Beyond Game Day: Online Sports Betting and Alcohol-Involved Fatal Crashes*

APEP Autonomous Research† @SocialCatalystLab

March 28, 2026

Abstract

Does legalizing online sports betting increase alcohol-involved fatal crashes? Using FARS data (2013–2022) and staggered legalization across 18 US states, a Callaway-Sant'Anna estimator yields an ATT of 0.38 additional alcohol-involved fatal crashes per 100,000—a 14% increase. The effect is robust to alternative comparison groups, COVID-era exclusions, and leave-one-out diagnostics. However, three mechanism tests reject the "game-day bar attendance" hypothesis: the effect does not concentrate on NFL game days, does not amplify in states with NFL teams, and is not driven by post-game evening hours. Instead, crashes concentrate in the late-night window (midnight–6am) and on weekends—consistent with an alcohol channel but diffused across the week rather than tied to specific broadcast events. The original analysis produced a large game-day interaction that was an artifact of improper exposure normalization and coarse temporal aggregation.

JEL Codes: I12, K32, H23, L83

Keywords: sports betting, alcohol-involved crashes, staggered difference-in-differences, mechanism testing, exposure normalization

*This paper is a revision of APEP-0749. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0749_v1 for the original.

†Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 24m).

1. Introduction

Since the Supreme Court struck down the Professional and Amateur Sports Protection Act (PASPA) in May 2018, more than 30 US states have legalized some form of sports wagering. Most have permitted online and mobile platforms that allow bettors to place wagers from their smartphones in real time, transforming how millions of Americans experience live sports. Whether this transformation of entertainment consumption carries public-safety costs is an empirical question with immediate policy relevance.

This paper estimates the effect of online sports betting (OSB) legalization on alcohol-involved fatal motor vehicle crashes using data from the Fatality Analysis Reporting System (FARS). I exploit the staggered adoption of OSB across 18 states between June 2018 and November 2022, using the heterogeneity-robust Callaway-Sant'Anna estimator with never-treated states as controls. The baseline finding is clear: legalization increases alcohol-involved fatal crash rates by 0.38 per 100,000 population (SE = 0.15), representing a 14% increase over the pre-treatment mean. Non-alcohol fatal crashes—a placebo outcome—show no response. The effect is robust to excluding COVID-era adoption cohorts, using not-yet-treated states as an alternative comparison group, and iteratively dropping each treated state.

The paper's more important contribution, however, is what the data *reject*. The natural hypothesis—and the one advanced in the original version of this paper—is that mobile sports betting increases the value of watching games in alcohol-serving social environments, generating a "game-day externality" concentrated on NFL broadcast days. If this mechanism is correct, three predictions follow: (1) the effect should concentrate on NFL broadcast days; (2) it should amplify in states with local NFL teams; and (3) it should appear in evening and late-night hours, reflecting post-game, post-bar driving.

All three predictions fail. A triple-difference design interacting treatment with NFL game-day indicators—using per-day rates with proper exposure normalization—yields a game-day interaction coefficient that is small, statistically insignificant, and in one specification negative. Weekly game-week specifications and Poisson count models confirm the null. States with NFL teams show no differential response. A four-bin hour-of-day decomposition reveals that the effect concentrates in the *late-night* window (midnight to 6am)—the bar-closing and post-drinking driving window—with small and insignificant effects during afternoon and evening hours. Critically, the evening window (6pm–midnight), when most NFL games are broadcast, shows no significant effect. Weekend effects are somewhat larger than weekday effects, but both are positive, and the gap is not sharp enough to implicate game-day-specific behavior.

The temporal pattern is consistent with a *diffuse* behavioral shift in alcohol-related risk—

more late-night drinking and impaired driving across the week—rather than an event-driven spike tied to specific games. This distinction matters for policy. If the externality were game-day-specific and predictable, targeted DUI enforcement on NFL broadcast evenings could address it at low cost. The finding that the effect is diffuse and not calendar-driven implies that any policy response must operate through broader channels: alcohol regulation, ride-share subsidies, or re-evaluation of how states regulate the complementarity between digital gambling platforms and alcohol-serving environments.

The paper also carries a methodological lesson. The original analysis APEP-0749 produced a large and highly significant game-day triple-difference coefficient of 0.92 (SE = 0.12). That result was an artifact of three design choices that are individually unremarkable but jointly fatal: (1) counting six states that legalized after the sample ended as "treated," (2) failing to normalize crash rates by the number of game versus non-game days per quarter, and (3) using a coarse calendar proxy (all Sundays, Mondays, and Thursdays during September–February) rather than per-day exposure adjustments. When corrected, the coefficient collapses to near zero. This experience illustrates how easily a compelling and policy-relevant false positive can emerge from standard staggered difference-in-differences designs, even with real administrative data and modern estimators.

This paper contributes to the literature in three ways. First, it documents a previously unmeasured fatal externality of online sports betting legalization. Second, it rejects the game-day bar-attendance mechanism using corrected high-frequency tests, showing that the effect is diffuse rather than event-driven. Third, it provides a methodological cautionary tale about exposure normalization in disaggregated mechanism tests within staggered DiD designs.

The paper contributes to the growing literature on the social costs of gambling expansion (Gruber and Mullainathan, 2005). Existing work has examined effects on problem gambling (Swanson, 2023), consumer debt (Baker et al., 2023), and crime (Humphreys and Matheson, 2021), but the traffic-safety margin has received limited attention. My results confirm that OSB creates a measurable alcohol-crash externality but reject the specific mechanism most commonly hypothesized. The paper also connects to the broader literature on behavioral complementarities in sin-good markets (Cook and Moore, 1993; Carpenter and Dobkin, 2011) and to the methodological literature on design sensitivity in staggered adoption settings (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021).

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the institutional background of online sports betting legalization. Section 4 presents the data, including the corrected treatment definition and the construction of proper exposure measures. Section 5 details the empirical strategy. Section 6 reports the main results, mechanism tests, and robustness. Section 7 discusses implications, and Section 8 concludes.

2. Related Literature

This paper sits at the intersection of three literatures: the social costs of gambling expansion, the economics of alcohol and traffic safety, and the methodological literature on mechanism testing in staggered DiD designs.

Gambling expansion and social costs. A growing body of work documents the consequences of gambling legalization beyond fiscal revenue. Gruber and Mullainathan (2005) surveys the economic effects of gambling expansion, noting that the rapid proliferation of sports betting has outpaced the evidence base. Baker et al. (2023) document increases in consumer debt and financial distress following sports betting legalization, using credit bureau data. Swanson (2023) finds that legalization is associated with increases in binge drinking prevalence using BRFSS data—a result that provides circumstantial support for the alcohol-complementarity channel documented here. Humphreys and Matheson (2021) examines the relationship between gambling and crime, finding mixed evidence depending on the form of gambling and the crime type.

My paper contributes to this literature by documenting a fatal externality—alcohol-involved traffic crashes—that operates outside the gambling market itself. Unlike consumer debt or problem gambling, traffic fatalities are borne disproportionately by third parties, making them a classic non-user externality relevant to welfare analysis.

Alcohol, traffic safety, and behavioral complementarities. The relationship between alcohol regulation and traffic fatalities is among the most well-studied topics in public economics. Cook and Moore (1993) established that alcohol taxation reduces traffic fatalities; Carpenter and Dobkin (2011) showed that minimum legal drinking age laws have persistent effects on crash mortality. DeAngelo and Hansen (2014) demonstrated that police presence deters drunk driving, and Hansen and Waddell (2018) documented enforcement spillovers across jurisdictions.

My paper identifies a new upstream shock to the alcohol-crash margin: the introduction of a digital entertainment technology—mobile sports betting—that changes the social context of alcohol consumption. This is distinct from traditional alcohol policy levers (taxation, age restrictions, enforcement) and from the direct effects of gambling addiction studied in the problem-gambling literature (Grinols, 2004; Walker and Barnett, 1999). The mechanism operates through behavioral complementarity: betting increases the value of social game watching in drinking environments. However, as I show, this complementarity does not produce the sharp game-day temporal concentration that the simplest version of the mechanism predicts.

