



Does drinking impair college performance? Evidence from a regression discontinuity approach ^{☆,☆☆}

Scott E. Carrell ^{a,*}, Mark Hoekstra ^b, James E. West ^c

^a UC-Davis and NBER, United States

^b University of Pittsburgh, 4714 WW Posvar Hall, Pittsburgh, PA 15260, United States

^c Department of Economics and Geosciences, United States Air Force Academy, CO 80840, United States

ARTICLE INFO

Article history:

Received 2 March 2010

Received in revised form 11 August 2010

Accepted 24 August 2010

Available online 21 September 2010

JEL classification:

I18 (Public Health)

I21 (Analysis of Education)

Keywords:

Public health

Education

ABSTRACT

This paper examines the effect of alcohol consumption on student achievement. To do so, we exploit the discontinuity in drinking at age 21 at a college in which the minimum legal drinking age is strictly enforced. We find that drinking causes significant reductions in academic performance, particularly for the highest-performing students. This suggests that the negative consequences of alcohol consumption extend beyond the narrow segment of the population at risk of more severe, low-frequency, outcomes.

© 2010 Elsevier B.V. All rights reserved.

Economists and other social science researchers have devoted significant effort toward understanding the effects of alcohol consumption. However, nearly all of this research has focused on low-frequency outcomes such as mortality and teen pregnancy (Carpenter and Dobkin, 2009; Dee, 1999, 2001; Fertig and Watson, 2009; Saffer, 1997). In contrast, little is known about the effect of drinking on the majority of users. In addition, identifying the causal effect of alcohol consumption is difficult since individuals would likely select into drinking based on unobserved characteristics that affect outcomes. Moreover, efforts to exploit state-level variation such as alcohol taxes or adoption of minimum legal drinking age laws have also been problematic. That is, there are concerns that the timing of the drinking age laws may be endogenous (Miron and Tetelbaum, 2007), while efforts to exploit tax differences across states have been limited by the relative weakness of taxes as an instrument (e.g., Chaloupka and Wechsler, 1996) and the corresponding imprecision of labor market outcome estimates (Dave and Kaestner, 2002).

In this paper, we examine the impact of drinking on academic performance in college. In doing so, this study contributes to the existing literature and the corresponding policy debate in several ways. First, examining the impact of drinking in a college context allows for us to gauge the impact of drinking on a wide range of students, rather than focusing on low-frequency outcomes that affect relatively few individuals. In addition, understanding the impact of drinking during college is itself particularly important for public policy, since college is a critical time in the human capital formation process that can significantly impact future labor market outcomes. Indeed, concern regarding the harmful effects of drinking in college is reflected by the Amethyst Initiative (2010), in which 135 university presidents and chancellors argue that current policy has resulted in “binge drinking” and thus ought to be reexamined.

Another distinct advantage of our approach is that by studying the impact of drinking in a college environment in which underage drinking is strictly prohibited, we can distinguish the effect of drinking from confounding factors. To do so, we utilize a regression discontinuity design in comparing the grades of students who turned 21 just before the final exams to those who turned 21 just afterward.¹ Under the plausible identifying assumption that other determinants of performance are smooth through age 21, any discontinuous changes in performance are properly interpreted as the causal impact of drinking on performance.

¹ This research design is similar to that used by Carpenter and Dobkin (2009) to estimate the effect of drinking on mortality.

[☆] We would like to thank Doug Miller and Hillary Hoynes for helpful comments and suggestions. Any errors are our own.

^{☆☆} The views expressed in this paper reflect those of the authors and do not necessarily reflect the official policy or position of the U.S. Air Force, Department of Defense, or the U.S. Government.

* Corresponding author. University of California-Davis, One Shields Ave., Davis, CA 95616, United States. Tel.: +1 530 752 5480; fax: 530 752 9382.

E-mail addresses: secarrell@ucdavis.edu (S.E. Carrell), markhoek@pitt.edu (M. Hoekstra), jim.west@usafa.edu (J.E. West).

We use administrative data on 3884 students at the United States Air Force Academy (USAFA) between 2000 and 2006. This educational setting offers two distinct advantages for our analysis. First, in contrast to many college campuses, the explicit ban on underage drinking at the USAFA is strictly enforced; violations can lead to expulsion.² As a result, in an anonymous survey of underage students in our sample, only 37 to 39% of students reported drinking any alcohol since arriving at the academy.³ By comparison, the National Institute on Alcohol Abuse and Alcoholism (NIAAA) reports that 83% of college students nationwide drink, and that 41% of college students reported consuming five or more drinks on a single occasion within the past two weeks (NIAAA, 2010).⁴ The second advantage of using USAFA data is a consistent measurement of academic achievement. Students at USAFA are required to take a core set of approximately 30 courses in mathematics, basic sciences, social sciences, humanities, and engineering. For these core courses, students have no discretion in choosing their professors. Furthermore, exams are standardized across different sections of the same course. These attributes lessen the scope for professor subjectivity and student selection into courses that might correlate with date of birth.

Results from our analysis indicate that alcohol consumption significantly reduces academic performance by nearly one-tenth of a standard deviation, or approximately as much as being assigned a professor whose quality is one standard deviation below average (Carrell and West, 2010). Strikingly, the negative effects are largest for the highest-performing students at the academy. This suggests that the negative consequences of alcohol consumption are not limited to a small fraction of users or even to those who might naturally struggle with academics.

