy Analysis 36*(3), 371–393.

Scott-Clayton, J. and O. Rodriguez (2015). Development, discouragement, or diversion? New evidence on the effects of college remediation policy. *Education Finance and Policy 10*(1), 4–45.

Scott-Clayton, J. and B. Sacerdote (2016). Access to higher education: Introduction to the special issue. *Economics of Education Review 51*, 1–3.

Sneyers, E. and K. D. Witte (2017). The effect of an academic dismissal policy on dropout, graduation rates and student satisfaction. evidence from the Netherlands. *Studies in Higher Education 42*(2), 354–389.

Valentine, J. C., A. S. Hirschy, C. D. Bremer, W. Novillo, M. Castellano, and A. Banister (2011). Keeping at-risk students in school. *Educational Evaluation and Policy Analysis 33*(2), 214–234.

Van Klaveren, C., K. Kooiman, I. Cornelisz, and M. Meeter (2018). The higher education enrollment decision: Feedback on expected study success and updating behavior. *Journal of Research on Educational Effectiveness 12*(1), 67–89.

Tables

Table 1: Student numbers VU Amsterdam.

	Students enrolled	
	First year	All years
2015	3,620	23,665
2016	3,691	23,452
2017	4,067	24,203
2018	5,045	25,398
2019	6,001	28,025
Total	22,424	124,743

Table 2: Dropout statistics VU Amsterdam.

	Dropout		Continue		Total	
	No.	%	No.	%	No.	%
2015	968	26.7	2,652	73.3	3,620	100
2016	925	25.1	2,766	74.9	3,691	100
2017	1,279	31.4	2,788	68.6	4,067	100
2018	1,362	27.0	3,683	73.0	5,045	100
2019	1,223	20.4	4,778	79.6	6,001	100
Total	5,757	25.7	16,667	74.3	22,424	100

Table 3: Descriptive statistics: Training set and 2019 cohort.

	Training set (N=11,493)			2019 cohort (N=6,001)				
	Mean	SD	Range	Mean	SD	Range	t-stat	p-val
Outcomes:								
EC Year 1	42.31	20.91	$\{0, 60\}$	41.82	20.99	$\{0, 60\}$	1.49	0.136
Predictors:								
EC Semester 1	20.25	10.15	$\{0, 54\}$	19.99	10.42	$\{0, 48\}$	1.61	0.107
Double degree	0.03	0.17	$\{0, 1\}$	0.03	0.17	$\{0, 1\}$	1.13	0.259
Female	0.55	0.50	$\{0, 1\}$	0.55	0.50	$\{0, 1\}$	0.25	0.803
International	0.11	0.31	$\{0, 1\}$	0.16	0.37	$\{0, 1\}$	-11.10	0.000
College	0.09	0.29	$\{0, 1\}$	0.11	0.31	$\{0, 1\}$	-2.85	0.004
High school	0.76	0.43	$\{0, 1\}$	0.69	0.46	$\{0, 1\}$	9.85	0.000
High school GPA	6.66	0.42	$\{4, 9\}$	6.65	0.41	$\{4, 9\}$	2.51	0.012
(missing)	0.15	0.36	$\{0, 1\}$	0.24	0.43	$\{0, 1\}$	-13.82	0.000
Math grade	6.75	0.88	$\{3, 10\}$	6.71	0.83	$\{3, 10\}$	2.75	0.006
(missing)	0.17	0.37	$\{0, 1\}$	0.25	0.43	$\{0, 1\}$	-13.52	0.000
Travel time	58.62	30.74	$\{15, 258\}$	58.73	30.71	$\{15, 254\}$	-0.23	0.817
(missing)	0.06	0.24	$\{0, 1\}$	0.06	0.24	$\{0, 1\}$	-0.22	0.825
Program dummies		Yes (45)			Yes (45)			
Grades semester 1		Yes (573)			Yes (573)			

Overall F-test: 34.5270 ($p < 0.000$).