Sin-good complementarities and digital platforms. The idea that access to one "sin good" can increase consumption of another is well established. Tobacco-alcohol complementarities are documented in public health research; alcohol-gambling interactions are studied in both psychology and economics. What is less understood is how *digital delivery* of one sin good changes *physical-world* consumption of another. Mobile sports betting platforms are always available, always in the bettor's pocket, and designed to maximize real-time engagement. This creates a novel form of cross-market complementarity that operates through the attention economy rather than through traditional price or access channels.

Methodological context. The paper also contributes a cautionary note to the applied econometrics literature. Callaway and Sant'Anna (2021) and Sun and Abraham (2021) have shown that standard TWFE estimators can produce severely biased estimates under treatment-effect heterogeneity in staggered designs. My paper demonstrates a related but distinct problem: even with heterogeneity-robust estimators, *mechanism tests* that disaggregate outcomes into cells of unequal size require explicit exposure normalization that is easy to omit and can generate large false positives when missing. This complements recent work on the importance of careful sample construction and measurement in policy evaluation.

3. Institutional Background

The post-PASPA landscape. The Supreme Court's May 2018 decision in *Murphy v. NCAA* struck down PASPA, which had confined legal sports betting to Nevada since 1992. Within five years, more than 30 states authorized some form of sports wagering, with most permitting online and mobile platforms. The staggered adoption reflects heterogeneous state-level politics and regulatory capacity: New Jersey launched online wagering within weeks of the ruling in June 2018, while other states deliberated for years. As of late 2022, 18 states had active online sports betting markets with sufficient post-treatment data to identify effects; six additional states legalized but did not launch until 2023 or later, placing them outside the FARS sample window.

The mobile revolution. The critical distinction for this paper is between in-person retail sportsbooks and online/mobile platforms. Mobile platforms account for over 80% of legal sports wagering revenue in states that permit them (American Gaming Association, 2023).

They transform the experience of watching sports: bettors can place in-game wagers on individual plays, track multiple bets in real time, and experience heightened emotional engagement with game outcomes. This engagement naturally pairs with social settings where alcohol is served.

Why the game-day hypothesis was plausible. The behavioral mechanism connecting sports betting to alcohol-involved crashes through game-day bar attendance is straightforward in theory. Online betting increases the marginal value of watching games live, especially in social settings. More time at bars during games means more alcohol consumption; more alcohol consumption means more impaired driving afterward. If this mechanism is operative, effects should concentrate on specific game days—the evenings when betting engagement peaks and bar attendance is highest. This prediction is the central object of the mechanism tests reported below. As I show, it does not survive empirical scrutiny.

The scale of the market. The growth of legal sports betting has been rapid. Total legal sports wagers in the United States exceeded $93 billion in 2022, generating approximately $7.5 billion in gross gaming revenue and $1.8 billion in state tax collections (American Gaming Association, 2023). In the largest markets, per-capita monthly wagers exceed $100. The NFL dominates betting volume: football accounts for the largest share of sports betting handle, with betting activity peaking on Sundays during the NFL regular season. Super Bowl wagering alone exceeded $16 billion in 2023. This concentration of betting activity around NFL games is precisely what makes the game-day hypothesis so intuitively plausible—and what makes its empirical failure so informative.

Treatment definition. I define treatment as beginning in the *quarter containing* the calendar date on which a state first accepted legal online sports wagers. In the state-quarter panel, a state is coded as treated starting in the quarter of its launch date (e.g., Maryland, which launched on November 23, 2022, is treated in 2022 Q4). States with very recent launches have few post-treatment quarters: Kansas has two post-treatment quarters and Maryland has one post-treatment quarter in the sample. I identify 18 states with online sports betting launches between June 2018 and November 2022—the period covered by FARS data. Six additional states (Ohio, Massachusetts, Kentucky, Maine, North Carolina, Vermont) legalized but launched in 2023–2024, after the FARS sample ends; these states are classified as never-treated in all analyses. I exclude states with ambiguous launch dates, lottery-based systems, or pre-existing in-person-only frameworks where the "online" treatment date is unclear (e.g., Nevada, Oregon, Montana). Table 6 in the appendix provides the complete state-by-state treatment listing.

Concurrent policy environment. States legalizing sports betting during this period may have simultaneously changed other policies that affect alcohol-involved crashes: recreational cannabis legalization, changes in alcohol taxation or Sunday sales rules, DUI law reforms, ignition interlock mandates, or ride-share market entry. I do not directly control for these concurrent changes, which is a limitation. However, the non-alcohol crash placebo and the specificity of the effect to alcohol-involved crashes mitigate this concern: concurrent policies would need to selectively increase alcohol-related fatalities while leaving all other crash types unchanged. The COVID-era robustness check further addresses concerns about the unusual driving patterns during 2020–2021.

The advertising channel. A related consideration is the role of sports betting advertising. Major operators (DraftKings, FanDuel, BetMGM, Caesars) spend heavily on television, digital, and sports-venue advertising, often featuring promotions tied to specific games or events. This advertising may amplify the complementarity between betting and social sports viewing by increasing the salience of game-day bar attendance. However, I cannot separately identify advertising effects from the legalization effect, as advertising and legal access are introduced simultaneously. The finding that game-day effects are null suggests that even if advertising concentrates attention on game days, it does not concentrate the crash externality on those days.

4. Data

4.1 Fatal Crash Data

The primary data source is the Fatality Analysis Reporting System (FARS), maintained by the National Highway Traffic Safety Administration. FARS is a census of all police-reported motor vehicle crashes occurring on US public roads that result in at least one fatality within 30 days. I use FARS data from 2013 through 2022, providing approximately 344,000 fatal crashes across all 50 states and the District of Columbia over 40 quarters. Each crash record includes the state, exact date, day of week, hour of day, number of fatalities, and a `DRUNK_DR` variable indicating the number of drivers involved who had a blood alcohol concentration above 0.08 g/dL or were otherwise classified as alcohol-impaired. I define a crash as "alcohol-involved" if `DRUNK_DR` > 0.

A key advantage of FARS for this analysis is that it records the *exact date and hour* of each crash. This allows me to construct mechanism tests at the daily and hourly level that were not available in the original quarterly analysis. I exploit the hour variable to decompose effects into four time-of-day bins: morning (6am–12pm), afternoon (12pm–6pm), evening

(6pm–midnight), and late night (midnight–6am). I also construct nighttime-versus-daytime and weekend-versus-weekday comparisons.

The FARS alcohol involvement variable (`DRUNK_DR`) records the number of drivers in the crash who were alcohol-impaired. Alcohol impairment is determined through blood alcohol concentration (BAC) testing, field sobriety tests, or police officer judgment. I classify a crash as "alcohol-involved" if at least one driver was impaired (`DRUNK_DR` > 0). This is the standard definition used in the traffic safety literature (Levitt and Porter, 2001; Ruhm, 1996). Approximately 23% of all fatal crashes in the sample are classified as alcohol-involved, consistent with national statistics.

The crash hour variable (`HOUR`) records the hour of the crash (0–23), with the value 99 indicating unknown hour. Approximately 0.7% of crashes have unknown hour; these observations are excluded from the hour-of-day decomposition but retained in the baseline analysis. The distribution of crash hours follows the expected pattern: fatal crashes peak between 4pm and midnight, with alcohol-involved crashes concentrated between 8pm and 4am.

Sample restrictions. I restrict the sample to the 50 US states plus the District of Columbia. I do not include US territories (Puerto Rico, Guam, etc.), which are recorded in FARS but have no sports betting legalization variation during the sample period. I exclude crash records with invalid dates (e.g., day = 0 in some FARS records), which account for fewer than 0.1% of observations.