1. Identification strategy

In order to estimate the causal impact of drinking on student performance, we apply a regression discontinuity design⁵ to exploit the sudden increase in drinking that occurs at age 21. This approach is similar to that used to examine the impact of Head Start (Ludwig and Miller, 2007) and, more recently, the effect of alcohol consumption on mortality (Carpenter and Dobkin, 2009). The identifying assumption of the design is that while other determinants of performance such as motivation or maturity vary smoothly over age 21, access to and consumption of alcohol varies discontinuously at that point. This smoothness assumption appears to be particularly reasonable when treatment depends on age because in contrast to other outcomes such as test scores, there is little scope for individuals to change their 21st birthday by a few days or weeks. Moreover, current evidence indicates that there is a discontinuous increase in alcohol consumption at age 21; Carpenter and Dobkin (2009) report a 21% increase in recent drinking days at age 21. Similarly, although we cannot directly test how much more students in our data drink once they turn 21 than those just about to turn 21, we can report that there is significant variation on the extensive margin: In a 2004 USAFA survey of students, nearly 90% of students 21 years or older reported drinking in the previous year. In contrast, only 37 to 39% of students under the age of 21 report having at least one drink since coming to the academy,

² Per the AFCWMAN36-3501, Section 1.6.1, two incidents of underage drinking result in the expulsion of the student. Additionally, some single incidents such as driving under the influence (DUI) result in expulsion.

³ This includes drinking over the summer and winter breaks, when students are at home and away from the watchful eyes of their superiors.

⁴ Additionally, in a self-reported sample of 10,424 college freshmen across 14 schools taken in the fall semester of 2003, roughly 55% reported drinking alcohol during a two-week time period. Among those who drank, there were on average 4.35 (sd = 2.99) drinking occasions for male students and 3.35 (sd = 2.32) drinking occasions for female students within that two-week period (White et al., 2006).

⁵ Imbens and Lemieux (2008) offer a thorough description of regression discontinuity design.

some of which presumably occurred off-campus during summer or winter breaks.

To formally test for discontinuities in academic performance at age 21, we estimate the following equation:

$$\text{CourseScore}_{ic} = \beta_0 + \beta_1 \text{Age21}_i + \beta_2 f(\text{Age})_i + \varepsilon_{ic}$$

where, CourseScore is the normalized final score earned in course *c* by student *i*, Age21 is an indicator variable equal to one if the student is at least 21 years old at the time of the exam, and *f*(Age) is a flexible polynomial of one's age at the time of the exam which we allow to take different forms on either side of the cutoff. Under the identifying assumption that other determinants of achievement are continuous at the age 21 cutoff, β_1 will be an unbiased estimate of the effect of the increase in drinking at age 21 on academic performance. Standard errors are clustered at the birthday level (month/day/year), which accounts for the fact that students are observed with scores in multiple courses.

In addition to the equation above, we also estimate specifications in which we additionally control for course by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT verbal, SAT math, academic composite, leadership composite, fitness score, and indicator variables for Black, Hispanic, Asian, recruited athlete, and preparatory school attendance. These additional controls account for any course, teacher, classroom, or cohort-level common shocks to student performance.

While the primary advantage of applying a regression discontinuity design to USAFA data is that it allows us to estimate cleanly the impact of drinking on performance, the tradeoff is that the resulting estimates may under- or over-state the impact of drinking in other contexts. For example, our estimates will reflect the effect of the type of drinking that occurs on and soon after one's 21st birthday, which may be closer to binge drinking than other types of drinking.⁶ To the extent that individuals may moderate drinking over time, our estimates may over-state the long-term effects of drinking. On the other hand, while our outcome of interest is *course* performance, the regression discontinuity estimates are primarily driven by changes to the final exam score. Given that the final exams count for as little as 35% of the overall course score, our estimates may underestimate the effect of this type of drinking by as much as a factor of three. Similarly, our regression discontinuity estimates capture the effect of drinking for a relatively short period of time prior to the final exams, which means that drinking could not have affected cognitive learning throughout the semester. Allowing for longer-term access to alcohol may well lead to larger negative effects. In addition, since there is limited scope for drinking to have affected classroom learning throughout the semester, our view is that the estimates here are best interpreted as the effect of drinking on performance, rather than the effect of drinking on learning.

2. Data

Our data include student-level observations for each course taken during the 2000–01 through the 2005–06 academic years. From this sample, we exclude physical education courses and observations on freshmen students and students who had attended military preparatory school.⁷ In addition, we exclude observations in which the final exam was taken more than 270 days before or after one's 21st birthday. This leaves a sample of 58,032 observations on 3884 students. Each observation measures the final score achieved by each student in each course.

⁶ We note, however, that Carpenter and Dobkin (2009) report a sustained increase in drinking at age 21 rather than a temporary spike.

⁷ We exclude freshmen because only a small fraction of students turn 21 during their freshman year and freshman at USAFA are prohibited from drinking alcohol, regardless of age. In addition, students who attended military preparatory school can be considerably older than average college students and are less representative of college students more generally. We note, however that our results are not sensitive to the exclusion of these students, as shown in columns 3 and 4 of Table 6.

Table 1
Descriptive statistics.

Variable	Mean (sd)	Range
Course score	82.32 (9.03)	0–117.73
Age (days from 21st birthday during finals week)	–8.25 (149.61)	–270–269
Black	0.03 (0.17)	0–1
Hispanic	0.04 (0.19)	0–1
Asian	0.05 (0.21)	0–1
Female	0.18 (0.38)	0–1
SAT verbal score	639.60 (60.80)	340–800
SAT math score	672.70 (57.86)	420–800
High school academic composite score	12.81 (2.04)	5.35–17.41
High school leadership score	17.52 (1.87)	9–24
High school fitness score	4.75 (0.92)	2.3–8.0
Recruited athlete	0.27 (0.44)	0–1

Notes: Figures come from 58,032 observations on 3884 students. This excludes freshman and students who attended a military preparatory school.

Summary statistics from our data are shown in Table 1. The average score achieved in academic courses was 82%, which corresponds to a mean grade point average of 2.9.⁸

The average combined SAT score of students in our sample is approximately 1300, which is similar to other undergraduate institutions such as UCLA, University of Michigan, University of Virginia, and UNC-Chapel Hill. Eighteen percent of the sample is female, 3% is Black, 4% is Hispanic, and 5% is Asian. Thus, the student body at USAFA is somewhat more White and male and less Black than four-year public universities nationwide.⁹

Students at the academy are also more likely to come from families in which their parents live with each other (84.1% versus 66.1%), and are more likely to have attended a religious service in the last year (90.2% versus 78.1%). Finally, and perhaps most importantly for this study, students at the USAFA report that they spent less time partying during their last year in high school than did students in general: 33.8% of students at the academy reported they spent no time partying (versus 25.3% for all public university students), while 15% reported spending at least 6 hours a week partying (versus 23.3%).

