Am I My Brother's Barkeeper? Sibling Spillovers in Alcohol Consumption at the Minimum Legal Drinking Age*

Eunju Lee
Louisiana State University

Geoffrey C. Schnorr
United States Military Academy at West Point

July 10, 2024

Abstract

We use data on siblings near the minimum drinking age to provide causal estimates of peer effects in alcohol consumption, exploiting the increase in consumption of the older sibling in a regression discontinuity design. We find no evidence for positive spillover effects of older sibling's legal access to alcohol on the younger sibling alcohol consumption. Although imprecise, preferred point estimates imply that younger sibling binge drinking *decreases* at the cutoff. These negative reduced form spillover effects are larger for siblings who are likely to spend more time together, for measures of excessive alcohol consumption, and in subgroups where the first stage discontinuity is largest. We argue that these patterns of heterogeneity are consistent with younger siblings learning from the costs of their older siblings' drinking behavior.

Keywords: Alcohol, Peer effects, Siblings

JEL Classification: I12, I18, J13

*Lee: eunjulee@lsu.edu. Schnorr (corresponding author): gcschnorr@gmail.com. We are thankful for guidance from Marianne Bitler, Scott Carrell, Marianne Page, Brendan Price, and Monica Singhal. We also thank Kitt Carpenter, Partha Deb, Chloe East, Emiliano Huet-Vaughn, Krzysztof Karbownik, Konstantin Kunze, Derek Rury, and Justin Wiltshire for helpful comments and conversations. Schnorr gratefully acknowledges financial support from the Bilinski Educational Foundation.

1 Introduction

Excessive alcohol consumption has been shown to harm young adults in a variety of dimensions including health, educational performance, criminal activity, and criminal victimization (Carpenter and Dobkin, 2009; Carrell et al., 2011; Lindo et al., 2013; Carpenter and Dobkin, 2015, 2017; Chalfin et al., 2023). Among the many potential determinants of adolescent alcohol use, peer effects have received substantial attention from academics and policymakers.[1]

Leading empirical work provides evidence for positive peer effects in alcohol consumption via quasi-random assignment to college roommates (Duncan et al., 2005; Eisenberg et al., 2014; Guo et al., 2015) or primary/secondary school cohort-mates (Argys and Rees, 2008) who consume more alcohol. While convincing, such designs cannot be replicated in many other peer groups of interest, notably siblings. A rapidly growing literature has used natural experiments to demonstrate the importance of sibling influences in a variety of important domains[2], and strong positive correlations in sibling alcohol consumption are well-established in the public health literature (Duncan et al., 2001; Fagan and Najman, 2005; Trim et al., 2006; Van Der Vorst et al., 2007; Whiteman et al., 2013). However, causally interpretable estimates of sibling spillovers in alcohol consumption are virtually non-existent.[3]

This paper provides causally interpretable estimates of peer effects in alcohol consumption between siblings. We focus on siblings residing in the same household and exploit a discontinuous increase in older sibling alcohol consumption at the minimum legal drinking age (MLDA) using a regression discontinuity design (RDD). Our results are, to our knowledge, the first quasi-experimental estimates of peer effects in alcohol consumption between siblings.

Attempts to estimate causal peer effects face four main difficulties. First, the "reflection" problem implies that a simple regression of peer A's behavior on peer B's cannot determine the direction of the effect (Manski, 1993). Second, peer groups are typically endogenous, which is a concern if individuals choose peers who have similar preferences. Third, peers are likely to experience unobservable common shocks which are correlated with their outcomes. Finally, since many research designs rely on variation in exposure to peers (and not peer behavior directly), researchers are often unable to determine whether a spillover effect is the result of the peer behavior of interest or some correlated peer characteristic—they are "contextual" as

[1]In its Underage Drinking Fact Sheet, the National Institute of Alcohol Abuse and Alcoholism lists peer pressure as one of three key causes of alcohol consumption among adolescents (https://pubs.niaaa.nih.gov/publications/UnderageDrinking/UnderageFact.htm).

[2]Including, for example, smoking (Harris and López-Valcárcel, 2008), health outcomes (Breining, 2014; Ho, 2017; Cawley et al., 2019; Daysal et al., 2019), fertility (Heissel, 2021), and education (Goodman et al., 2015; Joensen and Nielsen, 2018; Karbownik and Özek, 2023; Altmejd et al., 2021)

[3]The economic literature on this question that we are aware of is limited to Altonji et al. (2016), whose empirical approach (a joint dynamic probit model) relies on the assumption that any unobserved confounding variable correlated with the alcohol consumption of both siblings is time invariant.

opposed to "contagion" effects (Manski, 2000; Argys and Rees, 2008). Contagion effects create social multipliers in the behavior of interest (an intervention that reduces one person's alcohol consumption also reduces their peers'), while contextual effects may not.

Our setting and identification strategy address each of these difficulties. By focusing on siblings, a peer group which is naturally occurring, we avoid the potential for selection into the peer group. By restricting the sample to siblings with different ages, we ensure that the variation in alcohol consumption that we utilize is both exogenous (avoiding the problem of common shocks) and specific to one of the two siblings (avoiding the reflection problem). Since our identification strategy provides variation in the older sibling's alcohol consumption directly, we isolate contagion effects without any additional identification assumptions.

We utilize data from the 1997 National Longitudinal Survey of Youth (NLSY97), which has two key features for our analyses. First, the data contains unique "household roster" information on all individuals living with the NLSY97 respondent. These rosters include the birth months of each household member and their relationship to the respondent. Second, it contains several high frequency (past month) measures of alcohol consumption.

In the full sample of 5,220 sibling-pair-years, our estimates imply that there is no positive spillover effect of older sibling's legal access to alcohol on the younger sibling alcohol consumption. Since we typically observe alcohol consumption for only one sibling in a household, we focus on reduced form effects of the older sibling turning 21. Our preferred specification implies that the number of binge drinking days (5 or more drinks) in the past 30 days reported by younger siblings *decreases* by 0.25 days at the cutoff. Although imprecise, this is a substantial effect—the average younger sibling in the sample reports 1.25 binge drinking days in the past month. This estimate is also well into the left tail of a distribution of placebo discontinuities estimated at older sibling ages that are close, but not equal to, 21 years. Point estimates for other measures of more moderate alcohol consumption (e.g., number of days with any alcohol consumption) are sometimes positive, but always small in magnitude, and never statistically significantly different from zero.

It is important to emphasize that we only have statistical power to detect large effect sizes. However, our point estimates for binge drinking outcomes are consistently negative, economically meaningful, and precise enough to rule out relatively large positive effects in the related literature. We compare our estimates to those from Eisenberg et al. (2014) (henceforth EGW) since they utilize what is arguably the gold-standard research design (randomized college roommates) in the literature on peer effects in alcohol consumption, and their estimates are smaller than nearly all other estimates in the related literature.[4] Among 64 total specifications which can be compared to EGW, 58 of the 95% confidence intervals exclude the EGW point estimate.[5]

We interpret our results as providing suggestive evidence that the sibling spillover effects we estimate are negative. While a negative spillover effect is counterintuitive, a large literature has established that a series of negative outcomes spike at the MLDA (e.g., Carpenter and Dobkin, 2009; Carrell et al., 2011; Lindo et al., 2013; Carpenter and Dobkin, 2015, 2017; Chalfin et al., 2023) and it is plausible that younger siblings update their beliefs about the costs of alcohol consumption after observing their older siblings experience these negative consequences. Unfortunately, direct tests for this mechanism are not possible due to data constraints, but we do provide several pieces of evidence supporting a true negative effect.

First, we show that the negative spillover effects are larger in sibling pairs that we would expect to spend more time together (same gender siblings, older siblings that are not away at college) and those that experience larger first stage discontinuities in older sibling's own alcohol consumption at the MLDA (those with higher socio-economic status, non-Black and non-Hispanic race/ethnicity, or a male older sibling). We argue that this latter pattern provides evidence in favor of the exclusion restriction since it is unlikely that any other pathway from older sibling legal drinking status to younger sibling drinking behavior would exhibit this same variation across subgroups.

Next, we demonstrate that our estimates are especially negative for measures of excessive consumption (binge drinking). Prior work suggests that negative effects of alcohol consumption experienced at the MLDA are driven by these types of consumption (Carpenter et al., 2016), and it is reasonable to expect that any spillover effects would be most apparent on this margin as well.

We also show that these negative spillover effects do not occur in the other direction—i.e., older sibling binge drinking does not decrease when a younger sibling turns 21. This can be seen as a falsification test for the proposed mechanism since an older sibling is less likely to learn from a younger sibling's alcohol consumption. Point estimates in this sample are much closer to zero and more likely to be positive for binge drinking outcomes.

We address three specific concerns which may threaten our interpretation of these results. First, having an older sibling who can legally purchase alcohol may make it easier for the younger sibling to obtain it. However, this access effect would imply that our negative estimates are too high. Second, parents may prevent positive spillover effects when the older sibling turns 21

[4]Notably Gaviria and Raphael (2001), Duncan et al. (2005), Lundborg (2006), Fletcher (2012), Guo et al. (2015), and Altonji et al. (2016) report similar or larger peer effects in their preferred specifications.

[5]EGW's comparable result is that a roommate who binge drank immediately before move-in increases the respondent's probability of any binge drinking roughly 8 months later by 19%.

by monitoring younger siblings more closely, which would bias our results downwards. We are able to provide evidence against this concern using survey questions on parenting behaviors. Third, it is possible that younger siblings are not sufficiently aware of the alcohol consumption of their older siblings despite living in the same household. This might occur if the older sibling is temporarily away at college or if the siblings do not spend much time together for some other reason. To address the former concern, we show that many older siblings in our sample are not enrolled in school and that negative spillover effects are concentrated in this group. To address the latter, we use the American Time Use Survey (ATUS) to demonstrate that siblings similar to those in our sample spend substantial amounts of time together.

We make three main contributions to the literature. First, we provide what we believe to be the first quasi-experimental estimates of sibling spillovers in alcohol consumption. Despite the importance of siblings and the strong correlation in alcohol consumption between siblings (e.g., Duncan et al., 2001; Fagan and Najman, 2005; Trim et al., 2006; Van Der Vorst et al., 2007; Whiteman et al., 2013), prior studies on spillovers in alcohol consumption have focused on non-sibling peer groups (e.g., schoolmates (Argys and Rees, 2008), roommates (Duncan et al., 2005; Eisenberg et al., 2014; Guo et al., 2015), or spouses (Fletcher and Marksteiner, 2017)), or relied on much stronger identification assumptions (Altonji et al., 2016). By exploiting an exogenous shock in older sibling consumption, we provide causal estimates of sibling spillover effects and contribute to the literature on alcohol consumption spillovers between siblings. Our results suggest that, at least in the age groups that we study, contemporaneous positive correlations in sibling alcohol consumption are not contagion effects. This is important as it implies that policymakers and parents should not expect that interventions aimed at decreasing the alcohol consumption of adolescents would have beneficial spillovers on their siblings.

Second, we add to a limited number of studies which provide quasi-experimental estimates of contagion effects in alcohol consumption more generally (i.e., in any peer group), such as Argys and Rees (2008), Eisenberg et al. (2014), and Fletcher and Marksteiner (2017). The most similar work to ours in this respect is Fletcher and Marksteiner (2017), who estimate contagion effects of alcohol consumption with a different identification strategy (randomized assignment to an alcohol cessation program) in a very different population (adult spouses in which one spouse is a heavy drinker).

Third, we contribute to the wider literature on sibling spillovers, where a growing body of evidence has demonstrated the importance of sibling influences in a variety of important domains including health outcomes (Breining, 2014; Ho, 2017; Cawley et al., 2019; Daysal et al., 2019), fertility (Heissel, 2021), education (Goodman et al., 2015; Joensen and Nielsen, 2018; Karbownik and Özek, 2023; Altmejd et al., 2021), and military service (Bingley et al., 2021).