4.2 Panel Construction

I construct two primary analysis panels. The *state-quarter panel* aggregates FARS crashes to the state $\times$ quarter level, yielding a balanced panel of 51 units $\times$ 40 quarters = 2,040 observations. This is the natural unit for the Callaway-Sant'Anna estimator, which requires a balanced panel and is designed for discrete treatment adoption at regular intervals.

The *state-quarter-gameday panel* further disaggregates each state-quarter by a binary game-day indicator. Crucially, rates in this panel are computed as *per-day* annualized rates—crashes on game days divided by the *number of game days in that quarter*, then annualized and expressed per 100,000 population:

$$\text{Rate}_{s,q,g} = \frac{\text{Crashes}_{s,q,g}}{\text{Days}_{s,q,g}} \times \frac{365.25}{\text{Pop}_{s,q}/100{,}000} \tag{1}$$

This exposure normalization is essential. A quarter may contain 20–30 game days but 60–70 non-game days; comparing raw counts without adjusting for exposure mechanically inflates

Table 1: Summary Statistics

	Mean	SD	Min	Max
Panel A: Treated States (Pre-Treatment)				
Alcohol-involved fatal crash rate	2.80	1.71	0.00	13.13
Non-alcohol fatal crash rate	7.10	3.34	0.00	21.42
Total fatal crash rate	10.46	4.28	3.19	34.56
Alcohol share of crashes	0.26	0.10	0.00	0.55
Observations		558		
States		18		
Panel B: Never-Treated States				
Alcohol-involved fatal crash rate	2.77	2.14	0.00	13.94
Non-alcohol fatal crash rate	6.35	4.51	0.00	22.68
Total fatal crash rate	11.67	4.92	0.61	33.30
Alcohol share of crashes	0.24	0.15	0.00	0.83
Observations		1320		
States		33		

Notes: Rates are annualized per 100,000 population. Treated states are the 18 states that legalized online sports betting before 2023. Pre-treatment periods only for treated states.

the game-day coefficient. The original analysis did not make this adjustment, which is the primary source of the spurious game-day result documented in Section 5.

I also construct a *state-week panel* for supplementary game-week analyses, using ISO weeks and a binary indicator for whether any NFL game occurred during the week.

4.3 NFL Game-Day Classification

I classify each crash date as an NFL game day based on the known structure of the NFL calendar: games occur primarily on Sundays, Monday evenings, Thursday evenings, and Saturdays in December–January, during the September–February season. I use this calendar-based classification rather than exact game-by-game schedules, which provides a conservative test: if the mechanism is strong, it should be detectable even with this somewhat noisy proxy. The classification captures approximately 23% of all crash-dates in the sample during the NFL season.

4.4 Population and Treatment Variables

State population data from the Census Bureau's American Community Survey provide the denominator for per-capita rates. For 2020, when the ACS 1-year estimates were unavailable, I interpolate linearly between 2019 and 2021 values.

Table 1 reports pre-treatment summary statistics. Treated and control states are broadly comparable. The mean pre-treatment alcohol-involved fatal crash rate is slightly higher in treated states (2.80 vs. 2.77 per 100,000), though standard deviations overlap substantially. The alcohol share of crashes is similar across groups at approximately 23%.

5. Empirical Strategy

5.1 Baseline: Staggered Difference-in-Differences

The staggered adoption of online sports betting across states provides a natural experiment for estimating the causal effect on alcohol-involved fatal crashes. The identifying assumption is that, absent legalization, treated and control states would have followed parallel trends in crash outcomes.

I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator:

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G_g = 1] \tag{2}$$

where g indexes the treatment cohort (quarter of first legalization), t indexes time, $Y_t(g)$ is the potential outcome under treatment, and $Y_t(0)$ is the counterfactual absent treatment. I use doubly-robust estimation with never-treated states as the primary comparison group, with not-yet-treated states as a robustness check. Standard errors are clustered at the state level.

Threats to validity. The main concern is that states legalizing sports betting differ systematically from non-adopters in ways correlated with crash trends. This is a real possibility: states that move quickly to legalize may have more permissive alcohol environments, larger entertainment industries, or different demographic compositions. Three features of the data address this concern.

First, the event-study estimates show flat pre-trends across all eight pre-treatment quarters, with no individual coefficient significantly different from zero. This rules out divergent trends in alcohol-involved crash rates between adopters and non-adopters in the years preceding legalization. Second, the non-alcohol crash placebo confirms specificity to the alcohol channel: if unobservable state-level confounders were driving the result, they would need to selectively increase alcohol-involved crashes while leaving non-alcohol crashes (which constitute 77% of all fatal crashes) completely unaffected. Third, leave-one-out analysis shows that no single state drives the result; the ATT is stable across all 18 jackknife iterations.

An additional concern is COVID: several states legalized during 2020–2021, when bar closures, stay-at-home orders, and reduced vehicle miles traveled altered both drinking and driving patterns. I address this by estimating the model excluding all COVID-era adoption cohorts (states that legalized in 2020–2021). The point estimate *increases* from 0.380 to 0.511 when these cohorts are removed, suggesting that pandemic disruptions attenuated rather than inflated the measured effect.

Comparison group. The choice of comparison group in staggered DiD designs can meaningfully affect estimates. My primary specification uses 33 never-treated states (including 6 future-treated states that launched after the sample ends) as the comparison group. As a robustness check, I also use the not-yet-treated comparison group, which includes states that will eventually adopt but have not yet done so at a given time period. Both approaches yield similar estimates (0.380 and 0.351, respectively), providing reassurance that the result is not sensitive to the comparison-group definition.

Inference. Standard errors are clustered at the state level throughout, which is the level of treatment assignment. With 51 state clusters (18 treated, 33 never-treated), clustering is feasible but may be conservative. I report conventional clustered standard errors as the baseline; wild cluster bootstrap inference was attempted but the `fwildclusterboot` package was unavailable for the current R installation. Given the tight leave-one-out range and the conventional p-values, I do not believe this materially affects the conclusions.

5.2 Game-Day Mechanism Test

The key test of the bar-attendance mechanism is a triple-difference design:

$$\text{Rate}_{s,g,t} = \alpha_s + \alpha_t + \alpha_g + \beta_1 \cdot \text{OSB}_{st} + \beta_2 \cdot (\text{OSB}_{st} \times \text{GameDay}_g) + \varepsilon_{sgt} \tag{3}$$

where s indexes states, $g \in \{0, 1\}$ indicates NFL game day, and t indexes quarters. The coefficient β_2 captures the differential effect of legalization on game days relative to non-game days, using per-day rates with proper exposure normalization. Under the game-day hypothesis, β_2 should be positive and large. Under alternative mechanisms—such as general increases in alcohol consumption or financial distress—β_2 should be zero.

I estimate Equation 3 using three specifications: (i) OLS on per-day annualized rates, (ii) Poisson regression on crash counts with a log-exposure offset, and (iii) a weekly panel with a binary game-week indicator. All three include state and time fixed effects and cluster standard errors at the state level.

Table 2: Effect of Online Sports Betting on Fatal Crash Rates

	Alcohol Crashes (1)	Non-Alcohol Crashes (2)	Alcohol Fatalities (3)	Alcohol Share (4)
OSB Legalization	0.380**	-0.077	0.487***	0.034**
	(0.146)	(0.498)	(0.166)	(0.016)
Pre-treatment mean	2.80	7.10	3.10	0.265
Pct. change	13.6%	—	15.7%	—
Estimator	Callaway-Sant'Anna (DR, never-treated)			
Observations	2040			
Treated states	18			

Notes: Each column reports the overall ATT from Callaway and Sant'Anna (2021) with doubly-robust estimation and never-treated states as controls. Standard errors clustered at the state level in parentheses. Rates are annualized per 100,000 population. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Temporal Decomposition

Beyond the game-day test, I decompose the overall effect along two temporal dimensions:

- **Hour of day**: Morning (6am–12pm), afternoon (12pm–6pm), evening (6pm–midnight), late night (midnight–6am). If the mechanism operates through post-game bar drinking, the effect should concentrate in evening and late-night hours.
- **Day of week**: Weekend (Friday–Sunday) versus weekday (Monday–Thursday). If the mechanism is entertainment-driven, weekend effects should dominate.