Table 4: CMI-LASSO output (coefficients)

Dep. Var. Credits End Year 1	Estimated $\hat{\beta}_j^{lasso}$ (if $\hat{\beta}_j^{lasso}$ >0)
Non-penalized:	
International	-0.390
College	1.114
Double degree	.124
Math grade	817.459
Math grade^2	-417.012
Math grade^3	115.631
Math grade^4	-18.836
Math grade^5	1.8048
Math grade^6	-.094
Math grade^7	.002
Selected:	
EC Semester 1	1.779
Female	0.441
Math grade missing	-.0910
Individual grades	Included (18)
Program dummies	Included (17)
Constant	-707.259
λ	4936.187
N	11,493

Table 5: Machine learning performance indicators

(a) R-squared

	Training set	Test set	2019 cohort
	R^2	R^2	R^2
LASSO	0.80875828	0.82068926	0.80058939

(b) Accuracy

	Accuracy observed credits	Accuracy predicted credits	Ratio
2015-2018	0.845	0.727	1.162
Covid-19	0.766	0.667	1.148
Ratio	1.103	1.090	

Figures

Figure 1: Proportion continuation by accumulated credits in year 1.

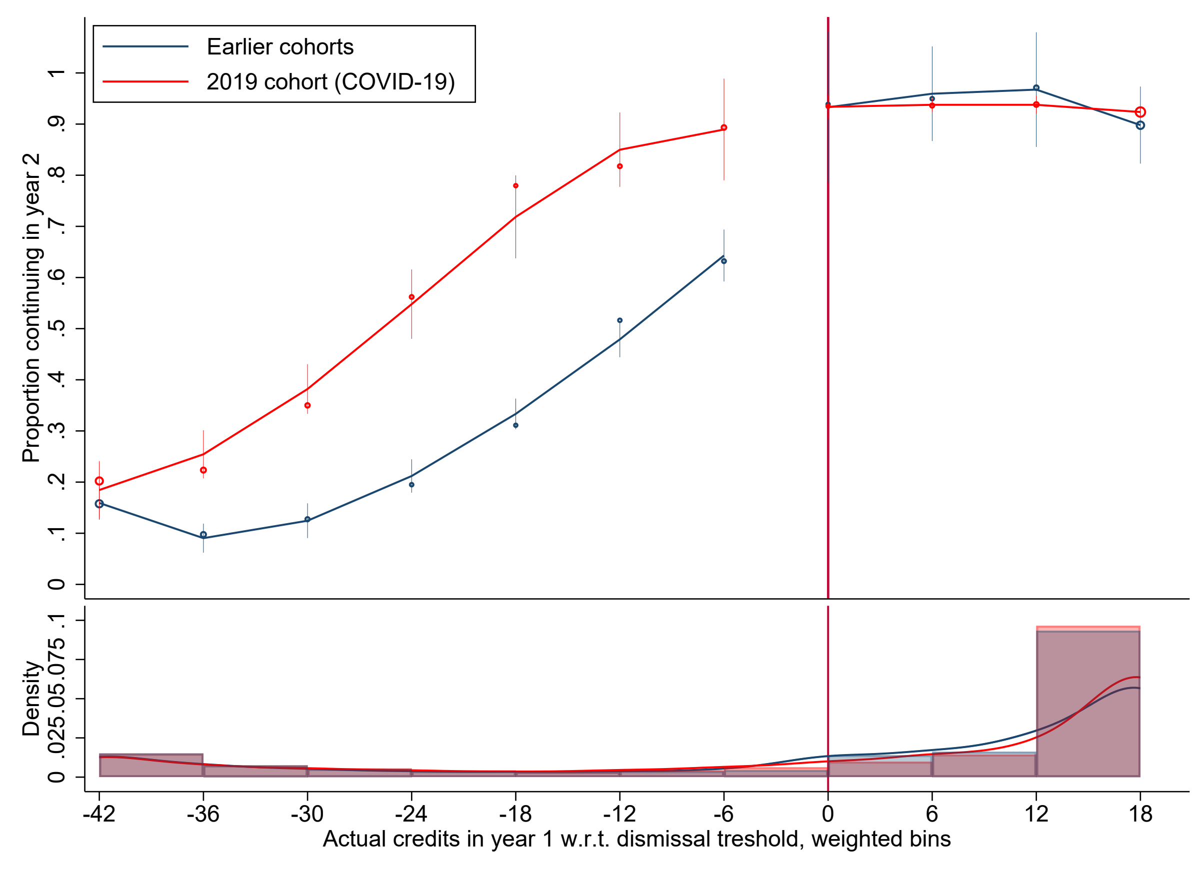

Figure 2: Proportion continuation by predicted credits.

