

“High” Achievers? Cannabis Access and Academic Performance

OLIVIER MARIE

Erasmus University Rotterdam, Maastricht University, Tinbergen Institute, Centre for Economic Performance, Institute for the Study of Labor (IZA), and Center for Economic Studies-Ifo

and

ULF ZÖLITZ

briq—Institute on Behavior and Inequality, IZA and Maastricht University

First version received March 2015; Editorial decision January 2017; Accepted March 2017 (Eds.)

This paper investigates how legal cannabis access affects student performance. Identification comes from an exceptional policy introduced in the city of Maastricht in the Netherlands that discriminated access via licensed cannabis shops based on an individual’s nationality. We apply a difference-in-difference approach using administrative panel data on course grades of local students enrolled at Maastricht University before and during the partial cannabis prohibition. We find that the academic performance of students who are no longer legally permitted to buy cannabis substantially increases. Grade improvements are driven by younger students and the effects are stronger for women and low performers. In line with how cannabis consumption affects cognitive functioning, we find that performance gains are larger for courses that require more numerical/mathematical skills. Our investigation of underlying channels using course evaluations suggests that performance gains are driven by an improved understanding of the material rather than changes in students’ study effort.

Key words: Marijuana, Legalization, Student performance.

JEL Codes: I18, I20, K42

1. INTRODUCTION

Public policy and opinion regarding the legalization of cannabis has reached a tipping point. As of 2016, 23 U.S. states and the District of Columbia have passed laws allowing the medical use of marijuana and four additional states have decriminalized consumption for recreational use.¹ Uruguay recently became the first nation in the world to fully legalize all aspects of the cannabis trade, including cannabis cultivation, wholesale, retail, and consumption. The Americas

1. In 2012, Colorado and Washington passed laws legalizing sale and possession of cannabis for recreational use after a popular vote. Alaska and Oregon have followed suit in 2014 and 2015, respectively. In addition to this, several other states are currently reconsidering their cannabis laws, which is likely to lead to more liberal regulations in the future. For example, Missouri is also scheduled to implement a decriminalization law on 1 January 2017. Updates on the fast-evolving U.S. situation can be found on the site of the Marijuana Policy Project (www.mpp.org).

are starting to “catch up” with the more liberal approach to soft drug policy in countries such as the Netherlands, where cannabis consumption has been decriminalized for almost four decades. Despite this development, little is known about many of the—perhaps unintended—consequences of legalization. This paper contributes to the ongoing legalization discussion by showing that a change in legal cannabis access can strongly affect students’ university performance.

Proponents of cannabis legalization have put forward the general failure of the long-running “war on drugs” and the huge costs that it imposes on the criminal justice system as an argument in favour of finding alternatives to drug prohibition (Donohue, 2013). They have also argued that legalization would undermine illegal markets and protect low-level users from associated risks such as contact with dealers who sell other types of drugs. Opponents of cannabis legalization often argue that making access to cannabis easier and more acceptable via legalization could push more individuals—especially youths—to become consumers. This could in turn lead to an increase in the number of individuals suffering from the adverse health, educational, and labour market outcomes associated with regular cannabis use (Cobb-Clark *et al.*, 2016; Hall, 2015; Van Ours and Williams, 2015).

While both sides of the legalization debate make plausible arguments, the actual effect of policies changing cannabis access on consumption decisions and outcomes influenced by cannabis use remains largely ambiguous. The lack of clear empirical evidence results from enduring identification problems, which mostly prevent a causal interpretation of most existing results. The principal issue is that drug policy changes are unlikely to be implemented exogenously and are usually the result of a longer process of societal change. When policy changes take place, they usually affect all individuals at the same time, thus making it impossible to fully disentangle treatment effects from underlying time trends in consumption: trends that may have caused the policy change in the first place. These issues cast doubt on the validity of results obtained from studies using cohort- or state-level variation, where the necessary *ceteris paribus* conditions for identification often do not hold.²

In this article, we utilize a unique natural experiment to obtain causal estimates of how changes in legal cannabis access affect students’ college performance. We exploit a temporary policy change in the city of Maastricht in the Netherlands, which locally restricted legal access to cannabis based on individuals’ nationality.³ After providing empirical evidence that the new policy essentially eliminated legal cannabis sales to the treated nationalities, we subsequently compare the achievements of university students who were affected by the cannabis ban against those of their peers who were unaffected. This unusual “partial prohibition” of legal marijuana access allows us to apply a difference-in-difference approach across nationality groups of students observed before and during the discriminatory policy. To eliminate any remaining concerns about unobserved individual heterogeneity, we exploit the panel nature of our data and apply student fixed effects to identify performance changes resulting from the prohibition policy using *within-individual* variation in outcomes.

From a medical perspective, there is substantial evidence on negative short-run effects of cannabis consumption on performance. Studies have repeatedly shown that cognitive functions

2. A review of this literature by Pacula and Sevigny (2014) discusses a number of recent articles using state-level difference-in-differences approaches to assess how the introduction of medical cannabis laws in the U.S. affects consumption behaviour and other outcomes. The review reports mixed findings and highlights multiple reasons (*e.g.* changes in police force behaviour) for why results might not have a causal interpretation. Anderson *et al.* (2013) show some evidence of increased self-reported cannabis use among adults following the introduction of medical marijuana laws in a study of their impact on traffic fatalities across U.S. states. In another context, Jacobi and Sovinsky (2016) develop a structural model that predicts, using Australian survey data, very large increases in cannabis consumption if legalization was implemented.

3. Importantly, students were not the intended target of the discriminatory policy, which was originally introduced to combat drug tourism in the city.

are strongly impaired by cannabis consumption in the short run.⁴ Therefore, we expect changes in cannabis consumption behaviour—brought about by the access restriction scheme studied here—to be reflected in the academic performance of the students affected. This rationalizes our reduced-form approach, which looks directly at student productivity—rather than (unavailable) individual changes in consumption—as an outcome.⁵

Our main finding is that the temporary restriction of legal cannabis access had a strong positive effect on course grades of the affected individuals. On average, students performed 10.9% of a standard deviation better and were 5.4% more likely to pass courses when they were banned from entering cannabis shops. Importantly, we do not detect a change in dropout probability, which could have created complex composition effects. Sub-group analysis reveals that these effects are somewhat stronger for women than men and that they are driven by younger and lower performing students. This can be explained by baseline differences in consumption rates or differences in marginal compliance with the prohibition. We also find some evidence for a social spillover of the cannabis restriction, whereby treated students in sections with a higher fraction of treated peers become marginally more likely to pass their courses. We can reject the notion that teachers' legal access to cannabis has an impact on their students' performance.

In order to assess whether the changes in performance that we detect genuinely stem from changes in students' cannabis consumption, we test whether our results are consistent with what is known about the impact of tetrahydrocannabinol (THC, the principal psychoactive constituent of cannabis that makes the user “high”) on human brain functioning and learning. First, previous research has documented that cannabis consumption most negatively influences quantitative thinking and math-based tasks (Block and Ghoneim, 1993; Pacula *et al.*, 2003). Therefore, we split all courses depending on whether they are described as requiring numerical skills or not and test whether such course grades are differentially affected. We find that the policy effect is 3.5 times larger for courses requiring numerical/mathematical skills: a result in line with the existing evidence on the association between cannabis use and cognitive functioning. Second, to provide some suggestive evidence on the underlying channels, we make use of evaluations that students are asked to complete for each course. In these evaluations, students report their own level of effort, overall understanding and the perceived quality of the course and teachers. We find no change in reported study hours, which suggests that we can eliminate effort adjustments as one channel of our results. We find some evidence of improved peer-to-peer interactions and an increase in the reported overall understanding of the course content when the policy was in place.

In order to test whether the legal cannabis access had any persistent effects on student outcomes, we investigate whether there are any detectable cohort-level treatment effects on longer-run measures of performance, namely final grade point average (GPA), graduation probabilities, and the type of elective courses chosen, whereby only the latter appears to be affected. Treated students exposed to the policy when they were about to make their specialization decisions choose relatively more courses with mathematical/statistical skills requirements. This is perhaps unsurprising considering the strong performance improvements in these types of courses that we document when the partial prohibition policy is in place.

4. Bossong *et al.* (2012a) and Bossong *et al.* (2012b) conduct randomized control trials where subjects have to carry out simple cognitive task and find that “performance was impaired after THC administration, reflected in both an increase in false alarms and a reduction in detected targets.”

5. This reduced-form approach avoids serious measurement problems with usual measures of drug consumption, since it does not have to rely on self-reported consumption or police seizures, which are likely to be correlated with changes in the legal status of this substance. Another highly relevant short-run outcome that might be affected by changes in soft drug access policy is criminal activity. This is perhaps not as “clean” an externality as productivity since it is the sum of changes in behaviour of all agents concerned, namely consumers, dealers, and the police. Adda *et al.* (2014) are the only authors to have attempted to disentangle the various channels from this complex relationship.

Finally, we put our main finding in perspective with respect to the estimated impact of other interventions on college student performance. It is most relevant that our change in legal cannabis access has almost exactly the same effect as students reaching the age when alcohol consumption is permitted in the U.S. (Carrell *et al.*, 2011; Lindo *et al.*, 2013). To gain a better understanding of the prevalence of cannabis consumption, we also carried out a survey among current students at Maastricht University, revealing that almost 60% had consumed cannabis in the past year. Using this as a proxy for the share of the potential population, we calculate the treatment effect on the treated under various potential compliance rates and argue that the prohibition policy had a large and positive impact on student performance.

To our knowledge, this paper presents the first solid causal evidence that a legal change in access to cannabis has had a strong short-run impact on productivity. However, it is important to note that we are only looking at a very specific outcome and that our results are only a small part of the multi-dimensional societal cost-benefit analysis that should drive drug policy decision-making.

The remainder of the paper is structured as follows. Section 2 provides general information on Dutch cannabis policy and presents the details of the particular change in cannabis access that occurred in Maastricht. Section 3 discusses the data on student performance that we collected at Maastricht University. Section 4 describes our empirical strategy and the various specifications that we will consider. Section 5 presents the main estimation results and carries out sensitivity analysis and placebo tests, before Section 6 explores underlying mechanisms. Section 7 looks at medium-run effects and puts the effect size in perspective and finally Section 8 provides concluding remarks.

