

Williams, Jenny; van Ours, Jan C.

Working Paper

Early Cannabis Use and School to Work Transition of Young Men

Tinbergen Institute Discussion Paper, No. 17-004/V

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Williams, Jenny; van Ours, Jan C. (2017) : Early Cannabis Use and School to Work Transition of Young Men, Tinbergen Institute Discussion Paper, No. 17-004/V, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/149518>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2017-004/V
Tinbergen Institute Discussion Paper



Early Cannabis Use and School to Work Transition of Young Men

*Jenny Williams*¹

*Jan C. van Ours*²

¹Department of Economics, University of Melbourne, Parkville, Australia and IZA

²Erasmus School of Economics, Erasmus University Rotterdam and Tinbergen Institute (Amsterdam/Rotterdam), The Netherlands; Department of Economics, University of Melbourne, Parkville, Australia and CEPR (London) and IZA (Bonn)

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Early Cannabis Use and School to Work Transition of Young Men

Jenny Williams ^{*} Jan C. van Ours ^{†‡}

January 10, 2017

Abstract

We study the impact of early cannabis use on the school to work transition of young men. Our empirical approach accounts for common unobserved confounders that jointly affect selection into cannabis use and the transition from school to work using a multivariate mixed proportional hazard framework in which unobserved heterogeneities are drawn from a discrete mixing distribution. Extended models account for school leavers' option of returning to school rather than starting work as a competing risk. We find that early cannabis use leads young men to accept job offers more quickly and at a lower wage rate compared to otherwise similar males who did not use cannabis. These effects are present only for those who use cannabis for longer than a year before leaving school. Overall, our findings are consistent with a mechanism whereby early non-experimental cannabis use leads to greater impatience in initial labor market decision-making.

Keywords: multivariate duration models, discrete factors, cannabis use, job search, wages

JEL-codes: C41, I12, J01

^{*}Department of Economics, University of Melbourne, Parkville, Australia and IZA; email: jenny.williams@unimelb.edu.au. Williams is grateful to the Department of Economics at Tilburg University and IZA, who supported visits while this research was in progress.

[†]Erasmus School of Economics, Erasmus University Rotterdam and Tinbergen Institute (Amsterdam/Rotterdam), The Netherlands; Department of Economics, University of Melbourne, Parkville, Australia and CEPR (London) and IZA (Bonn); vanours@ese.eur.nl.

[‡]The authors thank Mabel A. Andalón who set up a preliminary version of the data and Shannon Ward who set up the final version of the data used in this paper. Seminar participants at the IZA, especially Peter Kuhn, and the University of Queensland, and participants of the 2015 IAAE conference provided helpful feedback that has improved this research.

1 Introduction

The legal environment surrounding cannabis has changed dramatically in the US in the past few years. Medical marijuana pharmacies are becoming more common and four states, Alaska, Colorado, Oregon and Washington, have legalized cannabis. The shift towards a more liberal cannabis policy setting is likely to continue with several states, including California, Nevada, Maine, Arizona and Massachusetts, holding ballots to legalize cannabis in 2016. And while the full extent to which the new legal environment will impact on cannabis consumption is not yet clear (Pacula et al., 2015; Pacula and Sevigny, 2014), it is very likely that its effect will be greatest among young people. This follows from the fact that uptake of cannabis most often occurs during the teen years (van Ours and Williams, 2009) and from recent evidence showing an upward trend in cannabis use by young adults (Johnston et al., 2015). For example, monthly cannabis use among young adults (aged 19-28) has increased from 15.7% in 2006 to 20.1% in 2015 (Johnston et al., 2015). To put this in perspective, monthly cigarette use among young adults fell from 27.0% to 16.6% over the same time period. Given these changes in the policy environment and in cannabis use, understanding the consequences of cannabis use for youth and young adults is of increasing policy importance.

This paper focuses on the consequences of cannabis use in terms of early labor market experiences, investigating its impact on the school to work transition. Understanding the impact of youthful cannabis use on transitions into the labor market is important not only because of the potential for it to effect initial labor market outcomes. Early cannabis use is likely to have a more enduring impact on economic well-being because initial labor market positions are an important determinant of labor market trajectories (Altonji et al., 2014; Oreopoulos et al., 2012). With over 50% of young adults reporting to have ever used cannabis, and over 30% reporting use in the past 12 months (Substance Abuse and Mental Health Services Administration, 2014), understanding whether and to what extent cannabis use impacts on transitions from school to work is both a relevant and significant margin from an economic perspective.¹

Despite its potential to have long lived effects on the economic well-being of young adults,

¹Substance Abuse and Mental Health Services Administration (2014) report that, for the National Survey on Drug Use and Health, 51.9% of 18–25 year olds report having used cannabis in their lifetime and 31.6% report having used in the past year, while in Monitoring the Future, the figures are 53.3% (of 19-24 year olds) having used in their lifetime and 35.5% have used in the past year.

little is currently known about the impact of cannabis use on early labor market experiences. No previous research on the labor market effects of cannabis use specifically considers the school to work transition, and the findings on labor market outcomes in previous studies are ambiguous, inconclusive, and are now quite dated.² For example, some studies report that cannabis use increases wages, some find no effect, while others report that cannabis use decreases wages (Kaestner, 1994b; Register and Williams, 1992; van Ours, 2007). In light of the current literature and the changing policy environment, there is an urgent need for robust and reliable evidence on the early labor market consequences of youthful cannabis use.

In establishing whether, and to what extent, early cannabis use affects transitions into the labor market, we focus on males and consider two dimensions of this transition; the duration of time between leaving school and starting employment, and the wage rate received when starting the first post-school job. Within the context of a simple search-theoretic framework, joint consideration of these dimensions of labor market behavior provides useful predictions that help discern between potential mechanisms via which cannabis use may affect the school-to-work transition: by increasing an individual's preference for leisure, by increasing his discount rate (increasing impatience), by increasing his probability of being laid off, or by lowering his expected wage (through adversely affecting mental health, for example).

A central challenge to establishing the impact of cannabis use on labor market outcomes is that both cannabis use and labor market performance may be influenced by characteristics that are not observed. For example, a negative correlation between employment and cannabis use may reflect a stronger preference for leisure among cannabis users compared to non-users, rather than a causal effect of cannabis.³ We address this issue using a multivariate mixed proportional hazard framework in which cannabis use dynamics are modeled jointly with the duration of job search and the starting wage for the first post-school job. A discrete multivariate mixing distribution accounts for the potential correlation in unobserved heterogeneities affecting cannabis use dynamics and the outcomes of interest. Our model is

²See van Ours and Williams (2015) for an overview of studies on the labor market effects of cannabis use and a discussion on the strengths and weaknesses of the different empirical approaches and identification strategies used.

³Ideally, the effects of cannabis use should be studied in an experimental setting where cannabis use is the outcome of a lottery rather than the choice of an individual. However, as far as we know, Kagel et al. (1980) is the only paper that uses an experimental method to examine the effect of cannabis on labor market performance. The paper reports about a laboratory experiment of 98 days in which the productivity of subjects was measured after they were exposed to various intensities of cannabis use.

estimated using information on males from the very rich National Longitudinal Survey of Youth, 1997 cohort. Our main findings are as follows. After accounting for selection, early cannabis use reduces the time spent in job search before accepting a job offer and reduces the wage rate at which a job offer is accepted. These labor market effects of cannabis use are driven by those who used cannabis for more than a year before leaving school. We find no effect of cannabis use on either the rate at which jobs are found, or the wage rate received when starting work for those who use cannabis for a year or less before leaving school.

An extended sensitivity analysis investigates potential threats to the interpretation of our findings. We focus on the propensity of school leavers to return to school. This is an issue because, in our sample, a significant minority of school leavers exit post-school unemployment by returning to school, rather than finding a job. Moreover, returning to school is less likely among early cannabis users. This suggests a spurious correlation between returning to school and cannabis use as a plausible alternative explanation for our finding that cannabis users find jobs more quickly. We address this issue using a competing risk model in which school leavers exit unemployment either by returning to school or by finding a job, and allowing early cannabis use to affect the hazard of exiting via each option. This extended analysis provides new insights into the impact of cannabis use on completed schooling (and hence future labor market outcomes), by studying its effect on the decision to return to school. Additional sensitivity analyses seek to account for cannabis users' greater likelihood of working before returning to school, and examining the robustness of the findings to confining attention to full time employment rather than any employment, as well as examining their robustness to identifying assumptions. In all cases our findings are upheld, and typically strengthened. We conclude that early, non-experimental cannabis use leads young men to spend less time searching before accepting a job offer, and to accept job offers at a lower wage rate. In terms of mechanisms, our results are not consistent with mechanisms based on cannabis use increasing the value of leisure or adversely affecting the expected wage rate. They are, however, consistent with early non-experimental cannabis use affecting labor market outcomes through increasing the rate at which the future is discounted or increasing the rate of being laid off, with the former mechanism being more likely in the context of our study of initial labor market experiences.

The set-up of our paper is as follows. In the next section we provide an overview of the conceptual frameworks and findings of previous studies. We also introduce the conceptual framework underlying our own empirical investigation. In section 3 we provide information

about the NLSY97 data and present some stylized facts about cannabis use, job search, returning to school, and starting wages. In section 4 we discuss the set-up of our empirical analysis. Section 5 presents results from our baseline and extended analysis. Section 6 concludes.

2 Background and Conceptual Framework

Research into the labor market consequences of cannabis use tends to focus either on the effect of cannabis use on wages, or on the effect of cannabis use on labor supply. Following this convention, these two strands of the literature are discussed separately in this section.⁴ We end with a discussion of the framework used in this study, which takes a more integrated approach to analyzing the labor market effects of cannabis use.

2.1 Wages and Cannabis Use

The literature investigating the wage effects of cannabis use typically takes as its starting point a standard Mincer wage equation augmented with health capital. In this framework the (logged) wage rate, w is determined by the stock of (non-health related) human capital, H comprising education and labor market experience, and the stock of health capital, as measured by the addictive capital stock. As most research in this area is based on cross-sectional data, the stock of addictive capital (accumulated use) is not observed, and is replaced by a measure of current use, D :

$$w = \beta_D^w D + \beta_H^w H + \beta_x^w x + \epsilon^w \tag{1}$$

where x is a vector of other characteristics that impact on wages (such as age, race, sex and marital status), and $\beta = (\beta_D^w, \beta_H^w, \beta_x^w)'$ is a vector of parameters to be estimated and ϵ^w is a stochastic error term (Kaestner, 1991; Mullahy and Sindelar, 1993). The main parameter of interest, β_D^w , measures the impact of cannabis use on wages.

Since the wage rate is assumed to reflect an individual's labor market productivity, cannabis use reduces wages if it reduces productivity by adversely affecting an individual's physical and psychological well-being. However, identification of the impact of cannabis

⁴A more extensive review is provided by van Ours and Williams (2015).

use on wages is complicated by two issues. The first is reverse causality. This arises since cannabis use is a consumption good, and its demand depends on income and hence wages. Assuming that cannabis use is a normal good, higher wages (and hence income) will lead to increased cannabis use.⁵

The second complication in identifying the effect of cannabis use on wages is the presence of common confounders that induce correlation between error term ϵ^w and the indicator for cannabis use, D . For example, individuals with a high rate of time preference are more likely to use drugs since they place a lower value on the potential future adverse consequences (Becker and Murphy, 1988). They are also more likely to be in jobs with lower wages, since they place a higher value on a job offer at a slightly lower wage today than a job offer at a higher wage tomorrow. A final issue in identifying the causal effect of cannabis use on wages is that of endogenous selection into employment.⁶

The most common method of addressing the potential for reverse causality and unobserved common confounders in this literature is instrumental variable (IV) estimation. Some studies also address the issue of endogenous selection into employment using a Heckman selection type approach, which also requires instruments for identifying the selection term. For example, Kaestner (1991) implements an IV approach augmented with a Heckman selection procedure using cross-sectional data from the 1984 wave of the National Longitudinal Survey of Youth 1979 (NLSY79). He finds that a higher frequency of cannabis use in the past 30 days leads to higher wages, and that this positive relationship does not diminish with age. Using the same data for males, and also instrumenting current cannabis use and using a Heckman selection procedure to account for selection into employment, Register and Williams (1992) find that a higher frequency of cannabis use in the past 30 days increases wages while uninstrumented long term use and use on the job reduce wages. In a follow-up study based on longitudinal data from the 1988 and the 1984 waves of the NLSY79, Kaestner (1994b) shows that his previous findings are confirmed using the 1988 cross-section. However, using both the 1984 and 1988 waves and longitudinal techniques that account for time

⁵A further reason for reverse causality in the relationship between wages and cannabis use comes from the home production model of Michael and Becker (1973) and Becker and Stigler (1977), in which cannabis use is combined with time and purchased inputs in home production of the utility generating “good” (Kaestner, 1991).

⁶Specifically, those who are working are likely to differ in unobserved dimensions from those who are not, and these unobservables are likely to be correlated with wages, and potentially cannabis use. To the extent that these unobserved factors are time invariant or slow to evolve, they will also determine labor market experience, rendering it endogenous. Similar issues arise with respect to the education component of human capital.

invariant unobserved heterogeneity (in addition to IV and selection correction), the effect of cannabis use on wages is found to be insignificant and negative in all specifications for males and mostly negative and always insignificant in specifications for females. Burgess and Propper (1998) also use data on males from the NLSY, but consider the impact of cannabis use in 1980, when respondents are 16–22, on earnings 10 years later, on earnings growth over the period 1981-1992, and on mean earnings over the period 1981-1992. Using OLS and logit models, they find that early cannabis consumption has no harmful effect on economic prospects in later life.

