Journal of Health Economics 31 (2012) 112–121

Contents lists available at SciVerse ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/econbase

The effect of alcohol availability on marijuana use: Evidence from the minimum legal drinking age

Benjamin Crost[a,*], Santiago Guerrero[b]

[a] Department of Economics, University of Colorado Denver, Denver, CO 80127, United States
[b] Bank of Mexico, Direccion General de Investigación Económica, Mexico

ARTICLE INFO

Article history:
Received 16 June 2011
Received in revised form
12 December 2011
Accepted 16 December 2011
Available online 13 January 2012

JEL classification:
I18

Keywords:
Alcohol
Marijuana
Drug use
Minimum legal drinking age
Regression discontinuity

ABSTRACT

This paper exploits the discontinuity created by the minimum legal drinking age of 21 years to estimate the causal effect of increased alcohol availability on marijuana use. We find that consumption of marijuana decreases sharply at age 21, while consumption of alcohol increases, suggesting that marijuana and alcohol are substitutes. We further find that the substitution effect between alcohol and marijuana is stronger for women than for men. Our results suggest that policies designed to limit alcohol use have the unintended consequence of increasing marijuana use.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Economic theory suggests that when the cost of consuming a good increases, people will consume more of its substitutes and less of its complements. In the case of alcohol, the substitutes are likely to include other intoxicating substances. The minimum legal drinking age (MLDA), which restricts access to alcohol for those under 21, is therefore likely to affect the consumption of other drugs among that age group, as it sharply decreases the cost of consuming alcohol for individuals just over the MLDA. When assessing the costs and benefits of policies that aim to reduce alcohol consumption – like the MLDA or alcohol taxes – we need to take possible substitution behavior into account.

For example, proponents of the MLDA at age 21 argue that alcohol consumption in children and adolescents can cause long term and, sometimes, irreversible damages to the brain (AMA, 2008). In particular, adolescents who drink are more likely to develop smaller hippocampi, a part of the brain that controls learning and memory, and are more likely to show alterations in their prefrontal cortex (AMA, 2008). Alcohol consumption has also been shown to induce suicides and car accidents (Carpenter and Dobkin, 2009). However, if restricting access to alcohol causes people to switch to substitutes, such as marijuana or other illegal drugs, the benefits of reduced alcohol consumption need to be weighed against the cost of increased consumption of alcohol's substitutes. The potential alcohol substitute we analyze in this paper is marijuana, a substance made of a mixture of flowers, seeds and leaves of the hemp plant. The hemp plant contains tetrahydrocannabinol or THC, a psychoactive chemical that produces most of the intoxicating effects. Consumption of THC has been associated with cognitive deficits and changes in brain morphology and psychiatric disorders (Wilson et al., 2000; Pope et al., 2003; Hall and Degenhardt, 2009). In this paper we study the effects of an increase in the availability of alcohol on the consumption of marijuana.

Most previous studies of substitution between alcohol and marijuana (e.g. DiNardo and Lemieux, 2001; Chaloupka and Laixuthai, 1997; Pacula, 1998; Williams et al., 2004; Saffer and Chaloupka, 1999; Farrelly et al., 1999) are based on cross-sectional (usually between-state) variation in the prices of alcohol and marijuana, the MLDA, alcohol taxes, or laws that partially decriminalize marijuana. A problem for these approaches is that state-level prices of alcohol and marijuana and the policies governing their consumption are likely to be correlated with unobserved characteristics of the population living in those states, making it difficult to infer

* Corresponding author. Tel.: +1 303 315 2036.
E-mail address: ben.crost@gmail.com (B. Crost).

0167-6296/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jhealeco.2011.12.005

causality from cross-sectional comparisons (Carpenter and Dobkin, 2009).

We address the problem of causal identification that has plagued previous research through a regression discontinuity design. This approach exploits the sharply discontinuous nature of the minimum legal drinking age, the fact that a person cannot legally purchase alcohol up until the day before her 21st birthday, but can do so from her 21st birthday onwards. By comparing substance use in individuals just below and just above the age of 21,[1] we can therefore isolate the causal effect of the MLDA on alcohol and marijuana consumption. The identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just above and just below the age of 21 are similar in all characteristics that determine substance use. The regression discontinuity approach allows us to estimate the extent of substitution between alcohol and marijuana and identify the causal effect of changes in the MLDA on individuals close to 21 years of age.

Our results show that alcohol and marijuana are substitutes. At age 21, we observe a sharp increase in alcohol consumption but a decrease in marijuana consumption. This suggests that policies that restrict access to alcohol cause an increase in marijuana consumption. Our estimates suggest that the MLDA at age 21 decreases the probability of having consumed alcohol in the past 30 days by 16% and increases the probability of having consumed marijuana by 10%. Results from instrumental variables suggest an elasticity of substitution of approximately 0.7 for the probability of use and 0.4 for the frequency of use (defined as the number of days in which a substance was consumed). We further find that the substitution effect is substantially stronger for women than for men. Our results suggest that by restricting the age at which people can legally purchase alcohol, the MLDA causes an increase in the consumption of illicit drugs, especially by young women.

The next section reviews the existing literature on the MLDA and marijuana use. Section 3 describes the empirical strategy in more detail. Section 4 presents the data and results, Section 5 shows the robustness of our estimates and Section 6 concludes.

2. Literature review

Most of the previous literature on substitution between marijuana and alcohol exploits between-state variation in the minimum legal drinking age (MLDA) and marijuana decriminalization during the 1970s and 1980s. DiNardo and Lemieux (2001) estimate a structural model of alcohol and marijuana consumption to test the effect of increases in the MLDA. They analyze state-level percentages of high school seniors that reported having consumed alcohol/marijuana from the Monitoring the Future Surveys (MFS) during the period 1980–1989. Their find that alcohol and marijuana are substitutes and that increases in the MLDA lead to a decrease in alcohol consumption and an increase in marijuana consumption. In a similar study, Chaloupka and Laixuthai (1997) find that youths living in states where marijuana was decriminalized report having consumed less alcohol, providing some evidence of substitution between marijuana and alcohol consumption.

Other studies either find no evidence of substitution or evidence of complementarity. Using data from the National Longitudinal Survey of Youth (NLSY), Thies and Register (1993) do not find statistically significant evidence that state-level marijuana decriminalization affects consumption of alcohol or marijuana. Also using data from the NLSY, Pacula (1998) finds that state regulations that reduce the consumption of alcohol, such as state beer taxes and increases in the MLDA, are negatively correlated with marijuana consumption. Focusing on college students, Williams et al. (2004) analyze alcohol and marijuana consumption reported in the Harvard School of Public Health College Alcohol Study. They find that campus regulations banning the consumption of alcohol, and to a lesser extent state policies that restrict alcohol consumption, are negatively correlated with marijuana use. Using data from the National Household Surveys on Drug Abuse (NHSDA), Saffer and Chaloupka (1999) find that, controlling for the price of marijuana, county-level alcohol prices are negatively correlated with marijuana consumption. Also using data from the NHSDA, Farrelly et al. (1999) find that increases in state-level beer prices are negatively correlated with marijuana consumption for youths aged 12–20, but not for young adults aged 21–30.