2. BACKGROUND: CANNABIS ACCESS IN THE NETHERLANDS AND THE MAASTRICHT CASE

2.1. *The Dutch drug policy approach*

The sale and consumption of cannabis for recreational use has been legal in the Netherlands for over four decades. The 1976 Opium law—which forms the basis of the Dutch “tolerance” policy—was introduced to “minimize harm done to users and their environment” (Maccoun and Reuter, 1997). Practically, possession of up to 30 g of cannabis (1.06 ounces) has not been a prosecutable offense since this law was passed. The Dutch government still aims to reduce demand by means of preventive campaigns and by taking legal measures against any disturbance to public order caused by cannabis sale or consumption. Although personal recreational soft drug use is tolerated, all hard drug use is illegal. The production and illegal sale of hard *and* soft drugs are a severe offense and can result in custodial sentences. Cannabis is usually consumed mixed with tobacco and smoked in “joints” or pipes in the Netherlands. The average concentration of THC in the cannabis sold legally in the Netherlands in 2010 was around 16.7%, which is almost twice as strong as illegal marijuana confiscated in the U.S. However, following its legalization in certain U.S. states, the average potency of cannabis appears to have caught up with Dutch products, with the average strength of strains sold legally in Colorado recently being estimated to contain over 18% of THC.⁶

6. A monitoring survey of the strength of the strains sold in Dutch cannabis shops by Rigter and Niesink (2010) from the Netherlands Institute of Mental Health and Addiction (The Trimbos Institute) estimated that the average THC concentration was at about 16.7% in 2009–10. For the U.S., the UNODC: United Nations Office of Drugs and Crimes (2012) reports an average THC strength of 8.6% in confiscated (illegal) cannabis. Some recent evidence from preliminary laboratory tests on Colorado’s legally purchased marijuana revealed an average concentration level of 18.7% in 2015 (Lafraite and Armentano, 2015).

Through legal channels, cannabis in the Netherlands can be bought exclusively via cannabis shops, which are strictly regulated and can only function with a license granted by the municipality.⁷ Cannabis shops are not allowed to sell more than 5 g per person per day and they are not allowed to have more than 500 g at the shop premise. Furthermore, cannabis shops are not allowed to sell any hard drugs, advertise their products or sell their products to people under the age of 18 years. Cannabis shops can be shut down temporarily or permanently by the license issuing municipality if they fail to meet the regulation requirements or if they are perceived as being responsible for excessive public disturbance.

2.2. *The Maastricht situation*

Maastricht is the southern-most large city in the Netherlands. Due to its geographical proximity to Belgium, Germany, Luxembourg, and especially France, it has attracted a great deal of “drug tourists” for many years, coming to buy (and consume) cannabis legally. As a result, it has a high density of cannabis shops per population, second only to Amsterdam, a city infamous for international cannabis tourism. Figure 1 presents a map depicting the cannabis shop density of the 443 municipality districts of the Netherlands. Maastricht (circled) is located at the very south-east of the map in the region, encased between Belgium and Germany. In 2011, the city had thirteen cannabis shops across a population of about 122,000 inhabitants. A substantial part of the city’s population are students: overall, there are about 16,000 individuals studying at Maastricht University in any given year, more than half of whom are non-Dutch nationals. Figure 1 also shows that not all Dutch cities (only one-third) have cannabis shops and that the nearest one outside of Maastricht is more than 25 km away.

2.3. *The policy change in cannabis access*

Starting from 1 October 2011, the Maastricht association of cannabis shop owners (VOCM)—under pressure from local authorities—introduced a new policy that only allowed specific nationalities to buy cannabis on their premises. The aim of this policy was to reduce negative externalities arising from drug tourism, which the city argued constituted a public nuisance that could lead to the closure of most establishments. The policy targeted a specific nationality group of drug tourists, mostly individuals from France and Luxembourg, which the city council “identified” as the most nuisance-prone population and imposing the highest negative externalities on city residents. In a compromise, the VOCM convinced the municipality to maintain access to their cannabis shops not only exclusively to Dutch citizens but also to individuals from the two neighbouring countries—Germany and Belgium—in an attempt to solve the drug tourism problem. Retaining access rights for these three nationalities was crucial for the Maastricht establishments as these together represented the majority of their customers. The new policy was locally announced by retailers to inform users about two months before its official start.⁸

7. Cannabis shops are by far the most common place where users purchase the drug in the Netherlands. A recent survey carried out in the southern provinces—including Limburg, where Maastricht is located—asked consumers where they had obtained cannabis in the past 3 months (note that more than one answer was possible). It reveals that 91% reported they had bought cannabis from a cannabis shop, 25% from a street dealer, and 8% had used a product from home-grown plants (it is legal to have up to five marijuana trees at home in the Netherlands). See Van Ooyen-Houben *et al.* (2014) for details.

8. The policy was announced and implemented with a relatively short notice. Therefore, student application or enrolment decisions for the academic year 2011/2012 could not have been affected by the policy change. Since this information was not publicly available at the time when these decisions were taken, there is no reason to believe that the student composition of Maastricht University changed due to the policy change.

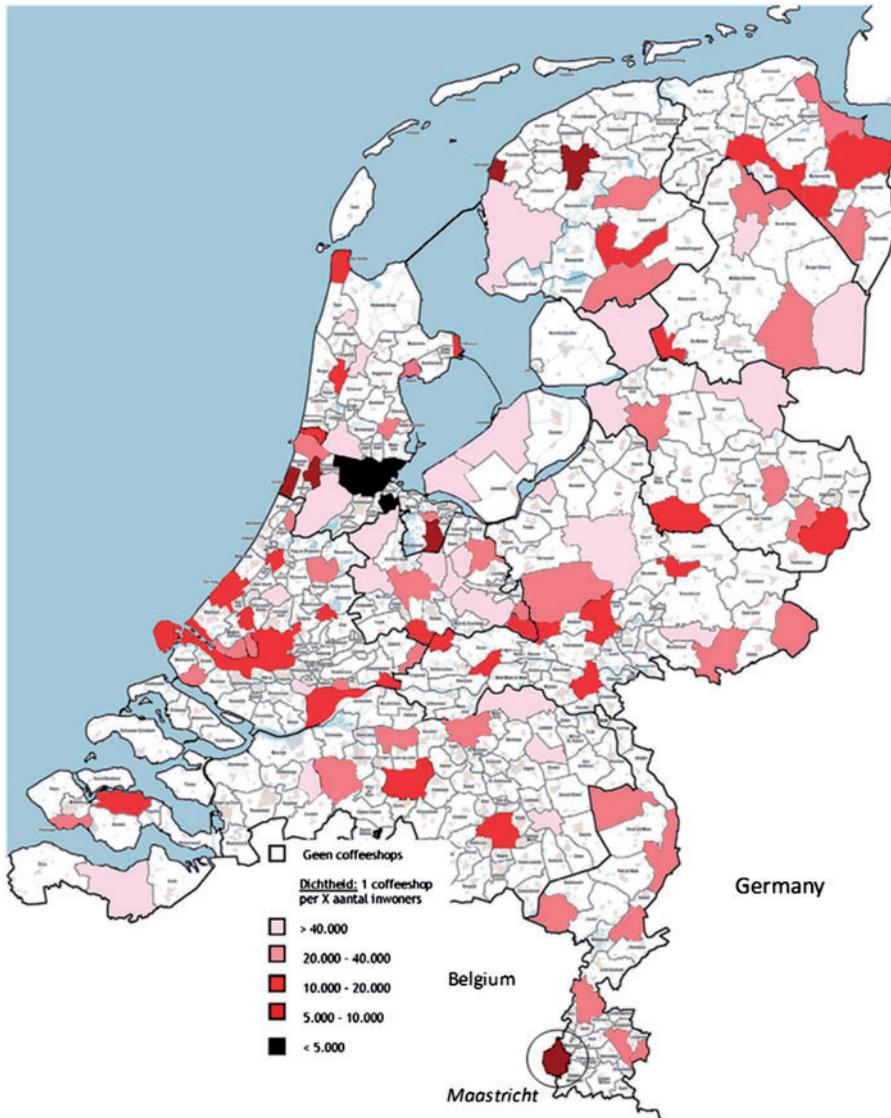


FIGURE 1

Number of cannabis shops per population across Dutch municipalities (Bielman *et al.*, 2013)

Figure 2 shows the (very discriminatory) poster announcing the policy change, which cannabis shops were required to put up on the front window of their premises. From 1 October 2011, anyone who was unable to present a valid Dutch, German, or Belgian form of identification was refused entry to cannabis shops. In Maastricht, all establishments have always been required to scan such documents when costumers enter to insure compliance with the minimum legal age requirement, which was now also used to enforce the nationality criteria.

To assess whether the access restriction indeed had an impact on the legal purchase of cannabis by nationality status, we obtained data on cannabis shop visitors in Maastricht collected before



FIGURE 2

Poster announcing the application of the “neighborhood country criterion” displayed in front of Maastricht Cannabis shops on 1 October 2011

and after the policy was introduced.⁹ Table 1 reports the composition of the customer population by nationality in September and October 2011. It first shows that before the policy almost one-fifth of costumers in Maastricht coffee shops were of another nationality than Dutch, German, or Belgian. In the following month, after the access restriction by nationality became effective on 1 October, the non-Dutch, German, and Belgian (DGB) population represented less than

9. This survey recording information on cannabis shop visitors in Maastricht was originally conducted by a local independent research institute (OWP <http://www.owp.nl/>) and took place during one week before (September 12–18) and one week after (October 10–16) the implementation of the partial prohibition policy on 1 October 2011. Customers to any of the city’s 13 cannabis shops were counted and asked to present an ID to record their nationality. These visitor counts were conducted for 10 min, four times a day (at noon, 4 p.m., 8 p.m., and 11 p.m.) on all 7 days of the week. All visitors were classified as Dutch, Belgian, German, French, Luxembourgish, and “Other Nationality”. During these 2 weeks, there were no domestic and foreign holidays or any major events in the city of Maastricht that could have affected the number of visitors from specific countries.