2 Data

The NLSY97 is a longitudinal survey of 8,984 American youths who were between the ages of 12 and 17 at the start of the survey in 1997 (Bureau of Labor Statistics, 2019). Survey Waves are annual from 1997-2011, and biennial beginning in 2013. Designed and administered by the Bureau of Labor Statistics, the NLSY97 includes a wide variety of questions related to family processes, education, employment, health, and family formation, among other topics.

The survey has two key features which make it uniquely suitable for our analysis. First, it includes detailed information on the composition of each respondent's household, including the relationship of each household member to the respondent and the age of each household member in months. Specifically, the public use version of the NLSY97 includes the exact date of the interview along with the month of birth for the respondent and all other members of the respondent's household. This allows us to observe the age in months of the respondent and all siblings who ever co-resided with the respondent during the survey, even if those siblings are not NLSY97 respondents themselves. Second, information on the alcohol consumption of respondents is obtained in all waves.

Since the NLSY97 sampled households and identified all age-eligible individuals in each household, a sample of siblings who are all NSLY97 respondents is available. Multi-respondent households are relatively rare, so we do not use this sample in our main analysis. However, we do use it to provide descriptive evidence of between-sibling correlation in alcohol consumption in the NLSY97.[6] This also implies that our main analyses cannot estimate the effect of older sibling alcohol consumption (e.g., via instrumental variables) and instead focus on the reduced form effect of an older sibling's MLDA status.

Each wave of the NLSY97 also includes a detailed set of covariates at the respondent and household levels which are also potentially related to alcohol consumption and therefore will serve as control variables in our analysis. These include educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, an indicator for whether the respondent worked in the past year, and dummies for the calendar month and year of the survey.

[6] Table A1 shows results from a regression of the younger sibling's past month alcohol consumption on the older sibling's past month alcohol consumption for multiple measures of consumption in a sample of sibling pairs that are both NLSY97 respondents. Households with 3 or more sibling NLSY97 respondents are excluded for simplicity and the sample is limited to individuals under the age of 23. These results are not causally interpretable and are not directly comparable to our sample of interest. However, they show that in this sample, much like in prior literature described previously, there is a very strong correlation in alcohol consumption between siblings.

We focus on measures of the excessive alcohol consumption in our main analyses: The number of binge drinking days (5+ drinks) in the past month, and an indicator for any binge drinking days in the past month. This helps make our results more comparable to the existing literature, which typically focuses on measures of binge drinking, and more policy relevant, since excessive drinking is more likely to result in the various negative outcomes that are associated with alcohol consumption. The focus on a small subset of the available outcomes also helps to reduce problems of multiple testing. However, to investigate robustness some analyses also use binary and count measures of drinking days (days in the past month on which any drinks were consumed) as outcomes.

For reasons that we will describe in the next section, we will primarily focus on 2,614 NLSY respondents who have only one older sibling in their household, and where the older sibling is between the ages of 19 years & 0 months (228 months of age) and 23 years & 0 months old (276 months of age). The *younger* siblings in this sample are on average 18 years (215.8 months) old, and 95.1% of them are above the age of 15 years (180 months). (See Figure A1 for the full younger sibling age distribution.) Summary statistics for both the full sample of sibling-pair-years in the data and those used in our main analyses are shown in Table 1.

NLSY97 household rosters are based on the respondent's permanent address and hence include siblings who spent most of the year away at college.[7] However, we note that nearly half of the observations in our main sample come from sibling-pair years where the older sibling is not enrolled in school, that not all college students live on campus, and that many college students who do live on campus attend a school that is near their permanent address (National Center for Education Statistics, 2000). According to the National Postsecondary Student Aid Study of Undergraduates, 42% of students were enrolled in public 2-year institutions in the 1999-2000 academic year (which are likely close-to-home community colleges), 50% of these public 2-year students still lived at home with their families, and 27% of all postsecondary students (including students attending private or 4-year institutions) lived with their parents.

Alternative samples drop either the requirement that the siblings *currently* live together,[8] the requirement that the two siblings are the oldest two siblings, or both. Key subgroups of interest are defined based on parental educational attainment, family income, race/ethnicity (recorded in the data as "Black or Hispanic", or "Non-Black, Non-Hispanic"), whether the siblings report the same gender, the gender of the older sibling, and whether the older sibling is enrolled in school.

[7]This is technically only the case for the 1997-2002 survey waves, but the 1997-2002 survey waves do account for 91.05% of our main sample.

[8]They must have lived together during at least one NLSY97 wave in order for us to observe the sibling and the sibling's age in months.

Despite the unique advantages of our data, readers should be aware that we have limited statistical power to detect even moderately sized effects. To make this clear, we performed power calculations based on our sample size and a minimum detectable effect size taken from the related literature. As described in Section 4.3 we rely on estimates from Eisenberg et al. (2014) as primary comparison to the prior literature, so we use this estimate as our minimum detectable effect size. We find that our dataset has a power of 0.681 to detect the effect size of the EGW estimate in absolute terms (a 0.086pp increase in the probability of any binge drinking) and a power of 0.305 in relative terms (an increase of 19% of the control mean). Readers should therefore exercise caution when interpreting our estimates that are not statistically significantly different from zero.

3 Methods

In this section we first explain our research design and the identification assumptions that it relies on. We then explain how we use the data and design to estimate both the reduced form effect of an older sibling's legal drinking status on the younger sibling's alcohol consumption, and the first stage effect of the older sibling's legal drinking status on their own alcohol consumption.

3.1 Estimating sibling spillover effects

In our main analyses, we implement a reduced form regression discontinuity design in which the older sibling's legal access to alcohol (an indicator for whether they are at least 21 years old) acts as an instrument for the older sibling's alcohol consumption. The MLDA was first used as an exogenous source of variation in alcohol consumption in an RDD by Carpenter and Dobkin (2009) to study the effect of alcohol consumption on mortality and has since been used to study a wide range of other outcomes.[9] So long as no other factors related to the outcome are changing discontinuously at the cutoff, this approach will provide causally interpretable estimates. Prior work using this design has provided convincing evidence of both a strong first stage (large, discontinuous increases in consumption at age 21) and of the credibility of the research design (e.g., establishing that observable covariates do not change discontinuously at the cutoff).

Many NLSY97 respondents have more than one older sibling, and this complicates the estimation of peer effects with an RDD since it is unclear how to define the running variable.

[9] Examples include criminal activity (Carpenter and Dobkin, 2015; Hansen and Waddell, 2018), crime victimization (Chalfin et al., 2023), morbidity (Carpenter and Dobkin, 2017), marijuana consumption (Yörük and Yörük, 2011; Crost and Guerrero, 2012; Crost and Rees, 2013; Yörük and Yörük, 2013), the consumption of other illegal drugs (Deza, 2015), risky sexual behavior (Yörük and Yörük, 2015), and mental health (Yörük and Yörük, 2012).

To avoid these complications, we define the peer group as siblings who currently reside in the same household and consider the effect of the oldest sibling in the peer group on the second oldest sibling. This ensures that we have a large sample of sibling pairs in which we observe the information necessary to implement the RDD (older sibling age in months and younger sibling alcohol consumption) and allows us to avoid the complications involved with defining the running variable in peer groups with more than two members.

The reduced form RD is then implemented by estimating the following equation via OLS:

$$alc_{2ht} = \gamma_1 1\{age_{1ht} \geq 21\} + f(age_{1ht}) + X'_{2ht}\gamma_2 + X'_{1ht}\gamma_3 + W'_{ht}\gamma_4 + \theta_{2h} + \mu_{2ht} \qquad (1)$$

where subscripts denote sibling (2 = younger, 1 = older), household (h), and the survey wave (t), the outcome is some measure of the younger sibling's past month alcohol consumption, the reduced form effect of older sibling's legal access to alcohol on the younger sibling's consumption is estimated by γ_1, $f(age_{1ht})$ is a flexible polynomial in the running variable (e.g., older sibling age fully interacted with the cutoff dummy $1\{age_{1ht} \geq 21\}$), X_{2ht} is a vector of younger sibling covariates, X_{1ht} is a vector of older sibling covariates, W_{ht} is a vector of household h level covariates, and θ_{2h} is a younger sibling fixed effect.

To construct the running variable, we use the birth month of the older sibling and the month of the relevant interview. This implies that our running variable is rounded *up*, and that the cutoff indicator will be mismeasured for some sibling pairs in which the older sibling is exactly 21 years (252 months) old. Following the recommendations in Dong (2015), we address this misclassification bias with a "donut RD" specification that excludes sibling pairs in which the older sibling is exactly 21 years (252 months) old. Following the existing literature on MLDA-based RDDs, the sample is limited to sibling pairs in which the older sibling is between the ages of 19 and 23—i.e., the youngest older sibling in our sample is 19 years and 0 months old (228 months of age) and the oldest is 23 years and 0 months old (276 months of age). Standard errors are cluster robust at the younger sibling level to account for correlation in the error term among observations from the same individual.

Additional models test the sensitivity of our results in several of the dimensions mentioned above (bandwidth, correction for rounding-induced bias, inclusion of covariates, inclusion of fixed effects, and order of the running variable polynomial). We also implement the "continuity-based" RD framework suggested by Cattaneo et al. (2019), including the use of a mean-squared error optimal bandwidth from Calonico et al. (2020), a triangular kernel, and bias-corrected robust confidence intervals as described in Calonico et al. (2014, 2018).

The main assumption required for a causal interpretation of γ_1 is that no unobserved confounding factors change discontinuously at the cutoff. We support this assumption by estimating models similar to equation 1, where outcomes are predicted values from separate regressions that predict the alcohol consumption measures using a range of covariates. A second required assumption is that the running variable is not manipulated. Although this is not technically a concern in our setting (since age is not manipulable), a related problem can arise if rates of nonresponse to alcohol consumption questions change discontinuously at the cutoff. This is primarily a concern in the first stage (described in the next subsection) where an older sibling may be more willing to report alcohol consumption once they reach age 21. We test this assumption by demonstrating visually that the density of older sibling age is smooth through the cutoff and by testing formally for a discontinuity in the density at the cutoff. Finally, to interpret γ_1 as the causal effect *of older sibling alcohol consumption* (as opposed to older sibling legal access), we require an exclusion restriction: the older sibling's legal access must influence younger sibling alcohol consumption *only* through an increase in the older sibling's consumption. We provide suggestive evidence in favor of this assumption in Sections 4.4 and 4.5.

3.2 First stage discontinuities in older sibling alcohol consumption

We separately estimate the increase in older sibling drinking at age 21 in a similar equation (where the outcome is the alcohol consumption of the older sibling, alc_{1ht}, and younger sibling characteristics are removed from the regression). Importantly, the sibling pairs used in that analysis are not the same as those used to estimate the reduced form effect. Sibling pairs in the reduced form sample are those in which the second oldest sibling in the household is an NLSY97 respondent. Sibling pairs in the first stage are those in which the oldest sibling is an NLSY97 respondent.

While we could rely on the prior literature as evidence for the strength of the first stage, estimating the first stage discontinuity ourselves is useful for two reasons. First, we are interested in the magnitude of the first stage specifically for older siblings. Since the increase in alcohol consumption at MLDAs has been shown to be heterogeneous (Ahammer et al., 2022; Carpenter et al., 2016) and 21-year-olds with younger siblings likely have different characteristics than those without (e.g., families with one child likely differ from those with two or more in various dimensions), the existing literature may not directly speak to this question. Second, as explained further in Section 4.5, we will use the correlation between the subgroup-level first stage and reduced form estimates to provide suggestive evidence in support of the exclusion restriction.

