Journal of Health Economics 32 (2013) 22–32

Contents lists available at SciVerse ScienceDirect

Journal of Health Economics

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

Alcohol and student performance: Estimating the effect of legal access

Jason M. Lindo[a,b,c], Isaac D. Swensen[a], Glen R. Waddell[a,b,*]

[a] University of Oregon, United States
[b] IZA, Germany
[c] NBER, United States

ARTICLE INFO

Article history:
Received 5 December 2011
Received in revised form 7 September 2012
Accepted 28 September 2012
Available online xxx

JEL classification:
I21
I18
K32

Keywords:
Alcohol
Post-secondary education
Minimum legal drinking age

ABSTRACT

We consider the effect of legal access to alcohol on student achievement. Our preferred approach identifies the effect through changes in one's performance after gaining legal access to alcohol, controlling flexibly for the expected evolution of grades as one makes progress towards their degree. We also report RD-based estimates but argue that an RD design is not well suited to the research question in our setting. We find that students' grades fall below their expected levels upon being able to drink legally, but by less than previously documented. We also show that there are effects on women and that the effects are persistent. Using the 1997 National Longitudinal Survey of Youth, we show that students drink more often after legal access but do not consume more drinks on days on which they drink.

1. Introduction

A large literature links alcohol consumption to adverse health and social outcomes.[1] Given long-standing and persistent efforts to restrict access to alcohol, it is no surprise that this topic has received considerable attention from researchers. However, relatively little is known about the effect of legal access to alcohol on the academic performance of students in college, where binge drinking is often cited as a serious and growing problem (DeSimone, 2007). That alcohol is associated with acute outcomes such as crime, mortality, and sexual activity gives cause for concern that the effect on student performance may be quite large.

In this paper, we assess the effect of legal access to alcohol on academic performance using two identification strategies. The first has been used to address this research question in a different setting and the second has not, but both approaches exploit the exogenous change in legality induced by the federally mandated minimum legal drinking age (MLDA). Our first identification strategy follows Carrell et al. (2011) who exploit the sharp change in legality that occurs at age 21 in a regression discontinuity (RD) framework to estimate the effect of legal access on student performance. While it is relatively straightforward to use an RD design to estimate effects of turning 21 on crime or traffic accidents, as in Carpenter and Dobkin's works, it is less straightforward as an approach to estimating effects on academic outcomes since they are not measured frequently. For this reason, the RD approach uses age from 21 *at the end of the academic term* as the running variable. As such, the estimates capture the effect of legal access to alcohol for students who obtain access near the end of the academic term. In the limit, the thought experiment compares the performance of students who turn 21 the day before their final exam to the performance of students who turn 21 on the day of their final exam. The resulting estimates can therefore be characterized as measuring a local average treatment effect, which may have limited external validity.

Our second and preferred identification strategy overcomes this limitation by making use of the longitudinal nature of the data. In particular, we identify the effect of legal access to

* Corresponding author at: University of Oregon, United States. Tel.: +1 5413461259.
E-mail addresses: jlindo@uoregon.edu (J.M. Lindo), isaac@uoregon.edu (I.D. Swensen), waddell@uoregon.edu (G.R. Waddell).

[1] In particular, quasi-experimental methods have been used to consider effects on mortality (Dee, 1999; Carpenter, 2004; Carpenter and Dobkin, 2009), crime (Markowitz and Grossman, 1998; Carpenter, 2005a, 2007; Carpenter and Dobkin, 2010), sexual activity (Chesson et al., 2000; Rees et al., 2001; Sen, 2002; Rashad and Kaestner, 2004; Carpenter, 2005b; Waddell, 2012), employment (Mullahy and Sindelar, 1996; Terza, 2002; Dave and Kaestner, 2002; MacDonald and Shields, 2004; Renna, 2008), and teenagers' educational outcomes (Cook and Moore, 1993; Dee and Evans, 2003; Chatterji and DeSimone, 2006), among others.

0167-6296/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jhealeco.2012.09.009

alcohol by comparing a student's post-21 academic performance to his own pre-21 academic performance with individual-fixed effects models—implicitly arguing that the best counterfactual for a student's post-21 performance is his own performance prior to turning 21. In addition, our regressions include fixed effects for the number of accumulated credits to account for the possibility that students may systematically improve, "slack off," or experience classes of different difficulty as they progress towards degree completion. Although it is not typical for researchers to be able to separately identify effects of experience (or accumulated credits in our case) and the effects of age (or an age-dependent treatment in our case), we can do so in our context by leveraging the variation in college starting ages.[2] As in the first approach, we use a student's course performance relative to their classmates' as our outcome variable, which will also serve to control for selection into courses.[3]

The data and institutional setting that we consider—transcript-level data from undergraduates at the University of Oregon—allow us to make several additional contributions to the literature. One of the unique features of Carrell et al.'s (2011) study using data from the U.S. Air Force Academy is that the prohibition on underage drinking is taken extremely seriously there—much more so than in most other institutional settings, where enforcement is more lax and punishment less severe.[4] As such, assuming Air Force Academy students are representative of the general student population, their RD estimates tell us about the local average treatment effect of prohibition in environments where enforcement and penalties are unusually strict. In contrast, our results are more likely to speak to the effect of minimum drinking age laws *as they are conventionally enforced* and, in turn, the effect of the changes in drinking behavior that are typically associated with legal access to alcohol. As we describe in the next section, the University of Oregon is also more representative of U.S. institutions, which leaves us more confident that estimates based on these data will have greater external validity. Further, our data include over four times the number of observations used in this earlier research, and approximately ten times the number of females which allows a more-precise consideration of heterogeneity across gender.

The results from our preferred approach indicate that students' grades fall below their expected levels by approximately 0.03 standard deviations upon being able to drink legally, a modest amount compared to the 0.06 to 0.13 standard-deviation effect estimated in earlier research. The effect is statistically significant, manifests in the term a student turns 21, and persists into later academic terms. In addition, we find that the effects on academic performance are especially large for females, low-ability males, and males who are most likely from financially disadvantaged backgrounds.

In order to shed light on the mechanisms underlying these effects, we also analyze the effects of legal alcohol access on various drinking-related behaviors among students enrolled in four-year colleges using data from the 1997 National Longitudinal Survey of Youth (NLSY97). This analysis reveals that college students drink more often upon gaining legal access but do not consume more drinks on days in which they drink.[5]

The rest of this paper is organized as follows. In Section 2 we discuss the data used in our analysis as well as the representative nature of the University of Oregon campus. In Section 3 we present an RD strategy, following existing literature, and our preferred longitudinal approach to identifying the influence of legal access to alcohol on academic performance. In Section 4 we present and discuss the empirical results, including our analysis of how college students' drinking-related behaviors change upon gaining legal access to alcohol based on the NLSY97. We conclude and discuss the implications of our results in Section 5.

2. Data

Our primary data are administrative student-course level data from the University of Oregon, spanning fall 1999 to winter 2007, for students entering at 18 or 19 years old. We focus on performance during the fall, winter, and spring terms.[6] Because our identification strategies use variation provided by the federal MLDA law, we also restrict the sample to students observed at least one term in which they are at least 21 years old. The resulting sample consists of 479,342 observations over 13,102 students.