TABLE 1
Nationality composition of Maastricht cannabis shop customers in the month before and after the policy change

Nationality	Visitors before the restriction of legal cannabis access	Visitors after the restriction of legal cannabis access
	(September 2011) (%)	(October 2011) (%)
Non-treated nationalities:		
Dutch	16.56	20.94
Belgian	58.22	70.19
German	6.82	7.44
Treated nationalities:		
French	9.90	0.29
Luxembourg	2.12	0.04
Other nationality	6.39	1.10
Sample size	4,955	4,145

Notes: This survey recording information on cannabis shop visitors in Maastricht was originally conducted by a local independent research institute (OPW <http://www.owp.nl/>) and took place during weeks before (September 12–18) and after (October 10–16) the implementation of the partial prohibition policy on 1 October 2011. Customers at any of the city's 13 cannabis shops were counted and asked to present an ID to record their nationality. These visitor counts were conducted for 10 min, four times a day (at noon, 4 p.m., 8 p.m., and 11 p.m.) on all 7 days of the week. All visitors were classified as Dutch, Belgian, German, French, Luxembourgish, and "Other Nationality". During these 2 weeks, there were no domestic or foreign holiday that could have affected the number of visitors from specific countries.

1.5% of costumers. These descriptive statistics show that the policy was quite strictly enforced with a very low level of non-compliance and confirm that legal sales to banned nationalities in Maastricht had almost completely stopped immediately after its introduction. Thus, even if we do not have information on individual smoking behaviour, we are quite convinced that it must have been affected by the policy, given that it so radically altered legal purchase behaviour.¹⁰

The "neighbourhood criterion" was in place for 7 months, from 1 October 2011 until 30 April 2012. From May 1 until around mid-June 2012, cannabis shops in Maastricht went on strike owing to the planned introduction of a new scheme by the municipality, called the weed-pass (wietpas). This was part of a new cannabis access policy that the central government wanted to introduce in all southern provinces of the Netherlands, which required anyone who wanted to maintain access to cannabis shops to register as a user at the local municipality.¹¹ Around mid-June, the cannabis shop strike ended, after which only residents with a valid weed-pass were allowed access to cannabis shops. In this study, we only consider the period up to the end of the cannabis shop strike, which means that we include three clear access policy periods in our analysis: all access, non-DGB restricted, and all restricted. Figure 3 graphically depicts the timing of the policy changes and puts it in perspective with teaching and examination periods at Maastricht University, which we discuss in detail in the following section.

10. Note that we will identify a lower bound of the actual treatment effect if some students do not comply with the policy change and continue using marijuana purchased from illegal sources, *e.g.* street-dealers or peers.

11. We do not study the period after the end of the strike when the weed-pass was introduced as it is impossible for us to identify which students decided to register at the municipality to obtain a "weed-pass", which granted them legal access to cannabis shops based purely on residency. More information about some of the effects of this policy introduced throughout the Southern Netherlands can be found in Van Van Ooyen-Houben *et al.* (2014).

Cannabis Access by Nationality	All Access												Non-DGB Restricted				All Restricted					
Month	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
Year	2009				2010				2011				2012									
Academic Year	2009 / 2010						2010 / 2011						2011 / 2012									
Teaching Periods in Each Academic Year	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5					
Total Periods with Course Outcome	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17					

FIGURE 3

Timing of changes to cannabis access in Maastricht and mapping to academic year/period with student course grades

3. DATA

3.1. Student performance data

The School of Business and Economics (SBE) is one of the largest schools of Maastricht University. On average, there are almost 5,000 students enrolled in the bachelor, master, and PhD programs of the SBE at any time. We obtained administrative information on all undergraduate students enrolled during the academic years 2009/2010, 2010/2011, and 2011/2012 from the SBE's examination office. Overall, we observe 57,019 course outcomes (including dropouts with no grade) from 4,419 different individuals in our main sample who are—over this period—in one of the three years that it takes to complete a bachelor's degree. Slightly more than one-third of students are female, 52% are German, 33% Dutch, 6% Belgian, and the remaining 8% have a different nationality.¹² The academic year at Maastricht University is divided into four regular teaching periods of 2 months each and 2 shorter skills periods of 2 weeks each. Therefore, there are six teaching periods per academic year for which we have course outcome information. On average, students take two courses at the same time in the regular periods and one course in the shorter skills periods. The SBE examinations office provided data on student grades, student course dropout, and some basic student characteristics, namely gender, age, and nationality. We also obtained data on students' enrolment duration, course choices, and bachelor graduation to explore some potential medium-run effects of the policy.

The Dutch university grading scale ranges from 1 to 10, with 5.5 usually being the lowest passing grade.¹³ The final course grade is often calculated as the weighted average of multiple graded components, such as the final examination grade, participation grade, presentation grade, and/or mid-term paper grade. The graded components and their respective weights differ by

12. This concretely means that there are overall 336 non-DGBs (non-Dutch/German/Belgian) students for which we observe 4,203 different course outcomes—course dropouts and grades/pass course—over the analysis period. More than two-third of these students, 236, are observed both before and after the introduction of the partial prohibition policy and we have 3,315 course outcomes (1,595 before and 1,720 after) for these individuals.

13. The Dutch university system is very comparable to all other higher education institutions that use the European Credit Transfer System (ECTS). In this system, sixty credits represent one year of study and one credit represents 28 h of study. One year of study comprises 42 weeks. For comparison with other countries, about the usual grading scale, EP-Nuffic, the organization in charge of the internationalization of education in the Netherlands, explains that “the grading system used in the Netherlands is on a scale from 1 (very poor) to 10 (outstanding). The lowest passing grade is 6 (5.5 rounded); 9 s are seldom given and 10 s are extremely rare. Grades 1–3 are hardly ever used.” For further details on this, including a comparison table to grading practices in the U.S. and the U.K. system, please see: <https://www.epnuffic.nl/en/study-and-work-in-holland/dutch-education-system>.

course, with most courses giving most of the weight to the final examination grade. We do not observe the individual components of the final grade separately. If the final course grade of a student after taking the final examination is lower than 5.5 (5 in the first year), the student fails the course and has the possibility to re-take the examination for a second time. We observe final grades after the first and second attempt separately. For our analysis, we only use first attempt grades since the second attempts take place about 2 months later than the original examinations and may not be comparable to the first examinations.¹⁴ From this data, we create three main performance measure outcomes for our analysis: standardized grades, course passing, and course dropout.¹⁵

3.2. Further data sources

3.2.1 Numerical versus non-numerical courses? The literature linking cannabis and cognitive performance has shown that numerical tasks are substantially more affected than non-numerical ones. In order to test this, we had to classify the 177 different courses available to students at the undergraduate level in terms of whether they required numerical skills or not. For this purpose, we looked into the description of every single course, which is publically available online (<http://code.unimaas.nl/>), classifying each as being numerical if the following words appeared in it: *math, mathematics, mathematical, statistics, statistical, theory focused*. This exercise resulted in fifty-six courses being classified as numerical and 121 as non-numerical. As courses requiring numerical skills are more often part of the compulsory curriculum of a degree, we end up with about 35% of course grade observations being categorized as numerical. In Section 6, we split our sample between this numerical and non-numerical course line to test whether we are indeed picking up the effect of cannabis consumption.

3.2.2. Student course evaluations. In addition to the scheduling and grade data, we also obtained data on students' course evaluations, which we match to the grade data using the individual student ID. We use these student course evaluations to provide additional suggestive evidence on some of the channels underlying our results. Two weeks before the examination, students are invited by email to evaluate the courses that they are currently taking in an online questionnaire. Students receive up to three email reminders and the questionnaire closes before the day of the examination. Students are assured that their individual answers will not be passed on to anyone involved in the respective course. Teaching staff receive no information about the evaluation before they have submitted the final course grades to the examination office.¹⁶ The exact length and content of the online questionnaires differ by course, although they typically contain between nineteen to twenty-five closed questions and two open questions. For our analysis, we use the nine core questions that are asked in most courses.¹⁷ These standard questions ask

14. In a robustness check, we estimated our model using "final course grades" and obtained similar results independent of whether we use first or second attempt grades. We prefer using first sit grades throughout as we know the timing of each examination while re-sit examinations can be taken at different times (which we do not always observe) and thus could fall either outside or inside the policy treatment period for certain students.

15. Grades are standardized to have mean zero and unit variance. Course dropouts are defined as students who registered for a course but decided to drop the course at some stage throughout the teaching period, did not fulfill their attendance requirements or did not show up for the final examination. From the data, it is not possible to distinguish between these types of dropouts.

16. This "double blind" procedure is implemented to avoid any of the two parties retaliating from negative feedback with lower grades or evaluations.

17. Table A1 in the Appendix shows the evaluation questions that we tried to group into different mechanism categories and which ones we group together to explore potential channels that explain changes in student performance.

TABLE 2
Student characteristics and education outcomes by nationality groups

	Mean	Standard deviation	Min value	Max value	Number of unique individuals	Total number of observations
Panel A: All students						
Female	0.353	0.478	0	1	4,419	57,019
Age	20.2	1.86	16.2	39.7	4,419	57,019
First sit grade	6.33	1.93	1	10	4,323	51,649
Passed course	0.819	0.385	0	1	4,323	51,649
Course dropout	0.094	0.292	0	1	4,419	57,019
Number of courses	2.01	0.544	1	5	4,419	57,019
Panel B: DGB students						
Female	0.349	0.477	0	1	4,083	52,816
Age	20.2	1.82	16.2	39.7	4,083	52,816
First sit grade	6.38	1.92	1	10	3,998	47,994
Passed course	0.825	0.380	0	1	3,998	47,994
Course dropout	0.091	0.288	0	1	4,083	52,816
Number of courses	2.00	0.536	1	5	4,083	52,816
Panel C: Non-DGB students						
Female	0.394	0.489	0	1	336	4,203
Age	20.3	2.29	16.3	31.3	336	4,203
First sit grade	5.74	2.04	1	10	325	3,655
Passed course	0.743	0.437	0	1	325	3,655
Course dropout	0.130	0.337	0	1	336	4,203
Number of courses	2.07	0.633	1	5	336	4,203

Notes: All the means presented are statistically different between Dutch, German, and Belgian (DGB) and non-DGB students at the 1% significance level.

students to evaluate different course aspects such as teacher performance, group functioning, course material, and general course organization, as well as stating the hours that they spent studying outside of the course. We group these questions into five main categories to explore the underlying mechanism that could explain our results: “hours worked”, “feel stimulated”, “functions well”, “understand better”, and “quality improved”.