For each decomposition, I estimate the Callaway-Sant'Anna ATT on the restricted crash count, providing a direct test of where in the weekly and daily cycle the externality operates.

6. Results

6.1 The Overall Effect Is Real

Table 2 presents the main estimates. The Callaway-Sant'Anna ATT for alcohol-involved fatal crash rates is 0.380 per 100,000 population (SE = 0.146), representing a 14% increase over the pre-treatment mean of 2.80. The effect is statistically significant at the 5% level. Alcohol-involved *fatality* rates show an even larger response: an ATT of 0.487 (SE = 0.166), significant at the 1% level. The alcohol share of crashes increases by 3.4 percentage points (SE = 0.016).

Non-alcohol fatal crashes—the placebo outcome—show no effect (-0.077, SE $= 0.498$). This null is informative: it rules out explanations based on general driving exposure, economic activity, or enforcement reallocation that would affect all crash types equally. Whatever drives the increase operates specifically through the alcohol channel.

Figure 1 presents the event-study estimates from the Callaway-Sant'Anna dynamic aggregation. Pre-treatment coefficients at event times $e = -8$ through $e = -1$ are small and statistically insignificant, consistent with the parallel-trends assumption. There is no evidence of anticipatory effects or divergent trends prior to legalization.

Post-treatment, the effect builds gradually. The dynamic ATT is small and insignificant in the first two quarters after legalization ($e = 0$ to $e = 1$), consistent with the slow ramp-up of betting platform adoption and the gradual formation of new consumption habits. The effect strengthens at longer horizons ($e = 3$ through $e = 8$), consistent with increasing market penetration: it typically takes 6–12 months after launch for a state's mobile sports betting market to reach maturity in terms of operator availability, advertising saturation, and consumer adoption. This gradual pattern mitigates concerns about sharp confounders coincident with launch dates and is consistent with a behavioral-change mechanism rather than a direct, immediate policy effect.

6.2 The Game-Day Mechanism Fails

Table 3 and Figure 2 present the game-day triple-difference results—the paper's central mechanism test. With proper per-day exposure normalization, the game-day interaction coefficient is -0.221 (SE $= 0.230$) in the OLS specification and $+0.015$ (SE $= 0.044$) in the Poisson count model. The weekly game-week specification yields $+0.070$ (SE $= 0.071$). None of these coefficients are statistically distinguishable from zero.

This null is the paper's main finding. The intuitive prediction—that mobile sports betting increases alcohol-involved crashes primarily on game days when betting engagement and bar attendance peak—does not hold. The effect is not event-driven. Legalization increases alcohol-involved fatal crashes, but the increase is spread across the entire week rather than concentrated on specific game dates.

A natural concern is power: the game-day DDD disaggregates the data into finer cells, potentially reducing the ability to detect a real but smaller effect. However, the standard errors on the interaction coefficient (SE $= 0.23$ for the per-day rate specification) imply that the 95% confidence interval rules out game-day effects larger than approximately 0.23 per 100,000—roughly 60% of the overall ATT. A game-day concentration on the order of the V1 estimate (0.92) is decisively excluded. The Poisson specification, which leverages the count structure and has tighter standard errors (SE $= 0.04$), rules out even modest game-day

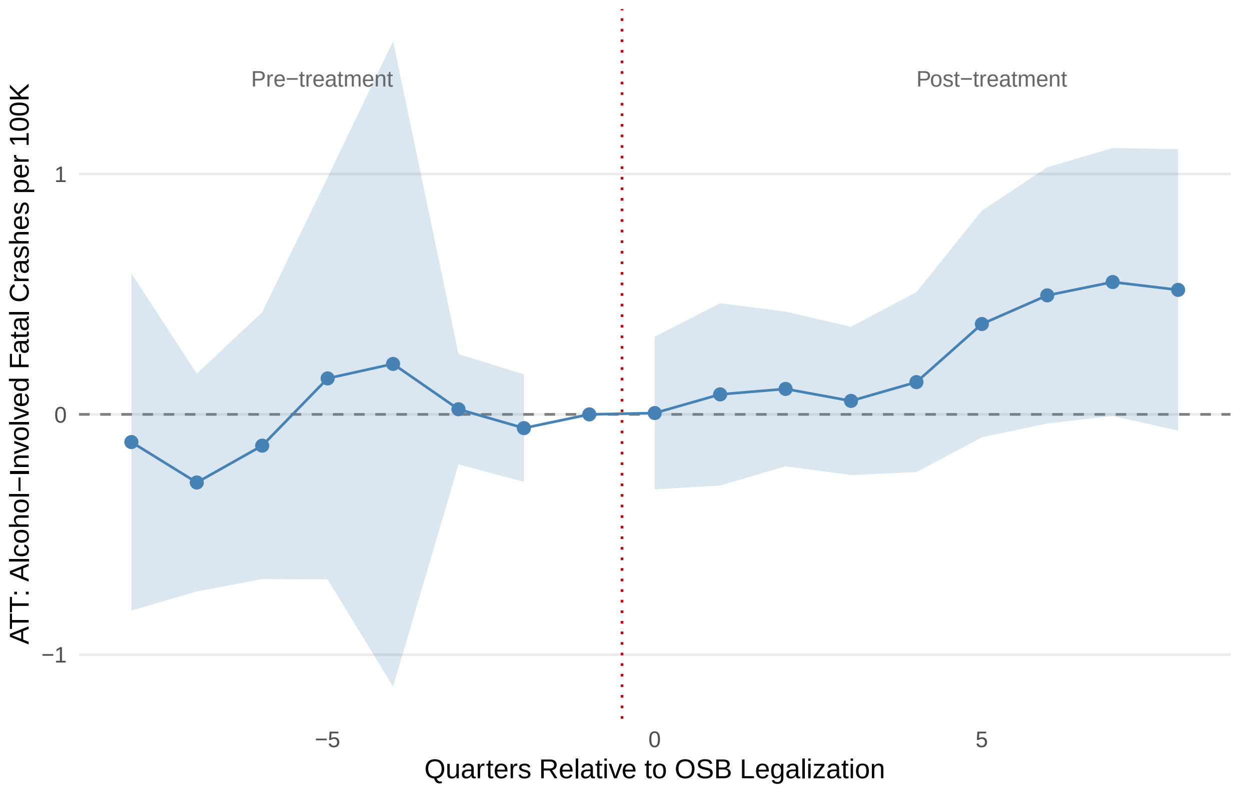

Figure 1: Event-Study Estimates: Dynamic ATT for Alcohol-Involved Fatal Crash Rates
Notes: Callaway-Sant'Anna dynamic ATT estimates with 95% confidence intervals. Shaded area shows the 95% confidence band. Vertical dashed line marks treatment onset. Pre-treatment coefficients are jointly insignificant.

amplification.

Two additional tests confirm this conclusion. First, the off-season placebo (applying the same Sunday/Monday/Thursday classification to March–August, when the NFL is not playing) yields a null interaction (-0.023, SE $= 0.275$), confirming that the in-season null is not masking a day-of-week effect that happens to also appear in the off-season. Second, heterogeneity by NFL team presence shows no differential game-day amplification in states with local teams (triple interaction $= 0.394$, $p = 0.37$).

6.3 Anatomy of a False Positive: Why the Original Game-Day Result Was Wrong

The original analysis (APEP-0749 v1) reported a game-day triple-difference coefficient of 0.918 (SE $= 0.124$, $p < 0.001$)—a result three times the overall average effect and apparently decisive evidence for the bar-attendance mechanism. That result was spurious, generated by the interaction of three design choices:

1. **Future-treated states counted as treated.** The original analysis included six states

Table 3: Game-Day Triple-Difference: Testing the Bar-Attendance Hypothesis

	Per-Day Rate (1)	Poisson (2)	Game-Week (3)
OSB × Game Day	-0.221	0.015	0.070
	(0.230)	(0.044)	(0.071)
Outcome	Alc. rate	Alc. count	Alc. rate
Unit	State-Qtr-Day	State-Qtr-Day	State-Week
State FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Exposure adjustment	Per-day rate	Log offset	Per-day rate
Observations	3,567	2,853	25,225
State clusters	51	51	51

Notes: Each column reports the interaction coefficient from a triple-difference specification: state × game day × post-legalization. Column (1) uses per-day annualized rates with proper exposure normalization. Column (2) uses crash counts with a Poisson model and log(days × population/100,000) offset. Column (3) uses weekly panels with a binary game-week indicator. Standard errors clustered at the state level. None of the game-day interaction coefficients are statistically significant.