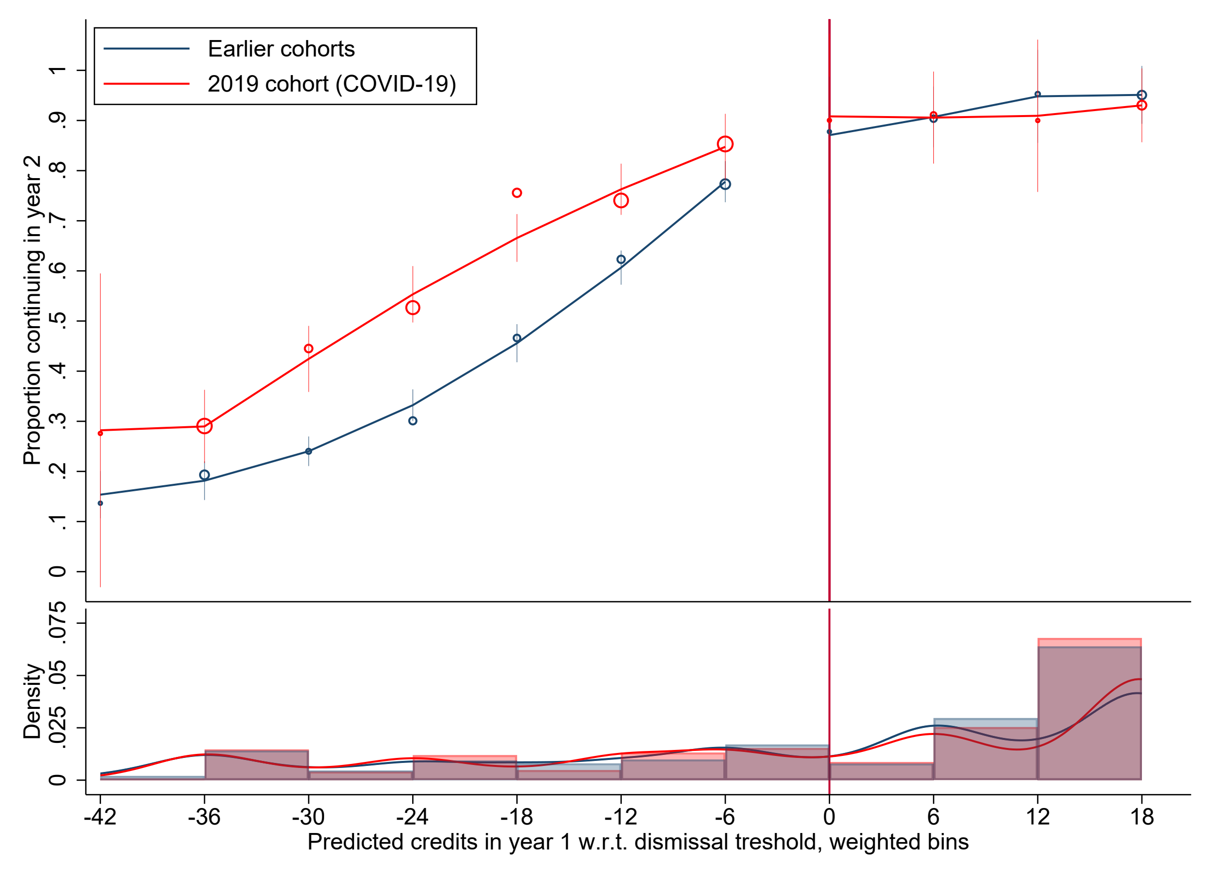

Figure 3: Accumulated credits in second year by predicted credits in first year.