A key concern with interpreting the findings from studies based on IV estimation as causal effects lies in the choice of instruments. This is especially so in the literature seeking to identify the causal effect of cannabis use on wages. For example, Kaestner (1991) uses non-wage income, the frequency of religious attendance, the number of dependent children, and the number of prior illegal acts to instrument for cannabis use. As discussed by Kaestner (1998), all of these variables may affect the wage rate, in which case they are not valid instruments.⁷

In contrast to these studies, van Ours (2007) uses an approach similar to the one we use, in which the decision to start using cannabis is jointly estimated with wages, assuming that unobserved heterogeneity for each process is drawn from a correlated discrete distribution. His analysis, which uses data on inhabitants of Amsterdam, finds that cannabis use has a negative effect on wages. The size of the effect is related to the age of onset, with earlier use of cannabis having a larger negative effect on wages. The types of assumptions this identification strategy relies on are discussed when we present our own estimation strategy.

⁷Instrumental variables for cannabis use employed in previous studies include attendance of religious services, parental education and previous divorce (Register and Williams, 1992), household composition at a young age, frequency of past religious attendance, number of illegal acts, and a measure of perceived self-esteem and of control over the world (Kaestner, 1994a), parent present at age 14, alcoholic parent and decriminalization of cannabis DeSimone (2002).

2.2 Labor Supply and Cannabis Use

Studies of the impact of cannabis use D , on labor supply h , typically start with a labor supply function along the lines of⁸

$$h = \beta_D^h D + \beta_w^h w + \beta_x^h x + \epsilon^h. \quad (2)$$

In practice, labor supply is often measured by an indicator for employment. Some studies use hours worked among those who are employed as the outcome, but this creates the problem of endogenous sample selection, since those who select into employment do so by comparing their market wage offer to their reservation wage. And the unobservables that determine whether one works are also likely to determine the extent to which one works, as measured by the hours of work. In terms of identifying the causal effect of cannabis use on employment, the challenges here are similar to those faced in identifying the impact of cannabis use on wages.⁹

In the empirical literature, an IV approach is typically employed to address the endogeneity of cannabis use, and sometimes also wages. For example, Kaestner (1994a) estimates the impact of frequency of cannabis use on hours worked using data from the 1984 and 1988 waves of the NLSY79. The endogeneity of cannabis use and the wage rate is addressed by instrumenting, and selection into employment is accounted for using a Heckman selection correction approach. Both cross-sectional and longitudinal estimates are provided, and both generally find negative but statistically insignificant effects of cannabis use on hours worked. In contrast, Register and Williams (1992) who also use data on males from the 1984 wave of the NLSY79 and correct for selection into cannabis use and employment, find that a higher frequency of past month cannabis use decreases the probability of being employed, whereas long term use and on the job use increase it. Burgess and Propper (1998) find that cannabis

⁸This equation can be interpreted as the compensated conditional labor supply function, derived assuming that consumers minimize the cost of achieving a certain level of utility, U conditional on a given quantity of drug consumption, D . An alternative justification for cannabis use directly entering the labor supply function is found in the household production framework. In this framework, individuals combine time, market inputs and cannabis use to produce “health”, where cannabis use has the effect of reducing the stock of health. Since health capital is a determinant of labor supply, cannabis use affects labor supply via its effect on health (Kaestner, 1994a).

⁹Specifically, unobserved confounders that determine hours worked, such as preference for leisure, are also likely to determine cannabis use. These common unobserved confounders render cannabis use endogenous to labor supply. Additionally, as discussed previously, these unobserved characteristics are likely to determine wages, and so wages are also endogenous to labor supply. Reverse causality whereby hours worked leads to cannabis use is unlikely to be an issue.

use in youth is significantly associated with a lower probability of employment 10 years later. DeSimone (2002) considers the impact of cannabis use on employment for males from the 1984 and 1988 waves of the NLSY79 separately. He finds that cannabis use has a substantial negative effect on the employment of males, reducing the probability of employment by 15–17 percentage points. Zarkin et al. (1998) use the 1991 and 1992 waves of the U.S. National Household Survey on Drug Abuse (NHSDA) data to study the relationship between hours worked and illicit drug use for young men (aged 18 to 24). Using an IV approach, they estimate that light cannabis use (1–3 joints in the past month) increases hours worked based on the 1991 wave, and decreases hours worked based on the 1992 wave of the NHDA, leading them to conclude that their analysis fails to find compelling evidence of a cannabis-labor supply relationship.¹⁰ Finally, van Ours (2006) uses data on inhabitants of Amsterdam to study the employment effects of cannabis use. Using a discrete factor approach in which the uptake of cannabis is estimated jointly with the probability of having a job, the author finds no evidence of a causal effect of cannabis use on employment.

2.3 A Search-theoretic Framework

As discussed above, previous research tends to focus on either the impact of cannabis use on wages, or the impact of cannabis use on labor supply.¹¹ This approach induces issues of endogenous sample selection (into employment when considering either hours worked or wages), or endogenous regressors (such as wages when modeling hours worked) so that reduced form and selection models must be estimated for the outcome that is not of primary interest.¹² These issues are largely overcome when we use a search-theoretic framework to study the impact of early cannabis use on the school to work transition. Search theory focuses on labor market dynamics, in particular the duration of time spent searching for a job and the wage rate at which an offer is accepted and search ceases. This framework is particularly appealing in our setting as we are concerned with young men’s first post-school job, for which the decision to accept a job and the wage rate at which this occurs are two aspects of a single decision.

¹⁰As instrumental variables they use respondents’ assessment of the risk associated with using illicit drugs and their assessment of the difficulty in obtaining illicit substances.

¹¹The few papers that have considered both labor market outcomes conduct separate analyses for each.

¹²Moreover, indirect effects of drug use on wages and labor supply, working through the accumulation of the human capital variables, labor market experience and education, are ignored as is their potential endogeneity.

If we allow for endogenous search effort both the reservation wage and job search effort determine the duration of time it takes for an individual to find a job (Cahuc et al., 2014). The basic setup is that an individual who is unemployed at time $t = 0$ decides on the amount of effort to exert on searching for employment s , where for simplicity, it is assumed that the probability of receiving a job offer is equal to $s \in [0, 1]$. Greater search effort increases the probability of finding a job but does so at cost $c(s)$, where $c'(s) > 0$ and $c''(0) > 0$. Job offers, characterized by a wage rate w , are realizations of a random variable W with a cumulative distribution function, F , where F is assumed to be known to the worker and independent of search effort. An important consequence of this assumption is that search effort determines the probability that a job offer is received by the individual, but not the distribution of wages from which the offer is drawn.

The individual accepts a job offer if it is at least as large as his reservation wage, w^r . If $w < w^r$ the individual will reject the job offer and continue searching. The reservation wage is equal to the flow value of being unemployed¹³, as this is the value for which the individual is indifferent between accepting the job offer and rejecting it. The reservation wage can be expressed as:

$$w^r = b - c(s) + \frac{s}{\delta + q} \int_{w^r}^{\infty} (x - w^r) dF(x) \quad (3)$$

where b is the utility associated with unemployment, including the value of leisure as well as the monetary value of unemployment benefits, δ is the discount rate where a higher discount rate implies greater impatience, and q is the rate of being laid off.

Optimal search effort is chosen such that the marginal cost of increasing the probability of finding a job is equal to the marginal benefit (expected present value of obtaining a job offer greater than the reservation wage):

$$c'(s) = \frac{1}{\delta + q} \int_{w^r}^{\infty} (x - w^r) dF(x) \quad (4)$$

Noting that job offers are accepted if the wage rate offered is at least as large as the reservation wage and that job search effort is chosen to equalize the marginal costs and benefits of search, a stationary job market is characterized by the exit rate out of unemployment θ , defined as

¹³The flow value of being unemployed is composed of the net current period benefit, given by the value of unemployment less search costs, plus the expected present value of obtaining a job offer greater than the reservation wage in the future.

the product of the rate at which job offers are received and the probability that a given job offer is accepted.

$$\theta = s[1 - F(w^r)]. \tag{5}$$

Comparative statics show that in equilibrium, search effort and the reservation wage are decreasing in δ . This implies that greater impatience lowers the reservation wage in addition to lowering the search effort. These two effects exert opposite influences on the job finding rate. DellaVigna and Paserman (2005) show that the wage effect dominates so that the net effect of an increase in impatience (increase in δ) is a lower reservation wage and a higher job finding rate, and therefore a shorter duration of unemployment.

In addition, using equations (3), (4) and (5) it can be shown that an increase in the value of unemployment b increases the reservation wage w^r while decreasing search effort s , each of which serve to decrease the job finding rate θ and hence increase the duration of time to find a job. An increase in the probability of being laid off q decreases both w^r and s , with the former effect dominating and leading to an increase in the job finding rate θ and hence a shorter duration of unemployment. Finally, a reduction in the mean of the distribution of wages, μ reduces both w^r and s , with the impact on search effort dominating leading to a lower job finding rate θ , and hence a longer duration of unemployment.

In our analysis, we jointly model the exit rate out of unemployment, θ and the wage offer at which search ceases and the job offer is accepted, w . In order to investigate the impact of cannabis use on the school to work transition, both the duration of time taken to find a job and the wage at which a job offer is accepted is allowed to depend upon cannabis use while at school. Predictions from the theoretical model provide a means of learning about potential mechanisms via which cannabis use may impact on school to work transitions. In particular, finding that cannabis use increased the job finding rate while decreasing the starting wage suggests that cannabis use increases impatience, or that it increases the job separation rate. If however, cannabis use is found to decrease the job finding rate and increase the starting wage, this would suggest that cannabis use increases the value of leisure. Finally, if cannabis use leads to lower wages and lower job finding rates, this suggests that cannabis use reduces the mean of the wage distribution from which job offers are drawn. This effect may operate through, for example, adverse mental health effects of cannabis use or a reduction in the

quality of human capital accumulated at school (by way of a lower grade point average).¹⁴

Common confounders are a potential source of endogeneity in studying the impact of early cannabis use on school to work transitions. We account for common time invariant confounders by modeling the dynamics of cannabis jointly with the duration of time taken to find a job and the starting wage, allowing for mixing of the distributions of (time invariant) unobserved heterogeneity terms affecting cannabis use dynamics and labor market outcomes. Importantly, our focus on the first post-school job obviates issues related to the potential endogeneity of labor market experience since school leavers have no such experience.

An important and novel aspect of our approach is that, in a sensitivity analysis, we account for the possibility that post-school durations of unemployment end with a transition back to education, rather than a transition into employment. As with the transition into employment, cannabis use while enrolled in school is permitted to affect the transition rate of returning to school. In this way, we account for the potential endogeneity of completed education as well as illuminating a yet unexplored mechanism via which cannabis impacts on educational attainment.

3 Data

3.1 NLSY97

Our analysis is based on males from the National Longitudinal Survey of Youth 1997 (NLSY97). The NLSY97 is a longitudinal survey representative of American youths who were 12–16 years old as of 31 December 1996. It is designed to capture educational and early labor market experiences of youth, but also collects detailed information on cannabis use. The first round of the survey was fielded in 1997 and 1998. Subsequent rounds of the survey have been conducted annually, and the retention rate for wave 15, fielded in 2011, is 85.6%.¹⁵ There were 8,984 individuals interviewed in wave 1 of the NLSY97, 4,599 of whom

¹⁴There is substantial evidence that heavy long term cannabis use is associated with poor mental health, especially among those who start using early, and use for a long period of time (Macleod et al., 2004; van Ours et al., 2014). There is also evidence that poor mental health is associated with lower wages (Frank and Gertler, 1991) and reduced employment, consistent with lower exit rates (Chatterji et al., 2011; Tekin and Markowiz, 2008).

¹⁵The interviews are conducted using a computer-assisted personal interview instrument, administered by an interviewer with a laptop computer. The mode of interview is in person or by telephone. When interviews were conducted in person, the information on drug use was self-administered. Personal interviews constitute the bulk of data.

were male.¹⁶ Our sample comprises 3,830 males for whom we observe information on both school leaving and whether or not cannabis is used for at least one round of the survey, as well as information on wages for those who report starting a post-schooling job.¹⁷ The age range in our sample, 12 to 18 in 1997 through to 26 to 32 in 2011, covers the periods in the life cycle during which cannabis uptake and the transition from school to work typically occur.

3.2 Outcomes of interest

Information from the baseline survey (round 1) and the fourteen annual follow-ups conducted up to and including 2011 (round 15) are used to construct individual histories on cannabis use, school leaving and first labor market experiences (the duration of time until starting the first post-school job and initial hourly wage in that first job).¹⁸ The starting point for our analysis is school leaving, where school is defined to include high school, college and post-graduate study. Individual's first school exit is defined to be the first month they report not being enrolled in school and not being on vacation, expelled or suspended. To ensure that we do not include summer vacations as school exits, we additionally require that a school exit has a length of time not enrolled in study of at least 3 months.