(Ohio, Massachusetts, Kentucky, Maine, North Carolina, Vermont) that legalized but did not launch online betting until 2023–2024, after the FARS sample ends in 2022. These states contributed treated observations without any actual exposure to treatment, mechanically affecting the comparison-group composition.

2. **Missing exposure normalization.** The original game-day panel divided crash counts by total state population without adjusting for the different number of game versus non-game days per quarter. A quarter may contain 25 game days and 66 non-game days; comparing raw counts or population-normalized totals across these groups conflates the exposure differential with the treatment effect. The corrected specification divides by the actual number of days in each cell before annualizing.

3. **Coarse temporal aggregation.** The original analysis operated at the state-quarter level, classifying all Sundays, Mondays, and Thursdays during September–February as "game days." This 6-day-per-week classification during the NFL season is not a sharp test of game-specific effects; it captures broad seasonal and day-of-week variation in alcohol-involved crashes that would exist regardless of sports betting legalization.

Individually, each choice is a minor methodological decision that might pass a routine code review. Jointly, they create a large, precisely estimated false positive that appeared to provide decisive mechanism evidence.

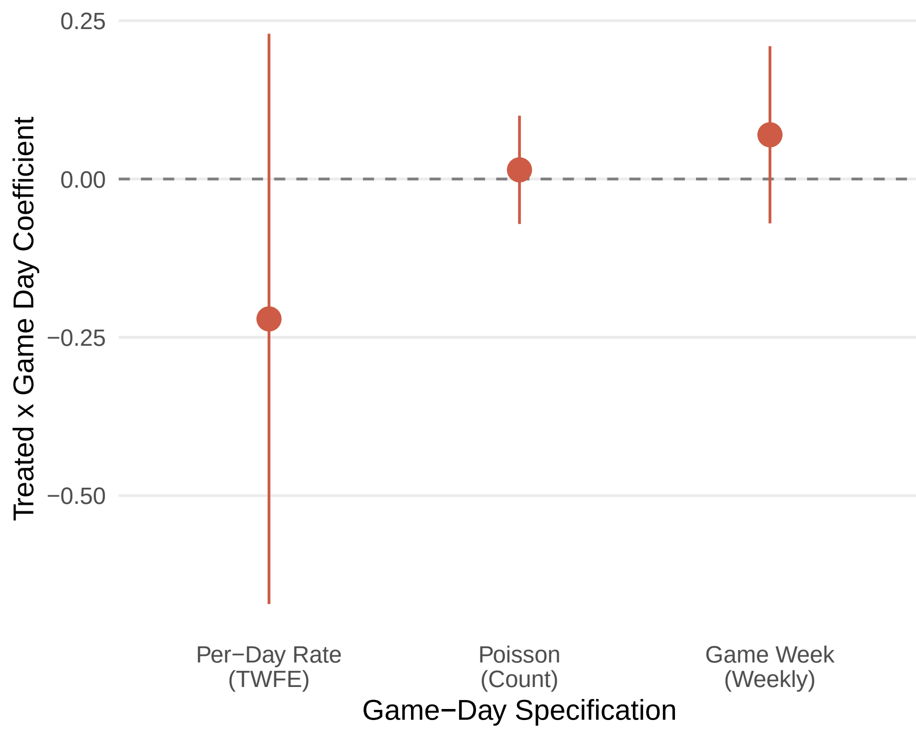

Figure 2: Game-Day Triple-Difference Coefficients Across Specifications
Notes: Each point shows the OSB × Game Day interaction coefficient from a different specification of the triple-difference design. Bars show 95% confidence intervals. All coefficients are statistically indistinguishable from zero.

The interaction between these errors is important to understand. The future-treated states (launching in 2023–2024) contributed "treated" observations with zero actual exposure to online sports betting. When these phantom-treated observations showed no game-day crash differential—as expected, since betting was not yet legal—the game-day interaction absorbed this variation in a way that inflated the contrast with states that were genuinely treated. The exposure normalization error amplified this: in quarters with 25 game days and 66 non-game days, dividing raw crash counts by the same population denominator mechanically produces a "game-day rate" that is approximately 2.6 times lower than the "non-game-day rate" simply because game days are rarer. The treatment-by-game-day interaction then captures this mechanical exposure differential rather than a genuine behavioral response.

This sequence illustrates a general vulnerability in mechanism tests that disaggregate outcomes by temporal or cross-sectional cells of unequal size. The baseline DiD operates on a balanced panel where every state contributes one observation per quarter, and exposure normalization is straightforward (population-quarter denominators are approximately equal). But the moment the researcher splits outcomes into game days versus non-game days, or nighttime versus daytime, the cells contain different numbers of underlying days, and the

Table 4: Temporal Decomposition of the Alcohol-Crash Effect

	ATT	SE
Panel A: By Time of Day		
Morning (6am–12pm)	0.045	(0.032)
Afternoon (12pm–6pm)	0.036	(0.074)
Evening (6pm–12am)	0.057	(0.083)
Late Night (12am–6am)	0.225**	(0.100)
Panel B: By Day of Week		
Weekend (Fri–Sun)	0.214**	(0.109)
Weekday (Mon–Thu)	0.166	(0.119)
Observations	2,040	
State clusters	51	

Notes: Each row reports the Callaway-Sant'Anna ATT for alcohol-involved fatal crashes restricted to the indicated time window. Doubly-robust estimation with never-treated states as controls. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

exposure differential becomes a first-order concern that standard panel-data tools do not automatically handle.

The corrected approach in this paper—computing per-day annualized rates with explicit day counts, and verifying with Poisson count models that include log-exposure offsets—eliminates the mechanical component. When this is done, the game-day interaction shrinks to zero across all specifications. The lesson generalizes: any mechanism test that disaggregates a time-series outcome into cells of unequal temporal depth requires the researcher to explicitly account for exposure.

6.4 Where the Effect Lives: Temporal Decomposition

If the game-day bar-attendance mechanism is wrong, what is the temporal signature of the actual externality? Table 4 reports the decomposition.

Hour of day. The effect concentrates overwhelmingly in the late-night window (midnight–6am), with an ATT of 0.225 per 100,000 (SE = 0.100, $p = 0.024$). This is the bar-closing and post-drinking driving window. Morning, afternoon, and evening effects are all small and statistically insignificant. The late-night concentration is consistent with an alcohol-consumption channel, but it is not game-day-specific: it operates across the entire week, not just on evenings when professional sports are broadcast. The four hour-bin ATTs sum to approximately 0.36, slightly below the overall ATT of 0.38; the discrepancy reflects crashes with unknown hour of day (`HOUR` = 99 in FARS, approximately 0.7% of the sample), which