3.3. *Student performance descriptive statistics*

Table 2 presents the main descriptive statistics for all students in Panel A, for DGB students in Panel B, and for all other nationality students (non-DGB) in Panel C. The non-DGB students display on average worse performance on all relevant indicators. They are somewhat younger and are more likely to be female than their DGB peers. These differences in terms of characteristics and grades are always statistically significant, which underlines the importance of applying a difference-in-differences approach rather than simply performing a naïve estimation that would not account for these baseline disparities in observable and potentially also unobservable differences.

Figure 4 provides a first visual hint at the existence of an effect of cannabis access restriction on course results. The figure shows course grades for treated and non-treated students over the seventeen time periods that we observe. To capture differences in levels between the two groups, we use two axes. The two vertical lines mark the start and the end of the discrimination policy that affected access of the non-DGB students exclusively (all are banned in the final period). We first note that there is evidence of some grade inflation over this period, as well as substantial cyclicity in examination results within years from one period to the next. Importantly, the examination results of both groups of students clearly trace each other up to the period when the

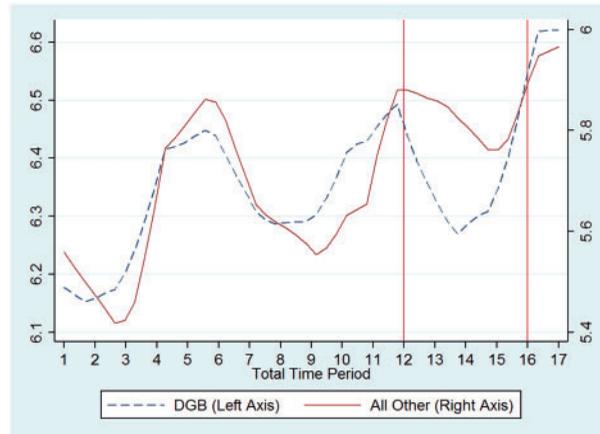


FIGURE 4

Course grades for DGB and all other nationality students

Notes: The solid line represents the grades of the students treated by the cannabis prohibition. The left axis refers to the average examination grades of Dutch–German–Belgian (DGB) students and the right axis refers to the grades of all other nationalities (non-DGB). The two vertical lines denote the start and end of the prohibition period when the “All Other” students had no access to cannabis shops.

non-DGB students are no longer allowed to buy cannabis in cannabis shops. The figure illustrates that the common pre-trend assumption—a necessary condition for our difference-in-differences approach to be valid—is likely to hold. After the policy introduction, non-DGB students appear to suddenly perform substantially better than their DGB peers, which is a first hint that the policy might have had a positive effect on the performance of those who could no longer buy cannabis legally. When Maastricht cannabis shops went on strike at the end of the period that we study, the grades of DGB student (who subsequently also lost access to legal cannabis) went back to trend very quickly. In the next section, we will present our empirical strategy and explain how we identify a causal relationship.

4. EMPIRICAL STRATEGY

In order to estimate the effect of legal cannabis access on student performance, we exploit a unique natural experiment that temporarily discriminated legal access to cannabis based on nationality. We apply a difference-in-differences approach across time and nationality groups. Accordingly, we obtain reduced form estimates of how the policy affects changes in student performance rather than (unavailable) individual changes in student consumption.¹⁸ However, the change in composition of Maastricht’s cannabis shop costumers shown in Table 1—which we discussed in the previous section—provides very strong evidence that the prohibition was de-facto effective and brought to an end to legal cannabis purchases by treated nationalities.¹⁹

18. The effect on performance is perhaps more policy-relevant since changes in cannabis consumption behaviour itself might be economically insignificant if they do not lead to any important negative externalities for society. This reduced-form approach also avoids serious measurement problems with usual measures of drug consumption since we do not have to rely on self-reported consumption or police seizures, which are likely to be correlated with changes in the legal status of this substance.

19. For students who still really wanted to consume cannabis, it might have been possible to obtain illegal access to the drug through peers with a different nationality who were not excluded from cannabis shops or through other illegal

The main outcome variables of interest to measure the impact of the cannabis access policy on student performance are standardized course grades and course passing rates. To test for compositional changes, we also assess whether the course dropout probability is affected. These outcomes are indicated by the dependent variable Y in Equation (1), which describes a simple difference-in-difference model:

$$Y_{it} = \alpha + \beta_1(\text{Nat}_i * \text{Discrim}_t) + \beta_2\text{Nat}_i + \beta_3\text{Discrim}_t + \varepsilon_{it}, \quad (1)$$

Subscript i and t denote, respectively, individual students and the seventeen time periods with a course outcome that we observe. Nat_i is a dummy equal to zero if a student is Dutch, German, or Belgian and equal to 1 if the student is of any other nationality. Discrim_t is a dummy equal to 1 for every period when cannabis access was restricted and zero otherwise. Note that this restriction applies to non-DGB students only from periods 13 to 16 and to all students in period 17 (cf Section 2.3 and Figure 3). The interaction term $(\text{Nat}_i * \text{Discrim}_t)$ enables us to derive an estimate of β_1 , the coefficient of interest of the policy impact. Finally, α is a constant, and ε an error term. To this basic specification we can also add gender and age in months to observe whether adding observable individual characteristics alters the results.²⁰

We can further improve upon this model by gradually adding a number of fixed effects layers to the estimation to account for any potential unobserved course and student heterogeneity. First, we include the total number of courses taken by a student in each period, $NCourses$, and course fixed effects γ_j for the $j=177$ different courses available to students at the bachelor level at the SBE. Second, we also exploit the panel nature of our data and replace the common intercept α with a student-specific fixed effect α_i .²¹ Third, as Figure 4 suggested that student grades were improving over the period studied and that there was some cyclicalality across the six study periods within academic years, our final model will also include study period and year dummies to account for time-varying patterns in student performance. This model is shown in Equation (2):

$$Y_{ijt} = \alpha_i + \beta_1(\text{Nat}_i * \text{Discrim}_t) + \beta_2\text{Discrim}_t + \beta_3\text{Age}_{it} + \beta_4\text{NCourse}_{it} + \gamma_j + \text{Period}_t + \text{Year}_t + \varepsilon_{it}. \quad (2)$$

This within individual estimation approach should take care of all remaining time invariant unobserved individual characteristics that could still not be accounted for in our previous models and could potentially bias our estimates of β_1 . Equation (2) will be our preferred specification to interpret our results concerning the impact of the discriminatory cannabis access policy on student performance measures. Using individual student fixed effects will rule out the possibility that the observed treatment effect of the policy is driven by a change in the student body composition. We later use two modified versions of model (2) to investigate the potential effect of other individuals' treatment status on own performance. We will achieve this by including an interaction between the main policy effect and the proportion of peers in the same class who are treated and a dummy of the teaching staff being DGB or not.

To test the robustness of our results, we will again use model (2) to run placebo tests to check that the estimated effects are indeed causal and not driven by spurious correlations. Our first placebo analysis will be a “placebo in time,” which switches the policy “on” one year

channels. If this was the case, our estimates would subsequently represent lower bounds of the effect of the policy change as we identify the intention-to-treat effect.

20. We will later also perform sub-group analyses along these dimensions to test whether responses to the policy differ depending on individual's observable characteristics.

21. Note that the inclusion of individual fixed effects will no longer make it possible to identify the gender or nationality group effects as students remain of the same sex and nationality throughout the period.

TABLE 3
Impact of restricted cannabis access on student examination scores

	Dependent variable = standardized grades				
	(1)	(2)	(3)	(4)	(5)
Non-DGB*Restriction	0.062* (0.025)	0.061* (0.024)	0.083** (0.031)	0.108** (0.017)	0.109** (0.017)
Non-DGB student	-0.275** (0.090)	-0.280** (0.087)	-0.263* (0.101)	-	-
Restriction period	0.044** (0.012)	0.045** (0.011)	0.042** (0.016)	-0.013* (0.006)	-0.014* (0.006)
Female dummy		0.132** (0.025)	0.108** (0.035)	-	-
Age in months		0.000 (0.000)	-0.002** (0.000)	0.019** (0.002)	0.019** (0.002)
Number of courses			-0.165* (0.065)	-0.054** (0.005)	-0.052** (0.006)
Course fixed effects	No	No	Yes	Yes	Yes
Individual fixed effects	No	No	No	Yes	Yes
Period and year fixed effects	No	No	No	No	Yes
Adjusted <i>R</i> -squared	0.005	0.009	0.151	0.505	0.505
Observations	51,649	51,649	51,649	51,649	51,649

Notes: Robust standard errors clustered at the nationality level reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively.

before it was actually put in place. We also run a “placebo in nationality,” where we consider Belgian students (who are statistically the most similar to non-DGB students) as those with restricted access to cannabis shops. We also obtain distinct policy effects for the numerical and non-numerical courses. Finally, we present further results using course evaluation surveys that follow this within-student difference-in-difference set up.

5. MAIN RESULTS

5.1. Average policy effect

Table 3 reports the estimates of how the policy change affected standardized student grades. We start with the most basic specification of Equation (1) in Column (1) of the table and then successively build up the model with additional controls and fixed effects in Columns (2)–(5). The main coefficient of interest on Nat*Restriction is always positive and statistically significant. The point estimate shows that students who could no longer buy cannabis legally obtained relatively better course grades than those who maintained access during the time when the policy was in place. The coefficients actually become slightly larger as we add more controls, reaching 0.109 of a standard deviation in Column (5) for our most demanding specification, which accounts for unobserved course and individual heterogeneity as well as period- and year-specific effects.

Table 4 reports the same point estimate for grades in the first column and subsequently extends the analysis to two further performance measures: “passing the course” and “course dropout”.²² Changes in the probability of passing a class are important since they indicate whether the grade effect is concentrated at the top or bottom end of the grade distribution. An effect on passing

22. These are discrete outcomes and the OLS used therefore estimates a linear probability model (LPM). We obtain very similar results in terms of marginal effects if we instead use probit models but use LPM throughout given the large number of fixed effects that have to be included in our preferred specification.