In our analysis of the dynamics of cannabis use, we model both the age at first use and the duration of use among users. The age at first use is self-reported for those who report having ever used cannabis in wave 1. For individuals who do not report using in wave 1, the age at first use is based on information provided in subsequent waves on whether the respondent had used cannabis since the date of the last interview (SDLI).¹⁹ The duration of use is

¹⁶Of the 8,984 originally surveyed, 6,748 composed a representative cross-sectional sample and 2,236 compose a supplement that over-sampled Hispanics and blacks who were born in the same time period as the representative sample. We use both in our analysis.

¹⁷We remove from the sample individuals for whom we observe no information on school leaving (N=220) or cannabis use (N=37), for whom we cannot determine if they used cannabis before leaving school (N=84), those for whom we cannot determine the highest grade completed before leaving school (N=14), those who report leaving school but leave the sample within 3 months without having found a job or returning to school (N=82), and those who report starting a job but have missing information on wages or whose hourly wage rate when they start their first post-school job is greater than \$50 (N=307).

¹⁸The information used in this paper was reported by the youth. The only exception is parental characteristics which were reported by the youths' parents in 1997.

¹⁹Harrison et al. (2007) investigate the the validity of self-reported drug use information in a general population survey in the US by comparing the self-reports of respondents with the results of drug tests of urine specimens. For cannabis, 93% of individuals self-reported use in the past 3 days agreed with their urine sample, 5.2% of those who reported no use tested positive and 1.8% reported using and tested negative. The authors also investigates whether there were differences in truthfulness of self-reports by age, comparing

calculated using information on the age at which the respondent last reported cannabis use combined with the age at first use.²⁰ The focus of our study is the impact of using cannabis while enrolled in school and the duration of use while at school. The indicator for cannabis use while enrolled in school is defined to be equal to one if initiation into cannabis use occurs prior to school leaving, and zero otherwise. The duration of cannabis use while at school is constructed using information on age at first use, age at last use, and the age at which the respondent exits from school (for the first time).

In terms of the school to work transition, we consider the first new job reported after individuals left school. Duration of time until employment is constructed by first calculating the number of months that elapsed from the month the respondent left school until the month he reports starting employment at a new job. Durations are censored for individuals who returned to school before finding an eligible job and individuals who had not found a job when they were last interviewed. The duration until first post-school job is defined in terms of quarters of years and constructed by dividing the months taken to find a job by 3. The wage rate at which a job is accepted is measured by the first post-school job's initial hourly wage excluding overtime pay or commissions. Information on the hourly wage for people who are paid by the hour is based on the question: "What was your hourly rate of pay when you first started this job?" For those who were non-hourly paid, weekly earnings were divided by hours worked per week to obtain the hourly wage rate.

Our extended analysis models the hazard rate of returning to school. For those who return to school rather than start working, the duration of time until returning to school is defined as the number of months elapsed from the month they first exited school to the month that they returned to school. Note that the minimum duration until returning to school is 3 months, since the definition of school exit requires individuals to not be enrolled in school for at least 3 months. We convert the duration until returning to school from months to quarters of years.

Note: The sample consists of 3830 males; ***, **, * indicates significant at a 1, 5, 10 percent level.

12-17 year olds with those aged 18-25. Overall, both age-groups reported their drug use accurately, with 3% under-reporting 3 day use among 12-17 year olds compared to 8% under-reporting 3 day use 18-25 year olds.

²⁰The 37 individuals who reported to have used cannabis before age 10 were coded as missing.

Table 1: Summary Statistics

	Full Sample	Used cannabis at school		
		No	Yes	Difference
Cannabis use ever	0.63	0.35	1.00	0.65***
Cannabis use before school leaving	0.43	0.00	1.00	–
Cannabis use before job finding	0.49	0.10	1.00	0.90***
Duration of cannabis use at school	1.63	0.00	3.82	3.82***
Duration of cannabis use lifetime	6.35	4.25	7.35	3.10***
Age left school	18.73	18.59	18.92	0.33**
Highest grade completed	12.13	12.04	12.24	0.20***
Started work	0.79	0.76	0.82	0.06***
Duration till found job (quarters)	4.26	4.38	4.09	-0.30
Duration till found job (quarters) work	3.64	3.82	3.42	-0.40*
Starting hourly wage work	9.14	9.12	9.17	0.05
Back to school	0.18	0.21	0.14	-0.07***
Duration back to school (quarters) back to school	4.12	4.23	3.92	-0.30

3.3 Descriptive Statistics

Table 1 provides descriptive statistics of the key variables for the full sample and according to whether or not cannabis was used before leaving school (where school is defined to include tertiary education). It shows that 63% of our sample have used cannabis in the sample period, and that 43% used cannabis before they left school. Among those who have ever used cannabis, the average age at first use is 16.6 years of age. The table also shows information on duration of cannabis use. For example, it shows that the average duration of use before leaving school among those who initiated use while in school is 3.8 years. The overall duration of use among those who initiated use while in school is 7.4 years, compared with 4.3 years for those whose uptake occurred after they left school. In terms of education, sample members leave school (for the first time) on average at around 18.7 years of age, with 12.1 years of education.

The table also shows that 49% of the sample used cannabis before finding a job, which means that 6% started cannabis use in the period between leaving school and finding their first job (or, for those who did not find a job, before leaving the sample or going back to school). Given that most of the people that use cannabis do so before leaving school for the first time, we model the job search process as a function of cannabis use before exiting school. Modeling job search as a function of cannabis use before finding a job would yield similar results because only a small percentage of people initiate into cannabis use after they

left school and before finding a job.

Finally, Table 1 provides information on the labor market outcomes, duration between leaving school for the first time and starting work (measured in quarters) and the hourly wage rate received when they start their first job. It shows that, among those who begin working, cannabis users take 0.4 fewer quarters (on average) to find a job and earn a slightly higher initial wage, although the difference in hourly wage rates is not statistically significant. Table 1 also shows that 18% of the men in our sample exit post-school unemployment by returning to school rather than finding a job. Cannabis users are less likely to go back to school (before finding a job) than non-cannabis users (14% compared to 21%). Conditional on going back to school the average duration until returning to school is 4.1 quarters, or just over one year. Although cannabis users who return to school (before finding a job) appear to do so more quickly than non-users who return to school (3.9 quarters compared to 4.2 quarters), the difference is not statistically significant.

We provide descriptive statistics for the rich set of control variables used in the analysis in Appendix A.

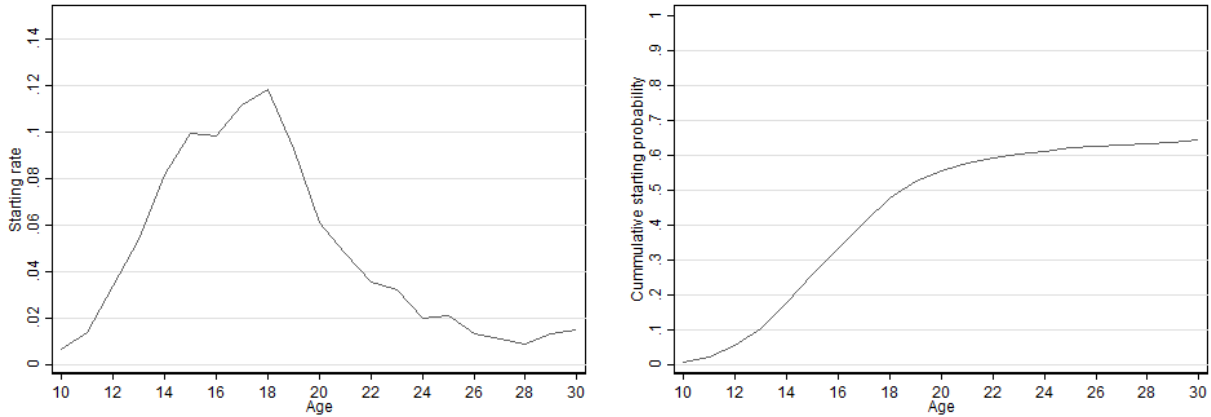
3.4 Stylized facts

As discussed above, in the self-administered portion of the NLSY97 questionnaire, respondents are asked about their cannabis use. In Round 1, respondents are asked whether they have ever used cannabis, and if so, the age at which they first used it. In subsequent rounds they are asked if they have used cannabis since the date of the last interview. This information is used to calculate age-specific starting rates and duration of use specific quit rates, which are graphed in Figure 1.

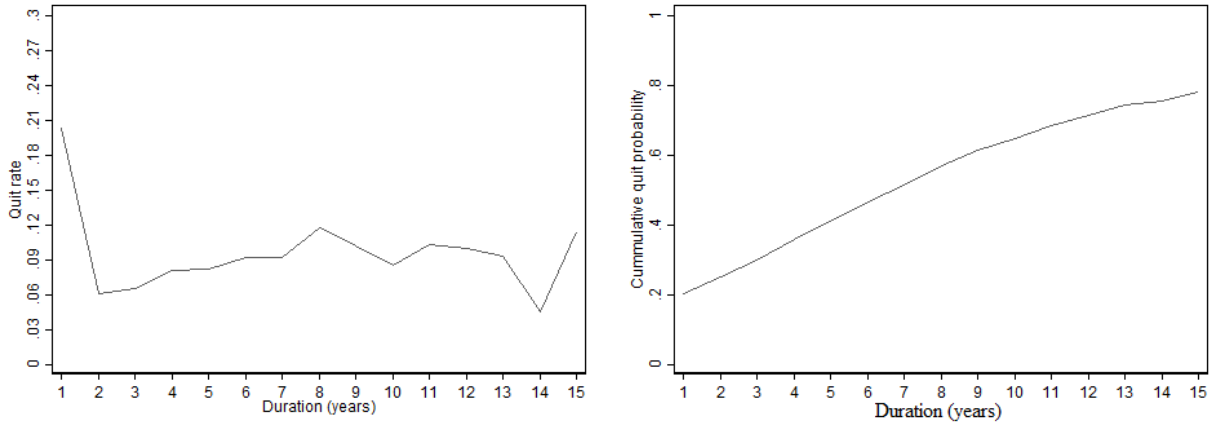
The top panel of Figure 1 contains graphs of the hazard rate for starting cannabis use (left) and the cumulative starting rate for cannabis use (right) at each age. The hazard rate is the probability of starting cannabis use at each age conditional on not having been a user up to that age and the cumulative probability of starting cannabis shows the proportion who have started using at each age. Individuals are assumed to be at risk of cannabis use from the age of 10. An individual who is never observed to use cannabis is considered to have an incomplete duration of non-use, i.e. is assumed to be a ‘right censored’ non-user. Figure 1 shows the starting rate for cannabis use has spikes in uptake at age 15 and 18. The top right panel of Figure 1 shows the cumulative starting probability of cannabis use for males

Figure 1: Cannabis Starting Rates and Quit Rates by Age

a. Starting rates and cumulative starting probability



b. Quit rates and cumulative quit probability



increases from about 20% at age 14 up to 60% at age 22. After that age, the cumulative starting probability hardly increases.

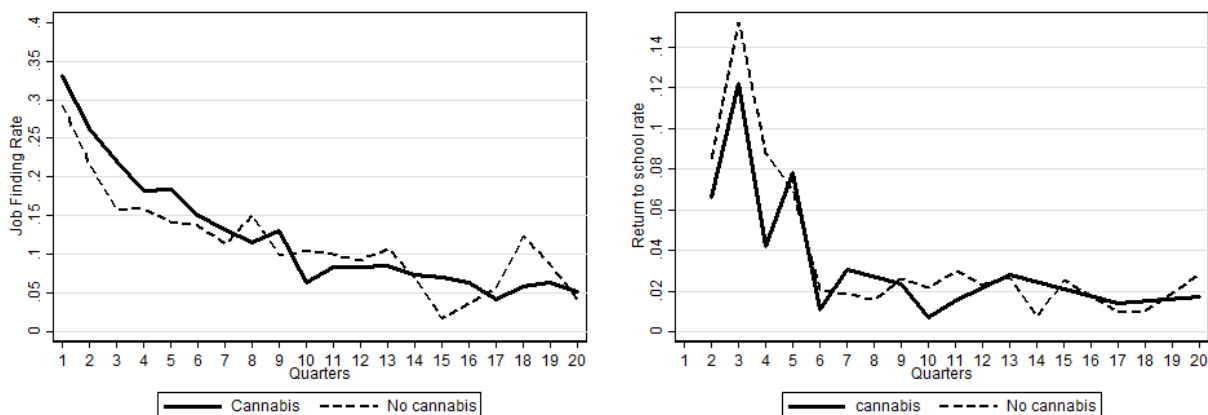
The bottom panel of Figure 1 provides information on the duration of use among those who have ever used cannabis. Similar to the top graphs of Figure 1, the bottom graphs show the hazard of quitting use (left) and the cumulative hazard of quitting use at each duration of use (right). The spike in the hazard for quitting after one year of use suggests that a large proportion of individuals experiment with cannabis use for a short time only. However, 15 years after first using, 20% are still using cannabis (80% have quit).