In summary, the literature on substitution between alcohol and marijuana finds contradicting results. DiNardo and Lemieux (2001) and Chaloupka and Laixuthai (1997) interpret their findings as reflecting substitution between alcohol and marijuana, while Pacula (1998), Williams et al. (2004), Saffer and Chaloupka (1999) and Farrelly et al. (1999) interpret their findings as reflecting complementarity. One possible reason for these mixed results is that different studies use different surveys and time periods, which prevents comparability. Another reason, perhaps more important, is that many of the previous studies are based on state-level (or in the case of Williams, 2004, campus-level) variations in prices of alcohol and marijuana and policies governing their consumption. While this approach can establish correlations between substance use, prices and policies, the correlations do not necessarily reflect causal effects, since state-level prices and policies governing alcohol and marijuana are likely to be correlated with unobserved population characteristics that determine alcohol and marijuana consumption (Carpenter and Dobkin, 2009). In this paper, we overcome this problem by exploiting the discontinuous nature of the MLDA, which creates an abrupt change in individuals' ability to legally purchase alcohol at age 21. The empirical approach, known as a regression discontinuity design, is described in detail in the next section.

3. Empirical strategy

This paper uses a regression discontinuity design (RDD) to identify the effect of the legal minimum drinking age on alcohol and marijuana use. The RDD approach exploits the sharply discontinuous nature of the minimum legal drinking age, the fact that a person cannot legally purchase alcohol up until the day before her 21st birthday, but can do so from her 21st birthday onwards. Individuals therefore switch from the control regime, in which they are legally prohibited from buying alcohol, to the treatment regime, in which they are allowed to do so, from one day to the next. We can therefore estimate the causal effect of the minimum legal drinking age by comparing individuals who have just turned 21 and individuals who are about to turn 21. Our identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just below and just above the age of 21 are similar in all characteristics that determine substance use, so that differences between the two groups can only be explained by the effect of the minimum drinking age.

Our estimates are based on the standard regression discontinuity estimator described by Imbens and Lemieux (2008):

$$\tau_{RD} = \lim_{x \uparrow 21}[Y_i | X_i = x] - \lim_{x \downarrow 21}[Y_i | X_i = x]$$

where Y_i and X_i denote individual i's substance use and age, respectively. That is, we estimate the limit of substance use on both sides of the age of 21. The difference between the limits is the regression discontinuity estimate of the effect of the minimum legal drinking

[1] For reasons of confidentiality, the data used for our empirical analysis is aggregated to the month-of-age level.

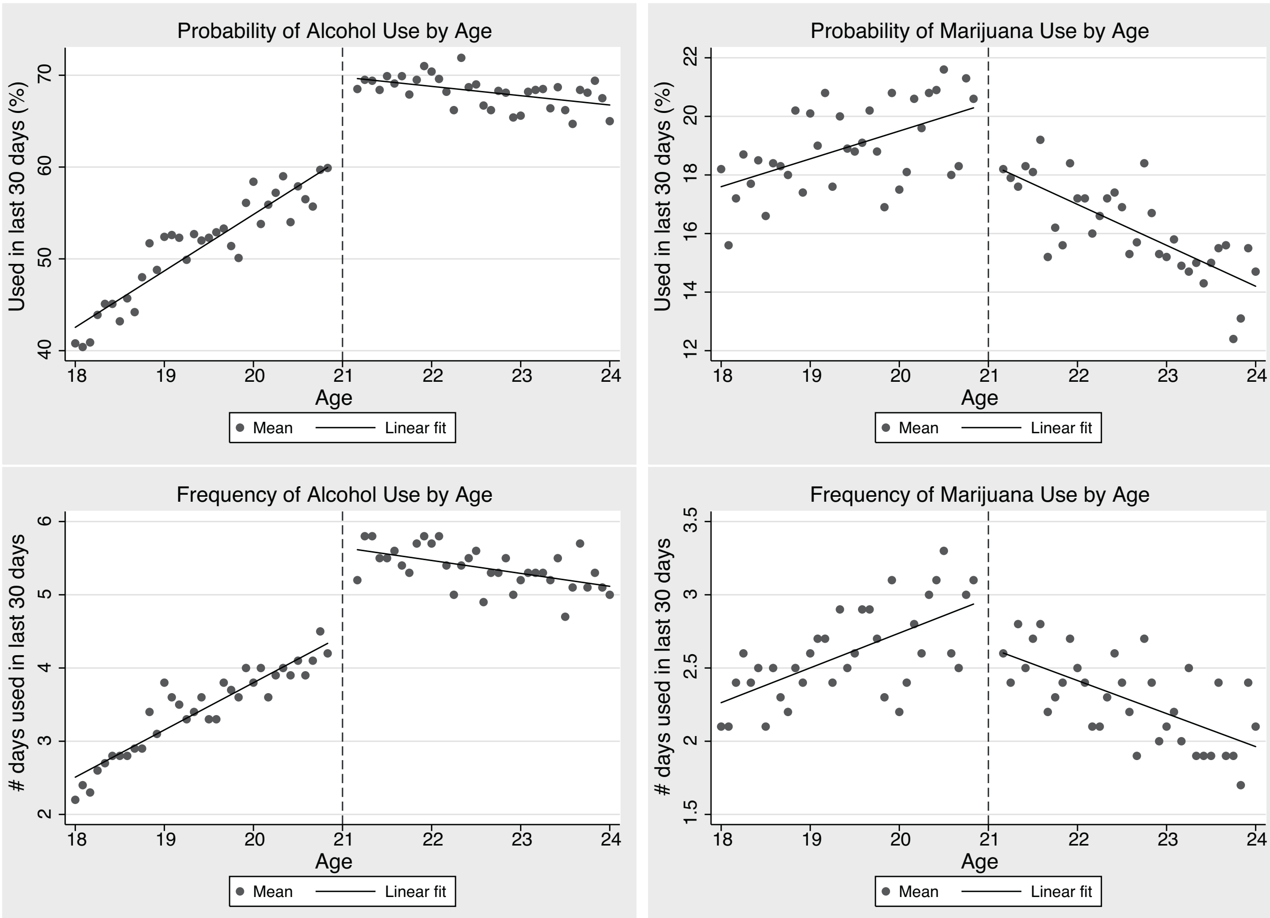

Fig. 1. Alcohol and marijuana use around age 21. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21. Data source: NSDUH 2002–2007.

age. We follow Carpenter and Dobkin (2009) and estimate the limits by local linear regression on both sides of the age of 21. In practice, this is equivalent to estimating a kernel-weighted regression of the following model (Imbens and Lemieux, 2008):

$$Y_i = \beta_0 + X_i\beta_1 + X_i * D_i\beta_2 + D_i\tau_{RD} + \varepsilon_i \quad (1)$$

As before, Y_i and X_i denote individual i's substance use and age, respectively. D_i is an indicator that takes the value 1 if individual i is 21 years old or older. The estimated coefficient τ_{RD} yields the (local) reduced form effect of the minimum legal drinking age on alcohol and marijuana use. This can be interpreted as the causal effect of a marginal increase (or decrease) of the MLDA on the alcohol/marijuana use of individuals who are exactly 21 years old.

3.1. *Instrumental variables estimates of the substitution between alcohol and marijuana*

In order to estimate the extent of substitution between alcohol and marijuana, we employ an instrumental variables approach that treats alcohol as the endogenous variable, which is instrumented

Table 1
Effect of the MLDA on alcohol and marijuana use: regression discontinuity estimates.

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	9.83	−2.01	1.30	−0.31
	(0.79)***	(0.54)***	(0.10)***	(0.11)***
Age	5.23	0.83	0.55	0.23
	(0.44)***	(0.30)***	(0.06)***	(0.06)***
Age × Over 21	−6.25	−2.15	−0.73	−0.47
	(0.62)**	(0.42)***	(0.08)***	(0.09)***
Constant	60.0	20.3	4.34	2.97
	(0.56)***	(0.38)***	(0.07)***	(0.08)***
Number of observations	68	68	68	68

Data source: NSDUH 2002–2007.
Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels.

Table 2
Effect of the MLDA on alcohol and marijuana use: men.

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	10.76	−1.47	1.48	−0.31
	(1.11)***	(0.92)	(0.15)***	(0.19)*
Age	6.11	1.39	0.78	0.35
	(0.62)***	(0.51)***	(0.08)***	(0.10)***
Age × Over 21	−7.04	−3.45	−0.87	−0.73
	(0.87)***	(0.72)***	(0.11)***	(0.15)***
Constant	64.5	25.2	5.43	4.01
	(0.79)***	(0.65)***	(0.10)***	(0.13)***
Number of observations	68	68	68	68

Data source: NSDUH 2002–2007, subsample of male respondents.
Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels.

Table 3
Effect of the MLDA on alcohol and marijuana use: women.

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	8.85	−2.62	1.08	−0.29
	(1.06)***	(0.68)***	(0.12)***	(0.10)*
Age	4.50	0.45	0.37	0.15
	(0.59)***	(0.38)	(0.06)***	(0.06)**
Age × Over 21	−5.65	−1.02	−0.61	−0.27
	(0.83)***	(0.53)*	(0.09)***	(0.08)***
Constant	55.9	15.5	3.30	1.92
	(0.75)***	(0.48)***	(0.08)***	(0.07)***
Number of observations	68	68	68	68

Data source: NSDUH 2002–2007, subsample of female respondents.
Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels.

by an indicator for being older than 21 years, and marijuana as the outcome of interest. As before, we control for linear time trends above and below the age of 21. This approach estimates the ratio of the changes in marijuana and alcohol use across the threshold at age 21: $\beta^{IV} = \Delta m/\Delta a$.

This estimate can be interpreted as the rate of substitution between alcohol and marijuana that is induced by the MLDA (or the elasticity of substitution, if a logarithmic functional form is chosen). However, when extrapolating from this estimate to the substitution induced by other policies that limit the use of alcohol, it should be kept in mind that the IV approach only estimates the local average treatment effect for the sub-population of *compliers*, those individuals who are exactly 21 years old and who (at least partly) comply with the MLDA, so that their consumption of alcohol increases discontinuously at the age of 21. Thus the external validity of the estimated substitution effect should be highest for policies that affect the alcohol use of individuals who are close to 21 years of age and likely to comply with regulations like the MLDA.

4. Data and results

Data on alcohol and marijuana use was obtained from the National Survey of Drug Use and Health (NSDUH), which is administered annually by the U.S. Department of Health and Human Services' Substance Abuse and Mental Health Services Administration (SAMHSA) and conducted by the Research Triangle Institute. The NSDUH provides estimates of alcohol and illicit substance use among persons aged 12 and older at the national and state-level using a randomly selected sample of approximately 70,000 people.

The period of observation for our analysis is 2002–2007. The NSDUH uses two measures of substance use, whether the respondent has used the substance within the past 30 days and the number of days on which the respondent has used it. For alcohol consumption, the question the survey asks is "Think specifically about the past 30 days, from [30 days before the interview date], up to and including today. During the past 30 days, on how many days did you drink one or more drinks of an alcoholic beverage?" For marijuana use, the question it asks is "Think specifically about the past 30 days, from [30 days before the interview date] up to and including today. What is your best estimate of the number of days you used marijuana or hashish during the past 30 days?". Since respondents' precise age is not available in the NSDUH's public-use files, we obtained data on the averages of the substance use measures by month of age from SAMHSA. We obtained these averages for the whole sample and separately for men and women. To maintain confidentiality of the data, SAMHSA only provided us with the average response by month of age but could not provide us with the number of individual responses that were used to calculate the average. For our baseline regressions we use a bandwidth of 3 years around age 21, so that we use data on individuals between the ages of 18 and 24. For this age-group there are 71 month-of-age cells, leading to 71 observations. To avoid measuring the effect of the (anticipated) birthday celebration itself, we drop the observations for the month of the 21st birthday as well as the preceding and following month.[2] We use a triangular kernel to estimate local linear regressions on each side of age 21.

[2] The NSDUH uses a recall period of 30 days, so that the observation in the month following the 21st birthday could still be affected by the birthday celebration. We drop the observation in the preceding month because the anticipation of the birthday celebration may lead people to consume fewer drugs than they normally would.

Table 4
Rate of substitution between alcohol and marijuana: instrumental variables estimates.

	Used MJ in last 30 days (%)			# of days used MJ in last 30 days		
	All	Men	Women	All	Men	Women
Alcohol use	−0.20	−0.14	−0.30	−0.23	−0.21	−0.27
	(0.06)***	(0.09)	(0.09)***	(0.09)**	(0.13)	(0.10)***
Age	1.90	2.22	1.78	0.36	0.52	0.25
	(0.59)***	(0.95)**	(0.78)**	(0.10)***	(0.18)***	(0.09)***
Age × Over 21	−3.43	−4.41	−2.69	−0.64	−0.91	−0.43
	(0.61)**	(0.96)***	(0.83)***	(0.11)***	(0.19)***	(0.11)***
Constant	32.6	34.0	32.0	3.98	5.16	2.81
	(4.0)***	(6.15)***	(5.58)***	(0.46)***	(0.80)***	(0.39)***
Number of observations	68	68	68	68	68	68

Data source: NSDUH 2002–2007.
Each observation is the average of substance use over a month-of-age cell. Estimates are the results of Instrumental Variables regressions in which alcohol use is instrumented by an indicator for age greater than 21 years. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels.

Table 5
Elasticity of substitution between alcohol and marijuana: instrumental variables estimates.