4 Main Results

4.1 Research design

We begin by demonstrating that older sibling drinking behavior changes discontinuously at age 21. Results using count and binary measures of drinking days and binge days in the past month are shown in Table 2. Similar to results from previous work using the same identification strategy in different data,[10] there is a large discontinuous increase in alcohol consumption at age 21. This effect is apparent in all models shown in Table 2. In preferred models (donut specification with a linear function of the running variable, controls,[11] and individual level fixed effects), the past month binge drinking and drinking days increase by 0.42 days and 1.36 days at the cutoff, respectively. Both increases are statistically significant at the 5% level. Increases in the extensive margin of consumption for binge drinking and drinking (from the same models) are 5.6pp and 8.1pp respectively, which are also both statistically significant at the 5% level. Figure 1 plots mean values of our four outcomes among older siblings in each month-of-age bin along with fitted lines and regression coefficients from the corresponding first stage regressions.

Figure A2 graphs the distribution of the older sibling's age in months, for the first stage sample. It is possible for nonresponse to change discontinuously at the cutoff, and such a response could potentially affect the interpretation of the results presented in Table 2. Figure A2 suggests that this is not occurring, given that the distribution is relatively smooth through the cutoff at 21 years of age.

In Table A2 and Figure 2 we support our assumption that unobservable sibling pair characteristics are smooth through the cutoff by showing that observables do not change discontinuously. Since there are many relevant observable characteristics, we use the following three-step procedure to test for the smoothness of many covariates in a single regression:

1. Regress an outcome of interest (measure of younger sibling alcohol consumption) on the relevant covariates.[12]

2. Use the point estimates from this regression to predict the outcome for each younger sibling in our sample.

[10] See the previous section for citations.

[11] The controls include month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year.

[12] Specifically, this regression includes the following covariates: age, educational enrollment, highest completed education, work status, indicators for whether the household lives in an urban area, census region dummies, household size, and interview month/year. (All variables refer to the younger sibling's information unless otherwise noted.)

3. Use our reduced from regression specification (equation 1) to test for the presence of a discontinuity in these predicted values at older sibling age = 21. Standard errors are bootstrapped.

Discontinuities for each of our four measures of younger sibling alcohol consumption are statistically indistinguishable from zero.

4.2 Reduced form

Estimates of equation 1 in Table 3 provide no evidence of a positive effect of older sibling legal access to alcohol on younger sibling alcohol consumption. Across models (columns) and outcomes (panels), point estimates are typically negative and often moderately sized. This is especially true for binge drinking outcomes, where our preferred model in the first column (which drops sibling pairs where the older sibling is exactly 21 years old and includes both controls and individual fixed effects) implies that the probability of any binge drinking in the past month decreases by 1.8pp at the cutoff and that the number of binge drinking days in the past month decreases by 0.245 days. Scaled by the average value of the outcome just before the cutoff ("control mean" in the Table 3) these point estimates imply economically significant reductions of 6.7% (any binge drinking) and 19.5% (number of binge drinking days), respectively. However, no estimates in Table 3 are statistically significantly different from zero, and point estimates for the number of drinking days in the past month are always positive, although very imprecise.

Corresponding binned scatter plots for all four outcomes are shown in Figure 3. These figures include coefficients from a regression of each outcome variable on the age of the older sibling (in months, centered at 21 years old), an indicator for older sibling age ≥ 21, and their interaction (i.e., they correspond to a specification *without* controls or fixed effects and which includes sibling pairs where the older sibling is exactly 21 years old, and therefore do not align with our preferred specifications). The figures show no significant change in the alcohol consumption of younger sibling when their older sibling turns 21, suggesting (as also shown in Table 3) that the size and precision of our negative estimates are somewhat sensitive to specification choices.

In Figure 4, we perform a series of placebo tests estimating discontinuities at older sibling age values near but not equal to 21. Specifically, for each monthly age bin up to 12 months on either side of the cutoff, we re-estimate equation 1, using the preferred donut-RD specification with a linear polynomial of the running variable, controls, and fixed effects. The results suggest that our negative point estimates for these outcomes are informative despite their imprecision—the true estimate "(vertical dark gray line)" (For online-only: "(vertical blue line)") is in the far left tail of both placebo distributions.

We view the results in this subsection as providing no evidence of a positive spillover effect and providing *suggestive* evidence of a negative spillover effect for binge drinking behavior. While we believe that the negative effects we estimate are plausible, they are imprecise and we are underpowered to detect even moderately sized negative effects. In light of this concern, we focus next on demonstrating that our preferred estimates rule out leading estimates of positive spillover effects in alcohol consumption from other peer groups.

4.3 Comparison to prior literature

Since the related literature consistently estimates positive spillover effects in alcohol consumption (e.g., Gaviria and Raphael, 2001; Duncan et al., 2005; Clark and Lohéac, 2007; Lundborg, 2006; Fletcher, 2012; Eisenberg et al., 2014; Guo et al., 2015; Fletcher and Marksteiner, 2017), a null result precise enough to rule out a meaningfully sized positive effect would still constitute an important contribution. To see if this is the case, we compare our results to similar estimates from Eisenberg et al. (2014) (EGW).

We chose EGW as our comparison for three reasons. First, although our focus is on siblings, the economic literature on alcohol consumption spillovers among siblings is very limited (one study, Altonji et al., 2016). Second, EGW's estimates are smaller than nearly all other estimates in the related literature. Therefore, if our null results are sufficiently precise to rule out the EGW point estimate, the same is true for a series of other leading estimates.[13] Third, EGW utilize what we see as the best existing identification strategy on this question: randomized assignment to college roommates.

Since spillover effects likely differ across peer groups we do not use this comparison to inform interpretation of EGW's estimates. Still, for the reasons mentioned above we consider the EGW estimate the most appropriate benchmark in the related literature. The conclusion that our estimates provide no evidence of positive effects would be made stronger if they could rule out effects comparable to EGW's.

EGW's key comparable result is that being assigned a roommate who binge drank in 30 days prior to a baseline survey fielded in August of 2009 (just before the students moved into their shared rooms) leads to a 19% increase in the probability of any binge drinking by the respondent in a second 30-day period roughly 8 months later. Although the EGW treatment (a roommate who did not binge drink vs. one who did) may seem larger than the effects of the MLDA on older-sibling alcohol consumption, timing implies that the difference is small.

[13] Notably Gaviria and Raphael (2001), Duncan et al. (2005), Lundborg (2006), Fletcher (2012), Guo et al. (2015), and Altonji et al. (2016) all report similar or larger peer effects of alcohol consumption in their preferred specifications.

The EGW treatment is a 100% increase in exposure to peer binge drinking measured roughly 8 months in the past, while we measure a contemporaneous spillover effect. Especially given the trajectory of alcohol consumption in this age range, it is likely that many of the "control" peers in EGW (no prior binge drinking) would have binge drank by the time the outcomes were measured 8 months later. Therefore, the contemporaneous treatment vs. control difference in peer binge drinking in the EGW sample is likely much smaller than 100%.

Our estimates are nearly always precise enough to rule out positive spillover effects as large as EGW's estimates. In Figure 5 we scale our preferred estimate for the any binge drinking outcome by the control (immediately pre-cutoff) mean, obtaining a 95% CI of [-27.6%,9.2%] in our preferred specification and ruling out the EGW point estimate (dashed horizontal line) in all five specifications. Additional results discussed further in Section 5.2 compare the EGW point estimate to a wide range of alternative specifications. The EGW point estimate falls outside of our 95% CIs in 58 of the 64 specifications we consider.

4.4 Potential mechanisms

An understanding of the potential mechanisms for sibling spillover effects in our context can help answer two key questions about our reduced form results. First, are we measuring the effect of older sibling alcohol consumption, or something else that changes when the older sibling turns 21? Second, why might an older sibling's legal access to alcohol lead to *a reduction* in younger sibling consumption?

In an influential review of the economic literature on social interactions, Manski (2000) suggests that economic peer effects can operate through one of three channels: constraints, preferences, or expectations. Here we use this taxonomy to organize a discussion of suggestive evidence in favor of the exclusion restriction and of potential mechanisms underlying our negative spillover effect estimates.

In our context, there are two clear constraints on a younger sibling's alcohol consumption that could be influenced by their older sibling's legal drinking status. First, there is the potential for an "access" effect. An older sibling's ability to legally purchase alcohol could ease the access constraint faced by the younger sibling—even if the older sibling's consumption remains unchanged. We cannot test for this effect in our data since we do not observe how (or how easily) NLSY97 respondents obtain alcohol. However, we note that an access effect would bias our negative reduced form estimates toward zero. Second, when an older sibling's alcohol consumption (or access to alcohol) increases, parents may respond by monitoring the younger sibling more closely. We can test for such responses in the NLSY97 data, and the results presented in Appendix C suggest that this does not occur.

Preferences act as a mechanism for peer effects when the utility that one person derives from a behavior depends on the behavior of their peer. Preferences are commonly hypothesized as mechanisms for positive peer effects in alcohol consumption—drinking together is more fun. It is more difficult to see how such a mechanism could explain a *negative* peer effect in alcohol consumption.

One plausible story related to preferences is that siblings enjoy consuming alcohol together but spend less time together once the older sibling turns 21. If preferences to consume together are strong enough, or the younger sibling's consumption relies on the presence of the older sibling for some other reason (say, access), then this could explain the reduction in younger sibling consumption that we observe. However, we can provide some suggestive evidence against this story using the ATUS—which allows us to observe a day's worth of minute-level activities for a large sample of young adults who live with younger siblings. In an analysis described further in Figure D2, we show that time spent with a younger sibling does not change noticeably at age 21.

A second plausible story related to preferences is that siblings regularly serve as designated drivers for each other so that more frequent alcohol consumption for one necessarily means less for the other. To explore this, we compare reduced form results for younger siblings with and without a driver's license. Estimates are similar between the two groups, suggesting that this does not explain the negative spillover effects in Section 4.2.[14]

Expectations are perhaps the most promising class of mechanisms in our context. A peer effect operates through expectations when there is uncertainty about the costs or benefits of some behavior, and peers can reduce that uncertainty by observing each other's experiences. Alcohol consumption is a costly behavior, and this is especially true for the types of excessive consumption that respond most strongly to the MLDA (Carpenter et al., 2016). Indeed, we know that several negative outcomes spike at the MLDA. It is reasonable to expect that adolescents are uncertain about these costs and that they would learn about them by observing role models like older siblings.

Many of the specific negative outcomes that have been studied in the MLDA literature (e.g., mortality (Carpenter and Dobkin, 2009), hospital admissions (Carpenter and Dobkin, 2017), criminal behaviors (Carpenter and Dobkin, 2015), etc.) are likely too rare to explain the negative spillover effects that we estimate. However, other less severe and more common

[14]Specifically, we estimate that binge drinking days among younger siblings with a driver's license (who we can reasonably assume are much more likely to serve as designated drivers) decrease by 0.37 days (se=0.21) while binge drinking days among younger siblings without a driver's license decrease by 0.38 days (se=0.29).

negative outcomes also likely spike at age 21. Notably, prior work provides evidence that academic performance decreases at the MLDA (Carpenter and Dobkin, 2015; Lindo et al., 2013), and poor academic performance is an example of a more commonplace negative outcome that could drive our results. To name a few other hypothetical examples, older siblings whose risky drinking behavior increases at age 21 might get in more trouble at home (or at work or school), get less sleep, or spend lots of time lying on the couch nursing a hangover.

Another important caveat to our proposed expectations mechanism is that it relies on the younger sibling's ability to observe their older sibling's behavior and its consequences. We provide three pieces of suggestive evidence for this. First, our main results focus on siblings who reside in the same household, and, as discussed further in Section 5, our reduced form estimates attenuate when we relax this restriction. Second, in Appendix D, we outline a descriptive analysis that uses data from the ATUS to demonstrate that sibling pairs similar to those in our NLSY97 sample spend substantial amounts of time together (Figure D1), and—as mentioned above—this does not change substantially when the older sibling turns 21 (Figure D2). Third, in Section 4.5 we show that the negative spillover effects are concentrated among subgroups of sibling pairs that are likely to spend more time together (same-gender sibling pairs, and sibling pairs where the older sibling is not enrolled in school).