TABLE 4
Impact of restricted cannabis access on various measures of academic performance

Educational outcomes	Standardized grade	Passed course	Dropout of course
	(1)	(2)	(3)
Non-DGB * Restriction	0.109** (0.017)	0.040** (0.004)	-0.014 (0.008)
Restriction period	-0.014* (0.006)	-0.006 (0.008)	-0.031** (0.005)
Age in months	0.019** (0.002)	0.008** (0.001)	-0.001 (0.000)
Number of courses	-0.052** (0.006)	-0.006 (0.004)	0.022** (0.003)
Mean of outcome	NA	0.739	0.143
Effect size at mean	NA	0.054	-0.097
All controls and FEs	Yes	Yes	Yes
Adjusted <i>R</i> -squared	0.503	0.315	0.312
Observations	51,649	51,649	57,019

Notes: All specifications include the same fixed effects and controls as in the last column of Table 3 (*i.e.* age in months, number of courses enrolled in, teaching period and year dummies, course-specific fixed effects, and student-specific fixed effects). Robust standard errors clustered at the nationality level are reported in parentheses. * and ** indicate significance at the 5% and 1% level, respectively.

probabilities might be economically more important than changes in grades since students who fail classes have to re-take the examination or course at a later time, which may result in delayed graduation or lead to failing to obtain a degree. We find a 4 percentage point increase in pass rates for non-DGB students when the policy is in place, a 5.4% improvement from the baseline pass rate of 73.9%. The coefficient on the probability of dropping out is small and not statistically significant. This is an important result since it indicates that treated individuals are as likely to complete courses during the policy period as before. It also simplifies the interpretation of our results as we can reject compositional effects that could arise if we would not observe the performance of the same individuals across time.

5.2. Sensitivity by sub-groups

One way to gain a better sense of where legal access to cannabis really “bites” is to consider differences in the policy impact on the outcomes of different population sub-groups. When interpreting coefficients for different sub-groups, it is important to keep in mind that these may not only differ in their baseline propensity to consume cannabis but also in their response and compliance to the policy. Table 5 reports estimates for the sample split by gender, age, and performance level in Panels A, B, and C, respectively. It also shows the sample size, number of unique individuals, and percentage treated for each of the sub-sample populations²³

A first intriguing finding in Table 5, Panel A is that the course grade effect seems to be stronger for female students (0.130 compared with 0.094 of a standard deviation). However, this difference

23. We have also estimated coefficients for these sub-groups using interactions of the difference-in-differences effects with dummy indicators for gender, age (younger), and performance (high achiever). Adding up the main policy coefficient to the extra sub-group effect estimated gave us almost exactly the same statistically significant point coefficients as when using the split-sample approach. Since we have enough statistical power in these split regressions we decided to report these results instead of triple interaction variant as they are much easier to interpret.

TABLE 5
Results by sub-groups

	Standardized grades	Passed course	Sample size Number of students Percentage treated
Panel A: Gender			
Male students	0.093** (0.018)	0.036** (0.005)	Number of observations = 32,968 Number of individuals = 2,861 % treated = 0.069
Female students	0.130** (0.031)	0.045** (0.008)	Number of observations = 18,681 Number of individuals = 1,558 % treated = 0.082
Panel B: Age			
Younger students	0.135** (0.019)	0.061** (0.006)	Number of observations = 25,961 Number of individuals = 2,300 % treated = 0.081
Older students	0.053 (0.034)	0.004 (0.013)	Number of observations = 25,733 Number of individuals = 2,520 % treated = 0.086
Panel C: Performance			
Low performers	0.129** (0.026)	0.063** (0.008)	Number of observations = 25,665 Number of individuals = 2,164 % treated = 0.089
High performers	0.094** (0.008)	0.012* (0.009)	Number of observations = 25,984 Number of individuals = 2,159 % treated = 0.062

Notes: Table reports coefficients on non-DGB* restriction for the same specification as Column (5) of Table 3 for each sub-group. Robust standard errors clustered at the nationality level reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively. For age, the sample is split between below and above the median age when last observed: 20.69 years. For performance, the sample is split between students below and above the median average examination score in the period before the introduction of the policy.

is relatively smaller for the probability of passing courses when the higher female baseline passing rate is taken into account. To rationalize this gender differences, one could consider previous evidence on differences in responses to legal status of substances across genders (Pacula, 1997) or the possible stronger residual effects of cannabis consumption on female test performance (Pope *et al.*, 1997). In our case, it is also probable that the marginal young women are more likely to comply with the legal change than you men and not switch to the illegal street market, where it might be a different experience to purchase drugs illegally compared with the previously legal cannabis shops.

The age sample split across the median age of 20.7 years (when the individual was last observed) in Panel B reveals that all of the detected policy impact comes from relatively younger students. As age almost perfectly maps with year of study in the three-year bachelor degree, this indicates that the performance improvements for no-access nationalities are only present in the first or second year of enrolment. This is indicative of a maturity effect, with individuals above a certain age threshold not changing consumption behaviour as a result of the cannabis prohibition. Another possible factor is that these individuals are in the third year of their degree and have mostly established networks of DGB students who could illegally provide them with cannabis if necessary.

Next, we test whether low performers—defined as students with a pre-treatment GPA below the median in Panel C—are affected differently than high performers (*i.e.* those above median pre-treatment GPA). We find that the cannabis ban has a significant effect on the grades of high

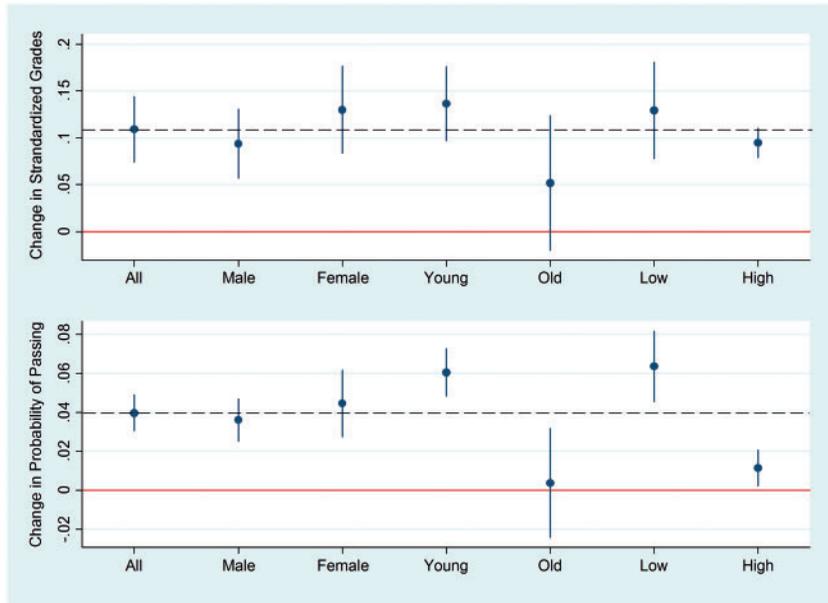


FIGURE 5

Main specification—point estimates for different sub-groups

Notes: This figure visualizes the point estimates and 95% confidence intervals of our main specification for different sub-groups of students. This figure is based on the estimation shown in Table 7. The horizontal dashed line marks the estimated effect size for the full sample and the horizontal solid line marks the zero or no policy effect. The categories “Young” and “Old” refer to below and above median age. The categories “Low” and “High” refer to below and above the median GPA.

performers but very little impact on their probability of passing a course, which is unsurprising as they already have an average passing rate of 96.5%. However, for low performers the grade effect is larger and crucially the policy also very strongly changed their likelihood of passing courses: a 6.4 percentage point increase from a relatively low baseline passing rate of 60%. This is not only a substantial 10% increase but also very policy-relevant since the affected sub-group comprises a student population that is more likely to drop out of university and take longer to graduate compared with the other students.

Figure 5 represents a graphical illustration of all the estimated coefficients that we reported in Table 5. To enable a simple visual comparison of how the sub-group effects differ from the main effect (dashed horizontal line), we show the point estimates and 95% confidence intervals for each of the sub-groups separately.

5.3. Spillovers from peers and teachers

In order to test for the presence of some social multiplier effects of the policy change and potentially affected university instructors, we also test whether classroom peer composition and teacher nationality had an impact on student performance.

To assess the effect of treated peers, we create a variable that calculates the fraction of other treated non-DGB students (from 0 to 1) in each teaching group within each course.²⁴ To test

24. Courses at Maastricht University are organized in multiple teaching sections called “tutorials”. One section usually contains about ten to fifteen students. Within courses, students are randomly assigned to sections (see Feld and Zölitz, 2017).

TABLE 6
Effect of share treated in class and nationality in class teacher

	Peer effects		Teacher effects	
	(1)	(2)	(3)	(4)
	Std. grade	Passed course	Std. grade	Passed course
Non-DGB * restriction period * share no-access nationality	0.172 (0.129)	0.209* (0.081)	–	–
Non-DGB * restriction period * nationality of class teacher	–	–	–0.023 (0.029)	–0.007 (0.014)
Non-DGB * restriction period (i.e. main policy effect)	0.093** (0.024)	0.021* (0.009)	0.122** (0.033)	0.052** (0.012)
All controls and FEs	Yes	Yes	Yes	Yes
Adjusted R-squared	0.505	0.316	0.504	0.320
Observations	51,620	51,620	34,897	34,897

Notes: The controls and FEs included in all specifications are as in the last column of Table 3 (i.e. age in months, number of courses enrolled in, teaching period and year dummies, course-specific fixed effects, and student-specific fixed effects). Additional included controls are “Share no-access nationality”, “Restriction time periods”, and “Restriction time periods * share no-access nationality”. Robust standard errors clustered at the nationality level are reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively.

whether the classroom composition during the time of restricted cannabis access had an impact on own performance, we interact the basic policy effect coefficient with the fraction of treated students in each section. This measure should capture the “extra” effect on performance of having more or fewer peers with cannabis access. The interaction and the main policy coefficient are reported in the first two columns of Table 6 for standardized grades and the probability of passing a course. The estimated impact of peer composition on grades is insignificant, although we detect a marginal improvement in passing rates as the fraction of treated peers in the section increases. The coefficient of 0.21 means that a 10% increase in the number of classmates who no longer have access to cannabis shops increases the chance of passing by 2 percentage points. Interestingly, this spillover effect only exists for students who were themselves affected by the policy change, which might reflect patterns of social interaction along nationality lines within and outside the classroom.

We also test whether student results improved because their section instructors now performed better due to their own cannabis access being restricted. In the administrative data that we obtained, we can observe the nationality of section instructors if it is a PhD student teaching the class. We use this information to form the same nationality groups (DGB versus non-DGB) that we applied for students and test whether student performance in those classes was affected by the treatment status of the teacher.²⁵ The last two columns of Table 6 report the interaction of this dummy with the main policy effect and the main difference-in-differences coefficient itself. We find no evidence of a teacher treatment effect, which is perhaps unsurprising considering that results by age group had already indicated that the performance of relatively older students was not affected by the drug policy change.