	Log. MJ use in last 30 days (%)			Log. # of days used MJ in last 30 days		
	All	Men	Women	All	Men	Women
Log. of alcohol use	−0.69	−0.39	−1.23	−0.41	−0.34	−0.52
	(0.23)***	(0.29)	(0.42)***	(0.19)**	(0.25)	(0.25)**
Age	0.11	0.10	0.14	0.14	0.16	0.15
	(0.04)***	(0.05)**	(0.06)**	(0.05)***	(0.07)**	(0.06)**
Age × Over 21	−0.20	−0.21	−0.22	−0.27	−0.28	−0.28
	(0.04)**	(0.05)***	(0.07)***	(0.05)***	(0.07)***	(0.07)***
Constant	5.83	4.84	7.67	1.69	1.96	1.27
	(0.96)***	(1.22)***	(1.72)***	(0.30)***	(0.46)***	(0.33)***
Number of observations	68	68	68	68	68	68

Data source: NSDUH 2002–2007.
Each observation is the average of substance use over a month-of-age cell. Estimates are the results of Instrumental Variables regressions in which the logarithm alcohol use is instrumented by an indicator for age greater than 21 years. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. *, ** and *** denote statistical significance at the 10%, 5% and 1% levels.

Fig. 1 displays average alcohol and marijuana use between the ages of 18 and 24. The individual observations are averages by month of age; the fitted lines are estimated by linear regressions of substance use on age on both sides of age 21. The top panels show that alcohol consumption increases drastically at age 21. The probability of having consumed alcohol in the last 30 days increases by about 10 percentage points from a baseline just under 60%. A similar result has previously been found by Carpenter and Dobkin (2009) and is consistent with the hypothesis that the cost of consuming alcohol decreases significantly at age 21. The frequency of alcohol consumption increases as well, from 4 to 5.5 days drinking out of the previous 30 days.

For marijuana, the effect goes in the opposite direction, though its size is smaller. At age 21, the probability of marijuana use decreases by about 2 percentage points from a baseline of about 20%. The frequency of marijuana use decreases by about 0.3 days out of a 30-day period, from a baseline of about 2.3 days.

Table 1 presents quantitative estimates from the local linear regression approach described in the previous section. The results reported in the table are for local linear regressions with a triangular kernel and a bandwidth of 3 years. Each observation is the mean of alcohol/marijuana consumption in the month of age. Robustness tests for different bandwidths are reported in Section 5. The regression results reinforce the visual impression gained from the graphs. There is a strong increase in consumption of alcohol (both probability and frequency of use) at age 21, while consumption of marijuana decreases. The changes are statistically significant and their sizes are similar to the changes visible on the graphs.

The results indicate that alcohol and marijuana are substitutes, at both the extensive and the intensive margin. The decrease in the probability of marijuana use at age 21 suggests that some individuals who use marijuana before the age of 21 stop using it (or at least use it less regularly) once they turn 21 and are able to legally consume alcohol. In absolute terms, the substitution effect on the probability of marijuana use is not very large – a 9.8 percentage point increase in the probability of alcohol consumption leads to a 2 percentage point decrease in marijuana use. However, the 9.8 percentage point increase in alcohol consumption constitutes a 16% increase from the estimated baseline consumption of 60% just below age 21, and the 2 percentage point decrease in marijuana use constitutes a 10% decrease from baseline use.

The estimated decrease in the frequency of marijuana use is 0.3 days per month which constitutes a 10% decline from the baseline of 3 days at age 21. Since the decline in the frequency of marijuana use is of similar size (in percentage terms) as the decline in the probability of use, it is unlikely that the estimated drop in the frequency of use is solely driven by the extensive margin, since people who stop using marijuana altogether probably used it with lower frequency than the average user to begin with. If the entire decline came from these 'lighter' users, we would therefore expect that the decline in the average frequency would be lower (in terms of percentage) than the decline in the probability of use. Hence, our estimates suggest that the decline in marijuana use at age 21 occurs on both the extensive and intensive margins.

4.1. Results by gender

In order to address the possibility of differentiated substitution effects by gender, we replicate the analysis for men and women separately. Figs. 4 and 5 show the trends in alcohol and marijuana

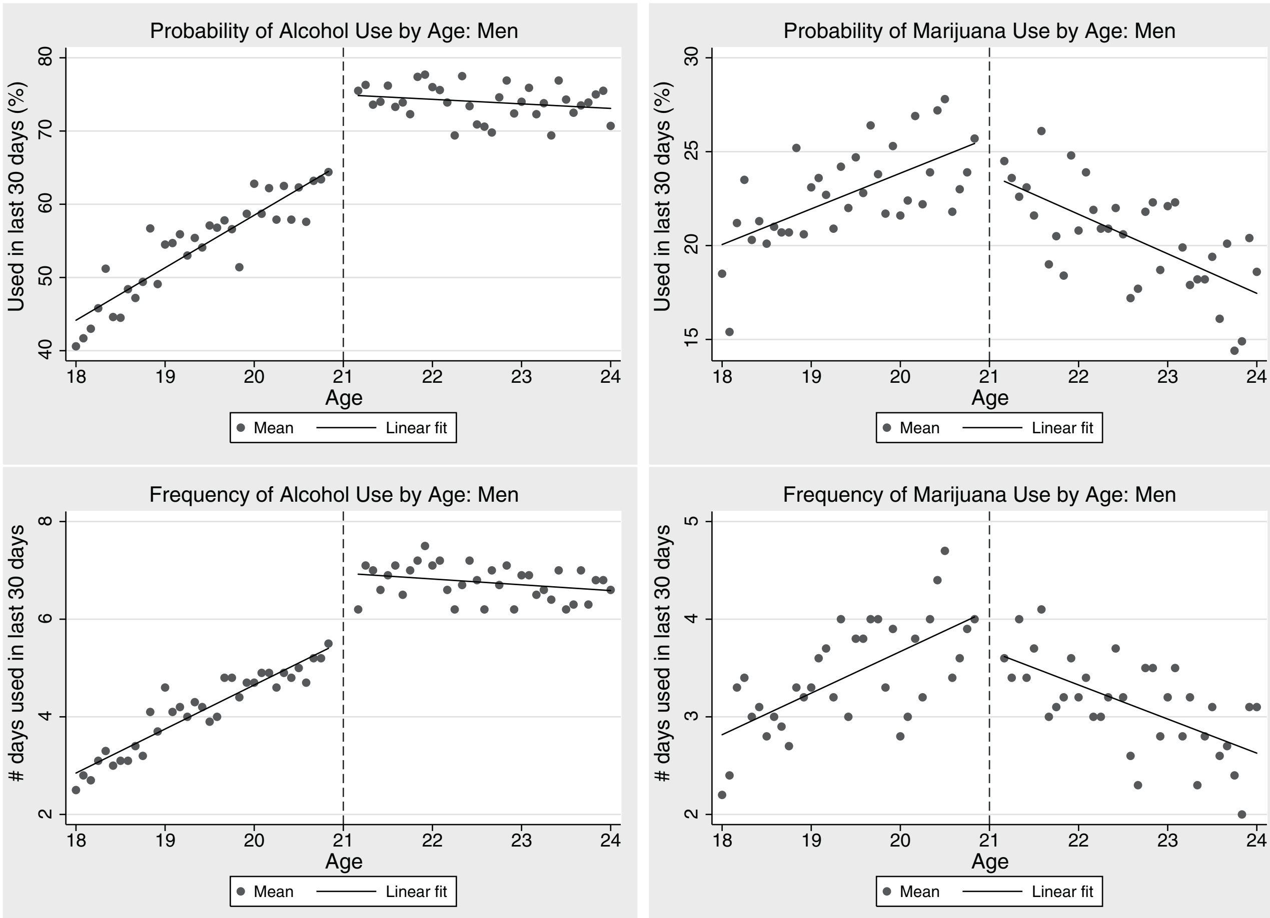

Fig. 2. Placebo tests: RD estimates by location of RD threshold. The vertical axis plots RD estimates of the change in marijuana use across the age-threshold specified on the horizontal axis. For details on the RD estimation, see the description in Table 1.
Data source: NSDUH 2002–2007.

use around age 21 for men and women, respectively. Tables 2 and 3 show the corresponding RD estimates. The results show that men have higher baseline levels of use of both alcohol and marijuana but that the effect of the MLDA on marijuana use is larger for women. For men the probability of marijuana use decreases by 1.5 percentage points at age 21, which corresponds to a 6% decrease compared to the baseline probability of 25%. For women, the probability of marijuana use decreases by 2.6 percentage points at age 21, which corresponds to a 17% decrease compared to the baseline probability of 15%. The effect of the MLDA on the *frequency* of marijuana use is also stronger for women. Women just above the age of 21 use marijuana on 0.29 days/month less than women just below the age of 21. Compared to baseline use, this effects represent a 15% decrease in the frequency of marijuana use. For men, the MLDA induces a decrease of 0.3 days/month, which represents a 7.5% decrease in the frequency of marijuana use.

4.2. IV estimates of substitution between alcohol and marijuana

This section discusses the estimates of the substitution between alcohol and marijuana that are generated by the instrumental variables approach described in Section 3.1. Table 4 presents rates of substitution that are estimated by a linear specification. For the entire population, we estimate that a one percentage point increase in the probability of using alcohol leads to a statistically significant 0.2 percentage point reduction in the probability of using marijuana. Similarly, a one-day increase in the frequency of alcohol use leads to a 0.23 day reduction in the frequency of marijuana use. While the estimated *rates* of substitution may appear small, the results of the logarithmic specifications in Table 5 show that, due to the smaller baseline use of marijuana, the corresponding *elasticities* are fairly large. For the entire population, we estimate statistically significant elasticities of substitution of 0.7 for the probability and 0.4 for the frequency of marijuana use.

In addition, the results in Tables 4 and 5 suggest that the substitution between alcohol and marijuana is larger for women than for men. For women, the estimated rates of substitution are 0.3 for the probability of use and 0.27 for the frequency of use, while the estimated elasticities are 1.2 for the probability and 0.5 for the frequency. For men, the estimates are substantially smaller – especially for the elasticities of substitution, which are estimated at 0.4 for the probability and 0.34 for the frequency of use – and not statistically significant. Overall, these results suggest that policies that limit access to alcohol are likely to lead to a substantially larger increase in the marijuana consumption of women than of men.

4.3. Can under-reporting of substance use explain the results?

Since our data are based on self-reported substance use, one concern is that our estimates are affected by under-reporting.