Unfortunately, data on adolescent expectations about the *relevant* consequences of alcohol consumption are hard to come by. The NLSY97 and some other surveys do contain survey questions on more severe and long-term risks of alcohol consumption, such as liver disease, heart disease, and alcoholism. However, it is unlikely that a younger sibling would update their beliefs about those risks based on their older sibling's *contemporaneous* experiences.[15] Instead, empirically testing the expectations mechanism in our context would require survey questions on expectations about smaller and more immediate costs of alcohol consumption. To our knowledge, these do not exist in a data set with sample size and information (sibling age in months) sufficient for our purposes.

However, the NLSY97 data does allow us to estimate spillover effects on older siblings when a younger sibling turns 21. This is useful because the expectations mechanism we have proposed is unlikely to be relevant when a younger sibling turns 21. Older siblings likely have much less uncertainty about the consequences of alcohol consumption by the time their younger siblings turn 21, and younger siblings are less likely to serve as role models. For these reasons, this

[15] Even if such risk perceptions were of interest, the NLSY97 only asks these questions in a small number of survey waves (restricting our sample size), and other potential data sources have similar issues. Notably, the National Health Interview Survey (NHIS) has few respondents under age 21, and the National Longitudinal Study of Adolescent Health (Add Health) lacks the detailed household information necessary to construct our running variable.

exercise can be seen as a falsification test for the proposed mechanism.

Table A3 implements this falsification test using several different specifications and outcomes. Since it is less common for sibling pairs in this age range to live together (here the younger sibling is between the ages of 19 and 23), sample sizes under our typical sample definitions are relatively small in this analysis. Therefore we also present these results with alternative samples that remove either the requirement that the siblings currently live together, the requirement that the siblings are the oldest two siblings in the household, or both. Point estimates, especially for binge drinking outcomes, are closer to zero and less precisely estimated. We view this as additional suggestive evidence in support of our main results.

It is worth emphasizing that these are not the same siblings used in our main analysis. Here the older sibling is the NLSY97 respondent (whereas in, e.g., Table 3 the younger sibling is), and here the older sibling is roughly 24 years old, on average (in, e.g., Table 3 the estimate applies at the cutoff where the older sibling is 21). Therefore, estimates between the siblings could be driven by these other differences, which may not be as informative about our proposed mechanism. One notable example is that siblings may spend less time together at these older ages. To provide some additional context on this possibility, in results not shown we extend the sample used in Figure D2 to include older siblings up to 25 years of age. We find that time spent with younger siblings does decrease as the older sibling ages from 21-25 years, but that siblings still spend meaningful amounts of time together at these ages. Specifically, 21-year-old older siblings in the ATUS sample spend 9.5% of their day with a younger sibling. At ages 22, 23, 24, and 25, the values are 8.5%, 9.3%, 6.9%, and 7.8%, respectively.

4.5 Heterogeneity

In this section we provide evidence that the reduced form effects in Section 4.2 vary substantially across subgroups of sibling pairs, and that the patterns of this heterogeneity support a causal interpretation of our results.

We distinguish broadly between two types of hypotheses about this reduced form heterogeneity. First, the nature of our setting suggests several subgroups of sibling pairs in which the treatment effect of older sibling alcohol consumption on younger sibling consumption should be larger. If first stage discontinuities in older sibling consumption are similar across these subgroups, then this heterogeneity in *treatment effects* should also show up in the reduced form (since treatment effect = reduced form/first stage). Second, estimates from the existing literature on the MLDA and from Section 4.1 above suggest several subgroups of sibling pairs in which the first stage should be larger. If treatment effects are similar across these subgroups, then this first stage heterogeneity should lead to comparable heterogeneity in the reduced form.

We expect that any treatment effect in our setting would be stronger among sibling pairs who are likely to spend more time together, such as same-gender sibling pairs (Dunifon et al., 2017) or sibling-pairs where the older sibling is not enrolled in school. As shown in Table 5, this is indeed the case. Among same-gender sibling pairs, younger sibling binge drinking days decrease by 0.54 days at the cutoff (statistically significant at the 5% level), while younger sibling binge drinking days increase by 0.05 days among different gender sibling pairs. Among sibling pairs where the older sibling is not enrolled in school, younger sibling binge drinking days decrease by 0.49 days (statistically significant at the 10% level), while younger sibling binge drinking days decrease by only 0.03 days if the older sibling is enrolled in school.

In other dimensions, where we instead would expect treatment effects to be relatively homogeneous, our negative reduced form estimates should be concentrated in the subgroups where the first stage is strongest. This follows from the exclusion restriction, which requires that an older sibling turning 21 influences the younger sibling's alcohol consumption only through an increase in the older sibling's alcohol consumption. If, for example, a subgroup with a very small first stage effect (little to no increase in older sibling consumption at age 21) also had a very large reduced form effect (large increase in younger sibling consumption when the older sibling turns 21), that would suggest a different causal pathway.

Prior work on the MLDA has uncovered consistent patterns of heterogeneity in the first stage. Specifically, the increase in (own) alcohol consumption at age 21 has been shown to be driven by excessive consumption among males (Carpenter et al., 2016). Similarly heterogeneous reduced form effects in our setting would provide support for our exclusion restriction. Here again we see the expected heterogeneity. Table 5 shows that younger sibling binge drinking days decrease by 0.487 at the cutoff (statistically significant at the 5% level) when the older sibling is male, and increase by 0.06 days at the cutoff when the older sibling is female. We also note that the negative reduced form effects are strongest for measures of excessive alcohol consumption where the first stage effect has been concentrated in the prior literature, as shown previously in Table 4.

As described in Section 4.1, we can estimate first stage effects in our data for similar samples of sibling pairs. This allows us to directly test for a meaningful correlation between first stage and reduced form estimates across subgroups observable in the NLSY97 data.[16]

The key takeaway from this exercise is immediately clear in Figure 6, which plots the

[16] Appendix B.1 follows Angrist et al. (2022) to demonstrate how this exercise follows from a causal model which asserts that all across-subgroup variation in the treatment effect is driven by variation in the strength of the first stage, leading to an overidentification test. We do not formally implement the overidentification test because our first stage and reduced form estimates come from separate samples.

subgroup-specific reduced form and first stage estimates against each other: Subgroups with large negative reduced form effects are typically the same subgroups that have large positive first stage effects.[17] More specifically, we find that *both* the increase in older sibling alcohol consumption and the decrease in younger sibling alcohol consumption that occur at the cutoff are concentrated among sibling pairs that are white, male, or from high socioeconomic status families. Again, this negative correlation is what we should expect to see if the exclusion restriction holds. If the correlation were instead positive, this would imply either that treatment effects are highly heterogeneous or that an older sibling turning 21 influences the outcome through some pathway other than the treatment.

Table 4 and Table 5 display the point estimates underlying Figure 6. The first stage increase in older sibling binge drinking is largest for males (0.613 days vs. 0.245 days among females), Non-Black and Non-Hispanics (0.610 days vs. 0.201 days among Blacks and Hispanics), families with high socioeconomic status (0.508 days for the high parental education group vs. 0.260 for the low parental education group and 0.884 days for the high household income group vs. 0.167 for the low household income group), older siblings not enrolled in school (0.466 days vs. 0.345 days for those enrolled), and siblings who have the same sex (0.540 days vs. 0.461 days). Differences are less stark in the reduced form, but are always more negative in subgroups with larger first stage discontinuities, and only positive among sibling pairs where the older sibling is female or the siblings have a different sex.

A common concern with subgroup-level analyses in RDDs is that any given subgroup dummy is likely correlated with other covariates, so that it is unclear whether heterogeneity in the treatment effect is driven by the subgroup dummy of interest or some other covariate. One could control for these other covariates (and their interaction with the treatment), but in an RDD these covariates would also need to be interacted with the running variable polynomial so that this approach quickly becomes untenable without very large sample sizes. Such differences present an additional reason for the first stage and reduced form estimates to diverge and threaten the logic underlying those analyses.

To address this we follow Carril et al. (2017) and Gerardino et al. (2017) and estimate a propensity score for each subgroup using a set of observable covariates,[18] and then weight the subgroup-level regressions by the inverse of that propensity score. Inverse propensity score

[17] To demonstrate robustness, we also provide the same results from regressions estimated without any covariates (i.e., the specification from column 4 in Table 3) in Table A4, Table A5, and Figure A3.

[18] Specifically, we regress educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, an indicator for whether the respondent worked in the past year, and dummies for the calendar month and year of the survey on the relevant subgroup dummy. In the reduced form specifications we also control for the older sibling's gender. In each case we omit the covariate that the score is estimated for from the set of controls as needed.

weighted results are shown in Table A6, Table A7, and panel (b) in Figure 1. Additional details on the estimation of these propensity scores are in Appendix B.2. Results are broadly consistent with the unweighted results.

Taken together, we believe that these patterns of heterogeneity support an interpretation of our reduced form estimates as the causal effect of older sibling alcohol consumption. It is unlikely that a spurious result or some violation of the exclusion restriction would show up in exactly the subgroups and margins of alcohol consumption where either the first stage result is strongest or where there is a strong reason to believe siblings are more exposed to each other.

5 Robustness

In this section, we explore the robustness of our results to different sample selection criteria and to various specification choices such as the inclusion of controls or FEs.

5.1 Alternative samples

Our main analyses focus on sibling pairs who currently reside in the same household, are the oldest two siblings in that household, and where the older sibling is between the ages of 19 and 23. Table 6 presents results that remove either the requirement that the siblings currently live together, the requirement that the siblings are the oldest two siblings in the household, or both. The first panel corresponds to Table 3, which uses our main sample. Point estimates for our main outcome are broadly consistent (small, negative, not always statistically significant) across these samples. For our preferred outcome (binge drinking days) estimates are closer to zero in the alternative samples. We view this as reassuring since a true spillover effect is likely to be larger when the siblings live together (siblings who do not live together likely spend less time with each other) and when the peer is the only older sibling present (if there are multiple older siblings the behavior of any one older sibling may become less influential).

5.2 Comparing many specifications

We visually summarize the overall robustness of our results in a plot of 64 specifications (most of which have been presented in other parts of this paper), which vary by the inclusion of fixed effects, the inclusion of controls, the order of the running variable polynomial, and the sample, in Figure A4 and Figure A5. Each regression in Figure A4 uses an indicator for any binge drinking as the outcome, and each coefficient is scaled by the corresponding "control" mean (mean of the outcome just below the cutoff)—facilitating comparisons with the EGW point estimate of +19%. Among 64 total specifications, 57 scaled estimates are negative, though the magnitudes vary and confidence intervals usually do not rule out a null effect. Among 64 total specifications, 58 of the 95% confidence intervals exclude the EGW point estimate. Figure A5 shows similar results with the count of binge drinking days as the outcome.

One additional reassuring pattern in this figure is that the handful of outlier CIs (in both figures) which include large positive effects *all* include quadratic polynomials in the running variable. Visual inspection of various binscatters in this paper suggests that the relationship between the outcome and running variable is linear. Further, higher-order polynomials in RD designs are known to increase the risk of detecting spurious effects (Gelman and Imbens, 2019).

5.3 Fixed effects

Figure A4 and (to a lesser extent) Figure A5 highlight that our results are somewhat sensitive to the inclusion of younger sibling FEs. Models with FEs generally have more negative point estimates and tighter confidence intervals. While the improved precision is expected, the sensitivity of the point estimates to the inclusion of FEs may appear surprising given that FEs should not be necessary for identification in our research design, and that adding FEs does not change the estimation sample (as shown in, e.g., Table 3). However, as explained in detail by Miller et al. (ming), FE estimates are identified only by groups that have variation in treatment (here, by younger siblings who show up in our sample at least once before their older sibling turns 21 and at least once after).

