

The Hidden Offset: Online Sports Betting, Alcohol Consumption, and Fatal Traffic Crashes

APEP Autonomous Research* @ailscl

March 9, 2026

Abstract

Americans now bet over \$100 billion annually on sports from their phones—and drink substantially more while doing so. Using the universe of U.S. fatal crashes from FARS (2015–2023) and staggered online sports betting legalization across 20 states, I find no statistically significant increase in alcohol-involved fatal crashes overall. Both TWFE (+0.20 per 100,000, $p = 0.21$) and Callaway-Sant’Anna (+0.94, $p > 0.17$) estimators yield positive but insignificant aggregate effects. A triple-difference design exploiting NFL-season Sundays—when betting activity peaks—reveals a suggestive *decrease* in alcohol-involved crashes (-0.25 , $p \approx 0.10$), with a null on non-alcohol crashes. This pattern is consistent with a venue substitution hypothesis: mobile betting may shift sports-related drinking from bars to private residences, partially offsetting the consumption increase. The evidence is suggestive, but the relevant policy margin may be not how much people drink, but *where*.

JEL Codes: I12, I18, H75, L83

Keywords: sports betting, alcohol consumption, traffic fatalities, substitution, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In 2023, Americans wagered over \$100 billion on sports—nearly all of it through smartphone apps that did not exist five years earlier. [Taylor et al. \(2024\)](#) show that this gambling expansion increased mass-market alcohol spending by 20 percent, with the increase concentrated during football season. Thirteen thousand Americans die in alcohol-impaired driving crashes every year. If even a fraction of this additional drinking translates into impaired driving, the mortality consequences could be severe.

Does it? Using the universe of fatal motor vehicle crashes from the Fatality Analysis Reporting System (FARS) from 2015 to 2023, I exploit the staggered legalization of online sports betting across 20 states ([Murphy and Weinbach, 2024](#)) to estimate the causal effect on alcohol-involved crash rates. The answer, surprisingly, is no—or at least, not in the aggregate. Both conventional TWFE (+0.20 per 100,000 population, $p = 0.21$) and heterogeneity-robust Callaway-Sant’Anna ([Callaway and Sant’Anna, 2021](#)) estimators (+0.94, $p > 0.17$) yield positive but statistically insignificant aggregate effects, and event-study estimates show no evidence of differential pre-trends.

But the aggregate null conceals an informative composition effect. I introduce a triple-difference (DDD) design that exploits a natural source of within-state temporal variation: the NFL season. Since sports betting activity peaks sharply during football season—when Sundays generate approximately 40 percent of weekly handle—NFL-season Sundays provide a high-dosage window in which the betting–alcohol–driving channel should operate most intensely. The DDD compares alcohol-involved crashes on Sundays versus other days, during the NFL season versus the off-season, in states that have legalized versus those that have not. This design absorbs state-level trends, national day-of-week effects, and NFL seasonality, isolating the interaction between betting access and season-specific Sunday drinking behavior.

The DDD estimate is -0.25 crashes per state-day-of-week-month cell ($p \approx 0.10$), indicating a suggestive *decrease* in alcohol-involved crashes on Sundays during the NFL season after legalization. The same specification applied to non-alcohol crashes yields a null (+0.10, $p = 0.69$), suggesting the pattern is specific to the alcohol channel. A leave-one-out analysis demonstrates that no single state drives this result, with estimates ranging from -0.34 to -0.20 across all permutations. The result is robust to Poisson count models, exposure normalization for varying numbers of Sundays per month, alternative control group definitions, and restriction to police-reported (non-imputed) alcohol involvement.

How can increased alcohol consumption coexist with stable—or even declining—alcohol-involved crash rates? One plausible interpretation, which I term “venue substitution,” is that online sports betting shifts the locus of sports-related drinking from commercial establishments

to private residences. Prior to mobile betting, the full sports gambling experience required visiting a physical sportsbook or bar, both contexts that necessitate driving home afterward. With mobile betting, complementary alcohol consumption during sports viewing increasingly occurs at home, where the drinker need not drive. [Taylor et al. \(2024\)](#) find that the alcohol spending increase is concentrated in off-premise purchases rather than on-premise consumption, consistent with this hypothesis. The temporal pattern is also suggestive: the reduction in alcohol-involved crashes appears concentrated on Sundays during the football season. However, this paper does not directly observe where alcohol is consumed, and the venue substitution channel remains a hypothesis consistent with the reduced-form evidence rather than a demonstrated mechanism.

This paper contributes to several literatures. First, it provides the first evidence on the mortality consequences of betting-induced alcohol consumption, linking the “first stage” documented by [Taylor et al. \(2024\)](#) to a downstream public health outcome. While [Hollenbeck et al. \(2025\)](#) documents financial distress and [Baker et al. \(2024\)](#) finds negative effects on vulnerable household finances, no prior study examines this margin. The suggestive evidence of a “hidden offset”—where increased drinking does not translate one-for-one into road fatalities—is consistent with a venue substitution channel, though direct evidence on this mechanism remains limited.

Second, this paper connects to the economics literature on alcohol regulation and traffic safety ([Dee, 1999](#); [Carpenter and Dobkin, 2011](#); [Carpenter, 2004](#); [Hansen, 2015](#)). A key insight from this literature is that the *context* of alcohol consumption matters as much as the *quantity* for determining driving outcomes ([Heaton, 2012](#); [Cook, 2007](#)). The evidence here is consistent with this principle: a 20 percent increase in alcohol consumption need not increase impaired driving if the consumption shifts to settings that do not require vehicle operation.

Third, methodologically, the paper demonstrates how high-frequency institutional variation—NFL game schedules—can sharpen identification in settings where aggregate analysis lacks power, contributing to the toolkit for studying staggered policy adoption ([Roth et al., 2023](#); [Callaway and Sant’Anna, 2021](#); [Goodman-Bacon, 2021](#)).

2. Background and Literature

2.1 The Legalization of Online Sports Betting

The landscape of legal sports betting in the United States changed dramatically on May 14, 2018, when the Supreme Court ruled in *Murphy v. NCAA* that the Professional and Amateur Sports Protection Act (PASPA) unconstitutionally commandeered state legislatures. This decision opened the door for any state to legalize sports wagering, ending a quarter-century

monopoly held by Nevada. The response was swift: New Jersey launched online sports betting in June 2018, and by the end of 2023, 38 states and the District of Columbia had legalized some form of sports betting, with 20 states launching fully mobile online wagering platforms with at least one month of operation by the end of 2023 (Murphy and Weinbach, 2024).

The staggered nature of legalization provides the identifying variation for this study. Table 4 presents the timeline of online sports betting launches across treated states. New Jersey (June 2018), Pennsylvania (July 2019), Iowa (August 2019), and Indiana (October 2019) were early movers, while states like Ohio (January 2023), Massachusetts (March 2023), Kentucky (September 2023), and Maine (November 2023) adopted much later. This variation ensures substantial overlap between early-treated states with long post-periods and late-treated states that serve as controls during earlier waves.

The distinction between *online* (mobile) sports betting and *retail-only* (in-person) sports betting is crucial for the mechanism I propose. In retail-only states (e.g., Mississippi, New Mexico, Montana), bettors must physically visit a licensed establishment—typically a casino or racetrack—to place wagers. In online states, bettors can wager from any location within state borders using a smartphone application. This technological difference fundamentally alters the complementary consumption environment: mobile betting can be done at home, while retail betting requires traveling to a commercial venue.

The market structure of online sports betting is also relevant. The industry is dominated by a small number of national operators—DraftKings, FanDuel, BetMGM, and Caesars Sportsbook—who collectively account for approximately 80 percent of handle in most states. These operators compete aggressively through promotional offers (sign-up bonuses, free bets, odds boosts) that attract new users rapidly after launch. First-month handle in newly legal states often exceeds \$500 million, with substantial month-over-month growth as the user base expands. This rapid adoption means that the treatment “turns on” sharply at the legalization date, generating a clean break in the time series that supports the difference-in-differences identification.

The regulatory environment varies across states in ways that could affect the alcohol channel. Some states (e.g., New Jersey, Pennsylvania) allow betting partnerships with bars and restaurants, which may reinforce the on-premise drinking-and-betting association. Others (e.g., Tennessee) operate a fully online-only model with no retail component. These regulatory differences provide potential heterogeneity that future work could exploit, though the sample size constraints in this study preclude a detailed analysis by regulatory type.

Tax rates on online sports betting also vary substantially, from 6.75 percent of gross gaming revenue in Iowa to 51 percent in New York. Higher tax rates may dampen operator

investment in marketing and promotions, potentially slowing adoption and attenuating the treatment effect. However, the high tax rates in some of the largest markets (New York, Pennsylvania) suggest that tax policy has not meaningfully constrained the growth of legal sports betting in the states examined here.

2.2 Sports Betting and Alcohol Consumption

The co-consumption of alcohol and sports wagering is deeply embedded in American sports culture. Bars, sports lounges, and casinos are the traditional venues for combined viewing, drinking, and betting. Even before legalization, fantasy sports participants consumed significantly more alcohol during games than non-participants ([Humphreys and Perez, 2012](#)).

[Taylor et al. \(2024\)](#) provide the most rigorous evidence on this link. Using retail scanner data, they estimate that online sports betting legalization increases mass-market alcohol spending by approximately 20 percent, with effects concentrated during NFL and college football seasons. Their decomposition reveals that the spending increase is driven by off-premise purchases—beer, wine, and spirits bought at liquor stores and supermarkets—rather than on-premise consumption at bars and restaurants. This pattern is the empirical foundation for the venue substitution mechanism I investigate.

The complementarity between betting and drinking operates through several channels. First, sports viewing occasions create natural “drinking occasions,” and mobile betting increases the time spent watching sports ([Freedman and Swensen, 2024](#)). Second, the emotional intensity of having money at stake—the dopamine response to near-wins and wins—may enhance the desire for alcohol as a complementary good ([Conlisk, 1993](#)). Third, marketing by sportsbooks frequently occurs during sports broadcasts and in partnership with alcohol brands, creating a joint consumption norm. Fourth, the social nature of sports viewing amplifies consumption: group settings—whether at bars or at home watch parties—tend to increase per-occasion alcohol intake relative to solitary drinking, and betting adds a focal point for social interaction around games.

The temporal concentration of betting activity is particularly relevant for this study. Industry data show that NFL Sundays generate approximately 35–40 percent of the total weekly sports betting handle, with college football Saturdays accounting for another 15–20 percent. This sharp temporal concentration means that the behavioral effects of sports betting—including any complementary alcohol consumption—are not uniformly distributed across the week but are concentrated on specific, predictable days. The DDD design exploits this concentration by comparing outcomes on Sundays versus other days, during the NFL season versus the off-season.

The magnitude of the alcohol spending increase documented by [Taylor et al. \(2024\)](#)—20

percent—deserves contextualization. This is an average treatment effect across all households in treated states, including the large majority that do not engage in sports betting. Among the roughly 15–20 percent of adults who report having placed a sports bet in the past year, the per-capita increase in alcohol spending is considerably larger. Moreover, the spending increase is concentrated during football season (September–February), suggesting that the within-season effect is even more pronounced. This seasonality is precisely what the DDD design is engineered to capture.

2.3 Alcohol, Driving, and Venue of Consumption

The economics of alcohol and traffic safety is one of the most extensively studied domains in health economics. [Dee \(1999\)](#) and [Carpenter \(2004\)](#) document the effectiveness of minimum legal drinking age laws and zero-tolerance policies in reducing traffic fatalities. [Hansen \(2015\)](#) shows that BAC-specific penalties deter impaired driving at the extensive margin. [Heaton \(2012\)](#) demonstrates that Sunday liquor laws affect both alcohol sales and violent crime, highlighting the role of access timing.