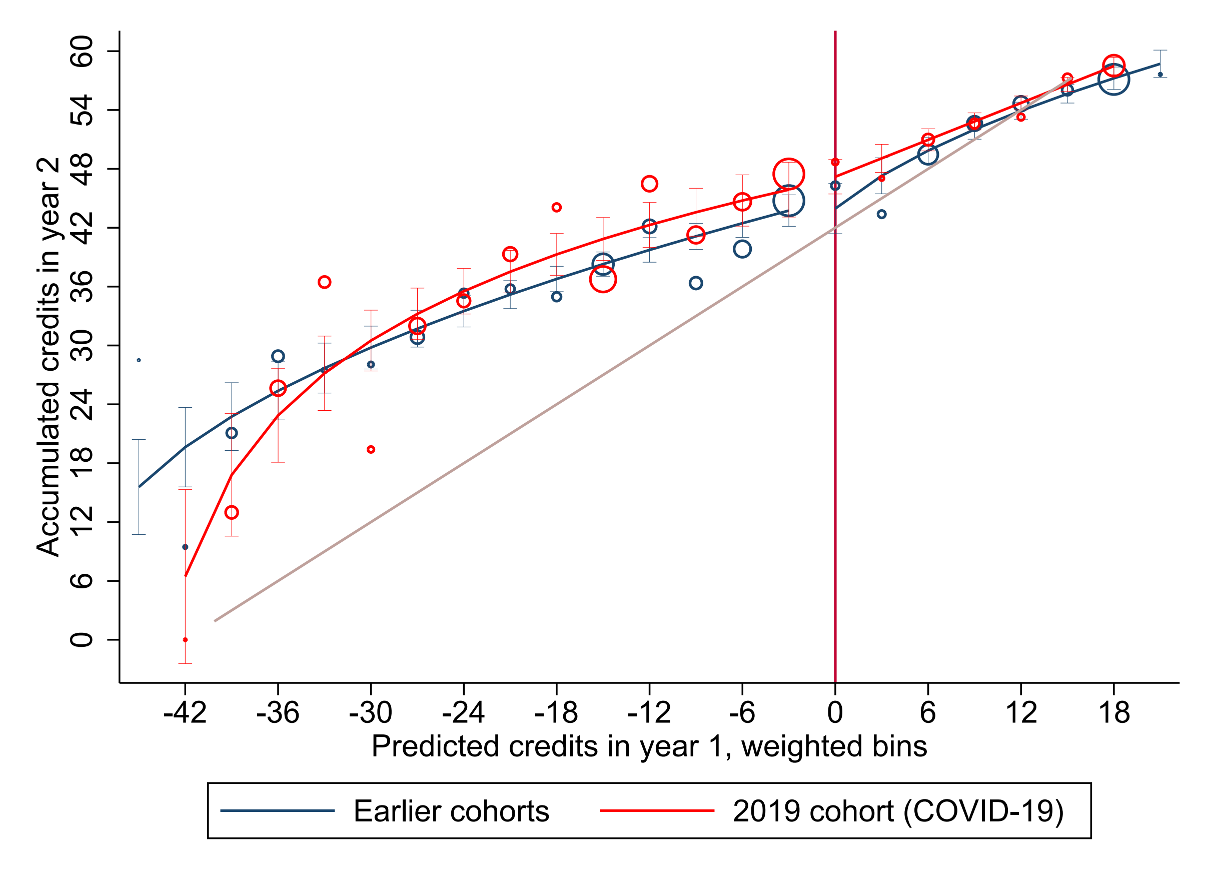

Figure 4: Total accumulated credits by period until graduation for Covid-19 and earlier cohorts, by predicted credit bins.

