



YOUTH IN FOCUS PROJECT DISCUSSION PAPER SERIES
No. 15, October 2014

**"High"-School: The Relationship between Early Marijuana Use
and Educational Outcomes**

Deborah A. Cobb-Clark
Sonja C. Kassenboehmer
Trinh Le
Duncan McVicar
Rong Zhang



Jointly funded by the Australian Government Department of Education, Employment and Workplace Relations, the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs and the Australian Research Council and carried out by The Australian National University

Youth in Focus Project Discussion Paper Series
ISSN 1835-4025

Youth in Focus
Research School of Economics
College of Business and Economics
H W Arndt Building 25A
The Australian National University
Canberra ACT 0200

"High"-School: The Relationship between Early Marijuana Use and Educational Outcomes*

DEBORAH A. COBB-CLARK^{†‡#}

SONJA C. KASSENBOEHMER^{‡‡}

TRINH LE[§]

DUNCAN MCVICAR[‡]

RONG ZHANG[†]

[†]Melbourne Institute of Applied Economic and Social Research,
The University of Melbourne

[‡]Institute for the Study of Labor (IZA)

[#]ARC Centre of Excellence for Children and Families over the Life Course

[‡]Centre for Health Economics, Monash University

[§]Motu Economic and Public Policy Research

[‡]Queen's University Management School,

Queen's University Belfast

30 October 2014

Abstract

We use unique survey data linked to nearly a decade of administrative income support data to examine the relationship between early marijuana use (at age 14 or younger) and young people's educational outcomes. We find evidence that early marijuana use is related to educational penalties that are compounded by high-intensity use and are larger for young people living in families with a history of income support receipt. The relationships between marijuana use and both high school completion and achieving a university entrance score appear to stem from selectivity into the use of marijuana. In contrast, early marijuana use is associated with significantly lower university entrance score for those who obtain one and we provide evidence that this effect is unlikely to be driven by selection. Collectively, these findings point to a more nuanced view of the relationship between adolescent marijuana use and educational outcomes than is suggested by the existing literature.

JEL codes: I20, I24, I10, I18

Keywords: marijuana, cannabis, educational achievement, educational attainment, socio-economic disadvantage, income support receipt

*This paper is based on research commissioned by the Australian Government Department of Employment under the Social Policy Research Services Agreement (2010–2012) with the Melbourne Institute of Applied Economic and Social Research. The paper uses data from the Youth in Focus Project which is jointly funded by DEEWR, the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs and the Australian Research Council (Grant Number LP0347164) and carried out by the Australian National University. The views expressed in this paper are those of the authors alone and do not necessarily represent the views of the Department of Employment, the Australian Government, or the Melbourne Institute. We would like to thank Anastasia Sartbayeva and Chris Ryan for sharing their programming code. We also thank Department of Employment staff, Barbara Hanel, Beau Kilmer, and seminar attendants at the Melbourne Institute, the Fifth Annual Meeting on the Economics of Risky Behaviors, the 2012 meeting of the European Association of Labour Economists, and the 2013 meeting of the International Association of Health Economists for their comments on previous drafts. Any remaining errors are those of the authors. Corresponding author: Sonja Kassenboehmer, Sonja.Kassenboehmer@monash.edu, telephone +61 3 9902 4015.

1. Introduction

Marijuana (cannabis) is one of the world's most frequently used illicit drugs. Worldwide, it has been estimated that 147 million individuals (2.5 per cent of world population) use marijuana (WHO, 2012). Marijuana is also often associated with popular 'youth culture' in that the starting age for marijuana is usually much lower than for other drugs (WHO, 2012). It is estimated that 12.7 per cent of Australians aged 14 to 26 suffer from a substance use disorder which is the most common of mental health disorders experienced by young people (National Youth Mental Health Foundation, 2012). Fifty per cent of Australians with a substance use disorder first engaged in substance abuse by the age of 20 (National Youth Mental Health Foundation, 2012). In the United States, 16 per cent of 8th graders and nearly 50 per cent of 12th graders have used marijuana (Johnston et al., 2011). A number of U.S. states (e.g. Washington and Colorado) have recently approved the full legalization of marijuana use for adults aged 21 years or over, despite concerns including the potential for 'leakage' to younger age groups.

There is a vast sociological, psychological and, to a lesser degree, economics literature linking youths' marijuana use to a range of negative outcomes including (i) adverse medical events (i.e. emergency room visits or death); (ii) poor labour market outcomes; (iii) lower educational attainment; (iv) criminal activity; and (v) subsequent drug use (see Pudney, 2003; Roebuck et al., 2004). At the same time, psychologists argue that adolescents' ability to regulate intense emotions, distinguish feelings from facts, reason, make decisions, and solve problems improves steadily as they transition to adulthood (Byrnes, 2003; Smetana and Turiel, 2003), implying that occasional experimentation does not necessarily lead to enduring problem behaviour (Steinberg and Morris, 2001). What, then, separates those young people who do not experience substantial penalties from marijuana use from their peers who do?

Our goal is to contribute to answering this question by focusing specifically on youths' socio-economic background and the age at which marijuana use first begins. Early use of illicit drugs, i.e. before age 15, is an especially strong predictor of later drug use (U.S. Department of Justice, 2001) suggesting that initiation into marijuana use in early adolescence may have more severe consequences than initiation into marijuana use that occurs later. We investigate this possibility by analysing the differential relationship of 'early' versus 'late' marijuana use with both educational attainment and educational achievement. Moreover, we account for socio-economic background using unique data from the Youth in Focus (YIF) Project which interviews a nationally representative sample of 20-year old Australians about their illicit drug use, school experiences, and family background. These survey data are then linked to nearly a decade of administrative data on the family's income support receipt while the young person was growing up. This allows us to assess whether the educational penalty associated with marijuana use is related to experiences of socio-economic disadvantage.

We make several contributions to the literature. First, unlike many other researchers, we can retrospectively identify the age at which marijuana use began. This allows us to extend the handful of economic studies that investigate any marijuana use in high school to explicitly consider whether disparity in the age of initiation is linked to heterogeneity in outcomes. Second, our administrative income support data allow us to construct measures of the timing and intensity of socio-economic disadvantage (i.e. income support receipt), not at a point in time, but over the course of a young person's childhood. This is important as outcomes for young people are linked to both the timing and duration of disadvantage (Duncan and Brooks-Gunn, 1997; Berzin et al., 2006). For the first time, we are able to investigate whether socio-economic disadvantage compounds the educational penalty associated with marijuana use. Third, we consider a broad range of educational outcomes

including educational attainment (i.e. high school completion) as well as relative academic achievement (i.e. achieving a university entrance score and the university entrance score obtained). This breadth of outcome measures allows us to paint a fuller picture of the link between marijuana use and educational outcomes.

We find a strong relationship between marijuana use that occurs at age 14 or younger and diminished educational attainment and achievement. Moreover, the educational penalties associated with early marijuana use are compounded by high-intensity use. The strong link between marijuana use and both high school completion and achieving a university entrance score appear to stem from selectivity into the use of marijuana. In contrast, early marijuana use is associated with university entrance scores that are at least 5 percentiles lower and we provide evidence that this effect is unlikely to be driven by selection. Finally, the magnitude of the negative association between early marijuana use and high school completion (achieving a university entrance score) is over twice (seven times) as large for young people living in families with a history of socio-economic disadvantage than for young people from more advantaged families. Collectively, these results point to a more nuanced view of the relationship between adolescent marijuana use and educational outcomes than is suggested by the existing literature.

The paper proceeds as follows. Section 2 briefly reviews the relevant literature placing and motivates our contributions in more detail. Section 3 provides details about the data, our selection criteria, and estimation sample. The estimation strategy is presented in Section 4, while our main findings and the results of our sensitivity analyses are presented and discussed in Sections 5 and 6. Section 7 concludes and draws out some potential policy implications.