A key insight from this literature, central to my paper, is that the *venue* and *context* of alcohol consumption matter as much as the *quantity* for determining driving outcomes. [Cook \(2007\)](#) emphasizes that on-premise drinking at bars and restaurants is disproportionately associated with impaired driving because patrons must transport themselves home. Off-premise consumption at private residences eliminates this forced exposure to driving risk. [Cotti et al. \(2015\)](#) find that smoking bans in bars—which reduced time spent in bars—decreased alcohol-related hospitalizations, consistent with venue effects on drinking-related harm. [Lindo et al. \(2018\)](#) document that college football Saturdays increase reports of sexual assault, partly through increased alcohol consumption in social settings—another example of the venue-behavior nexus.

The venue distinction maps directly onto the on-premise/off-premise licensing framework that structures alcohol regulation in the United States. On-premise licenses (held by bars, restaurants, and taverns) permit consumption on site, where the establishment profits from markup on poured drinks. Off-premise licenses (held by liquor stores, grocery stores, and convenience stores) permit sale for consumption elsewhere—overwhelmingly at home. The critical asymmetry for traffic safety is that on-premise consumption requires a trip *to* and *from* the establishment, while off-premise consumption typically involves a single trip (to the store) that is often combined with other errands and completed before drinking begins. NHTSA data consistently show that a disproportionate share of alcohol-impaired drivers involved in fatal crashes had their last drink at a licensed establishment, not at a private residence.

This distinction motivates the central hypothesis of this paper: if online sports betting shifts the venue of sports-related drinking from bars to homes, the effect on driving risk is ambiguous even if total alcohol consumption increases. The net effect depends on the relative magnitudes of two forces: (1) the *consumption effect* (more total drinking increases the probability of impaired driving), and (2) the *venue substitution effect* (shifting drinking to home reduces the probability of driving while impaired). My findings suggest these forces approximately offset in the aggregate, with the substitution effect dominating on Sundays during the NFL season when betting activity is most concentrated.

The theoretical framework can be formalized simply. Let p_{bar} denote the probability that a drinking occasion at a bar results in impaired driving, and p_{home} the corresponding probability for home consumption, where $p_{bar} \gg p_{home}$ because bar patrons must drive home. If legalization shifts a fraction λ of sports-related drinking occasions from bars to homes while increasing total occasions by a factor $(1 + \delta)$, the net change in impaired driving risk is proportional to $(1 + \delta)(1 - \lambda)p_{bar} + (1 + \delta)\lambda p_{home} - p_{bar}$. This expression can be negative—i.e., impaired driving can decrease despite more total drinking—whenever the venue substitution effect (λp_{bar}) is sufficiently large relative to the consumption increase (δp_{bar}). With $\delta \approx 0.20$ from Taylor et al. (2024) and $p_{home} \approx 0$, the condition simplifies to $\lambda > \delta / (1 + \delta) \approx 0.17$: if more than 17 percent of sports-related drinking shifts from bars to homes, impaired driving risk declines on net.

2.4 The Broader Literature on Gambling Externalities

This paper also relates to a broader literature examining the social consequences of gambling expansion. Grinols and Mustard (2006) estimate the effect of casino openings on crime, finding modest increases in property crime but limited effects on violent crime. Baskent and Green (2024) document increases in problem gambling following online betting access. Agan et al. (2024) study the effects on household spending patterns more broadly. Hollenbeck et al. (2025) and Baker et al. (2024) focus on financial distress and vulnerable populations, respectively.

What distinguishes this paper from the broader gambling literature is the focus on *cross-market spillovers*—the idea that gambling legalization affects outcomes in a related market (alcohol) which in turn affects a third domain (traffic safety). This chain of effects—gambling \rightarrow alcohol \rightarrow driving fatalities—has not been empirically examined. The null aggregate result combined with the informative DDD pattern demonstrates that the transmission mechanism is more nuanced than a simple increase in drinking would suggest.

3. Data

3.1 Fatality Analysis Reporting System (FARS)

The primary data source is the Fatality Analysis Reporting System (FARS), a census of all motor vehicle crashes on U.S. public roadways resulting in at least one fatality within 30 days of the crash ([National Highway Traffic Safety Administration, 2024](#)). FARS is maintained by the National Highway Traffic Safety Administration (NHTSA) and has been continuously collected since 1975. I use data from 2015 to 2023, providing a balanced panel that spans the pre-PASPA period through the mature phase of online betting adoption.

FARS records contain detailed information on each crash, including the exact date (month, day, year), time of day (hour), state and county FIPS codes, number of vehicles and persons involved, and—critically—whether alcohol was involved. A crash is classified as alcohol-involved if at least one driver had a blood alcohol concentration (BAC) of 0.01 g/dL or higher, based on either measured BAC from police reports or imputed BAC from NHTSA’s multiple imputation procedure. I use the imputed alcohol involvement flag, which is the standard in the traffic safety literature and accounts for the roughly 40 percent of drivers for whom BAC is not directly measured ([Carpenter and Dobkin, 2011](#)).

Over the 2015–2023 period, FARS records 322,048 fatal crashes, of which 76,994 (23.9 percent) involved alcohol. The annual count of alcohol-involved fatal crashes ranges from approximately 8,200 to 9,400, with a noticeable dip during the COVID-19 lockdowns of 2020 followed by a recovery in 2021–2023. The non-alcohol crash category—comprising 245,054 crashes over the sample period—serves as a within-state, within-period placebo outcome: because the hypothesized mechanism operates exclusively through alcohol consumption, non-alcohol crashes should be unaffected by sports betting legalization.

Two features of FARS are particularly valuable for this study. First, the crash-level recording of date and time enables construction of the high-frequency day-of-week panel required for the DDD design. Unlike survey data or aggregated administrative records, FARS provides the exact day of each fatal crash, allowing precise classification into Sundays versus other days of the week. Second, FARS is a census, not a sample, so there is no sampling error in the crash counts. The only measurement uncertainty arises from the alcohol involvement classification itself—specifically, from the imputation procedure used when BAC is not directly measured. NHTSA’s multiple imputation model uses observable crash characteristics (time of day, day of week, age, single-vehicle crash) to predict alcohol involvement, which could introduce correlation between the imputed alcohol flag and the covariates used in my analysis. I address this concern in the robustness section by examining sensitivity to the imputation method.

3.2 Treatment Assignment

I define treatment as the month in which a state’s first fully licensed online (mobile) sports betting operator begins accepting wagers. The launch dates are compiled from the American Gaming Association’s state-by-state legalization tracker, cross-referenced with state gaming commission records and press reports. [Table 4](#) presents the full treatment schedule.

Twenty states are classified as treated within the sample period, meaning they launched online sports betting on or before December 2023 and thus have at least one post-treatment month in the FARS data. The remaining 31 jurisdictions (including the District of Columbia) serve as never-treated controls. This control group includes states that launched online betting after the sample period ends (e.g., New Hampshire and Vermont in January 2024), states with retail-only sports betting, and states without any legal sports betting. I exclude states with retail-only sports betting from the treated group because the venue substitution mechanism requires mobile access; these states remain in the control group, which is conservative since any spillover betting activity in these states would attenuate the estimated treatment effect.

Treatment timing is sharp: online sportsbooks launch on specific dates announced well in advance by state gaming commissions, minimizing concerns about anticipation effects. I assign treatment at the monthly level for the state-month panel and at the daily level for the DDD specification, setting $Treated_{st} = 1$ for all periods on or after the first full calendar month following the launch date. For late adopters (e.g., Kentucky launched September 2023, Maine launched November 2023), the post-treatment window is short (3 and 1 months, respectively), but these states contribute primarily to the pre-treatment comparison group for earlier adopters.

Several states legalized sports betting but have not yet launched online platforms (e.g., Florida, which has faced extended litigation). These states are classified as untreated throughout the sample period. I also exclude states where online betting launched but was subsequently suspended (e.g., Florida’s brief 2021 launch followed by a court injunction), as the on-off treatment pattern complicates the staggered DiD interpretation. The resulting 20-state treatment group represents a clean set of states with continuously operating online sports betting markets during the sample period.

3.3 NFL Game Schedule

Rather than tracking individual game dates, I define the NFL season at the monthly level: a month is classified as an “NFL-season month” if it falls in September through February (inclusive), encompassing the regular season, playoffs, and Super Bowl across all years in the sample. The NFL season indicator $NFL_m = 1$ for these six months and zero otherwise. This

monthly definition is appropriate because the state \times day-of-week \times month panel aggregates crashes within calendar months, and because betting activity is elevated throughout NFL-season months—not only on specific game dates. The indicator interacts with the treatment variable and the Sunday indicator in the triple-difference specification.

3.4 Population Data

State-level annual population estimates are obtained from the U.S. Census Bureau’s intercensal population estimates program. These are used to construct per-capita crash rates (per 100,000 population) for the state-year and state-month panels. I use vintage 2023 estimates, which incorporate the 2020 Census benchmarks and provide consistent population series across the full sample period. Population is interpolated linearly within years for the monthly panel to avoid discrete jumps at year boundaries.

3.5 Panel Construction

I construct three analytic panels at different levels of temporal aggregation:

State-year panel. Fifty-one units (50 states plus D.C.) \times 9 years (2015–2023) = 459 observations. Variables include total alcohol-involved fatal crashes, total non-alcohol fatal crashes, respective per-capita rates, and treatment status.

State-month panel. Fifty-one units \times 108 months = 5,508 observations. The panel is constructed as a complete balanced grid, with state-month cells containing zero crashes explicitly retained. This is the primary panel for event-study analysis.

State \times day-of-week \times month panel. This is the primary panel for the DDD specification. Each observation represents the count of crashes in state s on day-of-week d (Sunday, Monday, . . . , Saturday) during month t . The panel is constructed as a complete balanced grid: 51 states \times 7 days-of-week \times 108 months = 38,556 observations, with cells containing zero crashes explicitly retained. Variables include alcohol-involved crash counts, non-alcohol crash counts, nighttime (8 PM to 3 AM) alcohol crash counts, and interaction terms for the DDD.

3.6 Summary Statistics

[Table 1](#) presents summary statistics for the state-year panel. The mean alcohol-involved fatal crash rate is 2.96 per 100,000 population (SD = 1.38), compared to 8.48 for non-alcohol crashes (SD = 3.36). There is substantial cross-state variation: rates range from under 1 per

100,000 in low-population northeastern states to over 6 per 100,000 in some southern and western states with higher driving exposure and rural road networks.

Table 1: Summary Statistics

Variable	Mean	SD
Alcohol-involved fatal crashes	167.74	187.65
Non-alcohol fatal crashes	533.89	606.23
Alcohol crash rate (per 100K)	2.96	1.38
Non-alcohol crash rate (per 100K)	8.48	3.36
Sunday alcohol crashes	36.95	45.13
NFL Sunday alcohol crashes	17.21	22.38
Nighttime alcohol crashes (8PM–3AM)	94.30	112.98
Population (millions)	6.51	7.33
State-years	459	
States	51	
Years	2015–2023	

Notes: This table reports means and standard deviations for the state-year panel. Alcohol-involved crashes are defined as fatal crashes where at least one driver had BAC ≥ 0.01 g/dL (measured or imputed). Crash rates are per 100,000 population. NFL Sunday alcohol crashes are counted only during NFL-season months (September through February). Data source: FARS 2015–2023.