As one contribution of this paper is to provide insight into the effects of MLDA laws in a "typical-college setting," Table 1 compares characteristics of students at the University of Oregon to those at other U.S. public four-year institutions. While Column 1 provides summary statistics based on our sample, Column 2 considers a more comprehensive set of characteristics based on data from the Integrated Postsecondary Education Data System (IPEDS). Similarly, Column 3 shows statistics on other public four-year institutions, also using data from IPEDS.[7]

Table 1 largely supports that the University of Oregon provides a representative-college setting. While it is twice the size and has higher admission rates than the average public four-year

[2] As an alternative to our preferred models that control for accumulated-credit fixed effects, we have also estimated models that instead control for the age profile of student grades. Whereas the former relies on the assumption that the effects of progress are similar for those with and those without legal access, this alternative approach relies on the assumption that the age profile of grades is well captured by the chosen parametric specification. The estimates based on both approaches are statistically significant. The estimates based on the accumulated-credit fixed effects models, however, are more conservative.

[3] In related studies, Williams et al. (2003) and Powell et al. (2004) consider the effect of alcohol consumption on college GPAs using data from the Harvard School of Public Health's College Alcohol Study. These studies involve cross-institution comparisons of student GPAs, with measures of alcohol costs serving as instruments for drinking intensity among those who drink. Kremer and Levy (2008) consider a different-but-related question, exploiting the random assignment of roommates at a large state university in order to identify the effect of having a roommate who drinks. Finally, Lindo et al. (2012) consider how non-athlete academic performance varies with collegiate football success and provide survey evidence of increased alcohol consumption and partying associated with football success.

[4] Carrell et al. (2011) highlight this feature, pointing to the fact that two incidents of underage drinking at the Air Force Academy resulted in expulsion and that some related incidents (e.g., driving under the influence) have also resulted in expulsion.

[5] Though the NLSY97 also includes information on GPAs in college, these data have several shortcomings that preclude an analysis similar to that which we perform on our transcript data. Most crucially, the NLSY97 does not offer transcript data on academic performance but instead has retrospective self-reports of GPA by term, leaving much opportunity for measurement error. In particular, in each survey, the students are asked about their GPAs in all academic terms since they were last surveyed. This time period most frequently spans one year but sometimes spans several. Possibly as a result of this survey design, there are also many cases in which it appears as if the performance in a given term may have been reported more than once (in different survey years). Moreover, several features of these data make it difficult to accurately assign whether the student had legal access to alcohol in a given term. Specifically, the month and year in which the term began can be recorded but information on term length is more often missing than not. In our best attempts to impute the data where necessary, NLSY97 data yield a point estimate that is negative (suggesting a negative effect of legal access on GPAs) but not statistically significant.

[6] One reason for excluding summer terms is that summer enrollment could itself be considered an outcome variable. In addition, summer terms tend to be fundamentally different from other terms in class size, course offerings, student attendance, teacher and student attributes, and term structure.

[7] In comparing across institutions we have used variables that provide a snapshot of school admissions and graduation rates, general academic standards, undergraduate student demographics, and student financial costs and aid. The statistics reported in Columns 2 and 3 are based primarily on the 2003–2004 academic year, which is close to the median year for our data.

Table 1
Summary statistics.

	Oregon (Sample)	Oregon (IPEDS)	Four-year Public U.S. Institutions (IPEDS)	NLSY97 (Sample)
SAT I Verbal 25th percentile score, incoming students	500	490	464	470
SAT I Verbal 75th percentile score, incoming students	620	610	568	600
SAT I Math 25th percentile score, incoming students	500	500	472	470
SAT I Math 75th percentile score, incoming students	620	610	578	620
Number of undergraduates	13,102	15,983	8674	2298
Fraction female	0.55	0.53	0.55	0.54
Fraction white	0.79	0.75	0.67	0.81
Fraction black	0.02	0.02	0.11	0.11
Fraction Hispanic	0.03	0.03	0.08	0.07
Fraction Asian	0.08	0.12	0.11	0.00
Total price for in-state students living on campus		14,734	13,272	
Total price out-of-state students living on campus		26,170	20,022	
Fraction receiving any financial aid		0.70	0.75	
Fraction receiving federal-grant aid		0.18	0.34	
Fraction receiving student-loan aid		0.40	0.45	

Notes: Data used in the first columns consists of University of Oregon undergraduates from 1998 through 2007. Financial aid statistics shown in the subsequent two columns are calculated using 2004 IPEDS data, while all other statistics in the same columns are calculated using 2003 IPEDS data. The number of institutions used to calculate the means in the third column range from 352 to 653. NLSY97 sample statistics use the last observed sampling weight for each respondent.

institution, it is similar in terms of enrollment rates and in the ability of enrolled students as measured by SAT scores. It is also similar to the average college in terms of costs and financial aid. Like most other institutions, the University of Oregon is over half female and predominately white, although at seventy-five percent it has a larger share of white students than average.[8]

In contrast, the U.S. Air Force Academy, the only other institution where this research question has been addressed, offers a relatively unique setting. In addition to being highly selective, it is very different from most schools in terms of its students' objectives. In particular, all students at the Air Force Academy are given full scholarships but are expected to serve a five-year commitment as a commissioned officer in the U.S. Air Force following graduation. Moreover, females comprise only eighteen percent of its student body, which stands in stark contrast to the nation-wide average of fifty-five percent. As mentioned in the introduction, it is also important to note that the Air Force Academy is an outlier in strongly enforcing the MLDA law. That students at the Air Force Academy are such a select group from the distribution of all students, in both ability and preferences, and that they are in an environment that is unusually strict with respect to underage drinking, gives cause for concern about the external validity of earlier estimates and highlights the importance of considering the research question in different contexts.

Our secondary data source is a subsample of respondents from the 1997 National Longitudinal Survey of Youth (NLSY97), which allows us to speak to the sensitivity of alcohol consumption to legality among college goers. To the extent possible, our sample restrictions on these data reflect those above. In particular, we restrict the NLSY97 sample to individuals who report being currently enrolled in a four-year college, who were first observed in college at age 18 or 19, and also observed in college after turning 21.[9] We also balance the sample by removing observations when any of the outcome variables we consider (described below) are missing. Ultimately, the sample spans 1998 to 2009, includes 2298 person observations and 9023 person-year observations. Summary statistics for the NLSY97 data are provided in the last column of Table 1. Like the University of Oregon, the NLSY97 offers a sample that is similar in ability to the broader group of students attending four-year public institutions and its gender composition, but differs in some important ways in its racial composition. In particular, like the University of Oregon, it provides a sample that is disproportionately white; however, whereas the University of Oregon has very relatively few black students, the NLSY97 has relatively few Asian students.

3. Empirical strategy

3.1. RD-based approach

Following the prior literature, we begin by estimating the effect of having had one's 21st birthday before the end of the academic term on academic performance using the following regression equation:

$$G_{ijt} = \alpha_0 + \alpha_1 1\{\mathrm{AGE}_{it} \geq 0\} + f(\mathrm{AGE}_{it}) + \epsilon_{ijt}, \tag{1}$$

where G_{ijt} is the normalized grade for student i in class j in term t. Normalized grades are calculated as the deviation in a student's grade from the class mean divided by the class standard deviation. AGE_{it} is the student's age at the end of the term in days, centered on 21 years. For example, in the comparison of means as estimates approach the treatment threshold from each side, a bandwidth of 90 days would put weight on all students who had their 21st birthday in the range 90 days prior to the end of the term (i.e., AGE = 90) through 90 days after the end of the term (i.e., AGE = −90). Last, $f(\mathrm{AGE}_{it})$ controls for a student's age at the end of the term in a flexible manner. In practice, we estimate a variety of models, including those that do not control for age at the end of the term, those that control for age at the end of the term with a linear specification flexible on each side of the cutoff, and those that control for age at the end of the term with a quadratic specification flexible on each side of the cutoff. Further, we consider a variety of bandwidths between 20 days and 240 days.