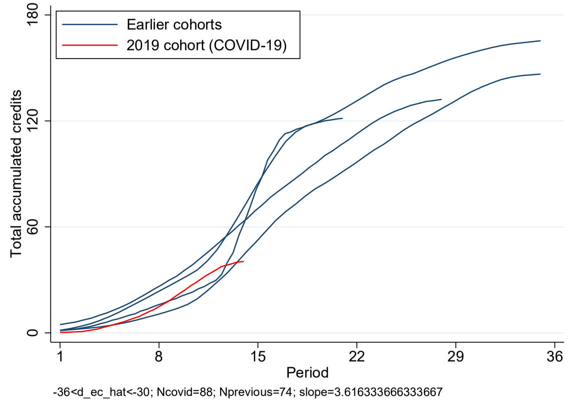

(a)

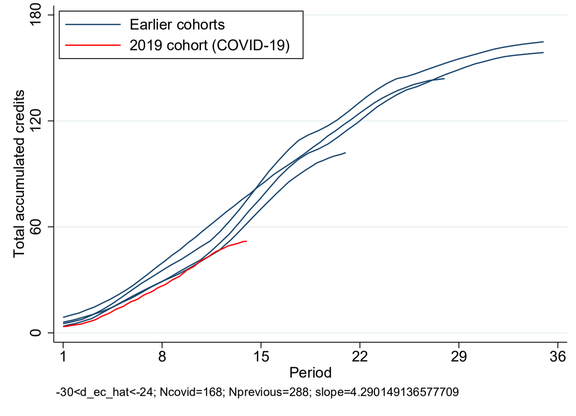

(b)

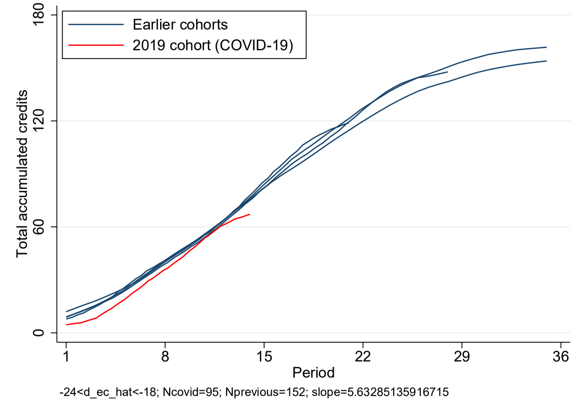

(c)

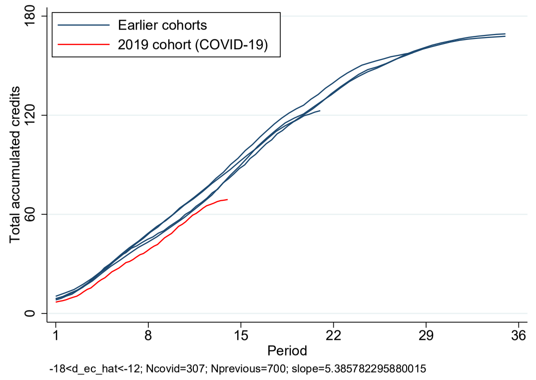

(d)

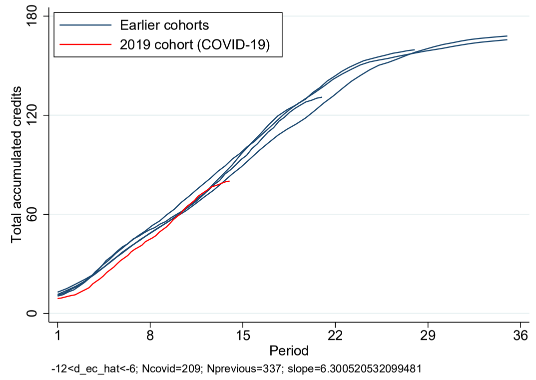

(e)

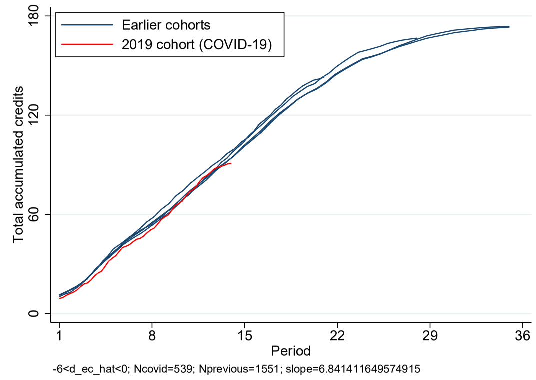

(f)