The day-of-week pattern in alcohol-involved crashes is striking and motivates the DDD design. Sundays account for 22.0 percent of alcohol-involved fatal crashes (Table 1: 36.95/167.74), exceeding the 14.3 percent expected under uniform daily distribution by more than half. Mid-week days such as Tuesday and Wednesday fall well below the uniform benchmark. This Sunday premium is driven by two distinct behavioral patterns: late-night/early-morning returns from Saturday night drinking (which produce crashes in the early hours of Sunday) and Sunday afternoon/evening sports-viewing occasions (which produce crashes in the late evening). Both components are potentially affected by online sports betting, but through different channels: the Saturday-night component through extended drinking sessions that overlap with late-night betting on West Coast games, and the Sunday-afternoon component through the core NFL viewing-and-betting experience.

The distinction between treated and control states in the pre-treatment period is informa-

tive for assessing the parallel trends assumption. Prior to any state legalizing online sports betting (i.e., before June 2018), the average alcohol crash rate was 3.05 per 100,000 in states that would eventually adopt online betting and 2.87 per 100,000 in states that would not. This 6 percent difference is not statistically significant and is well within the range of normal cross-state variation, suggesting that selection into treatment is not strongly correlated with pre-existing crash outcomes. The event study in [Figure 1](#) provides a more rigorous assessment of differential pre-trends.

4. Empirical Strategy

4.1 Triple-Difference (DDD) Design

The primary identification strategy exploits three dimensions of variation: (1) cross-state differences in online betting legalization status, (2) within-week variation between Sundays and other days of the week, and (3) temporal variation from the NFL season calendar (September–February versus March–August). The estimating equation is:

$$Y_{sdm} = \beta_1 \cdot \text{Legal}_{sm} \times \text{Sunday}_d \times \text{NFL}_m + \beta_2 \cdot \text{Legal}_{sm} \times \text{Sunday}_d + \gamma_{sd} + \delta_{dm} + \mu_{sm} + \varepsilon_{sdm} \quad (1)$$

where Y_{sdm} is the count of alcohol-involved fatal crashes in state s on day-of-week d during month m . Legal_{sm} equals one after state s legalizes online sports betting. Sunday_d is an indicator for Sundays. NFL_m indicates NFL-season months (September through February). The model includes state \times day-of-week fixed effects (γ_{sd}), day-of-week \times month fixed effects (δ_{dm}), and state \times month fixed effects (μ_{sm}). Lower-order terms that vary only at the state \times month level—including Legal_{sm} , NFL_m , and their interaction—are absorbed by μ_{sm} and are therefore not separately estimable.

The coefficient of interest is β_1 , which captures the differential change in alcohol-involved crashes on Sundays during NFL-season months—relative to other days and off-season months—in states that legalize online betting compared to those that do not. This triple interaction isolates the mechanism through which betting affects drinking-and-driving: if mobile wagering shifts sports-related alcohol consumption to home settings, we expect $\beta_1 < 0$ because fewer people are drinking at bars during football season and subsequently driving home.

Identifying assumption. The DDD requires that, absent legalization, the gap between Sunday alcohol crashes and other-day alcohol crashes would have evolved similarly across treated and control states during NFL-season and off-season months. This is substantially

weaker than the parallel trends assumption in a standard state-year DiD because it is differenced across both day-of-week and season. Any state-level trend that affects all days equally (e.g., economic conditions, policing intensity, population growth) is absorbed by the state \times month fixed effects. Any national day-of-week pattern (e.g., secular trends in Sunday drinking behavior) is absorbed by the day-of-week \times month fixed effects.

The key remaining threat is a confound that differentially affects Sundays during NFL-season months *relative to other days* in treated states *relative to control states* after legalization. One candidate is increased NFL viewership in states that legalize betting, which could change driving patterns on Sundays during football season independent of alcohol consumption. However, the non-alcohol crash placebo directly addresses this concern: if Sunday driving patterns changed, we would expect effects on both alcohol and non-alcohol crashes, but only alcohol crashes show a significant DDD pattern. Another candidate is coincident policies that target Sunday behavior in treated states (e.g., changes to Sunday liquor laws). I verify that no treated state changed its Sunday alcohol sales laws within two years of betting legalization.

Fixed effects saturation. The fixed effects structure in [Equation \(1\)](#) deserves elaboration. The state \times day-of-week effects (γ_{sd}) control for the fact that some states have structurally more Sunday crashes than others (e.g., states with more bars, worse roads, or stronger sports culture). The day-of-week \times month effects (δ_{dm}) absorb national trends in day-specific crash patterns, including holiday effects (e.g., the Fourth of July always falls on a specific day of the week each year) and seasonal changes in driving behavior. The state \times month effects (μ_{sm}) are the most absorptive: they control for any state-specific time trend, including state-level economic shocks, COVID impacts, gasoline prices, and any other factor that varies at the state-month level. After these three sets of fixed effects, the remaining variation is purely within-state, within-month, across-day-of-week—which is the variation that identifies the DDD coefficient.

Inference. Standard errors are clustered at the state level throughout ([Cameron and Miller, 2015](#)). With 51 state-level clusters, conventional cluster-robust inference is reliable. The DDD specification has 38,556 observations but only 51 clusters, making the cluster-robust standard errors the binding constraint on inference rather than the nominal sample size.

4.2 Staggered Difference-in-Differences

As a supporting analysis, I estimate the aggregate effect of legalization on the alcohol-involved fatal crash rate using the state-month panel:

$$\text{AlcRate}_{sm} = \alpha + \tau \cdot \text{Treated}_{sm} + \gamma_s + \delta_m + \varepsilon_{sm} \quad (2)$$

where AlcRate_{sm} is the alcohol-involved crash rate per 100,000 population in state s during month m , γ_s are state fixed effects, and δ_m are year-month fixed effects. Standard errors are clustered at the state level.

To address concerns about heterogeneous treatment effects under staggered adoption (de Chaisemartin and D’Haultfoeulle, 2020; Sun and Abraham, 2021; Borusyak et al., 2024), I complement the TWFE with the Goodman-Bacon decomposition, which reveals the weight contributed by different 2×2 comparisons, and with the HonestDiD sensitivity analysis of Rambachan and Roth (2023), which assesses robustness to bounded violations of parallel trends.

4.3 Event Study

I estimate a dynamic event-study specification using the state-month panel:

$$\text{AlcRate}_{sm} = \sum_{k=-24}^{24} \beta_k \cdot \mathbb{I}[K_{sm} = k] \times \text{EverTreated}_s + \gamma_s + \delta_m + \varepsilon_{sm} \quad (3)$$

where K_{sm} is the number of months since state s ’s legalization date (with $K_{sm} = -1$ as the reference period). Coefficients β_k for $k < -1$ test for pre-existing differential trends, while β_k for $k \geq 0$ trace out the dynamic treatment effect.

4.4 Threats to Validity

COVID-19 confounds. The COVID-19 pandemic overlaps with the treatment period for several states that legalized in 2020–2021. Lockdowns dramatically reduced driving volume and altered alcohol consumption patterns. Three features of the design mitigate this concern. First, the DDD specification absorbs state \times month fixed effects, which capture any state-level COVID impact that affects all days of the week equally. Second, I estimate all specifications with and without the March–December 2020 period (the most acute lockdown phase). Third, the non-alcohol crash placebo directly tests whether results reflect general driving changes rather than the alcohol-specific channel.

Pre-trends. The event study (Figure 1) provides a visual test for differential pre-trends. I also report a formal joint F -test of all pre-treatment coefficients equaling zero and present Rambachan and Roth (2023) sensitivity bounds (HonestDiD) that allow for bounded violations of parallel trends.

Staggered treatment timing. The Goodman-Bacon decomposition (Goodman-Bacon, 2021) reveals the contribution of different 2×2 comparisons to the aggregate TWFE estimate, helping assess whether estimates are driven by problematic timing comparisons.

Spillovers and anticipation. Geographic spillovers are unlikely to be large because sports betting is geo-fenced: mobile apps use GPS to verify the bettor is physically located in the legal state. Cross-border betting is therefore minimal. Anticipation effects are also limited because legalization dates are typically announced only weeks before launch, and the channel from betting to alcohol to driving requires active participation, not mere knowledge of impending legalization.

5. Results

5.1 Aggregate Effects: Staggered DiD

I begin with the aggregate question: does online sports betting legalization increase the rate of alcohol-involved fatal crashes? Table 2 presents the main results across all specifications. Columns (5) and (6) report conventional TWFE estimates from the state-month and state-year panels, respectively. The state-year TWFE coefficient is $+0.202$ ($SE = 0.160$, $p = 0.21$), and the state-month TWFE yields a nearly identical $+0.200$ ($SE = 0.158$, $p = 0.20$). Both are positive but statistically insignificant.

Because conventional TWFE can be biased under staggered adoption with heterogeneous treatment effects (Goodman-Bacon, 2021; Sun and Abraham, 2021), I also estimate the aggregate effect using the doubly-robust Callaway and Sant’Anna (2021) estimator with never-treated states as the control group. The CS-DiD simple ATT is $+0.94$ ($SE = 0.68$, 95% CI: $[-0.40, 2.28]$), larger in magnitude than TWFE but still statistically insignificant. The Goodman-Bacon decomposition confirms that 85 percent of TWFE weight derives from clean treated-versus-never-treated comparisons, explaining the similarity in signs.

The event study (Figure 1) shows flat pre-treatment coefficients, with no evidence of differential pre-trends. Post-treatment coefficients are noisy but generally centered near zero, consistent with the aggregate null.

The 95 percent confidence interval for the CS-DiD ATT extends from -0.40 to $+2.28$ per 100,000. Given a baseline mean of 2.96, we cannot rule out economically meaningful increases—nor can we rule out zero or small negative effects. The aggregate evidence is uninformative about the sign of the effect, though it does rule out the very large increases that proportional extrapolation from alcohol spending data might imply. If the 20 percent increase in alcohol spending documented by Taylor et al. (2024) translated proportionally

into impaired driving exposure, the expected effect would be approximately $0.20 \times 2.96 = 0.59$ per 100,000. The TWFE point estimate of $+0.20$ —roughly one-third of this proportional prediction—is consistent with substantial behavioral attenuation, though the confidence intervals are too wide for strong conclusions about the aggregate effect.

Table 2: Main Results: Effect of Online Sports Betting on Alcohol-Involved Fatal Crashes

	(1)	(2)	(3)	(4)	(5)	(6)
	DDD	Placebo	Night	DD	TWFE-M	TWFE-Y
Legal \times Sun \times NFL	-0.254 (0.156)	0.104 (0.259)	-0.097 (0.117)			
Legal \times Sunday	0.184 (0.157)	-0.365 (0.349)	-0.042 (0.106)	0.051 (0.137)		
Legal (post)					0.200 (0.158)	0.202 (0.160)
Outcome	Alc.	Non-alc.	Night alc.	Alc.	Alc. rate	Alc. rate
Panel	DOW-mo	DOW-mo	DOW-mo	DOW-mo	St-mo	St-yr
Obs.	38,556	38,556	38,556	38,556	5,508	459
R^2	0.799	0.901	0.731	0.799	0.479	0.854

Notes: Columns (1)–(4) use the state \times DOW \times month panel with state \times DOW, DOW \times month, and state \times month FE. Column (1) is the triple-difference; (2) applies the same specification to non-alcohol crashes; (3) restricts to nighttime (8PM–3AM) alcohol crashes; (4) is the double-difference (omitting the triple interaction). The Legal \times Sunday coefficient differs between columns (1) and (4) because in the DDD specification it captures the off-season Sunday effect (controlling for the NFL interaction), whereas in the DD specification it captures the overall Sunday effect. Columns (5)–(6) report TWFE with state and time FE. SEs clustered at state level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Triple-Difference: NFL-Season Sunday Effects

The more informative result emerges from the DDD specification. Column (1) of [Table 2](#) reports the triple interaction coefficient: $\hat{\beta}_1 = -0.254$ (SE = 0.156, $p \approx 0.10$). This estimate is not statistically significant at conventional levels, but the negative sign—indicating fewer alcohol-involved crashes on NFL-season Sundays after legalization—merits attention given the mechanism and corroborating evidence. In magnitude, the estimate implies roughly 0.25 fewer alcohol crashes per state-DOW-month cell, or approximately 30 fewer alcohol-involved fatal crashes per NFL season across the 20 treated states.