In Table A8 we demonstrate that such "switchers" make up roughly two-thirds of the corresponding regression sample used in our main analyses (siblings who live together and are the two oldest siblings in the household), and that these switchers differ meaningfully in some dimensions from the remainder of the sample ("non-switchers"). Given prior results (Table 5) showing that treatment effects in this context are heterogeneous, we take this as suggestive evidence that the sensitivity of our main results to the inclusion of FEs is driven by heterogeneity in the strength of the first stage, the treatment effect, or both. Since FEs also buy us improvements in precision, and we have no *a priori* reason to prefer estimates for the combined sample over the switcher sample, we consider the FE estimates to be preferred.

Table A9 shows characteristics of similar switcher and non-switcher subsets of the alternative sample that *does not* require the siblings to live together at the time of the sample.[19] A common reason for sibling pairs in our main sample to appear only when the older sibling is under 21 (i.e., to be non-switchers) is that the older sibling moves out too soon. It is therefore not surprising

[19]The sample is used in the bottom panel of Table 6 and denoted as "Sample 4" in Figure A4 and Figure A5.

that switchers make up a much larger proportion of this alternative sample (without the "same household" restriction). Reassuringly, Figure A4 and Figure A5 also demonstrate that results in this alternative sample are less sensitive to the inclusion of FEs.

5.4 Bandwidth and estimation

All results discussed thus far have relied on a bandwidth of 24 months (i.e., restricting the sample to sibling pairs in which the older sibling is between the ages of 19 years and 0 months and 23 years and 0 months). This follows prior literature utilizing MLDA RDDs, which nearly universally chooses this bandwidth (e.g., Carpenter and Dobkin, 2009; Deza, 2015; Carpenter et al., 2016).

We demonstrate the robustness of our main results to the choice of bandwidth in two ways. First, Figure A6 graphs point estimates and confidence intervals from the preferred model[20] for 36 separate bandwidths (single month increments from 12 to 48 months). Results are stable across bandwidth choices, although (expectedly) less precise as bandwidth shrinks. Second, we verify that our results persist when using the framework for estimation and inference in RDDs developed by Cattaneo et al. (2019). This approach involves local polynomial estimation of the discontinuity using a mean-squared error (MSE) optimal bandwidth (Calonico et al., 2020), a triangular kernel, and bias-corrected robust confidence intervals (Calonico et al., 2014, 2018). Table A10 presents results from these specifications. Columns (1) and (3) use MSE-optimal bandwidths, while columns (2) and (4) use our ad-hoc 24-month bandwidth. All columns include the MSE-optimal estimates and bias-corrected robust 95% confidence intervals where the standard errors are cluster robust at individual level. Consistent with Table 3, all coefficients are negative with moderate magnitudes. For models with individual fixed effects, the coefficients are large and statistically significant compared to models without fixed effects, a pattern also present in Table 3.

6 Conclusions

We focus on a population that allows for the estimation of causally interpretable contagion effects in alcohol consumption under relatively weak assumptions: sibling pairs close to the MLDA in the United States. This setting is helpful for two reasons. First, it allows us to avoid empirical problems that commonly plague attempts to estimate causal peer effects. Our research design estimates contagion effects (the effect of older sibling alcohol consumption) rather than

[20]The preferred model is a donut RD specification including a vector of controls, a linear polynomial in the running variable fully interacted with the cutoff dummy, and individual level fixed effects.

contextual effects (the effect of exposure to a peer), avoids common shocks and the reflection problem, and focuses on a peer group that is not selected. Second, despite growing evidence on the importance of sibling influences, causally interpretable estimates of sibling spillovers in alcohol consumption are rare. While somewhat imprecise, we are consistently able to rule out large positive effects reported in the prior literature on non-sibling peer groups. We also provide suggestive evidence that younger siblings engage in less binge drinking when the older sibling gains legal access to alcohol.

These results have several important implications for policymakers, parents, and other stakeholders interested in reducing excessive consumption of alcohol by adolescents. First, they suggest that well-known contemporaneous correlations in sibling alcohol consumption are not causal, at least among the types of sibling pairs in our data. Second, our results speak indirectly to recent debates about the desirability of a lower MLDA (Carpenter and Dobkin, 2011; Wechsler and Nelson, 2010; DeJong and Blanchette, 2014). Some stakeholders have argued that a lower MLDA will help adolescents learn to consume alcohol responsibly via earlier exposure but recent empirical work (Ahammer et al., 2022) has pushed back against this claim.[21] We believe that our results provide suggestive evidence for a related hypothesis—that adolescents may learn more responsible drinking behavior by observing the experiences of their older siblings. Given the extensive evidence that alcohol consumption and its harms rise substantially at the MLDA, we find it unlikely that this indirect benefit (decreased lower sibling consumption) would outweigh the direct costs of a lower MLDA (increased older sibling consumption). However, our results point to a potential strategy for other interventions that attempt to improve adolescent behavior around alcohol—foster interaction with older siblings (or other role models) who have more direct experience with the substance *and* its negative consequences.

Several limitations are worth mentioning. First, as in any RDD, external validity is a concern. Our results apply only to a certain peer group, adolescent siblings residing in the same household where the older sibling is near the MLDA. However, our results are less limited in this way than a typical MLDA-based RDD, given that the running variable and the outcome are taken from different individuals (ages range from 12-20 years for younger siblings in our sample whose older sibling is within one month of the cutoff).

Second, an older sibling's legal drinking status could *potentially* affect younger sibling consumption indirectly, e.g., via parental responses or access effects. However, one key set of confounders (parental responses) is, to some extent, ruled out based on observable information

[21] Specifically, Ahammer et al. (2022) shows that the spike in alcohol consumption at the MLDA is actually larger in a lower-MLDA setting (Austria, where the MLDA is 16 years) and provides additional evidence which suggests that obtaining legal access actually makes adolescents perceive excessive consumption as *less* risky.

in the data, and another (access effects) is likely to bias our results towards zero.

An important goal for future work will be to understand the differences between the growing set of well-identified results on peer effects in alcohol consumption. The research designs used in our work, Fletcher and Marksteiner (2017), Argys and Rees (2008), and the series of papers utilizing randomly assigned college roommates (Duncan et al., 2005; Eisenberg et al., 2014; Guo et al., 2015) all differ in two key dimensions: the peer group studied and the assumptions necessary to disentangle contagion effects from contextual effects. Future work should aim to understand to what extent these two factors explain the differences in these results and to understand the nature of spillovers in alcohol consumption in more generalizable contexts.

References

Ahammer, A., S. Bauernschuster, M. Halla, and H. Lachenmaier (2022). Minimum legal drinking age and the social gradient in binge drinking. Journal of health economics 81, 102571. (Cited on pages 10 and 23.)

Altmejd, A., A. Barrios-Fernández, M. Drlje, J. Goodman, M. Hurwitz, D. Kovac, C. Mulhern, C. Neilson, and J. Smith (2021). O brother, where start thou? sibling spillovers on college and major choice in four countries. The Quarterly Journal of Economics. (Cited on pages 2 and 6.)

Altonji, J. G., S. Cattan, and I. Ware (2016). Identifying sibling influence on teenage substance use. Journal of Human Resources, 0714–6474R1. (Cited on pages 2, 4, 5, and 13.)

Angrist, J., D. Autor, and A. Pallais (2022). Marginal effects of merit aid for low-income students. The Quarterly Journal of Economics 137(2), 1039–1090. (Cited on page 18.)

Argys, L. M. and D. I. Rees (2008). Searching for peer group effects: A test of the contagion hypothesis. The Review of Economics and Statistics 90(3), 442–458. (Cited on pages 2, 3, 5, and 24.)

Bingley, P., P. Lundborg, and S. V. Lyk-Jensen (2021). Brothers in arms spillovers from a draft lottery. Journal of Human Resources 56(1), 225–268. (Cited on page 6.)

Breining, S. N. (2014). The presence of adhd: Spillovers between siblings. Economics Letters 124(3), 469–473. (Cited on pages 2 and 5.)

Bureau of Labor Statistics (2019). National longitudinal survey of youth 1997 cohort, 1997-2017 (rounds 1-18). U.S. Department of Labor. (Cited on page 6.)

Calonico, S., M. D. Cattaneo, and M. H. Farrell (2018). On the effect of bias estimation on coverage accuracy in nonparametric inference. Journal of the American Statistical Association 113(522), 767–779. (Cited on pages 9 and 22.)

Calonico, S., M. D. Cattaneo, and M. H. Farrell (2020). Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. The Econometrics Journal 23(2), 192–210. (Cited on pages 9 and 22.)

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. Econometrica 82(6), 2295–2326. (Cited on pages 9 and 22.)

Carpenter, C. and C. Dobkin (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. (Cited on pages 2, 4, 8, 15, and 22.)

Carpenter, C. and C. Dobkin (2011). The minimum legal drinking age and public health. The Journal of Economic Perspectives 25(2), 133–156. (Cited on page 23.)

Carpenter, C. and C. Dobkin (2015). The minimum legal drinking age and crime. The Review of Economics and Statistics 97(2), 521–524. (Cited on pages 2, 4, 8, 15, and 16.)

Carpenter, C. and C. Dobkin (2017). The minimum legal drinking age and morbidity in the united states. The Review of Economics and Statistics 99(1), 95–104. (Cited on pages 2, 4, 8, and 15.)

Carpenter, C. S., C. Dobkin, and C. Warman (2016). The mechanisms of alcohol control. Journal of Human Resources 51(2), 328–356. (Cited on pages 4, 10, 15, 18, and 22.)

Carrell, S. E., M. Hoekstra, and J. E. West (2011). Does drinking impair college performance? evidence from a regression discontinuity approach. Journal of Public Economics 95(1), 54–62. (Cited on pages 2 and 4.)

Carril, A., A. Cazor Katz, M. P. Gerardino, S. Litschig, and D. Pomeranz (2017). Rddsga: Stata module to conduct subgroup analysis for regression discontinuity designs. (Cited on page 19.)

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2019). A practical introduction to regression discontinuity designs: Foundations. Cambridge University Press. (Cited on pages 9 and 22.)

Cawley, J., E. Han, J. Kim, and E. C. Norton (2019). Testing for family influences on obesity: The role of genetic nurture. Health economics 28(7), 937–952. (Cited on pages 2 and 5.)

Chalfin, A., B. Hansen, and R. Ryley (2023). The minimum legal drinking age and crime victimization. Journal of Human Resources 58(4), 1141–1177. (Cited on pages 2, 4, and 8.)

Clark, A. E. and Y. Lohéac (2007). "it wasn't me, it was them!" social influence in risky behavior by adolescents. Journal of health economics 26(4), 763–784. (Cited on page 13.)

Crost, B. and S. Guerrero (2012). The effect of alcohol availability on marijuana use: Evidence from the minimum legal drinking age. Journal of health economics 31(1), 112–121. (Cited on page 8.)

Crost, B. and D. I. Rees (2013). The minimum legal drinking age and marijuana use: New estimates from the nlsy97. Journal of health economics 32(2), 474–476. (Cited on page 8.)

Daysal, N. M., M. Simonsen, M. Trandafir, and S. Breining (2019). Spillover effects of early-life medical interventions. The Review of Economics and Statistics, 1–46. (Cited on pages 2 and 5.)

DeJong, W. and J. Blanchette (2014). Case closed: research evidence on the positive public health impact of the age 21 minimum legal drinking age in the united states. Journal of Studies on Alcohol and Drugs, Supplement (s17), 108–115. (Cited on page 23.)

Deza, M. (2015). The effects of alcohol on the consumption of hard drugs: regression discontinuity evidence from the national longitudinal study of youth, 1997. Health economics 24(4), 419–438. (Cited on pages 8 and 22.)

Dong, Y. (2015). Regression discontinuity applications with rounding errors in the running variable. Journal of Applied Econometrics 30(3), 422–446. (Cited on page 9.)

Duncan, G. J., J. Boisjoly, and K. M. Harris (2001). Sibling, peer, neighbor, and schoolmate correlations as indicators of the importance of context for adolescent development. Demography 38(3), 437–447. (Cited on pages 2 and 5.)

Duncan, G. J., J. Boisjoly, M. Kremer, D. M. Levy, and J. Eccles (2005). Peer effects in drug use and sex among college students. Journal of abnormal child psychology 33(3), 375–385. (Cited on pages 2, 4, 5, 13, and 24.)