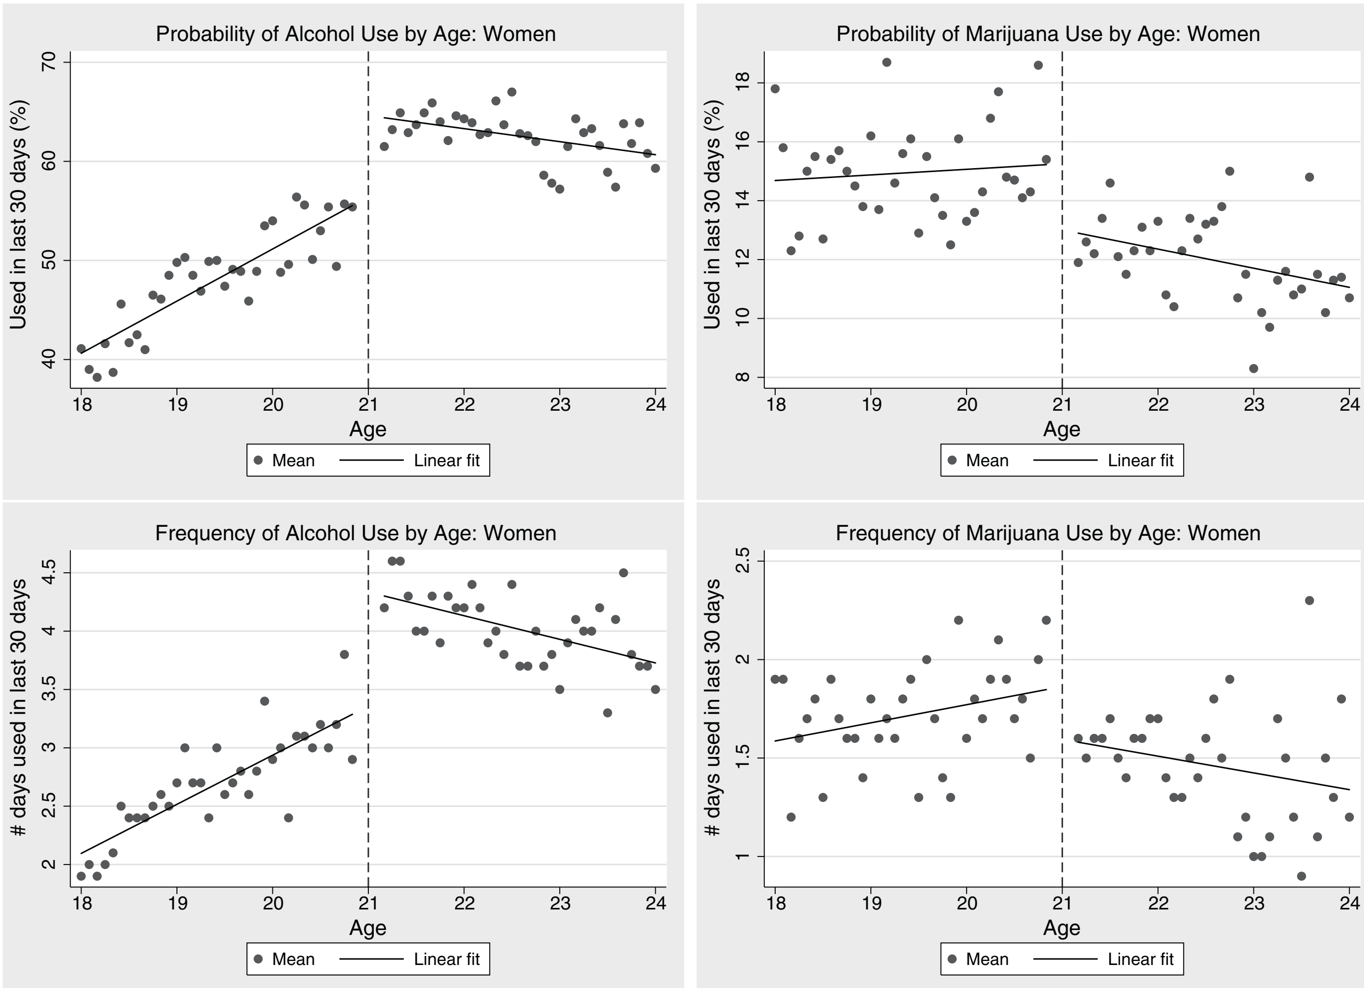

Fig. 3. Robustness of results to choice of bandwidth. In the top 4 panels, the vertical axis plots RD estimates of the effect of the MLDA on marijuana use. In the bottom 4 panels, the vertical axis plots IV estimates of the rate/elasticity of substitution between marijuana and alcohol use. The bandwidth is specified on the horizontal axis. Intervals are 95% confidence intervals. For details on the RD estimation, see the description in Table 1.

Data source: NSDUH 2002–2007.

However, for the regression discontinuity (RD) estimates, under-reporting is only a concern if individuals just under the age of 21 are more (or less) likely to under-report substance use than individuals just above the age of 21. This is likely to be the case for alcohol use, since drinking is illegal under the age of 21 and individuals may be unwilling to admit illegal behavior. The self-reported data are therefore likely to underestimate the alcohol use of individuals under 21, which leads to an upward bias in the RD estimates of the effect of the minimum drinking age on alcohol use.[3] However, under-reporting is unlikely to affect our estimates of the effect of the MLDA on marijuana use. Since marijuana use is illegal at all ages, there is no reason to expect that individuals just under the age of 21 would under-report marijuana use more or less strongly than individuals just above the age of 21. The changes in marijuana use at age 21 are therefore likely to be driven by changes in alcohol availability and not by changes in misreporting. It should, however, be kept in mind that our estimates of the substitution between alcohol and marijuana are likely to be biased downward, since the effect of the MLDA on alcohol use, which enters in the denominator, is most likely biased upward.

5. Robustness tests

5.1. Placebo tests for location of the discontinuity

In order to make sure that our estimated decline in marijuana use is driven by the change in alcohol accessibility at age 21 and not than merely due to a time trend that follows an inverted u-shape, we conduct placebo tests for the location of the discontinuity. For these tests, we estimate the same regression as in Table 1, but vary the location of the threshold. If there really is a discontinuity at age 21 the estimated change in substance use should be largest if the regression's threshold is located close to 21 years.

Fig. 2 displays the estimated change across the threshold as a function of the threshold's location. The top panels show the estimates for alcohol use, the bottom panels show the estimates for marijuana use. The figure shows that for alcohol, the estimates discontinuity is in fact largest when the threshold at age 20.9 rather than 21, which may be due to random fluctuations of alcohol use

[3] It is, however, unlikely that the effect of the MLDA on alcohol use is entirely an artifact of the survey methodology. Carpenter and Dobkin (2009) find a large spike in mortality, particularly from car accidents, at the age of 21, which suggests that there is a real increase in alcohol use at that age.

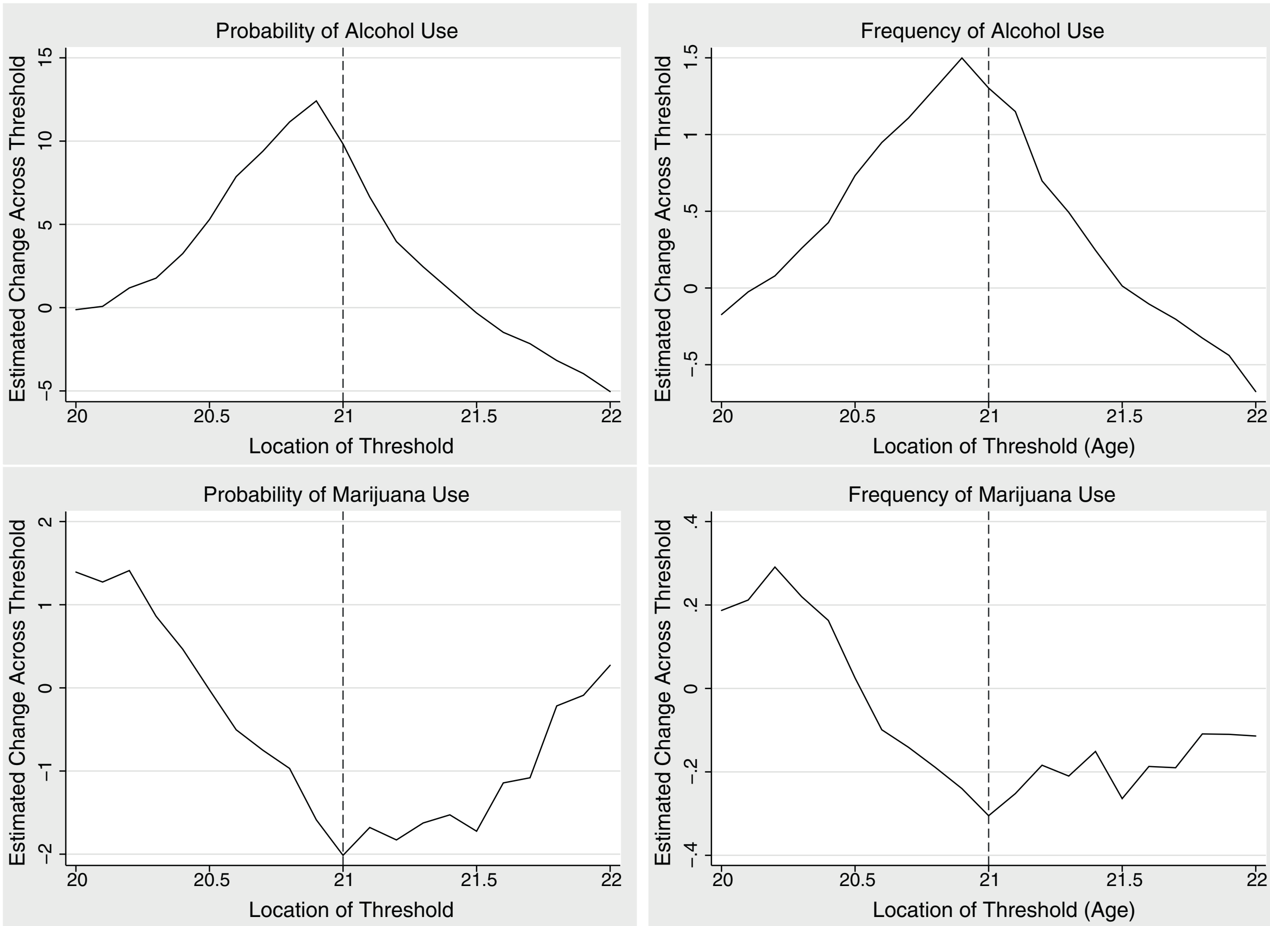

Fig. 4. Alcohol and marijuana use around age 21: men. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21. Data source: NSDUH 2002–2007, subsample of male respondents.