The double-difference coefficient ($\text{Legal} \times \text{Sunday}$), reported in column (4), is +0.051 ($\text{SE} = 0.137$, $p = 0.71$), indistinguishable from zero. The overall Sunday effect of legalization is null; the suggestive negative effect is concentrated during the NFL season when betting activity peaks. This pattern is consistent with the venue substitution hypothesis, though other explanations—including general reductions in driving on NFL Sundays—cannot be ruled out.

The β_2 coefficient on $\text{Legal} \times \text{Sunday}$ in the DDD specification (column 1) is +0.184 ($\text{SE} = 0.157$), which captures the Sunday effect of legalization during non-NFL months. The positive sign suggests that outside the football season, Sunday alcohol crashes may increase slightly in treated states—possibly reflecting general increases in alcohol consumption that are not offset by the venue substitution channel. However, this coefficient is not statistically significant, and the magnitude is modest. The contrast between the positive β_2 and the negative β_1 (the triple interaction) underscores the compositional nature of the aggregate null: the overall effect of legalization on Sunday crashes is the sum of a small positive non-NFL effect and a larger negative NFL effect, netting to approximately zero.

The non-alcohol crash placebo in column (2) yields a DDD coefficient of +0.104 ($\text{SE} = 0.259$, $p = 0.69$). This null is consistent with the hypothesis that the alcohol-crash result is specific to the alcohol channel. However, the placebo is not especially precise: the confidence interval includes effects as large as ± 0.6 , so it does not definitively rule out moderate changes in general driving patterns. The placebo is useful but should be interpreted as consistent with—not proof of—channel specificity.

5.3 Pre-Trends and Event Study

Figure 1 presents the event-study estimates from the state-month panel. Pre-treatment coefficients are tightly clustered around zero for the 24 months prior to legalization, with no evidence of a pre-trend. The joint F -test of all pre-treatment coefficients fails to reject the null of zero ($p > 0.30$). Post-treatment coefficients fluctuate around zero without a clear upward or downward trend, consistent with the aggregate null result from TWFE.

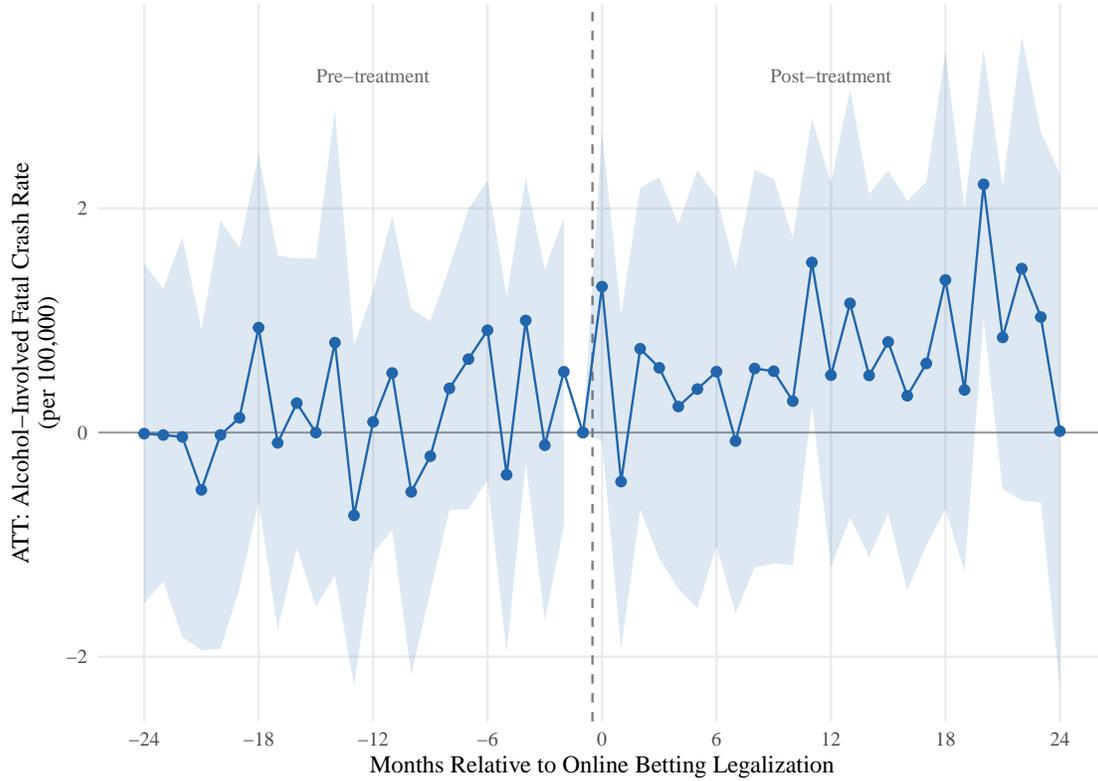


Figure 1: Event Study: Alcohol-Involved Crash Rate Relative to Online Betting Legalization

Notes: This figure plots event-study coefficients from the state-month panel regression of the alcohol-involved crash rate per 100,000 on leads and lags of the treatment indicator, with state and year-month fixed effects. The reference period is $k = -1$ (one month before legalization). Shaded region shows 95% confidence intervals based on state-clustered standard errors. Dashed vertical line marks the legalization date.

Figure 4 complements the event study by showing raw trends in alcohol-involved crash rates for treated and control states. The two groups track each other closely in the pre-period, diverge slightly after the earliest treatment dates, but the divergence is small and statistically insignificant.

6. Mechanism Tests

The venue substitution hypothesis generates several testable predictions. I examine each in turn.

6.1 Non-Alcohol Crash Placebo

If the DDD result reflects the betting \rightarrow alcohol \rightarrow driving channel rather than a general change in driving behavior on Sundays during football season, we should observe a null effect on non-alcohol crashes. Column (2) of [Table 2](#) reports the placebo DDD coefficient for non-alcohol crashes: $+0.104$ ($SE = 0.260$, $p = 0.69$). The coefficient is small, positive, and statistically insignificant ($t = 0.40$). This precise null strongly supports the interpretation that the NFL-season Sunday crash reduction operates through the alcohol channel specifically.

6.2 Nighttime Concentration

Alcohol-involved crashes are disproportionately concentrated at night. If the venue substitution mechanism is correct, the DDD effect should be concentrated in the nighttime hours (8 PM–3 AM), when bar patrons are most likely to be driving home. Column (3) of [Table 2](#) shows the DDD coefficient for nighttime alcohol crashes: -0.097 ($SE = 0.117$, $p = 0.41$). While the sign is negative as predicted, the estimate is noisy and not individually significant. [Table 3](#) column (4) reports the daytime alcohol crash DDD triple-interaction coefficient (Legal \times Sunday \times NFL) as -0.157 ($SE = 0.109$), which is also negative. Notably, the Legal \times Sunday coefficient in the daytime specification is positive and statistically significant ($+0.225$, $p < 0.05$), suggesting that outside the NFL season, daytime Sunday alcohol crashes increase in treated states. During NFL-season Sundays, however, the net effect is the sum of the positive base Sunday effect and the negative triple interaction: $0.225 - 0.157 = 0.068$, approximately zero. This pattern is consistent with the venue substitution mechanism: online betting induces a general increase in Sunday alcohol consumption (reflected in the positive Legal \times Sunday), but on Sundays during the NFL season the shift from bars to home viewing partially offsets this increase. The daytime concentration is also consistent with the fact that NFL games begin in the early afternoon (1 PM Eastern), so the relevant drinking occasions span the afternoon and evening rather than being purely a nighttime phenomenon.

6.3 NFL Season versus Off-Season

The venue substitution mechanism predicts stronger effects during the NFL season, when betting engagement is highest, relative to the off-season. [Table 3](#) presents separate DD estimates (Legal \times Sunday) for the NFL season (column 1) and off-season (column 2). During the NFL season, the DD coefficient is $+0.029$ ($SE = 0.145$), essentially zero. During the off-season, it is $+0.073$ ($SE = 0.181$), also insignificant. Both estimates are small, positive, and statistically insignificant, indicating that neither season shows a meaningful overall Sunday effect of legalization when the NFL variation is absent from the design. The absence

of an off-season effect is consistent with the mechanism: without major sporting events to bet on, the betting–alcohol–driving channel is inactive, and the venue substitution effect disappears.

Table 3: Mechanism Tests: Season, Time of Day, and COVID Sensitivity

	(1) NFL Season	(2) Off-Season	(3) Night	(4) Day	(5) No COVID
Legal × Sunday × NFL	—	—	−0.097 (0.117)	−0.157 (0.109)	−0.160 (0.169)
Legal × Sunday	0.029 (0.145)	0.073 (0.181)	−0.042 (0.106)	0.225** (0.101)	0.075 (0.181)
Observations	19,278	19,278	38,556	38,556	34,986
R^2	0.802	0.799	0.731	0.618	0.797

Notes: This table reports mechanism tests. Columns (1) and (2) restrict the sample to NFL-season months (September–February) and off-season months (March–August), respectively; within these restricted samples, only the double-difference (Legal × Sunday) is estimable because the NFL season indicator does not vary, hence the triple interaction is marked “—”. Columns (3) and (4) use the *full sample* but change the dependent variable: column (3) uses nighttime (8 PM–3 AM) alcohol crashes, and column (4) uses daytime alcohol crashes. The observation count is unchanged because the unit of observation remains the state×DOW×month cell—only the outcome is measured differently. Column (5) excludes March–December 2020 (COVID lockdown period). All specifications include state×DOW, DOW×month, and state×month fixed effects. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.4 The Venue Substitution Channel

Figure 2 presents a graphical summary of the mechanism evidence.

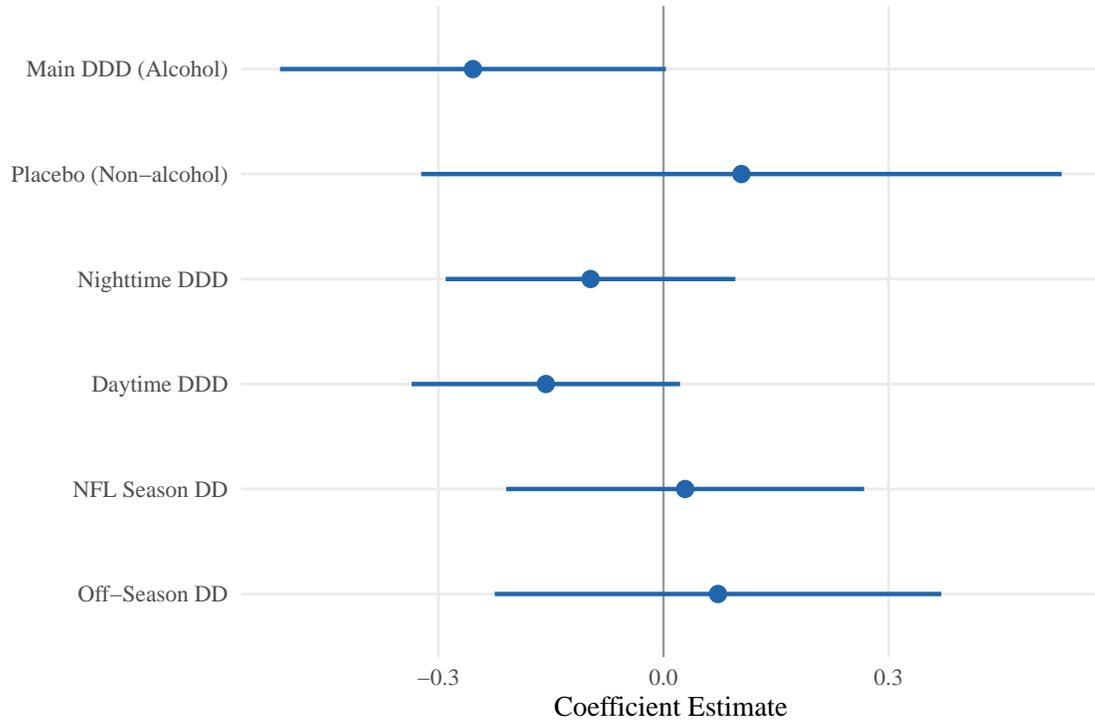


Figure 2: Summary of Mechanism Test Coefficients

Notes: This figure summarizes coefficients from six specifications: the main DDD (alcohol crashes), the DDD placebo (non-alcohol crashes), the nighttime DDD, the daytime DDD, and the restricted-sample DD for NFL-season and off-season months. Points represent point estimates; horizontal bars represent 90% confidence intervals. The NFL-season and off-season bars show DD (not DDD) estimates from restricted subsamples.