Dunifon, R., P. Fomby, and K. Musick (2017). Siblings and children's time use in the united states. Demographic Research 37, 1611–1624. (Cited on page 18.)

Eisenberg, D., E. Golberstein, and J. L. Whitlock (2014). Peer effects on risky behaviors: New evidence from college roommate assignments. Journal of health economics 33, 126–138. (Cited on pages 2, 3, 5, 8, 13, 24, and 33.)

Fagan, A. A. and J. M. Najman (2005). The relative contributions of parental and sibling substance use to adolescent tobacco, alcohol, and other drug use. Journal of Drug Issues 35(4), 869–883. (Cited on pages 2 and 5.)

Fletcher, J. and R. Marksteiner (2017). Causal spousal health spillover effects and implications for program evaluation. American Economic Journal: Economic Policy 9(4), 144–66. (Cited on pages 5, 13, and 24.)

Fletcher, J. M. (2012). Peer influences on adolescent alcohol consumption: evidence using an instrumental variables/fixed effect approach. Journal of Population Economics 25(4), 1265–1286. (Cited on pages 4 and 13.)

Gaviria, A. and S. Raphael (2001). School-based peer effects and juvenile behavior. The Review of Economics and Statistics 83(2), 257–268. (Cited on pages 4 and 13.)

Gelman, A. and G. Imbens (2019). Why high-order polynomials should not be used in regression discontinuity designs. Journal of Business & Economic Statistics 37(3), 447–456. (Cited on page 21.)

Gerardino, M. P., S. Litschig, and D. Pomeranz (2017). Can audits backfire? evidence from public procurement in chile. NBER Working Papers (23978). (Cited on page 19.)

Goodman, J., M. Hurwitz, J. Smith, and J. Fox (2015). The relationship between siblings' college choices: Evidence from one million sat-taking families. Economics of Education Review 48, 75–85. (Cited on pages 2 and 5.)

Guo, G., Y. Li, C. Owen, H. Wang, and G. J. Duncan (2015). A natural experiment of peer influences on youth alcohol use. Social science research 52, 193–207. (Cited on pages 2, 4, 5, 13, and 24.)

Hansen, B. and G. R. Waddell (2018). Legal access to alcohol and criminality. Journal of health economics 57, 277–289. (Cited on page 8.)

Harris, J. E. and B. G. López-Valcárcel (2008). Asymmetric peer effects in the analysis of cigarette smoking among young people in the united states, 1992–1999. Journal of health economics 27(2), 249–264. (Cited on page 2.)

Heissel, J. A. (2021). Teen fertility and siblings' outcomes evidence of family spillovers using matched samples. Journal of Human Resources 56(1), 40–72. (Cited on pages 2 and 5.)

Ho, C. Y. (2017). Estimating sibling spillovers in health: Evidence on symptoms. Economics & Human Biology 27, 93–101. (Cited on pages 2 and 5.)

Joensen, J. S. and H. S. Nielsen (2018). Spillovers in education choice. Journal of Public Economics 157, 158–183. (Cited on pages 2 and 5.)

Karbownik, K. and U. Özek (2023). Setting a good example?: Examining sibling spillovers in educational achievement using a regression discontinuity design. Journal of Human Resources 58(5), 1567–1607. (Cited on pages 2 and 5.)

Lindo, J. M., I. D. Swensen, and G. R. Waddell (2013). Alcohol and student performance: Estimating the effect of legal access. Journal of health economics 32(1), 22–32. (Cited on pages 2, 4, and 16.)

Lundborg, P. (2006). Having the wrong friends? peer effects in adolescent substance use. Journal of health economics 25(2), 214–233. (Cited on pages 4 and 13.)

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. The Review of Economic Studies 60(3), 531–542. (Cited on page 2.)

Manski, C. F. (2000). Economic analysis of social interactions. The Journal of Economic Perspectives 14(3), 115–136. (Cited on pages 3 and 14.)

Miller, D. L., N. Shenhav, and M. Z. Grosz (forthcoming). Selection into identification in fixed effects models, with application to head start. Journal of Human Resources. (Cited on page 21.)

National Center for Education Statistics (2000). National postsecondary student aid survey: 2000 undergraduates (npsas:ug). United States Department of Education. Institute of Education Sciences https://nces.ed.gov/surveys/npsas (accessed August 26, 2022). (Cited on page 7.)

Trim, R. S., E. Leuthe, and L. Chassin (2006). Sibling influence on alcohol use in a young adult, high-risk sample. Journal of studies on alcohol 67(3), 391–398. (Cited on pages 2 and 5.)

Van Der Vorst, H., R. C. Engels, W. Meeus, M. Deković, and J. Van Leeuwe (2007). Similarities and bi-directional influences regarding alcohol consumption in adolescent sibling pairs. Addictive behaviors 32(9), 1814–1825. (Cited on pages 2 and 5.)

Wechsler, H. and T. F. Nelson (2010). Will increasing alcohol availability by lowering the minimum legal drinking age decrease drinking and related consequences among youths? American journal of public health 100(6), 986–992. (Cited on page 23.)

Whiteman, S. D., A. C. Jensen, and J. L. Maggs (2013). Similarities in adolescent siblings' substance use: Testing competing pathways of influence. Journal of studies on alcohol and drugs 74(1), 104–113. (Cited on pages 2 and 5.)

Yörük, B. K. and C. E. Yörük (2011). The impact of minimum legal drinking age laws on alcohol consumption, smoking, and marijuana use: Evidence from a regression discontinuity design using exact date of birth. Journal of health economics 30(4), 740–752. (Cited on page 8.)

Yörük, B. K. and C. E. Yörük (2013). The impact of minimum legal drinking age laws on alcohol consumption, smoking, and marijuana use revisited. Journal of health economics 32(2), 477–479. (Cited on page 8.)

Yörük, C. E. and B. Yörük (2015). Alcohol consumption and risky sexual behavior among young adults: evidence from minimum legal drinking age laws. Journal of Population Economics 28(1). (Cited on page 8.)

Yörük, C. E. and B. K. Yörük (2012). The impact of drinking on psychological well-being: Evidence from minimum drinking age laws in the united states. Social Science & Medicine 75(10), 1844–1854. (Cited on page 8.)

Figures and Tables

Figure 1: Discontinuity in Older Sibling Alcohol Consumption

(a) Count of Drinking Days

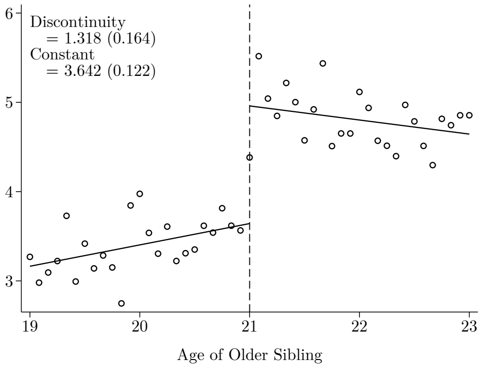

(b) Prob. of Any Drinking

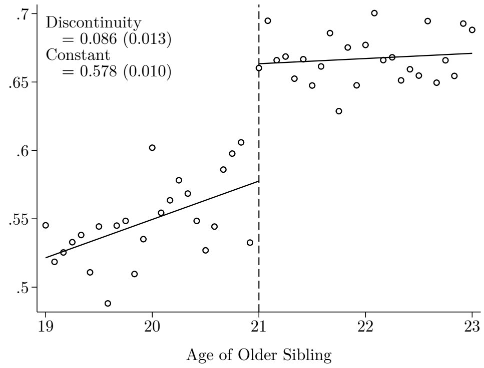

(c) Count of Binge Drinking Days

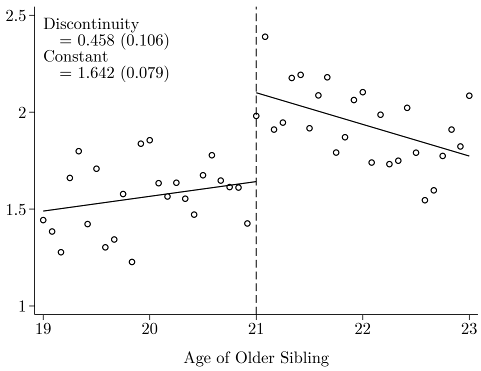

(d) Prob. of Any Binge Drinking

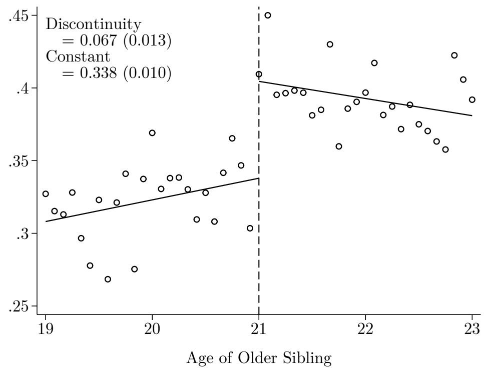

Note: Each panel shows mean alcohol consumption in each older-sibling-age bin (in months), with linear fits estimated separately on each side of the cutoff. Point estimates and standard errors at the top left of each panel are from the corresponding OLS regressions of the relevant alcohol consumption measure on age (centered at 21), an age-21+ indicator, and their interaction (estimated with individual-level data). The sample includes all NLSY97 respondents in the age range who have 1 or more younger siblings and are the oldest siblings in their household.

Figure 2: Smoothness of Covariates at Cutoff

(a) Predicted Count of Drinking Days

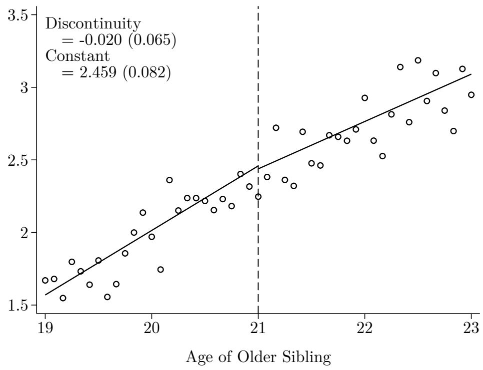

(b) Predicted Prob. of Any Drinking

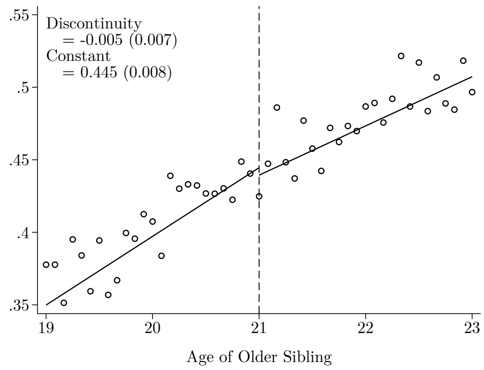

(c) Predicted Count of Binge Drinking Days

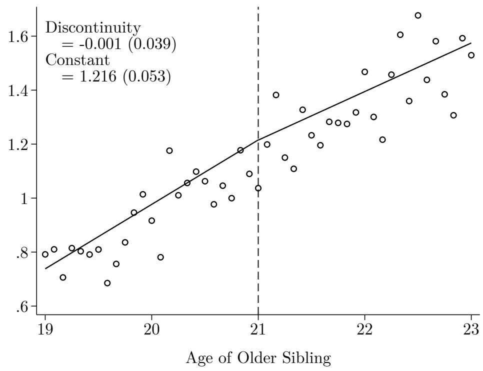

(d) Predicted Prob. of Any Binge Drinking

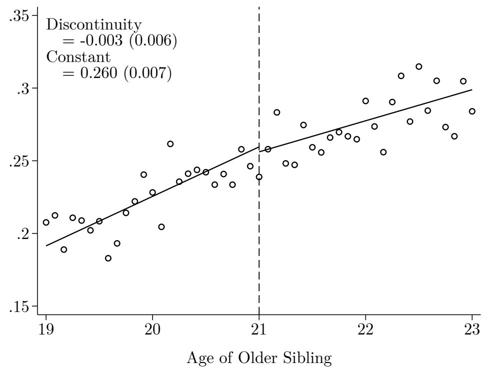

Note: Each panel shows mean predicted alcohol consumption in each older-sibling-age bin (in months), with linear fits estimated separately on each side of the cutoff (using the binned means). Predicted alcohol consumption measures are fitted values from a regression of the corresponding alcohol consumption measure on age, educational enrollment, highest completed education, work status, indicators for whether the household lives in an urban area, census region dummies, household size, and interview month/year. (All variables refer to the younger sibling's information unless otherwise noted.) Point estimates and standard errors at the top left of each panel are from corresponding OLS regressions of the prediction on the running variable, the cutoff indicator, and their interaction (estimated with individual-level data). Standard errors are calculated via bootstrap. The sample includes all NLSY97 respondents in the age range who are the second oldest siblings in their household.