Overall, this descriptive analysis suggest there are at least three types of individuals (with respect to cannabis use) present in our data. Those who never use cannabis, those who do

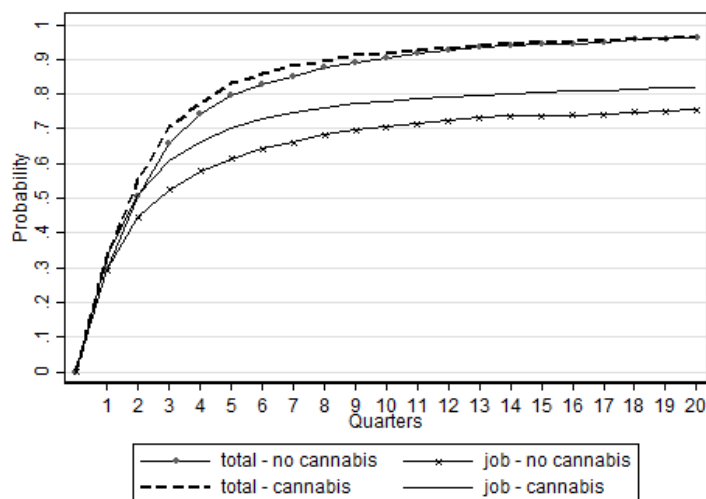
use cannabis but only for a short time (experimenters), and those who are more persistent (non-experimental) users.

Figure 2: Job Finding Rates and Back to School Rates

a. Job finding rates and back to school rates



b. Cumulative job finding and back to school probability



The job finding rate as a function of duration since leaving school and entering unemployment (measured in quarters) is shown in the top left panel of Figure 2. The job finding rate is graphed separately for those who used cannabis before leaving school and for those who did not. These graphs show that in the first 7 quarters after school leaving, those who used cannabis while in school find jobs more quickly than those who did not. The top right

part of Figure 2 shows the rates by which men return to school.²¹ Reaching a peak of only 0.15 for non-cannabis users and 0.12 for cannabis users, the hazard rates of returning to school are substantially lower than the job finding rates, and after six quarters they are close to zero. Having left school, it seems that men either return to school within a year and a half or they are very unlikely to do so.

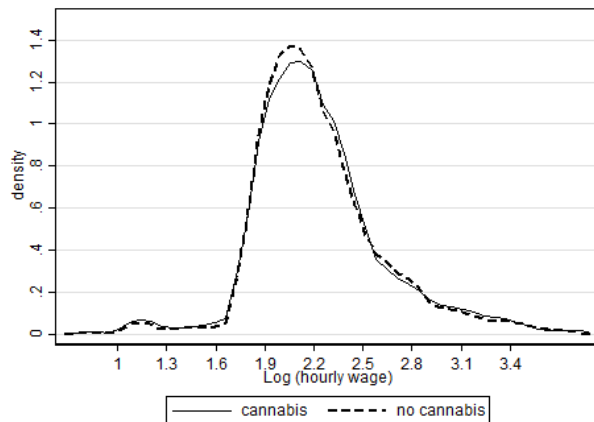
The bottom panel of Figure 2 shows the cumulative probabilities of finding a job and returning to school as a function of duration since leaving school and entering unemployment (in quarters). These are graphed separately for those who used cannabis before leaving school and for those who did not. The top lines show the overall cumulative probability of transitioning out of unemployment to either a job or school, while the lower lines show the cumulative probability of transitioning from unemployment to employment. The difference between the overall cumulative transition probability and the cumulative job finding probability reflects the cumulative probability of returning to school. As is shown in the figure, within a year (4 quarters) of leaving school, around 77% of cannabis users and 74% of non-users have transitioned out of unemployment, with 66% of cannabis users and 58% of non-users finding a job, and 11% of cannabis users and 16% of non-users returning to school. After 4 years (16 quarters), there is no difference in the overall cumulative transition rate out of unemployment between users and non-users of cannabis, with 95% of men having exited unemployment. There are differences by cannabis user status on the composition of the exits, with 81% of cannabis users and 74% of non-users exiting to employment while 14% and 21% of cannabis users and non-users, respectively returning to school. Clearly after some time most men find a job and a smaller proportion return to school. However, the proportion of men returning to school is sufficiently large that if ignored, it may result in misleading conclusions regarding the impact of early cannabis use on school to work transitions. This issue will be returned to in our sensitivity analysis.

Figure 3 plots the kernel estimate of the density and the cumulative distribution of (log) hourly starting wages separately for those who used cannabis while at school and those who did not. The graphs reveals no obvious evidence of a difference in the starting wages for those who used cannabis while in school compared to those who did not.

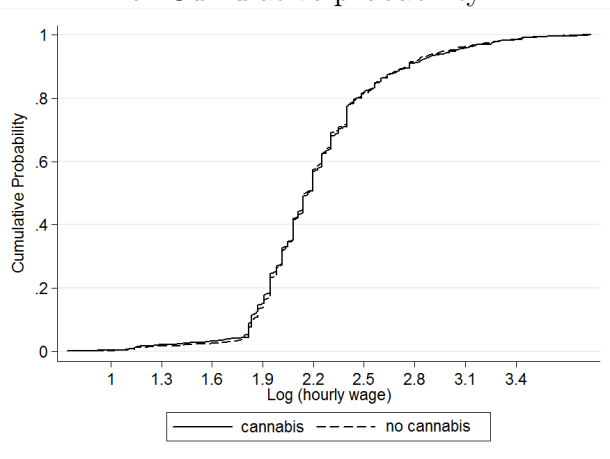
Overall, this descriptive analysis suggests that those who use cannabis while enrolled in school find jobs more quickly and are less likely to return to school than those who do not use

²¹Recall that the definition of school leaving requires an individual to not be enrolled for at least 13 weeks. For this reason, the minimum duration until returning to school is 13 weeks.

Figure 3: Kernel Plots of Log Hourly Wages
a. Density



b. Cumulative probability



cannabis. There does not appear to be any difference in the first post-school wage between cannabis users and non-users on the basis of the raw data. With these stylized facts in mind, we next describe the empirical framework we use to account for common unobserved confounders as well as observed characteristics of our sample members, in order to determine the causal effect of early cannabis use on the school to work transition.

4 Set-up of the empirical analysis

The main challenge we face in establishing the impact of cannabis use on transitions from school to work is that both selection into cannabis use while in school and the post-school

transition to employment may be affected by circumstances and characteristics that are not observed. These common unobserved ‘confounding’ factors may be a source of spurious association and must be accounted for in order to identify the causal impact of cannabis use. In this paper, we use a non-parametric maximum likelihood estimation (NPMLE) approach in which the joint distribution of unobserved heterogeneity on which selection is based is accounted for non-parametrically using a discrete factor approximation (Abbring and Van den Berg, 2003; Gaure et al., 2007; Heckman and Singer, 1984; Mroz, 1999). The main advantages of the NPMLE discrete factor approach in our application are that it provides reliable estimates in models with endogenous variables without imposing arbitrary distributional assumptions on the unobserved heterogeneity, it retains the efficiency of maximum likelihood estimators, and it does not rely on instrumental variables which are difficult to find.²² In a Monte-Carlo evaluation of discrete factor approximations in simultaneous models Mroz (1999) found that when the true distribution of unobserved heterogeneity is not known, and in the absence of instruments, the discrete factor approximation outperformed two-stage least squares and MLE estimators that attempted to control for endogeneity.

Our baseline specification uses NPMLE in the context of a four equation system. This system comprises the hazard of starting cannabis use, the hazard of quitting cannabis use, the hazard of starting a job, and a log linear wage equation. In order to account for endogenous selection into cannabis use and the inter-related nature of the decision to accept a job offer and the wage rate at which this occurs, each equation includes a time invariant unobserved heterogeneity term that is drawn from a joint discrete distribution. In addition to being time invariant, the unobserved heterogeneity terms are assumed to be uncorrelated with variables measuring observed heterogeneity. This implies that the sole source of endogenous selection into cannabis use while at school is time invariant unobserved heterogeneity. In addition, in the case of job finding, cannabis uptake and quitting, we assume a multivariate mixed proportional hazard structure. Similar to Mroz (1999), identification of unobserved heterogeneity in the wage equation relies on its linear functional form (see for a recent example Arni et al. (2013)). Given that our approach controls for unobserved heterogeneity that is time invariant, any unobserved characteristics that affect both cannabis use dynamics and labor market transitions must be present and unchanging from the time individuals are first at risk of cannabis uptake (age 10 in our sample), whereas the labor

²²This approach has previously been used in a wide variety of applications in health and labor economics (see for example, Cutler (1995); Bray (2005); van Ours (2006, 2007); Yang et al. (2009)).

market impacts of cannabis use occur subsequently, after leaving school. Since we study the impact of a treatment (cannabis use while at school) that occurs before entering post-school unemployment on transitions out of unemployment, we are unable to leverage off dynamic selection to assist in identifying causal effects from selection. As a consequence, our identification strategy relies on the mixed proportional hazard assumption (along with linear functional form and normality of idiosyncratic shocks assumptions for the wage equation).²³ However, we do not impose functional form assumptions on duration dependence or on the distribution of unobserved heterogeneity for the hazard rates for cannabis uptake, quitting, and job finding, or for unobserved heterogeneity in the wage equation. We examine the robustness of the results to these identifying assumptions in a sensitivity analysis in Section 5.5.

In the following sections, we build up our empirical model in three steps. First, we model cannabis uptake and quitting, collectively referred to as cannabis use dynamics. Next, we model the hazard of job finding and the starting wage, ignoring selection into cannabis use and the jointness of the decision to accept a job and the wage at which an offer is accepted. We then bring cannabis use dynamics together with the equations for labor market outcomes in order to account for endogenous selection in estimating the causal effect of early cannabis use on the school to work transition.

4.1 Cannabis use dynamics

The transition rate into cannabis use is modeled as a mixed proportional hazard with a flexible baseline hazard. Differences between individuals in the rate at which they start using cannabis are assumed to depend upon observed characteristics, elapsed duration of time they are potentially exposed to use and unobserved characteristics. Age 10 is assumed to be the time at which individuals are first exposed to cannabis. The starting rate for cannabis at age t^s conditional on observed characteristics x and unobserved characteristics u^s is specified as (omitting a subscript for individual)

$$\theta^s(t^s | x, u^s) = \lambda^s(t^s) \exp(x'\beta + u^s) \tag{6}$$

²³Functional form and distributional assumptions are commonly relied upon in structural modeling. See for example, Alford (2015).

where the observable characteristics that we control for are characteristics of the individual: race (with indicators for race is black and race is Hispanic, non-black and non-Hispanic as the as the reference category), ability (standardized schooling-corrected CAT-ASVAB score)²⁴, and religiosity (with an indicator equal to one if the respondent reported a religious affiliation in 1997); characteristics of the household (log of household income and log of household size in 1997); presence of parents (with separate indicators for presence of mother and presence of father); parental education (with separate indicators for mother’s and father’s educational attainment is high-school graduate and separate indicators for mother’s and father’s educational attainment is greater than high school graduate, less than high school graduate is the reference category for each parent); mother’s age when she gave birth; and an indicator for mother’s parenting style is authoritarian; and year of birth.²⁵ $\lambda^s(t^s)$ represents individual age dependence, and the superscript, s refers to starting cannabis use. We model age dependence flexibly by using a step function:

$$\lambda^s(t) = \exp(\sum_k \lambda_k^s I_k(t)) \tag{7}$$

where k ($= 1, \dots, N$) is the subscript for age-interval and $I_k(t^s)$ are time-varying dummy variables for consecutive age-intervals. Age intervals are specified to be one year from age 10 up until age 25, and the last interval refers to ages over 25. Because we estimate a constant term, λ_1^s is normalized to 0.

The conditional density functions for the completed durations until starting cannabis use can be written as

$$f^s(t^s | x, u^s) = \theta^s(t^s | x, u^s) \exp\left(-\int_0^{t^s} \theta^s(s | x, u^s) ds\right) \tag{8}$$

²⁴The standardized schooling-corrected Armed Services Vocational Aptitude Battery (CAT-ASVAB) score has a mean of zero and a standard deviation of one in 1999. The CAT-ASVAB score is a summary of four tests taken: Mathematical Knowledge, Arithmetic Reasoning, Word Knowledge and Paragraph Comprehension. The score is likely to be affected by schooling in the year of the test. To correct the CAT-ASVAB score for schooling, we use a method based on Hansen et al. (2004). This procedure exploits the randomness in the amount of schooling undertaken at the test date, since all the respondents took the CAT-ASVAB in the same year. Our schooling-corrected test score is the residual of a regression of the CAT-ASVAB score on years of education completed in the year of the test, holding completed schooling at the time of first school leaving constant (Carneiro et al., 2011, 2013)

²⁵Variables that indicate personal characteristics such as marital status and presence of children, are not very useful because, in addition to being potentially endogenous, they do not reflect circumstances at the time individuals first face the decision of whether or not to use cannabis (or, conditional on using cannabis, whether or not to stop using. As we do not have access to information on state of residence, we are not able to control for the policy regime faced in modeling the uptake and quitting transitions.

Individuals who have not used cannabis by the time they are last observed are assumed to have a right-censored duration of non-use.