25. About one-third of the university instructors have a non-DGB nationality.

TABLE 7
Placebo in policy timing and treated group

	Placebo specification			
	Policy 1 year earlier		Belgians treated group	
	Std. grade	Passed	Std. grade	Passed
Placebo policy effect	-0.054 (0.057)	-0.014 (0.026)	0.023 (0.050)	0.028 (0.023)
All controls and FEs	Yes	Yes	Yes	Yes
Adjusted <i>R</i> -squared	0.522	0.328	0.500	0.308
Observations	33,533	33,498	47,994	47,994

Notes: Robust standard errors clustered at the nationality level are reported in parentheses. * and ** indicate significance at the 5% and 1% level, respectively. The placebos report coefficients on non-DGB * restriction for the same specification as Column (5) of Table 3 when, respectively, the time period for treatment is changed to -1 year, and the group treated is changed to Belgians.

5.4. *Time and nationality placebos*

We report the results from two falsification exercises in Table 7 that test for a potential non-policy-related impact on student performance if we change the time of its introduction or the nationality of the individuals treated.

For the first falsification test, we generate a placebo policy by estimating Equation (2) with the treatment period artificially placed one year earlier than when cannabis access restriction was actually introduced (dropping the policy period from the sample). The coefficients on both grades and course passing for this “placebo in time” are very small and statistically insignificant. This confirms that we were not picking up some period-specific effect unaccounted for in our previous specifications.

Next, we consider Belgian students (instead of non-DGB) as those who are prohibited from entering and buying cannabis at cannabis shops (dropping the other non-Dutch nationalities actually treated from the sample), given that students from Belgium are the closest in terms of observable characteristics to the treated non-DGB. Again, in this second falsification test, the coefficients on both measures of performance of this “placebo in nationality” are small and non-significant. This further supports our claim that it is the policy limiting legal cannabis access and not another unobserved event affecting certain types of students during this period that improved student performance in the short run.

6. MECHANISMS DRIVING THE FINDINGS

Our results quite clearly show that students who lost the right to buy cannabis legally experienced important performance improvements relative to their peers who could still enter cannabis shops. Results from the sub-group analysis further reveal that these effects are mostly driven by those individuals whom we would expect to be affected by the temporary cannabis prohibition. In what follows, we conduct two additional exercises with the administrative data that we have available to test whether our findings are consistent with the particular manner in which THC consumption affects cognitive functioning.

6.1. *Numerical versus non-numerical courses*

We first propose a very simple extension to our analysis of the student performance data inspired by Block and Ghoneim (1993) and Pacula *et al.* (2003), who find that numerical skills are more

TABLE 8
Differences between courses requiring more and less numerical skills

	Standardized grades		Passed course	
	Numerical	Non-numerical	Numerical	Non-numerical
	(1)	(2)	(3)	(4)
No-access nationality *	0.227**	0.065**	0.073**	0.025**
Restriction time periods	(0.028)	(0.015)	(0.010)	(0.004)
Restriction time periods	-0.183**	0.066**	-0.055**	0.015**
	(0.030)	(0.014)	(0.019)	(0.005)
All controls and FE	Yes	Yes	Yes	Yes
Mean of outcome	NA	NA	0.663	0.785
Effect size	NA	NA	0.110	0.032
Adjusted <i>R</i> -squared	0.578	0.467	0.363	0.286
Observations	18,092	33,557	18,092	33,557

Notes: All courses available to students at the undergraduate level were classified on whether they required math/numerical skills or not based on their online course descriptions. These were classified as “Numerical” if the following words appeared in this description (and “Non-numerical” otherwise): math, mathematics, mathematical, statistics, statistical, theory focused. Robust standard errors clustered at the nationality level are reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively.

impaired by cannabis use than non-numerical skills. Consequently, if the increase in performance detected is more pronounced for courses that require more skills in mathematics or statistics, we can more confidently attribute it to a change in students’ cannabis consumption. We should not expect to observe such a disparity in effects if the results were driven by a change in alcohol consumption caused by the policy change. If students reduce (or increase) their alcohol consumption owing to complementarities (substitution) between cannabis and alcohol we would expect numerical and non-numerical courses to be affected in a similar way. This has recently been confirmed by Carrell *et al.* (2011), who show that access to alcohol and its consumption affect both numerical and non-numerical skills equally. Apart from cannabis use, it is very difficult to come up with any other plausible explanation why performance in these two types of courses would be affected in a systematically differential way.

Table 8 reports results for our main specification split by the numerical versus non-numerical categorization of courses.²⁶ The dependent variable in the first two columns is the standardized course grade. The estimates reveal that the policy effect is about 3.5 times larger for numerical rather than non-numerical courses. Since there might also be differences in the average difficulty of courses that require more or less numerical skills driving the grade differences, we also estimated the effects on passing rates. These are reported in Columns (3) and (4) of Table 8, confirming that numerical courses are on average more difficult: only two-thirds of students pass these compared with the almost 79% passing non-numerical courses on average. Despite these baseline differences, the difference in the estimated policy effect remains much stronger for the probability of passing math-oriented courses (11%), which is still 3.6 times that of passing non-mathematical courses (3.2%). This difference is statistically significant and a strong indicator that the improvement in performance that we observe is driven by non-DGB students altering their cannabis consumption as a result of the changes in the legal access to cannabis.

26. For further details regarding the classification of courses, see our description and discussion in Section 3.2.1.

TABLE 9
Exploration of potential channels—student course evaluations

Mechanism categories	Non-DGB *Restriction	Mean (non-standardized)	Adjusted R^2	Number observations
Hours worked	0.022 (0.019)	13.1	0.510	15,987
Feel stimulated	0.057* (0.026)	7.1	0.268	15,937
Functions well	0.041 (0.025)	7.8	0.176	15,997
Understand better	0.215** (0.027)	7.1	0.342	13,520
Quality improved	0.137** (0.025)	7.8	0.267	17,546

Notes: All specifications include student fixed and course effects, teaching period and year dummies. See Table A1 in the Online Appendix for details about the original questions and how we categorize them into the five main mechanisms reported here. All questions, except Hours worked, were standardized to mean zero and unit variance, then averaged within each mechanism category and again standardized to mean zero and unit variance. Robust standard errors clustered at the nationality level are reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively.

6.2. Evidence from student evaluations

We now exploit additional data from students' online course evaluation surveys which they are asked to complete at the end of every course. The participation rate for student course evaluations is not very high with 37% out of all surveys requests sent out being completed, although at least two-thirds of students respond to at least one of these questionnaires. In any case, since we investigate within-individual changes using student fixed effects, we do not believe that selection into survey response is a serious threat to the interpretation of our results, given that identification will only come from those who have answered multiple times and at least once before and once after the policy change, which is the case for over half of the observed students. We match the evaluation data to students' nationality and course grades at the individual level. For the analysis, we grouped the nine most common survey questions into five potential mechanism categories: "Hours worked", "Feel Stimulated", "Functions Well", "Understand Better", and "Quality Improved".²⁷

Table 9 reports the coefficients of the estimated difference-in-differences policy effect on each of the potential mechanisms. A first observation is that the effect on hours per week spent studying for a course outside of the classroom is extremely small and not statistically significant. This suggests that changes in the study effort of students is not the main driver of our results and that the performance increase that we observe is not driven by changes in students' time use outside the classroom. The most significant change is an almost 0.22% of a standard deviation increase in the student reported subjective understanding of course material and lectures after the policy introduction. There is also a—perhaps unsurprising—significant increase in the perceived overall quality of the courses/teacher following from this improved understanding. These underlying channels would be consistent with clinical evidence suggesting that the main effect of cannabis on human functioning is worsening one's memory of things learned while "high", or as Ranganathan and D'Souza (2006) put it in their review of the clinical literature: "THC ... impairs immediate and delayed free recall of information presented after, but not before, drug administration." We take the fact that treated students report improvements with

27. Table A1 in the Appendix provides more details and descriptive statistics on the original survey questions and shows how they are grouped together to form the mechanism categories we investigated.

TABLE 10
Longer-run student performance effects: final GPA; graduation; and proportion math electives

	Final GPA			Graduation			Proportion math electives		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Non-DGB student *	-0.120			-0.009			-0.032		
Restriction in 1st year of BA	(0.152)			(0.058)			(0.025)		
Non-DGB student *		0.058			-0.007			0.058*	
Restriction in 2nd year of BA		(0.163)			(0.062)			(0.027)	
Non-DGB student *			0.077			0.016			-0.020
Restriction in 3rd year of BA			(0.161)			(0.061)			(0.026)
Non-DGB student	-0.534**	-0.598**	-0.604**	-0.164*	-0.165*	-0.173**	0.021	-0.008	0.015
	(0.096)	(0.089)	(0.090)	(0.036)	(0.034)	(0.034)	(0.016)	(0.015)	(0.015)
Female	0.282**	0.282**	0.282**	0.080*	0.080*	0.080**	-0.037**	-0.037*	-0.037*
	(0.041)	(0.041)	(0.041)	(0.016)	(0.016)	(0.016)	(0.007)	(0.007)	(0.007)
Age	0.008	0.008	0.008	0.014*	0.014*	0.014**	0.001	0.001	0.001
	(0.011)	(0.011)	(0.011)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)	(0.002)
Study cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.107	0.107	0.107	0.019	0.019	0.019	0.006	0.007	0.006
Observations/individuals	4,258	4,258	4,258	4,415	4,415	4,415	4,415	4,415	4,415

Notes: Robust standard errors clustered at the nationality level are reported in parenthesis. * and ** indicate significance at the 5% and 1% level, respectively. Note that the respective reference group changes depending in which year of the Bachelor the non-DGB students were enrolled in when they were exposed to the restriction policy. The reference group in Columns (1), (4), and (7) consists of all DGB student plus non-DGB students who enrolled in 2009 or 2010. The reference group in Columns (2), (5), and (6) consist of all DGB student plus non-DGB students who enrolled in 2009 or 2011. The reference group in Columns (3), (6), and (9) consist of all DGB student plus non-DGB students who enrolled in 2010 or 2011.

respect to the understanding as additional suggestive evidence that the observed performance improvements indeed stem from a decrease in cannabis consumption caused by the legal access restriction policy.