2. Existing Literature on Adolescent Marijuana Use and Educational Outcomes

During adolescence, young people make critical (and increasingly independent) decisions regarding their own health, education, employment, and family arrangements (e.g. Gruber, 2000). So it is not surprising that research in disciplines such as psychology, economics, sociology and public health consistently outlines the critical importance of adolescence in shaping the life chances of individuals. The decisions made during the period between early childhood and adolescence known as ‘early adolescence’ (roughly ages 10-14) are potentially even more important, although this key stage in the life course is often overlooked. During early adolescence, students typically move from small, self-contained primary school classrooms to larger, more integrated learning environments in secondary schools. Young people also face the biological transformations of puberty and the psychological shifts that accompany the emergence of sexuality. This time can be problematic for many young people as they undergo so many changes simultaneously (Eccles, 1999). Hence, it seems reasonable to expect that young people who make poor choices at this time, e.g. using marijuana, may face severe and long-lasting penalties on their educational attainment and achievement, labour market outcomes, health, and well-being.

There are a number of mechanisms via which marijuana use in adolescence might affect educational outcomes, including through diminished cognitive efficacy and psychological health (e.g. Scheier and Botvin, 1995; Patton et al., 2002), delinquent peer affiliations (e.g. Fergusson and Horwood, 1997), and substitution of time spent under the influence of drugs for time spent studying (Van Ours and Williams, 2009). These mechanisms are usually discussed in the literature in the context of educational attainment, although effects on educational achievement seem at least as likely as effects on educational attainment, at least through the cognitive efficacy and time use mechanisms. This in part motivates our examination of educational achievement as well as attainment.

Adolescent marijuana use and subsequent outcomes might also be correlated for other, non-causal, reasons. For example, individuals with less academic inclination, low motivation, poor mental health or less supervision at home may be more likely to try marijuana and will tend to have poorer educational outcomes whether or not they are involved in drug use. The extent to which negative associations between adolescent marijuana use and educational outcomes reflect causal impacts of marijuana use is therefore an empirical question, and one which is, arguably, yet to be conclusively answered in the existing empirical literature.

Several studies quantify the negative association between marijuana use and subsequent years of schooling or related outcomes such as labour market earnings, while stopping short of establishing the extent to which these associations can be interpretable as capturing causal effects of marijuana use. Examples include Brook et al. (1999) who use longitudinal U.S. data for approximately 1,200 African American and Puerto Rican youths to estimate the association between early-adolescent marijuana use and high school graduation. They find marijuana users to be twice as likely as non-users to drop out of high school. Bray et al. (2000) find similar results using data from four longitudinal surveys of students in a south-eastern U.S. public school system. Burgess and Propper (1998) examine the association between (late) adolescent marijuana use and earnings measured at age 28, using data drawn from the National Longitudinal Survey of Youth. They find no association between 'light' marijuana use and earnings, but strong, negative associations between 'heavy' marijuana use and earnings, particularly for black men.

Another group of studies, again largely using U.S. longitudinal data, has attempted more explicitly to identify causal relationships between adolescent marijuana use and later educational outcomes, although typically without identifying the particular mechanism underpinning the relationship. Register et al. (2001), for example, use data from the National

Education Longitudinal Study (NELS) to show that adolescent drug use reduces subsequent years of schooling by around one year, on average. Their identification strategy exploits information on the religious affiliation of the young people to instrument for marijuana use. The validity of their approach is difficult to assess, however. Chatterji (2006) also uses an instrumental variables (IV) approach with NELS data – exploiting information on school and state-level drugs policies – to try to identify the causal relationship between adolescent marijuana use and years of schooling. Unfortunately, the instrument lacks power so in the end identification relies on a conditional independence assumption with results suggesting that adolescent marijuana use leads to a reduction of around 0.2 years in schooling. Van Ours and Williams (2009) not only provide a more credibly identified estimate of the causal impact of adolescent marijuana use on educational attainment, but also examine whether the magnitude of the impact of adolescent marijuana use on educational attainment depends on the age of initiation into marijuana use. These authors exploit cross-sectional survey data, including retrospective information on marijuana use, from the Australian National Drug Strategy Household Survey. They estimate a bivariate duration model for transitions into marijuana use and out of formal education, allowing the unobservable determinants of these transitions to be correlated. They find evidence of a negative impact of adolescent marijuana use on educational attainment that is stronger the younger is the age of initiation.

This paper builds on Van Ours and Williams (2009) in several important ways which allow us to paint a uniquely detailed picture of heterogeneity in the relationship between adolescent marijuana use and educational outcomes. First, we exploit information on the intensity and duration of marijuana use to disentangle the effect of age-at-initiation from the

effect of intensity or duration of use, both of which are correlated with earlier initiation.¹ Second, we are the first to use income support records to directly investigate the extent to which the educational penalty associated with early marijuana use is compounded by socio-economic disadvantage, measured not at a point in time, but over the course of a young person's childhood. This is important as outcomes for young people are linked to both the timing and duration of disadvantage (Duncan and Brooks-Gunn, 1997; Berzin et al., 2006). Third, we investigate the effect of early marijuana use on educational achievement -- using nationally standardized achievement rankings -- as well as educational attainment. This breadth of outcome measures allows us to paint a fuller picture of the link between marijuana use and educational performance. Finally, we control for a wide range of personal and family background characteristics, including parental education and employment status, family income support history, etc. which minimizes the potential for selectivity bias associated with marijuana use to confound our results. We then adopt the approach suggested by Altonji et al. (2005) to assess the potential for any remaining selectivity bias to be driving our estimates.

3. The Youth in Focus Survey

3.1 Estimation Sample

We use unique data from the Youth in Focus (YIF) Project to investigate the link between marijuana use and both educational attainment and achievement. The YIF Project uses Australian administrative social security records to identify all young people born between October 1987 and March 1988 who ever had contact with the social security system between 1993 and 2005 (Breunig et al., 2007; 2009). Some young people are captured in the

¹ Van Ours and Williams (2009) do not have data on intensity of use but instead use a variable for use of other drugs to proxy for more intense use of marijuana.

administrative data because they receive a payment -- typically student or unemployment benefits -- in their own right. More commonly, they are in the administrative data because a family member (usually a parent) received a payment between 1993 and 2005 which depended in part on his or her relationship to the youth. A comparison of the number of young adults in these administrative data to census data indicates that over 98 per cent of young people born between October 1987 and March 1988 are represented in the administrative data (Breunig et al., 2007). This expansive coverage of the administrative data occurs because the Australian social security system is nearly universal for families with children with some payments such as the Child Care Benefit having no income test at all and others, such as the Family Tax Benefit, being denied only to families in the top quintile of the income distribution.² At the other extreme, income support payments that are targeted towards low-income parents (mainly single parents) or unemployed individuals and which are subject to income, asset and/or activity tests are also captured in the administrative data. A stratified random sample of young people was selected from the administrative data for phone interview in late 2006 (Wave 1) when they were approximately 18 years of age. Respondents were interviewed again in late 2008 (Wave 2). Respondents in both waves were also asked to undertake a self-completion survey. With permission, these survey data can be linked to the administrative records. The Wave 1 response rate was 36.1 per cent with 73.1 per cent of respondents also completing the self-completion questionnaire and more than 96 per cent of respondents consenting to having their survey data linked to their administrative records.

² The Family Tax Benefit is essentially an income tax credit to families with children. At the time of writing, families with two children (aged 0 to 12) receive a Family Tax Benefit for incomes up to \$112,400 AUD (Centrelink, 2013).

Our analysis relies on Wave 2 data when the young people are approximately 20 years of age. A total of 3,623 youths were interviewed in Wave 2. As the data on marijuana use come from the self-completion questionnaire, we drop the 31.5 per cent of respondents (1,141 youths) who failed to return the self-completion questionnaire. We drop an additional 192 young people with missing data on the variables of interest leaving us with an estimation sample of 2,290 youths.

We investigated the potential for non-response bias to influence our results and found that young people who complete the self-completion questionnaire have somewhat better educational outcomes. Specifically, they are 6 percentage points more likely to complete high school, 4 percentage points more likely to obtain a university entrance score, and have a university entrance score that is 1.5 percentiles higher. These differences are statistically significant at the 5 per cent level and suggest that our results may understate the educational penalties associated with marijuana use (results available upon request).