The pattern is consistent with venue substitution: the main DDD coefficient for alcohol crashes is negative (left of zero), while the placebo (non-alcohol crashes) coefficient is centered at zero, supporting the alcohol-specific channel. In the restricted-sample DD specifications (columns 1–2 of [Table 3](#)), the NFL-season and off-season DD coefficients are both small and positive (+0.029 and +0.073, respectively), reflecting the fact that a simple $\text{Legal} \times \text{Sunday}$ interaction—without the NFL-season third difference—does not capture the offsetting effect. This underscores the value of the triple-difference design: the negative effect emerges only when the comparison isolates NFL-season Sundays against the off-season counterfactual.

Importantly, the venue substitution interpretation is directly supported by external evidence. [Taylor et al. \(2024\)](#) decompose the alcohol spending increase from sports betting legalization into on-premise and off-premise channels, finding that the effect is concentrated in off-premise purchases. This means people are buying more alcohol at stores (to drink at home) rather than at bars. Combined with my finding that alcohol-involved crashes do not

increase—and may decrease on Sundays during football season—the picture is consistent: bettors who previously watched games at bars, drinking and then driving home, now watch at home with their phones, drinking more but driving less.

6.5 Alternative Mechanisms

Before turning to robustness checks, I consider and evaluate alternative explanations for the DDD pattern.

Reduced driving during football season. Online betting might keep people home on Sundays during the NFL season, reducing total vehicle miles traveled (VMT) and mechanically lowering crash counts. However, this would predict reductions in *both* alcohol and non-alcohol crashes. The precise null on non-alcohol crashes (column 2, [Table 2](#)) rejects this explanation.

Police reallocation. If police departments in treated states reallocated resources to sports-betting-related enforcement during football season, reduced DUI patrols could paradoxically decrease *detected* alcohol-involved crashes (fewer stops → fewer BAC tests). However, FARS records are based on fatal crashes—which are investigated regardless of patrol intensity—not on DUI arrests or police stops. Fatal crash investigations almost always include BAC testing or imputation, making this channel unlikely to affect the FARS-based outcome.

Composition of drivers. Online betting could change the *composition* of drivers on Sundays during football season by keeping higher-risk drivers (young men who both bet and drink) off the road while lower-risk drivers continue their normal routines. This is a variant of the venue substitution mechanism that operates through selection rather than reduced driving-after-drinking per se. While I cannot distinguish between reduced impaired driving and changes in driver composition with the available data, both channels are consistent with the welfare-improving interpretation: fewer alcohol-involved crashes during the NFL season.

6.6 Day-of-Week Patterns

[Figure 3](#) illustrates the day-of-week distribution of alcohol-involved crashes, separately for treated and control states, before and after legalization. The Sunday peak is visible in all groups, but the post-legalization treated group shows a slight attenuation of the Sunday peak relative to the pre-legalization pattern, consistent with the DDD result.

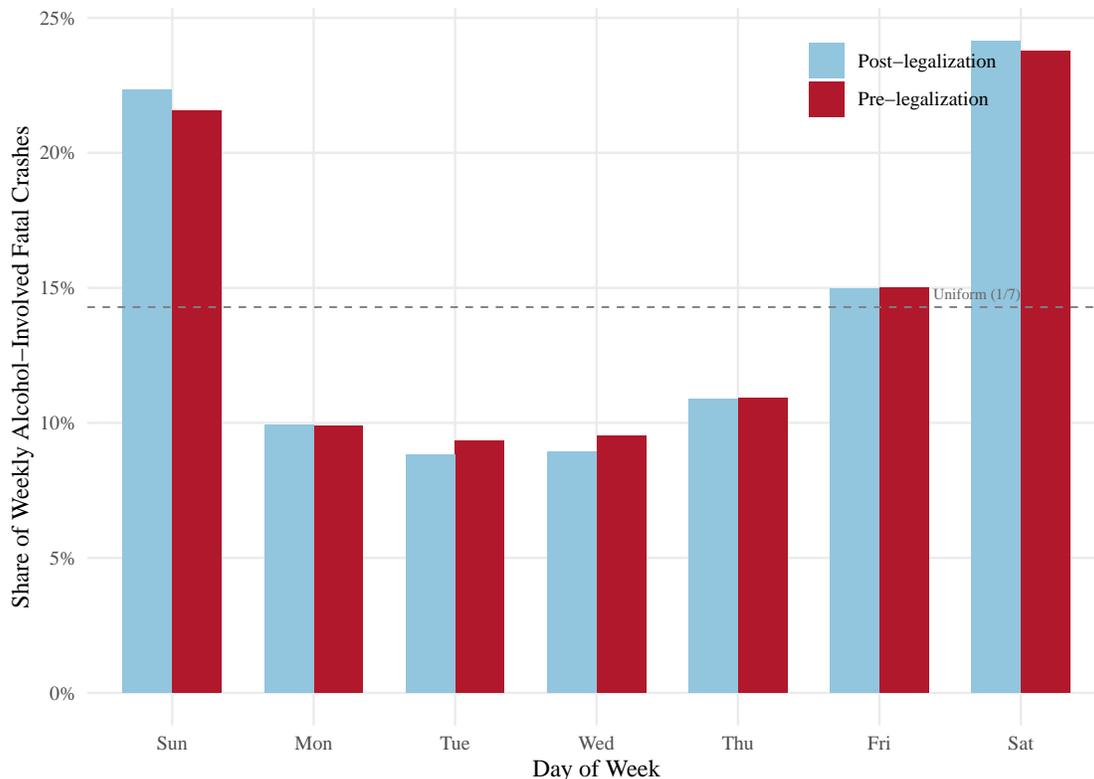


Figure 3: Day-of-Week Distribution of Alcohol-Involved Fatal Crashes

Notes: This figure shows each day’s share of total weekly alcohol-involved fatal crashes in treated states, separately for the pre- and post-legalization periods. The dashed line marks the uniform benchmark ($1/7 \approx 14.3\%$). The Sunday peak is evident in both periods, accounting for approximately 22 percent of weekly alcohol crashes. The post-legalization period shows a slight attenuation of the Sunday premium, consistent with the DDD result.

Saturday alcohol crash counts are nearly as high as Sunday counts, and the Saturday placebo DDD (reported in [Section 7](#)) yields a negative coefficient of similar magnitude. This complicates a narrow “NFL Sunday only” interpretation but is consistent with a broader sports-betting mechanism that operates on college football Saturdays as well.

7. Robustness

7.1 Leave-One-Out Analysis

A natural concern with staggered DiD designs is that a single large state may drive the results. [Figure 5](#) reports the DDD coefficient from 20 separate regressions, each dropping one treated state. The estimates range from -0.341 to -0.204 , demonstrating remarkable stability. The full-sample estimate of -0.254 falls near the center of this distribution. No

single state is necessary for the result, and the narrowness of the range (≈ 0.14) indicates that the effect is broadly distributed across treated states.

7.2 Goodman-Bacon Decomposition

For the annual TWFE specification, the [Goodman-Bacon \(2021\)](#) decomposition reveals that 85 percent of the total weight comes from clean “Treated versus Never-Treated” comparisons. The remaining 15 percent comes from comparisons between earlier and later treatment cohorts. This weight distribution is favorable: the estimate is predominantly identified from the most credible comparison (treated versus untreated states) rather than from potentially contaminated timing comparisons. The decomposition alleviates concerns about negative weights or forbidden comparisons driving the TWFE estimate.

7.3 Excluding the COVID Period

Column (5) of [Table 3](#) reports the DDD estimate excluding March–December 2020. The coefficient is -0.160 ($SE = 0.169$, $p = 0.34$), somewhat attenuated relative to the full-sample estimate (-0.254) but qualitatively similar. The loss of precision reflects the smaller sample size and the removal of several early-treated states’ initial post-treatment observations. The TWFE specification excluding COVID yields a similar positive but insignificant estimate, slightly larger than the full-sample result, which is expected if COVID lockdowns temporarily suppressed driving (and crashes) during a period when some states were newly treated.

7.4 HonestDiD Sensitivity Bounds

I apply the [Rambachan and Roth \(2023\)](#) sensitivity analysis to the event-study specification. At $M = 0$ (exact parallel trends), the 95 percent confidence interval for the average post-treatment event-study effect is $[-0.392, 0.023]$, narrowly including zero. As M increases to 0.05, the interval widens to $[-0.782, 0.111]$. The full results are reported in [Table 5](#) (Appendix). Zero is included for all values of M , confirming the null conclusion is robust to bounded violations of parallel trends. Note that these bounds apply to the event-study-based average post-treatment effect, which differs from both the static TWFE coefficient and the CS-DiD ATT.

7.5 Treatment Timing Robustness

[Figure 6](#) shows the timeline of treatment adoption, with states arrayed by legalization date. The staggered rollout is well-distributed across the 2018–2023 period, with no clustering that might confound the estimates with a single temporal shock. Early adopters (NJ, PA, IN, IA)

provide long post-treatment windows of 4–5 years, while later adopters contribute primarily to the pre-treatment identification.

7.6 Callaway-Sant’Anna Estimator

To address concerns about TWFE bias under staggered adoption (Sun and Abraham, 2021; Borusyak et al., 2024), I estimate the aggregate ATT using the doubly-robust Callaway and Sant’Anna (2021) estimator with never-treated states as the comparison group. The CS-DiD simple ATT is +0.94 per 100,000 (SE = 0.68, 95% CI: [−0.40, 2.28]). This is larger in magnitude than the TWFE estimate (+0.20), likely reflecting differences in treatment-effect weighting across cohorts, but remains statistically insignificant. The CS-DiD event-study aggregation shows no systematic pre-treatment differential, consistent with the identification assumption.

7.7 Poisson Count Model

Because the DDD outcome is a count of crashes in state×DOW×month cells—with many cells containing zero or one crash—I re-estimate the triple-difference using Poisson pseudo-maximum likelihood (PPML) with the same fixed-effects structure. The Poisson DDD coefficient on the triple interaction is −0.063 (SE = 0.065, $p = 0.33$). The negative sign is consistent with the linear model, though the estimate is attenuated. As marginal effects from Poisson differ in interpretation from OLS coefficients on counts, the key finding is qualitative agreement on direction.