7. PERSISTENCE OF EFFECTS AND INTERPRETATION OF FINDINGS

7.1. *Persistence of effects: longer-run performance*

Although our results support the presence a short-run effect of the policy change on student performance, we also want to explore whether the restriction in cannabis access had any persistent medium- or longer-run effects on the treated students. We analyse this by following differently treated cohorts of DGB and non-DGB students through their first, second, and third year at Maastricht University. The student outcomes that we are able to consider in this context are final GPA, graduation probability, and the proportion of elective courses chosen that have mathematical/numerical skills requirements. Since these outcomes do not vary over time at the student level, we cannot include course, period, or individual fixed effects, unlike in all previous specification. This implies that the variation that we can exploit for identification “only” stems from *across-cohort variation* between nationalities and not from *within-individual* changes in outcomes. Therefore, we consider that the following results should be interpreted with some caution as the identifying assumptions are much less restrictive than in our previous estimated models.

Table 10 reports cohort level difference-in-difference estimates for changes in the longer-run outcomes of non-DGB students depending upon the academic year in which they were exposed to the policy (*i.e.* in which year they were enrolled in 2011–2012) relative to all other students. These inform us on whether the students who experienced restricted cannabis in either the 1st (Columns (1), (4), and (7)), 2nd (Columns (1), (5), and (8)), or 3rd (Columns (3), (6), and (9)) year of their bachelor exhibit different outcomes by the end of their studies compared with never treated DGB

student cohorts and non-DGB student cohorts treated in another year.²⁸ The first six columns of Table 10 reveal that neither the final GPA nor the graduation probability of students who were either treated in their first, second, or third year improves relative to the other students. All point estimates are small, close to zero and statistically insignificant.²⁹ However, in Columns (7)–(9), we observe that students affected by the policy change in their second year of study—which is the time when students make their elective choices—select a higher proportion of electives with math-related content.³⁰ The estimated impact, although relatively small, is positive and statistically significant at the 5% level. This is consistent with our previous finding that performance gains were mostly driven by courses that require more numerical/mathematical skills. Despite the perhaps suggestive nature of this finding, it could imply that students—who obtain better grades in courses with math content after they no longer have access to legal cannabis—update their beliefs about the type of courses that they will perform well in and consequently choose more courses with math content. Taken together, these results present some tentative evidence of a small yet significant effect of the policy on this medium-run outcome.

7.2. Interpretation of findings

7.2.1. Relative size of estimated effect. The main finding from our most restrictive specification—which uses both student fixed effects and course fixed effects—shows that the temporary restriction of legal cannabis access increased performance on average by 0.109 standard deviations and raised the probability of passing a course by 5.4% (Columns 1 and 2 of Table 4). These point estimates suggest that restricting legal access to cannabis resulted in a substantial increase in student performance. To assess the relative size of such an effect, it is perhaps useful to put it in perspective with other treatments known to affect the performance of college students, particularly including the effect of legal alcohol access.

Our reduced form estimates of the short-run effect on performance are roughly the same size as the effect as having a professor whose quality is one standard deviation above the mean (Carrell and West, 2010) or the effect of being taught by a non-tenure track faculty member (Figlio *et al.*, 2013). It is about twice as large as having a same gender instructor (Hoffmann and Oreopoulos, 2009) and of similar size as having a roommate with a one standard deviation higher GPA (Sacerdote, 2001). The effect of the cannabis prohibition that we find is slightly smaller than the effect of starting school 1 h later and thus being less sleep-deprived (Carrell *et al.*, 2011).

28. Note that as the study year cohort of treatment changes, the control group used also changes and includes not only cohorts of DGB students never treated during their bachelor but also non-DGB students in other cohorts who were treated during another study year. Concretely this means that the reference group in Columns (1), (4), and (7) consists of all DGB student plus non-DGB students who enrolled in 2009 or 2010. The reference group in Columns (2), (5), and (6) consists of all DGB student plus non-DGB students who enrolled in 2009 or 2011. The reference group in Columns (3), (6), and (9) consists of all DGB student plus non-DGB students who enrolled in 2010 or 2011.

29. The fact that we do not detect any effect on the first two longer-run measures of performance, despite the strong direct effects identified, is perhaps not surprising for two main reasons. First, the cohort-level approach adopted here yields only one outcome per student and the policy will only have affected less than a third of all grades over the whole seventeen periods used in the analysis. Second, the students affected by the partial prohibition in their first or second year (when most of the effect on almost all the effect was detected, see Section 5.2) were then treated by a different restriction policy no longer dependent on nationality (the “wietpas”, see Section 2.3), which will have affected student consumption of cannabis in a very uncertain way. By contrast, the longer-run policy effect on choosing specific courses at the end of the second bachelor year is thus perhaps more likely to be detected in this context.

30. In Column (7) of Table 10, we would perhaps not expect any effect for students in their first year since these students cannot chose any elective courses yet and are required to take a curriculum of first year compulsory courses. Students in their third year (Table 10; Column (7)) have already made their elective courses choices in the past and thus are unlikely to be affected by the policy that restricted access.

A perhaps more relevant benchmark for the comparison of our reduced form estimates is in relation to recent findings concerning how legal alcohol access has been found to impair college achievement. Lindo *et al.* (2013) use an identification strategy akin to ours and show that legal alcohol access reduces course grades by 0.033–0.097 standard deviations when including student fixed effects. Exploiting a discontinuity in the legal drinking age for students at the United States Air Force Academy (USAFA), Carrell *et al.* (2011) estimate that alcohol access causes course grades to drop on average by 0.092 standard deviations. This is remarkably close to the impact of legal cannabis access that we estimate here. The reduced form point estimates of both of these studies suggest that the legal status of cannabis affects overall student achievement in a similar way to the legal status of alcohol.³¹

7.2.2. Treatment effect on the treated and price concerns. The final policy-relevant exercise that we attempt here is to interpret our results in view of the proportion of individuals who actually responded to the change in legal status of cannabis, that is, the treatment effect on the treated. The first step towards understanding the underlying treatment effect on the treated is to gain an idea of baseline consumption rates for the particular group of individuals who were affected by the policy. To obtain rough estimates of these rates, we carried out an anonymous survey among currently enrolled students at Maastricht University.³² To make the question about cannabis consumption less salient, we embedded it in a more general questionnaire on risky behaviour. Overall, 206 students answered the various surveys, which was over 97% of the students present in the various lectures where it was distributed. The survey question that we focused on asks students if they “have ever smoked cannabis or hashish” and—if so—whether they have done so “in the last 12 months”, “in the last 30 days”, or “in the last 7 days”.³³

Interestingly, the baseline consumption rates that we obtain from the survey are very similar across the treated and non-treated population, with 59% of non-DGB and 61% of DGB students reporting having smoked cannabis in the past year (detailed results available in Table A2 of the Appendix). We can consider these individuals as the potentially treated group, as the others are unlikely to change a behaviour that they did not participate in before the prohibition. Assuming full compliance to the policy, the treatment effect on the treated would be about 0.19 standard deviations ($=0.109/0.59$) in course grades and a 9.2% increase in the pass rate ($=5.4/0.59$). Taking a perhaps more reasonable assumption of a 38% compliance rate to the cannabis prohibition (using the Jacobi and Sovinsky [2016] estimate in consumption change among individuals under 30 years from removing accessibility barriers) would roughly translate into a policy impact on the treated of a 0.49 of a standard deviation’s improvement in course grades ($=0.109/0.598/0.38$) and an 18.6% increase in the pass rate of potential cannabis consumers

31. Our results suggest that most of the legal access change effect comes from younger students (see Section 5.2) but are perhaps not applicable to the debate in some countries, including the Netherlands, about whether an age restriction of 21 would be effective in removing some of the negative spillover effects of cannabis use. The main reason is that these larger findings for younger students are likely to partially stem from the potential stronger network effects for older students who may find it easier to keep on purchasing cannabis via non-banned friends during the partial prohibition.

32. Although these are different students to those for which we have performance data that we use in the rest of the analysis, their baseline consumption rates are relevant for two reasons: first, their demographic characteristics (age, gender, and nationality) are extremely similar to the students we previously studied; and second, since the discriminatory policy was no longer in place at the time we conducted our survey, they enjoy the same legal access to cannabis as the Dutch, German, and Belgian students as only some proof of residence was needed to enter coffee shops when the survey was conducted.

33. Despite the fact that we guaranteed strict anonymity, it is still possible that the baseline consumption rates obtained from this survey may *underestimate* the baseline consumption rates since students may not report honestly and understate their consumption levels while sitting next to their peers and in front of a guest lecturer.

(= 5.4/0.58/0.38). Even if this treatment effect on the treated is somewhat overestimated due to student under-reporting baseline consumption or de-facto higher compliance rates, the effects that we identify here are large and economically significant.

One potential remaining concern for the interpretation of our findings is that the drug access limitation may have had an effect on cannabis prices. As the partial prohibition reduced demand, one could expect prices to have decreased during this period. In turn, this could have led to an increase in cannabis consumption for the nationalities who are still allowed to buy the substance legally (income effect). In this case, our results would thus overstate the true policy effect and capture the aggregate effect of non-DGB smoking less owing to prohibition *and* DGBs smoking more due to the drop in the legal price. To rule out this mechanism, we collected prices for 10 types of cannabis strains sold in five of the most popular coffee shops in Maastricht around the time of the policy introduction. We extracted this information from historical postings in online forums where cannabis consumers exchanged information on prices and the “perceived” quality of different cannabis products. An average price per gram calculated from this data was found to be €9.60 before and €9.70 during the period of restrictive legal access. This suggests that the legal cannabis price was unaffected by the introduction of the policy and thus is not a factor affecting the interpretation of our results.

8. CONCLUSION

In this paper, we have investigated how restricting cannabis access affects student achievements, finding that the performance of students who lose legal access to cannabis substantially improves. Our analysis of underlying channels suggests that the effects are specifically driven by an improvement in numerical skills, which existing literature has found to be particularly impaired by cannabis consumption. This article provides the first causal evidence that restricting legal access to cannabis affects college students’ short-term study performance. We believe that our findings also imply that individuals change their consumption behaviour when the legal status of a drug changes.