3.2 Educational Attainment and Achievement

Upon high school completion, those students who meet certain minimum coursework requirements (e.g. with respect to minimum credit hours, English requirements, etc.) are assigned a percentile ranking (ranging from 1 to 100) based on their academic performance in grades 11 and 12.³ The calculation of this ranking however varies by state and territory with some relying on standardized, state-wide exams and others deriving rankings from students' results in specific subjects. A national conversion formula produces comparisons across students educated in different jurisdictions resulting in the Australian Tertiary Admission Rank (ATAR). Students who wish to attend university register their preferences (in rank

³ Percentile scores under 30 are reported as 30 to the student. We set 19 cases in which students claimed that their score was below 30 to missing.

order) for the specific degree programs offered at various universities. This ranking serves as a university entrance score, as university placement offers are made centrally on the basis of students' rankings once they are known (see Marks et al., 2001). Programs in fields such as law or medicine are highly competitive and often require rankings in the 99th percentile, while most degree programs at Australia's top-tier universities accept only those students in the top quartile of the distribution. Students with rankings towards the bottom of the scale are often not offered any university place at all.

The YIF data provide us with information on three educational outcomes: (i) an indicator of high school (secondary school) completion; (ii) an indicator of obtaining a university entrance rank (ATAR); and (iii) the continuous ATAR rank itself. These are our outcomes of interest.

3.3 Marijuana Use, Socio-Economic Disadvantage, and Other Control Variables

Two unique features of the YIF data make it particularly ideal for our purposes. The first is that YIF data contain extensive information on the timing, intensity and duration of marijuana use. This allows us to specifically consider the potential interactions between age of initiation, duration, and intensity of use, all of which may be related. We create two indicators of age at initiation: (i) ever versus never used by wave 2 (i.e. typically age 20); and (ii) used versus not used by age 14. We then create a series of indicators designed to capture the intensity of marijuana use interacted with age at initiation: (i) low intensity (≤ 8 times/year) usage and initiation at age 14 or younger; (ii) medium/high intensity (9+ times/year) usage and initiation at age 14 or younger; (iii) low (≤ 8 times/year) intensity usage and after age 14; and (iv) medium/high intensity (9+ times/year) usage and initiation after age

14; (compared to never used marijuana).⁴ These later indicators allow comparisons to be made with those who at age 20 had never used marijuana.

Second, young people's survey responses can be linked to their families' social assistance records allowing us to account for the family's complete income support history. This is important in reducing unobserved heterogeneity, and provides a unique opportunity to investigate the differential consequences of marijuana use by socio-economic disadvantage (as proxied by income support history). The Australian government does not consider either the Family Tax Benefit or the Child Care Benefit to be income support payments and we adopt this convention here. We also follow Cobb-Clark et al. (2012) in creating four indicators of socio-economic disadvantage based on the intensity and timing of income support receipt: (i) no history of income support; (ii) less than six years of income support after 1998 when the respondent was aged 10 or more (late moderate income support receipt); (iii) less than six years of income support, some of which occurred before 1998 when the respondent was younger than 10 (early moderate income support receipt); and (iv) more than six years of income support receipt (intense income support receipt).

In addition to accounting for each respondent's family income support history, we also control for demographic characteristics (gender, indigenous status, country of birth) and family background when the respondent was 14 years of age (whether the youth lived with both parents, employment status of mother and father, educational attainment of mother and father and country of birth of the parents). We also check the sensitivity of our results to the inclusion of other childhood experiences (number of schools attended and whether the respondent was ever suspended from school). This detailed information helps us to draw a

⁴ Specifically, respondents are asked at what age they started using marijuana and at what age they stopped using marijuana. From this the duration of use is derived. Respondents are also asked how many times they have used marijuana in their whole life time. Dividing this number by the duration of use generates the average intensity of use. Our conclusions are robust to the particular cut-off adopted, but a cut-off value of 8 maximizes explanatory power in the regression for educational attainment discussed in the following section.

number of tentative conclusions regarding possible causal mechanisms and the likely extent to which the estimated associations could be driven by selection into marijuana use.

3.4 Descriptive Statistics

Table 1 provides descriptive statistics for the variables used in the analysis for the estimation sample as a whole (column 1) and separately for those who did and did not use marijuana by age 14 (columns 2 and 3).

The results indicate that 41 per cent of all young people have used marijuana by the age of 20, with 7 per cent having used marijuana for the first time by the age of 14. Only 62 per cent of early marijuana users completed high school in comparison to 88 per cent of those who either (i) began to use marijuana after the age of 14 or (ii) by age 20 had never used marijuana at all. Of those who completed high school, 78 per cent of early marijuana users also met the requirements to obtain an entrance score for university. In contrast, 86 per cent of late- and non-users achieved a score for university entrance. Finally, the entrance score, for those that achieve one, is on average 6 percentiles lower for early marijuana users than for their peers.

Early marijuana users also differ from late- and non-starters in their observable characteristics. Most importantly, early marijuana users are more likely (39 per cent) to live in families with a long history of income support receipt relative to late- and non-users (23 per cent). Nearly half (46 per cent) of young people who had not used marijuana by age 14 live in families with no history of income support receipt. The same is true of only a third (32 per cent) of early marijuana users. Early marijuana users also appear to be more disadvantaged than their late- or non-using peers according to most other observed characteristics.

<Table 1>

Detailed information about the distribution of age of marijuana initiation, intensity of use, and duration of use are presented in Figures 1-2. Specifically, Figure 1 depicts kernel density estimates of the starting age of marijuana use (top panel) and duration of use (bottom panel). The average age of initiation is just above 16 years of age. Of those young people who by age 20 report having used marijuana, 21 per cent reported starting at age 16.

<Figure 1>

Initiation ages and duration of use are highly (significantly) correlated with a correlation coefficient of -0.77. The mean duration of marijuana use by age 20 for those young people who started at age 14 or younger is 6.1 years, whereas for those who started after age 14 it is only 2.4 years. In fact, 82 (68) per cent of early marijuana users are still using marijuana when they are aged 18 (20).

Figure 2 shows kernel density estimates of the intensity of use, i.e. average number of times per year, for those who began using marijuana at age 14 or younger and those who started after the age of 14 but before age 20. Early marijuana use is correlated with more intensive use. Specifically, those young people who begin using marijuana after the age of 14 use it less intensively (10.2 times a year) than those who begin earlier (12.7 times per year). Among late users, the majority (56 per cent) use marijuana less than four times per year compared to only 30 per cent of early users.

<Figure 2>

Given this relationship between age at initiation on the one hand and duration and intensity of use on the other, it is important to control for the overall pattern of marijuana use so that the effect of initiation age is not confounded by either the intensity or duration of use.

4. Estimation Strategy

We begin by estimating the effect of (early) marijuana use on our three educational outcomes (Y) using the following reduced-form model:

$$Y^* = \alpha M + \mathbf{X}'\boldsymbol{\gamma} + \boldsymbol{\varepsilon} \quad (1)$$

where M is a binary indicator of marijuana use (or alternatively, marijuana use interacted with intensity of use) and \mathbf{X} is a vector of controls for own, family background and parental characteristics. A probit model is estimated (and the average marginal effects derived) when the outcome of interest (Y^*) is the propensity to either complete high school or obtain a university entrance score. We do not observe these latent propensities, but observe instead binary indicators (Y) of completions and obtaining entrance scores such that

$$\begin{aligned} Y &= 1 \text{ if } Y^* > 0, \\ Y &= 0 \text{ if } Y^* \leq 0. \end{aligned}$$

OLS estimation is used to analyse the determinants of the continuous university entrance score itself. In this case: $Y^* = Y$.⁵ It is important to note that the estimated determinants of the probability of obtaining a university entrance score and the entrance score itself are only representative of the sample of high school completers and high school completers who obtain an entrance score, respectively, and cannot be generalized to the whole population of 20-year olds as these subsamples are unlikely to be randomly drawn from the larger population. As a result, it is likely that our models underestimate the negative association between marijuana use and educational achievement since it is reasonable to expect that those who fail to complete high school (or fail to achieve an entrance score upon completion)

⁵ As a sensitivity check, we also estimated a tobit model in order to account for the fact that ATARs are censored at 30 and 100. The results are virtually identical to the OLS estimates as data censoring affects only seven respondents.

have unobserved characteristics that are positively correlated with both dropping out and marijuana use.⁶

Reduced-form estimates such as those from equation (1) can be interpreted as causal only if marijuana use is exogenous with respect to educational attainment and achievement. There are many reasons, however, to suspect that young people who choose to use marijuana early differ in unobserved ways to those who do not. Some previous researchers have addressed this endogeneity problem using an IV approach (see French and Popovici 2011 for a review). Arguably, the most credible IV approach has been to use exogenous variation in state decriminalization laws (e.g. Yamada et al., 1996 and Chaterij, 2006) or in drug prices (e.g. Chaterij, 2006), yet such variation can often be challenging to find. Many researchers must rely instead on maintained, i.e. untestable, assumptions about the excludability of the particular individual or family characteristics used as instruments to achieve identification. Moreover, IV models typically provide estimates of local -- rather than total -- average treatment effects which can be viewed as causal only if the maintained assumptions hold.