in the sample. Another explanation is that alcohol use is steeply upward sloping below the age of 21 and downward sloping above it. Therefore, moving the threshold to the left leads to a big reduction in the estimated alcohol use right below the threshold and only a small (if any) reduction of the estimated use right above the threshold. Either way, the bottom panels of the figure show that for both probability and frequency of marijuana use, the estimated decrease in use is largest if the threshold is located at exactly 21 years. While the results for alcohol use suggest that this placebo test is not always completely precise, it does increase our confidence that the estimates reported above are driven by the discontinuous change in alcohol accessibility at age 21 rather than by a time trend of marijuana use that follows an inverted u-shape.

5.2. Robustness tests for choice of bandwidth

A crucial parameter for local linear regressions like the ones reported above is the choice of bandwidth. By choosing a bandwidth that is too small we reduce the effective sample size and obtain estimates of low precision. By choosing a bandwidth that is too large we increase the risk of mis-specification if the relationship between age and substance use is non-linear. Though some authors have suggested rules-of thumb for bandwidth choice (e.g. Fan and Gijbels, 1996), no rule-of-thumb guarantees an optimal choice of bandwidth. Imbens and Lemieux (2008) therefore suggest robustness tests for different choices of bandwidth.

The results of these tests are reported in Fig. 3. For the tests, we gradually decrease the bandwidth until we reach half of the initial bandwidth of 3 years. The point estimates and 95% confidence intervals are plotted on the vertical axes of the graphs, against the bandwidth on the horizontal axes. If the estimates based on larger bandwidths suffer from specification bias due to a non-linear relationship between age and substance use, we would expect the point estimates to change substantially as the bandwidth becomes smaller, since the linear functional form better approximates the true relationship over smaller intervals. If there is no specification bias, we would expect the point estimates to have small fluctuations and hence be robust to the choice of bandwidth. Since the estimates for smaller bandwidths are based on smaller effective samples, we naturally expect the confidence intervals to increase as the bandwidth becomes smaller.

The results in Fig. 3 show that the estimates are robust to the choice of bandwidth. The point estimates differ very little for bandwidths between 18 months and 3 years. As expected, the size of the confidence intervals increases for smaller bandwidths, since fewer observations are used for estimation. Nevertheless, the estimated effect on the probability of marijuana use is statistically significant for all tested bandwidths; the effect on the frequency of marijuana use is significant at the 10% level for all tested bandwidths (not shown) and at the 5% level for bandwidths of 2 years and larger.