The quit rate is also assumed to have a mixed proportional hazard specification. The quit rate for cannabis at duration of use, t^q conditional on observed characteristics, z and unobserved characteristics, u^q is specified similarly as

$$\theta^q(t^q | z, u^q) = \lambda^q(t^q) \exp(z'\gamma + u^q) \quad (9)$$

where z contains the age at which the individual started cannabis use in addition to the variables contained in x , $\lambda^q(t^q)$ represents individual duration dependence and the superscript, q refers to quitting cannabis use.²⁶ Duration dependence is again modeled as piece-wise constant,

$$\lambda^q(t^q) = \exp(\sum_m \lambda_m^q I_m(\tau)) \quad (10)$$

where m ($= 1, \dots, M$) is the subscript for duration of use interval and $I_m(t^q)$ are time-varying indicator variables that are equal to one in consecutive duration intervals from 1 year to 18 years, and the last interval is for 19 years or more. The conditional density functions for the completed durations of cannabis use can be written as

$$f^q(t^q | z, u^q) = \theta^q(t^q | z, u^q) \exp\left(-\int_0^{t^q} \theta^q(s | z, u^q) ds\right) \quad (11)$$

Individuals who have started cannabis use and are still using at the time they are last observed have a right-censored duration of use.

In order to allow for correlation across uptake and quitting decisions we specify the joint density function of the duration until cannabis uptake and the duration until quitting cannabis use conditional on z and x as

$$f^{sq}(t^s, t^q | x, z) = \int_{u^q} \int_{u^s} f^s(t^s | x, u^s) f^q(t^q | z, u^q) dG(u^s, u^q) \quad (12)$$

where $G(u^s, u^q)$ is assumed to be a discrete distribution. In theory, this discrete distribution has an unknown number of points of support. In practice, we are able to identify three

²⁶Note that quits are assumed to be permanent. Once individuals have decided to quit cannabis use they don't return to use again in this model.

points of support in the joint distribution, (u_1^s, u_1^q) , (u_1^s, u_2^q) , (u_2^s) , with $u_2^s = -\infty$ implying a zero starting rate. The specification of the distribution of unobserved heterogeneity implies that conditional on the observed personal characteristics (including age and duration of use) there are three types of individuals. The first type represents the “experimenters” who have a positive starting rate and a high quit rate. The second type represents the “persistent users” who have a positive starting rate and a low quit rate. Individuals in this group use cannabis for longer than those in the first group. The third type are “abstainers”. They have a zero starting rate, and therefore their quit rate is non-existent. The associated probabilities are denoted as

$$\Pr(u^s = u_1^s, u^q = u_1^q) = p_1 \quad \Pr(u^s = u_1^s, u^q = u_2^q) = p_2 \quad \Pr(u^s = u_2^s) = p_3 \quad (13)$$

and are assumed to have a multinomial logit specifications with $p_n = \frac{\exp(\alpha_n)}{\sum_n \exp(\alpha_n)}$, with $n = 1, 2, 3$ and α_3 normalized to zero. Details of the specification of the likelihood are given in Appendix A1.

4.2 Job finding and Starting Wages

We focus on two inter-related aspects of the transition from school to work: (1) the time it takes from leaving school to find a job and (2) the initial wage rate for the first post-school job.

Beginning with the transition into employment, differences in the rate at which individuals who have left school find a job are assumed to depend on observed characteristics, x_j , the elapsed duration of time since they left school, t^j and unobserved characteristics, ν^j . The starting rate for job finding at time t^j conditional on observed characteristics, x^j and unobserved characteristics, ν^j is specified as

$$\theta^j(t^j | x_j, \nu^j) = \lambda^j(t^j) \exp(x_j' \beta_j + \nu^j) \quad (14)$$

where x_j includes an indicator for cannabis use while enrolled in school (equal to one if age of initiation occurs before first school exit), the educational attainment of the respondent (indicators for educational attainment is high school graduate, and educational attainment is greater than high school graduate, with educational attainment is less than high school graduate as the reference category), and information on the location of the respondents

residence (indicators for urban location when the respondent left school for the first time, and indicators for geographic region of residence when the respondent left school for the first time: west, north east, north central parts of the US), and the observed characteristics controlled for in the cannabis use dynamics. $\lambda^j(t^j)$ represents duration dependence, which is modeled flexibly using a step function

$$\lambda^j(t^j) = \exp(\sum_k \lambda_k^j I_k(t^j)) \quad (15)$$

where k ($= 1, \dots, N$) is the subscript for the number of elapsed time intervals since leaving school and $I_k(t)$ are time-varying indicators variables for consecutive intervals. These intervals are specified in terms of individual quarters (3 month periods) for quarters 1 through to 14, an interval from quarter 15 through to quarter 25, and the last interval is for quarter 26 onward. Because a constant term is estimated, λ_1^j is normalized to 0.

The conditional density function for the completed durations of unemployment until finding a job is

$$f^j(t^j | x_j, \nu^j) = \theta^j(t^j | x_j, \nu^j) \exp\left(-\int_0^{t^j} \theta^j(j | x_j, \nu^j) dj\right)$$

Individuals who have not found employment by the end of the observation period are assumed to have a right-censored duration of time until finding a job. Integrating out the unobserved heterogeneity component, we obtain the duration until job finding conditional on x_j as:

$$f^j(t^j | x_j) = \int_{\nu^j} f^j(t^j | x_j, \nu^j) dH(\nu^j)$$

where $H(\nu^j)$ is assumed to be a discrete distribution with g points of support. In the empirical application, we are able to identify two points of support for this distribution, implying that (conditional on observed characteristics) there are two types of individuals: the first type have a high rate of job finding (and so leave unemployment relatively quickly), while the second type have a lower rate of job finding (and so spend longer in unemployment). The associated probabilities, denoted $\Pr(\nu^j = \nu_1^j) = p_4$ and $\Pr(\nu^j = \nu_2^j) = 1 - p_4$, are modeled using a logit specification.

Individuals who accept a job offer do so at the initial wage offer. The log (hourly) wage rate received upon accepting a job offer and starting employment is modeled using the

standard human capital approach as follows:

$$w = x'_w \beta_w + \nu^w + \epsilon, \quad \text{where } \epsilon \sim iid N(0, \sigma^2) \quad (16)$$

where the error term is composed of the unobserved heterogeneity term ν^w and the idiosyncratic shock, $\epsilon \sim iid N(0, \sigma^2)$, and x_w consists of observed determinants of wages, including an indicator for whether cannabis was used while at school, educational attainment, the industry in which the respondent is employed, the calendar year in which he started work, and the observable characteristics that we control for in the hazard for starting work. We assume ν^w is drawn from a discrete distribution with 2 points of support. The associated probabilities, $Pr(\nu^w = \nu_1^w) = p_5$ and $Pr(\nu^w = \nu_2^w) = 1 - p_5$, are modeled using a logit specification which allows us to integrate out the unobserved heterogeneity from the wage equation.

Note that we have described the three components: (1) cannabis use dynamics; (2) the hazard for starting work; and (3) the wage equation, as independent of one another. Our empirical strategy begins by mirroring this description, estimating each component as an independent process. We then relax the assumption of independence between the components, allowing for correlated unobserved heterogeneity. Conflating the three types of individuals (based on their unobserved heterogeneity) found in the cannabis dynamics, the two types found in the job finding and the two types found in the wage equation produces a distribution of unobserved heterogeneity with up to 12 points of support in the fully correlated system of four equations.

5 Parameter estimates

The parameter estimates are obtained using the method of non-parametric maximum likelihood, and account for the fact that our duration information relates to intervals rather than to exact durations.²⁷

²⁷For example, a young man who starts using cannabis at age 16 may start on his 16th birthday or on the day before his 17th birthday. For this individual, we model that he did not yet start at age 15, but started before turning 17.

Table 2: Parameter estimates cannabis use dynamics

	Start Rate		Quit Rate	
Black	-0.37	(5.0)**	-0.07	(0.7)
Hispanic	-0.19	(2.5)**	0.13	(1.2)
Ability	-0.11	(3.3)**	-0.41	(7.6)**
Religious	-0.37	(4.9)**	-0.01	(0.1)
Household income	0.02	(0.5)	0.02	(0.5)
Household size	-0.26	(2.9)**	0.08	(0.6)
Mother education high school	0.01	(0.2)	-0.09	(0.8)
Mother education > high school	0.14	(1.6)	-0.42	(3.4)**
Father education high school	-0.16	(2.2)**	0.23	(2.1)**
Father education > high school	-0.19	(2.2)**	0.14	(1.1)
Mother age at birth/10	-0.09	(1.7)*	0.07	(0.9)
Mother authoritarian	-0.11	(2.0)**	-0.05	(0.7)
Mother present	-0.26	(2.7)**	0.12	(0.8)
Father present	-0.17	(2.6)**	0.12	(1.3)
Age first use/10			2.28	(13.7)**
Unobserved heterogeneity				
Mass-point 1	-3.26	(10.7)**	-7.37	(13.8)**
Mass-point 2	$-\infty$		2.31	(14.7)**
α_1		0.33	(2.6)**	
α_2		0.35	(2.9)**	
-Loglikelihood				13304.0

Note: Cannabis starting rate contains 15 age categories (each age from 10 to 24 years and age 25 or older; the second mass-point is estimated as a difference from the first mass-point; cannabis quit rate contains 19 duration dependence categories (each year from 1 to 18 and 19 or more years). Also included are 4 dummy variables for year of birth. Absolute t-statistics in parentheses; ** (*) indicates significant at a 5 (10) percent level.

Distribution of unobserved heterogeneity		
Starting rate	Quit rate	Prob (%)
Positive	High	37.2
Positive	Low	36.6
Zero	–	26.2
Total		100

5.1 Cannabis use dynamics

The parameter estimates of the cannabis dynamics model are reported in Table 2.²⁸ The general picture that emerges from these estimates is that transition rates into cannabis use are lower for blacks and Hispanics compared to non-black, non-Hispanics (whites), for those who have higher ability, those who are religious, are from larger families, whose mother and father are present, whose mother was older when the respondent was born, whose father's level education is at least a high school graduate, and whose mother's parenting style is authoritarian. In terms of quitting, the analysis confirms that uptake at older ages leads to a higher hazard of quitting and hence shorter duration of use. We also find that, all else being equal, individuals of high ability, who have a mother with more than a high school education, and a father who is a high school graduate have lower hazards of quitting and therefore longer durations of cannabis use.

The results in Table 2 also show that unobserved heterogeneity is important and that three types of individuals can be distinguished. Conditional on observed characteristics, the estimates imply that 37% of males are of the type who have a positive starting rate and a high quit rate (type 1 - the experimenters); 37% of males are of the type who have a positive starting rate and a low quit rate (type 2 - persistent users); and 26% of males are of the type who have a zero starting rate (type 3 - never users). Note that we tried to identify additional mass-points in the discrete distribution of unobserved heterogeneity but did not succeed.²⁹

5.2 Job finding

The parameter estimates for the duration until finding a job in which cannabis use is assumed to be exogenous and unobserved heterogeneity determining the wage rate is assumed to be independent are reported in the first column of Table 3. They show that cannabis use while in school is associated with a job finding rate that is 14% higher than that of an otherwise similar male who did not use cannabis while at school.³⁰ In terms of other characteristics, we find that those with at least a high school education take less time to find a job than those whose educational attainment is less than a high school graduate. Black men take

²⁸The parameters on the age dependence terms in the starting rate and the duration dependence variables in the quit rates are not reported in the table but are available upon request.

²⁹We were not able to improve the likelihood by including an additional mass-point. Gaure et al. (2007) report that the most reliable criteria for determining the number of mass-points is the likelihood itself.

³⁰ $14\% = (100 \exp(0.13) - 1)$

Table 3: Parameter estimates labor market experience

	Job finding		Starting wage	
Cannabis use at school	0.13	(3.0)**	-0.019	(1.8)*
High school graduate	0.20	(3.8)**	0.025	(1.7)*
Education > high school	0.26	(3.2)**	0.214	(8.9)**
Black	-0.14	(2.4)**	-0.035	(2.2)**
Hispanic	-0.09	(1.4)	-0.022	(1.3)
Ability	0.05	(1.8)*	0.009	(1.2)
Religious	-0.06	(0.8)	-0.002	(0.1)
Household income	0.00	(0.1)	0.021	(3.3)**
Household size	0.10	(1.4)	-0.026	(1.4)
Mother education high school	-0.02	(0.3)	0.000	(0.0)
Mother education > high school	0.10	(1.5)	-0.007	(0.4)
Father education high school	-0.02	(0.4)	0.010	(0.6)
Father education > high school	-0.08	(1.1)	0.005	(0.2)
Mother age at birth/10	0.03	(0.7)	-0.025	(2.3)**
Mother authoritarian	0.02	(0.4)	0.000	(0.0)
Mother present	-0.01	(0.1)	0.026	(1.2)
Father present	0.00	(0.0)	0.018	(1.3)
Suburban residence	-0.07	(1.2)	0.043	(2.6)**
σ^2			0.259	(99.1)**
Unobserved heterogeneity				
Mass-point 1	-1.09	(5.3)**	2.50	(27.8)**
Mass-point 2	-4.00	(5.9)**	0.93	(32.5)**
α_1	1.55	(15.2)**	3.26	(26.2)**
-Loglikelihood	7199.78		597.94	

Note: The parameter estimates for the job finding rate are from a duration model; also included are 4 dummy variables for year of birth, 4 dummy variables for region of residence. The estimates for the job finding rate included duration dependence parameters for every quarter up to year 10, 11-15 quarters, 16-25 quarters and more than 25 quarters. The wage estimates also include 15 dummy variables for industry, and 13 dummy variables for year of entering the labor market. Absolute t-statistics in parentheses; ** (*) indicates significant at a 5 (10) percent level.