3. Results

3.1. The discontinuity in drinking at age 21

Although we cannot match student-level data on drinking to our sample, there is evidence of a large increase in drinking at age 21 at the USAFA. In the anonymous USAFA Climate Survey, relatively few students under the age of 21 report having a drink since coming to USAFA (2003: 39%, 2004: 38%, and 2005: 37%) while in 2004, approximately 90% of students 21 years or older reported having a drink in the previous year. Moreover, this difference likely understates the true increase in drinking at age 21 at the USAFA for several reasons. First, these figures only represent the increase in drinking on the extensive margin. Indeed, while we do not have data on the frequency or intensity of drinking at the academy, Carpenter and Dobkin (2009) report a 21% increase in the proportion of days drinking at age 21. In addition, it is likely that the majority of the underage drinking that does occur by students at the USAFA is done while off-campus during holiday breaks, when students are no longer under the watchful eye of their superiors. Finally, once students at the USAFA turn 21 years of age, social drinking is quite accessible because a campus bar is located adjacent to

⁸ These statistics represent the academic achievement of students in our sample within 180 days of their 21st birthday.

⁹ According to The Freshman Survey, administered to entering freshmen nationwide by the UCLA Cooperative Institutional Research Program (CIRP, 2007), 74.5% of students attending all public four-year colleges are White, 13.7% are Black, and 3.8% are Asian.

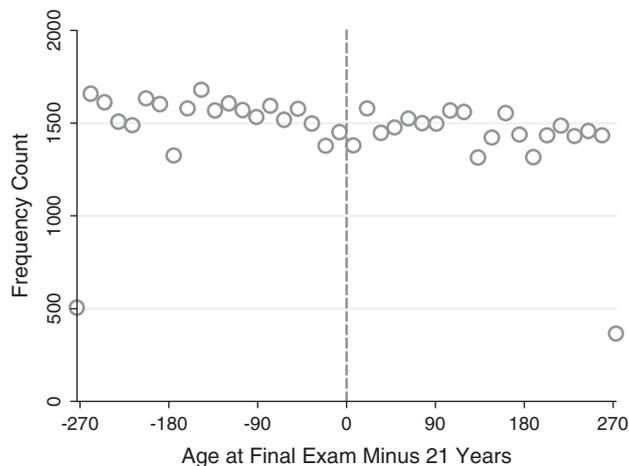


Fig. 1. Distribution of final exam observations near age 21.

the student dormitories.¹⁰ Thus, the evidence suggests that there is a significant increase in drinking at age 21 at the USAFA.

3.2. Tests of the validity of the RD design

The primary threat to identification in a RD design comes from the possibility of nonrandom sorting of students to either side of the cutoff. While the scope for such sorting seems limited, if not impossible, given both the inflexibility of age and the timing of the final exam schedule at the USAFA, we can nonetheless test the identifying assumption that other determinants of achievement are smooth through the age 21 threshold.

We first examine whether there is evidence of sorting graphically. Fig. 1 shows the distribution of observations around the age 21 cutoff, which shows the density is smooth across the minimum drinking age cutoff. Graphically there is little evidence of nonrandom sorting near age 21.¹¹

In addition, we also test whether a variety of predetermined characteristics vary smoothly across the age 21 cutoff, including math and verbal SAT scores, academic composite scores, leadership composite scores, and fitness scores. Results are shown in Fig. 2, in which the open circles represent local averages and the line is the predicted score. Corresponding regression discontinuity estimates are shown in Table 2. Both the visual evidence and regression estimates reveal no evidence of discontinuities in predetermined characteristics at age 21. Collectively, these results are supportive of the identifying assumption that while drinking is discontinuous at age 21, other determinants of performance are continuous.

3.3. The effect of drinking on academic performance

Fig. 3 provides a graphical representation of the effect of drinking on academic performance.¹² The results indicate that students who turn 21 just prior to taking the final exams score approximately one-tenth of a standard deviation lower than students who turn 21 just after finals.

¹⁰ Hap's Place Lounge is located in the campus student union, called Arnold Hall, adjacent to the student dormitories. Students at the USAFA are required to reside in the dormitories. A student-made video of Hap's Place can be found at: <http://www.youtube.com/watch?v=LHWNN9-c0ho>.

¹¹ To further test this result, we performed chi-squared goodness of fit tests by semester. For the first ten semesters of data, we fail to reject the null hypothesis of uniformity at a 5% level. In the final three semesters of data, the frequency of birthdates is skewed due to the absence of data from the rising class. The main results of our paper are robust to the exclusion of these final three semesters of data.

¹² Scores are not centered at zero because scores were normalized for the full sample, though as described earlier we exclude students who had attended military preparatory school prior to enrolling at the academy. Those students perform worse academically, on average, than other students.

This drop in performance is both statistically and economically significant; it is roughly the same effect as having a professor whose quality is one standard deviation below the mean (Carrell and West, 2010).

Corresponding regression discontinuity estimates are shown in Table 3, which includes observations on all courses in which the final exam was taken within 180 days of the student's 21st birthday. Columns 1 and 2 show the estimates from specifications without and with controls, respectively, while Column 3 shows the estimate when allowing for quadratic functions of age on either side of the age 21 cutoff. All three specifications show that drinking causes a statistically significant reduction in performance on the order of approximately one-tenth of a standard deviation.

To test the robustness of these results, we estimated specifications using different bandwidths, different polynomials of age, and with and without individual controls. Results are shown in Table 4, where Panel A shows results from specifications without controls and Panel B shows results with controls. Columns 1 through 6 show results for bandwidths of 9, 8, and 7 months, where odd-numbered columns include a linear polynomial of age and even-numbered columns include a quadratic. Columns 7 through 10 control for linear functions of age for bandwidths of 150, 120, 100, and 80 days. Columns 11 through 14 show results for bandwidths of 80, 60, 40, and 20 days. In those specifications no controls for age are included, under the assumption that among those sufficiently

close to their 21st birthday, whether one can drink is essentially random. Consistent with the visual evidence in Fig. 1 as well as the results of Table 2, estimates indicate that drinking causes a statistically significant performance drop of 0.06 to 0.10 standard deviations in the course.