Unfortunately, our data do not provide us with plausible instruments. Neither is the bivariate duration model approach of Van Ours and Williams (2009) appropriate for modelling educational outcomes that are not expressible as transitions as is the case for both our educational achievement measures and arguably also for our measure of educational attainment. We therefore adopt two alternative strategies for dealing with the potential endogeneity problem generated by selective marijuana use. First, we estimate a series of models increasing in controls in order to assess how stable our results are to the inclusion of individual, parental, and family background characteristics. This sheds light on whether the

⁶ We do not apply a Heckman selection model as it is not feasible to find a valid exclusion restriction that influences one educational outcome (such as high school completion) but not another closely related educational outcome (obtaining an ATAR given high school completion).

conditional independence assumption is likely to hold. Second, in Section 6 we assess whether our results can be explained by the duration of marijuana use, educational experiences potentially linked to marijuana use, or selection on unobservable characteristics. The advantage of our approach is that no exclusion restrictions are needed. The disadvantage is that, at best, we learn only whether or not it is reasonable to expect any component of the estimated effect to be causal.

5. Estimation Results

5.1 Educational Attainment: High School Completion

The estimated marginal effect (and standard error) of marijuana use on the probability of high school completion is shown in Table 2.⁷ Estimates in the first column control only for the effects of respondents' own characteristics. Estimates in columns 2 and 3 additionally control for the effects of family income support history and family characteristics at age 14, respectively. Our preferred specification is in column 3.

The first panel shows the marginal effect of ever (versus not) using marijuana by age 20 on high school completion. The estimated marginal effects are small, but statistically significant, and interestingly, very stable to the inclusion of further controls. This is true even though the controls have predictive power themselves with the Pseudo R-squared increasing from 0.078 to 0.160. This supports the zero conditional mean assumption for the error terms.

⁷ Average marginal effects are calculated using Stata's 'margins' command. For every person i in the dataset and every covariate, the change in the probability that $Y = 1$ is calculated for a change of each covariate x_z from 0 to 1, leaving all other covariates at their original values. For example for the effect of early marijuana use: $\Pr(Y_i = 1|M = 1) - \Pr(Y_i = 1|M = 0) = \Phi(1\hat{\alpha}_i + X_i'\hat{\gamma}) - \Phi(0\hat{\alpha}_i + X_i'\hat{\gamma})$. The average marginal effect is the average of all individual marginal effects for each covariate. The prediction depends on the normal density evaluated at $X_i'\hat{\gamma}$ which varies with the number of control variables. The standard errors were calculated using the delta method (explained in Feiveson, 1999). As we compare alternative specifications with different number of controls, we also checked whether the normal density function evaluated at each set of control variables remains stable. We found that this was indeed the case as $\Phi(X_i'\hat{\gamma})$ only changed at the third or fourth decimal point.

<Table 2>

The relationship between marijuana use and high school completion is clearly linked to the age of initiation. Specifically, the results in panel 2 highlight the enormous disparity in the marginal effect of marijuana use for those who begin early (age 14 or younger) versus late (older than 14) relative to those who by age 20 have never tried marijuana. The large differential in the probability of high school completion associated with early marijuana use is readily apparent and appears to have been mostly driving the small average effect of marijuana use in general. The effect is again very stable to the inclusion of parental and family background controls despite the fact that these additional controls have predictive power in their own right. Our preferred estimates (column 3) indicate that having used marijuana by the age of 14 is associated with a probability of high school completion that is 20.6 percentage points lower. This effect is stronger than the effect of any other control variable in the regression. The effect of marijuana use after the age of 14 (-4.5 percentage points) is small in comparison. This is reflected in the similarity of the marginal effects on early use in panel 3 (in which late use and no use are combined) compared to panel 2.

Finally, young people who use marijuana early also tend to use it more intensively on average (see Figure 2). In contrast to Van Ours and Williams (2009), however, we can exploit information on the intensity of marijuana use to disentangle the pure age-at-initiation effect from intensity-of-use effects. The results in panel 4 of Table 2 show the key marginal effects from an extended model interacting an indicator of low (1-8 times/year) and medium/high intensity usage (9+ times/year) with our indicator for early versus late initiation.⁸ The results indicate that the educational penalties associated with early initiation into marijuana use are

⁸ We initially estimated a model in which we simply added a control for intensity of use. This slightly reduces the absolute magnitude of the coefficient for marijuana use by the age of 14, but the effects are not significantly different from one another. Results, including for a variety of other simple non-linear specifications, are available upon request.

compounded by high-intensity use. Specifically, young people who begin using marijuana at age 14 or younger, but use it with low intensity, have a probability of high school completion that is 11.5 percentage points lower than those who have not tried marijuana by the time they turn 20. This is virtually identical to the difference in high school completion rates for those who begin later (i.e. after age 14), but then move on to use marijuana intensively. While late, low-intensity marijuana use has no discernible effects on high school completion, young people who both begin marijuana use at a young age and use it intensively have a probability of dropping out of high school that is 28.2 percentage points higher. The bottom line is that both intensity of use and age of initiation matter for educational attainment.

5.2 Educational Achievement: University Entrance Scores

We turn now to consider whether (i) marijuana use is associated with the probability that high school graduates in Australia have met the curriculum requirements necessary to be awarded a university entrance score; and (ii) whether entrance scores are lower for marijuana users. In particular, the marginal effect (and standard error) of marijuana use on the probability of obtaining an entrance score -- conditional on completing high school -- are presented in Table 3. The estimated effect of marijuana use on continuous entrance scores are presented in Table 4 for those who obtain them. Tables 3 and 4 have the same structure as Table 2.

Our results indicate that it is only the combination of early initiation into and intensive use of marijuana that is related to high school graduates failing to obtain an entrance score that would permit them to apply for university (Table 3, panel 4). In our preferred specification (column 3) young people who begin using marijuana in early adolescence (age 14 or younger) and then use it with medium/high intensity have an 18.8 percentage point lower probability of obtaining a university entrance score, despite completing high school. In

all other cases, high school graduates are equally likely to obtain an entrance score irrespective of their history of marijuana use. These results are robust across specifications.

<Table 3>

Corresponding OLS results for entrance scores themselves are presented in Table 4. Although there is no significant effect of marijuana use in general on young people's university entrance scores (see panel 1), early users of marijuana have entrance scores that are between 4 and 5 percentiles lower than their peers who either began to use marijuana after the age of 14 or who had never used it by age 20 (panel 2). This latter effect is completely driven by those early users who also use marijuana intensively. Although many young people who begin marijuana use early and use it intensively successfully complete high school and obtain a university entrance score, their scores are on average 5 percentiles lower than their peers. This represents a substantially lower level of educational achievement, despite having the same level of educational attainment. Again note the robustness of the estimates across specifications.

<Table 4>

5.3 Educational Penalties and Socio-economic Disadvantage

One of our objectives is to assess whether the relationship between marijuana use and young people's educational outcomes varies by the socio-economic circumstances of their families. Our conjecture is that the educational penalty associated with marijuana use may be larger for young people living in disadvantaged households. Children living in households with a low socio-economic status, who are disproportionately from non-conventional families, tend to have less parental supervision than children in households with a high socio-economic status (e.g. Zick and Allen, 1996), and parental supervision has been shown to have important positive impacts on adolescent development through a number of mechanisms (e.g. Aizer,

2004). Marijuana use among disadvantaged young people may therefore be more likely to lead to reduced study time and engagement with delinquent peers, for example, than is the case for young people from more advantaged families. Identifying any disparities associated with socio-economic disadvantage is important in identifying the group of young people that might benefit the most from policy interventions, including interventions aimed at reducing adolescent marijuana use or supporting educational attainment and achievement.