Distribution of unobserved heterogeneity		
	Job finding Prob (%)	Starting wage Prob (%)
High	82.5	3.7
Low	17.5	96.3
Total	100.0	100.0

longer to find a job than white men, and higher ability is associated with a higher hazard rate of finding a job and hence shorter durations of unemployment. In terms of unobserved heterogeneity, conditional on observed characteristics we are able to identify two mass-points, corresponding to two types of individuals in terms of job finding. The larger group of 83% has a high predisposition for job finding and the smaller group of 17% has a lower predisposition for job finding.

5.3 Initial wages

The second column of Table 3 shows the parameter estimates for the starting wage equation, assuming that cannabis use is exogenous and that the unobserved heterogeneity determining the wage rate is independent of the unobserved heterogeneity determining job finding. The estimates show that cannabis use at school is associated with an initial wage that is 1.9% lower than the wage received by an otherwise similar male who did not use cannabis while at school. Wages are also found to be increasing in educational attainment, household income in 1997, and living in a urban location during the year that they leave school. All else being equal, being black and having an older mother is associated with a lower starting wage in the first post-school job. As shown in the lower part of Table 3 the distribution of discrete unobserved heterogeneity in the wage equation is highly skewed; conditional on observed characteristics, 96% of the sample are of the type who have a low wage relative to the other 4%, who belong to the type who enjoy a relatively high wage.

5.4 Accounting for correlation in unobserved heterogeneity

Table 4 presents estimates on the effect of cannabis use while in school on job finding and initial wages for specifications in which the correlation between the unobserved heterogeneity across equations for cannabis dynamics and labor market outcomes is accounted for. For ease of reference, panel *a* repeats the results assuming that job finding, wages and cannabis dynamics are independent.

Panel *b* presents estimates in which we permit the unobserved heterogeneity determining cannabis use dynamics and labor market outcomes to be drawn from a joint distribution. The lower part of Table 4 shows the distribution of unobserved heterogeneity for this specification. As indicated, we identify 12 points of support in the distribution of unobserved heterogeneity. An LR-test rejects the null hypothesis that the unobserved heterogeneity is

Table 4: Parameter estimates of the effect of cannabis use

	Job Finding		Wages		-LogL	Mass-points
a1. Exogenous cannabis	0.13	(3.0)**	-	-	7199.8	2
a2. Exogenous cannabis	-	-	-0.019	(1.8)*	597.9	2
b. Correlated cannabis	0.08	(1.7)*	-0.017	(1.5)	21075.1	12
c. Age at uptake					21074.2	12
By age 15	0.06	(1.0)	-0.015	(1.1)		
At age 16 or 17	0.12	(1.7)*	-0.028	(1.7)*		
At age 18 or older	0.09	(1.0)	0.001	(0.04)		
d. Duration of use at school					21072.1	12
Up to 1 year	-0.07	(0.9)	-0.004	(0.2)		
More than 1 year	0.12	(2.4)**	-0.020	(1.7)*		

Note: All estimates contain the same explanatory variables as in Table 2 (for cannabis) and Table 3 (for job finding, back to school and wages). Absolute t-statistics in parentheses; ** (*) indicates significant at a 5 (10) percent level.

Distribution of Unobserved Heterogeneity panel (b)				
Cannabis Starting Rate	Cannabis Quit Rate	Job Finding Rate		
		High	Low	
High Wages				
Positive	High	0.6	1.1	1.6
Positive	Low	1.2	0.3	1.5
Zero	-	0.7	0.7	1.4
Low Wages				
Positive	High	30.1	5.4	35.5
Positive	Low	30.5	4.3	34.8
Zero	-	18.7	6.4	25.1
Total		81.7	18.3	100.0

independent.³¹ Comparing the results in panels *a* and *b* shows that accounting for the endogenous selection into cannabis use reduces the magnitude (and statistical significance) of the impact of cannabis use in school on the duration of time it takes to find the first post-school job, and the starting wage in the first post-school job. Specifically, cannabis use while at school is estimated to increase the job finding rate by 14% when selection into cannabis is ignored, and increase the job finding rate by 8% when it is accounted for. Similarly, the magnitude of the estimated effect of cannabis use while in school on the starting wage of the first post-school job is reduced from 1.9% when selection into cannabis use is ignored, to 1.7% when selection is accounted for. Moreover, after accounting for selection into cannabis use, the effect of its use at school on job finding remains weakly significant (at the 10% level) while the effect on wages becomes statistically insignificant at conventional levels (with a t-statistic of 1.5).

Inspection of the distribution of unobserved heterogeneity for this specification confirms that most individuals are of a type that is characterized by being susceptible to having a relatively low wage (35.5%+ 34.8%+ 25.1%=95.4%). In terms of cannabis types, 37.1% (=1.6% + 35.5%) are susceptible to experimenting with cannabis, 36.3%(=1.5%+34.8%) are susceptible to longer term use, and 26.6% (=1.4% +25.2%) are not susceptible to using cannabis. Among those vulnerable to having a low wage type, the most commonly represented types are those who have a high susceptibility to finding a job and a vulnerability to cannabis use, who account for 60.6% (+30.1% +30.5%) of the sample. In terms of cannabis user types, of this 60.6%, experimenters and persistent users are equally represented (30% are experimenters and 31% are persistent users). The balance of this low wage high job finding group, 19% of the full sample, are not susceptible to cannabis use. Among the low wage and low job finding group, experimenters, persistent cannabis users and non-users are fairly equally represented, contributing 5.4%, 4.3% and 6.4%, respectively, to the 16% of the sample in this group. Similarly, experimenters, persistent users and non-users are equally likely among the 5% who are inclined to be in the high wage group, with respective probabilities of 1.6%, 1.5% and 1.4% of the full sample.

Panels *c* and *d* of Table 4 seek to provide a more nuanced understanding of for whom cannabis use at school matters by investigating whether there exist differences in its impact by age at first use (panel *c*) and by whether use at school is experimental (no more than 1

³¹The LR-test comparing the results in Table 2 and Table 4 panel a1 and a2 with Table 4 panel b has a value of 53.2, which is significant at the 1% level; the critical χ^2 -value with seven degrees of freedom is 18.5.

year of use before leaving school) or non-experimental (more than a year while at school) in nature (panel *d*). The results in panel *c* suggest that the detrimental impact of cannabis use on the school to work transition is driven by uptake at ages 16 to 17. However, the LR test comparing specifications *b* and *c* fails to reject the null hypothesis of homogeneous effects by age of uptake.³² In contrast, the LR test comparing specifications in panel *d* and *b* does reject the null hypothesis that the impact of experimental (no more than a year) and non-experimental cannabis use (greater than a year) while in school has equal effects on labor market transitions. In particular, the results from estimating the model allowing for differential effects finds that using cannabis for no more than a year while at school has no impact on either the rate at which school leavers find their first post-school job, or the wage rate at which they accept their first job. In contrast, using cannabis for longer than a year increases the job finding rate by 13% compared to an otherwise similar individual who had not used cannabis at school (or did so for no more than a year), and reduces the wage rate received when starting their first job by 2% compared to their non-cannabis using counterparts.

Overall our findings are that, after accounting for endogenous selection into cannabis use and the potential correlation in unobserved heterogeneity determining job finding and initial wages, non-experimental cannabis use at school (using for more than a year) reduces both the duration of time to find the first post-school job and the starting wage. Experimenting with cannabis while enrolled at school (using for no more than a year) has no impact on school to work transitions. Relating these findings back to the discussion of potential mechanisms in section 2.3, we are able to rule out an increased value of leisure and lower expected wages as the mechanism via which cannabis impacts on young men’s transition to the labor market since these mechanisms imply a greater length of time spent searching for a job, not a shorter length, as we find. Our results are consistent with cannabis use impacting labor market outcomes via a higher discount rate, suggesting that cannabis use increases long run impatience in the labor market, or through increasing the risk of being laid-off. However, given that our analysis focuses on the first post-school job, the former mechanism is more likely than the latter.³³

³²The LR test statistic is 1.8, which is smaller than the critical value for $\chi^2_2(\alpha = 0.05) = 5.99$.

³³An alternative interpretation of our findings is that longer term cannabis users are more likely to be arrested, and as a result of their limited job options, take the first job they are offered and accept lower wages. Unfortunately we are unable to explore this possibility since, although the NLSY97 does contain information on whether an respondent has been arrested, it does not include the offense for which he is

5.5 Sensitivity Analysis

Our first sensitivity analysis examines the issue of identification. As discussed in Section 4 above, in the absence of a dynamic treatment effect, we rely on the mixed proportional hazard functional form in order to identify the causal effect of early cannabis use on the job finding rate, and the log linear function form and normally distributed idiosyncratic shocks to identify the causal effect of early cannabis use on the starting wage. We now revisit this issue, relaxing these identifying assumptions by way of exclusion restrictions.

Specifically, we exclude the indicator for the respondent’s mother having an authoritarian parenting style and the indicator for the respondent expressing a religious preference (each are measured in 1997) from the hazard for job finding and the wage equation. As shown in Table 2, each of these variables are strong predictors of the hazard of starting cannabis (conditional on observed characteristics). Table 3 shows that neither are significantly related to the duration until finding a job or the starting wage (conditional on the other observed characteristics).³⁴ We choose these exclusion restrictions because the variables are not empirically associated with the labor market outcomes (conditional on the other control variables) in our sample while at the same time they are relevant in explaining cannabis uptake. Nonetheless, it is worth noting that religiosity has been used to identify the effect of cannabis use on labor market outcomes in several previous studies (see for example Register and Williams (1992), Kaestner (1994a)).

The purpose of the exclusion restrictions used here is to examine the robustness of our findings to relaxing identification through functional form. The value of the (log) likelihood for the restricted model (imposing the exclusion restrictions) is -21075.6 compared to -21075.1 when no variables are excluded from job finding hazard and the wage equation, with the Likelihood Ratio test statistic indicating that the exclusion restrictions are supported by the data (the chi-squared test statistic with 4 degrees of freedom has a value of 1.0, and corresponding p-value of 0.91). It is therefore unsurprising that the coefficient estimates are

arrested. Nonetheless, we think it unlikely that arrest for cannabis related offenses explains our findings since job offers are less likely for those with a criminal record, and for this reason we expect that cannabis related arrests would increase the time until finding a job, not decrease it.

³⁴In fact, other than cannabis use at school, education, race and ability, there is little evidence that household and other background characteristics are significantly related to the initial labor market outcomes. Exceptions are found in the wage equation, where household income in 1997 and living in a suburban area at the time of leaving school are positively associated with the initial wage rate and having a young mother is associated with a lower starting wage. This suggests that males from higher SES backgrounds may be better informed or have higher expectations regarding initial wages.

unchanged. These results are not reported in the paper but are available upon request.³⁵ The robustness of our estimates to the use of exclusion restrictions to identify the model provides some confidence that our identification strategy is delivering reliable estimates of the causal effect of early cannabis use on school to work transitions.

We next examine the robustness of our results to accounting for school leavers' option to return to school. As discussed in Section 3, a significant proportion of young men in our sample exit post-school unemployment not through finding a job, but by returning to school. Moreover, those who use cannabis while enrolled in school are less likely than non-users to return to school. This poses a potential threat to the interpretation of our findings, since the duration of time spent in job search until finding a job is censored for those who return to school. Thus, the estimated impact of cannabis use on the duration until job finding may not reflect a causal effect, but simply reflect the negative correlation between cannabis use and the decision to return to school. We investigate this issue by replacing the job finding hazard in the four equation system with a competing risk model. In the competing risk model, there are two ways in which the period of unemployment after first leaving school can end. Individuals can either find a job or they can return to school. The competing risk model is specified as a special case of a bivariate mixed proportional hazard model with flexible duration dependence.

Transitions to employment are modeled as previously described. Transitions back to school are modeled symmetrically, with the rate at which males return to school, θ^r , given by

$$\theta^r(t^r | x_j, \nu^r) = \lambda^r(t^r) \exp(x_j' \beta_r + \nu^r) \quad (17)$$

where $\lambda^r(t^r)$ represents duration dependence, which is modeled flexibly using the same step function as for the duration dependence of the job finding rate.

The conditional density function for the completed duration until finding a job or returning to school is

$$f^{jr}(t | x_j, \nu^j, \nu^r) = ((\theta^j(t | x_j, \nu^j) + \theta^r(t | x_j, \nu^r)) \exp\left(-\int_0^t ((\theta^j(y | x_j, \nu^j) + \theta^r(y | x_j, \nu^r)) dy\right)$$

³⁵In the hazard for job finding in the restricted model, the point estimate (t-stat) for cannabis use before leaving school is 0.08 (1.7) which is the same as for the unrestricted model. For the starting wage the coefficient estimate on cannabis use at school in the restricted model is -0.017 (1.49) compared to -0.017 (1.5) in the unrestricted model.

(18)

Similar to the cannabis use dynamics, we integrate out the unobserved heterogeneity component assuming that it follows a discrete distribution with four points of support with the following associated probabilities

$$Pr(\nu^j = \nu_1^j, \nu^r = \nu_1^r) = p_6, Pr(\nu^j = \nu_1^j, \nu^r = \nu_2^r) = p_7 \quad (19)$$

$$Pr(\nu^j = \nu_2^j, \nu^r = \nu_1^r) = p_8, Pr(\nu^j = \nu_2^j, \nu^r = \nu_2^r) = p_9 \quad (20)$$

which, as before, are modeled using a multinomial logit specification, $p_n = \frac{\exp(\alpha_n)}{\sum_n \exp(\alpha_n)}$, with $n = 6, 7, 8, 9$ and α_9 is normalized to zero. The duration of stay in post-school unemployment is determined by the sum of the two transition rates $\theta^j(t|x_j, \nu^j)$ and $\theta^r(t|x_j, \nu^r)$. Identification of the separate transition rates in the competing risk model comes from the nature of the transition, to a job or back to school. Appendix A2 provides details of the specification of the likelihood.