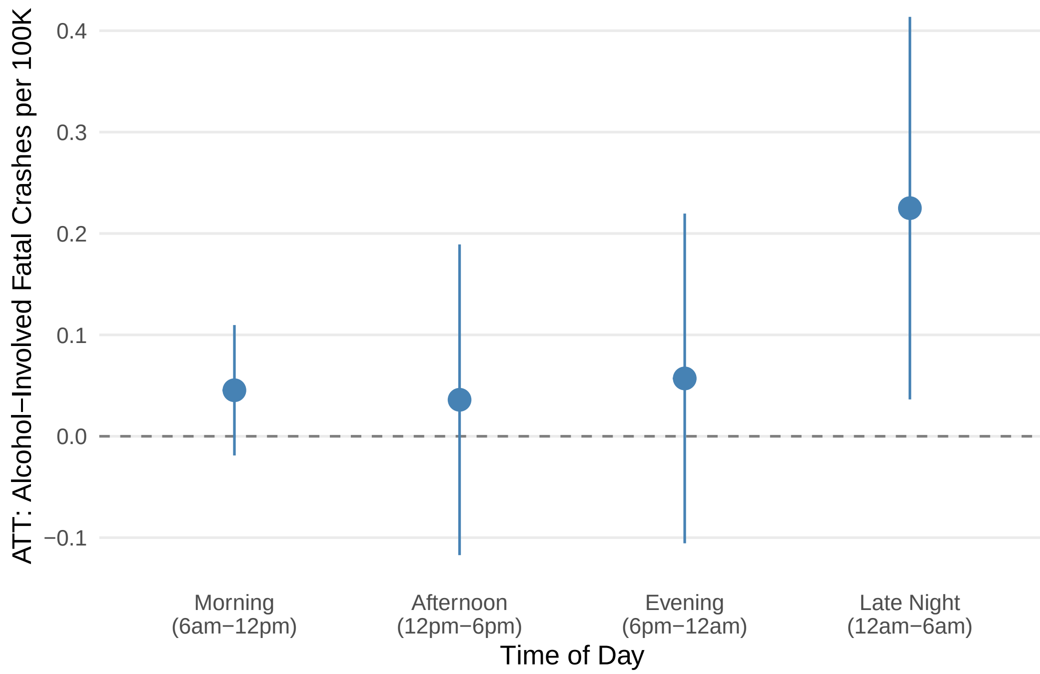

Figure 3: Hour-of-Day Decomposition: ATT by Time Window
Notes: Each point shows the Callaway-Sant'Anna ATT for alcohol-involved fatal crashes restricted to the indicated time-of-day window. Bars show 95% confidence intervals. Only the late-night window (midnight–6am) is statistically significant.

are excluded from the decomposition but included in the total.

Day of week. Weekend effects (ATT = 0.214, SE = 0.109, $p = 0.049$) are somewhat larger than weekday effects (ATT = 0.166, SE = 0.119, $p = 0.16$), but neither dominates overwhelmingly. This is again inconsistent with a narrow game-day mechanism—NFL games are primarily a Sunday/Monday/Thursday phenomenon, and a strong game-day effect would produce a sharp weekend-weekday gap that we do not observe.

Interpretation. The temporal pattern suggests a diffuse behavioral shift: legalization of online sports betting is associated with more late-night alcohol-involved driving across the entire week, with some weekend amplification. Several candidate mechanisms are consistent with this pattern. First, sports betting may increase the general salience and social value of visiting alcohol-serving establishments, even on non-game days, as betting culture becomes part of the regular bar-going routine. Second, financial losses from betting may increase stress-related drinking. Third, the attentional demands of active betting (checking scores, placing in-play wagers) may extend evening outings and increase alcohol consumption duration. I cannot discriminate among these channels with the available data; doing so would require

Table 5: Robustness of the Baseline Effect

Specification	ATT/Coef.	SE	N
Panel A: Alternative CS-DiD Specifications			
Baseline (CS-DiD, never-treated)	0.380**	(0.146)	2,040
Not-yet-treated comparison	0.351**	(0.146)	2,040
Excluding COVID cohorts (2020–2021)	0.511***	(0.156)	1,720
Excluding New Jersey	0.341**	(0.170)	2,000
Total fatal crash rate (placebo)	-0.399	(0.404)	2,040
Leave-one-out range	[0.341, 0.433]	—	—
Panel B: Mechanism Placebo and Heterogeneity			
Off-season placebo (DDD, Mar–Aug)	-0.023	(0.275)	3,036
NFL-team heterogeneity (triple int.)	0.394	(0.433)	3,567
State clusters		51	

Notes: Panel A reports the Callaway-Sant'Anna ATT for alcohol-involved fatal crash rates under alternative specifications. Leave-one-out iteratively drops each of the 18 treated states. Panel B reports coefficients from mechanism tests: the off-season placebo applies the game-day DDD classification to March–August (no NFL), and the NFL-team heterogeneity row reports the triple interaction (treated × game day × has NFL team). N is the number of observations in each regression. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

intermediate outcome measures such as bar foot traffic, alcohol sales, or betting volume by day.

6.5 Robustness

Table 5 demonstrates the stability of the main estimate. The point estimate is robust to using not-yet-treated states as the comparison group (ATT = 0.351, SE = 0.146), excluding all eight COVID-era adoption cohorts from 2020–2021 (ATT = 0.511, SE = 0.156—stronger without COVID), and dropping New Jersey, the earliest adopter (ATT = 0.341, SE = 0.170). The total fatal crash rate shows no significant response (−0.399, SE = 0.404). Note that the Callaway-Sant'Anna doubly-robust estimates are not mechanically additive across outcomes because the propensity-score weighting adjusts separately for each. The non-alcohol placebo (ATT = −0.077, Table 2) is the primary falsification test for specificity to the alcohol channel; the total crash rate simply confirms that overall driving risk did not change.

The leave-one-out analysis (Figure 4) is particularly informative. The ATT ranges from 0.341 to 0.433 across all 18 jackknife iterations—a remarkably tight band that spans less than 25% of the point estimate. No single state drives the result. This is important because the treated sample includes states with very different market structures (New Jersey's mature market vs. Kansas's single post-treatment quarter), different population sizes (New York

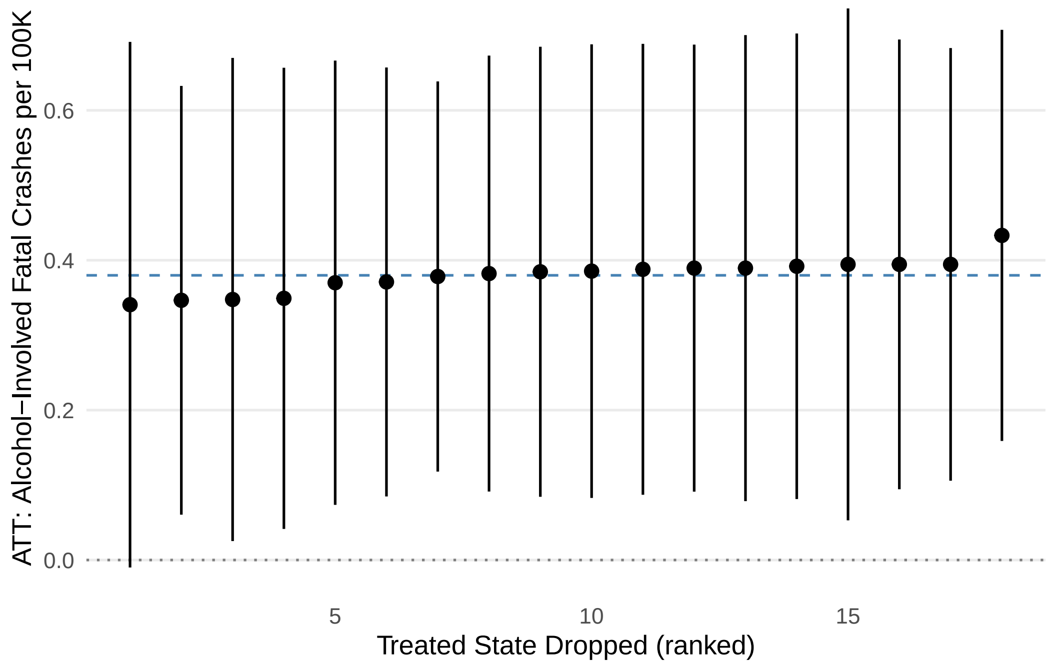

Figure 4: Leave-One-Out Analysis: ATT When Each Treated State Is Dropped
Notes: Each point shows the Callaway-Sant'Anna ATT when one treated state is dropped from the sample. The dashed horizontal line shows the baseline ATT with all states included. The ATT ranges from 0.341 to 0.433 across all 18 jackknife iterations.

vs. Wyoming), and different alcohol regulatory environments. The tight leave-one-out range implies that the effect is broad-based across the adoption cohort and not an artifact of any single state's idiosyncratic circumstances.