An important question is whether our estimates represent temporary or permanent reductions in performance due to the increase in drinking that occurs after turning 21. While Fig. 3 appears to show a lasting impact of age 21-induced drinking, we note that definitively estimating the permanent effect of drinking is difficult for two reasons. First, it requires us to estimate how well students would have performed months after turning 21 had they not been able to drink. This is considerably more tenuous than making an assumption regarding performance immediately after turning 21, which is required for the regression discontinuity analysis. The second reason relates to the fact that we do not directly observe alcohol consumption in the months after turning 21. Thus, even if we were to observe a rebound in performance in the months after turning 21, it would be difficult to know whether it is due to a reduction in drinking intensity or due to the students' ability to adjust to a lifestyle of drinking.

However, with those caveats in mind, we can test whether the performance drop we observe at age 21 is temporary. To do so, we estimate average academic performance 0 to 2 months, 2 to 4 months, 4 to 6 months and 6 to 8 months after turning 21 relative to

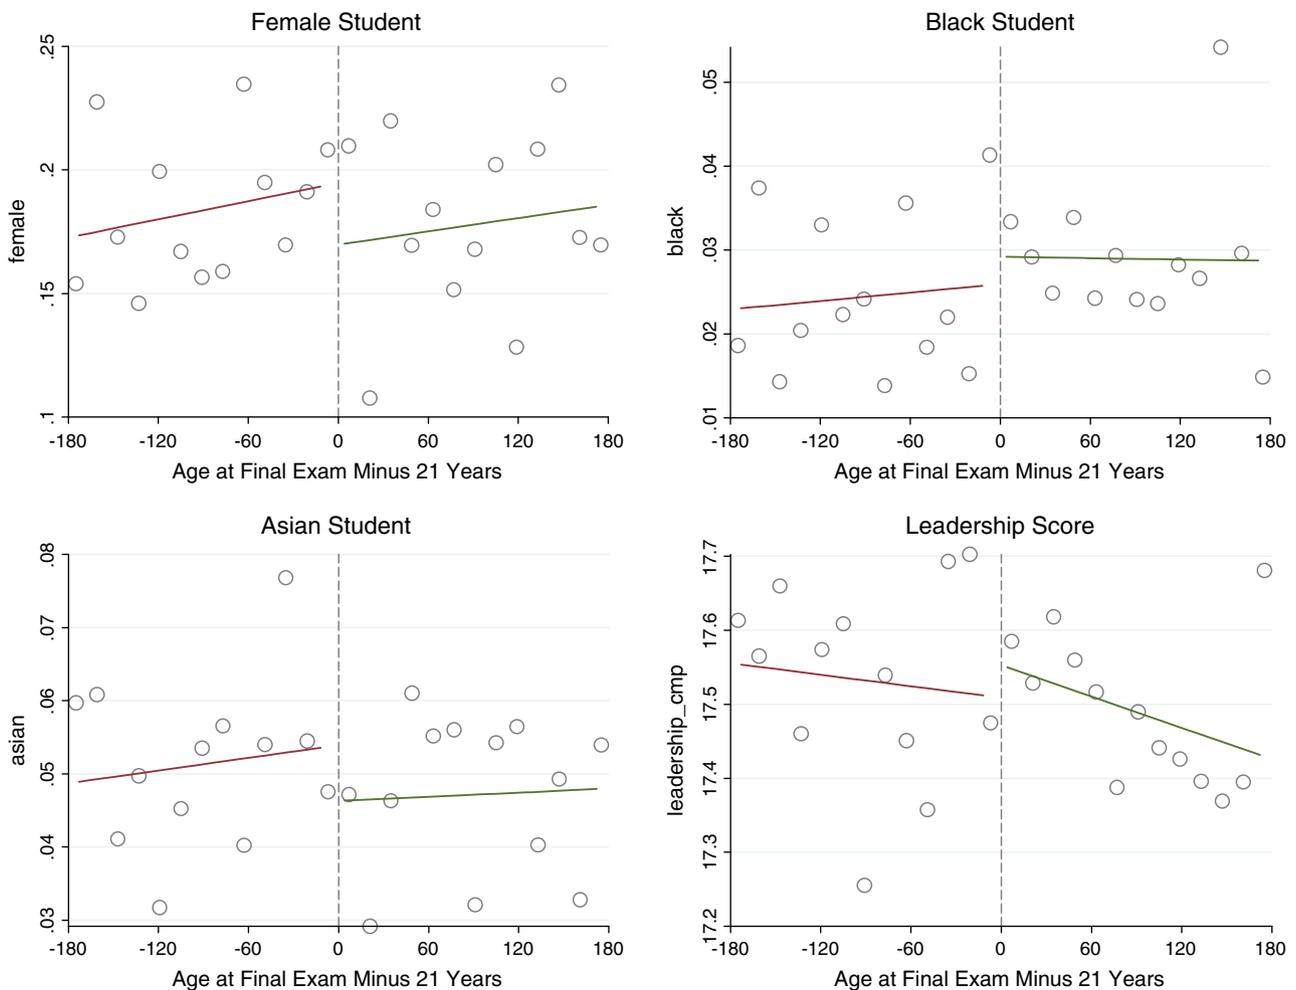


Fig. 2. Tests of the regression discontinuity design.

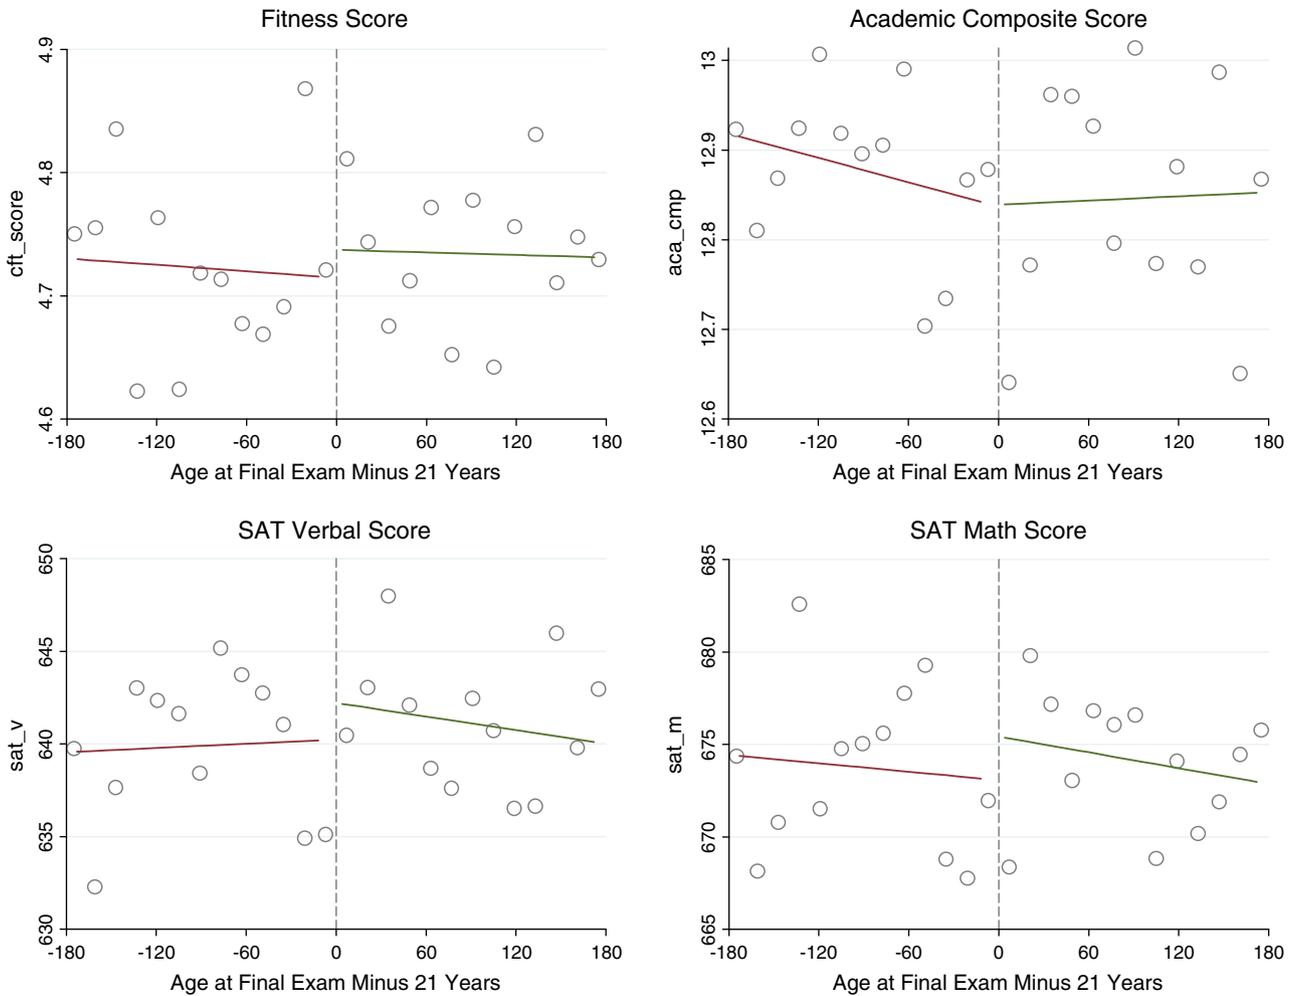


Fig. 2 (continued).

performance in the 2 months prior to turning 21.¹³ Consistent with Fig. 3, the results suggest a performance drop that persists through at least 8 months.¹⁴ While we cannot be certain whether the persistence we observe is due to continued alcohol consumption, “senioritis”, or some other effect, we do find a similar lasting reduction in performance when we limit the sample to only students who turn 21 during their junior year.

Next, we examine which students are most affected by drinking. Table 5 shows results when we split the sample into two groups by cumulative GPA at the end of the freshman year. Strikingly, we find that the largest negative effects of drinking are for the students who were in the top half of their class by GPA at the end of their freshman year. Point estimates in Column 2 indicate that these students experience a drop in course performance of 0.10 to 0.15 standard deviations, or more than the effect of being assigned a professor whose quality was one standard deviation below average. This result is shown graphically in Fig. 4, which shows the discontinuous drop in performance for students whose freshmen GPAs placed them in the

¹³ Thus, the model implicitly assumes that in the absence of age 21-induced drinking, performance would have remained at levels observed in the two months prior to turning 21. This appears to be a conservative assumption, since performance appears to be increasing, if anything, prior to turning 21.

¹⁴ Estimates are -0.101 , -0.097 , -0.050 , and -0.113 after 0 to 2, 2 to 4, 4 to 6, and 6 to 8 months, respectively. All estimates except for that between 4 and 6 months are statistically significant at the 1% level.

top half of their class. In contrast, we find little effect for students in the bottom half of their class.

Columns 4 and 5 of Table 5 show results for men and women, respectively. Estimates are negative for both groups, though the estimates for women are imprecisely estimated due to the limited number of women at the academy. Finally, Columns 5 and 6 show estimates separately for math and science courses versus social sciences and humanities. Results indicate that student performance drops similarly in both types of courses.

Table 2
Tests of the identifying assumption of the RD analysis.

Outcome	SAT math	SAT verbal	Academic composite	Leadership composite	Fitness score
	1	2	3	4	5
Discontinuity at age 21	2.371 (2.81)	1.932 (2.79)	0.009 (0.10)	0.045 (0.08)	0.025 (0.04)
Observations	38,782	38,782	38,782	38,782	38,782

Notes: Each cell represents results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for Age 21. Standard errors clustered by age are in parentheses. All specifications control for a linear function of distance from Age 21 in which the slope is allowed to vary on either side of the cutoff. The bandwidth of the data is 180 days on either side of Age 21.

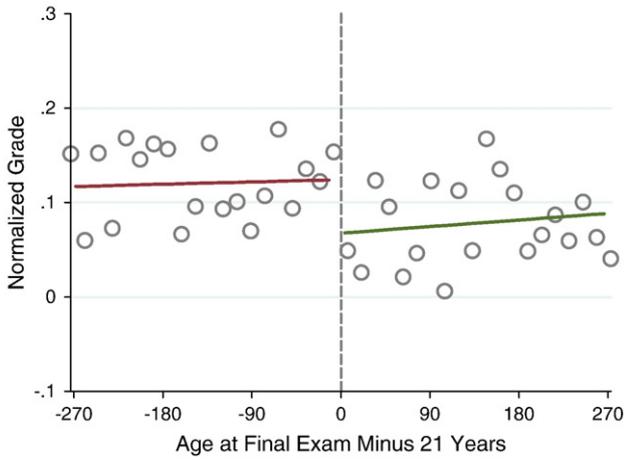


Fig. 3. Regression discontinuity estimates of the effect of drinking on achievement.

In summary, our results yield two notable findings. We find that there is a large and statistically significant discontinuous drop in college performance at age 21 that is robust across various bandwidths and functional forms. The drop in performance from the increase in drinking in the weeks prior to the final exams is economically significant, as it is approximately the same as the effect of being assigned to a professor whose quality is one standard deviation below the mean in quality for the entire semester. We also find that the largest negative effects of drinking are on the high-ability students.

3.4. Robustness tests

To test the robustness of our findings, in Table 6, Column 2, we show that our results are virtually unchanged when including freshman students. In Column 3, we show similar results when we include students who attended military preparatory schools prior to entering the USAFA. As a third robustness check, in Column 4 we restrict our observations to the required core courses taken by all students at the USAFA. These courses have the advantage of common examinations for all students taking the course in a given semester and eliminate any possible concerns of self-selection of courses during the semester in which a student turns 21 years of age. Again, our results remain virtually unchanged compared to our main specification.

Table 3
Regression discontinuity estimates of the effect of drinking on academic performance.

Specification	1	2	3
Discontinuity at age 21	-0.092*** (0.03)	-0.114*** (0.02)	-0.106*** (0.03)
Observations	38,782	38,782	38,782
Age polynomial	Linear	Linear	Quadratic
Control variables	No	Yes	Yes

Notes: Each cell contains results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for age 21. Standard errors clustered by age are in parentheses. All specifications control for a flexible polynomial of age in which the slope is allowed to vary on either side of the cutoff. Data include all observations on student performance within 180 days of their 21st birthday. Controls include course by semester by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT math and verbal scores, academic composite score, leadership composite score, fitness score, and indicator variables for Black, Hispanic, Asian, and recruited athlete. The bandwidth of the data is 180 days on either side of Age 21.

*** Significant at the 0.01 level.

Table 4
Regression discontinuity estimates for different bandwidths and specifications.

Bandwidth	270 days		240 days		210 days		150 days		120 days		100 days		80 days		60 days		40 days		20 days	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	12	13	14	14	14	
Panel A: no controls																				
Discontinuity at age 21	-0.057** (0.03)	-0.057** (0.04)	-0.057** (0.03)	-0.104** (0.04)	-0.062** (0.03)	-0.117*** (0.04)	-0.088** (0.04)	-0.086** (0.04)	-0.095** (0.04)	-0.067*** (0.03)	-0.081* (0.05)	-0.067*** (0.03)	-0.078*** (0.03)	-0.062* (0.04)	-0.079* (0.04)					
Panel B: with controls																				
Discontinuity at age 21	-0.066*** (0.02)	-0.114*** (0.03)	-0.070 (0.02)	-0.137*** (0.03)	-0.073*** (0.02)	-0.131*** (0.03)	-0.119*** (0.03)	-0.125*** (0.03)	-0.131*** (0.04)	-0.122*** (0.04)	-0.087*** (0.04)	-0.086** (0.04)	-0.087** (0.04)	-0.057 (0.05)	-0.085 (0.11)					
Observations	58,032	58,032	51,431	51,431	45,092	45,092	32,321	25,930	21,410	17,121	17,121	17,121	12,670	8240	4080					
Age polynomial	Linear	Quadratic	Linear	Quadratic	Linear	Quadratic	Linear	None	None	None	None	None	None	None						

Notes: Each cell represents results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for age 21. Standard errors clustered by age are in parentheses. Controls include course by semester by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT math and verbal scores, academic composite score, leadership composite score, fitness score, and indicator variables for Black, Hispanic, Asian, and recruited athlete.

* Significant at the 0.10 level.

** Significant at the 0.05 level.

*** Significant at the 0.01 level.

Table 5
Regression discontinuity estimates by subgroup.

Sample	Main results	Top half freshman GPA	Bottom half freshman GPA	Male	Female	Math and science courses	Social sciences and humanities courses
	1	2	3	4	5	6	7
<i>Panel A: no controls</i>							
Discontinuity at age 21	−0.092*** (0.03)	−0.095** (0.04)	−0.01 (0.04)	−0.101*** (0.04)	−0.037 (0.08)	−0.094** (0.04)	−0.090** (0.04)
<i>Panel B: controls</i>							
Discontinuity at age 21	−0.106*** (0.02)	−0.145*** (0.03)	0.024 (0.04)	−0.115*** (0.03)	−0.109 (0.09)	−0.125*** (0.03)	−0.075** (0.03)
Observations	38,782	22,180	16,602	31,766	7016	19,783	18,999

Notes: Each cell represents results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for Age 21. Standard errors clustered by age are in parentheses. All specifications control for a linear function of distance from Age 21 in which the slope is allowed to vary on either side of the cutoff. Controls include course by semester by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT math and verbal scores, academic composite score, leadership composite score, fitness score, and indicator variables for Black, Hispanic, Asian, recruited athlete, and preparatory school attendance. The bandwidth of the data is 180 days on either side of Age 21.