Table 5 panel *b* reports nonparametric maximum likelihood estimates of the impact of experimental and persistent cannabis use on the school to work transition accounting for the competing risk of returning to school. Panel *a* repeats the estimates from Table 4 panel *d*, in which the option of returning to school is ignored. As can be seen from the reported estimates, accounting for the fact that a significant proportion of the sample exits unemployment by returning to school produces stronger effects of non-experimental cannabis use both in terms of magnitude and significance. The point estimate of the coefficient on non-experimental cannabis use (use for more than a year) in the hazard for job finding increases from 0.12 to 0.17, implying an increase in the estimated job finding rate from 13% to 19% relative to non-cannabis users.³⁶ The higher rate of job finding by non-experimental users comes at the cost of a wage rate that is 2.0% lower based on estimates that ignore the option of returning to school, and 2.3% lower if returning to school is accounted for. In addition, for the specification that incorporates the competing risk of returning to school, the coefficient estimates suggest that both experimental and non-experimental cannabis use at school reduces the likelihood of returning to school. Overall, these findings are consistent with the previous finding that non-experimental cannabis use increases the impatience of young male school leavers or it increases their risk of being laid off from their job.

³⁶13%=100(exp(0.12)-1), 19%=100(exp(0.17)-1)

Table 5: Parameter estimates of the effect of cannabis use: competing risk for return to school or find a job

	Job finding		Back to school		Starting wage		-LogL	Mass points
a. No return to school							21072.1	12
Up to 1 year	-0.07	(0.9)	-	-	-0.004	(0.2)		
More than 1 year	0.12	(2.4)**	-	-	-0.020	(1.7)*		
b. With return to school							23427.9	13
Up to 1 year	-0.09	(1.0)	-0.38	(1.8)*	-0.001	(0.4)		
More than 1 year	0.17	(2.9)**	-0.47	(3.1)**	-0.023	(2.0)**		
c. Ignore job if return to school							23730.2	12
Up to 1 year	-0.01	(0.1)	-0.42	(2.4)**	-0.006	(0.3)		
More than 1 year	0.28	(4.7)**	-0.46	(3.8)**	-0.026	(2.0)**		
d. Full time job							24143.5	12
Up to 1 year	0.09	(1.0)	-0.56	(2.9)**	-0.015	(0.7)		
More than 1 year	0.26	(4.4)**	-0.58	(4.3)**	-0.028	(2.2)**		

Note: All estimates contain the same explanatory variables as in Table 2 (for cannabis) and Table 3 (for job finding, back to school and wages). Absolute t-statistics in parentheses; ** (*) indicates significant at a 5 (10) percent level.

Distribution of Unobserved Heterogeneity panel b						
Cannabis Starting	Cannabis Quit	Job Finding Rate				
Rate	Rate	High	Low	Low	Back to School Rate	
		Low	High	Low		
High Wages						
Positive	High	-	5.5	-	5.5	
Positive	Low	1.0	1.4	-	2.4	
Zero	-	0.6	-	0.6	1.2	
Low Wages						
Positive	High	29.4	-	3.2	32.6	
Positive	Low	28.4	4.2	1.0	33.6	
Zero	-	17.0	4.7	3.0	24.7	
Total		76.4	15.8	7.8	100.0	

While those who use cannabis are more likely to exit unemployment by getting a job compared to non-cannabis users, a more careful examination of the behavior of young men following school leaving reveals that this is partly because cannabis users are more likely to be employed before returning to school compared to non-cannabis users. For example, among those who return to school within a year, 44% of cannabis users and 38% of non-cannabis users first had a job. It is questionable whether these temporary “gap year” jobs are part of a school to work transition. Moreover, since there are differences in this behavior between those who use cannabis before leaving school and those who do not, failing to address this issue may lead to misleading and unreliable estimates of the impact of cannabis use on the school to work transition. Therefore, to put cannabis users and non-cannabis users who return to school (relatively quickly) on a more equal footing, we re-estimate the model that accounts for returning to school, but this time we ignore jobs for individuals who return to school within a year (of leaving school). The results from this estimation are reported in Table 5 panel *c*. They show that, by focusing on jobs that are not short-term “gap year” jobs, the impact of non-experimental cannabis use on the school to work transition is in fact strengthened, with non-experimental cannabis use estimated to increase the job finding rate by 32% ($\exp(0.28) - 1$), and decrease the starting hourly wage rate by 2.6%. Once again, these results confirm those of the baseline estimates, suggesting that non-experimental cannabis use leads school leavers to be more impatient in their first post-school labor market decision making, or makes them more vulnerable to being laid off from work.

A final issue we examine is whether our findings are sensitive to defining the school to work transition in terms of any employment or full-time employment only. Up until this point, any job that had a start month that occurred after the month the respondent left school has been considered a school to work transition. The only exception to this is found in panel *c* of Table 5, for which “gap year” jobs were ignored. As a final robustness check we examine the sensitivity of our findings to focusing on the transition from school to full-time employment, rather than any employment. In particular, we restrict our definition of employment to jobs in which the respondent worked 35 hours or more per week on average. The results from estimation, reported in panel *d* of Table 5, produce coefficient estimates for the indicator for using cannabis for more than a year before leaving school that are similar to those reported in Table 5 panel *c*, and are consistent with earlier findings that experimental cannabis use has no impact on the school to work transition, while non-experimental cannabis use leads to faster job finding and lower wages. As a final point, it is worth noting that also common

to competing risk specifications is the finding that cannabis use before leaving school (either experimentally or non-experimentally) reduces the likelihood of returning to school.

6 Discussion

We investigate the effect of cannabis use while in school on young men’s transition from school to work. The two dimensions of this transition that we focus on are the duration of time between leaving school and finding employment for the first time, and the wage rate received when starting the first job. Our empirical approach accounts for selectivity into cannabis use as well as common unobserved confounders that jointly affect finding a job and the initial wage rate if a job is found. After accounting for selectivity, we find that using cannabis at school leads to shorter durations of post-school unemployment. This comes at the cost of lower wages, with those who use cannabis while at school accepting a slightly lower hourly wage compared to those who did not use cannabis while at school. However, the effect of cannabis on the rate at which young men find jobs and the initial wages in the first job is present only for those who used for a period of more than a year before leaving school. Experimenting with cannabis use, for up to a year while at school, does not seem to affect young men’s early labor market experience.

In an extended sensitivity analysis, we account for school leavers’ option to exit unemployment by returning to school rather than finding a job within a competing risk framework. Doing so provides new and significant insights in terms of understanding the effects of early cannabis use on initial labor market experiences and in terms of the mechanisms through which cannabis use affects completed education. Our findings with regard to the impact of cannabis use on school to work transitions are strengthened by accounting for school leavers’ option to return to school. In addition, we show that compared to non-users, non-experimental cannabis users are more likely to exit post-school unemployment by finding a job, and both experimental and non-experimental users are less likely to exit post-school unemployment by returning to school.

If we take our preferred parameter estimates presented in panel *b* of Table 5 at face value they provide an indication of the magnitude of the effects. Non-experimental cannabis use at school (i.e. more than 1 year) increases the job finding rate by almost 19% ($100 \cdot \exp(0.17) - 1$) while the rate of returning to school is reduced by almost 38% ($100 \cdot \exp(-0.47) - 1$). If we assume the job finding rate in the first year to be on average equal to 20% per quarter for

non-cannabis users, our estimates suggest that it would be 24% for cannabis users. Similarly if we assume the rate of return to school in the first year to be on average equal to 8% per quarter for non-cannabis users, it would be 5% for cannabis users. So, although the overall exit rate from the post-schooling state (of unemployment) after 1 quarter is not very different for non-experimental cannabis users and non-cannabis users (29% versus 28%), the composition of the exit is different. In addition, non-experimental cannabis users accept job offers with starting wages a 2.3% lower than otherwise similar young men who have not used cannabis while at school.

All in all, we conclude that non-experimental cannabis users find a job more quickly but at a lower wage, and are also less likely to return to school. Our finding that cannabis use may have a negative effect on educational attainment is not new (see for example van Ours and Williams (2009)). However, we identify a new mechanism through which cannabis use leads to lower educational attainment after initial school leaving. School leavers who have used cannabis at school are less likely to return to school. In addition to that, conditional on their observed characteristics, those who used for more than a year at school are also more likely to find a job quickly and accept job offers at lower initial wage rates. Apparently, cannabis users have changed their trade-off between search costs and search benefits in the sense that their search costs are perceived to be higher and/or the perceived benefits from search are lower. The question is, why is this the case? What is the mechanism through which cannabis users are affected? We are able to eliminate cannabis use leading to an increased value of leisure or lower expected wages as credible mechanisms through which cannabis affects initial labor market outcomes, since these mechanisms each imply a longer duration of search.

A plausible avenue through which cannabis use may affect the school to work transition is by increasing the rate at which people discount the future, making them more impatient. This is consistent with cannabis users accepting a job offer after a shorter search and at a slightly lower wage rather than continuing to search for a job which pays a slightly higher wage. This mechanism is also consistent with the reduced likelihood of cannabis users returning to school (and investing in human capital) at the expense of current earnings. A second mechanism consistent with our findings is that using cannabis while at school, especially for longer durations, leads to a higher probability of being laid off. However, this mechanism would seem unlikely given that we are considering young men's first post-school job. Further, it is not immediately obvious that having a higher probability of being laid

off is consistent with our finding that cannabis use reduces the probability of returning to school. For this reason, we interpret our findings as suggesting that using cannabis for more than a year before leaving school increases young men’s impatience in their first labor market experience.

An important question arising from this research is what the findings imply for the long run economic well-being of young men. The answer to this question depends on whether, and under what circumstances, the increase in impatience in initial labor market decision making induced by non-experimental cannabis use persists. For example, a permanent increase in impatience with respect to the job market may lead to reduced on the job investment and flatter earnings profiles (Munasinghe and Sicherman, 2006). Discerning whether the effect of early non-experimental cannabis use persists, and in the event it does, its effect on the long run labor market outcomes of young men, are important avenues for future research.

Appendix A: Likelihood specifications

A1. Cannabis use dynamics

In terms of cannabis use there are three possible situations:

1. An individual starts using cannabis and quits using cannabis, in which case we have a completed duration until use and a completed duration of use (*PC1*).
2. An individual starts using cannabis and does not quit using, in which case we have a completed duration until use and an incomplete duration of use (*PC2*).
3. An individual has not started using cannabis, in which case we have a incomplete duration until use and no duration of use (*PC3*).

When specifying the likelihood we have to take the interval nature of the data into account. For example, we know at which age an individual started using cannabis but not precisely when. So, if an individual started using at age 16, we model that this individual did not start before his 16th birthday but did start before he turned 17. The same holds for the completed duration of use. Here too, we observe the number of years but not the exact timing within the year.

As indicated in the text in section 3.2 we have three types of individuals, with different combinations of unobserved heterogeneity:

1. Type 1: (u_1^s, u_1^q) , which occurs with probability p_1 .
2. Type 2: (u_1^s, u_2^q) , which occurs with probability p_2 .
3. Type 3: $u_2^s = -\infty$, which occurs with probability $p_3 = 1 - p_1 - p_2$.

Clearly Type 3 individuals can only be in Situation 3, but Type 1 and Type 2 can be in any of the three situations.

We define S^s and S^q as the survivor functions related to f^s and f^q . And, we use the following notation: $S_z^s = S^s(t^s|x, u_z^s)$ and $\Delta S_z^s = S^s(t^s - 1|x, u_z^s) - S^s(t^s|x, u_z^s)$ for $z = 1, 2$. S_z^q and ΔS_z^q are defined similarly.

Then, the loglikelihood consists of three parts:

1. $X_1 = \sum_{PC1} \Delta S_1^s \Delta S_1^q + \sum_{PC2} \Delta S_1^s S_1^q + \sum_{PC3} S_1^s$
2. $X_2 = \sum_{PC1} \Delta S_1^s \Delta S_2^q + \sum_{PC2} \Delta S_1^s S_2^q + \sum_{PC3} S_1^s$
3. $X_3 = \sum_{PC3} S_2^s$

and the loglikelihood = $\log(p_1 X_1 + p_2 X_2 + p_3 X_3)$

A2. Job finding and returning to school

In terms of job finding and returning to school there are three possibilities:

1. An individual finds a job, in which case we have a completed duration to job finding ($PJ1$).
2. An individual returns to school, in which case we have a completed duration to return-to-school ($PJ2$).
3. An individual neither finds a job nor returns to school, in which case we have two incomplete durations ($PJ3$).

We define S^{jr} as the survivor function of f^{jr} , $S_{wz}^{jr} = S^{jr}(t|x, v_w^j, v_z^r)$ and $\Delta S_{wz}^{jr} = S^{jr}(t - 1|x, v_w^j, v_z^r) - S^{jr}(t|x, v_w^j, v_z^r)$ for $w = 1, 2$ and $z = 1, 2$.