The results presented in Table 5 come from a model in which we interact the effect of early marijuana use with an indicator of whether the young person's family has a history of intensive (as opposed to moderate or no) income support receipt. We interpret the latter as a proxy for socio-economic disadvantage generally. Marginal effects are calculated so that the sum of the cross-partial interaction effect with the marginal effect of the (standalone) early marijuana use indicator is interpretable as the impact of early marijuana use for young people from disadvantaged families, as would be the case for OLS.⁹

<Table 5>

Unfortunately, these results exhibit a general lack of estimation precision due to the small number of individuals in some cells. Despite this, we find that there is an economically and statistically significant interaction between socio-economic disadvantage and early marijuana use on the probability of high school completion. The effect of early marijuana use and high school completion for those young people growing up in income support-intensive families (28.4 percentage points) is 16 percentage points larger than that for their

⁹ Throughout this paper, we calculate the cross-partial interaction effect as described in Ai and Norton (2003) and Karaca-Mandic et al. (2012). We use Stata's 'margins' command (as explained in footnote 7) to calculate the effect of early marijuana use at low SES (variable low SES=1 (a)) and high SES (variable low SES=0 (b)). The interaction effect is the difference between (a) and (b), and the corresponding standard errors are calculated according to the usual formula: $\sqrt{\text{var}(a - b)} = \sqrt{\text{var}(a) + \text{var}(b) - 2\text{cov}(a, b)}$. The marginal effect of each of 'low SES' and 'marijuana use' is evaluated when the other term is 0.

more advantaged peers (12.4 percentage points). This more than doubles the estimated penalty on high school completion associated with early marijuana use that advantaged young people face. The interaction between early marijuana use and socio-economic disadvantage is also large in magnitude, yet not statistically significant, in the model for obtaining a university entrance score (column 2). In this case the educational achievement penalty associated with early marijuana use is over seven times as large for young people from disadvantaged families compared to their more advantaged peers. In contrast, the interaction between early marijuana use and socio-economic disadvantage in the model for university entrance score obtained is statistically insignificant and small in magnitude.¹⁰ On balance, there is clear evidence that socio-economic disadvantage compounds at least some of the educational disparities associated with early initiation into marijuana use.

6. Discussion: Understanding the Potential Mechanisms

In order to understand the potential mechanisms linking marijuana use and educational attainment and achievement, we investigated the potential for our results to be confounded by three issues: (i) the duration of marijuana use; (ii) educational experiences, i.e. suspension, expulsion, and school change, potentially linked to marijuana use; and (iii) selection on unobservable characteristics.

One possibility is that our indicator of early marijuana use simply identifies those young people who, by age 20, have used marijuana for a long time. In other words, the estimated effect of early use is not the effect of a young age of initiation, but rather the effect of a long duration of use for those who started early. We test this proposition by estimating a

¹⁰ The effect of socio-economic disadvantage on educational outcomes remains much the same when we also control for the intensity of marijuana use. Thus it does not appear that socio-economic disadvantage simply captures the effect of intensity of use. Results are available upon request.

model including interaction terms for ever used marijuana with starting age and duration of use. The results are presented in Table 6.

<Table 6>

We find that -- for those who used marijuana by age 20 -- delaying the age of initiation into marijuana by one year is associated with a 2.8 percentage point increase in the probability of completing high school, a 2.8 percentage point increase in the probability of obtaining a university entrance score; and an entrance score that is 1 percentile higher. In contrast, there is very little difference in the marginal effect of marijuana use attributable to duration of use. Thus, it appears that the age-of-initiation effect dominates the duration-of-use effect.¹¹

A second possibility is that the educational penalties associated with early marijuana use stem from disruptions in education linked to suspension, expulsion, and school changes associated with schools enforcing anti-drugs policies. Specifically, young people who use marijuana before the age of 14: (i) are four times as likely to ever be suspended from school; (ii) are nearly three times as likely to ever have to repeat a year; and (iii) attend more schools on average than their peers who do not use marijuana early (see Table 1). All of these differences are statistically significant. Given the potential endogeneity of these outcomes, our preferred specification does not include them as controls. Nonetheless, it is interesting to compare results that do and do not control for these other schooling experiences in order to explore the potential mechanisms underlying the link between marijuana use and educational outcomes at age 20.

We find that the educational attainment penalty associated with marijuana use (4.5 percentage points) is only somewhat smaller when we account for previous suspensions,

¹¹ Note that the apparently very large marginal effects of the indicator of marijuana use correspond to the notional case of marijuana use starting at age 0.

expulsions, and school changes which themselves might be the outcome of marijuana use (see Table 7). Similarly, we find that accounting for these educational disruptions reduces, but does not eliminate, the association between early marijuana use and either obtaining a university entrance score or the university entrance score itself. Taken together, these results suggest that while early marijuana use may influence educational outcomes in part through increased school disruptions, this does not provide a complete explanation.

<Table 7>

Third, we turn to the sensitivity analysis developed by Altonji et al. (2005) to assess the extent to which the estimated relationships between early marijuana use and later educational outcomes is likely to stem from selection bias as a result of unobserved differences between those who do and do not use marijuana by the age of 14 which are also correlated with educational outcomes.¹² Specifically, we exploit the fact that we control for a rich set of observed characteristics and follow Altonji et al. (2005) by assuming that the degree of selection on unobservable characteristics is no greater than the degree of selection on observable characteristics. Assuming that selection on observables is equal to selection on unobservables, the degree of selectivity bias B in our estimate of the impact of early marijuana use on educational attainment and achievement in equation (1) is then given by:

$$B = \frac{[E(X'\gamma|M=1)-E(X'\gamma|M=0)]}{Var(X'\gamma)} * Var(\varepsilon) * \frac{Var(M)}{Var(\vartheta)} \quad (2)$$

where ϑ is the error term in an OLS model of the determinants of selection into marijuana use M . This implies that the ‘true’ impact of marijuana use on young people’s educational outcomes (α) is given by:

¹² We focus here only on the selection bias associated with early versus late marijuana use because the Altonji et al (2005) method is applicable when there is one discrete endogenous variable. Our results in Tables 2 to 4 show that the effect of early marijuana use does not change significantly if we group late users with non-users or retain them as a separate group justifying this focus.

$$\hat{\alpha} = \alpha + B \quad (3)$$

where $\hat{\alpha}$ results from estimating equation (1).

For convenience, we calculate the ratio R of the estimated impact of marijuana use to the approximated selection bias as follows:

$$R = \frac{\hat{\alpha}}{B} \quad (4)$$

In effect, R shows how strong selection on unobservable characteristics would have to be (relative to selection on observable characteristics) such that the complete effect of $\hat{\alpha}$ could potentially be attributed completely to the effects of selection bias. That is, selection on unobservable characteristics would produce the estimated coefficient in equation (1) even if the true effect of early marijuana use were zero. Altonji et al. (2005) argue that a ratio of estimate to bias of less than 1 is indicative of an association that is likely to be explained by selection bias.

Table 8 shows the implied bias as given in equation (2) and the ratio of estimate to bias as given in equation (4). We find that our estimated effect of early marijuana use on high school completion (see Table 2, column 3) could be the result of selection bias if selection on unobservable characteristics were only 69 per cent as strong as selection on observables. We draw a similar conclusion with respect to the probability of obtaining a university entrance score conditional on graduating from high school (see column 2). This implies that it is very unlikely that marijuana use has a causal effect on high school completion or on the probability that high school graduates meet the requirements to obtain an entrance score for university. Rather, the relationship stems from unobserved differences in the characteristics of early marijuana users.

<Table 8>

In contrast, the ratio of estimate to bias is negative when we consider the effect of marijuana use on university entrance scores themselves. This occurs because the implied bias given by equation (2) is positive rather than negative. This is true only if:

$$[E(X'\gamma|M = 1) - E(X'\gamma|M = 0)] > 0. \quad (6)$$

which indicates that early marijuana users who graduate from high school and obtain a university entrance score have ‘better’ observable characteristics than their non- or late-using peers who also graduate from high school and obtain university entrance scores. We investigated this in detail and found that this is indeed the case in our data. In particular, early users who obtain a university entrance score are less likely to have experienced socio-economic disadvantage and are more likely to have highly educated mothers.¹³ To the extent that selection on unobservable characteristics is of a similar sign and magnitude, it appears that those unobservable characteristics that are responsible for increasing early marijuana use are positively correlated with university entrance scores themselves. It therefore seems reasonable to conclude that our estimates of the percentile reduction in university entrance scores associated with early marijuana use (Table 4) are unlikely to be driven entirely by selection on unobservables. Nor are they easily explained by reverse causality. While it is possible that there are other threats to causality, we believe that these results provide strong evidence that there is a negative causal relationship between early marijuana use and educational achievement as measured by entrance scores themselves.