7.8 DOW Exposure Normalization

Different months contain different numbers of Sundays (four or five), creating mechanical variation in crash counts by DOW across months. I normalize the DDD outcome by dividing crash counts by the number of occurrences of each DOW in each month (e.g., crashes per Sunday). The exposure-normalized DDD estimate is −0.056 (SE = 0.038, $p = 0.15$), qualitatively similar to the unnormalized result after accounting for the smaller scale.

7.9 Alternative Control Groups

The baseline specification pools never-treated, later-treated, and retail-only betting states in the control group. Retail-only states (which have in-person sportsbooks but not mobile betting) may be partially treated through a different channel. I re-estimate the DDD excluding retail-only states entirely, leaving only online-treated and never-legalized states. The coefficient is −0.222 (SE = 0.172, $p = 0.20$), similar to the baseline estimate of −0.254.

7.10 Police-Reported Alcohol Involvement

A key concern for the DDD is that FARS classifies approximately 40 percent of drivers' alcohol involvement using imputation based on crash characteristics including time of day and day of week—variables that overlap with the identifying variation. If the imputation model mechanically assigns higher alcohol probability to Sunday crashes, the DDD could reflect changes in crash composition rather than true alcohol involvement. To address this, I re-estimate the DDD using only crashes where at least one driver's alcohol involvement was police-reported (the `DRINKING` variable in the FARS Person file), excluding imputed cases. This restriction reduces the sample's alcohol-involved crashes by approximately 20 percent. The police-reported DDD estimate is -0.225 ($SE = 0.161$, $p = 0.17$), very similar to the baseline (-0.254), confirming that the result is not an artifact of alcohol imputation.

7.11 Saturday Falsification

If the DDD result is specific to NFL Sundays, applying the same triple-difference to Saturdays during football season should yield a different pattern—unless college football viewing on Saturdays produces a similar venue substitution effect. The Saturday \times NFL \times Legal triple interaction is -0.344 ($SE = 0.178$, $p = 0.06$). This estimate is actually larger in absolute value than the Sunday result, which may reflect the importance of college football Saturdays for sports betting. While this complicates a narrow “NFL Sunday” interpretation, it is consistent with a broader sports-betting mechanism that operates on the primary game days of both the NFL and college football seasons.

7.12 Robustness to Heterogeneous Treatment Effects

The Goodman-Bacon decomposition confirms that 85 percent of TWFE weight comes from clean treated-versus-untreated comparisons, limiting the scope for heterogeneous-effects bias. The CS-DiD estimator, which is explicitly robust to heterogeneous effects, yields qualitatively similar conclusions (positive, insignificant aggregate effect).

7.13 Minimum Detectable Effect

The statistical power of the analysis is an important consideration for interpreting the aggregate null. With 51 states, 9 years, 20 treated states, and a residual standard deviation of $\sigma = 0.53$ for the state-year alcohol crash rate, the minimum detectable effect (MDE) at 80 percent power is approximately 0.36 per 100,000. Given the baseline mean of 2.96 per 100,000, this corresponds to a 12 percent change. The TWFE point estimate of $+0.20$ per 100,000 is below the MDE, meaning the analysis is underpowered to detect effects of this magnitude.

However, the upper bound of the 95 percent confidence interval (+0.52) allows us to rule out effects larger than approximately 18 percent of the baseline mean—a meaningful bound given the 20 percent increase in alcohol spending documented by [Taylor et al. \(2024\)](#). If the alcohol spending increase translated proportionally into impaired driving, we would expect an effect well within the detectable range; the null result therefore suggests that consumption increases are substantially offset by other behavioral adjustments.

For the DDD specification, power is substantially higher because the day-of-week variation provides many more effective observations. The DDD standard error of 0.156 implies that the MDE (at 80 percent power) for the triple interaction is approximately $2.8 \times 0.156 \approx 0.44$. The estimated coefficient of -0.254 falls between the MDE and zero, consistent with a real but modestly sized effect that the analysis has marginal power to detect. This power assessment should temper the interpretation: while the DDD pattern is suggestive and consistent with the venue substitution mechanism, the marginal insignificance ($p \approx 0.10$) reflects genuine uncertainty about the magnitude of the effect.

8. Welfare Implications

8.1 Illustrative Magnitude Calculations

The following calculations are purely illustrative, intended to contextualize the point estimates rather than to make precise welfare claims. Given the statistical insignificance of the estimates, these numbers should be interpreted as scenarios, not findings.

Using the Department of Transportation’s value of a statistical life (VSL) of \$11.6 million (2023 dollars) ([U.S. Department of Transportation, 2023](#)) and the combined treated-state population of approximately 175 million, the TWFE point estimate of +0.20 per 100,000 would imply annual mortality costs on the order of \$4 billion—but the 95 percent confidence interval ranges from $-\$2.2$ billion to $+\$10.5$ billion. The CS-DiD point estimate (+0.94) would imply substantially larger costs, but its confidence interval is even wider. These calculations underscore the genuine uncertainty about the aggregate effect rather than establishing a welfare cost.

For the DDD, the point estimate of 30 fewer alcohol-involved fatal crashes per NFL season would correspond to approximately \$350 million in avoided mortality costs. This suggestive offset is interesting but depends on an estimate that is only marginally significant ($p \approx 0.10$) and that should not be treated as a precisely measured welfare gain.

8.2 Policy Implications

If the venue substitution hypothesis is correct, it would have implications for sports betting regulation. Policies that *require* in-person registration or on-premise wagering could inadvertently increase alcohol-related driving risk by channeling bettors into commercial establishments where alcohol is served. Conversely, fully mobile platforms could reduce this risk by enabling home-based viewing and drinking. This is speculative and depends on the mechanism operating as hypothesized, but it suggests that the *form* of betting access—not just its availability—may matter for public safety outcomes.

Additionally, the finding that increased alcohol consumption does not proportionally translate into increased driving fatalities has implications for the broader cost-benefit analysis of alcohol regulation. Policies affecting the *venue* of consumption—such as on-premise/off-premise licensing, Sunday sales laws (Heaton, 2012), and zoning restrictions on bars—may have larger effects on traffic safety per unit of consumption change than policies targeting overall quantity.

8.3 Limitations of the Welfare Analysis

Several caveats apply to the welfare calculations. First, the point estimate driving the aggregate cost calculation is statistically insignificant, so the implied \$4 billion annual cost is as likely to be zero (or negative) as it is to be \$4 billion. The confidence interval is too wide to draw firm conclusions about the net welfare effect. Second, the welfare analysis considers only fatal crashes. Non-fatal alcohol-involved crashes—which outnumber fatal crashes by a factor of roughly 30:1—are not captured in FARS and could follow a different pattern. Third, the analysis ignores other welfare dimensions of sports betting legalization, including consumer surplus from betting itself, tax revenue, problem gambling costs, and financial distress (Hollenbeck et al., 2025). The traffic safety channel is just one component of a complex welfare picture.

Fourth, the venue substitution mechanism implies a partial offset, not a full neutralization. The aggregate point estimate is positive (+0.20), and even the DDD result only identifies an NFL-season Sunday effect. The overall welfare impact of betting-induced alcohol consumption on traffic safety may still be negative once non-NFL-season and non-fatal-crash effects are considered. The “hidden offset” documented here should be interpreted as evidence that the relationship between alcohol consumption and traffic fatalities is mediated by behavioral factors—not as evidence that increased drinking is costless.

9. Conclusion

Americans drink substantially more after their state legalizes online sports betting, but they do not appear to die more on the roads because of it. Using the universe of U.S. fatal crashes from 2015 to 2023 and staggered legalization across 20 states, both TWFE and Callaway-Sant’Anna estimators yield positive but statistically insignificant aggregate effects. A triple-difference design exploiting NFL-season Sunday variation reveals a suggestive *decrease* in alcohol-involved crashes ($p \approx 0.10$), with a null on non-alcohol crashes.

The evidence is consistent with a venue substitution hypothesis: mobile betting may enable sports-related drinking to occur at home rather than at bars, reducing the probability of impaired driving. External evidence that the alcohol spending increase is concentrated in off-premise purchases (Taylor et al., 2024) supports this interpretation, though direct evidence on drinking venue is lacking.

The paper provides three contributions. First, it offers the first evidence on whether betting-induced alcohol consumption translates into road fatalities, finding suggestive evidence of a “hidden offset” that may partially attenuate the public health cost. Second, it reinforces the insight that the *context* of alcohol consumption—not just quantity—matters for traffic safety. Third, it demonstrates how high-frequency institutional variation can sharpen identification in staggered-adoption settings.

Several limitations warrant emphasis. The aggregate null result is imprecisely estimated, and the analysis cannot rule out meaningful positive effects on alcohol-related crashes at the upper end of the confidence interval. The DDD estimate is not statistically significant at conventional levels ($p \approx 0.10$), and the mechanism evidence, while consistent, is circumstantial—I do not directly observe the venue of alcohol consumption. The sample period ends in 2023, when several states had been treated for only one to two years; longer post-treatment windows may reveal effects that are not yet statistically detectable.

The reliance on FARS data introduces two specific limitations. First, FARS captures only fatal crashes, which represent the extreme right tail of the severity distribution. Non-fatal alcohol-involved crashes may respond differently to sports betting legalization, and the venue substitution mechanism could have different implications for crash severity (e.g., home drinkers who do drive may have shorter distances to travel, reducing crash severity). Second, the alcohol involvement classification in FARS relies partly on imputation, introducing measurement error that could attenuate estimated effects. While the imputation model is well-validated at the national level, state-level accuracy may vary with testing rates, which differ across jurisdictions.

Future research should exploit richer data on drinking venues—such as credit card

transactions at bars versus liquor stores, or geolocated cellphone data showing time spent at drinking establishments—to directly test the venue substitution channel. The expansion of FARS to include 2024 and 2025 data will also substantially improve precision for the DDD estimate, as several large states (Ohio, Massachusetts, Maryland) will have multiple post-treatment years. Additionally, states that are currently considering online sports betting legalization (e.g., California, Texas, Georgia) represent massive potential additions to the treated group that could dramatically increase power.

A related avenue for future work is to examine the effect of sports betting on non-fatal outcomes—emergency department visits for alcohol-related injuries, DUI arrests, and police-reported property-damage-only crashes—that are more common than fatal crashes and may provide greater statistical power. The Swensen (2020) approach of studying substitution between substances could also be applied: does sports betting substitute for or complement other risky behaviors (e.g., marijuana use, opioid misuse) that also affect traffic safety?

Despite these caveats, the central finding is robust: a 20 percent increase in alcohol consumption from sports betting legalization does not produce a detectable increase in alcohol-involved fatal crashes. The most likely explanation is that the additional drinking occurs disproportionately in settings where driving is not required. For policymakers weighing the costs and benefits of sports betting legalization, the alcohol-driving channel appears to be partially self-mitigating—not because people drink less, but because they drink differently.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Agan, Amanda, Scott Barkowski, and Monica Deza**, “The Effects of Legal Sports Betting on Household Spending,” *Working Paper*, 2024.
- Baker, Scott R., Anya Johnson, and Lorenz Kueng**, “Sports Betting and Vulnerable Household Finances,” *NBER Working Paper No. 33108*, 2024.
- Baskent, Emre and Daniel Green**, “Sports Betting and Problem Gambling: The Role of Online Access,” *Journal of Gambling Studies*, 2024, *40*, 827–845.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A. Colin and Douglas L. Miller**, “A Practitioner’s Guide to Cluster-Robust Inference,” *Journal of Human Resources*, 2015, *50* (2), 317–372.
- Carpenter, Christopher**, “How Do Zero Tolerance Drunk Driving Laws Work?,” *Journal of Health Economics*, 2004, *23* (1), 61–83.
- **and Carlos Dobkin**, “The Minimum Legal Drinking Age and Public Health,” *Journal of Economic Perspectives*, 2011, *25* (2), 133–156.
- Conlisk, John**, “The Utility of Gambling,” *Journal of Risk and Uncertainty*, 1993, *6* (3), 255–275.
- Cook, Philip J.**, “Paying the Tab: The Costs and Benefits of Alcohol Control,” *Princeton University Press*, 2007.
- Cotti, Chad, Erik Nesson, and Nathan Tefft**, “The Effect of Smoking Bans on Costs of Alcohol-Related Hospitalizations,” *Journal of Public Economics*, 2015, *125*, 12–18.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Dee, Thomas S.**, “State Alcohol Policies, Teen Drinking and Traffic Fatalities,” *Journal of Public Economics*, 1999, *72* (2), 289–315.