We must note here that this paper only assesses the impact on one particular outcome for a specific group of individuals. The impact on examinations that require skills in math and statistics might be different from the effects on individuals in environments where performance requires different skills or is measured differently. Our estimates perhaps represent an upper bound because the THC concentration in Dutch cannabis is relatively high compared with that of the strength in products available in most other countries. However, it could also be argued that our estimates are lower bounds because the policy that we study did not restrict access to all students who study in Maastricht, and it may have been possible to obtain illegal access to the drug through peers with different nationalities who were not excluded from cannabis shops or through other illegal channels. From the results of this article, it is unclear whether restricting cannabis may have other severe negative consequences on—for example—crime, since it is likely to increase demand through illegal channels. It should also be noted that it is not obvious from our results whether the effects of legalization and prohibition are symmetric.

After taking these caveats into account, we maintain that our findings have potentially important policy implications for countries that are considering relaxing drug laws. Observing that student achievement is affected by cannabis regulations is perhaps more policy-relevant than documenting changes in cannabis consumption itself, since it might be irrelevant how much cannabis individuals smoke if it does not lead to important negative externalities for society. The effects that we estimate and the change in consumption behaviour that they imply should thus be taken into account along with other pro and con arguments of drug legalization. Accordingly,

these new findings should become integrated in the complex and multi-dimensional societal cost-benefit analysis that should drive any drug policy decision-making.

APPENDIX

TABLE A1
Student course evaluation questions

Nr.	Question wording	Answering scale	Mean [Number of observations]	Categorization
1	How many hours per week on the average (excluding contact hours) did you spend on self-study (presentations, cases, assignments, studying literature, etc.)?	Open question (0–70 h)	13.1 [15,987]	Hours worked
2	The learning materials stimulated discussion with my fellow students.	1–5	3.5 [16,005]	Feel stimulated
3	The learning materials stimulated me to start and keep on studying.	1–5	3.6 [16,010]	Feel stimulated
4	Evaluate the overall functioning of your tutor in this course with a grade.	1–10	7.7 [16,121]	Functions well
5	My tutorial group has functioned well.	1–5	4.0 [16,231]	Functions well
6	The lectures contributed to a better understanding of the subject matter of this course.	1–5	3.1 [13,600]	Understand better
7	Working in tutorial groups with my fellow students helped me to better understand the subject matters of this course.	1–5	4.0 [16,118]	Understand better
8	The tutor sufficiently mastered the course content.	1–5	4.3 [16,135]	Quality improved
9	Please give an overall grade for the quality of this course.	1–10	7.1 [17,546]	Quality improved

TABLE A2
Smoking Propensity of DGB and Non-DGB Students

Latest cannabis consumption	DGB		Non-DGB	
	%	Cumulative	%	Cumulative
Last 7 days	25.00	25.00	21.75	21.75
Last 30 days	10.63	35.63	13.04	34.79
Last 12 month	25.00	60.63	23.91	58.70
Over 12 months	18.75	79.39	17.39	76.09
Never	20.63	100.00	23.91	100.00

Notes: Author's estimation from survey on smoking behaviour from 206 bachelor students (160 DBG and 46 non-DGB) who were enrolled in either first, second, or third year classes in May and December 2014 and May 2015.

Acknowledgments. We thank the editor and three anonymous referees for useful suggestions and comments. We also thank Mark Anderson, Dan Black, Daniel Hammerness, Brian Jacob, Paul Lemmens, Daniel Mejia, Rosalie Pacula, Daniel Rees, Nicolas Salamanca, participants at the 2014 NBER-SI, the 4th AL CAPONE meeting, the SEA annual meeting and seminar participants at CPB, the Hague, IZA Bonn, NED Amsterdam, Northwestern University, RAND and the University of Chicago for useful comments and suggestions. We further thank Joël Castermans, Sanne Klasen, and Kim Schippers from the SBE Scheduling Department, Sylvie Kersten from the SBE Exams Office, and Paul Jacobs from the Educational Research and Development Department for providing data and valuable background information. We thank Raginee Baruah, Hugo Nóbrega, and Sophia Wagner for excellent research assistance. Marie is grateful for financial support for this research from the Netherlands Organization for Scientific Research (NWO Veni 451.12.005).

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- ADDA, J., MCCONNELL, B. and RASUL, I. (2014), "Crime and the Depenalization of Cannabis Possession: Evidence from a Policing Experiment", *Journal of Political Economy*, **122**, 1130–1202.
- ANDERSON, D. M., HANSEN, B. and REES, D. I. (2013), "Medical Marijuana Laws, Traffic Fatalities, and Alcohol Consumption", *Journal of Law and Economics*, **56**, 333–369.
- BIELMAN, B., NIJKAMP, R., REIMER, J. *et al.* (2013), *Coffeeshops in Nederland 2012*, WODC Report, IntraVal, Groningen-Rotterdam.
- BLOCK, R. I. and GHONEIM, M. M. (1993), "Effects of Chronic Cannabis Use on Human Cognition", *Psychopharmacology*, **110**, 219–228.
- BOSSONG, M. G., JANSMA, J. M., VAN HELL, H. H. *et al.* (2012a), "Effects of δ 9-Tetrahydrocannabinol on Human Working Memory Function", *Biological Psychiatry*, **71**, 693–699.
- BOSSONG, M. G., JANSMA, J. M., VAN HELL, H. H. *et al.* (2012b), "Default Mode Network in the Effects of δ 9-Tetrahydrocannabinol (THC) on Human Executive Function", *PLoS One*, **8**, e70074.
- CARRELL, S. E. and WEST, J. E. (2010), "Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors", *Journal of Political Economy*, **118**, 409–432.
- CARRELL, S. E., MAGHAKIAN, T., and WEST, J. E. (2011), "A's from Zzzz's? The Causal Effect of School Start Time on the Academic Achievement of Adolescents", *American Economic Journal: Economic Policy*, **3**, 62–81.
- CARRELL, S. E., HOEKSTRA, M. and WEST, J. E. (2011), "Does Drinking Impair College Performance? Evidence from a Regression Discontinuity Approach", *Journal of Public Economics*, **95**, 54–62.
- COBB-CLARK, D. A., KASSENBOEHMER, S. C., LE, T. *et al.* (2016), "High-School: The Relationship between Early Marijuana Use and Educational Outcomes", *Economic Record*, forthcoming.
- CREAN, R. D., CRANE, N. A., and MASON, B. J. (2011), "An Evidence Based Review of Acute and Long-Term Effects of Cannabis Use on Executive Cognitive Functioning", *Journal of Addiction Medicine*, **5**, 1–8.
- DONOHUE, J. J. III. (2013), "Drug Prohibition and Its Alternatives", in Cook, P. J., Machin, S., Marie, O. and Mastrobuoni, G. (eds), *Lessons from the Economics of Crime: What Reduces Offending?* (Cambridge, MA: MIT Press).
- FELD, J. and ZÖLITZ, U. (2017), "Understanding Peer Effects—On the Nature, Estimation and Channels of Peer Effects", *Journal of Labor Economics*, **35**, 387–427.
- FIGLIO, D. N., SCHAPIRO, M. O. and SOTER, K. B. (2013), "Are Tenure Track Professors Better Teachers?" (No. w19406). National Bureau of Economic Research.
- HALL, W. (2015), "What Has Research over the Past Two Decades Revealed about the Adverse Health Effects of Recreational Cannabis Use?", *Addiction*, **110**, 19–35.
- HOFFMANN, F. and OREOPOULOS, P. (2009), "A professor like me the influence of instructor gender on college achievement", *Journal of Human Resources*, **44**, 479–494.
- JACOBI, L. and SOVINSKY, M. (2016), "Cannabis on Main Street? Estimating the Demand in Markets with Limited Access", *American Economic Review*, **106**, 2009–2045.
- LAFRATE, A. and ARMENTANO, P. (2015), Presentation at the Annual Meeting of the American Chemical Society, Denver, Colorado, March 23, 2015.
- LINDO, J. M., SWENSEN, I. D. and WADDELL, G. R. (2013), "Alcohol and Student Performance: Estimating the Effect of Legal Access", *Journal of Health Economics*, **32**, 22–32.
- MACCOUN, R. and REUTER, P. (1997), "Interpreting Dutch Cannabis Policy: Reasoning by Analogy in the Legalization Debate", *Science*, **278**, 47–52.
- PACULA, R. L. (1997), "Women and Substance Use: Are Women Less Susceptible to Addiction?", *American Economic Review*, **87**, 454–459.
- PACULA, R. L., ROSS, K. E. and RINGEL, J. (2003), Does Cannabis Use Impair Human Capital Formation? (No. w9963), National Bureau of Economic Research.
- PACULA, R. L. and SEVIGNY, E. L. (2014), "Cannabis Liberalization Policies: Why We Can't Learn Much from Policy Still in Motion", *Journal of Policy Analysis and Management*, **33**, 212–221.
- POPE, H. G. Jr, JACOBS, A., MIALET, J. P. *et al.* (1997), "Evidence for a Sex-Specific Residual Effect of Cannabis on Visuospatial Memory", *Psychotherapy and psychosomatics*, **66**, 179–184.
- POPE, H. G., GRUBER, A. J., HUDSON, J. I. *et al.* (2003), "Early-Onset Cannabis Use and Cognitive Deficits: What Is the Nature of the Association?", *Drug and Alcohol Dependence*, **69**, 303–310.
- RANGANATHAN, M. AND D'SOUZA, D. C. (2006), "The Acute Effects of Cannabinoids on Memory in Humans: a Review", *Psychopharmacology*, **188**, 425–444.
- RIGTER, S. and NIESINK, R. J. M. (2010), *THC-concentraties in wiet, nederwiet en hasj in Nederlandse coffeeshops (2009–2010)* (Utrecht: Trimbo-instituut).
- SACERDOTE, B. (2001), "Peer Effects With Random Assignment: Results for Dartmouth Roommates", *Quarterly Journal of Economics*, **116**, 681–704.

- UNODC: United Nations Office of Drugs and Crimes. (2012), World Drug Report, 2012. Available at <http://www.unodc.org/unodc/en/data-and-analysis/WDR-2012.html>.
- VAN OUYEN-HOUBEN, M., BIELEMAN, B. AND KORF, D. J. (2014), Coffeeshop, Toeristen en Lokale Markt, WOCD Cahier 2014–12.
- VAN OURS, J. C. and WILLIAMS, J. (2015). “Cannabis Use and Its Effects on Health, Education and Labor Market Outcomes”, *Journal of Economic Surveys*, **29**, 993–1010.