Finally, in addition to the Altonji et al. (2005) approach, we also follow McCaffrey et al. (2010) in including an indicator for tobacco smoking (equal to one for those that report

¹³ Marijuana users are more likely to live in families with a history of income support receipt (46.9 per cent) than are non-users (31.9 per cent). However for the group who obtain a university entrance score, income support receipt is less common among early users (18.8 per cent) than among non- and later-users (24.5 per cent). Similarly, early marijuana users who obtain a university entrance score are more likely to have mothers who also completed high school (59.4 per cent) than are other young people who also obtain university entrance scores (56.6 per cent).

having ever smoked and zero otherwise) and its interaction with marijuana use in our main specifications. The intuition for this is that tobacco smoking may capture some of the selection process underpinning the decision to smoke marijuana. The smoking indicator is significantly negative in all cases indicating that smoking tobacco is also associated with an educational penalty even among those young people who do not smoke marijuana. This educational penalty is significantly larger, however, for those who smoke both marijuana and tobacco. Our substantive conclusions regarding the educational penalties associated with marijuana use are unchanged, however, as the average marginal effect of marijuana use is similar in size and statistical significance irrespective of whether or not we control for smoking (results available upon request).

7. Conclusions

Adolescent marijuana use is widespread across much of the developed world. Previous research has documented that marijuana users tend to have poorer educational attainment than non-users. This negative relationship is particularly pertinent right now as Uruguay and additional U.S. states consider moves towards full legalization of marijuana. In this paper, we examine the relationship between marijuana use and measures of educational achievement as well as educational attainment. We also examine the extent to which the strength of these associations varies with age of initiation, intensity of use, duration of use, and socio-economic status. We show that the negative relationship between marijuana use and educational outcomes is particularly strong for some young people.

In particular, we find that early marijuana use is strongly related to diminished educational attainment and achievement and that the educational penalties associated with early marijuana use are compounded by high-intensity use. We also find that the strong link between marijuana use, on the one hand, and high school completion and achieving a

university entrance score, on the other, are likely to be driven by the selectivity associated with the use of marijuana. In contrast, our analysis suggests that early marijuana users who complete the necessary requirements obtain university entrance scores that are 5 percentiles lower, a relationship that is unlikely to be entirely explained by unobserved heterogeneity associated with marijuana use. Finally, we show that the magnitude of the negative relationship between early marijuana use and high school completion (achieving a university entrance score) is over twice (seven times) as large for young people living in families with a history of socio-economic disadvantage than for young people from more advantaged families.

Governments across the developed world are under intense pressure to improve educational attainment and achievement, often in the face of tight funding constraints and powerful vested interests. This paper has potentially important implications for policy in this space. First, the strong links between early marijuana use and educational outcomes at age 20 suggest that those who have begun to use marijuana by age 14 – particularly those from disadvantaged family backgrounds – may be an important group to target for intervention. If these relationships are not driven entirely by selection, then interventions to reduce or delay marijuana use, or to ameliorate its detrimental effects, may impact positively on educational outcomes. Moreover, our evidence that the negative educational penalty associated with early marijuana use is larger and more broadly-based for young people from socio-economically disadvantaged backgrounds -- coupled with the fact that early use is more common among those from disadvantaged backgrounds -- suggests that disadvantaged young people should be the focus of such interventions. Making progress in this area could disproportionately improve the educational outcomes of disadvantaged young people and might contribute to efforts aimed at closing the socio-economic gap in educational performance.

References

- Ai, C. and Norton, E. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1):123–129.
- Aizer, A. (2004). Home alone: Supervision after school and child behavior. *Journal of Public Economics*, 88(9–10):1835–1848.
- Altonji, J. G., Elder, T. E. and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of Political Economy*, 113(1):151–184.
- Berzin, S. C., De Marco, A. C., Shaw, T. V., Unick, G. J. and Hogan, S. R. (2006). The effect of parental work history and public assistance use on the transition to adulthood. *Journal of Sociology and Social Welfare*, 33(1):141–162.
- Bray, J., Zarkin, G., Ringwalt, C. and Qi, J. (2000). The relationship between marijuana initiation and dropping out of high school. *Health Economics*, 9(1):9–18.
- Breunig, R., Cobb-Clark, D., Gorgens, T. and Sartbayeva, A. (2007). User's guide to the Youth in Focus Data, version 1.0. Youth in Focus Project Discussion Paper Series 1, Research School of Economics, Australian National University.
- Breunig, R., Cobb-Clark, D., Gorgens, T., Ryan, C. and Sartbayeva, A. (2009). User's guide to the Youth in Focus data, version 2.0. Youth in Focus Project Discussion Paper Series 8, Research School of Economics, Australian National University.
- Brook, J., Balka, E., and Whiteman, M. (1999). The risks for late adolescence of early adolescent marijuana use. *American Journal of Public Health*, 89(10):1549–1554.
- Byrnes, J. P. (2003). Cognitive development during adolescence. In *Blackwell Handbook of Adolescence*, G. R. Adams and M. D. Berzonsky (eds), Malden, MA: Blackwell Publishing, 227–246.
- Burgess, S., Gardiner, K. and Propper, C. (2001). Growing up: School, family and area influences on adolescents' later life chances. CASE papers 249, Centre for Analysis of Social Exclusion, London School of Economics.
- Centrelink (2013). A guide to Australian government payments. Available at: <http://www.humanservices.gov.au/corporate/publications-and-resources/a-guide-to-australian-government-payments> (accessed March 2013).
- Chatterji, P. (2006). Illicit drug use and educational attainment. *Health Economics*, 15(5):489–511.
- Cobb-Clark, D., Ryan, C. and Sartbayeva, A. (2012). Taking chances: The effect of growing up on welfare on the risky behaviour of young people. *Scandinavian Journal of Economics*, 114(30):729–755.

- Duncan, G. J. and Brooks-Gunn, J. (eds) (1997). *Consequences of Growing Up Poor*. NY: Russell Sage Foundation.
- Eccles, J. S. (1999). The development of children ages 6 to 14. *The Future of Children*, 9(2): 30–44.
- Feiveson, A. H. (1999). Explanation of the delta method. Available at: <http://www.stata.com/support/faqs/statistics/delta-method/> (accessed October 2014).
- Fergusson D. M. and Horwood, L. J. (1997). Early onset cannabis use and psychosocial adjustment in young adults. *Addiction*, 92(3):279–296.
- French, M. T. and Popovici, I. (2011). That instrument is lousy! In search of agreement when using instrumental variables estimation in substance use research. *Health Economics*, 20(2):127–146.
- Gruber, J. (2000). *Risky Behavior among Youths: An Economic Analysis*. University of Chicago Press, Chicago.
- Johnston, L. D., O'Malley, P. M., Bachman, J. G. and Schulenberg, J. E. (2011). *Monitoring the Future national survey results on drug use, 1975-2010. Volume I: Secondary school students*. Institute for Social Research, The University of Michigan.
- Karaca-Mandic, P., Norton, E. and Dowd, B. (2012). Interaction terms in nonlinear models. *Health Services Research*, 47(1):255–275.
- Marks, G. N., McMillan, J. and Hillman, K. (2001). Tertiary entrance performance: The role of student background and school factors. *Research Report Number 22*, Australian Council for Educational Research.
- McCaffrey, D. N., Pacula, R. L., Han, B. and Ellickson, P. (2010). Marijuana use and high school dropout: The influence of unobservables. *Health Economics*, 19(11):1281–1299.
- National Youth Mental Health Foundation (2012). Substance use disorders. Available at: <http://www.headspace.org.au/what-works/research-information/substance-use>. (accessed December 2011).
- Patton, G. C., Coffey, C., Carlin, J. B., Degenhardt, L., Lynskey, M. and Hall, W. (2002). Cannabis use and mental health in young people: Cohort study. *British Medical Journal*, 325(7374):1195–1198.
- Pudney, S. (2003). The road to ruin: Sequences of initiation into drugs and crime in Britain. *Economic Journal*, 113(486):C182-C198.
- Register, C. A., Williams, D. R. and Grimes, P. W. (2001). Adolescent drug use and educational attainment. *Education Economics*, 9(1):1–18.