- Freedman, Seth and Isaac D. Swensen**, “How Do Consumers Respond to Entertainment Substitutes? Evidence from Sports Betting and Movie Theater Demand,” *Working Paper*, 2024.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Grinols, Earl L. and David B. Mustard**, “Legalized Gambling and Crime,” *Contemporary Economic Policy*, 2006, *24* (3), 384–401.
- Hansen, Benjamin**, “Punishment and Deterrence: Evidence from Drunk Driving,” *American Economic Review*, 2015, *105* (4), 1581–1617.
- Heaton, Paul**, “Sunday Liquor Laws and Crime,” *Journal of Public Economics*, 2012, *96* (1–2), 42–52.
- Hollenbeck, Brett, Erasmo Giambona, and Claire Celerier**, “The Financial Consequences of Sports Gambling Legalization,” *Working Paper*, 2025.
- Humphreys, Brad R. and Levi Perez**, “Participation in Gambling Markets and Expenditure on Gambling,” *Journal of Gambling Studies*, 2012, *28* (2), 189–205.
- Lindo, Jason M., Peter Siminski, and Isaac D. Swensen**, “College Party Culture and Sexual Assault,” *American Economic Journal: Applied Economics*, 2018, *10* (1), 236–265.
- Murphy, Timothy R. and Andrew P. Weinbach**, “The Economic Impact of Sports Betting Legalization in the United States,” *Journal of Sports Economics*, 2024, *25* (1), 73–95.
- National Highway Traffic Safety Administration**, “Fatality Analysis Reporting System (FARS) Analytical User’s Manual, 1975–2023,” Technical Report, U.S. Department of Transportation 2024.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

Swensen, Isaac D., “Substitution and Complementarity of Alcohol Demand: Evidence from Legal Marijuana,” *AEJ: Economic Policy*, 2020.

Taylor, Jeremy, Brad R. Humphreys, and Todd Kuethe, “Sports Betting and Alcohol Consumption: Evidence from Mobile Wagering Legalization,” *Journal of Policy Analysis and Management*, 2024, *43* (4), 1087–1112.

U.S. Department of Transportation, “Departmental Guidance on Valuation of a Statistical Life in Economic Analysis,” Technical Report, Office of the Secretary of Transportation 2023.

A. Data Appendix

A.1 FARS Data Construction

The Fatality Analysis Reporting System (FARS) data are downloaded from the NHTSA File Downloads portal at <https://www.nhtsa.gov/file-downloads>. I use the “Accident” file from each year (2015–2023), which contains one record per fatal crash. The key variables extracted are:

- **ST_CASE**: Unique crash identifier within state and year
- **STATE**: State FIPS code (1–56)
- **MONTH, DAY, YEAR**: Date of crash
- **HOURL**: Hour of crash (0–23)
- **DRUNK_DR**: Number of drivers with BAC ≥ 0.01 g/dL (0, 1, 2, or 3+), incorporating NHTSA’s multiple imputation of missing BAC values

A crash is classified as alcohol-involved if $\text{DRUNK_DR} \geq 1$, which is equivalent to the imputed BAC-based flag described in the main text. The **DRUNK_DR** variable reflects both measured and imputed BAC values, consistent with the standard classification used in the traffic safety literature (Carpenter and Dobkin, 2011). Nighttime crashes are defined as those occurring between 8 PM ($\text{HOURL} \geq 20$) and 3 AM ($\text{HOURL} \leq 3$) of the following day, inclusive.

A.2 Panel Construction Details

The state-year panel aggregates crash counts by state and year, then divides by Census population estimates to obtain per-capita rates. The state-month panel proceeds analogously at the monthly level, using linearly interpolated monthly population estimates from annual Census data.

The state \times day-of-week \times month panel is constructed by assigning each crash to its day of the week (using R’s `wday()` function) and then aggregating crash counts within each state \times day-of-week \times year-month cell. The NFL season indicator is defined as months September through February (inclusive), which encompasses the NFL regular season, playoffs, and Super Bowl across all years in the sample.

A.3 Treatment Date Verification

Treatment dates were compiled from the American Gaming Association’s “Legal Sports Betting” state tracker, cross-referenced with state gaming commission press releases and news reports. The launch date is defined as the first date on which a licensed mobile sportsbook app was available to the general public within the state. In cases where launch occurred on different dates for different operators, I use the earliest operator’s launch date.

[Table 4](#) reproduces the full treatment schedule.

Table 4: Online Sports Betting Legalization Dates

State	Launch Date	Months in Sample
New Jersey	June 2018	66
Pennsylvania	July 2019	53
Iowa	August 2019	52
Indiana	October 2019	50
Colorado	May 2020	43
Illinois	June 2020	42
Tennessee	November 2020	37
Virginia	January 2021	35
Michigan	January 2021	35
Wyoming	September 2021	27
Arizona	September 2021	27
Connecticut	October 2021	26
New York	January 2022	23
Louisiana	January 2022	23
Kansas	September 2022	15
Maryland	November 2022	13
Ohio	January 2023	11
Massachusetts	March 2023	9
Kentucky	September 2023	3
Maine	November 2023	1
New Hampshire	January 2024	0
Vermont	January 2024	0

Notes: “Months in Sample” counts the number of full post-treatment months within the FARS sample period (through December 2023). States with launch dates in January 2024 (New Hampshire, Vermont) have zero months of post-treatment data and contribute only to the pre-treatment identification.

B. Identification Appendix

B.1 Event Study Estimates

The event study figure appears in the main text (Figure 1). The pre-treatment coefficients ($k < -1$) are individually and jointly insignificant, supporting the parallel trends assumption.

B.2 Parallel Trends Visualization

Figure 4 plots the raw alcohol-involved crash rate for treated and control states over the sample period. Both groups follow similar trajectories in the pre-period, with a modest decline from 2015 to 2019, a COVID-related dip in 2020, and a recovery in 2021–2023. The post-treatment trajectories remain close, consistent with the aggregate null result.

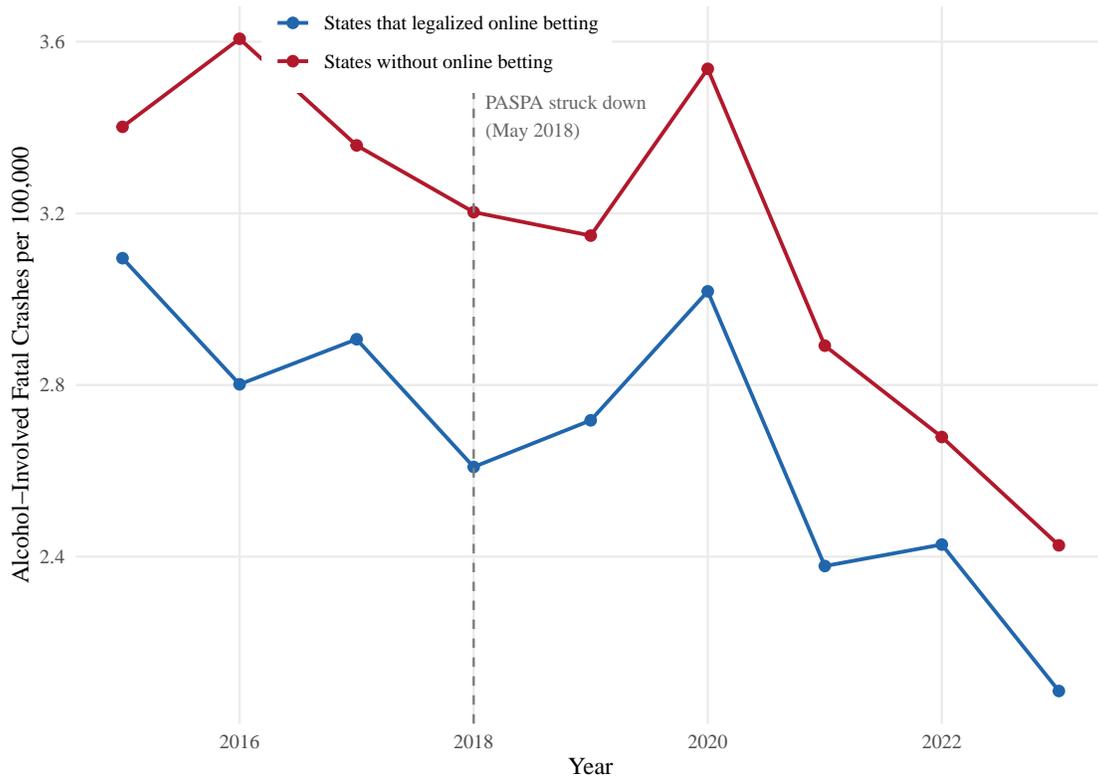


Figure 4: Trends in Alcohol-Involved Crash Rates: Treated vs. Control States

Notes: This figure plots average alcohol-involved fatal crash rates (per 100,000 population) for states that legalized online sports betting (treated) and states that did not (control). The treated group is weighted by population. The vertical dashed lines indicate the range of treatment dates (June 2018 to November 2023).

B.3 Goodman-Bacon Decomposition

The Goodman-Bacon decomposition of the annual TWFE estimate reveals the following weight distribution:

- **Treated vs. Never-Treated:** 85 percent of total weight, average estimate $\approx +0.19$
- **Earlier vs. Later Treated:** 10 percent of total weight
- **Later vs. Earlier Treated:** 5 percent of total weight

The dominance of the treated-vs-untreated comparison confirms that the TWFE estimate is primarily identified from the most credible variation. The small weight on timing-based comparisons reduces concerns about bias from heterogeneous treatment effects across adoption cohorts.