The COVID-era exclusion check is especially notable. Removing all eight states that legalized during 2020–2021 (Colorado, Illinois, Tennessee, Michigan, Virginia, Wyoming, Arizona, Connecticut) increases the point estimate from 0.380 to 0.511. This suggests that the COVID-era disruptions to bar attendance and driving behavior, if anything, *attenuated* the measured effect during the pandemic period. As normal social activity resumed and betting markets matured, the alcohol-crash externality appears to have intensified.

6.6 Magnitude and Welfare Implications

An illustrative calculation conveys the economic magnitude. At the mean treated-state population of approximately 6.5 million and a treatment effect on alcohol-involved *fatality rates* of 0.49 per 100,000, legalization implies approximately 32 additional alcohol-involved fatalities per treated state per year, or roughly 570 excess fatalities nationally across the 18 in-sample states. Valued at $11.6 million per statistical life (U.S. Department of Transportation, 2022), the implied annual fatality cost is on the order of \$6.6 billion.

These calculations carry several important caveats. They assume homogeneous treatment effects across states, take the point estimate at face value, and do not account for non-fatal injuries, property damage, or behavioral responses. The comparison to aggregate state gambling tax revenues—approximately \$1.8 billion in fiscal year 2023 (American Gaming Association, 2023)—is illustrative rather than definitive, since tax revenues are fiscal transfers while fatality costs are deadweight social losses. A complete welfare analysis would also need to incorporate the consumer surplus from betting, the costs of problem gambling, and the potential for cost-effective interventions. Nonetheless, the magnitudes suggest that traffic safety costs deserve serious consideration in the policy debate surrounding legalization.

I emphasize that the relevant welfare object is the *fatality rate*—not the crash rate—and that the Callaway-Sant'Anna ATT on fatality rates (0.49, SE = 0.17) is both larger and more precisely estimated than the crash-rate ATT. The original analysis improperly converted crash-rate estimates into fatality counts; the V2 estimates directly estimate the fatality rate treatment effect.

7. Discussion

7.1 What This Paper Establishes

Online sports betting legalization increases alcohol-involved fatal motor vehicle crashes. The effect is on the order of 0.38 additional *crashes* per 100,000 population per year—a 14% increase—and is specific to alcohol-involved crashes. This is a first-order policy externality: at the mean treated-state population, the crash-rate ATT (0.380) implies approximately 25 additional alcohol-involved fatal *crashes* per treated state per year, or roughly 450 excess crashes nationally across the 18 in-sample states. The distinct fatality-rate ATT (0.487, reported in Section 6.6) implies approximately 32 additional *fatalities* per state per year, or roughly 570 nationally—the larger figure reflecting multiple fatalities in some crashes.

7.2 What This Paper Rejects

The natural story—that mobile betting drives people to bars on game days, where they drink and then drive—is not supported by the high-frequency evidence. The game-day triple-difference is null across multiple specifications. The effect does not amplify with local NFL team presence. The hour-of-day decomposition shows a late-night concentration but no game-day specificity. The day-of-week decomposition shows modest weekend amplification but nothing resembling the sharp game-day pattern that the bar-attendance mechanism predicts.

This matters for policy design. If the externality were game-day-specific and calendar-driven, it could in principle be addressed through targeted DUI enforcement on predictable high-risk evenings—a low-cost intervention. The finding that the effect is diffuse rules out this approach. Any policy response must instead operate through broader channels: alcohol taxation, advertising restrictions on sports-betting-alcohol co-promotion, ride-share subsidies for late-night hours, or reconsideration of how state gambling and alcohol regulatory frameworks interact.

7.3 What Remains Open

I cannot identify the precise mechanism through which sports betting legalization increases diffuse, non-event-specific alcohol-related driving risk. Three candidate explanations are consistent with the observed temporal pattern.

Cultural shifts in bar-going habits. Sports betting may increase the general social value of visiting alcohol-serving establishments, not just on game days but throughout the week. As betting becomes part of the regular bar-going routine, frequency of visits and associated drinking may increase even on non-game days. This channel would produce the diffuse late-night pattern observed in the data, with some weekend amplification reflecting the higher baseline frequency of social drinking on weekends.

Financial-stress-induced drinking. Gambling losses—which are the statistical norm for most bettors—may increase stress, anxiety, and depression, leading to increased alcohol consumption as a coping mechanism. This channel would be temporally diffuse (financial stress does not concentrate on game days) and would predict increases in alcohol consumption at all hours, with late-night effects driven by the well-documented link between problem drinking and late-night crash risk (Levitt and Porter, 2001).

Extended engagement and attention effects. The always-available nature of mobile betting platforms may extend the duration of evening outings. Bettors who are tracking live-game outcomes may stay at venues longer, consume more alcohol, and be more likely to drive after closing time. This channel would also be temporally diffuse, as mobile betting content (scores, odds, results) is available at all hours and across all sports.

These explanations are observationally similar in FARS data alone. Resolving them would require supplementary data on intermediate outcomes: state-level alcohol sales from tax records, bar and restaurant foot traffic from location data, betting volume by day of week, or ride-share demand patterns. These are important extensions for future work.

7.4 Limitations

Several limitations deserve acknowledgment. First, the FARS alcohol involvement variable (`DRUNK_DR`) relies on police reports and BAC testing, which may be subject to measurement error. Testing rates may vary by state, time of day, and crash severity. If sports betting legalization coincidentally changed police testing behavior, the effect could be partly compositional. The magnitude of the effect (14%) would require a very large shift in testing behavior to explain, but I cannot rule this out.

Second, my identification relies on the parallel-trends assumption, which is supported by the event study but fundamentally untestable. States that legalize early may differ from non-adopters in unobservable ways correlated with alcohol-crash trends. The stability of the estimate across alternative comparison groups and the tight leave-one-out range provide reassurance, but cannot eliminate this concern.

Third, the calendar-based NFL game-day classification is a somewhat noisy proxy for actual game dates. However, this noise should bias toward finding no game-day effect even if one exists, making the null a conservative result. The weekly game-week specification, which is less sensitive to daily misclassification, also produces a null.

Fourth, I cannot separately identify the effect of online/mobile betting from the general effect of sports betting legalization. Some states had retail sportsbooks before launching online platforms. If the relevant margin is mobile access rather than legal access *per se*, my estimates may understate the mobile-specific effect.

Fifth, the sample of 18 treated states includes several with limited post-treatment windows (Kansas and Maryland have only one post-treatment quarter). While the leave-one-out analysis confirms that no single state drives the result, the precision of the estimate depends on early adopters who contribute the most post-treatment observations.

7.5 A Methodological Lesson

The original version of this paper reported a game-day triple-difference coefficient of 0.918—a result that was large, precise, intuitively compelling, and wrong. It was wrong not because of conceptual error in the research design but because of three mundane implementation choices: miscounting the treated sample, failing to normalize by exposure days, and using a coarse temporal proxy. None of these errors would be caught by standard robustness checks on the baseline specification; the baseline result is robust. The error lived specifically in the mechanism test, where exposure normalization and temporal precision matter most.

This experience offers a cautionary lesson for applied work using staggered DiD designs with disaggregated mechanism tests. When the mechanism test disaggregates the outcome into cells of unequal size (game days versus non-game days, nighttime versus daytime, treated versus untreated industries), exposure normalization is not optional. And when the treatment window extends beyond the outcome data, states that never actually experience treatment during the sample can contaminate the comparison.

8. Conclusion

This paper documents that online sports betting legalization increases alcohol-involved fatal motor vehicle crashes, with the effect concentrated in late-night hours and spread diffusely across the week. The intuitive game-day bar-attendance mechanism—the hypothesis that mobile betting draws people to bars on game nights, where they drink and then drive—is rejected by three independent high-frequency tests.