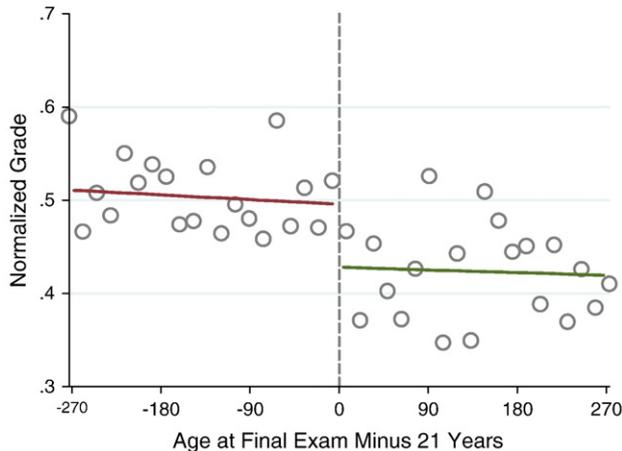
** Significant at the 0.05 level.

*** Significant at the 0.01 level.

3.5. Birthday effects

An alternative interpretation of the performance decline at age 21 is that it represents the effect of birthdays generally, rather than the effect of drinking. To test this hypothesis, we perform a similar regression discontinuity analysis for students turning 20. If the performance declines estimated earlier were the result of a birthday

A. Top Half



B. Bottom Half

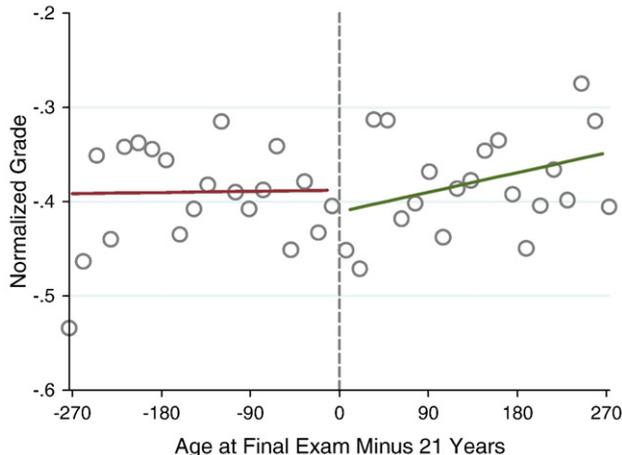


Fig. 4. Regression discontinuity estimates for students in the top and bottom halves of their freshman class by GPA.

Table 6
Robustness checks.

Sample	Main results	Includes freshmen	Includes preparatory school students	Mandatory core courses
	1	2	3	4
<i>Panel A: no controls</i>				
Discontinuity at age 21	−0.092*** (0.03)	−0.098*** (0.03)	−0.068** (0.03)	−0.088* (0.05)
<i>Panel B: controls</i>				
Discontinuity at age 21	−0.106*** (0.02)	−0.108*** (0.02)	−0.091*** (0.02)	−0.074** (0.04)
Observations	38,782	39,508	46,271	12,680

Notes: Each cell represents results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for Age 21. Standard errors clustered by age are in parentheses. All specifications control for a linear function of distance from Age 21 in which the slope is allowed to vary on either side of the cutoff. Controls include course by semester by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT math and verbal scores, academic composite score, leadership composite score, fitness score, and indicator variables for Black, Hispanic, Asian, and recruited athlete. The bandwidth of the data is 180 days on either side of Age 21.

* Significant at the 0.10 level.

** Significant at the 0.05 level.

*** Significant at the 0.01 level.

effect rather than the increase in drinking at age 21, we would expect to observe a similar decline at age 20. Results are shown in Fig. 5, with corresponding regression discontinuity estimates presented in the first 4 columns of Table 7. The results indicate relatively small declines in performance, and only one of eight estimates is statistically significant at the 10% level. More importantly, the underlying data shown in Fig. 5 do not appear indicative of a drop in performance at age 20.

We also test whether performance declines discontinuously at age 22. In contrast to age 20, it is possible that drinking increases discontinuously at age 22. However, we would expect any such increase to be considerably smaller than the increase at age 21. Results are shown in Fig. 6, with corresponding regression discontinuity estimates presented in columns 5 through 8 of Table 7. In order to ensure comparability of the sample to the left and right of the age cutoff, we restrict the sample to grades received in the fall semester.¹⁵ As with age 20, there is little compelling graphical or statistical evidence of a meaningful drop in performance; none of the 8 estimates is statistically significant at the 10% level.

¹⁵ That is, since a significant proportion of students do not turn 22 until after the final exams in the spring of their senior year, we would never observe these students on the right-hand side of the age 22 cutoff in the spring semester.

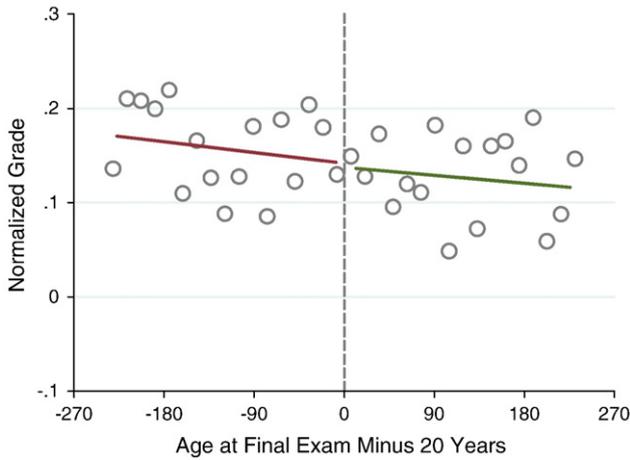


Fig. 5. The effect of turning 20 years old on academic achievement.