[8] Also like most other large public institutions, a majority freshman at the University of Oregon live on campus (87 percent) while sophomores, juniors, and seniors do not often live on campus (7, 3, and 2 percent, respectively).

[9] An individual's age at the time of the interview is based on the interview date along with the individual's month and year of birth. We impute the fifteenth as the day of birth. Though this will lead some person-year observations to be misclassified in terms of legal access, few observations are on the margin for this to make a difference. 3790 observations are classified as having legal access when the fifteenth of the month is imputed as the day of birth; this count becomes 3827 and 3719, respectively, if the first or twenty-eighth of the month are used instead.

It is important to note that this identification strategy departs from the usual RD exercise. Typically, we observe—or know as a result of institutional detail—the extent to which the treatment of interest jumps on the "treatment side" of the threshold. For example, in DiNardo and Lee's (2004) unionization study, all elections with union support greater than fifty percent lead to unionization while elections with less support do not. Similarly, in Angrist and Lavy's (1999) class-size study, we observe class-size reductions above multiples of forty enrolled students. Our example is similar insofar as all students on the "treatment side" of the threshold have had the opportunity to drink alcohol legally prior to the conclusion of the academic term. However, because the underlying effects on drinking behavior is unobserved, the magnitude of any estimated effect will be somewhat difficult to interpret. Even though drinking tends to increase when a student turns 21, as we show in Section 4.3, we do not know to what extent this holds true for students who turn 21 near the end of an academic term, which this identification strategy pre-supposes. As such, the comparison involved with this RD approach is informative about the effect of drinking on college performance but its "local" nature (close to 21 *and* close to the end of the term) introduces additional interpretive challenges.[10]

In the absence of estimated effects on drinking behavior for such a sample, the results are appropriately characterized as intent-to-treat effects, measuring the reduced-form effect of the minimum-drinking-age law. While this is certainly of interest in itself, the RD design only provides an estimate of a very local intent-to-treat effect, corresponding to students gaining legal access to alcohol at the end of the academic term. This will remain a disadvantage of the RD approach in this setting, something that we improve on with the identification strategy presented in the next section where we exploit the longitudinal nature of the data.

3.2. Individual fixed effects approach

Our preferred approach to estimating the effects of legal access to alcohol focuses on within-student variation over time. Although we first present estimates from more-parsimonious models, we ultimately arrive at estimates derived from the following regression:

$$G_{ijt} = \theta \text{AGE21}_{it} + \beta X_{ijt} + \alpha_i + u_{ijt}, \tag{2}$$

where G_{ijt} is the normalized grade for student i in class j in term t, AGE21_{it} is an indicator variable that takes a value of one if the student could drink legally at any time during term t and zero otherwise, X_{ijt} can include term- or class-varying individual characteristics, α_i are a set of individual fixed effects, and u_{ijt} is a random error term. In our analysis of transcript data from the University of Oregon, we always include "experience controls" in X_{ijt}, i.e., fixed effects for the number of accumulated credits (in four-credit intervals) to control for grade changes that are expected as a student progresses toward degree completion.[11] For example, these variables are intended to control for phenomena such as "senioritis." As such, the estimation strategy compares a student's grades after turning 21 to what would be expected based on his average prior performance and accumulated experience.[12] This approach relies on the assumption that accumulated experience has a similar effect on grades before and after a student turns 21. As an alternative to fixed effects for accumulated credits, one might anticipate that we would model the age profile of academic performance, and consider deviations from the predicted grade upon turning 21 as a reflection of the causal effect of legal access to alcohol. Each of these two approaches offers conceptual advantages and both support our conclusions. We limit the presentation and discussion of the results to the former approach, which yields the more-conservative estimates.

Our approach to estimating the effect of legal alcohol access on drinking behavior in the NLSY97 is quite similar with the exception of the controls we include in the model. In particular, since the data are at the individual-by-survey-year level (asking about drinking behavior "in the past 30 days") as opposed to the term-level, we estimate the effects using the model that controls for age at the time of the interview with a quadratic as opposed to the number of credits that the student has acquired in college.[13]

4. Results

In this section, we first present RD-based estimates of the effects of legal access on academic grades using Oregon transcript data. We then we report the results from our preferred approach—fixed-effect based estimates—of the same. Finally, we report on the NLSY97 analyses, where we estimate the response of alcohol consumption and drinking behavior to legalization.

4.1. RD-based estimates using transcript data

Panel A of Table 2 presents RD-based estimates of the effect of legal access to alcohol at the end of a term on academic performance. Across the fourteen columns, the table shows estimates based on a wide range of bandwidths and functional-form choices. While the upper portion of this panel reports unadjusted estimates, the lower portion reports estimates that controls for course-by-quarter-by-year fixed effects, birth-year fixed effects, accumulated-credits fixed effects, gender, math and verbal SAT scores, high-school GPA, and indicator variables for university athletes, private high school attendance, race, and ethnicity.[14]

Overall, the set of results in Panel A of Table 2 suggests that turning 21 before the end of a term has a negative impact on a student's grades. While the point estimates vary somewhat from specification to specification and are sensitive to control variables, they are routinely negative and suggest that students who turn 21 prior to the end of the quarter score roughly 0.03 to 0.05 standard deviations lower than those who turn 21 after the quarter ends. However, the sensitivity of RD estimates to the inclusion of controls—primarily the inclusion of individual characteristics and accumulated credits—casts doubt on the validity of this strategy in our setting.

[10] We note that all RD-based studies that consider the effect of being able to drink legally are local in the first (close to 21) sense but that the second sense is specific to this application, driven by the fact that outcomes are not measured daily.

[11] While it would be attractive to also include fixed effects for the number of terms a student has been at the university, doing so is likely to introduce problems of multicollinearity in conjunction with the individual fixed effects and accumulated-credits fixed effects since there is little variation in credits attempted each term. For example, such a model would be impossible to estimate if all students earned 12 credits each term. We have explored models that include fixed effects for the number of terms a student has been at the university *instead* of the accumulated-credits fixed effects and the results are quite similar.

[12] We also estimate models that control for course characteristics.

[13] One could conceivably control for accumulated credits using the NLSY97, but the number of credits a student has taken in each term is often missing. Moreover, lacking the ability to track transfers from school to school and to determine how many classes have been passed or failed, we would not expect such a measure to reflect true progress towards degree completion. Thus, we rely on the assumption that the age profile of drinking behavior is well captured by the chosen parametric specification.

[14] Race and ethnicity controls consists of a set of indicator variables for being black, Hispanic, or Asian.

Table 2
RD-based estimates of the effect of turning 21 (and other ages) at the end of term on grades.