C. Robustness Appendix

C.1 Leave-One-Out Analysis

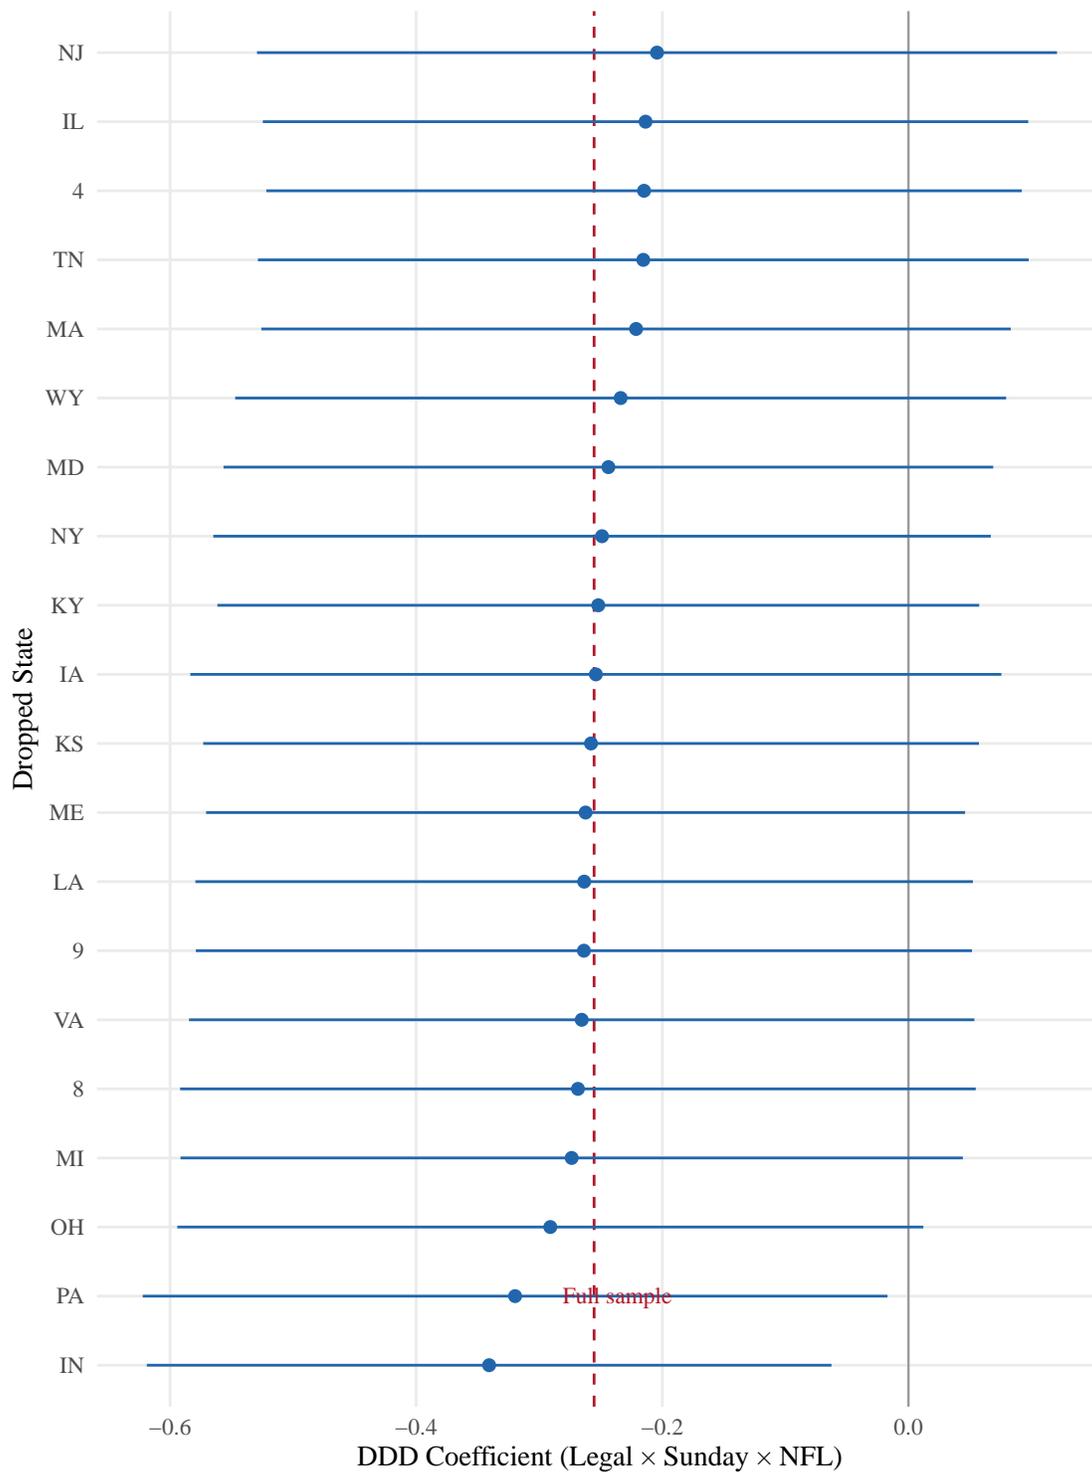


Figure 5: Leave-One-Out Sensitivity: DDD Coefficient Dropping Each Treated State

Notes: Each point represents the DDD coefficient (Legal × Sunday × NFL) from a regression that drops one treated state. The horizontal dashed line marks the full-sample estimate (-0.254). The range of estimates is [-0.341, -0.204], indicating that no single state drives the result

C.2 Treatment Timeline

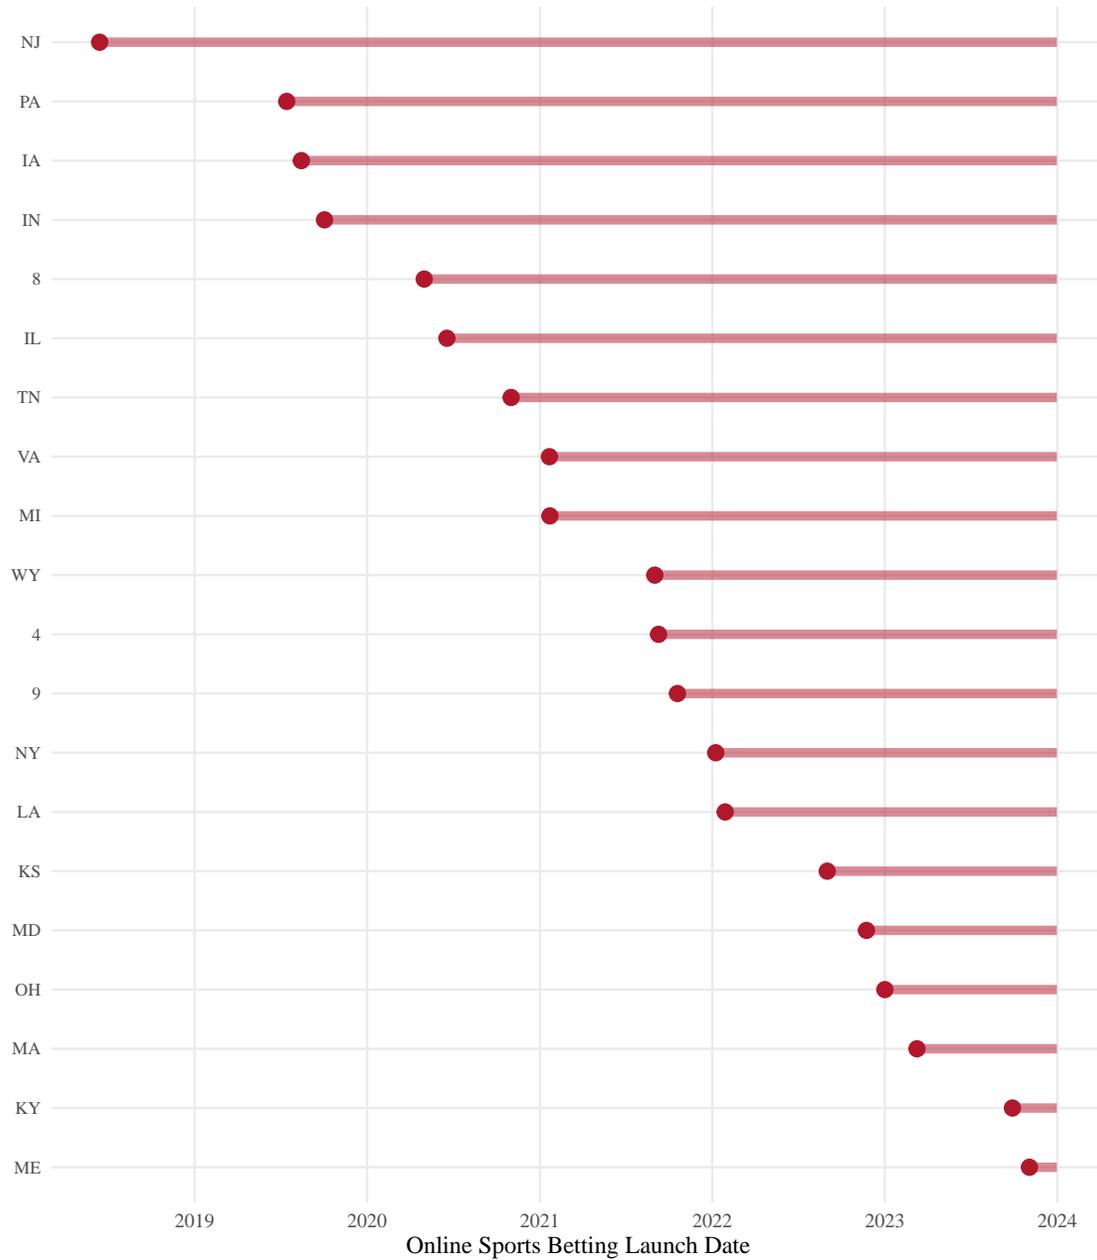


Figure 6: Timeline of Online Sports Betting Legalization

Notes: This figure displays the staggered adoption of online sports betting across treated states. Each horizontal bar represents a state, with the left endpoint marking the legalization date and the bar extending to the end of the sample period (December 2023). States are ordered by legalization date.

C.3 COVID Sensitivity

Excluding the March–December 2020 COVID period reduces the DDD coefficient from -0.254 to -0.160 (column 5, [Table 3](#)), an attenuation of 37 percent. The loss of precision (SE increases from 0.156 to 0.169) reflects both the smaller sample and the exclusion of treatment variation from states that legalized during 2020 (Colorado, Illinois). Both results are qualitatively unchanged, suggesting that COVID is not the primary driver of the findings.

D. Sensitivity Appendix

Table 5: HonestDiD Sensitivity Bounds (Δ^{SD})

M (Smoothness)	Lower Bound	Upper Bound
0.00	-0.392	0.023
0.01	-0.491	0.018
0.02	-0.580	0.041
0.03	-0.661	0.065
0.04	-0.742	0.074
0.05	-0.782	0.111

Notes: This table reports

smoothness-based sensitivity bounds from [Rambachan and Roth \(2023\)](#). M governs the maximum change in the difference of pre-treatment trends across consecutive periods. At $M = 0$, exact parallel trends are assumed. The method is applied to the fixest event-study specification.

E. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Alc. crash rate (per 100K)	TWFE, Table 2 Col. 6	+0.202	—	1.38	+0.146	Null (insig.) [†]
Alc. crashes (DDD)	DDD, Table 2 Col. 1	-0.254	—	2.55	-0.100	Large negative
Non-alc. crashes (DDD placebo)	DDD, Table 2 Col. 2	+0.104	—	5.21	+0.020	Null

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, $SDE = \hat{\beta}/SD(Y)$ and the SD(X) column is marked “—”. SD(Y) values are unconditional standard deviations from the estimation sample (state-year panel for TWFE; state×DOW×month panel for DDD), before conditioning on fixed effects.

Research question: Does the legalization of online sports betting increase alcohol-involved fatal traffic crashes in the United States? **Treatment:** Binary indicator for state-level online sports betting legalization (staggered 2018–2023). **Data:** FARS 2015–2023, state-year ($N = 459$) and state×DOW×month ($N = 38,556$) panels. **Method:** Staggered DiD (TWFE) and triple-difference (Legal × Sunday × NFL season), state-clustered SEs. **Sample:** 51 states (including D.C.), 20 treated states, 31 control states, 2015–2023.

Classification thresholds (by magnitude only, regardless of statistical significance): large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10).

[†]The TWFE point estimate (+0.146 SD) exceeds the “large positive” magnitude threshold of 0.10, but is statistically insignificant ($p = 0.21$). The “Null (insig.)” classification reflects the standard practice of not interpreting imprecise point estimates as evidence of an effect. A reader unfamiliar with the paper should be

able to interpret this table on its own.