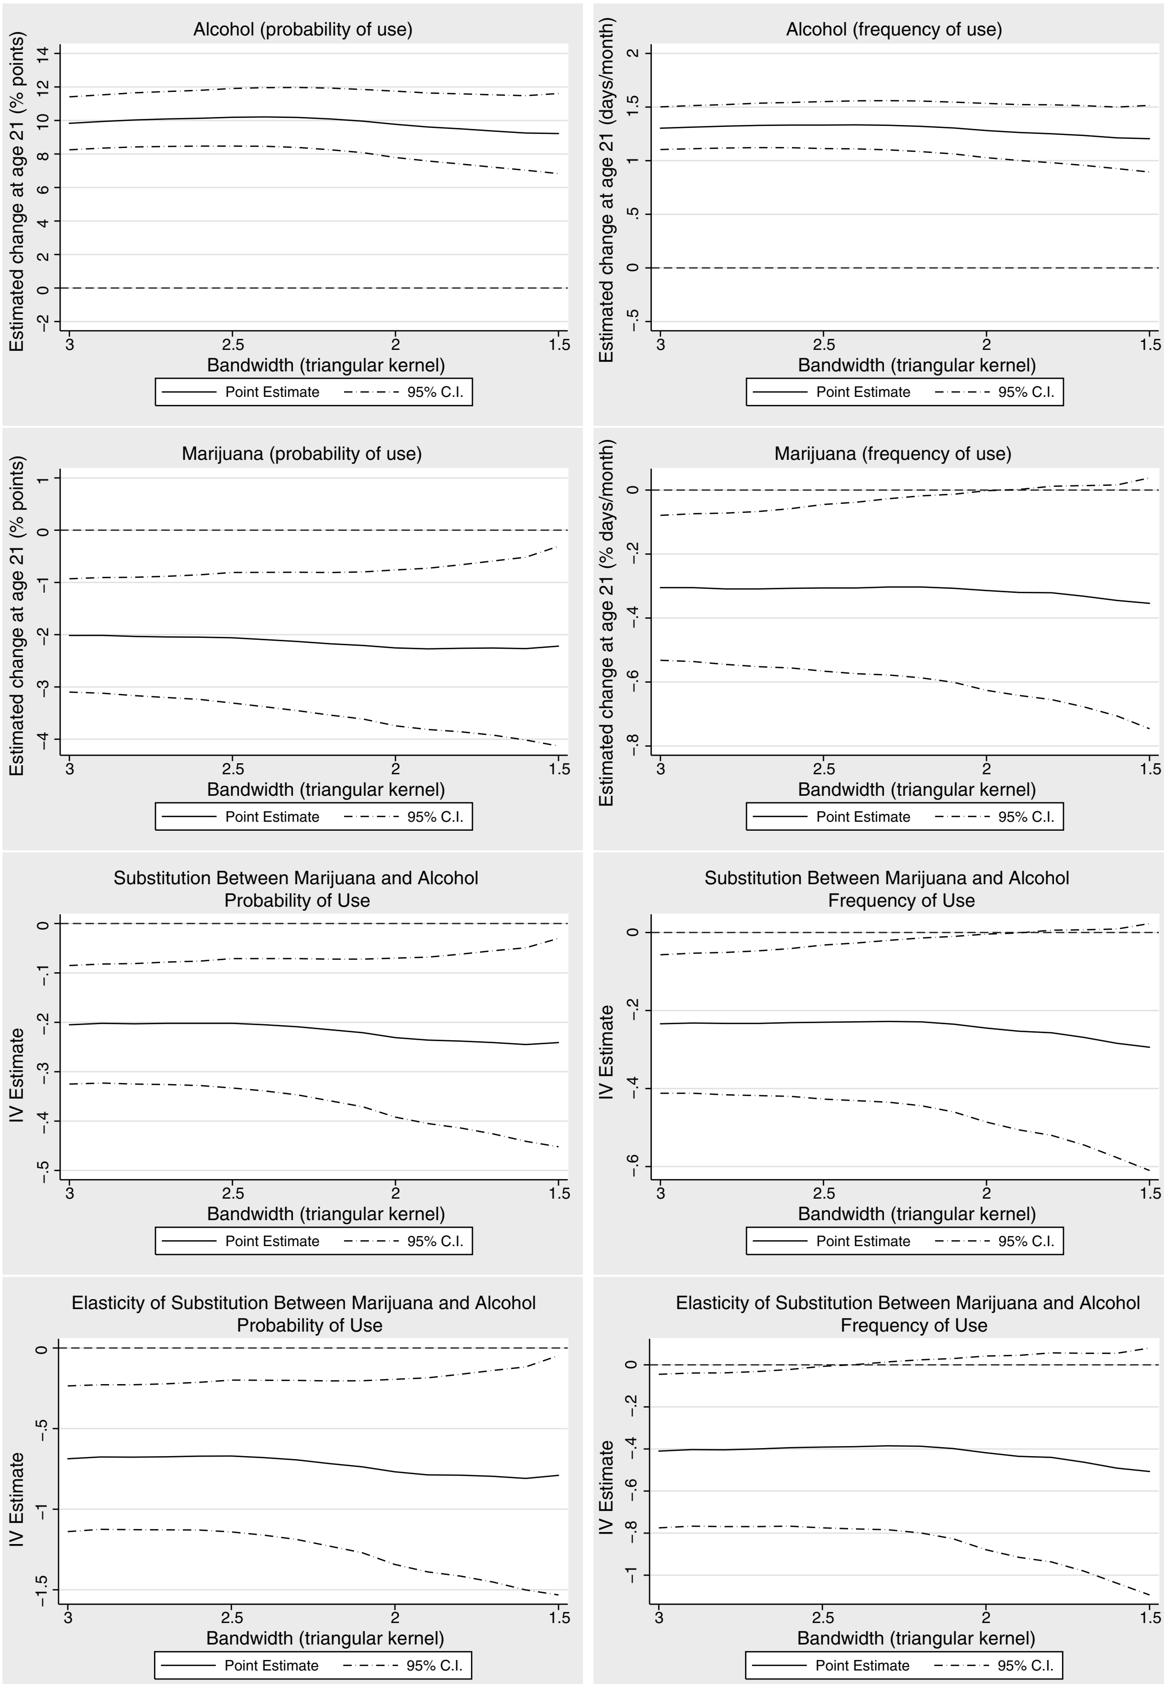

Fig. 5. Alcohol and marijuana use around age 21: women. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21. Data source: NSDUH 2002–2007, subsample of female respondents.

6. Conclusions

By exploiting the sharp decrease in the effective cost of alcohol consumption induced by the minimum legal drinking age (MLDA) at age 21, this paper estimates the causal effect of legal access to alcohol on marijuana consumption. Our identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just above and just below the age of 21 are similar in all characteristics that determine substance use. Compared to previous research (e.g. DiNardo and Lemieux, 2001; Chaloupka and Laixuthai, 1997; Pacula, 1998; Williams et al., 2004; Saffer and Chaloupka, 1999; Farrelly et al., 1999), this approach has the advantage of not having to rely on cross-sectional (often state-level) variation in alcohol and marijuana prices and related policies, which are likely to be correlated with unobserved characteristics of the population. This allows us to cleanly identify the causal effect of the MLDA in a way that is not afflicted by omitted variable bias.

Our results show that legal access to alcohol causes a significant decrease in marijuana use among young adults close to the age of 21. The point estimates suggest that marginally lowering the MLDA would decrease the probability of marijuana consumption in the affected age group by about 10%. The substitution effect is substantially larger for women than for men. Our results suggest that marijuana and alcohol are substitutes, so that a decrease in the ‘full’ price of alcohol (including the cost of access) leads to a decrease in marijuana use. The main implication of our study is that policies, such as the MLDA, that are aimed at restricting alcohol consumption among young adults are likely to have the unintended consequence of increasing the use of illegal drugs such as marijuana. When assessing the net benefits of alcohol-related policies these substitution effects need to be taken into account in order to assess the trade-off between the positive health effects from reduced alcohol consumption and the negative effects of increased use of other substances.

References

American Medical Association, 2008. Harmful Consequences of Alcohol Use on the Brains of Children, Adolescents, and College Students. American Medical Association.

Carpenter, C., Dobkin, C., 2009. The effect of alcohol consumption on mortality: regression discontinuity evidence from the minimum drinking age. American Economic Journal: Applied Economics 1 (1), 162–182.

Chaloupka, F.J., Laixuthai, A., 1997. Do youths substitute alcohol and marijuana? Some econometric evidence. Eastern Economics Journal 23 (3), 253–276.

DiNardo, J., Lemieux, T., 2001. Alcohol, marijuana, and American youth: the unintended consequences of government regulation. Journal of Health Economics 20 (6), 991–1010.

Fan, J., Gijbels, I., 1996. Local Polynomial Modeling and its Applications. Chapman and Hall, London.

Farrelly, M.C., Bray, J.W., Zarkin, G.A., Wendling, B.W., Pacula, R.L., 1999. The effects of prices and policies on the demand for marijuana: evidence from the National Household Surveys on Drug Abuse. NBER Working Paper Series, 6940.

Hall, W., Degenhardt, L., 2009. Adverse health effects of non-medical cannabis use. Lancet 374 (9698), 1383–1391.

Imbens, G.W., Lemieux, T., 2008. Regression discontinuity designs: a guide to practice. Journal of Econometrics 142 (2), 615–635.

Pacula, R.L., 1998. Does increasing the beer tax reduce marijuana consumption. Journal of Health Economics 17 (5), 557–586.

Pope, H.G., Gruber, A.J., Hudson, J.I., Cohane, G., Huestis, M.A., Yurgelun-Todd, D., 2003. Early-onset cannabis use and cognitive deficits: what is the nature of the association? Drug and Alcohol Dependence 69 (3).

Saffer, H., Chaloupka, F.J., 1999. Demographic differentials in the demand for alcohol and illicit drugs. NBER Working Paper Series, 6432.

Thies, C.F., Register, C.A., 1993. Decriminalization of marijuana and the demand for alcohol, marijuana, and cocaine. The Social Science Journal 30 (4), 385–399.

Williams, J., Pacula, R.L., Chaloupka, F.J., 2004. Alcohol and marijuana use among college students: economic complements or substitutes? Health Economics 13 (9), 825–843.

Wilson, W., Mathew, R., Turkington, T., Hawk, T., Edward Coleman, R., Provenzale, J., 2000. Brain morphological changes and early marijuana use: a magnetic resonance and positron emission tomography study. Journal of Addictive Diseases 19 (1), 1–21.