Bandwidth Age polynomial	240 days Linear (1)	240 days Quadratic (2)	210 days Linear (3)	210 days Quadratic (4)	180 days Linear (5)	180 days Quadratic (6)	150 days Linear (7)	120 days Linear (8)	100 days Linear (9)	80 days Linear (10)	80 days None (11)	60 days None (12)	40 days None (13)	20 days None (14)
Panel A: estimated effect of turning 21														
Estimated effect	−0.035***	−0.026	−0.033***	−0.030	−0.036**	−0.024	−0.031**	−0.024	−0.038*	−0.037	−0.038***	−0.048***	−0.031*	−0.027
	(0.011)	(0.019)	(0.012)	(0.021)	(0.014)	(0.021)	(0.014)	(0.017)	(0.020)	(0.024)	(0.009)	(0.012)	(0.016)	(0.023)
Estimated effect including controls	−0.026**	−0.031*	−0.023**	−0.040**	−0.030**	−0.029	−0.023	−0.018	−0.032	−0.045*	−0.075***	−0.080***	−0.072***	−0.047
	(0.011)	(0.018)	(0.012)	(0.020)	(0.014)	(0.020)	(0.015)	(0.018)	(0.021)	(0.025)	(0.013)	(0.016)	(0.021)	(0.036)
Observations	156,956	156,956	138,574	138,574	119,608	119,608	100,344	81,589	68,903	54,963	54,963	41,473	27,655	14,239
Panel B: estimated effect of turning 20														
Estimated effect	−0.017	−0.014	−0.013	−0.023	−0.021	−0.014	−0.009	−0.012	−0.027	−0.047**	−0.011	−0.020*	−0.027*	−0.040*
	(0.010)	(0.018)	(0.011)	(0.019)	(0.014)	(0.019)	(0.014)	(0.015)	(0.019)	(0.023)	(0.009)	(0.012)	(0.016)	(0.021)
Estimated effect including controls	−0.009	−0.010	−0.007	−0.014	−0.013	−0.003	0.008	0.006	−0.011	−0.036	−0.027**	−0.029**	−0.032*	−0.035
	(0.010)	(0.016)	(0.011)	(0.018)	(0.013)	(0.018)	(0.013)	(0.015)	(0.018)	(0.022)	(0.011)	(0.014)	(0.019)	(0.027)
Observations	163,568	163,568	144,184	144,184	123,830	123,830	103,701	83,931	70,338	55,546	55,546	41,762	27,790	13,985
Panel C: estimated effect of turning 22														
Estimated effect	−0.012	−0.016	−0.010	−0.025	−0.018	−0.016	−0.015	−0.014	−0.037	−0.034	−0.028**	−0.045***	−0.039**	0.002
	(0.013)	(0.022)	(0.014)	(0.025)	(0.016)	(0.025)	(0.017)	(0.020)	(0.024)	(0.029)	(0.011)	(0.014)	(0.019)	(0.028)
Estimated effect including controls	−0.001	−0.012	0.001	−0.022	−0.013	−0.005	−0.013	−0.008	−0.018	−0.025	−0.052***	−0.053***	−0.037	0.041
	(0.014)	(0.023)	(0.015)	(0.025)	(0.017)	(0.027)	(0.019)	(0.022)	(0.027)	(0.033)	(0.016)	(0.020)	(0.028)	(0.049)
Observations	114,397	114,397	102,009	102,009	88,277	88,277	74,321	60,664	51,108	40,810	40,810	30,747	20,670	10,444

Notes: The dependent variable is a student's normalized course grade. Controls include course-by-quarter-by-year fixed effects, birth-year fixed effects, accumulated-credits fixed effects, gender, math and verbal SAT scores, high-school GPA, and indicator variables for university athlete, private-school attendance, Black, Hispanic, and Asian. Standard errors (in parentheses) are corrected for clustering at the date-of-birth level.

* Significant at 10%.

** Significant at 5%.

*** Significant at 1%.

Table 3
Fixed-effects-based estimates of the effect of legal access to alcohol on grades.

	(1)	(2)	(3)	(4)
Age >21 during term	−0.146***	−0.097***	−0.033***	−0.033***
	(0.005)	(0.004)	(0.006)	(0.006)
Individual fixed effects	No	Yes	Yes	Yes
Accumulated-credits fixed effects	No	No	Yes	Yes
Course-specific controls	No	No	No	Yes
Number of students	13,102	13,102	13,102	13,102
Observations	479,342	479,342	479,342	479,342

Notes: The dependent variable is equal to the student's normalized course grade. Accumulated-credits fixed effects are fixed effects for a student's accumulated credits at the beginning of a term. Course-specific controls include subject-by-level fixed effects and term fixed effects. Standard errors (in parentheses) are corrected for clustering at the individual level. *Significant at 10%; **Significant at 5%.
*** Significant at 1%.

As a further robustness check, Panels B and C of Table 2 report the results from a similar exercise that instead considers the effect of turning 20 and 22, respectively, before a quarter ends. In particular, these results test for a more-general "birthday effect" which would raise the concern that the estimates in Panel A might reflect a "21st birthday effect" that cannot be separated from the effect of gaining legal access to alcohol. Although the estimates are usually not significant, the fact that 51 of the 56 point estimates are negative casts further doubt on the validity of this strategy in our setting. If we believe that the estimated "20th birthday effect" or "22nd birthday effect" provides a good estimate for the "21st birthday effect independent of legal access," then the difference between Panel A and either Panel B or Panel C can be interpreted as the effect of legal access at the end of the term. Given the marginal significance of most of the estimates in Panel A and the consistently negative estimates in Panels B and C, this approach would not produce convincing evidence that legal access at the end of the academic term significantly reduces grades.[15]

4.2. Fixed-effects estimates using transcript data

In this section, we present our main results followed by a consideration of the possible dynamic response to being able to drink legally (Section 4.2.2), and the heterogeneity of the established results by gender, ability, and financial need (Section 4.2.3).

4.2.1. Main results

In Table 3 we present our main results, making use of the longitudinal nature of the data. In Column 1, we show the estimated effect based on a regression of a student's normalized grade on an indicator for whether a student could drink legally at any time during the term. Because we anticipate that relatively low ability students will be observed more often at older ages (as they take longer to complete their degrees), we anticipate that this approach will overstate the negative effect of legal access to alcohol. After we control for ability and other unobservable characteristics with the inclusion of individual fixed effects, the estimate is indeed much smaller (falling from −0.146 to −0.097 from Column 1 to Column 2). However, estimates in Column 2 may still suffer from bias due to the potential for grades to fall as students progress towards their degrees while they become increasingly likely to be 21 years old. As anticipated, the magnitude of the estimate is even smaller when we remove this source of bias by controlling for a student's accumulated credits with fixed effects. That said, the point estimate (shown in Column 3) remains statistically significant at the one-percent level, indicating that a student's course-normalized grades fall by 0.033 standard deviations after they gain legal access to alcohol relative to what we would expect based on their prior performance and accumulated experience. The estimated effect is identical when we add controls for subject-by-level fixed effects and term fixed effects in Column 4, which is not surprising since our outcome variable is normalized at the class level.[16]

In order to better understand our main results in Table 3, in Table 4 we explore the effects of legal access on additional academic outcomes. To begin, we consider the distributional effects of legal alcohol access on grades. In particular, we use linear probability models (with the same controls used in the richest specification of Table 3) to separately estimate the effect of legal alcohol access on the probability that a student earns an A grade, a B grade, a C grade, and a D or F grade, respectively. These estimates suggest that the negative effect of legal alcohol access on grades overall is driven by its negative effect on the probability that a student earns an A grade and its positive effect on the probability that a student earns a C grade. The estimated effects on B grades and failing grades are negligible.[17]

Although the estimates in Table 3 address omitted variable bias that might be induced by effects on course-taking behavior by normalizing students' grades relative to their classmates and by controlling for course characteristics, any effect on course selection is of interest itself. We explore this issue Columns 5 and 6 of Table 4, considering the effect of legal alcohol access on course difficulty and course loads. This analysis is identical to the preceding analysis, except that it is conducted at the student-by-quarter level rather than the student-by-quarter-by-course level and, as such, omits course-level controls.

As our measure of course difficulty, we focus on a student's unconditional "expected" GPA. This is calculated based on the aver-

[15] Carrell et al. (2011) conduct a similar analysis using alternative age cutoffs and find no evidence of 20th or 22nd birthday effects at the U.S. Air Force Academy. In an alternative attempt to separate the birthday effect from that of a potentially-persistent effect of legal access to alcohol we have also explored the use of a donut-RD approach (Carpenter and Dobkin, 2009; Barreca et al., 2011a,b). In particular, we have conducted a similar analysis after dropping observations 1, 2, 3, 10, and 15 days to either side of the cutoff. This analysis continued to show similar estimates when considering the effect of turning 20 and 21.

[16] For these fixed effects, subjects correspond to economics, english, and mathematics. Levels correspond to either 100-, 200-, 300-, or 400-level classes. As summer terms are not considered as part of our analysis, terms are fall, winter, and spring.

[17] We have also estimated the effect on the probability that a student earns a quarterly GPA weakly above 2.0 (i.e., the GPA that would imply probation) and the probability that a student earns a quarterly GPA above 3.75 (i.e., the GPA that would imply Dean's List membership). These results indicate that legal access does not affect the probability that a student earns a 3.75 but increases the probability that a student is placed on academic probation, with point estimates (and standard errors) of −0.003 (0.002) and 0.019 (0.002), respectively. In addition, we have separately estimated the effects on grades in classes taken within a student's major and outside of a student's major. These results reveal larger effects for courses taken outside a student's major, with point estimates (and standard errors) of −0.058 (0.007) and −0.023 (0.009), respectively. Finally, we do not find any evidence that legal access affects the probability that students repeat classes.

Table 4
Fixed-effects-based estimates of the effect of legal access to alcohol on the grade distribution, course difficulty, and course load.

	(1) A grade	(2) B grade	(3) C grade	(4) D or F grade	(5) Expected term GPA	(6) Course load
Age >21	−0.008***	0.000	0.007***	0.000	0.009***	0.053
	(0.003)	(0.003)	(0.002)	(0.001)	(0.003)	(0.035)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Accumulated-credits fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Course-specific controls	Yes	Yes	Yes	Yes	–	–
Number of students	13,102	13,102	13,102	13,102	13,102	13,102
Observations	479,342	479,342	479,342	479,342	146,730	146,730

Notes: The analysis in Columns 1–4 is based on data at the at the student-by-course level whereas the analysis in Columns 5–6 is based on data at the student-by-term level (and thus does not control for course characteristics). The outcome variable for Column 5, a student's expected term GPA, is calculated based on the average grades in the previous offering of each course a student is taking in a given term. The outcome variable for Column 6, course load, is the number of credits taken in a term. Accumulated-credits fixed effects are fixed effects for a student's accumulated credits (in four-credit intervals) at the beginning of a term. Standard errors (in parentheses) are corrected for clustering at the individual level. *Significant at 10%; **Significant at 5%.
*** Significant at 1%.

age grades in the most-recent offerings of the courses a student has enrolled in for their current quarter. The estimated effect on this outcome is 0.009, which suggests that students take slightly easier classes upon turning 21—a small and statistically significant effect on course-taking behavior. However, as we show in the final column of Table 4, there is no evidence that students take more or fewer credits upon gaining legal access to alcohol.

4.2.2. Treatment-effect dynamics

In order to consider the dynamic effect of being able to drink legally, we return to our preferred outcome variable, students' normalized grades, and our preferred specification but replace the post-21 indicator variable with a set of indicator variables corresponding to the number of terms following the term in which a student gains legal access to alcohol. In particular, we include separate indicator variables for the term in which the individual turns 21, one term after a student turns 21, . . ., five terms after a student turns 21, and six-or-more terms after a student turns 21. The omitted category, essential for identifying individual fixed effects and trends, is being in a term prior to turning 21.[18]

Although it is possible to include indicator variables for terms prior to turning 21 to verify that grades do not fall below their expected levels in anticipation of gaining legal access—which we do in a series of falsification tests—our preferred estimates of the dynamic effects do not take this approach. We make this choice out of consideration for the general tradeoff involved with including pre-treatment indicator variables when using panel data approaches to estimation. Specifically, as one includes more indicator variables for pre-treatment periods, the counterfactual for the post-treatment periods becomes worse and worse as fewer observations contribute to the estimate of the individual fixed effects. For example, if we were to include indicators for one, two, three, and four terms prior to turning 21, our model would be projecting a student's future performance using observations from when he was under the age of 20. As such, our estimates of interest corresponding to post-21 terms would be noisier and less reliable than estimates that do not include these indicator variables and instead use all pre-21 terms to form counterfactuals.

Our preferred estimates of the treatment effect dynamics, shown in Column 1 of Table 5, indicate that grades fall significantly below their expected levels—by 0.036 standard deviations—in the term a student turns 21. This suggests an immediate negative effect

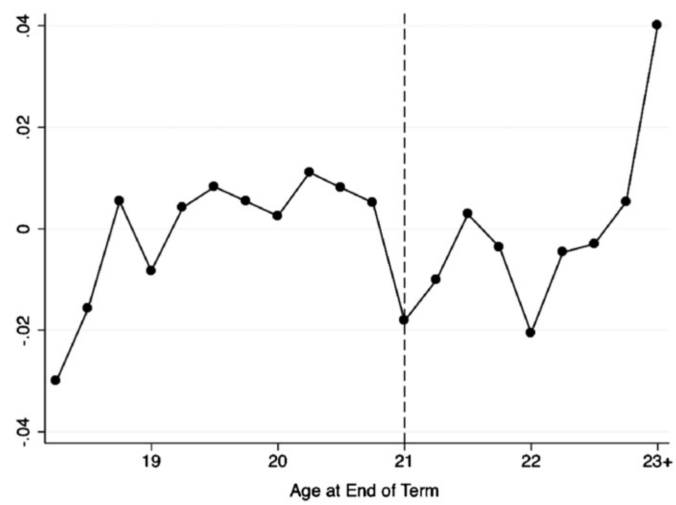

Fig. 1. Normalized GPAs by age adjusted for individual, accumulated credits, and course-type fixed effects. *Notes*: This figure plots average residuals from a regression of students' normalized GPAs on individual fixed effects, fixed effects for a student's accumulated credits at the beginning of a term, subject-by-level fixed effects, and term fixed effects.

of legal access to alcohol on academic performance. Further, the estimated coefficients corresponding to subsequent terms are usually significant and of similar magnitude, which indicates that the effect persists. We do note, however, that the coefficient on having turned 21 four terms ago (−0.055) is somewhat higher than the rest, which may reflect a 22nd-birthday effect.[19]

In Fig. 1, we present a graphical analogue to this analysis. In particular, we plot average adjusted normalized GPAs by students' ages in quarters. The normalized GPAs have been adjusted by taking the residuals from a regression on individual fixed effects, accumulated-credits fixed effects, and the course-specific fixed effects described above. Like the estimates in Column 1 of Table 5, this figure shows clearly that student GPAs fall below their expected levels when students turn 21 and, further, they stay below their below their expected levels for several subsequent quarters.

In Columns 2 through 5 of Table 5, we subject our estimation strategy to a series of specification tests. In particular, we add to our

[18] Note that although summer terms do not contribute to out analysis, such terms are considered in defining the term-based proximity to the term in which a student turns 21. As such, when the "turned 21 four terms ago" indicator variable is equal to one we are considering an individual in the term he turns 22.

[19] We have also examined whether the timing of a student's 21st birthday during the quarter is related to its impact on grades. These results, discussed in detail in Lindo et al. (2011), demonstrate that there are effects of gaining legal access to alcohol at any time during a given quarter. Further, we cannot reject that the effect is the same for students gaining legal access to alcohol at different times during the quarter.

Table 5
Fixed-effects-based estimates of the dynamic effects of legal access to alcohol on grades.

	(1)	(2)	(3)	(4)	(5)
Turns 21 in 4 terms					0.002
					(0.006)
Turns 21 in 3 terms				0.008	0.009
				(0.006)	(0.008)
Turns 21 in 2 terms			0.004	0.008	0.009
			(0.006)	(0.007)	(0.009)
Turns 21 in 1 term		0.002	0.003	0.008	0.009
		(0.007)	(0.007)	(0.009)	(0.010)
Term of 21st birthday	−0.036***	−0.035***	−0.033***	−0.028***	−0.027**
	(0.006)	(0.008)	(0.009)	(0.010)	(0.012)
Turned 21 1 term ago	−0.030***	−0.029***	−0.027**	−0.021*	−0.020
	(0.008)	(0.009)	(0.011)	(0.012)	(0.014)
Turned 21 2 terms ago	−0.026***	−0.025**	−0.023*	−0.017	−0.015
	(0.009)	(0.011)	(0.012)	(0.014)	(0.016)
Turned 21 3 terms ago	−0.031***	−0.030**	−0.027**	−0.021	−0.020
	(0.011)	(0.013)	(0.014)	(0.015)	(0.017)
Turned 21 4 terms ago	−0.055***	−0.054***	−0.051***	−0.044***	−0.043**
	(0.012)	(0.014)	(0.015)	(0.017)	(0.019)
Turned 21 5 terms ago	−0.038***	−0.037**	−0.035**	−0.027	−0.026
	(0.014)	(0.016)	(0.017)	(0.019)	(0.021)
Turned 21 6+terms ago	−0.021	−0.020	−0.017	−0.010	−0.008
	(0.017)	(0.019)	(0.020)	(0.022)	(0.023)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Accumulated-credits fixed effects	Yes	Yes	Yes	Yes	Yes
Course-specific controls	Yes	Yes	Yes	Yes	Yes
Number of students	13,102	13,102	13,102	13,102	13,102
Observations	479,342	479,342	479,342	479,342	479,342

Notes: The dependent variable is equal to the student's normalized course grade. Accumulated-credits fixed effects are fixed effects for a student's accumulated credits at the beginning of a term. Course-specific controls include subject-by-level fixed effects and term fixed effects. Standard errors (in parentheses) are corrected for clustering at the individual level.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

model indicator variables for terms preceding the term in which a student turns 21. Simply put, it would be a threat to the validity of the research design if similar effect evident in terms before a student turns 21. In order to maximize power, we take an incremental approach to adding indicator variables for terms preceding an individual's 21st birthday. In Column 2 we add an indicator for being one term prior to turning 21, in Column 3 we add an indicator for being two terms prior to turning 21, in Column 4 we add an indicator for being three terms prior to turning 21, and in Column 5 we add an indicator for being four terms prior to 21. Ultimately, we have ten "placebo tests" across these four columns where we do not anticipate any effects. Of these ten estimates, none are significant, which provides support for our preferred identification strategy. We also note that the estimates shown in Column 5 are what one would get if they were estimating the effect of turning 20 on student performance. Unlike the RD-approach above, where a 20th birthday effect is evident, we find no evidence that performance declines when a student turns 20 using our preferred approach.

4.2.3. Treatment-effect heterogeneity

In Table 6 we explore the extent to which there are heterogeneous effects of legal alcohol access on student achievement. Motivated by prior research documenting gender differences in educational performance and in tendencies to engage in risky behaviors, these tables report separate estimates for males and females. We also consider heterogeneity by ability and financial-aid eligibility to determine whether our main results are driven by individuals more likely to struggle with coursework or those from particular economic backgrounds.

In Panel A we stratify the sample by student gender and ability, with "high ability" students defined as those with SAT scores above the sample median of 1120 and "low ability" students defined as those at or below the sample median. The results in Columns 1 and 2 suggest that the effect of being able to drink legally is larger for females on average than it is for males. The point estimates remain small, however, with legal access reducing female grades by 0.045 and male grades by 0.024 standard deviations. In Columns 3 and 4, point estimates also suggest that the effect on low-ability students may be slightly greater than the effect on high-ability students.

In Columns 5 through 8, we separately consider the effects for low-ability males, high-ability males, low-ability females, and high-ability females. These estimates reveal substantial heterogeneity among males. Although there is a significant effect on low-ability males whose grades fall 0.047 standard deviations below their expected level after they gain legal access to alcohol, there appears to be no effect on high-ability males. On the other hand, our point estimates suggest that there are negative effects for both high- and low-ability females, although the estimated effects are greatest for low-ability females.

In Panel B we stratify the estimates by financial-aid eligibility and gender for the seventy percent of students who submitted a Free Application for Federal Student Aid (FAFSA). In so doing, we define a student as "high-eligibility" if Pell Grant eligibility is above the sample median and "low-eligibility" if Pell Grant eligibility is below the sample median. In Column 1, we show that the estimated effect for this sample of students (−0.042) is somewhat larger than the estimated effect based on the full sample (−0.033). However, the set of estimates suggests that, among males, the effect is concentrated among those who are likely to be from disadvantaged backgrounds. In contrast, the estimated effect is similar across differing levels of financial-aid eligibility among females.

Table 6
Heterogeneity across gender, ability, and financial-aid eligibility.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: gender and ability	Male	Female	High ability	Low ability	Male high ability	Male low ability	Female high ability	Female low ability
Age >21 during term	−0.024**	−0.045***	−0.021**	−0.046***	−0.006	−0.047***	−0.039***	−0.051***
	(0.010)	(0.008)	(0.009)	(0.009)	(0.013)	(0.014)	(0.012)	(0.011)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Accumulated-credits fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Course-specific controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of students	5903	7199	6332	6770	3221	2682	3111	4088
Observations	218,479	260,863	234,099	245,243	119,946	98,533	114,153	146,710
Panel B: gender and financial aid	Male	Female	High eligibility	Low eligibility	Male high eligibility	Male low eligibility	Female high eligibility	Female low eligibility
Age >21 during term	−0.026**	−0.058***	−0.051***	−0.040***	−0.045***	−0.015	−0.057***	−0.063***
	(0.011)	(0.009)	(0.010)	(0.010)	(0.017)	(0.015)	(0.013)	(0.013)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Accumulated-credits fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Course-specific controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of students	3900	5213	4556	4557	1887	2013	2669	2544
Observations	145,471	190,444	166,504	169,411	69,764	75,707	96,740	93,704

Notes: The dependent variable is equal to the student's normalized course grade. Accumulated-credits fixed effects are fixed effects for a student's cumulative credits at the beginning of a term. Course-specific controls include subject-by-level fixed effects and term fixed effects. Standard errors (in parentheses) are corrected for clustering at the individual level. The high-ability group consists of students with SAT scores above the sample median (1120) while the low-ability group consists of those with SAT scores at or below the sample median. *Significant at 10%.