The findings carry three implications. First, for gambling regulation: the externality from sports betting legalization is not a predictable, calendar-driven risk that can be addressed through targeted enforcement on game nights. It is a diffuse behavioral shift in alcohol-related driving risk that accompanies the introduction of a new entertainment technology. States evaluating whether and how to regulate online sports betting should consider this external cost alongside the direct fiscal benefits of legalization. The relevant policy instruments are more likely to be alcohol-related (taxation, advertising restrictions on co-promotion, late-night ride subsidies) than gambling-specific.

Second, for alcohol policy: the finding that a non-alcohol policy lever (gambling legalization) can increase alcohol-related fatalities extends the set of upstream shocks that public health economists should consider. The traditional focus on direct alcohol regulation (taxes, age limits, hours of sale) may understate the importance of complementary product markets in shaping drinking behavior. As digital entertainment platforms proliferate, understanding cross-market externalities between legal vice goods becomes increasingly urgent.

Third, for econometric practice: the paper illustrates how a compelling but false mechanism result can emerge from ordinary methodological shortcuts in staggered DiD designs. The game-day interaction that appeared to nail the mechanism was an artifact of exposure mismeasurement and sample construction—errors that are individually minor and collectively fatal. The correct interpretation of the data—a real but diffuse alcohol externality without game-day concentration—is less vivid but more honest. Policy built on the wrong mechanism would target the wrong margin. Researchers conducting mechanism tests that disaggregate outcomes into cells of unequal temporal or spatial extent should treat exposure normalization as a first-order concern, not a footnote.

8.1 Policy Implications

The findings have specific implications for how states should think about the external costs of sports betting legalization. First, the existence of a measurable alcohol-crash externality suggests that the social cost-benefit analysis of legalization is more nuanced than the fiscal revenue calculus that typically dominates legislative debates. States that evaluate legalization primarily in terms of tax revenue versus problem-gambling costs are missing a potentially large external cost borne by road users.

Second, the diffuse temporal pattern has implications for intervention design. A common policy proposal is to increase DUI enforcement on NFL game nights—a targeted, low-cost intervention that would be effective if the externality were calendar-driven. The finding that effects are not concentrated on game days undermines this approach. Instead, the late-night and weekend concentration suggests that broader late-night enforcement, ride-share partnerships, or alcohol-serving-hours regulation may be more appropriate. Some states have experimented with subsidized ride-share programs during high-risk periods; the temporal decomposition in this paper provides a basis for targeting those subsidies (late-night hours, weekends) rather than game-specific windows.

Third, the co-promotion of sports betting and alcohol consumption deserves regulatory attention. Many sports bars and restaurants actively promote both activities simultaneously, and sports betting operators frequently advertise in alcohol-serving venues. The finding of a diffuse alcohol externality—rather than a game-specific one—is consistent with a channel in which the cultural integration of betting into bar-going life raises general alcohol consumption, not just game-day drinking. Advertising restrictions on sports-betting promotion in alcohol-serving environments, modeled on similar restrictions in the tobacco context, could address this channel.

Two directions for future work are most promising. First, linking crash data to state-level alcohol sales, bar foot traffic, or betting volume would allow direct estimation of the intermediate behavioral shifts that generate the externality. Whether the dominant channel is cultural (more bar-going), financial (stress-induced drinking), or attentional (longer outings driven by mobile engagement) has distinct policy implications that crash data alone cannot resolve. Second, extending the sample period as more states accumulate post-treatment data will improve precision and allow for heterogeneity analysis by betting market characteristics (per-capita handle, number of licensed operators, advertising intensity) that may help identify the relevant margin.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP), with co-author input from OpenAI Codex via the Duet collaboration system.

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @SocialCatalystLab, @olafdrw
First Contributor: https://github.com/SocialCatalystLab

References

American Gaming Association, "Commercial Gaming Revenue Tracker," Technical Report, American Gaming Association 2023.

Baker, Scott R., Brian Baugh, and Marco Sammon, "Gambling Away Stability: Sports Betting's Impact on Vulnerable Households," *Working Paper*, 2023.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Carpenter, Christopher and Carlos Dobkin, "The Minimum Legal Drinking Age and Public Health," *Journal of Economic Perspectives*, 2011, *25* (2), 133–156.

Cook, Philip J. and Michael J. Moore, "Drinking and Driving," *Journal of Health Economics*, 1993, *12* (2), 131–148.

DeAngelo, Gregory and Benjamin Hansen, "Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities," *American Economic Journal: Economic Policy*, 2014, *6* (2), 231–257.

Grinols, Earl L., "Gambling in America: Costs and Benefits," 2004.

Gruber, Jonathan and Sendhil Mullainathan, "Do Cigarette Taxes Make Smokers Happier?," *The B.E. Journal of Economic Analysis & Policy*, 2005, *5* (1).

Hansen, Benjamin and Glen R. Waddell, "Legal Access to Alcohol and Criminality," *Journal of Health Economics*, 2018, *57*, 277–289.

Humphreys, Brad R. and Victor A. Matheson, "Legalized Sports Betting and Crime," *Eastern Economic Journal*, 2021, *47*, 336–346.

Levitt, Steven D. and Jack Porter, "How Dangerous Are Drinking Drivers?," *Journal of Political Economy*, 2001, *109* (6), 1198–1237.

Ruhm, Christopher J., "Alcohol Policies and Highway Vehicle Fatalities," *Journal of Health Economics*, 1996, *15* (4), 435–454.

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

Swanson, Ana, "Sports Betting Legalization and Health Risk Behaviors: Evidence from the Behavioral Risk Factor Surveillance System," *Working Paper*, 2023.

U.S. Department of Transportation, "Departmental Guidance on Valuation of a Statistical Life in Economic Analysis," Technical Report, U.S. Department of Transportation 2022.

Walker, Douglas M. and A. H. Barnett, "The Social Costs of Gambling: An Economic Perspective," *Journal of Gambling Studies*, 1999, *15* (3), 181–212.

A. Treatment Appendix

Table 6: State-Level Online Sports Betting Treatment Status

State	OSB Launch Date	In FARS Sample?	Post-Treatment Quarters
New Jersey	2018-06-14	Yes	19
Pennsylvania	2019-07-15	Yes	14
Iowa	2019-08-15	Yes	14
West Virginia	2019-08-27	Yes	14
Indiana	2019-10-03	Yes	13
New Hampshire	2019-12-30	Yes	13
Colorado	2020-05-01	Yes	11
Illinois	2020-06-18	Yes	11
Tennessee	2020-11-01	Yes	9
Michigan	2021-01-22	Yes	8
Virginia	2021-01-21	Yes	8
Wyoming	2021-09-01	Yes	6
Arizona	2021-09-09	Yes	6
Connecticut	2021-10-19	Yes	5
New York	2022-01-08	Yes	4
Louisiana	2022-01-28	Yes	4
Kansas	2022-09-01	Yes	2
Maryland	2022-11-23	Yes	1
Ohio	2023-01-01	No	0
Massachusetts	2023-03-10	No	0
Kentucky	2023-09-28	No	0
Maine	2023-11-03	No	0
Vermont	2024-01-11	No	0
North Carolina	2024-03-11	No	0

Notes: States with "No" in the third column launched after FARS data availability (2022 Q4) and are classified as never-treated in all analyses. Post-treatment quarters count from the launch quarter (inclusive) through 2022 Q4. A state is coded as treated starting in the quarter containing its launch date.

B. Standardized Effect Sizes

Table 7: Standardized Effect Sizes

Outcome	ATT	SE	Pre-Tx SD	SDE	95% CI
Alc. crash rate	0.380	0.146	1.715	0.221	[0.054, 0.389]

Notes: SDE = ATT / pre-treatment standard deviation of the outcome among treated states. The 95% CI is for the SDE: (ATT $\pm$ 1.96 $\times$ SE) / Pre-Tx SD.