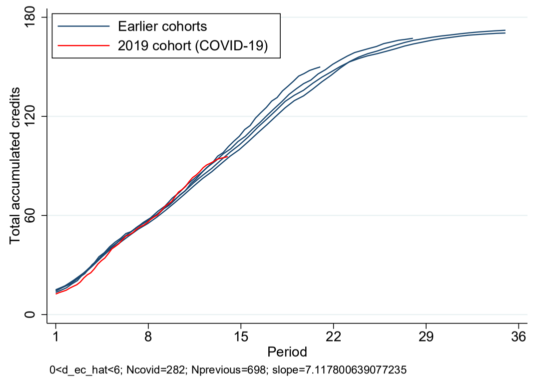

(g)

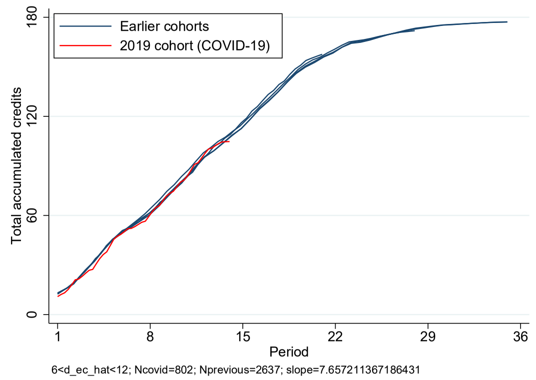

(h)

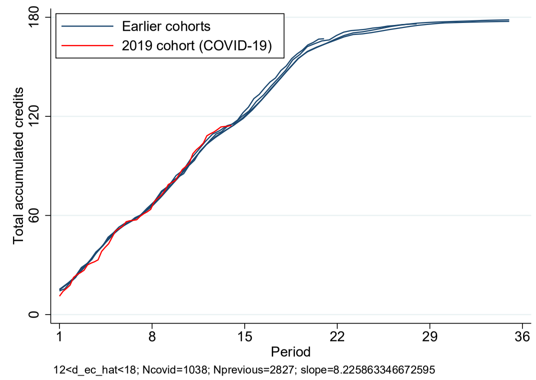

(i)

Appendix

Table A1: Descriptive statistics: Training set cohort comparison

	2015 (N=2,547)		2016 (N=2,546)		2017 (N=2,828)		2018 (N=3,572)			
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD	F-st.	p-val
Outcomes:										
EC Year 1	42.53	20.94	42.50	20.69	42.51	20.81	41.88	21.13	0.76	0.517
Predictors:										
EC Semester 1	20.25	10.21	20.38	10.07	20.42	10.17	20.02	10.14	1.02	0.382
Double degree	0.03	0.18	0.04	0.20	0.03	0.16	0.03	0.16	4.04	0.007
Female	0.54	0.50	0.53	0.50	0.56	0.50	0.56	0.50	2.11	0.096
International	0.06	0.24	0.08	0.27	0.13	0.33	0.14	0.35	46.84	0.000
College	0.11	0.31	0.07	0.26	0.07	0.26	0.11	0.32	16.47	0.000
High school	0.79	0.41	0.81	0.39	0.75	0.43	0.71	0.46	35.87	0.000
High school GPA	6.67	0.43	6.67	0.44	6.66	0.42	6.66	0.41	0.56	0.645
(missing)	0.10	0.30	0.10	0.31	0.18	0.38	0.21	0.41	67.81	0.000
Math	6.72	0.91	6.73	0.91	6.80	0.88	6.73	0.84	4.45	0.004
(missing)	0.11	0.32	0.11	0.32	0.19	0.39	0.22	0.42	66.91	0.000
Travel time	57.39	30.75	57.76	30.60	59.84	30.91	59.13	30.67	3.84	0.009
(missing)	0.02	0.15	0.04	0.19	0.08	0.28	0.08	0.27	47.22	0.000