- Roebuck, M. C., French, M. T. and Dennis, M. L. (2004). Adolescent marijuana use and school attendance. *Economics of Education Review*, 23(2):133–141.
- Scheier, L. and Botvin, G. (1995). Effects of early adolescent drug use on cognitive efficacy in early-late adolescence: A developmental structural model. *Journal of Substance Abuse*, 7(4):379–404.
- Smetana, J. G. and Turiel, E. (2003). Emotional development during adolescence. In *Blackwell Handbook of Adolescence*, G. R. Adams and M. D. Berzonsky (eds), Malden MA: Blackwell Publishing, 229–289.
- Steinberg, L. and Morris, A. S. (2001). Adolescent development. *Annual Review of Psychology*, 52:83–110.
- U.S. Department of Justice (2001). Substance Abuse: The Nation's Number One Health Problem. OJJDP Fact Sheet 17, Office of Justice Programs.
- van Ours, J. and Williams, J. (2009). Why parents worry: Initiation into cannabis use by youth and their educational attainment. *Journal of Health Economics*, 28(1):132–142.
- World Health Organization (WHO) (2012). Management of substance abuse. Available at: http://www.who.int/substance_abuse/facts/cannabis/en/ (accessed December 2011).
- Yamada T., Kendrix M. and Yamada, T. (1996). The impact of alcohol consumption and marijuana use on high school graduation. *Health Economics*, 5(1):77–92.
- Zick, C. D. and Allen, C. R. (1996). The impact of parents' marital status on the time adolescents spend in productive activities. *Family Relations*, 45(1):65–71.

Figures

Figure 1: Age of Initiation into Marijuana Use and Duration of Use

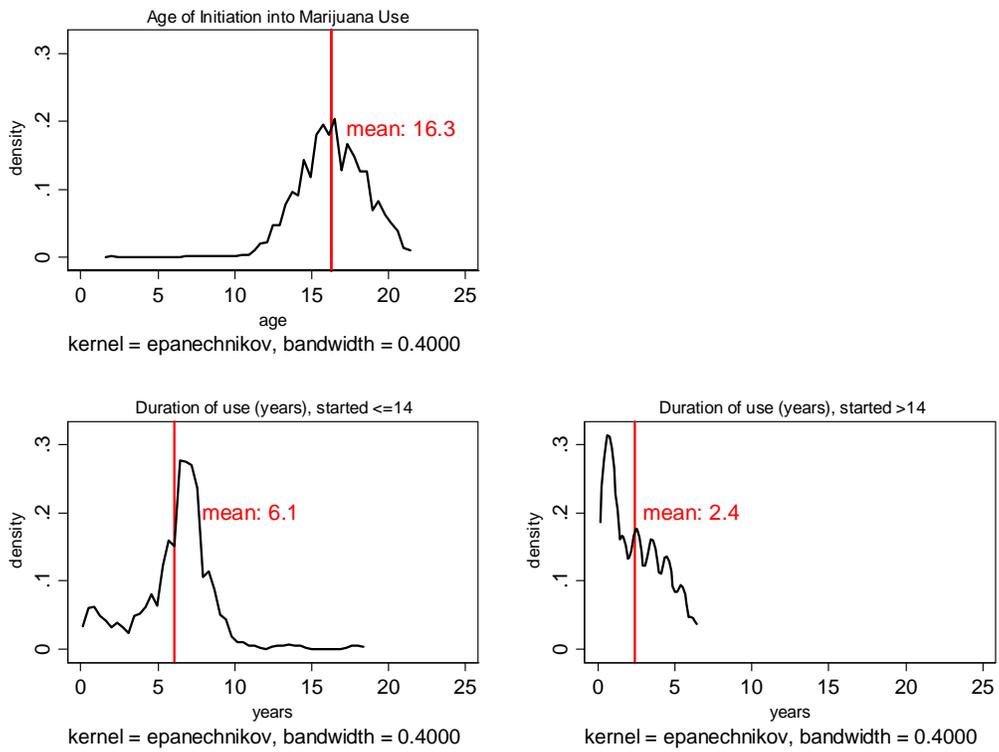
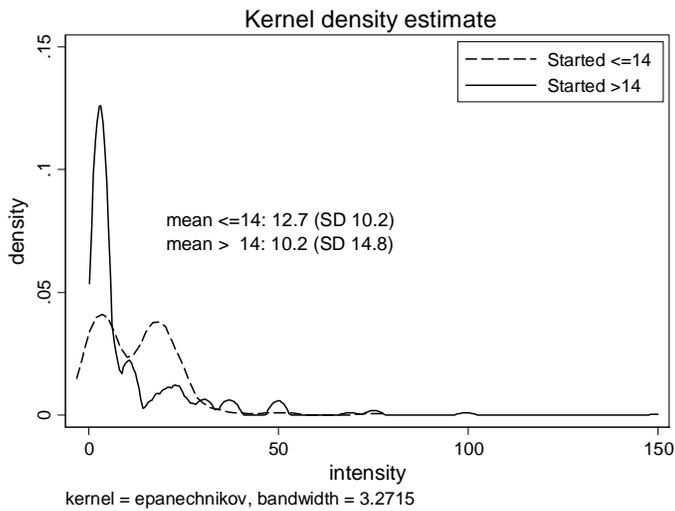


Figure 2: Intensity of Marijuana Use



Tables

Table 1: Means of Key Regression Variables

	Total Sample mean	Used marijuana <=14 (total sample)	
		No mean	Yes mean
Used marijuana	0.406	0.359	1.000 ***
Used marijuana <= 14	0.073	0.000	1.000 ***
Outcomes			
High School Completion	0.860	0.879	0.617 ***
Obtained a university entrance score	0.854	0.858	0.775 *
University entrance score (scale: 0-100)	76.521	76.793	70.885 ***
Own characteristics			
Female	0.543	0.545	0.522
Indigenous Australian	0.022	0.019	0.061 **
Born in a non-English speaking (NES) country	0.079	0.084	0.008 ***
Metropolitan residence	0.689	0.691	0.657
Family characteristics at age 14 (own report)			
Lived with both parents at 14	0.779	0.796	0.558 ***
Mother employed at 14	0.710	0.709	0.724
Mother's education: Year 12 at 14	0.530	0.532	0.497
Mother's education: Year 12 at 14: n/a	0.058	0.058	0.051
Father's education: Year 12 at 14	0.481	0.487	0.402 **
Fathers's education: Year 12 at 14: n/a	0.134	0.126	0.227 ***
At least one parent born in a NES country	0.254	0.262	0.158 ***
At least one parent born in a NES country: n/a	0.006	0.005	0.026 *
Family income support history			
Intensive income support	0.245	0.233	0.393 ***
Moderate (early) income support	0.228	0.231	0.194
Moderate (late) income support	0.080	0.079	0.090
Number of observations	2290	2108	182
Schooling experiences (own report)			
Suspended from school	0.142	0.116	0.477 ***
Number of schools	2.948	2.892	3.656 ***
Repeated a year	0.085	0.076	0.200 ***

Notes: All statistics are weighted. *, ** and *** denote sample means that are significantly different from the column to the left at the 10%, 5% and 1% level respectively.