We have four types of individuals:

1. Type 1: (v_1^j, v_1^r) , which occurs with probability p_4

2. Type 2: (v_1^j, v_2^r) , which occurs with probability p_5
3. Type 3: (v_2^j, v_1^r) , which occurs with probability p_6
4. Type 4: (v_2^j, v_2^r) , which occurs with probability $p_7 = 1 - p_4 - p_5 - p_6$

Furthermore, we use the following notation: $\theta_w^j = \theta(t|x_j, v_w^j)$ for $w = 1, 2$ and $\theta_z^r = \theta^r(t|x_r, v_z^r)$ for $z = 1, 2$. Then, taking the interval nature of the data into account, the loglikelihood consists of four parts:

1. $X_4 = \sum_{PJ1} \frac{\theta_1^j}{\theta_1^j + \theta_1^r} \Delta S_{11}^{jz} + \sum_{PJ2} \frac{\theta_1^r}{\theta_1^j + \theta_1^r} \Delta S_{11}^{jz} + \sum_{PJ3} S_{11}^{jz}$
2. $X_5 = \sum_{PJ1} \frac{\theta_1^j}{\theta_1^j + \theta_2^r} \Delta S_{12}^{jz} + \sum_{PJ2} \frac{\theta_2^r}{\theta_1^j + \theta_2^r} \Delta S_{12}^{jz} + \sum_{PJ3} S_{12}^{jz}$
3. $X_6 = \sum_{PJ1} \frac{\theta_2^j}{\theta_2^j + \theta_1^r} \Delta S_{21}^{jz} + \sum_{PJ2} \frac{\theta_1^r}{\theta_2^j + \theta_1^r} \Delta S_{21}^{jz} + \sum_{PJ3} S_{21}^{jz}$
4. $X_7 = \sum_{PJ1} \frac{\theta_2^j}{\theta_2^j + \theta_2^r} \Delta S_{22}^{jz} + \sum_{PJ2} \frac{\theta_2^r}{\theta_2^j + \theta_2^r} \Delta S_{22}^{jz} + \sum_{PJ3} S_{22}^{jz}$

and the loglikelihood = $\log(p_4 X_4 + p_5 X_5 + p_6 X_6 + p_7 X_7)$

Table A.1: Summary Statistics for Control Variables

	Full	Used cannabis at school		
	Sample	No	Yes	Difference
<i>Individual and family background controls</i>				
Black	0.26	0.28	0.22	-0.06**
Hispanic	0.21	0.23	0.19	-0.04***
Non_black	0.53	0.48	0.58	0.10***
Ability	-0.04	-0.09	0.04	0.13**
Religious (1997)	0.87	0.88	0.85	-0.04***
Household income (log; 1997)	3.54	3.49	3.60	0.11***
Household size (log; 1997)	1.46	1.49	1.43	-0.05***
Father education < high school	0.24	0.25	0.23	-0.02
Father education = high school	0.42	0.44	0.39	-0.04*
Father education > high school	0.34	0.32	0.38	0.04**8
Mother education < high school	0.24	0.26	0.20	-0.06***
Mother education = high school	0.36	0.37	0.35	0.02
Mother education > high school	0.40	0.37	0.45	0.08***
Mothers age at respondent's birth	25.62	25.49	25.79	0.30*
Mothers parenting style authoritarian	0.55	0.56	0.54	-0.02
Mother present	0.93	0.94	0.92	-0.02**
Father present	0.69	0.69	0.69	-0.01
Year born is 1980	0.18	0.17	0.19	0.02
Year born is 1981	0.21	0.21	0.20	-0.01
Year born is 1982	0.21	0.20	0.22	0.02
Year born is 1983	0.20	0.21	0.20	-0.01
School urban	0.84	0.84	0.85	0.01
School in NE	0.17	0.16	0.18	0.02
School in NC	0.23	0.21	0.24	0.03
School in W	0.23	0.20	0.25	0.05***
<i>Additional controls: job search and wage eqn</i>				
Education less than high school	0.32	0.31	0.34	0.03*
Education is high school graduate	0.40	0.43	0.35	-0.08***
Education is greater than high school graduate	0.28	0.26	0.31	0.05**
<i>Additional controls: wage eqn</i>				
Year started job is 1997	0.01	0.01	0.01	0.00
Year started job is 1998	0.06	0.05	0.06	0.01
Year started job is 1999	0.13	0.13	0.12	-0.02
Year started job is 2000	0.14	0.14	0.15	0.02
Year started job is 2001	0.14	0.15	0.13	-0.02
Year started job is 2002	0.14	0.14	0.14	-0.01
Year started job is 2003	0.12	0.12	0.11	-0.01
Year started job is 2004	0.07	0.07	0.07	0.00
Year started job is 2005	0.06	0.06	0.06	0.00
Year started job is 2006	0.05	0.05	0.05	0.00
Year started job is 2007	0.04	0.04	0.04	0.00
Year started job is 2008	0.02	0.02	0.03	0.01*
Year started job is 2009	0.01	0.01	0.01	0.01
Year started job is 2010	0.01	0.01	0.01	0.00
Year started job is 2011	0.01	0.01	0.01	0.00
Industry 1	0.01	0.01	0.01	0.00
Industry 2	0.00	0.00	0.00	0.00
Industry 3	0.10	0.10	0.11	0.02*
Industry 4	0.09	0.10	0.08	-0.01
Industry 5	0.23	0.24	0.21	-0.03**
Industry 6	0.04	0.04	0.03	-0.01*
Industry 7	0.02	0.02	0.02	0.00
Industry 8	0.04	0.04	0.04	0.00
Industry 9	0.12	0.12	0.12	0.00
Industry 10	0.07	0.07	0.06	-0.01
Industry 11	0.20	0.18	0.22	0.04***
Industry 12	0.05	0.04	0.05	0.00
Industry 13	0.02	0.02	0.02	0.01
Industry 14	0.01	0.01	0.01	0.00

References

- Abbring, J. H. and G. J. Van den Berg (2003). The non-parametric identification of treatment effects in duration models. *Econometrica* 71, 1491–1518.
- Alford, C. (2015). *High Today Versus Lows Tomorrow: Substance Use, Education, and Employment Choices of Young Men*. Doctoral Dissertation. University of Virginia.
- Altonji, J., L. Kahn, and J. Speer (2014). Cashier or consultant? Entry labor market conditions, field of study, and career success. *NBER working paper 20531*.
- Arni, P., R. Lalive, and J. C. van Ours (2013). How effective are unemployment benefit sanctions? looking beyond unemployment exit. *Journal of Applied Econometrics* 28, 1153–1178.
- Becker, G. and K. Murphy (1988). A theory of rational addiction. *Journal of Political Economy* 96(4), 675–700.
- Becker, G. and G. Stigler (1977). De gustibus non est disputandum. *American Economic Review* 67, 76–90.
- Bray, J. W. (2005). Alcohol use, human capital, and wages. *Journal of Labor Economics* 23, 279–312.
- Burgess, S. M. and C. Propper (1998). Early health related behaviours and their impact on later life chances: Evidence from the US. *Health Economics* 7(5), 381–399.
- Cahuc, P., S. Carcillo, and A. Zylberberg (2014). *Labor Economics, second edition*. MIT Press. Cambridge, Mass.
- Carneiro, P., J. J. Heckman, and E. J. Vytlačil (2011). Estimating marginal returns to education. *American Economic Review* 101(6), 2754 – 2781.
- Carneiro, P., C. Meghir, and M. Pary (2013). Maternal education, home environments, and the development of children and adolescents. *Journal of the European Economic Association* 11, 123–160.
- Chatterji, P., M. Alegria, and D. Takeuchi (2011). Psychiatric disorders and labor market outcomes: Evidence from the national comorbidity survey-replication. *Journal of Health Economics* 30, 858–868.
- Cutler, D. (1995). The incidence of adverse medical outcomes under prospective payment. *Econometrica* 63, 29–50.
- DellaVigna, S. and M. D. Paserman (2005). Job search and impatience. *Journal of Labor Economics* 23(3), 527–587.
- DeSimone, J. (2002). Illegal drug use and employment. *Journal of Labor Economics* 20(4), 952–977.
- Frank, R. and P. Gertler (1991). An assessment of measurement error bias for estimating the effect of mental distress on income. *The Journal of Human Resources* 26, 154–164.

- Gaure, S., K. Roed, and T. Zhang (2007). Time and causality: A Monte Carlo assessment of the timing-of-events approach. *Journal of Econometrics* 141(2), 1159 – 1195.
- Hansen, K. T., J. J. Heckman, and K. J. Mullen (2004). The effect of schooling and ability on achievement test scores. *Journal of Econometrics* 121, 39 – 98.
- Harrison, L. D., S. S. Martin, T. Enev, and D. Harrington (2007). Comparing drug testing and self-report of drug use among youths and young adults in the general population. Technical report, (DHHS Publication No. SMA 07-4249, Methodology Series M-7). Rockville, MD: Substance Abuse and Mental Health Services Administration, Office of Applied Studies.
- Heckman, J. and B. Singer (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica* 52, 271–320.
- Johnston, L. ., P. OMalley, J. Bachman, J. Schulenberg, and R. Miech (2015). Monitoring the future national survey results on drug use, 1975-2015; volume 2 college students & adults ages 1955. Technical report, Institute for Social Research, The University of Michigan.
- Kaestner, R. (1991). The effect of illicit drug use on the wages of young adults. *Journal of Labor Economics* 9(4), 381–412.
- Kaestner, R. (1994a). The effect of illicit drug use on the labor supply of young adults. *Journal of Human Resources* 29(1), 126–155.
- Kaestner, R. (1994b). New estimates of the effect of marijuana and cocaine use on wages. *Industrial and Labor Relations Review* 47(3), 454–470.
- Kaestner, R. (1998). Illicit drug use and labor market outcomes: A review of economic theory and its empirical implications. *Journal of Drug Issues* 28(3), 663–680.
- Kagel, J. J., R. C. Battalio, and C. G. Miles (1980). Marihuana and work performance: Results from an experiment. *Journal of Human Resources* 15(3), 373–395.
- Macleod, J., R. Oakes, A. Copello, I. Crome, M. Egger, M. Hickman, T. Oppenkowski, H. Stokes-Lampard, and G. Smith (2004). Psychological and social sequelae of cannabis and other illicit drug use by young people: A systematic review of longitudinal, general population studies. *The Lancet* 363, 1579–1588.
- Michael, R. and G. Becker (1973). On the new theory of consumer behavior. *Swedish Journal of Economics* 75, 378–396.
- Mroz, T. A. (1999). Discrete factor approximations in simultaneous equation models: Estimating the impact of a dummy endogenous variable on a continuous outcome. *Journal of Econometrics* 92, 233–274.
- Mullahy, J. and J. Sindelar (1993). Alcoholism, work, and income. *Journal of Labor Economics* 11, 494–520.
- Munasinghe, L. and N. Sicherman (2006). Why do dancers smoke? Smoking, time preference, and wage dynamics. *Eastern Economic Journal* 32, 595–616.

- Oreopoulos, P., T. von Wachter, and A. Heisz (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics* 4, 1–29.
- Pacula, R. L., D. Powell, P. Heaton, and E. L. Sevigny (2015). Assessing the effects of medical marijuana laws on marijuana use: The devil is in the details. *Journal of Policy Analysis and Management* 34(1), 7–31.
- Pacula, R. L. and E. L. Sevigny (2014). Marijuana liberalizations policies: why we can't learn much from policy still in motion. *Journal of Policy Analysis and Management* 33, 212–221.
- Register, C. and D. Williams (1992). Labor market effects of marijuana and cocaine use among young men. *Industrial and Labor Relations Review* 45(3), 435–448.
- Substance Abuse and Mental Health Services Administration (2014). Results from the 2013 national survey on drug use and health: Summary of national findings. Technical report, NSDUH Series H-48, HHS Publication No. (SMA) 14-4863. Rockville, MD: Substance Abuse and Mental Health Services Administration, 2014.
- Tekin, E. and S. Markowiz (2008). The effects of suicidal behavior on productive activities of young adults. *Southern Economic Journal* 75, 300–331.
- van Ours, J. C. (2006). Cannabis, cocaine and jobs. *Journal of Applied Econometrics* 21, 897–917.
- van Ours, J. C. (2007). The effect of cannabis use on wages of prime age males. *Oxford Bulletin of Economics and Statistics* 69, 619–634.
- van Ours, J. C. and J. Williams (2009). Why parents worry: Initiation into cannabis use by youth and their educational attainment. *Journal of Health Economics* 28, 132–142.
- van Ours, J. C. and J. Williams (2015). Cannabis use and its effects on health, education and labor market success. *Journal of Economic Surveys* 29(5), 993–1010.
- van Ours, J. C., J. Williams, D. Fergusson, and L. J. Horwood (2014). Cannabis use and suicidal ideation. *Journal of Health Economics* 32(3), 524–537.
- Yang, Z., D. Gilleskie, and E. Norton (2009). Health insurance, medical care and health outcomes: A model of elderly dynamics. *Journal of Human Resources* 44, 47–114.
- Zarkin, G. A., T. A. Mroz, J. W. Bray, and M. T. French (1998). The relationship between drug use and labour supply for young men. *Labour Economics* 5(4), 385–409.