** Significant at 5%.

*** Significant at 1%.

4.3. Fixed-effects estimates using the NLSY97

The statistics shown in Section 2 support the notion that the University of Oregon provides a fairly representative college setting. However, we lack data on drinking behavior at the University of Oregon. Thus, in this section we use data from the National Longitudinal Survey of Youth (NLSY97) to analyze how college students' drinking-related behaviors change upon gaining legal access to alcohol. While several prior studies have analyzed the effects of legal alcohol access on drinking-related behaviors (see Carpenter and Dobkin, 2011 for a review), to our knowledge we are the first to focus on the effects among college students.

Again utilizing the longitudinal nature of the data, we estimate the effect of legal alcohol access (being at least 21 years old at the time of the interview) on drinking-related behaviors during the past 30 days with an individual fixed effects model that controls for a quadratic in age. As such, the estimated effect is identified based on the discrete changes in behaviors that occur after an individual turns 21, adjusted for the gradual changes that are expected as individuals grow older. All regressions use NLSY97 sampling weights and cluster the standard-error estimates on the individual.

In Column 1 of Table 7 we summarize the results of our analysis for the full sample. Across the Panels A through D, we report the estimated effects on the probability that an individual has drank alcohol during the past 30 days, the number of days in which an individual has drank alcohol in the past 30 days, the number of alcoholic drinks an individual has had on average when drinking during the past 30 days, and the number of days an individual has drank five or more drinks during the past 30 days.[20] These results demonstrate that legal access causes college students to drink more often, as it increases the probability that an individual reports drinking by 6.5 percentage points, the number days drinking by 1.4, and the number of days drinking five or more drinks by 0.4. However, there is no evidence that legal access causes college students to drink more intensely on occasions during which they drink. This may reflect the likelihood that college students are more likely to be drinking in public after they turn 21 where social pressures (and bartenders) may limit the amount that an individual consumes.[21]

In Columns 2 through 9 we consider heterogeneity across gender, ability, and their interaction. High and low ability groups are defined based on the same SAT cutoff used to define the two groups in our analysis of University of Oregon transcript data. However, because SAT scores are missing for approximately two-thirds of the sample where Armed Services Vocational Aptitude Battery (ASVAB) test scores are not, we use ASVAB scores to impute predicted SAT scores where actual SAT scores are not available.[22]

These estimates illustrate that the link between legal access and drinking behaviors is complex. In particular, even putting imprecision aside, we usually cannot say that the effects are greater for one group than another. For example, the point estimates suggest that access has a bigger effect on the probability of drinking for females but has a bigger effect on the number of days drinking for males.

[20] We impute 15 drinks per day for the 0.9 percent of individuals who report that they drink more than 15 drinks per drinking occasion.

[21] These results are consistent with Carpenter and Dobkin (2009) who report that legal access leads to a 21 percent increase in drinking days (versus our estimate of is 33 percent), an insignificant 20-percent increase in heavy drinking days (versus our significant estimate of 19 percent), and find little evidence of any effect on the number of drinks consumed when drinking.

[22] Specifically, predicted SAT scores are based on the conditional expectation function implied by a regression of SAT scores on a quartic in ASVAB scores for the sample where both scores are available.

Table 7
Estimated effects on drinking behaviors during the previous 30 days using NLSY97 data.

	(1) Full sample	(2) Male	(3) Female	(4) High ability	(5) Low ability	(6) Male high ability	(7) Male low ability	(8) Female high ability	(9) Female low ability
Panel A: drank									
Age > 21	0.065***	0.051**	0.077***	0.088***	0.054***	0.084**	0.026	0.091**	0.072***
	(0.014)	(0.021)	(0.017)	(0.024)	(0.017)	(0.033)	(0.031)	(0.036)	(0.021)
Students	2298	997	1301	693	1419	348	522	313	818
Observations	9023	3943	5080	2789	5521	1422	2029	1253	3203
Pre-21 mean	0.66	0.67	0.65	0.66	0.66	0.68	0.66	0.66	0.65
Panel B: days drank									
Age > 21	1.394***	1.643***	1.173***	1.727***	1.265***	1.597***	1.695***	2.027***	0.976***
	(0.184)	(0.300)	(0.227)	(0.293)	(0.255)	(0.443)	(0.471)	(0.401)	(0.304)
Students	2298	997	1301	693	1419	348	522	313	818
Observations	9023	3943	5080	2789	5521	1422	2029	1253	3203
Pre-21 mean	4.27	4.80	3.83	3.95	4.42	4.42	5.04	3.51	3.97
Panel C: drinks when drank									
Age > 21	−0.190*	−0.234	−0.155	−0.189	−0.147	−0.384	−0.026	0.034	−0.243
	(0.111)	(0.178)	(0.139)	(0.169)	(0.153)	(0.255)	(0.256)	(0.227)	(0.201)
Students	1905	841	1064	607	1147	308	427	271	654
Observations	6163	2745	3418	2023	3662	1045	1360	908	2100
Pre-21 mean	4.54	5.35	3.86	4.31	4.67	4.91	5.60	3.61	4.03
Panel D: days drank 5+drinks									
Age > 21	0.409***	0.520***	0.302**	0.377**	0.440***	0.350	0.403	0.478*	0.319
	(0.118)	(0.186)	(0.151)	(0.187)	(0.165)	(0.278)	(0.274)	(0.268)	(0.211)
Students	2298	997	1301	693	1419	348	522	313	818
Observations	9023	3943	5080	2789	5521	1422	2029	1253	3203
Pre-21 mean	2.11	2.71	1.61	1.97	2.19	2.39	2.99	1.54	1.66

The sample is restricted to individuals currently enrolled in a four-year college between 1998 and 2009. All regressions are weighted using sampling weights and control for individual fixed effects and a quadratic in age. Standard errors (in parentheses) are corrected for clustering at the individual level. Consistent with the definitions used in our analysis of University of Oregon transcript data, the high-ability group consists of students with actual (or ASVAB-predicted) SAT scores above 1120 while the low-ability group consists of those with scores at or below the 1120.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

This mixed pattern also appears when making comparisons across gender within ability groups. Comparing the effects across ability groups, it appears as if the effect of legal access is greater for high-ability females than low-ability females. We do not, however, see the same pattern among males.

Given this mixed set of results, we are inclined to view these estimates as evidence of "proof of concept." Specifically, that we find effects of legal alcohol access on grades among students at the University of Oregon is only compelling if legal access affects the behaviors of college students, which we demonstrate using the NLSY97. However, we are not inclined to extend our interpretation of the results to estimate the effect any particular measure of "drinking" on college performance given the variety of ways in which one can measure drinking. As we alluded to in the introduction, legal access may affect achievement through its impact on whether a student drinks, the frequency with which a student drinks, and the intensity with which a student drinks, in addition to the wide array of social activities that are associated with drinking. Moreover, there are generally not any specific groups in the NLSY97 that we can point to as being particularly responsive to legal access in their drinking activity, except perhaps high-ability females.

That said, we do find it interesting that we find no significant effect on the grades of high-ability males at the University of Oregon despite finding evidence in the NLSY97 that they drink more often with legality. This may suggest that high-ability students are particularly adept at changing their drinking-related behaviors without compromising their grades but, of course, we cannot rule out the possibility that high-ability males at the University of Oregon do not alter their drinking behaviors upon gaining legal access. That said, the results for females also support the notion that high-ability students are better able to maintain their grades when they change their drinking activity. In particular, the NLSY97 data suggest that legal access has a greater impact on the drinking behaviors of high-ability females than on low-ability females; yet transcript data from the University of Oregon suggests that legal access has a greater impact on the grades of low-ability females (though the difference is not statistically significant).

5. Conclusion

As a whole, our analysis suggests that legal access to alcohol does affect the drinking behavior of college students and, in turn, affects student performance. At the University of Oregon, legal access to alcohol reduces grades by 0.03 standard deviations, or the equivalent of causing a student to perform as if his or her SAT score were 20 points lower. As such, the effect we identify is smaller than Carrell et al. (2011) who find that gaining legal access at the end of the academic term reduces grades by approximately 0.10 standard deviations. Given the more conventional enforcement of MLDA at large public universities, this difference might exist because legal access has a different effect on alcohol-related behavior across the two settings. We also find substantial heterogeneity across gender and ability, in ways that diverge meaningfully from the prior research. In particular, given that the U.S. Air Force Academy is more selective and has a much larger fraction of men than the University of Oregon, it is perhaps surprising that we find no evidence of an effect among high-ability males. In addition, we identify a significant effect on the performance of females.

Acknowledgements

We are grateful to Scott Carrell, Jeff DeSimone, Ben Hansen, and Mark Hoekstra for thoughtful comments and suggestions.

References

Angrist, J.D., Lavy, V., 1999. Using Maimonides' rule to estimate the effect of class size on scholastic achievement. Quarterly Journal of Economics 114 (2), 533–575.

Barreca, A., Guldi, M., Lindo, J.M., Waddell, G.R., 2011a. Saving babies? Revisiting the effect of very low birth weight classification. Quarterly Journal of Economics 126 (4), 2117–2123.

Barreca, A., Lindo, J.M., Waddell, G.R., 2011. Heaping-induced bias in regression-discontinuity designs. NBER Working Paper No. 17408.

Carpenter, C., 2004. Heavy alcohol use and youth suicide: evidence from tougher drunk driving laws. Journal of Policy Analysis and Management 23 (4), 831–842.

Carpenter, C., 2005a. Heavy alcohol use and the commission of nuisance crime: evidence from underage drunk driving laws. American Economic Review 95 (2), 267–272.

Carpenter, C., 2005b. Youth alcohol use and risky sexual behavior: evidence from underage drunk driving laws. Journal of Health Economics 24 (3), 613–628.

Carpenter, C., 2007. Heavy alcohol use and crime: evidence from underage drunk-driving laws. Journal of Law and Economics 50, 539–781.

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), 164–182.

Carpenter, C., Dobkin, C., 2010. The drinking age, alcohol consumption, and crime. Mimeo.

Carpenter, C., Dobkin, C., 2011. The minimum legal drinking age and public health. Journal of Economic Perspectives 25 (2), 133–156.

Carrell, S.E., Hoekstra, M., West, J.E., 2011. Does drinking impair college performance? Evidence from a regression discontinuity approach. Journal of Public Economics 95 (1–2), 54–62.

Chatterji, P., DeSimone, J.S., 2006. High school alcohol use and young adult labor market outcomes. NBER Working Paper.

Chesson, H., Harrison, P., Kassler, W.J., 2000. Sex under the influence: the effect of alcohol policy on sexually transmitted disease rates in the United States. Journal of Law and Economics 43 (1), 215–238.

Cook, P.J., Moore, M.J., 1993. Drinking and schooling. Journal of Health Economics 12 (4), 411.

Dave, D., Kaestner, R., 2002. Alcohol taxes and labor market outcomes. Journal of Health Economics 21 (3), 357–371.

Dee, T.S., 1999. State alcohol policies, teen drinking and traffic fatalities. Journal of Public Economics 72 (2), 289–315.

Dee, T.S., Evans, W.N., 2003. Teen drinking and educational attainment: evidence from two-sample instrumental variables estimates. Journal of Labor Economics 21 (1).

DeSimone, J., 2007. Fraternity membership and binge drinking. Journal of Health Economics 26 (5), 950–967.

DiNardo, J., Lee, D.S., 2004. Economic impacts of new unionization on private sector employers: 1984–2001. Quarterly Journal of Economics 119 (4), 1383–1441.

Kremer, M., Levy, D., 2008. Peer effects and alcohol use among college students. The Journal of Economic Perspectives 22 (3), 189–3A.

Lindo, J., Swensen, I., Waddell, G., 2011. Alcohol and student performance: estimating the effect of legal access. NBER Working Paper No. 17637.

Lindo, J., Swensen, I., Waddell, G., 2012. Are big-time sports a threat to student achievement? American Economic Journal: Applied Economics 4 (4).

MacDonald, Z., Shields, M.A., 2004. Does problem drinking affect employment? Evidence from England. Health Economics 13 (2), 139–155.

Markowitz, S., Grossman, M., 1998. Alcohol regulation and domestic violence towards children. Contemporary Economic Policy 16 (3), 309–320.

Mullahy, J., Sindelar, J., 1996. Employment, unemployment, and problem drinking. Journal of Health Economics 15 (4), 409–434.

Powell, L.M., Williams, J., Wechsler, H., 2004. Study habits and the level of alcohol use among college students. Education Economics 12 (2), 135–149.

Rashad, I., Kaestner, R., 2004. Teenage sex, drugs and alcohol use: problems identifying the cause of risky behaviors. Journal of Health Economics 23 (3), 493–503.

Renna, F., 2008. Alcohol abuse, alcoholism, and labor market outcomes: looking for the missing link. Industrial and Labor Relations Review 62 (1), 92–103.

Rees, D.I., Argys, L.M., Averett, S.L., 2001. New evidence on the relationship between substance use and adolescent sexual behavior. Journal of Health Economics 20 (5), 835–845.

Sen, B., 2002. Does alcohol-use increase the risk of sexual intercourse among adolescents? Evidence from the NLSY97. Journal of Health Economics 21 (6), 1085–1093.

Terza, J.V., 2002. Alcohol abuse and employment: a second look. Journal of Applied Econometrics 17 (4), 393–404.

Waddell, G.R., 2012. Gender and the influence of peer alcohol consumption on adolescent sexual activity. Economic Inquiry 50 (1), 248–263.

Williams, J., Powell, L.M., Wechsler, H., 2003. Does alcohol consumption reduce human capital accumulation? Evidence from the college alcohol study. Applied Economics 35 (10), 1227–1239.