Figure 3: Discontinuity in Younger Sibling Alcohol Consumption

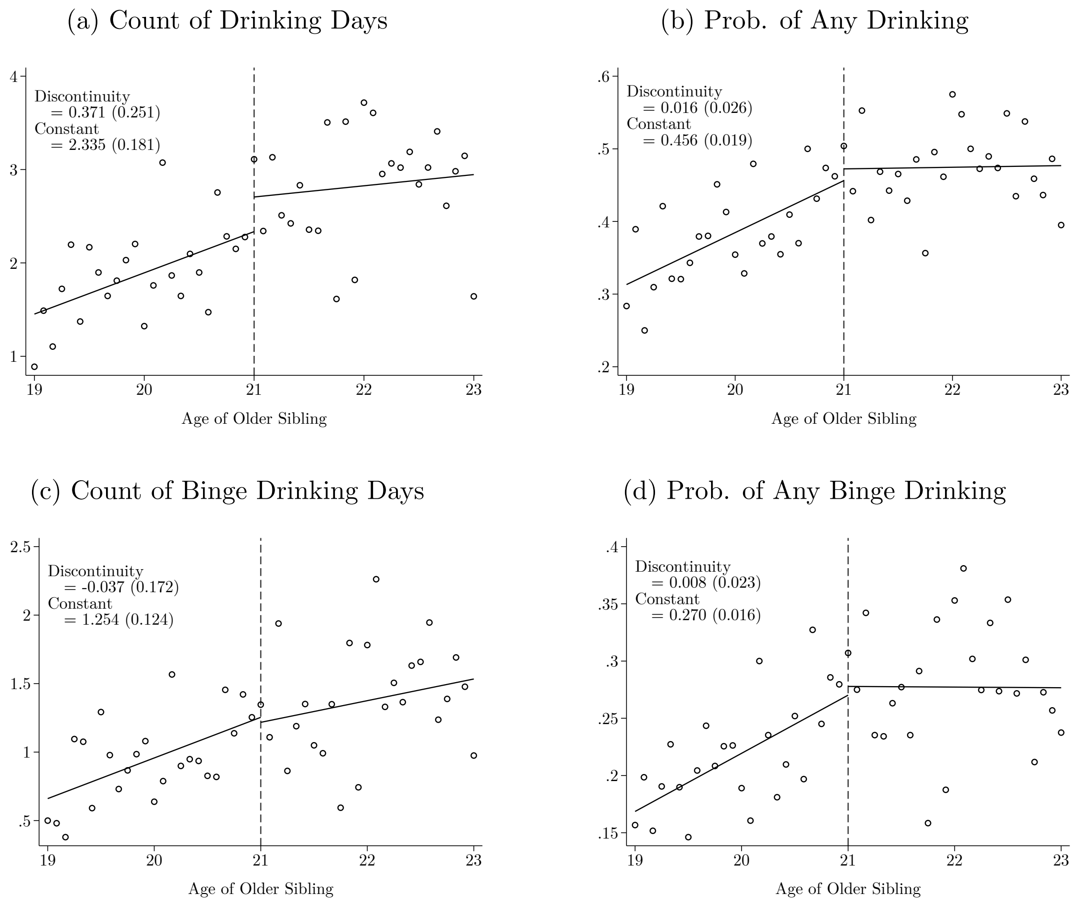

Note: Each panel shows mean alcohol consumption in each older-sibling-age bin (in months), with linear fits estimated separately on each side of the cutoff. Point estimates and standard errors at the top left of each panel are from corresponding OLS regressions of the younger sibling's alcohol consumption on the running variable, the cutoff indicator, and their interaction (estimated with individual level data). The sample includes all NLSY97 respondents in the age range who are the second oldest siblings in their household.

Figure 4: Placebo Discontinuities

(a) Count of Binge Drinking Days (b) Prob. of Any Binge Drinking

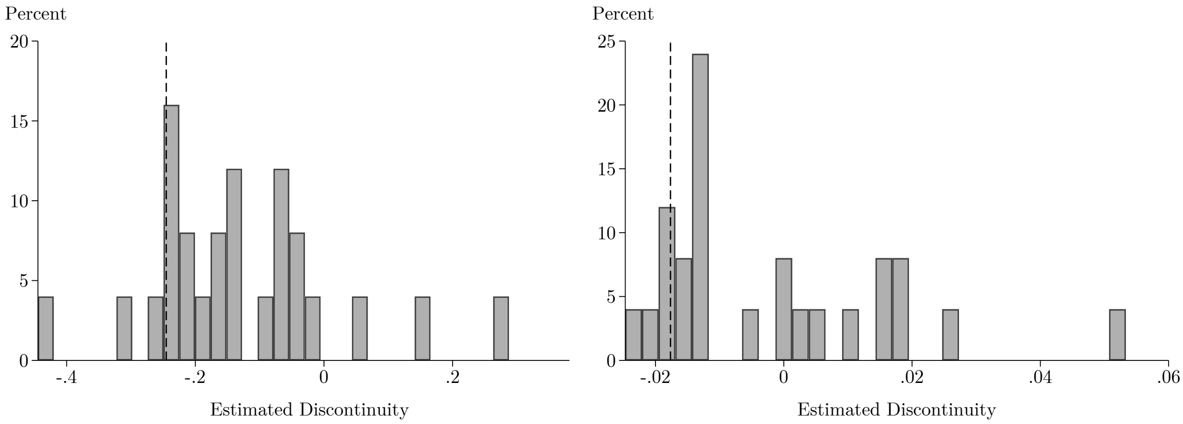

Note: Figure shows the distribution of point estimates from a series of placebo discontinuity estimates, at running variable values near but not equal to the cutoff. Each underlying regression uses the preferred specification in Table 3 (donut, linear, controls, FE) and a different cutoff age (in months) which varies from one year below to one year above age 21. The actual estimate is shown with a vertical dashed line.

Figure 5: Comparison with Eisenberg et al. (2014)

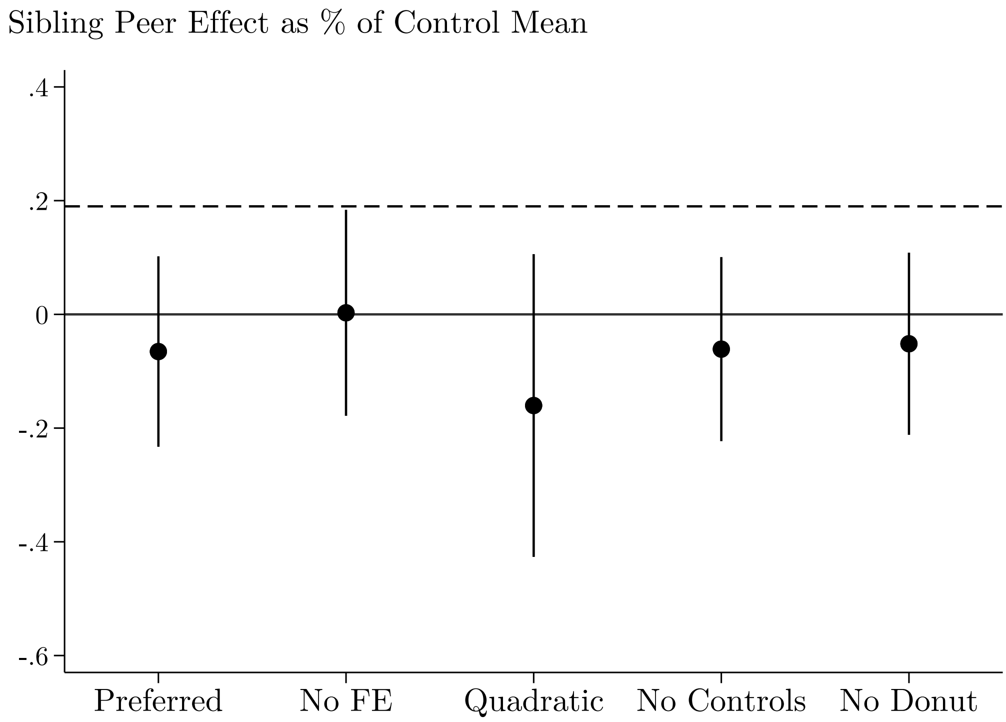

Note: Figure shows point estimates and confidence intervals from various specifications in Table 3 alongside the main estimate from EGW (dashed horizontal line). The outcome in each regression (including EGW's) is an indicator for any binge drinking days in the past month.