Table 2: (Marginal) Effects of Marijuana Use on the Probability of High School Completion

	(1)	(2)	(3)
(1)			
Marijuana Use	-0.086*** (0.016)	-0.078*** (0.015)	-0.076*** (0.015)
Observations	2290	2290	2290
Pseudo R-squared	0.078	0.116	0.160
(2)			
Marijuana Use <=14	-0.251*** (0.036)	-0.218*** (0.035)	-0.206*** (0.034)
Marijuana Use >14	-0.046** (0.016)	-0.044** (0.016)	-0.045** (0.016)
Observations	2290	2290	2290
R-squared	0.096	0.130	0.173
(3)			
Marijuana Use <=14	-0.233*** (0.036)	-0.201*** (0.034)	-0.187*** (0.033)
Observations	2290	2290	2290
Pseudo R-squared	0.092	0.126	0.169
(4)			
Marijuana Use <=14, low intensity	-0.116* (0.047)	-0.108* (0.045)	-0.115** (0.045)
Marijuana Use <=14, med/high intensity	-0.361*** (0.050)	-0.312*** (0.049)	-0.282*** (0.048)
Marijuana Use >14, low intensity	-0.008 (0.017)	-0.010 (0.017)	-0.012 (0.017)
Marijuana Use >14, 9-150 med/high intensity	-0.129*** (0.029)	-0.115*** (0.028)	-0.114*** (0.027)
Observations	2290	2290	2290
Pseudo R-squared	0.111	0.142	0.183
Other control variables			
Own characteristics	Yes	Yes	Yes
Family income support history	No	Yes	Yes
Family characteristics at age 14	No	No	Yes

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: (Marginal) Effects of Marijuana Use on the Probability of Obtaining a University Entrance Score

	(1)	(2)	(3)
(1)			
Marijuana Use	-0.017 (0.018)	-0.017 (0.018)	-0.024 (0.018)
Observations	1799	1799	1799
Pseudo R-squared	0.027	0.054	0.080
(2)			
Marijuana Use <=14	-0.073* (0.043)	-0.073* (0.043)	-0.080* (0.042)
Marijuana Use >14	-0.007 (0.019)	-0.008 (0.018)	-0.015 (0.019)
Observations	1799	1799	1799
R-squared	0.029	0.055	0.082
(3)			
Marijuana Use <=14	-0.070* (0.042)	-0.071* (0.042)	-0.074* (0.042)
Observations	1799	1799	1799
Pseudo R-squared	0.029	0.055	0.081
(4)			
Marijuana Use <=14, low intensity	0.016 (0.047)	0.013 (0.046)	0.003 (0.047)
Marijuana Use <=14, med/high intensity	-0.189*** (0.073)	-0.189*** (0.072)	-0.188*** (0.071)
Marijuana Use >14, low intensity	0.010 (0.020)	0.006 (0.020)	0.000 (0.020)
Marijuana Use >14, med/high intensity	-0.055 (0.033)	-0.044 (0.032)	-0.054 (0.033)
Observations	1799	1799	1799
Pseudo R-squared	0.035	0.060	0.087
Other control variables			
Own characteristics	Yes	Yes	Yes
Family income support history	No	Yes	Yes
Family characteristics at age 14	No	No	Yes

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: (Marginal) Effects of Marijuana Use on the University Entrance Score

	(1)	(2)	(3)
(1)			
Marijuana Use	-0.218 (0.914)	-0.290 (0.908)	-0.757 (0.915)
Observations	1385	1385	1385
R-squared	0.014	0.029	0.063
(2)			
Marijuana Use <=14	-4.459** (2.101)	-4.741** (2.089)	-5.010** (2.069)
Marijuana Use >14	0.389 (0.951)	0.347 (0.946)	-0.154 (0.951)
Observations	1385	1385	1385
R-squared	0.018	0.033	0.067
(3)			
Marijuana Use <=14	-4.597** (2.073)	-4.863** (2.062)	-4.951** (2.036)
Observations	1385	1385	1385
R-squared	0.018	0.033	0.067
(4)			
Marijuana Use <=14, low intensity	-1.912 (2.561)	-2.040 (2.545)	-2.370 (2.515)
Marijuana Use <=14, med/high intensity	-9.342*** (3.497)	-9.919*** (3.481)	-10.108*** (3.448)
Marijuana Use >14, low intensity	1.077 (1.054)	0.899 (1.049)	0.426 (1.050)
Marijuana Use >14, med/high intensity	-1.580 (1.595)	-1.243 (1.587)	-1.850 (1.577)
Observations	1385	1385	1385
R-squared	0.022	0.036	0.070
Other control variables			
Own characteristics	Yes	Yes	Yes
Family income support history	No	Yes	Yes
Family characteristics at age 14	No	No	Yes

Notes Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: (Marginal) Effects of Marijuana Usage, by Socio-economic Status

	High School Completion	Obtaining a University Entrance Score	University Entrance Score
Low SES = intensive family welfare receipt			
Marijuana Use <= 14	-0.124 *** (0.04)	-0.025 (0.043)	-5.151 ** (2.26)
Low SES	-0.050 *** (0.017)	-0.071 *** (0.022)	-3.000 *** (1.118)
Marijuana Use <= 14 × low SES	-0.160 ** (0.068)	-0.164 (0.101)	1.109 (5.179)
Observations	2290	1799	1385
(Pseudo) R-squared	0.168	0.077	0.054
Number of low SES marijuana users	85	28	12

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. All regressions control for own characteristics and family characteristics at 14.

Table 6: Sensitivity of (Marginal) Effects of Marijuana Use to Starting Age and Duration of Use

	Year 12		Obtained a University Entrance Score		University Entrance Score	
	(1) Probit	(2) Probit	(3) Probit	(4) Probit	(5) OLS	(6) OLS
Ever Used Marijuana	-0.790 *** (0.058)	-0.702 *** (0.157)	-0.539 *** (0.162)	-0.732 *** (0.143)	-17.942 ** (5.935)	-13.869 (9.700)
Ever Used x Start Age	0.028 *** (0.002)	0.031 *** (0.002)	0.027 *** (0.004)	0.027 *** (0.004)	1.019 ** (0.348)	0.819 (0.514)
Ever used x Duration		-0.006 (0.009)		0.009 * (0.005)		-0.262 (0.493)
Observations	2290	2290	1799	1799	1385	1385
R-squared/Pseudo R-squared	0.182	0.182	0.086	0.087	0.069	0.069
Other control variables						
Own characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Family welfare receipt history	Yes	Yes	Yes	Yes	Yes	Yes
Family characteristics at age 14	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: (Marginal) Effects of Marijuana Use on Educational Achievement and Attainment after Controls for Suspension/Expulsion and Number of School included

	(1) Year 12	(2) Obtained a University Entrance Score	(3) University Entrance Score
(1)			
Marijuana Use	-0.045*** (0.015)	-0.018 (0.018)	-0.348 (0.921)
Observations	2290	1799	1385
R-squared	0.194	0.083	0.075
(2)			
Marijuana Use <=14	-0.130*** (0.032)	-0.063 (0.042)	-3.842* (2.090)
Marijuana Use >14	-0.027 (0.015)	-0.010 (0.019)	0.115 (0.953)
Observations	2290	1799	1385
R-squared	0.200	0.084	0.077
(3)			
Marijuana Use <=14	-0.115*** (0.031)	-0.059 (0.041)	-3.890* (2.051)
Observations	2290	1799	1385
R-squared	0.198	0.084	0.077
(4)			
Marijuana Use <=14, low intensity	-0.077* (0.042)	0.014 (0.046)	-1.935 (2.513)
Marijuana Use <=14, med/high intensity	-0.182*** (0.046)	-0.167** (0.071)	-7.746** (3.504)
Marijuana Use >14, low intensity	-0.003 (0.017)	0.004 (0.020)	0.603 (1.050)
Marijuana Use >14, med/high intensity	-0.082*** (0.026)	-0.049 (0.032)	-1.379 (1.579)
Observations	2290	1799	1385
R-squared	0.206	0.089	0.079
Other control variables			
Own characteristics	Yes	Yes	Yes
Family income support history	Yes	Yes	Yes
Family characteristics at age 14	Yes	Yes	Yes
Suspension/expulsion	Yes	Yes	Yes
No. Schools attended	Yes	Yes	Yes

Notes: Standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Amount of Selection on Unobservables Relative to Selection on Observables Required to Attribute the Entire Effect of Marijuana Usage to Selection Bias

	High School Completion	Obtaining a University Entrance Score	University Entrance Score
Unconstrained estimate of marijuana use	-0.726***	-0.294**	-4.951**
Standard error	(0.108)	(0.151)	(2.036)
Marginal effect	[-0.187]	[-0.074]	[-4.951]
Implied bias	-1.050	-0.416	2.299
Ratio of estimate to bias	0.689	0.707	-2.154
Sample size	2290	1799	1385
Number of marijuana users	182	97	64

Notes: Standard errors are in parentheses, marginal effects in square brackets. *** p<0.01, ** p<0.05, * p<0.1. Regressions control for own characteristics, family income support history and family characteristics at 14.