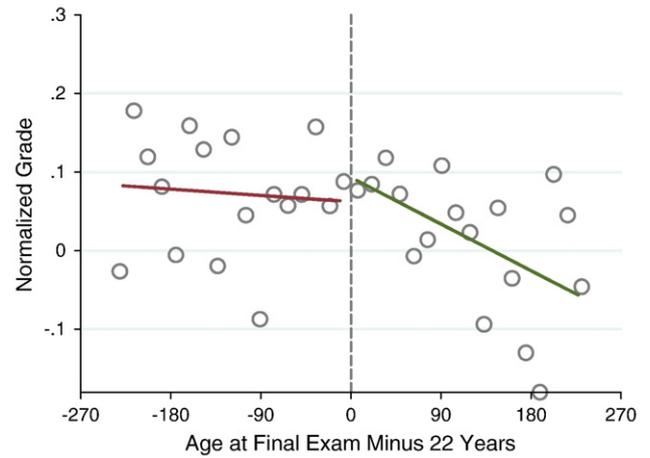


Fig. 6. The effect of turning 22 years old on academic achievement.

4. Conclusion

While there has been a considerable amount of research on the consequences of alcohol consumption, existing studies have faced one of two significant limitations. First, attempts to exploit state-level laws have resulted in questions regarding both the relevance and the validity of the instruments. Second, the majority of the existing research has focused on low-frequency outcomes such as fatalities and out-of-wedlock childbearing, so it is unclear whether drinking has consequences on a broader set of individuals.

This study addresses the consequences of drinking on academic performance in college, which allows us to examine a high frequency individual-level outcome. To overcome the identification problem caused by selection into drinking, we exploit the discontinuity in drinking that occurs at age 21 at a unique educational institution in which the minimum legal drinking age is strictly enforced. Results show that drinking prior to and during final exam week causes a statistically and economically meaningful reduction in academic performance. The performance drop is of approximately the same magnitude as being assigned to a professor whose quality is one standard deviation below average.

Moreover, we find that these effects are largely driven by the highest-performing students. This suggests that the negative con-

sequences of drinking are not limited to the narrow segment of the population at risk of more severe, low-frequency outcomes.

References

Amethyst Initiative, www.amethystinitiative.org, accessed February 23, 2010.
 Carpenter, Christopher, Dobkin, Carlos, 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.
 Carrell, Scott E., West, James E., 2010. Does professor quality matter? Evidence from random assignment of students to professors. *Journal of Political Economy* 118 (3), 409–432.
 Chaloupka, Frank J., Wechsler, Henry, 1996. Binge drinking in college: the impact of price, availability, and alcohol control policies. *Contemporary Economic Policy* 14, 112–124.
 Cooperative Institutional Research Program, 2007. The Freshman Survey. Higher Education Research Institute at the University of California, Los Angeles.
 Dave, Dhaval, Kaestner, Robert, 2002. Alcohol taxes and labor market outcomes. *Journal of Health Economics* 21, 357–371.
 Dee, Thomas S., 1999. State alcohol policies, teen drinking, and traffic fatalities. *Journal of Public Economics* 72, 289–315.
 Dee, Thomas S., 2001. The effects of minimum legal drinking ages on teen childbearing. *The Journal of Human Resources* 36 (4), 823–838.
 Fertig, Angela R., Watson, Tara, 2009. Minimum drinking age laws and infant health outcomes. *Journal of Health Economics* 28, 737–747.
 Imbens, Guido W., Lemieux, Thomas, 2008. Regression discontinuity designs: a guide to practice. *Journal of Econometrics* 142 (2), 615–635.

Table 7

Falsification checks: age 20 and age 22 birthday effects.

Bandwidth	Age 20 effect				Age 22 effect			
	240 days	180 days	120 days	80 days	240 days	180 days	120 days	80 days
	1	2	3	4	5	6	7	8
<i>Panel A: no controls</i>								
Discontinuity at age 21	–0.004 (0.03)	–0.029 (0.04)	–0.032 (0.04)	–0.005 (0.05)	0.031 (0.05)	0.027 (0.05)	–0.040 (0.06)	0.009 (0.08)
<i>Panel B: controls</i>								
Discontinuity at age 21	–0.018 (0.02)	–0.031 (0.02)	–0.053* (0.03)	–0.030 (0.04)	0.031 (0.06)	0.020 (0.06)	0.006 (0.06)	–0.021 (0.08)
Age polynomial	Linear	Linear	Linear	Linear	Linear	Linear	Linear	Linear
Observations	43,237	34,072	23,862	16,125	17,717	14,000	10,172	6972

Notes: Each cell represents results for separate regression where the dependent variable is normalized course grade and the key independent variable is an indicator for Age 20 or Age 22. Standard errors clustered by age are in parentheses. All specifications control for a linear function of distance from Age 20 or Age 22 in which the slope is allowed to vary on either side of the cutoff. Controls include course by semester by section fixed effects, graduating class by semester by year at USAFA fixed effects, birth year fixed effects, SAT math and verbal scores, academic composite score, leadership composite score, fitness score, and indicator variables for Black, Hispanic, Asian, and recruited athlete. The Age 22 effects are for the fall semester only.

* Significant at the 0.10 level.

- Ludwig, Jens, Miller, Doug, 2007. Does head start improve children's outcomes? Evidence from a regression discontinuity design. *Quarterly Journal of Economics* 122 (1), 159–208.
- Miron, Jeffrey A., Tetelbaum, Elina, 2007. Does the minimum legal drinking age save lives? *Economic Inquiry* 47 (2), 317–336.
- National Institute on Alcohol Abuse and Alcoholism, 2010. Research Findings on College Drinking and the Minimum Legal Drinking Age. <http://www.niaaa.nih.gov/NR/rdonlyres/F099FF68-37B3-4EBC-8573-93CC768A74AA/0/CollegeDrinkingMLDA.pdf>, accessed February 24, 2010.
- Saffer, Henry, 1997. Alcohol advertising and motor vehicle fatalities. *The Review of Economics and Statistics* 79 (3), 431–442.
- White, Aaron M., Kraus, Courtney L., Swartzwelder, Harry S., 2006. Many college freshmen drink at levels far beyond the binge threshold. *Alcoholism, Clinical and Experimental Research* 30 (6), 1006–1010.