Figure 6: First Stage vs. Reduced Form by Subgroup, Count of Binge Drinking Days

(a) Unweighted

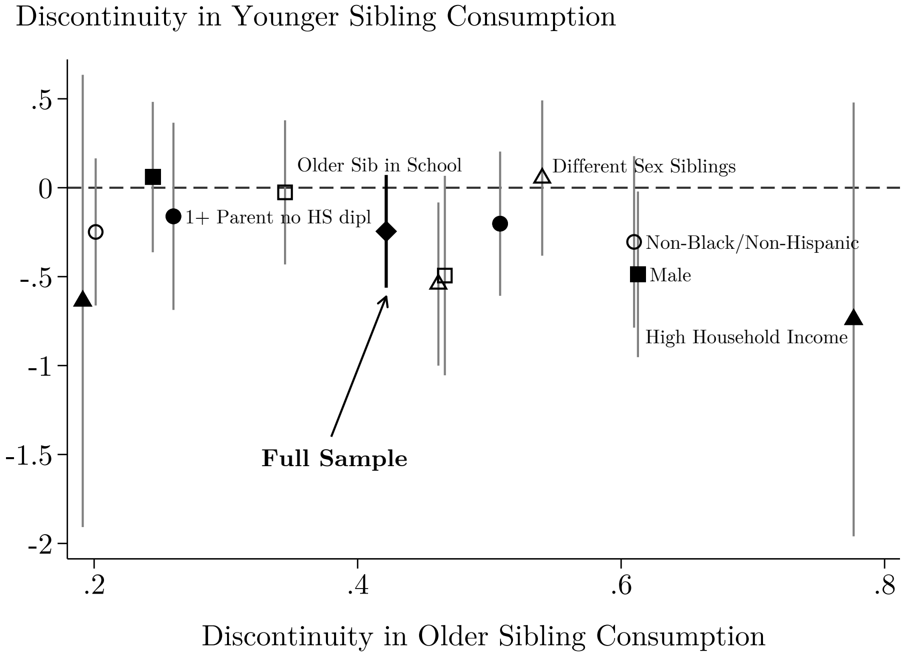

(b) Inverse Propensity Score Weighted

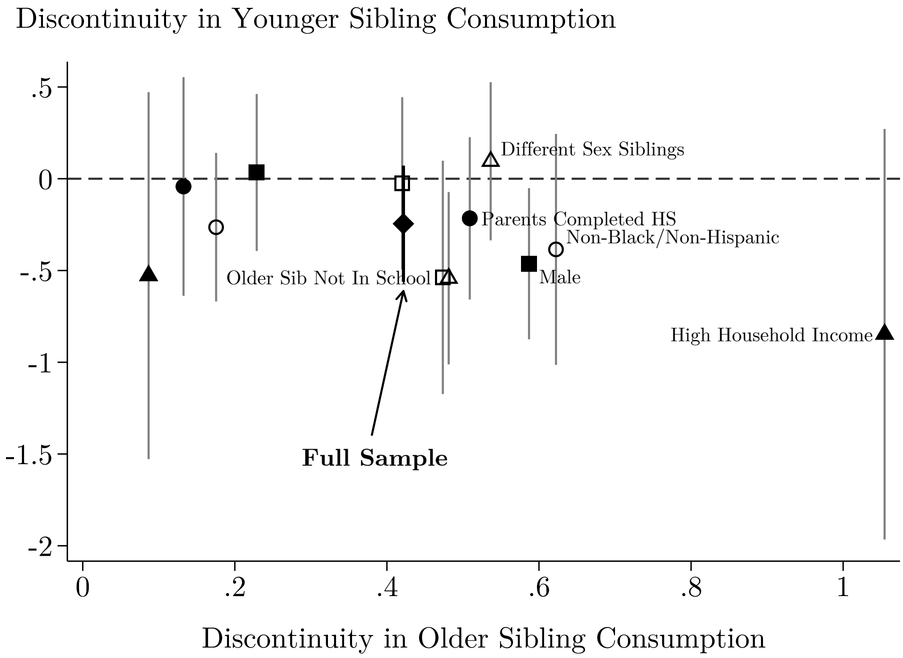

Note: Panel (a) plots the subgroup-level first stage effects (estimates from Table 4) against the reduced form effects for the same subgroups (estimates from Table 5). Panel (b) repeats this for the inverse propensity score weighted subgroup level estimates from Table A6 and Table A7. Whiskers mark 95% confidence intervals. Only one subgroup in each pair is labeled, e.g., the labeled hollow square in the top panel corresponds to sibling pairs where the older sibling is not enrolled in school, and the unlabled hollow square in the top panel corresponds to the sibling pairs where the older sibling is enrolled in school. The same symbols are used for all subgroups in both panels. The black diamonds and bolded CIs display estimates from the preferred specification in the full sample (Table 2 and Table 3).

Table 1: NLSY97 Sample Summary Statistics

	Full Sample mean/sd	RDD Sample mean/sd
Younger Sibling's Past Month Drinking:		
Drinking Days	2.41	2.31
	(4.98)	(4.76)
Binge Days	1.02	1.14
	(3.03)	(3.25)
Any Drinking Days	0.41	0.42
Any Binge Days	0.23	0.24
Younger Sibling's Characteristics:		
Age	18.72	17.99
	(3.97)	(1.72)
Household Income	$63,484.53	$61,559.45
	(58,841.17)	(55,247.96)
AFQT score	42,842.18	44,616.99
	(28,720.84)	(29,150.67)
Female	0.47	0.48
Race: Black	0.28	0.26
Race: Hispanic	0.26	0.24
In High School	0.53	0.55
In College	0.15	0.20
Urban	0.78	0.77
Worked Past Yr	0.66	0.71
Older Sibling's Characteristics:		
Age	22.15	20.87
	(5.29)	(1.17)
Female	0.47	0.47
Worked Past Yr	0.38	0.39
Currently Enrolled In School	0.53	0.50
N	15,230	5,702

Note: The full sample consists of all NLSY respondents with older siblings living in the same household. Our analysis sample ("RDD") consists of co-resident siblings who are the oldest two siblings in the household, where the younger sibling is a NLSY97 respondent and the older sibling is between the ages of 19 and 23.

Table 2: Discontinuity in Older Sibling Alcohol Consumption

	Count of Drinking Days				
	(1)	(2)	(3)	(4)	(5)
Age 21+	1.357***	1.579***	1.683***	1.339***	1.221***
	(0.139)	(0.169)	(0.218)	(0.133)	(0.133)
Control Mean	3.642	3.642	3.642	3.642	3.642
Observations	20,239	16,682	20,239	21,376	20,680
	Any Drinking Days				
Age 21+	0.081***	0.089***	0.096***	0.083***	0.075***
	(0.011)	(0.013)	(0.018)	(0.011)	(0.011)
Control Mean	0.578	0.578	0.578	0.578	0.578
Observations	20,239	16,682	20,239	21,376	20,680
	Count of Binge Drinking Days				
Age 21+	0.422***	0.582***	0.588***	0.404***	0.383***
	(0.091)	(0.111)	(0.144)	(0.089)	(0.086)
Control Mean	1.642	1.642	1.642	1.642	1.642
Observations	20,134	16,601	20,134	21,268	20,571
	Any Binge Drinking Days				
Age 21+	0.056***	0.067***	0.063***	0.058***	0.051***
	(0.011)	(0.013)	(0.018)	(0.011)	(0.011)
Control Mean	0.338	0.338	0.338	0.338	0.338
Observations	20,134	16,601	20,134	21,268	20,571
FE	X		X	X	X
Quadratic			X		
Controls	X	X	X		X
Donut	X	X	X	X	

Note: Controls include the month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year. All models are estimated in a sample of NLSY97 respondents who are the oldest siblings currently residing in their household and are between the ages of 19 and 23. All models include cluster robust standard errors at individual level. Age is centered at 21 years. +, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively.

Table 3: Discontinuity in Younger Sibling Alcohol Consumption

	(1)	(2)	(3)	(4)	(5)
	Count of Drinking Days				
Sibling 21+	0.083	0.101	0.087	0.156	0.186
	(0.233)	(0.249)	(0.396)	(0.229)	(0.226)
Control Mean	2.335	2.335	2.335	2.335	2.335
Observations	5,225	4,282	5,225	5,527	5,344
	Any Drinking Days				
Sibling 21+	-0.000	-0.009	-0.024	0.004	0.000
	(0.026)	(0.028)	(0.042)	(0.025)	(0.024)
Control Mean	0.456	0.456	0.456	0.456	0.456
Observations	5,225	4,282	5,225	5,527	5,344
	Count of Binge Drinking Days				
Sibling 21+	-0.245	-0.175	-0.197	-0.195	-0.177
	(0.162)	(0.169)	(0.266)	(0.161)	(0.155)
Control Mean	1.254	1.254	1.254	1.254	1.254
Observations	5,220	4,278	5,220	5,521	5,339
	Any Binge Drinking Days				
Sibling 21+	-0.018	0.001	-0.043	-0.017	-0.014
	(0.023)	(0.025)	(0.037)	(0.022)	(0.022)
Control Mean	0.270	0.270	0.270	0.270	0.270
Observations	5,220	4,278	5,220	5,521	5,339
FE	X		X	X	X
Quadratic			X		
Controls	X	X	X		X
Donut	X	X	X	X	

Note: Controls include the month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year. All models are estimated in a sample of NLSY97 respondents who have exactly one older sibling in their household, and where that older sibling is between the ages of 19 and 23. All models include cluster robust standard errors at individual level. Sibling age is centered at 21 years. +, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively.

Table 4: Discontinuity in Older Sibling Binge Drinking Days by Subgroup

	Parental Education		Household Income	
	< HS	≥ HS	< Median	≥ Median
Age 21+	0.260^{+}	0.508***	0.167	0.884***
	(0.151)	(0.144)	(0.133)	(0.211)
Control Mean	1.120	1.992	1.524	1.895
Observations	5,241	9,333	10,359	6,104
	Older Sibling's School Enrollment		Older Sibling's Sex	
	Enrolled	Not Enrolled	Male	Female
Age 21+	0.345*	0.466***	0.613***	0.245*
	(0.143)	(0.126)	(0.157)	(0.096)
Control Mean	1.729	1.590	2.360	0.930
Observations	7,626	12,508	9,894	10,240
	Race		Sibling Sex Composition	
	Black or Hispanic	Non-Black & Non-Hispanic	Same	Different
Age 21+	0.201^{+}	0.610***	0.461*	0.540*
	(0.117)	(0.136)	(0.207)	(0.219)
Control Mean	1.055	2.136	1.626	1.757
Observations	9,268	10,866	4,257	4,381

Note: Controls include the month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year. All models are estimated in the corresponding subgroup of a sample of NLSY97 respondents who are the oldest siblings currently residing in their household who are between the ages of 19 and 23. The median household income is $51,500. All models include cluster robust standard errors at individual level. Age is centered at 21 years. +, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively.

Table 5: Discontinuity in Younger Sibling Binge Drinking Days by Subgroup

	Parental Education		Household Income	
	< HS	≥ HS	< Median	≥ Median
Sibling 21+	-0.161	-0.202	-0.637	-0.740
	(0.268)	(0.207)	(0.648)	(0.621)
Control Mean	1.128	1.477	0.958	2.321
Observations	1,396	2,566	1,061	1,062
	Older Sibling's School Enrollment		Older Sibling's Sex	
	Enrolled	Not Enrolled	Male	Female
Sibling 21+	-0.026	-0.494^{+}	-0.487*	0.060
	(0.207)	(0.286)	(0.237)	(0.215)
Control Mean	1.259	1.214	1.344	1.143
Observations	2,585	2,608	2,795	2,425
	Race		Sibling Sex Composition	
	Black or Hispanic	Non-Black & Non-Hispanic	Same	Different
Sibling 21+	-0.249	-0.305	-0.542*	0.054
	(0.211)	(0.246)	(0.234)	(0.222)
Control Mean	0.903	1.613	1.348	1.138
Observations	2,608	2,612	2,848	2,372

Note: Controls include the month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year. All models are estimated in the corresponding subgroup of a sample of NLSY97 respondents who have exactly one older sibling in their household, and where that older sibling is between the ages of 19 and 23. The median household income is $53,515 in this sample. All models include cluster robust standard errors at individual level. Sibling age is centered at 21 years. +, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively.

Table 6: Alternative Samples

	Drinking Days	Any Drinking Days	Binge Drinking Days	Any Binge Drinking Days
	Two Oldest Siblings Living in the Same Household			
Sibling 21+	0.083	-0.000	-0.245	-0.018
	(0.233)	(0.026)	(0.162)	(0.023)
Control Mean	2.335	0.456	1.254	0.270
Observations	5,225	5,225	5,220	5,220
	Two Siblings Living in the Same Household			
Sibling 21+	0.190	-0.003	-0.170	-0.019
	(0.210)	(0.023)	(0.143)	(0.021)
Control Mean	2.305	0.447	1.215	0.267
Observations	6,298	6,298	6,290	6,290
	Two Oldest Siblings			
Sibling 21+	-0.104	-0.003	-0.164	-0.023
	(0.173)	(0.018)	(0.115)	(0.017)
Control Mean	2.472	0.472	1.271	0.290
Observations	8,439	8,439	8,425	8,425
	Two Siblings			
Sibling 21+	0.028	-0.013	-0.055	-0.018
	(0.147)	(0.015)	(0.092)	(0.014)
Control Mean	2.384	0.456	1.184	0.274
Observations	12,125	12,125	12,106	12,106

Note: The first panel is estimated in a sample of NLSY97 respondents who have exactly one older sibling in their household, and where that older sibling is between the ages of 19 and 23. The second panel is estimated in a sample of NLSY97 respondents who are the second oldest siblings, where their older sibling is between the ages of 19 and 23, and the older sibling may not currently reside in the same household as the respondent. The third panel is estimated in a sample of NLSY97 respondents who have one or more older siblings in their household where the closest older sibling to the respondent is between the ages of 19 and 23. The last panel is estimated in a sample of NLSY respondents who have one or more older siblings where the closest older sibling to the respondent is between the ages of 19 and 23, and the closest older sibling may not currently reside in the same household as the respondent. Controls include the month and year of the survey, the age of the respondent, educational attainment and enrollment of respondents, geography (urban/rural, census region), household size, an indicator for whether the respondent has children, and an indicator for whether the respondent worked in the past year. All models include cluster robust standard errors at individual level. Sibling age is centered at 21 years. +, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively.