Online Market Launch and the Distribution of Gambling Risk: Evidence from Online Sports Betting

Daniel M. McCarthy
University of Maryland
dmccar@umd.edu

Wayne J. Taylor
Airbnb
wayne.j.taylor@gmail.com

Kenneth C. Wilbur
University of California, San Diego
kcwilbur@ucsd.edu

April 9, 2026

Abstract

When a legally restricted, predominantly offline product category becomes digitally accessible, who begins participating and how is spending distributed relative to consumers' financial capacity? We study online sports betting market launches using Generalized Synthetic Control methods applied to financial panel data covering 1.2 million individuals in 11 treatment and 12 control states between 2019 and 2023. We measure spending levels alongside the Income-Adjusted Gambling Expenditure Ratio (IGER), which flags individuals whose rolling three-month gambling spending exceeds a threshold fraction of income. Online market launch increases monthly observed spending per individual by \$3.64 and increases the share of individual-months above the 1%-of-income threshold by 0.78 percentage points. The central finding is distributional: dollar spending increases with income but less than proportionally, so

IGER increases substantially more for lower-income consumers — a disparity that widens at higher thresholds. Two decompositions reveal dynamics that the roughly stable aggregate effects mask. First, new users account for roughly one-fifth of the participation increase but less than 12% of spending; their contributions decay within months while returning-user spending persists and grows. Second, about 42% of the IGER treatment effect reflects repeated threshold crossing in consecutive non-overlapping three-month windows; this persistent component grows throughout the observation period and concentrates among lower-income and pre-launch-active consumers, while one-off crossings attenuate. These estimates capture observed activity within regulated digital channels. The results suggest that market expansion and income-relative risk exposure follow different paths across consumer segments, with implications for segment-specific consumer protection policy.

1 Introduction

In 2018, the U.S. Supreme Court struck down the Professional and Amateur Sports Protection Act (PASPA), a federal law that had prohibited most states from passing new sports betting regulations since 1992. States subsequently launched regulated online sports-betting markets, and online and mobile channels now account for the large majority of legal U.S. sports wagering activity (American Gaming Association 2026, Legal Sports Report 2025, Reynolds 2025). Because daily fantasy sports (DFS) already operated in the vast majority of U.S. states before these launches, what these launches represent is broader market opening: the simultaneous arrival of legal access, new operators, and expanded product offerings in a category where online sports betting had been federally prohibited.

This setting lets us study a question relevant to both marketing and policy: when a legally restricted, predominantly offline category becomes digitally accessible through an online market launch, who begins participating, and how do gambling spending and gambling spending as a proportion of income vary across consumer segments? Importantly, the consumers who generate the most incremental dollars need not be the consumers whose activity rises the most relative to income. We address this question using a large individual-level financial panel covering 1.2 million U.S. consumers' spending activity across thousands of merchants from 2019 to 2023, including online sports-gambling (OSG) operators,[1] as well as observed incomes. This lets us measure how much spending changes following online market launch, as well as how those changes relate to each consumer's financial capacity.

Our central finding is that the consumers who most drive spending growth are not the same consumers who most increase their spending relative to income. To measure the latter, we construct the Income-Adjusted Gambling Expenditure Ratio (IGER), which flags individuals whose rolling three-month OSG spending exceeds a threshold fraction of income (e.g., 1% for IGER_1). Applying Generalized Synthetic Control (GSC) methods to construct counterfactual trends from states that did not launch regulated online sports betting during our observation window, we estimate that online market launch increases observed monthly OSG spending per individual by \$3.64 (from \$0.97 to \$4.61, an approximately 375% increase; this corresponds to \$27.12 per month among the 13.4% of individuals who ever spend) and increases IGER_1 by 0.78 percentage points (from about 0.17% to about 0.95%, an approximately 456% increase). Dollar spending increases rise with income but less than proportionally, so IGER_1 rises substantially more for lower-income consumers. A new-versus-returning- user decomposition shows that first-observed users account for roughly one-fifth of the participation increase but less than 12% of the spending increase, and their contributions decay within months, while returning-user activity persists and continues to grow. The roughly stable aggregate spending effect therefore masks a composition shift from front-loaded first-observed activity toward a larger share of spending volume coming from previously observed users. A persistence decomposition unpacks the IGER_1 result further: 42% of the treatment-effect increase reflects repeated threshold crossing across consecutive non-overlapping three-month windows, and the treatment effect on this persistent component grows throughout the observation period. The remaining 58% reflects one-off threshold crossings whose treatment effect spikes at launch and then attenuates. Beneath the roughly stable aggregate IGER_1 effect, then, there are two diverging trajectories, a fading transitory component and a growing persistent one, that would not be apparent from the aggregate result alone. We briefly examine a small set of contextual outcomes in supplementary analyses.

These findings arrive amid growing concern about gambling-related financial harm. National problem-gambling helpline contacts, including texts and calls longer than 60 seconds, grew by 351% from 30,465 in 2018 to 137,323 in 2025.[2] The share of U.S. adults who say that legal sports betting is bad for society rose from 34% in 2022 to 43% in 2025, with especially large increases among younger men (Gramlich 2025). Against this backdrop, understanding how market expansion distributes both spending growth and elevated spending-to-income ratios across consumer segments is central to evidence-based regulatory policy.

To construct IGER, for each individual and month, we classify whether their OSG spending over the current and prior two months exceeds a given percentage of their income over the same three-month window (e.g., 1%, 5%, or 10%). We then aggregate to the state-month level as the share of individuals above that threshold. We use a rolling three-month window because it increases the economic importance of threshold crossing relative to a single-month measure and smoothes transitory fluctuations in both income and spending. To illustrate, an individual earning $70,000 per year, which is near the low-income tercile boundary (and thus at the upper end of the low-income group) in several treatment states (Web Appendix Table 2), would cross the IGER_1 threshold by spending more than $175 cumulatively over a three-month window. For individuals near the middle-income tercile boundary (approximately $130,000, at the lower end of the high-income group), the corresponding 1% threshold is approximately $325 over three months.

A large body of research shows that gambling-harm incidence increases disproportionately with the percentage of income spent on gambling (Young et al. 2024). Cross-national studies covering Canada, Norway, Sweden, and the U.K. consistently identify 1% of gross income as the threshold at which gambling-related harms, including financial distress, relationship problems, and adverse health outcomes, begin to accelerate sharply (Canadian Centre on Substance Use and Addiction 2021, Jonsson et al. 2022a, Muggleton et al. 2021b). We follow this literature and define our primary measure, IGER_1, using the 1% threshold. The income-based disparities we document become more pronounced at higher thresholds (5% and 10%), indicating that our results are conservative with respect to the threshold choice.

Because our income measure is a vendor-estimated gross-income figure inferred from direct-deposit patterns rather than observed on tax forms, IGER should be interpreted as a tractable transaction-data operationalization of an income-based benchmark, not as a clinical diagnosis or a direct measure of realized harm. It captures observed spending-to-income ratios within regulated, digitally tracked channels but does not measure net welfare; any broader welfare assessment would also require quantities we do not observe, including consumer surplus and harms outside the measured channel. We present new evidence on channel substitution in Section 7. We use IGER alongside three decompositions—first-observed-versus-returning activity, income group, and persistence of threshold crossing—to study how market expansion is distributed.

Estimating the causal effects of online market launch is challenging. Randomized experimentation is neither feasible nor ethical. Sports gambling fluctuates over time with strong seasonality tied to major sporting events, and launch timing varies across states. Data limitations also matter: historically, much sports betting was cash-based and difficult to track. Today, however, most legal sports wagers are placed through online and mobile platforms (Fisher 2025), so our focus on regulated, digitally observed channels that require electronic payments captures the large majority of regulated sports betting activity. We do not observe informal, illegal, or cash-based gambling activity, so the estimates should be interpreted as effects within regulated digital channels and not total gambling across all channels.

The financial panel dataset is sourced from Consumer Edge and tracks card spending and account transfers within a balanced panel of 1,201,571 regular card users ("individuals") over the observation period. The data source can attribute multiple credit and debit cards, across multiple card networks, to the same individual. Further details about the panel construction appear in Web Appendix A. We observe panel members' spending at the 20 largest online sports-gambling operators, along with a vendor-estimated gross-income measure that Consumer Edge infers from observed direct-deposit patterns, and the state in which each transaction occurred, by individual and month. The spending-related outcome variables sum gambling behaviors occurring within each state and month, divided by the number of panelists that lived in that state in that year, to express gambling activity per state-assigned panelist; Web Appendix A1 details the construction and its rationale. Our analysis proceeds at the “full panel level”: we analyze outcomes across all individuals observed within the panel in treatment and control states, regardless of whether they gambled. The estimand therefore applies to all active cardholders in our panel, rather than only to the self-selected gamblers in the population.[3]

The panel has several attractive features for this setting. The distribution of annual household incomes observed within the panel resembles the U.S. Census Bureau’s Current Population Survey (CPS), and Kim and McCarthy (2023) show that merchant-level revenue measures from the same card spending panel closely track firms’ disclosed revenues. Because regulated online sports betting requires digital payment methods, the panel’s composition, regular users of tracked credit and debit cards, aligns with the population that can access the market we study. That said, like other major consumption panels, panel membership may not fully represent the broader population. For more details regarding panel construction and external validity, see Web Appendix A and Web Appendix A2.

We use the GSC model to estimate causal effects within a principled quasi-experimental framework. The GSC approach generates counterfactual predictions for states that launched online sports betting by comparing them to weighted combinations of control states that did not launch regulated online sports betting during the sample period. This method is particularly well-suited for settings with small numbers of treated units with sufficiently long pretreatment time series, such as ours, because it constructs a tailored counterfactual prediction for each treated unit, enhancing the accuracy of the comparisons. Staggered difference-in-differences estimators yield point estimates that are close in magnitude and identical in sign across all outcomes, showing that the primary findings are robust across estimators.

The rest of the paper is organized as follows. We provide an overview of related literature, discuss the data and empirical context, present our model specification, analyze the results and robustness checks, then conclude with a discussion of implications and future work.

2 Related Literature

Our paper contributes to research on how policy and market design shape consumption in categories with meaningful downside risk. Recent marketing literature documents how public policy has impacted marketing strategies and consumer behaviors in product categories with downside health risks, including tobacco and sweetened beverages (Goli et al. 2023, Anand and Kadiyali 2024, Seiler et al. 2020, Bollinger and Sexton 2023). We study a different but related margin: a policy change that expands digital access to a risky category while simultaneously making behavior observable through electronic payments. We estimate policy effects on observed spending and spending-to-income exposure within regulated digital channels, with supplementary contextual evidence on related outcomes. Our work is informed by and contributes to three distinct literatures.

First, recent work has studied how gambling policy decisions lead to important downstream financial consequences for consumers. Using a large consumer-credit panel, Hollenbeck et al. (2026) show that sports betting legalization reduces credit scores and increases bankruptcies, debt sent to collections, use of debt consolidation loans, credit card delinquencies, and auto loan delinquencies, with especially large effects in states that introduced online or mobile access. Using household transaction data, Baker et al. (2026) show that legalization expands betting participation and frequency, crowds out savings and brokerage deposits, and increases debt stress among financially constrained households. Most recently, Goss and Mangrum (2026) document large increases in sportsbook spending, cross-border spillovers, and deteriorating credit performance after legalization. Using internet search trends, Yeola et al. (2025) show that sports betting legalization correlates with significant increases in gambling-addiction help-seeking, with online launches driving larger and more persistent effects than retail-only openings. We complement this literature by directly measuring gambling intensity relative to income within regulated markets, as the gamblers who spend the most relative to income are those at greatest risk of experiencing gambling-related harms. In particular, spending rises broadly across income groups, but elevated gambling spending-to-income behavior rises disproportionately among lower-income consumers and remains especially persistent among individuals with observed pre-launch OSG activity. We also decompose market expansion into contributions from first-observed and previously observed users, showing that launch-period adoption is front-loaded while persistent spending growth is driven by previously observed users.

Second, the public health literature has carefully documented how gambling spending relates to gambling-related harms. Independent researchers have converged on the sensible convention that gambling-related harm incidence is strongly predicted by gambling spending relative to gambler income, and that particular percentages of income are associated with accelerating risks of harms. Muggleton et al. (2021a) analyzed a bank transaction dataset from 6.5 million U.K. individuals and found that gambling expenditure as a proportion of monthly income predicted financial distress, unemployment, disability, and at the highest levels, substantially elevated mortality. Hodgins et al. (2023) linked gambling spending to income ratios to gambling harm incidence in 11 representative survey datasets from eight countries, finding that expenditure as a percentage of gross household income consistently discriminated between lower-risk and higher-risk gambling across national contexts. This analysis formed the empirical basis for Canada's Lower-Risk Gambling Guidelines, which recommend that individuals gamble no more than 1% of household income per month. Jonsson et al. (2022b) replicated these risk-curve analyses using objective spending data from over 35,000 customers of Norway's state gambling monopolist, confirming that the income-relative spending thresholds identified in representative surveys generalize to operator-recorded gambling behavior. Langeland et al. (2022) further showed that optimal expenditure cutoffs for classifying moderate-risk and problem gamblers differ across income groups, with lower-income individuals reaching high-risk thresholds at lower absolute spending levels, underscoring the importance of measuring gambling intensity relative to income rather than in absolute terms. We offer the first study to report how much online sports betting legalization increased gambler spending above the one-percent-of-income threshold, and also report changes at higher thresholds of five and ten percent of income. We want to emphasize that IGER is not a direct measure of realized harm; rather, it operationalizes a harm-relevant spending benchmark and therefore enables us to study how online market launch changes the share of consumers whose gambling spending relative to income places them in categories that prior research has found significantly more likely to experience gambling harms.

Finally, our findings connect to the literature on lotteries, much of which links gambling demand to consumer income. Lockwood et al. (2024) show that cross-sectional lottery spending decreases with income. Hansen et al. (2026) study a large national lottery price change, from $2 to $5, which led to 72% lower ticket quantity and 30% lower revenue. Lower-income consumers were the most price-sensitive and the least responsive to the larger prize structure, so lottery spending as a share of income fell most in the poorest neighborhoods. Larsen et al. (2025) show that online sports gambling legalization cannibalizes state lottery sales by 5–7%, driven by intensive-margin reductions among frequent lottery buyers, with the most substitution occurring in lower-income neighborhoods. We add to this literature by reporting that online sports betting market launch increases gambling-to-income ratios most among lower-income consumers, even though it increases average spending most among high-income consumers. The distinction is policy-relevant because market growth and risk exposure differ between income strata, suggesting that the highest-spending gamblers may not be at the most risk of gambling harms, and that different groups may present different cost-benefit tradeoffs and need different consumer-protection policies.

3 Data and Context

3.1 Treatment and Control States, Treatment Timing and Types

The Supreme Court's 2018 PASPA decision enabled states to launch new sports betting markets. States' legalization policies varied in their timing, licensees, and online/offline modalities.

We focus our analysis on the largest group of states with the most comparable policies for which we can estimate treatment effects. Among states that launched online sports betting during the observation window (Jan 2019–Sep 2023), we include all states with (i) at least 12 months of stable observed pre-launch betting data, to ensure a reliable pre-launch seasonal baseline against which post-launch outcomes may be measured; (ii) licensed commercial operators with national scope, as opposed to state-owned entities, which enables observation of gambler spending in the dataset; and (iii) online sports betting launched without concurrent online casino gambling (also known as iGaming), allowing us to isolate the online sports betting launch effect. These criteria yield the 11 treatment states in Table 1; Web Appendix C details each exclusion.

Table 1: States and Launch Timings

Treatment State	**Offline Launch**	**Online Launch**	**Treatment Category**
Illinois	March 2020	June 2020	OWO
Colorado	June 2020	May 2020	OWO
Virginia	January 2021	January 2021	OWO
Wyoming	September 2021	September 2021	OWO
Arizona	September 2021	September 2021	OWO
Louisiana	October 2021	January 2022	OWO
Kansas	September 2022	September 2022	OWO
Ohio	January 2023	January 2023	OWO
Arkansas	July 2019	March 2022	OAO
New York	July 2019	January 2022	OAO
Tennessee	*n.a.*	November 2020	OO

Control states: Alabama, Alaska, California, Georgia, Idaho, Minnesota, Missouri, Oklahoma, South Carolina, Texas, Utah, Vermont

Note. ‘Offline’ represents offline sports betting; ‘online’ represents online sports betting. ‘Treatment Category’ abbreviations: ‘OWO’ represents online-with-offline (states that launched both online and offline sports betting within 3 months of each other), ‘OAO’ represents online-after-offline (states that launched online sports betting more than 3 months after offline sports betting had launched), and ‘OO’ represents online-only (states that launched online sports betting without offline sports betting). Control states used in the analysis did not launch sports betting during the observation period. The balanced panel contains 1,201,571 individuals distributed across 23 states.

The treatment states include three legalization policy types: states that launched online and offline (in-person) sports betting nearly concurrently—within three months (“online-with-offline”); states that launched online after offline betting was already legal (“online-after-offline”); and states that launched online only (“online-only”).[4] We primarily use the treatment classifications to estimate how different online launch policy types relate to observed card spending at offline casinos.

Interpretation is most straightforward as category expansion in online-only and online-with-offline states, where new observed online activity cannot reflect substitution from an already-legal in-state offline market.[5] Interpretation is more composition-sensitive in online-after-offline states because online launch may partly reallocate gambling from existing legal retail betting. Overall, the online channel dominates retail gambling, as approximately 95% of the $147.9 billion in sports betting was gambled in online channels in 2024 (Legal Sports Report 2025, Reynolds 2025).

State treatment timing is determined by online market launch dates, that is, the dates when consumers could begin placing sports bets online. The estimates should therefore be interpreted as the effect of regulated online market launch and rollout as implemented, not the effect of statutory legalization alone. The estimated launch effect bundles legal access, operator entry, product availability, promotional offers, and post-launch marketing activity.

The post-2018 variation in launch timing appears to be driven by idiosyncratic differences in legislative calendars and implementation schedules. Recent work finds no evidence that employment growth, wage growth, total private income, weekly wages, or COVID-19 case rates predict the timing of online gambling legislation (Baker et al. 2026, Hollenbeck et al. 2026). While Hollenbeck et al. (2026) find that states legalizing sports gambling tend to have more generous unemployment insurance policies, they argue this would likely lead to understating rather than overstating any negative financial impacts. We extend this analysis by examining additional potential confounders not previously studied in this context, including state financial health indicators and political composition. In Web Appendix D, we find no evidence that online market launch timing correlates with state financial deficits, debt, political affiliations, wage levels, or unemployment rates.

We compare the 11 treatment states to all 12 control states that did not allow sports betting and did not change gambling policies during the observation period.[6] Figure 1 illustrates the geographical distribution of treatment and control states in our analysis. The 11 treatment states are geographically dispersed.

It is important to note that control-state residents may physically travel to treatment states to

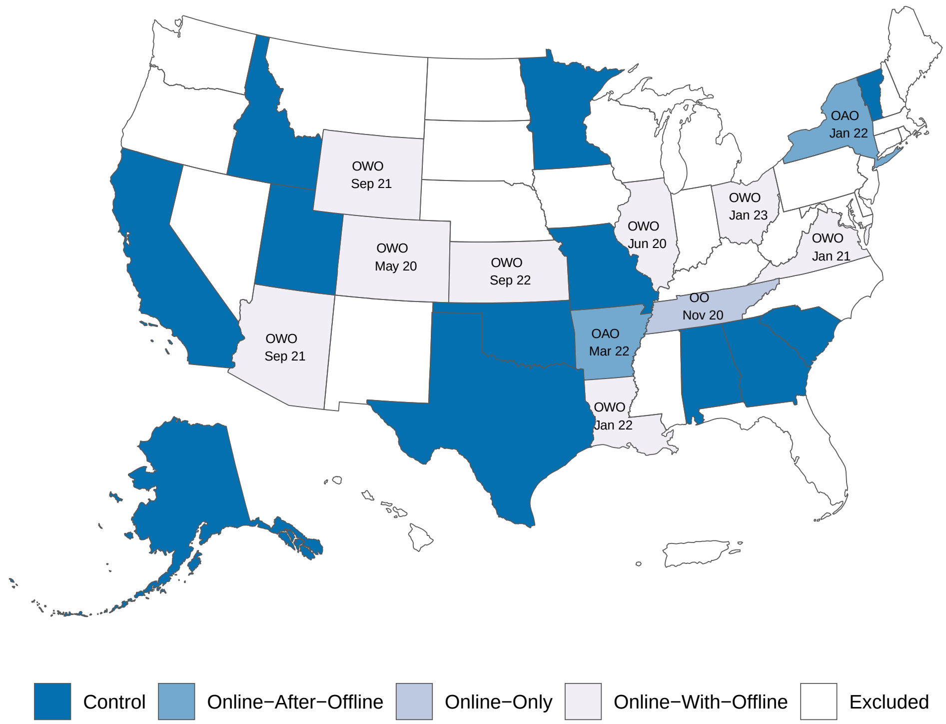

Figure 1: Geographic Distribution of Treatment and Control States by Market Launch Pattern.

Note. Map shows the 23 states included in our analysis, categorized by their sports-betting market-launch pattern during the observation period (January 2019 – September 2023). Online-with-offline ('OWO') states launched both online and offline (retail) sports betting within three months of each other. Online-after-offline ('OAO') states launched online sports betting more than three months after legalizing retail sports betting. Online-only ('OO') states legalized only online sports betting without retail sports betting. Control states did not legalize sports betting during the observation period. Excluded states are ignored in the analysis.

gamble. As explained in Web Appendix A1, the spending-related measurements reflect spending in states where gambling occurred, rather than states of gambler residence. Hence, if a control-state resident gambles while physically located in a treatment state, that activity enters the treatment state's spending measures, not the resident's home-state series, leaving control-state measurements unaffected. A related concern is that control-state individuals could use Virtual Private Network (VPN) services to gamble online in treated states. However, major gambling operators prohibit cross-state VPN usage, use a variety of sophisticated techniques to proactively detect illicit VPN-based gambling, and warn bettors that they will not honor illegal winning bets, as paying out such bets would reduce operator profits and may imperil their operating licenses.

3.2 Descriptive Trends

We first examine model-free evidence on how online market launch corresponds to observed gambling activity over time, focusing on our two key outcome metrics: monthly OSG spending per individual and IGER_1. For IGER_1, each individual-month is classified as above threshold if cumulative OSG spending over the current and prior two months exceeds 1% of cumulative income over the same three-month window; we then average this indicator across individuals within each state-month. For both metrics, we compare trends between states that launched regulated online sports betting (treatment states) and those that did not (control states). To align observations across states with different launch dates, we index time relative to each treatment state's online market launch, with month zero representing the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month zero may reflect partial exposure. We also include a line reflecting all control states, with month zero reflecting the midpoint of the sample period, to show the typical variation across a two-year time series in non-launch states.

The leftmost panel of Figure 2 shows the evolution of monthly gambling spending per individual within treated and control states. We observe a stark contrast: while control states maintain relatively stable spending levels throughout the observation window, treatment states exhibit a pronounced increase in gambling activity immediately following online market launch. This elevated spending persists well after the initial implementation, suggesting a lasting shift in behavior.

The rightmost panel of Figure 2 presents the corresponding comparison for IGER_1. Here too, we observe a clear divergence: treatment states show a substantial and sustained increase in IGER after online market launch, while control states maintain consistent levels throughout. These descriptive plots are illustrative only. Because treated states are aligned in event time whereas the control series is shown in calendar time, the figure is not the basis for our pretreatment-fit assessment. Pretreatment fit is evaluated using the dynamic ATT estimates in Figures 4 and 5 and the state-level diagnostics in Web Appendix F.

Taken together, these descriptive patterns motivate a causal analysis but do not identify treatment effects on their own. Pre-existing DFS activity, seasonality around major sporting events, and

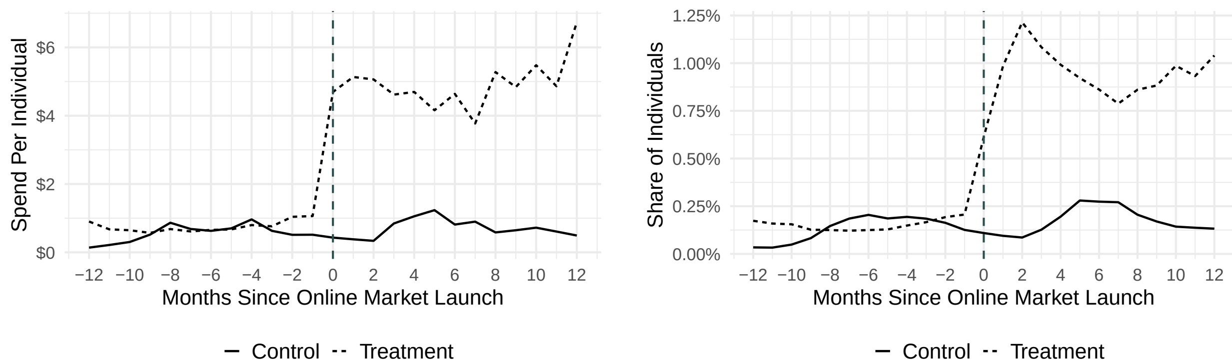

(a) OSG

(b) IGER_1

Figure 2: Raw monthly online sports gambling (OSG) spending per individual (left) and IGER_1 (right)

Note: Left panel represents average observed OSG spending per individual across all individuals within the credit/debit card panel, for individuals with transactions in states that launched regulated online sports betting (dashed line) versus in states that did not (solid line). For each treatment state, period 0 is the time of online market launch. For control states, we define period 0 as the midpoint of the sample period to show the typical variation in non-launch states over a two-year period. Right panel represents the corresponding figure for IGER_1, defined as the proportion of individuals whose cumulative OSG spending over months t, $t-1$, and $t-2$ exceeded 1% of cumulative income over the same three-month window. Web Appendix Figure W5 presents analogous descriptive trends for IGER_5 and IGER_{10}.

time-varying state conditions could all influence post-launch outcomes. We therefore use GSC to construct counterfactuals that flexibly account for these factors.

4 Model Development

4.1 Generalized Synthetic Control

We estimate causal effects using the Generalized Synthetic Control (GSC) approach. This method offers three key advantages: it incorporates interactive fixed effects that capture both state-specific unobservable factors (such as cultural attitudes toward gambling) and their interactions with time-varying influences (such as macroeconomic conditions or major sporting events); it replaces the parallel trends assumption required by traditional difference-in-differences methods with the more flexible interactive fixed effects structure; and it provides a framework for leveraging data from multiple treated states with varying implementation timing. This flexibility is crucial in our setting, where states enacted similar policies at different times.

Following Bai (2009), Xu (2017), and Li and Sonnier (2023), we model the outcome Y_{it} for each state i, month t as:

$$Y_{it} = \delta_{it}D_{it} + X'_{it}\beta + \lambda'_i F_t + \mu_i + \mu_t + e_{it}$$

Here, D_{it} is a binary treatment indicator equal to 1 for treated states in post-treatment periods and 0 otherwise. δ_{it} represents the treatment effect, X_{it} is a vector of observed time-varying state-level controls with coefficients β, F_t is a vector of time-varying latent factors common to all states, with λ_i being the state-specific factor loadings. μ_i and μ_t are state- and time-specific fixed effects.

To account for the potential confounding effects of the COVID-19 pandemic, which occurred during our observation period, we also include the following time-varying state-level controls: monthly unemployment rates from the Bureau of Labor Statistics; Google Community Mobility data capturing percent changes from baseline across six categories (retail and recreation, grocery and pharmacy, parks, transit stations, workplaces, and residential mobility); monthly COVID-19 case counts per capita from The New York Times COVID-19 database; and monthly COVID-19 death counts per capita from the CDC's Provisional COVID-19 Deaths dataset. These controls help isolate the effects of online market launch from concurrent pandemic-related disruptions to economic activity, consumer mobility patterns, and public health conditions that could independently affect gambling behaviors.

The estimation proceeds in the usual three steps:

1. Estimate common factors and fixed effects, and control states' factor loadings, using control state data alone.

2. Estimate treated states' factor loadings using pretreatment data.

3. Calculate treatment effects as post-treatment differences between treated states' observed outcomes and predicted counterfactuals.

We select the number of latent factors using cross-validation on control-state data, following Xu (2017) and Li and Sonnier (2023). Treatment states launch at different times, so we use each treated state's pretreatment period to estimate its factor loadings $\hat{\lambda}_i$. Importantly, the GSC estimator never uses any treated state as a control for any other treated state in any period, ensuring that our counterfactual predictions for each state are constructed entirely from never-treated units. We estimate state/time treatment effects as differences between each observed and predicted outcome variable within each treated period, and then estimate average treatment effects by averaging across treated states and time periods. Results are similar when we use panel size-weighted averages instead (see the Robustness Checks section and Web Appendix E12). Under both weighting procedures, we follow Li and Sonnier (2023) and construct 95% confidence intervals from the factor-robust asymptotic-normal estimator, which accounts for both latent-factor estimation error and serially correlated idiosyncratic shocks.

5 Results

5.1 Impact on Gambling Spending

Figure 3a shows that online market launch increased monthly observed OSG spending per individual by an estimated \$3.64, from an estimated \$0.97 in the absence of launch to \$4.61 after launch. The point estimates show that high-income individuals increased spending most (by \$4.02), followed by middle-income (\$3.82) and low-income individuals (\$2.64). The 95% confidence intervals for middle- and high-income groups overlap, but the low-income estimate is significantly below both (see Web Appendix Table W7). Because gambling spending increases with income but less than proportionally, the income-adjusted effect is largest for low-income consumers, which is a central finding we expand upon in the IGER analysis below.

These estimates are population averages that include the large majority of individuals who never gamble online. To put the spending effect in perspective, 13.4% of individuals in treated states spent with an OSG operator at least once during the post-launch window.[7] If the entire causal spending lift accrues to this group, the \$3.64 population-average ATT corresponds to approximately \$27 per month, or roughly \$324 per year, per ever-active gambler.

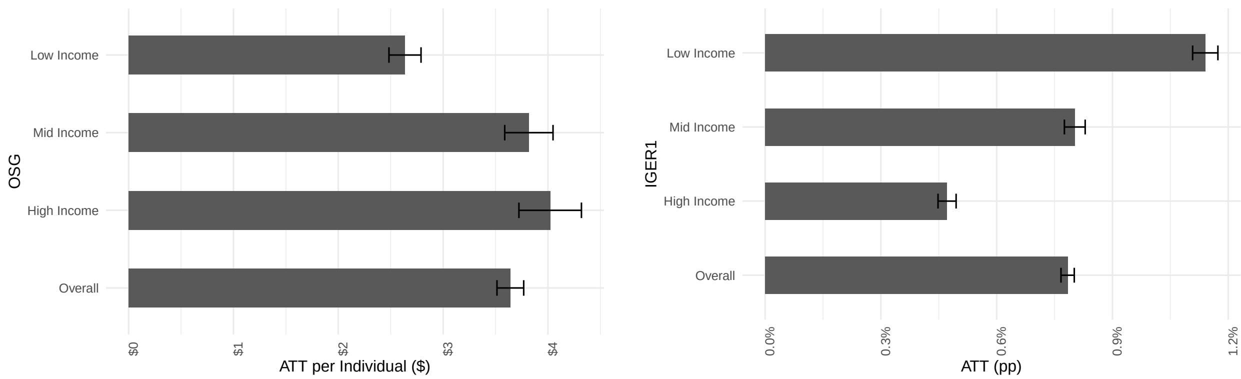

(a) OSG Spending

(b) IGER_1

Figure 3: **OSG Spending and IGER_1 ATT by Income Group**

Note. Panel (a) shows the average treatment effect on monthly OSG spending per individual; panel (b) shows the effect on IGER_1, the share of individual-months in which rolling three-month OSG spending exceeded 1% of rolling three-month income. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Income groups are defined by state-specific terciles.

Dynamic treatment effects (Figure 4) show no significant pretreatment trends, with confidence intervals including zero in every pretreatment month (in Web Appendix F we show that parallel pretrends hold at the state-specific level as well). Post-treatment, spending increases sharply to $5.47 in the first full month after launch (month 1), then stabilizes at approximately $3.55 per month, indicating a large, persistent effect.

5.2 Impact on Income-Adjusted Gambling Expenditure Ratio (IGER)

Figure 3b shows the impact of online market launch on IGER_1, the measure of elevated observed gambling spending-to-income over a rolling three-month period; see Web Appendix Table W8 for the estimates in tabular form. Online market launch increases IGER_1 by 0.78 percentage points, from about 0.17% to about 0.95% of individual-months. IGER_1 indicates gross spending at regulated operators exceeding 1% of consumer income, so these estimates capture money committed to gambling in observed channels rather than realized net losses; Web Appendix E11 shows directionally similar results using net spending, with an overall IGER_1 effect of 0.57 percentage points rather than 0.78.[8]

The increase is larger for lower-income consumers in both absolute and relative terms: IGER_1

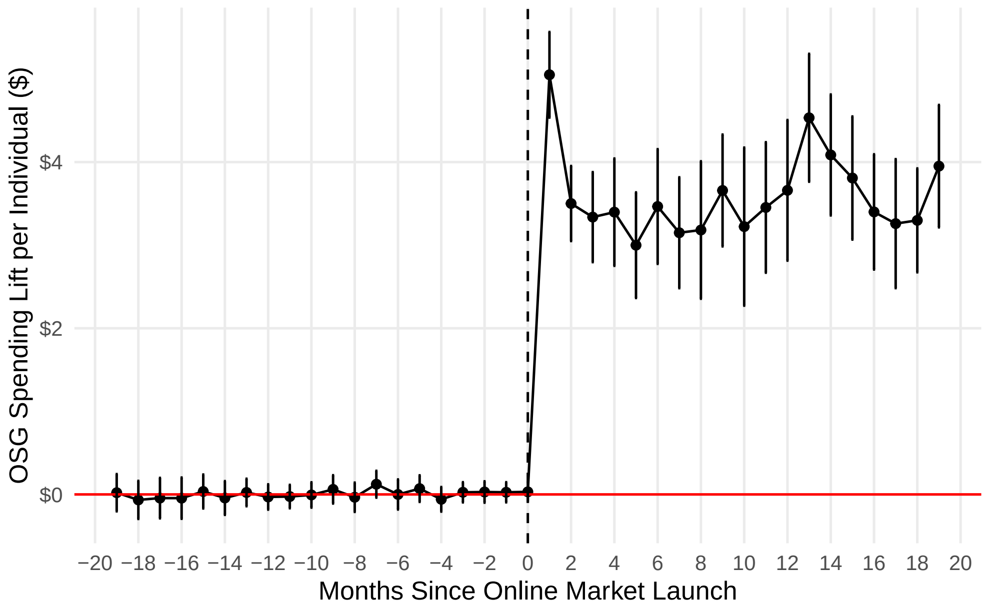

Figure 4: **Online Sports Gambling Spending Dynamic ATT**

Note. Dynamic treatment effects over time for online sports gambling (OSG) spending per individual. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Post-treatment observation spans 9 months for Ohio and 13 months for Kansas; all other states have at least 19 months post-treatment.

rises by 1.14, 0.80, and 0.47 percentage points for low-, middle-, and high-income groups. These disparities widen at higher IGER thresholds. At the 5% threshold, the low-income ATT is 1.8–4.4 times the middle- and high-income ATTs; at the 10% threshold, the corresponding multiples are 1.9–4.6 (Web Appendix E4, Figure 6). The income gradient in spending-to-income exposure thus intensifies at thresholds associated with more severe gambling harm.

The increase is also greatest in relative terms for low-income individuals. Measured as a percentage of what $IGER_1$ would have been without online market launch, the increase is 475% (from 0.24% to 1.38%), 421% (from 0.19% to 0.99%), and 336% (from 0.14% to 0.61%) for low-, middle-, and high-income individuals, respectively.

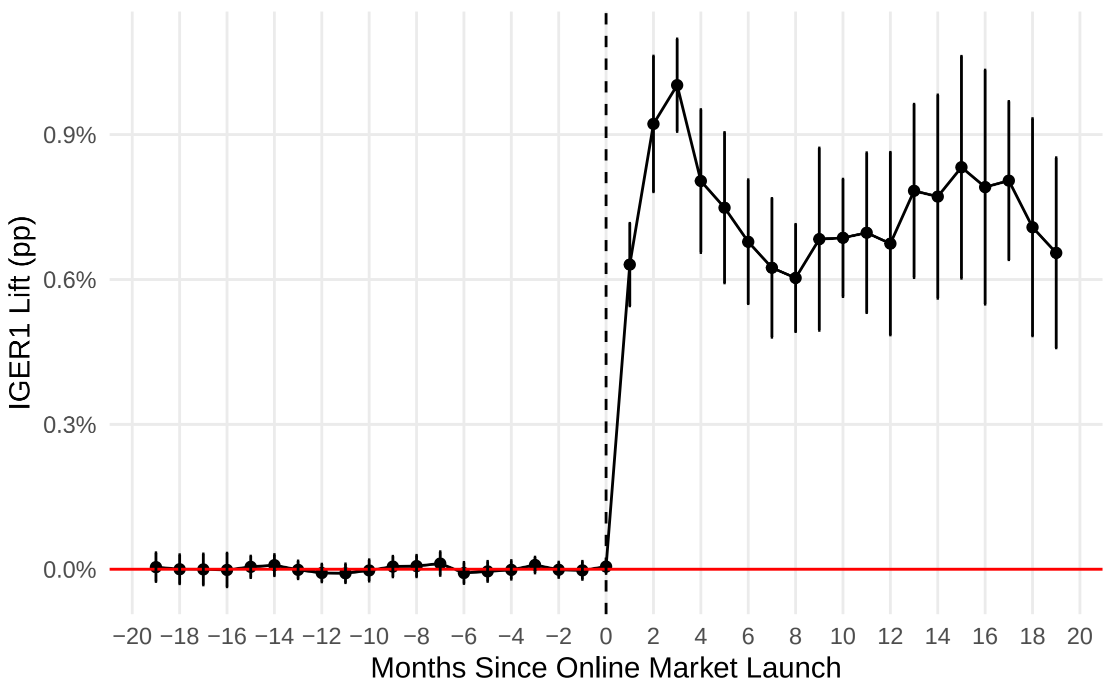

Figure 5: **$IGER_1$ Dynamic ATT**

Note. Dynamic treatment effects on $IGER_1$ over time. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Post-treatment observation spans 9 months for Ohio and 13 months for Kansas, all other states have at least 19 months post-treatment.

Dynamic effects (Figure 5) show no evidence of significant pretreatment effects. Post-treatment, $IGER_1$ rises sharply, peaking at approximately 1.00% around month 3 before settling at approximately 0.76% from month 13 onward. The rolling three-month window mechanically smooths the initial ramp-up, but the persistence well beyond the window length reflects sustained behavioral change. We observe similar time-varying patterns across income groups, although at different levels.

6 Decomposing Market Growth and Spending-to-Income Exposure

The average treatment effects above establish that online market launch increases mean OSG spending and spending-to-income ratios. This section decomposes those effects, in the sense of partitioning an outcome variable into mutually exclusive, collectively exhaustive components defined by an observable criterion; then estimating the causal effect on each component separately using the same GSC identification strategy, in order to reveal heterogeneous dynamics that the aggregate treatment effect masks.

We decompose the average treatment effects using three variables. The first two—a new-versus-returning-user decomposition separating newly active from returning users, and an income decomposition attributing the spending lift across income terciles—characterize who accounts for the growth in OSG spending. The third, a persistence decomposition, shifts the spending-to-income ratio and distinguishes repeated from one-off elevated IGER exposure. Each component is estimated in a separate GSC model, so the component ATTs need not sum exactly to the unified estimate; we report both for transparency. In practice, component sums are within a few percent of the direct estimates.

6.1 Effects on Spending: Newly Active vs. Returning Users

We classify each active individual-month as "newly active" (equivalently, first-observed) if the current month contains the individual's first observed OSG spend, and as "returning" (previously observed) if the individual spent in any prior month. Figure labels use the shorter "new" and “returning” terminology for readability. Online market launch increased the share of individuals with any OSG spending by 1.61 percentage points, from 0.51% to 2.12%. Newly active users account for 22% of this participation lift; returning users account for the remaining 78%. The mean spending effect decomposition shows an even more lopsided split: 88% of the incremental spending comes from returning users.

The dynamic effects in Figure 6 reveal how these contributions evolve over time. New-user spending spikes sharply in the first post-treatment month and decays toward zero within roughly eight months. Returning-user spending rises immediately and continues growing throughout the observation window. The roughly stable aggregate spending ATT therefore masks a composition shift: as new-user spending fades, returning-user spending increases, so the total appears flat even though the underlying mix is changing. A corresponding decomposition of participation rates, showing the same front-loaded acquisition and persistent returning-user growth, is presented in Web Appendix E2.

6.2 Effects on Spending: Contribution by Income Tercile

We next ask which income groups account for the aggregate spending lift. Each individual is assigned to a state-specific income tercile based on average gross income, and each group’s spending ATT is expressed per panelist so the three components sum to the total.

High-income individuals account for the largest share of the incremental spending lift (44.0%), followed by middle-income (32.8%) and low-income (23.3%). The composition is similar for participation: high-income individuals account for 40.3% of the active-user lift, middle-income for 34.8%, and low-income for 24.9% (Figure 7). High-income individuals account for the largest share of both incremental spending volume and active-user growth. Because the component shares are estimated in separate models, we interpret this pattern descriptively rather than as a precise disproportionality statement. Dynamic spending effects by income group, showing broadly similar temporal profiles across terciles, are presented in Web Appendix E3.

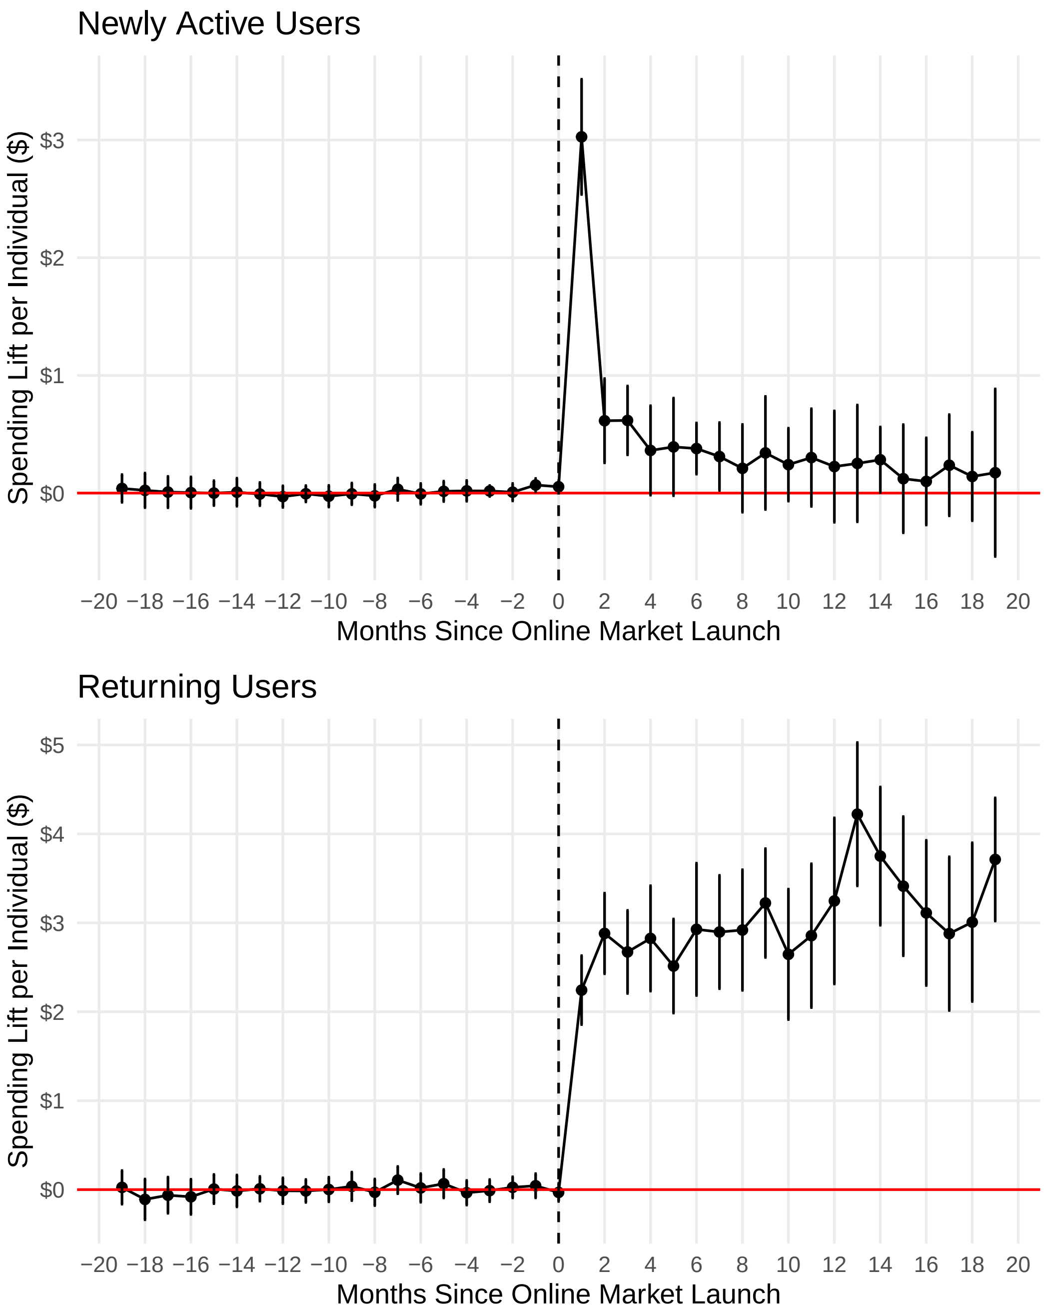

Figure 6: Dynamic Spend Decomposition: Newly Active vs. Returning Users

Note: Top panel shows the dynamic treatment effect on OSG spending per individual attributable to newly active users (first observed OSG spending in the current month). Bottom panel shows the corresponding effect for returning users (prior observed OSG spending in an earlier month). Both panels share the same y-axis scale. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Post-treatment observation spans 9 months for Ohio and 13 months for Kansas; all other states have at least 19 months post-treatment.

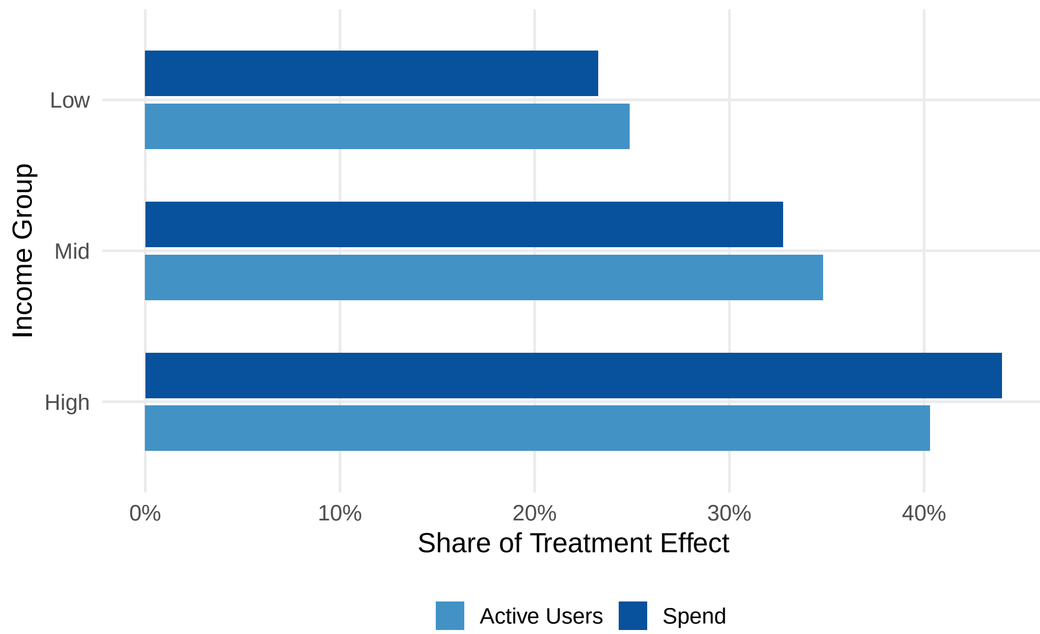

Figure 7: Income Group Shares of Incremental OSG Activity and Spend

Note: Bars show each income tercile's share of the total treatment effect on active users (light) and OSG spending (dark). Income terciles are defined within each state using time-invariant average gross income. Each tercile comprises approximately one-third of the panel. Shares are computed from separately estimated GSC models; because estimation is independent, shares may not sum to exactly 100%.

6.3 Persistence: Repeated vs. One-Off Threshold Crossing

The new-versus-returning-user and income decompositions characterize who spends and how much. We now ask whether the increase in elevated spending-to-income exposure reflects repeated high-ratio behavior or one-off launch-period spikes. Because IGER_1 uses rolling three-month windows, a single large spending amount can mechanically trigger the indicator for up to three months. We therefore define persistence using consecutive non-overlapping three-month windows:

$$\text{Persistent IGER}_{1,it} = \mathbf{1}\left[\text{IGER}_{1,it} = 1 \text{ and } \text{IGER}_{1,i,t-3} = 1\right],$$

$$\text{Non-persistent IGER}_{1,it} = \mathbf{1}\left[\text{IGER}_{1,it} = 1 \text{ and } \text{IGER}_{1,i,t-3} = 0\right].$$

Because the windows $\{t, t-1, t-2\}$ and $\{t-3, t-4, t-5\}$ are disjoint, a single spending spike cannot mechanically generate Persistent $\text{IGER}_1 = 1$; the same individual must cross the threshold in two distinct three-month windows.

Decomposing the $IGER_1$ treatment effect into its persistent and non-persistent components, we find that 42% reflects persistent threshold crossing while the other 58% reflects non-persistent threshold crossing (the component ATTs are 0.35 and 0.49 percentage points, respectively, summing to 0.84 percentage points; the difference from the directly estimated 0.78 percentage points total reflects independent factor selection across models). The dynamic treatment effects in Figure 8 reveal why this distinction matters. Non-persistent $IGER_1$ spikes sharply in the first three post-treatment months and then attenuates, consistent with launch-period trial activity that does not sustain. Persistent $IGER_1$ remains near zero through month 2, begins rising around month 3, when the lagged non-overlapping window first includes any post-launch exposure in month 0 (both non-overlapping windows lie entirely in the post-treatment period only in month 5 after treatment in month 0), and continues growing through the end of the observation window with no sign of attenuation. The persistent component thus mirrors the returning-user spending trajectory from the new-versus-returning-user decomposition, as both build gradually and show sustained growth rather than a "launch-then-fade" pattern.

Subgroup estimates reinforce the earlier findings. Pre-launch OSG users exhibit the largest persistent $IGER_1$ increase (3.94 percentage points), and when expressed per full-panel member, the persistent component contributes 0.15, 0.12, and 0.08 percentage points for the low-, middle-, and high-income terciles, respectively (Web Appendix E13). A descriptive analysis of repeat rates tells a similar story: among individuals who first cross the $IGER_1$ threshold after their state's market launch, 39% cross again within three months, with returning users repeating at higher rates (47%) than newly active users (37%).

In summary, these three decompositions paint a consistent picture. Online market launch grows the market through acquisition at launch but shifts over time toward spending by previously observed users. Higher-income users account for the largest share of incremental spending volume. Lower-income users bear larger increases in spending-to-income exposure, and over 40% of that exposure is persistent and even growing across contiguous non-overlapping three-month windows. This reinforces the new-versus-returning-user finding that the post-launch impact represents a sus-

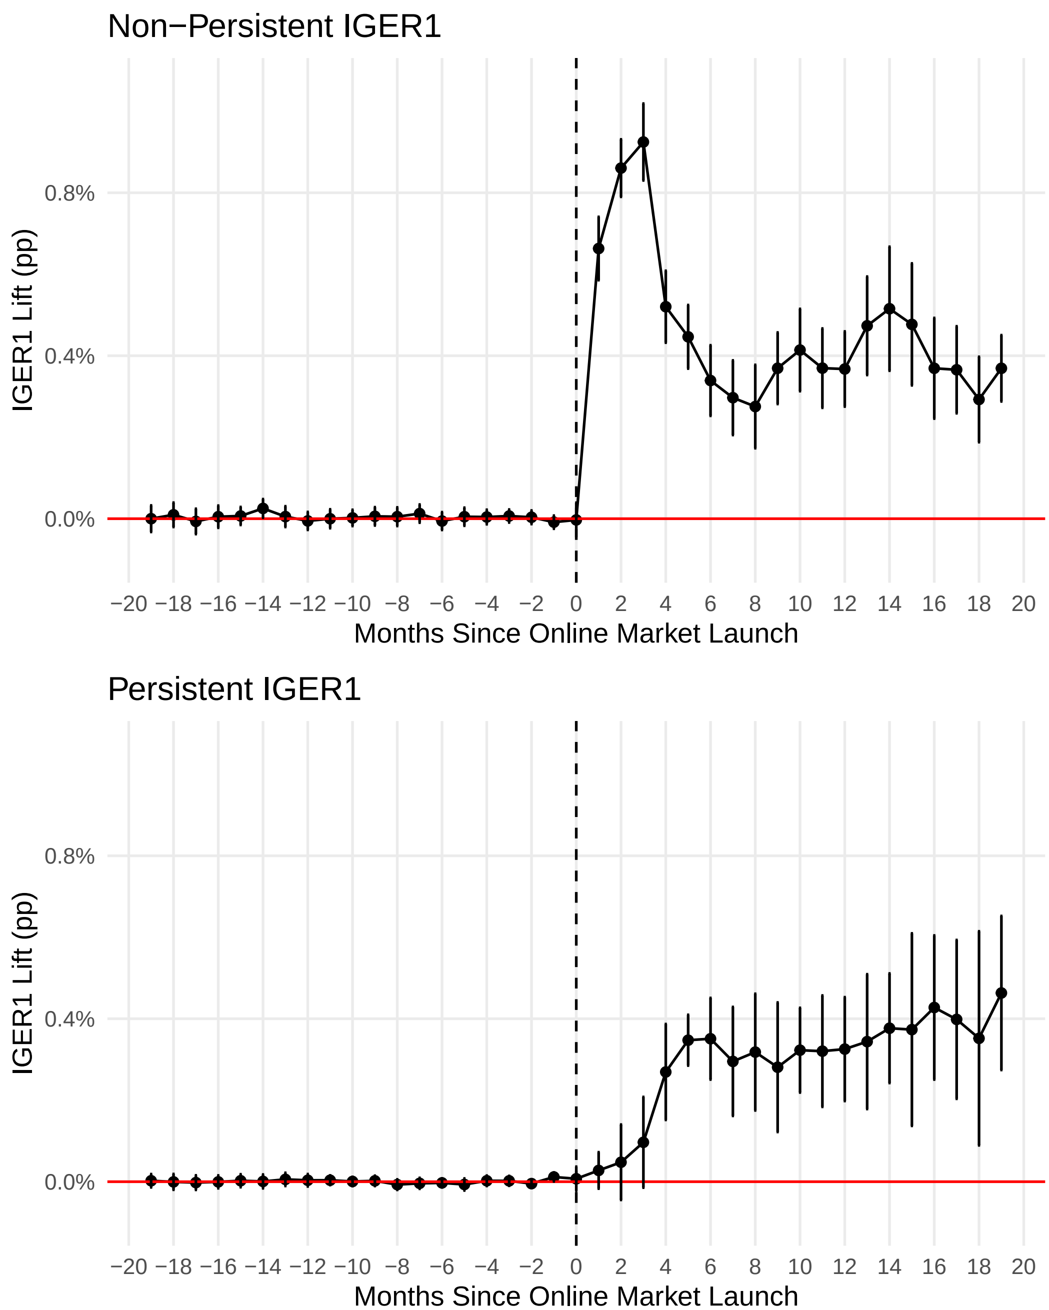

Figure 8: Dynamic $IGER_1$ Decomposition: Non-persistent vs. Persistent Threshold Crossing

Note: Top panel shows the dynamic treatment effect on the share of individual-months with a non-persistent $IGER_1$ crossing (above threshold now but not three months earlier). Bottom panel shows the corresponding effect for persistent crossings (above threshold in both the current and prior non-overlapping three-month window). Both panels share the same y-axis scale. Month 0 is the first calendar month in which legal online sports betting was available. The persistent component is mechanically near zero in post-treatment months 0–2 because the lagged window references only the pretreatment period in months 0–2 and begins to include post-launch exposure starting in month 3; the persistent effect is cleanly identified from month 5 onward, when both non-overlapping windows lie entirely in the post-treatment period. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023).

tained composition shift, not merely a transitory launch spike. The same persistence decomposition applied at the IGER_5 and IGER_{10} thresholds yields a strikingly similar persistent share of approximately 42–44%, confirming that this pattern is robust to the choice of threshold (Web Appendix E14).

7 Observed Casino Spending

To what degree does online sports betting divert spending from retail casinos? States launched online sports betting with differing levels of existing casino infrastructure, making channel substitution a natural concern. Consumers use spending cards to buy meals, accommodations, and tickets at retail casinos, so we estimate policy effects on consumer card spending at offline casinos. We treat this analysis as supplementary context for the paper's main contribution on market expansion and segment-level financial exposure.

We estimate online sports betting legalization treatment effects on casino-related spending across three distinct market-launch patterns: states that launched online sports betting only (Tennessee), states that launched online and offline sports betting nearly concurrently (Illinois, Colorado, Virginia, Wyoming, Arizona, Louisiana, Kansas, and Ohio), and states that launched online sports betting after previously legalizing offline sports betting (Arkansas and New York). We also distinguish between in-state casino spending (card spending at casino merchants located in the individual's modal state) and out-of-state casino spending (card spending at casino merchants in any other state) to capture potential geographic redistribution effects and to match the policy concern most directly relevant to states with local casino industries (see Web Appendix A1). These splits let us assess whether online sports betting market launch reduces in-state or out-of-state casino spending, and whether any such pattern differs when retail betting already existed.

The results are shown in Table 2. Across all three online market launch patterns, online sports gambling spending increases substantially and significantly, though the magnitude differs meaningfully: online-with-offline states show the largest increase at \$4.21 per individual-month, fol-

Table 2: Changes in OSG and Observed Casino Card Spending by Market Launch Pattern
($ Per Individual-Month)

Market launch pattern	Δ In-State Offline		Δ Out-of-State Offline		Δ OSG	
	ATT	Std. Error	ATT	Std. Error	ATT	Std. Error
Online-only	NM		-0.014	(0.021)	3.16	(0.028)
Online-with-offline	0.199	(0.030)	0.034	(0.053)	4.21	(0.058)
Online-after-offline	-0.007	(0.004)	0.042	(0.029)	1.61	(0.098)

Note. 'Online-only' represents Tennessee, which only legalized online sports gambling, and did not legalize offline sports gambling. 'Online-with-offline' represents states that launched both online and offline sports gambling within three months of one another, consisting of Illinois, Colorado, Virginia, Wyoming, Arizona, Louisiana, Kansas, and Ohio. 'Online-after-offline' represents states that launched online sports gambling more than three months after legalizing offline sports gambling, consisting of Arkansas and New York. Control states are Alabama, Alaska, California, Georgia, Idaho, Minnesota, Missouri, Oklahoma, South Carolina, Texas, Utah, and Vermont, which legalized neither form of sports gambling. 'Δ In-State Offline' represents the change in card spending at offline casinos located in the individual's modal state per individual-month, while 'Δ Out-of-State Offline' represents the change in card spending at offline casinos outside the individual's modal state per individual-month, relative to the estimated counterfactual with no intervention. 'Δ OSG' represents the change in online sports gambling spending per individual-month, presented for comparison purposes. Estimates are average treatment effects on monthly spending per individual-month with standard errors in parentheses. 'NM' stands for Not Measured, as the online-only state has no in-state casinos by definition.

lowed by online-only at \$3.16, and online-after-offline states at \$1.61. Because online-only and online-after-offline settings contain one and two states, respectively, these pathway-specific estimates are best read as informative decompositions rather than equally powered standalone designs.

The casino results provide no evidence of reduced observed casino card spending. In the online-only category (Tennessee), out-of-state casino card spending falls by \$0.014 per individual-month, a decline that is economically small relative to the \$3.16 increase in observed OSG spending and is not statistically significant ($p = 0.503$). In online-with-offline states, in-state casino card spending rises by \$0.199 per individual-month ($p < 0.001$), while out-of-state casino spending is not significantly changed. In online-after-offline states, casino spending estimates are economically small and statistically nonsignificant in both in-state and out-of-state channels despite the increase in online sports gambling activity. Overall, while we cannot directly measure cash gambling at offline casinos, we do not find evidence that online market launch reduces observed card spending at major casino operators.

8 Contextual Outcomes and Robustness Checks

We examine a small set of contextual outcomes as supplementary analyses that situate the main findings. We estimate treatment effects on alcohol spending, helpline contacts, and suicides using the same GSC methodology reported previously. We measure these outcomes using card panel data for spending effects, and state-level administrative data for help-seeking total suicides. We also report mean monthly per-capita state tax revenue paid by licensed online sports betting operators in treated states, which does not require quasi-experimental identification as the untreated tax revenue is always zero. Table 3 summarizes the results.

Table 3: Supporting Spending, Help-Seeking, and Fiscal Indicators Following Online Market Launch

Category	ATT	Std. Error	p-value	% Change
Related Consumption (Per Individual-Month)				
Mass-Market Alcohol	+\$0.654	(0.130)	$<$0.001	+19%
Premium Alcohol	+\$0.015	(0.027)	0.567	+1%
Public Health Indicators (Per Million Capita-Month)				
NCPG Helpline Calls	+7.00	(1.09)	$<$0.001	+79%
Suicides	-0.472	(0.746)	0.527	-3%
State Financial Impact (Per Capita-Month)				
Tax Revenue	+\$0.78	(0.081)	$<$0.001	NM

Notes: 'Mass-Market Alcohol' includes mass-market retailers (BevMo, Total Wine, Spec's Liquor), alcohol delivery services (Drizly, Minibar, Saucey), online wine retailers (Wine.com, Winc), and state-controlled stores (Virginia ABC). 'Premium Alcohol' includes premium wine subscription services (Bright Cellars, Wine of the Month Club), specialty wine merchants (KL Wine Merchants, Scout Cellar, Naked Wines, Underground Cellars), and tasting rooms. 'NCPG Helpline Calls' include all calls to National Council on Problem Gambling helplines. 'Suicides' represents all suicides at the state level, as reported by the Centers for Disease Control through 2022 (states with treatment dates in 2023 and beyond are not included in this estimate). 'Tax Revenue' represents monthly tax revenue per capita generated from sports betting activities. Standard errors in parentheses. Percent changes are with respect to counterfactual estimates, with the exception of Tax Revenue, which is reported as a per-capita dollar amount rather than a percent change.

8.1 Alcohol Spending

We find that online sports betting legalization increases card spending at mass-market alcohol merchants (+\$0.65 per individual-month, +19%), but find no significant change at premium-oriented merchants (\$0.02, $p = 0.567$). The merchant-category definitions are detailed in Table 3. This pattern is directionally consistent with complementary spending around sports viewing and accessible retail channels.

8.2 Public Health Impacts

Help-seeking measures provide an additional signal of the broader implications of online market launch. We examine texts and calls to National Council on Problem Gambling (NCPG) helplines as an important indication of consumer help-seeking behavior. As shown in Table 3, online market launch is associated with an increase of 7.00 calls per million residents per month. We interpret this as meaningful evidence that may reflect both increased distress and increased awareness of helpline resources, especially because operators often display helpline information prominently in advertising and on betting platforms.

We also examine population-level suicide rates, given prior research documenting associations between gambling behaviors and suicide. The results in Table 3 do not reveal a statistically significant change in overall suicide rates following online market launch. Because suicide has many determinants and administrative mortality records do not identify which individual deaths were linked to gambling, any gambling-specific effect would need to be substantial to register in this aggregate measure.[9] As such, this analysis lacks the precision to rule out a meaningful gambling-specific effect and should not be read as evidence that no such risk exists.

8.3 Fiscal Context

Online market launch also generates new tax revenue streams for states, providing a measurable fiscal input for policy evaluation. As shown in Table 3, launch generates \$0.78 of incremental tax revenue to the state per capita per month. Because statutory tax rates and market structures differ sharply across states, we treat this pooled per-capita tax estimate as contextual rather than as a state-specific projection.

8.4 Robustness Checks

We consider a variety of robustness checks to assess sensitivity to key assumptions, parameters, and research design choices:

- **IGER Threshold**. We re-estimate our main models using thresholds of 5% and 10% of rolling three-month income. As detailed in Web Appendix E4, absolute incidence is lower at higher thresholds, but the income gradient becomes more pronounced. This confirms that our substantive conclusions are not an artifact of threshold selection.

- **Simultaneous Launch**. Our main analysis includes all 11 treatment states in Table 1. As a robustness check, we restrict the sample to the five states that launched both retail and online betting in the exact same month (Arizona, Kansas, Ohio, Virginia, and Wyoming). The results in Web Appendix E5 remain substantively unchanged. OSG spending increases by $4.10 per individual-month overall, and IGER_1 increases by 0.66 percentage points, with low-income groups again showing the largest increases.

- **Alternative Estimators**. Our main analysis uses Generalized Synthetic Control (GSC), which is attractive both because it suits settings with limited treated units and because it never uses any treated unit as a control for another treated unit. To assess dependence on estimator family, we re-estimate our main effects using two widely used staggered difference-in-differences estimators: Sun and Abraham (2021, "SA") and Callaway and Sant'Anna (2021, "CS"). Because the SA and CS estimators do not incorporate latent factors, we use zero-factor GSC (i.e., two-way fixed effects) as the common baseline, isolating the question of estimator dependence from the separate question of latent-factor specification. As shown in Web Appendix E6, Table 9 presents estimates for all 13 outcomes across all four specifications, including our main GSC specification. The results are consistent in sign and comparable in magnitude across estimator families, suggesting that our findings are not strongly sensitive to the choice of estimator.

- **COVID-19 Pandemic**. The COVID-19 pandemic is the most important potential confounding factor during our study period, and our main GSC specification includes unemployment rates, Google mobility metrics across six categories, and COVID-19 case and death rates per capita to account for pandemic-related disruptions. We further assess this threat in two complementary ways. First, we exclude states that launched around the onset of the pandemic (Colorado, Illinois, Tennessee, and Virginia), removing approximately 36% of our treatment observations. As shown in Web Appendix E7, this produces similar results: OSG spending increases by \$3.48 per individual-month and IGER_1 increases by 0.59 percentage points. Second, we remove the initial COVID period (March–December 2020) from the sample and shift treatment dates for affected states to January 2021. As shown in Web Appendix E8, this approach also yields consistent findings, with OSG increases of \$3.68 and IGER_1 increases of 0.60 percentage points.

- **Placebo Treatment Timing**. A key testable implication of the identifying assumptions is that treatment effects should appear only after the true online market launch date. We test this by moving each state's treatment date earlier by 25% of its pretreatment period length. As shown in Web Appendix E9, the placebo estimates behave as expected: confidence intervals contain zero in 93% of cases, close to the expected 95% benchmark.

- **Placebo Categories**. To assess whether our design spuriously detects effects in unrelated domains, we examine treatment effects on large spending categories that should have no systematic relationship to online market launch. These placebo categories include air travel, big-box retailers, cable/internet providers, health insurance premiums, home and auto insurance, lodging/accommodation, and online travel agencies. As shown in Web Appendix Table W 10, all estimated treatment effects are statistically insignificant and economically small relative to pretreatment spending levels. For example, the largest estimated effect is a \$2.57 increase in big-box retailer spending ($p = 0.204$) on a base of \$342.71, representing a change of less than 1%.

- **Net Spending**. Our primary analysis uses gross OSG spending, reflecting the total amount of money individuals commit to gambling. Web Appendix E11 redefines IGER using net OSG spending (spending minus withdrawals). The results are directionally similar: the overall $IGER_1$ ATT is 0.57 percentage points rather than 0.78 percentage points, and the income gradient persists.

- **Population Weighting**. Our main analysis averages treatment effects across treated states, treating each state's policy experiment as equally informative. To assess sensitivity to this choice, we re-estimate the main results using weights proportional to each state's panel representation. As shown in Web Appendix E12, the population-weighted results closely mirror the unweighted estimates.

- **Merchant Set**. Our main analysis includes all 20 tracked OSG operators. To verify that results are not driven by ambiguous non-sportsbook operators (horse racing, poker, online casino, DFS-heavy platforms), we re-estimate after excluding four such merchants (TVG Network, PartyPoker, Golden Nugget Online Gaming, and Underdog Sports), which together account for approximately 6.2% of deposits. As shown in Web Appendix E15, the narrow-merchant estimates are virtually identical to the broad estimates.

- **Income Measure Reliability**. Our income measure is a Consumer Edge gross-income estimate derived from direct-deposit activity, which varies over time as Consumer Edge updates its estimate. To assess whether the regressive $IGER_1$ pattern depends on income measurement quality, we restrict to members whose monthly income coefficient of variation does not exceed 0.50 and exclude two control states with fewer than 200 post-filter members; all 11 treatment states are retained. As shown in Web Appendix E16, the ATTs attenuate moderately (OSG spending: \$3.33; $IGER_1$: 0.67 pp), but the regressive income gradient—low $>$ middle $>$ high—persists.

Across a wide range of alternative specifications, our substantive conclusions remain unchanged. In the baseline specification, online market launch increases monthly observed OSG spending by

$3.64 per individual and increases the share of individual-months above the 1% IGER threshold by 0.78 percentage points. These findings remain similar when we use alternative IGER thresholds, restrict the treatment sample to same-month online and retail launches, change estimator family, exclude COVID-period states, remove the initial COVID months from the sample, redefine IGER using net rather than gross spending, compute panel-weighted rather than unweighted averages, restrict to a narrower sportsbook-only merchant set, and restrict to members with stable income over time. Placebo analyses also support the design: placebo treatment dates produce confidence intervals containing zero 93% of the time, and placebo spending categories yield no statistically significant effects. Taken together, these checks indicate that online market launch robustly increases observed regulated OSG activity and observed spending-to-income exposure, with disproportionate increases among lower-income individuals in the tracked channel.

9 Discussion and Conclusions

Regulated online market launch causally increases observed gambling spending in tracked digital channels and the incidence of elevated spending relative to observed income in those channels. Using a large financial panel and Generalized Synthetic Control methods, we estimate that online market launch increases monthly observed OSG spending per individual by $3.64 (from $0.97 to $4.61) and increases the share of individual-months above the 1% IGER threshold by 0.78 percentage points (from about 0.17% to about 0.95%). These effects appear quickly at launch and persist in the post-treatment period.

These findings should be interpreted with scope conditions in mind. The outcomes measure regulated, digitally observed activity and do not capture illegal, informal, and cash-based retail gambling. Treatment is defined at online market launch, so the estimates reflect the rollout of regulated online sports betting, including operator entry, promotions, and marketing, and not statutory bill passage alone. The balanced panel improves visibility into spending over time but represents users of active credit and debit cards, not the full population.

Within this scope, the central substantive finding is distributional. Dollar spending increases rise with income but less than proportionally, so higher-income individuals account for the largest share of the spending lift while lower-income individuals experience larger increases in observed spending-to-income exposure, a divergence that becomes more pronounced at higher IGER thresholds (at 10% of income, the low-income ATT exceeds the high-income ATT by a factor of nearly five). A new-versus-returning-user decomposition further shows that the market grows through acquisition at launch but shifts over time toward spending by previously observed users, with first-observed users accounting for roughly one-fifth of the participation increase but less than 12% of spending volume, and their contributions decay within months, while previously observed users' spending lift persists and, if anything, continues to grow over time. The roughly stable aggregate spending effect therefore masks a composition shift: a larger share of post-launch spending volume comes from previously observed users. A persistence decomposition reinforces this conclusion, uncovering that 42% of the IGER_1 treatment effect comes from individuals who exceed the threshold in consecutive non-overlapping three-month windows, with this persistent component's treatment effect growing over time and being concentrated among pre-launch users and lower-income consumers. The non-persistent component, by contrast, spikes at launch and fades, mirroring the transitory new-user adoption pattern.

Supplementary contextual analyses on alcohol spending, helpline contacts, casino-card spending, and tax revenue are directionally consistent with the main findings but are not central to the paper's contribution.

The treatment effect estimates might help policymakers in non-legalizing states or countries to predict the effects of online sports betting legalization. The results may also help to quantify costs and benefits of various outcomes in the ongoing jurisdictional debate about whether prediction markets offer sports betting, and if so, whether it circumvents state regulations. Further, the decompositions suggest that persistent returning-user behaviors and sustained IGER_1 behaviors may motivate particular consumer protection instruments. The new-versus-returning-user decomposition shows that first-observed users generate a front-loaded participation spike that decays within roughly eight months, while returning-user spending activity persists and continues growing throughout the observation window. The persistence decomposition tells a parallel story: non-persistent IGER_1 crossings spike at launch and fade, while persistent crossings build gradually and show no sign of attenuating. Together, these patterns motivate sustained interventions targeting returning users — such as ongoing spending-to-income monitoring or self-exclusion prompts triggered by repeated threshold crossing — that would potentially reduce the component of elevated spending-to-income exposure that is both more durable and more concentrated among lower-income and pre-launch-active consumers.

Our data do not identify total-gambling effects across all channels; they identify changes in observed tracked-channel activity after online market launch. A natural question is how much the observed gambling spending reflects substitution from unobserved gambling channels, including cash gambling and black-market gambling. Like all similar analyses, we are unable to directly observe gambling spending in non-regulated channels. Several pieces of evidence weigh against a pure substitution account. We find no decline in observed card spending at major casino operators across any of the three market-launch patterns (Section 7), and in Tennessee, which launched online betting without concurrent retail authorization, the \$3.16 per individual-month OSG increase is accompanied by no measurable change in out-of-state casino card spending. Baker et al. (2026) find that legalization crowds out savings and brokerage deposits, not spending on other gambling products, suggesting that at least some of the observed increase reflects newly allocated dollars. We cannot rule out partial substitution from cash-based or black-market channels, but even if this were the case, the result would be that lower-income consumers now carry elevated spending-to-income ratios in an observable, regulated channel where responsible-gambling tools and consumer protection mechanisms operate and where regulatory intervention is feasible. That finding is policy-relevant whether the underlying behavior is new or newly visible.

The heterogeneity and new-versus-returning-user results jointly suggest that different policy margins may call for different interventions: launch-period protections address a transient surge of new adoption, while ongoing monitoring targets the persistent and growing spending activity among returning users, particularly among lower-income consumers whose spending-to-income exposure rises most. Future work could examine whether similar distributional patterns arise in other digitally intermediated forms of event-based speculation. As jurisdictions continue to legalize and refine regulatory frameworks, linking variation in regulatory design to the distributional and new-versus-returning-user outcomes documented here offers a promising path toward evidence-based consumer protection policy.

References

American Gaming Association (2026). Commercial gaming revenue tracker. Report, American Gaming Association. Accessed: 2026-03-03.

Anand, P. and Kadiyali, V. (2024). Frontiers: Smoke and Mirrors: Impact of E-cigarette Taxes on Underage Social Media Posting. *Marketing Science*.

Bai, J. (2009). Panel Data Models With Interactive Fixed Effects. *Econometrica*, 77(4):1229–1279.

Baker, S. R., Balthrop, J., Johnson, M. J., Kotter, J. D., and Pisciotta, K. (2026). Gambling away stability: Sports betting's impact on vulnerable households. Working Paper 33108, National Bureau of Economic Research.

Bollinger, B. and Sexton, S. E. (2023). Local excise taxes, sticky prices, and spillovers: evidence from Berkeley's soda tax. *Quantitative Marketing and Economics*, 21(2):281–331.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Canadian Centre on Substance Use and Addiction (2021). Developing lower-risk gambling guidelines. [Online; accessed 21. Apr. 2024].

Fisher, E. (2025). Americans bet $148 billion on sports in 2024, up 23.6%. *Front Office Sports*. [Online; accessed 4 Jun. 2025].

Goli, A., Mummalaneni, S., and Chintagunta, P. K. (2023). Making a Smooth Exit? Menthol Bans and Cigarette Sales in Massachusetts. *Marketing Science*.

Goss, A. and Mangrum, D. (2026). Sports betting, spending, and financial outcomes. Working paper.

Gramlich, J. (2025). Americans increasingly see legal sports betting as a bad thing for society and sports. Pew Research Center. Accessed 22 Feb. 2026.

Hansen, K., Misra, K., and Singh, V. (2026). When "selling hope" backfires: A lottery price hike reduces revenue. Working paper.

Hodgins, D. C., Young, M. M., Currie, S. R., Abbott, M., Billi, R., Brunelle, N., Campell, T., Dufour, M., Flores-Pajot, M.-C., Green, R., Griswold, M., Lund, I., Nadeau, L., and Pallesen, S. (2023). Lower-risk gambling limits: Linked analyses across eight countries. *International Gambling Studies*, 23(2):274–290.

Hollenbeck, B., Larsen, P., and Proserpio, D. (2026). The financial consequences of legalized sports gambling. *Available at SSRN: https://ssrn.com/abstract=4903302*.

Jonsson, J., Hodgins, D. C., Lyckberg, A., Currie, S., Young, M. M., Pallesen, S., and Carlbring, P. (2022a). In search of lower risk gambling levels using behavioral data from a gambling monopolist. *Journal of Behavioral Addictions*, 11(3):890–899.

Jonsson, J., Hodgins, D. C., Lyckberg, A., Currie, S., Young, M. M., Pallesen, S., and Carlbring, P. (2022b). In search of lower risk gambling levels using behavioral data from a gambling monopolist. *Journal of Behavioral Addictions*, 11(3):890–899.

Kaunitz, L., Zhong, S., and Kreiner, J. (2017). Beating the bookies with their own numbers – and how the online sports betting market is rigged. *arXiv preprint arXiv:1710.02824*.

Kim, K. and McCarthy, D. M. (2023). Wheels to Meals: Measuring the Impact of Micromobility on Restaurant Demand. *Journal of Marketing Research*, 61(1):128–142.

Langeland, E., Johnsen, I. F., Sømme, K. K., Morken, A. M., Erevik, E. K., Kolberg, E., Jonsson, J., Mentzoni, R. A., and Pallesen, S. (2022). One size does not fit all. Should gambling loss limits be based on income? *Frontiers in Psychiatry*, 13:1005172.

Larsen, P., Vijayaraghavan, S., and Ananthakrishnan, U. M. (2025). Cross-channel demand for addictive goods: Evidence from online sports gambling and state lotteries. Working paper.

Legal Sports Report (2025). Turning the page on a record 2024 for US sports betting. Accessed April 4, 2026.

Li, K. T. and Sonnier, G. P. (2023). Statistical Inference for the Factor Model Approach to Estimate Causal Effects in Quasi-Experimental Settings. *Journal of Marketing Research*, 60(3):449–472.

Lockwood, B. B., Allcott, H., Taubinsky, D., and Sial, A. (2024). What Drives Demand for State-Run Lotteries? Evidence and Welfare Implications. *Review of Economic Studies*, page rdae086.

Muggleton, N., Parpart, P., Newall, P., Leake, D., Gathergood, J., and Stewart, N. (2021a). The association between gambling and financial, social and health outcomes in big financial data. *Nature Human Behaviour*, 5(3):319–326.

Muggleton, N., Parpart, P., Newall, P., Leake, D., Gathergood, J., and Stewart, N. (2021b). The association between gambling and financial, social and health outcomes in big financial data. *Nat. Hum. Behav.*, 5:319–326.

Purdum, D. (2024). Sportsbooks defend practice of limiting how much sharp customers can bet. ESPN.

Reynolds, C. (2025). AGA reports record growth for sector. Accessed April 4, 2026.

Seiler, S., Tuchman, A., and Yao, S. (2020). The Impact of Soda Taxes: Pass-Through, Tax Avoidance, and Nutritional Effects. *Journal of Marketing Research*, 58(1):22–49.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Xu, Y. (2017). Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models. *Political Analysis*, 25(1):57–76.

Yeola, A., Allen, M. R., Desai, N., Poliak, A., Yang, K. H., Smith, D. M., and Ayers, J. W. (2025). Growing health concern regarding gambling addiction in the age of sportsbooks. *JAMA Internal Medicine*, 185(4):382–389.

Young, M. M., Hodgins, D. C., Currie, S. R., Brunelle, N., Dufour, M., Flores-Pajot, M.-C., and Nadeau, L. (2024). Not too much, not too often, and not too many: The results of the first large-scale, international project to develop lower-risk gambling guidelines. *International Journal of Mental Health and Addiction*, 22(1):666–684.

Notes

[1]Our merchant-level data do not reliably separate daily fantasy sports (DFS) from sportsbook activity within all tracked operators. We therefore use "online sports gambling" (OSG) to denote observed deposits to the tracked operator set. The main estimates should be interpreted as effects on observed tracked-channel gambling deposits, not necessarily on pure sportsbook activity alone. Web Appendix E15 shows that results are similar when we restrict to a narrower set of unambiguous sportsbook operators.

[2]`https://public.tableau.com/app/profile/national.council.on.problem.gambling.ncpg/viz/NationalProblemGamblingHelplineDashboard-IncomingTraffic/IncomingTraffic`, accessed March 2026.

[3]Our data contain realized, regulated transactions—merchant, amount, date, linked credit and debit card activity, income, and state—but do not contain the marketing environment surrounding launch, such as advertising exposure, welcome bonuses, promotion eligibility, or platform-specific onboarding frictions. Accordingly, our estimates identify the total effect of online market launch as implemented, rather than isolating the separate channels through which launch affects behavior.

[4]Market-launch timing is defined by the dates when gamblers in each state could start placing legal wagers on sports events, according to the American Gaming Association. Web Appendix E5 shows that the 'online-with-offline' treatment effect estimates are similar in the smaller group of five states that launched offline and online betting in the same month (Arizona, Kansas, Ohio, Virginia, and Wyoming).

[5]Offline sports betting takes place in physical establishments and is often cash-based, making it infeasible to observe comprehensively in our data.

[6]Note that some control states contained tribal lands with casino gambling activities. Tribes do not offer sports betting in any of the 12 control states. We do not know of any tribal gambling policy changes in control states during the observation period.

[7]This is a descriptive share of individuals who made at least one observed OSG spending transaction at any point during the post-launch period, not itself a causal estimate.

[8]We focus on gross spending because it indicates ex-ante loss exposure and gambling harm potential. Net spending adjusts gross spending for winning amounts withdrawn from gambling operators. The two measures usually align closely because gambling operators throttle or limit consistent winners' bet sizes (Kaunitz et al. 2017, Purdum 2024).

[9]State coroners cannot typically access decedents' financial records.

Online Market Launch and the Distribution of Gambling Risk: Evidence from Online Sports Betting

Web Appendix

Daniel M. McCarthy • Wayne J. Taylor • Kenneth C. Wilbur

Web Appendix A. Panel Data Details

In this section, we document how the panel was constructed. The full panel reports credit and debit card payment records from several million anonymous U.S. adults to about 5,000 large merchants. Each observation resembles a row on a monthly card statement, reporting an anonymous panelist ID, merchant name, transaction state, amount and date. The dataset also includes a vendor-estimated gross-income measure that Consumer Edge infers from observed direct-deposit patterns. We observe transactions at each included merchant for each individual.[10]

We construct a balanced panel by filtering the full dataset to individuals whose tracked cards remained active, defined as having at least one transaction categorized as a primary expense (e.g., groceries) in every year from 2019 through 2023. This filter ensures that each panelist's cards are consistently in use over the sample period, so that observed changes in spending reflect actual behavioral shifts and not the mechanical effect of an individual starting or discontinuing use of a card. The resulting analysis is conducted at the full-panel level and not only among gamblers, so the estimand pertains to stable, active cardholders in treatment and control states.

These data observe realized transactions – merchant, amount, date, total linked credit and debit card activity, income, and state – but do not observe advertising exposure, welcome bonuses, promotion eligibility, media salience, or platform-specific onboarding frictions. As a result, the paper identifies the total effect of online market launch as implemented, and not the separate roles of intermediate actions or a pure effect of statutory legalization.

Web Appendix A1. State Assignment and Outcome Construction

Our main outcome variables are ratios constructed at the state-month level. For example, OSG spending per individual divides total OSG deposits observed in a state-month by the number of individuals assigned to that state and year, and IGER_1 divides the count of individuals whose rolling three-month gambling-to-income ratio exceeds 1% by the same denominator. Constructing these ratios requires two geographic assignments: a transaction state for the numerator and an individual-to-state assignment for the denominator. Because individuals can transact in multiple states, these two assignments serve different purposes and are constructed differently.

Transaction state (numerator). Each transaction in the raw data carries a transaction state recorded by the data provider, representing the state associated with the cardholder's location at the time of the transaction. We use this field without modification. All outcome numerators—OSG deposits, IGER threshold counts, new gambler counts, category spending, and casino card spending—are aggregated by transaction state. Each state's outcome series therefore reflects activity that occurred within its borders.

Individual-to-state assignment (denominator). All per-individual outcome variables divide state-level aggregates by a count of individuals assigned to each state and year. We construct this assignment using total card spending across all approximately 5,000 tracked merchants—groceries, fuel, retail, restaurants, and all other categories—not gambling merchants alone. For each individual and year,[11] we assign the individual to the single state where they had the highest total dollar spending, which we refer to as their modal state. Because everyday local purchases (groceries, fuel, etc.) dominate total spending, this assignment proxies for the individual's state of residence.

This assignment has several properties relevant to the estimand. First, each individual is assigned to exactly one state per year, so there is no double-counting across states. Second, the assignment updates annually, so it captures residential moves in the year the individual's spending pattern shifts. Third, the assignment is unconditional on gambling: the denominator includes all approximately 1.2 million balanced-panel members, regardless of whether they ever gambled. The resulting estimand is therefore a sample-average treatment effect across all stable cardholders, not an effect conditional on gambling participation.

The numerator uses transaction state and the denominator uses the modal state assignment. This means that out-of-state activity—for example, a Maine resident gambling while visiting Maryland—enters Maryland's outcome series naturally. This allocation is sensible because Maryland's per-individual gambling spending reflects actual gambling activity within Maryland's borders, which is divided by the number of individuals whose primary spending activity places them in Maryland. States with more out-of-state visitor activity will thus naturally have higher spending per capita, analogous to how states with tourism-driven economies would have higher tourist spending per state resident.

Income terciles are also computed using the modal state assignment. For each state, we pool individuals across years to obtain stable income cutoffs, then classify each individual as low-, middle-, or high-income based on their annual gross income relative to their assigned state's tercile boundaries. Table 2 reports these cutoffs for each treatment state.

Casino in-state vs. out-of-state classification. The casino substitution analysis in Section 7 classifies each casino transaction as in-state or out-of-state by comparing the transaction state to the individual's modal state for that year. A casino transaction counts as "in-state" if it occurs in the same state as the individual's modal state, and "out-of-state" otherwise. This is consistent with the allocation decisions presented above, allowing us to assess whether online market launch shifts where individuals spend at casinos relative to their modal state.

Web Appendix A2. Financial Panel External Validity

Panel privacy protections limit our knowledge about panel membership. Some U.S. consumers do not have access to digital payments, so no digital-payments panel can be fully representative of the U.S. population. However, digital-payments access is a necessary participation condition for regulated online gambling. We proceed with presenting several indicators of external validity while treating the panel as strongest for within-panel causal inference.

Figure 1 compares distributions of annual incomes observed among panelists in our study and the Current Population Survey (CPS). Overall, as one would expect, study panelists earn more than the U.S. average: the CPS mean annual household income is $106,400, whereas the financial panel mean annual household income is $130,096, but panelists are observed in all income categories.

Table W 1: Online Gambling Merchants

Merchant Name	**% of Total Spend**	**% of Total Deposits**
Barstool Sportsbook	1.95	1.81
Bet365	0.19	0.39
Betfair	-8.03	0.28
BetMGM	19.35	12.34
BetRivers	1.33	1.1
Betway	0.06	0.05
Caesars Online Deposits	-0.3	0.16
Caesars Sportsbook	-0.08	0.5
DraftKings	42.15	43.3
FanDuel	29.36	32.03
FOX Bet	0.15	0.2
Golden Nugget Online Gaming	1.49	0.65
Hard Rock Digital	0.97	0.43
PartyPoker	-0.01	0
Pointsbet	0.15	0.64
theScore Bet	0.1	0.05
TVG Network	5.55	3.17
Underdog Sports	5.34	2.33
Unibet	0	0.1
WynnBET	0.28	0.49

Note. "% of Total Spend" represents net spend at each merchant and "% of Total Deposits" represents total deposits at each merchant. Both are expressed as a percentage of total net spend across all merchants. Total spend amounts can be negative if withdrawals exceed deposits.

Table 2 shows the income tercile levels associated with our treatment states.

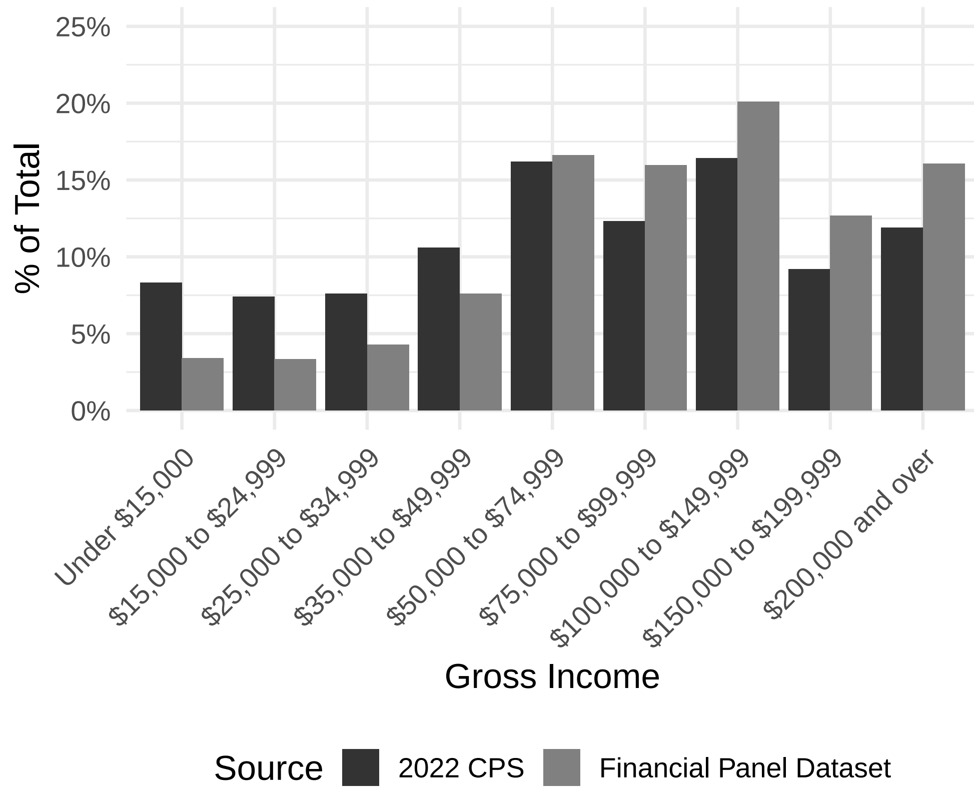

Figure W 1: **Annual Income Distribution Comparison**

Table W 2: Annual Gross Income Tercile Cutoffs

State	**Lower**	**Upper**
Arizona	$68,136	$123,192
Arkansas	$72,996	$125,700
Colorado	$73,188	$139,560
Illinois	$71,400	$138,708
Kansas	$72,528	$131,628
Louisiana	$69,264	$133,812
New York	$70,344	$135,084
Ohio	$68,208	$130,896
Tennessee	$68,136	$127,872
Virginia	$84,252	$167,400
Wyoming	$64,080	$116,064

Kim and McCarthy (2023) checked panel representativeness by correlating panel-based revenue measures from the same credit card panel dataset with public revenue disclosures for 36 restaurant brands over four years. They found a median correlation of 0.98, and an interquartile range of 0.92 to 0.99, suggesting that panel-based revenue metrics accurately tracked intertemporal changes in population-level revenue figures disclosed by public companies in their filings with the Securities and Exchange Commission.

We do not claim that the financial panel fully represents the population of all consumers. Instead, we view its strengths as passive measurement, stable longitudinal observation, and direct visibility into regulated digital transactions. Those features make the data well suited for within-channel causal estimation in the main paper, while extrapolation beyond users of active, tracked payment methods should be done cautiously.

Web Appendix B. Market Launch Dates

Table 3 lists online and offline sports-betting markets launch dates by state.

Table W 3: Online Market Launch Dates by State and Modality

State	Offline Launch	Online Launch
Arizona	9/9/21	9/9/21
Arkansas	7/1/19	3/6/22
Colorado	6/17/20	5/1/20
Connecticut	9/30/21	10/12/21
Delaware	6/5/18	12/27/23
District of Columbia	8/31/20	5/28/20
Florida	11/7/23	12/7/23
Illinois	3/9/20	6/18/20
Indiana	9/1/19	10/3/19
Iowa	8/15/19	8/15/19
Kansas	9/1/22	9/1/22
Kentucky	9/7/23	9/28/23
Louisiana	10/6/21	1/28/22
Maine	11/3/23	11/3/23
Maryland	12/9/21	11/23/22
Massachusetts	1/31/23	3/10/23
Michigan	3/11/20	1/22/21
Mississippi	8/1/18	8/1/18
Montana	3/11/20	3/11/20
Nebraska	6/22/23	—
Nevada	1/1/1949	10/1/13
New Hampshire	8/12/20	12/30/19
New Jersey	6/14/18	8/6/18
New Mexico	10/16/18	—
New York	7/16/19	1/8/22
North Carolina	3/18/21	3/11/24
North Dakota	12/1/21	—
Ohio	1/1/23	1/1/23
Oregon	8/27/19	10/15/19
Pennsylvania	11/15/18	5/31/19
Rhode Island	11/26/18	9/4/19
South Dakota	9/9/21	—
Tennessee	—	11/1/20
Vermont	—	1/11/24
Virginia	1/21/21	1/21/21
Washington	9/9/21	—
West Virginia	9/1/18	12/27/18
Wisconsin	11/30/21	—
Wyoming	9/1/21	9/1/21

Note. States without any retail or online legalization are not shown. Dates based on "Market Launch" dates as of 10/22/24 from the American Gaming Association: https://www.americangaming.org/research/state-gaming-map/

Web Appendix C. Treatment State Selection

Several states that legalized online sports betting during the observation window are excluded from the treatment sample. First, GSC estimation requires sufficient pre-treatment data for cross-validated factor selection, which in practice requires approximately 12 pre-treatment months to adequately control for seasonality. This requirement excludes states that launched before teh observation window (West Virginia, New Jersey, Mississippi, and Nevada) and states with online launches in the first year of the observation window (Indiana, New Hampshire, Oregon, Rhode Island, Iowa, and Pennsylvania). It also excludes Maryland because Maryland legalized retail sports betting 11 months before legalizing online sports betting, so the pre-treatment period is insufficient to establish the 12-month baseline within a single policy regime (with and without legal retail sports betting). Second, we exclude states whose online sports betting operates through lottery or government platforms rather than through any of the 20 commercial operators tracked in our data. This excludes Montana (Montana Lottery/Intralot), the District of Columbia (GambetDC/Intralot), and Oregon (Oregon Lottery/SBTech). Third, we exclude states where online casino gambling (iGaming) launched simultaneously with online sports betting, preventing isolation of the online sports betting effect. This excludes Michigan and Connecticut. Finally, we exclude Massachusetts, which has only 6.5 months of post-treatment data and exhibits poor pre-treatment matching in the GSC estimation. Applying these criteria results in the 11 treatment states that we analyze to obtain our main results. Table 3 lists all states' launch dates.

Web Appendix D. Treatment Assignment

An open question is whether states that launch online sports betting may differ systematically from states that do not. For example, states that experience budget shortfalls may be more likely to launch online sports betting and they may adopt more permissive market structures that encourage more gambling. Baker et al. (2026) found no evidence that employment growth, wage growth, or total private income predict the timing of state legislation, and Hollenbeck et al. (2026) compare treated and control states across various social assistance programs and COVID-19 fiscal responses and find no significant differences except in unemployment insurance, which they argue could understate financial health declines due to states' more generous policies. They also test launch timing against weekly wages, the quarterly unemployment rate, and the number of COVID cases and find no significant relationship.

We provide two additional checks: state financial health and political affiliation. First, we use the 2017-2018 Annual Survey of State Government Finances Summary compiled by the US Census and calculate three annual ratios to reflect a state's financial health: absolute deficit to revenue, absolute debt to revenue, and cash and securities to revenue. These three variables could predict whether states might have self-selected into treatment or treatment timing based on an urgent need to improve their fiscal situation. The results in Table 4 indicate that none of these variables have a significant relationship on either the decision to legalize or the time until treatment.

Incorporating political affiliation into our analysis allows us to explore whether states' party leanings correlate with their likelihood of legalizing online sports betting, potentially reflecting broader ideological differences, as political affiliation may serve as a proxy for regulatory tendencies. We measure political affiliation by state using the 2018 Gallup tracking data.[12] Table 5 shows that the percentage of residents that associate with the Republican or Democratic party does not significantly associate with whether the state is in our treatment or control group. In Table 6 we show the distribution of Gallup's party classification based on the difference between Republican and Democratic residents; both Treated and Control groups show a similar distribution across states, and a χ-squared test finds no association between state party classification and treated or control status (p-value = 0.805).

Table W 4: Legalization Treatment by State Financial Health Ratios

	Dependent variable:	
	Treated	Years to Treatment
	(1)	(2)
Deficit / Revenue	−19.365	6.445
	(13.099)	(13.682)
Debt / Revenue	5.742	−1.943
	(4.149)	(3.033)
Cash and Securities / Revenue	0.503	−0.061
	(0.568)	(0.353)
Constant	15.618	−2.077
	(10.959)	(12.076)
Observations	23	11
R^2		0.060
Log Likelihood	−14.309	

Note. Data from the 2017-2018 Annual Survey of State Government Finances Summary compiled by the US Census.

Table W 5: Legalization Treatment by Political Affiliation

	Dependent variable:	
	Treated	
	(1)	(2)
% Republican/Lean Republican	−0.018	
	(0.055)	
% Democratic/Lean Democratic		0.014
		(0.055)
Constant	0.715	−0.645
	(2.446)	(2.259)
Observations	23	23
Log Likelihood	−15.865	−15.889
Akaike Inf. Crit.	35.730	35.778

Note. Data from 2018 Gallup tracking data.

Table W 6: Distribution of Gallup political classification by Treated status

	Control	**Treated**
Competitive	2	2
Democratic	3	4
Republican	7	5

Note. Data from 2018 Gallup tracking data. A χ-squared test finds no association between state party classification and treated or control status (p-value = 0.805).

Web Appendix E. OSG and IGER Treatment Effects

This section presents the primary treatment effects of online market launch on monthly observed OSG deposits per individual and on the income-adjusted gambling expenditure ratio (IGER_1), both overall and disaggregated by income terciles.

For each individual i in month t, we define rolling three-month deposits $D_{it}^{(3)} = \sum_{s=t-2}^{t} D_{is}$ and rolling three-month income $Y_{it}^{(3)} = \sum_{s=t-2}^{t} Y_{is}$, and set $IGER_{it,1} = \mathbb{1}\left[D_{it}^{(3)}/Y_{it}^{(3)} > 0.01\right]$. The state-month outcome $IGER_{st,1}$ is the share of individuals in state s with $IGER_{it,1} = 1$. We use a rolling three-month window (months t, $t-1$, and $t-2$) for two reasons. First, on a single-month basis, 1% of monthly income is a low absolute threshold that could be exceeded by even modest deposit activity; aggregating over three months raises the effective bar so that threshold-crossing reflects more economically meaningful gambling expenditure relative to income. Second, the wider window smooths transitory month-to-month fluctuations in both income and spending. IGER should therefore be interpreted as an observed regulated-channel deposit-to-income indicator, not as total gambling loss, total household gambling burden, or realized harm.

Table W 7: OSG Spend

Income Segment	ATT	Std. Error	t-stat	p-value
Overall	3.64	0.07	55.99	$<.001$
Low-income tercile	2.64	0.08	33.79	$<.001$
Middle-income tercile	3.82	0.12	32.49	$<.001$
High-income tercile	4.02	0.15	26.43	$<.001$

Note. Positive values indicate higher monthly online-sports-gambling (OSG) spending after online market launch in treated states relative to control states. OSG spending here refers to gross deposits to regulated online gambling operators. Standard errors computed using the procedure from Li and Sonnier (2023)

Table W 8: IGER_1

Income Segment	ATT	Std. Error	t-stat	p-value
Overall	0.784%	0.009%	89.3	$<.001$
Low-income tercile	1.140%	0.017%	68.3	$<.001$
Middle-income tercile	0.802%	0.014%	58.4	$<.001$
High-income tercile	0.471%	0.012%	39.4	$<.001$

Note. Positive values indicate a higher monthly IGER_1 after online market launch in treated states relative to control states. IGER_1 should be interpreted as a proxy for elevated gamble-to-income risk within observed regulated channels.

Web Appendix E1. Dynamic $IGER_1$ By Income Group

Figure 2 shows the dynamic ATT estimates of $IGER_1$ by income group.

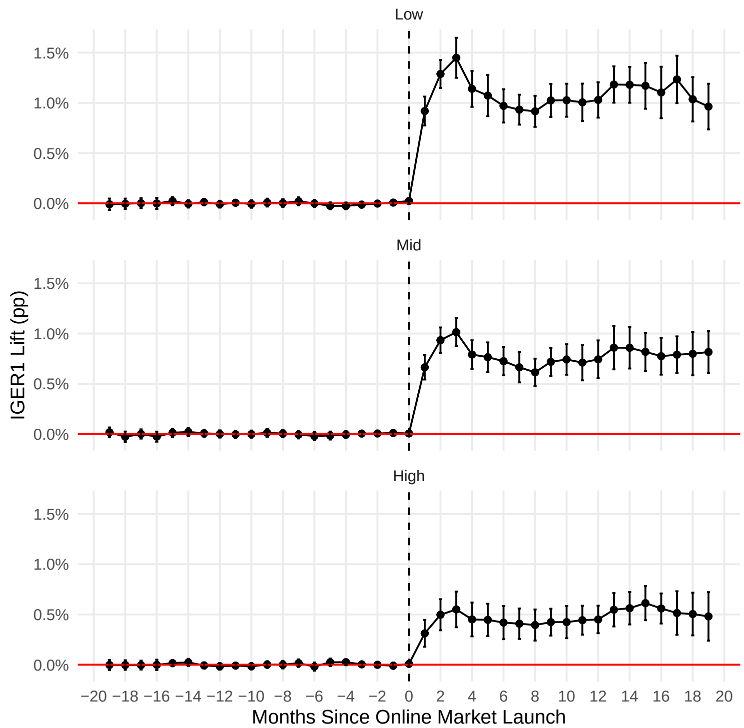

Figure W 2: **$IGER_1$ Dynamic ATT Estimates by Income Group**

Note. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Post-treatment observation spans 9 months for Ohio and 13 months for Kansas, all other states have at least 19 months posttreatment.

Web Appendix E2. Dynamic Participation Rate Decomposition

Figure 3 decomposes the dynamic treatment effect on the OSG participation rate into contributions from newly active users (top panel) and returning users (bottom panel). The pattern mirrors the deposit decomposition in Figure 6: new-user participation spikes at launch and attenuates within approximately eight months, while returning-user participation rises immediately and remains persistently elevated.

Web Appendix E3. Dynamic Deposit Effects by Income Tercile

Figure 4 shows the dynamic treatment effects on OSG deposits per individual separately for low-, middle-, and high-income terciles. All three panels share the same y-axis scale. The temporal pattern is broadly similar across groups: deposits rise at launch and remain elevated, with high-income individuals showing the largest level effects throughout the posttreatment period.

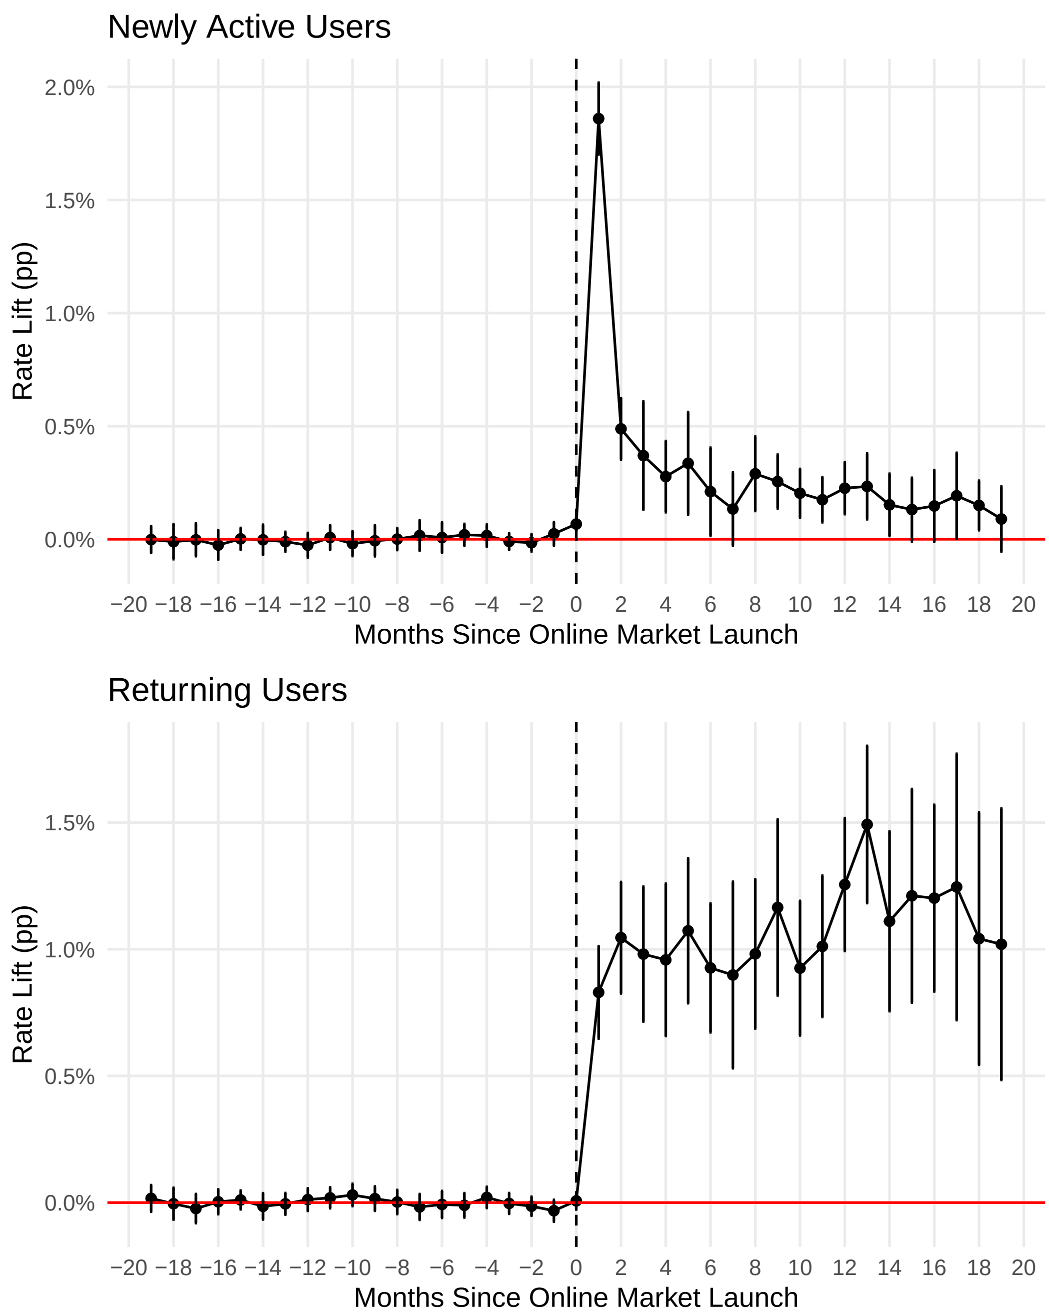

Figure W 3: Dynamic Participation Rate Decomposition: Newly Active vs. Returning Users

Note: Top panel shows the dynamic treatment effect on the share of individuals making their first observed OSG deposit in the current month. Bottom panel shows the effect on the share of individuals with a prior observed OSG deposit who deposit again. Both panels share the same y-axis scale. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Posttreatment observation spans 9 months for Ohio and 13 months for Kansas; all other states have at least 19 months posttreatment.

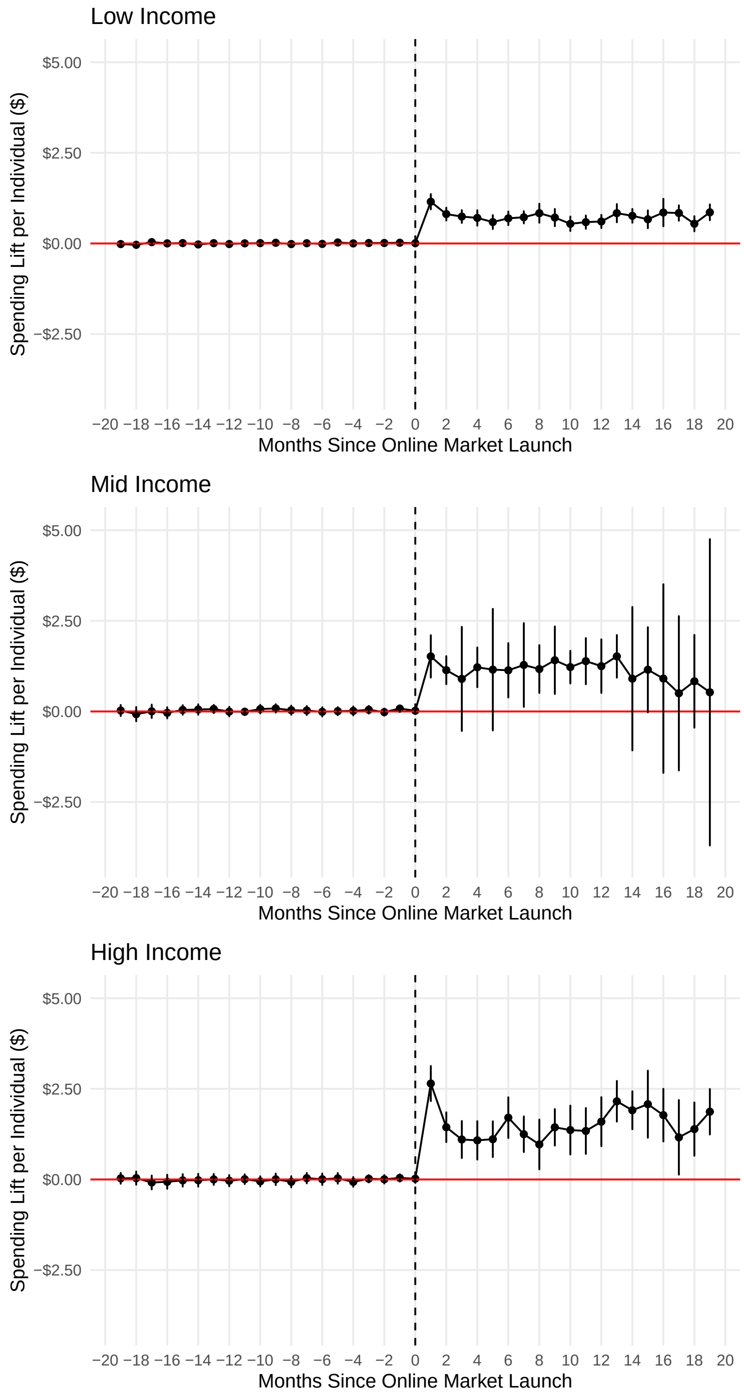

Figure W 4: Dynamic Spending Effects by Income Tercile

Note: Each panel shows one income tercile's contribution to the overall per-capita OSG deposit effect. Because the effects are expressed per full-panel member rather than per tercile member, the three panels sum to the aggregate per-capita effect. Terciles are defined within each state using time-invariant average gross income. All three panels share the same y-axis scale. Error bars are 95% confidence intervals computed using the procedure of Li and Sonnier (2023). Posttreatment observation spans 9 months for Ohio and 13 months for Kansas; all other states have at least 19 months posttreatment.

Web Appendix E4. Alternative IGER Thresholds

In this section, we provide supplementary results when using thresholds of 5% and 10% for the IGER measure, instead of 1%.

Figure 5 presents descriptive trends for IGER_5 and IGER_{10}, constructed identically to Figure 2(b) in the main paper, except using the higher 5% and 10% spending-to-income thresholds. In both panels, treatment and control states track closely during the pretreatment period and diverge sharply at online market launch, with treatment states exhibiting sustained elevated shares of individuals above the respective thresholds. The levels are lower than for IGER_1, as expected, since fewer individuals exceed these higher thresholds.

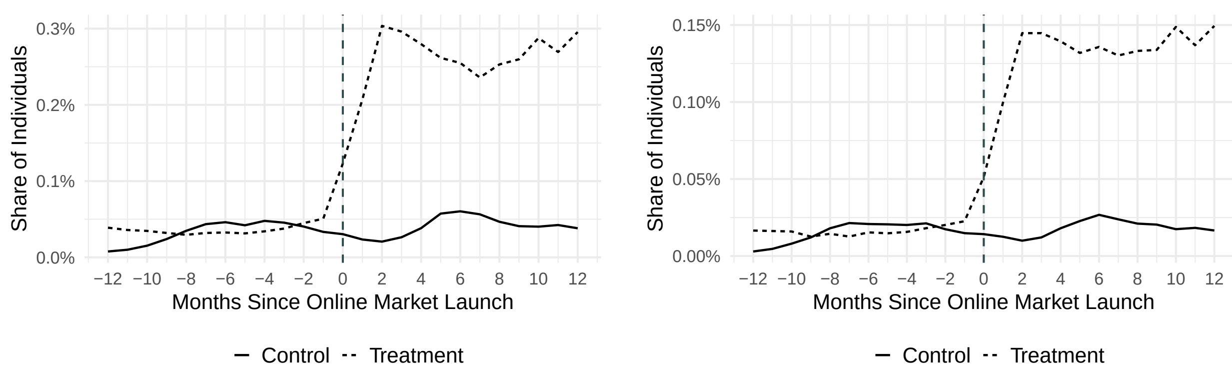

(a) IGER_5 (b) IGER_{10}

Figure W 5: Descriptive Trends for IGER_5 (left) and IGER_{10} (right)

Note. Each panel is constructed identically to Figure 2(b) in the main paper, using the 5% and 10% thresholds in place of 1%. The y-axis shows the share of individuals whose cumulative OSG deposits over months t, $t-1$, and $t-2$ exceeded the indicated percentage of cumulative income over the same three-month window. For each treatment state, period 0 is the time of online market launch. For control states, period 0 is the midpoint of the sample period.

Figure 6 compares the treatment effect of online market launch on IGER_5 and IGER_{10} to IGER_1, overall and by income group. Note that the horizontal axis scales change across the three IGER thresholds in the figure. The results assuming outcome measures of IGER_5 and IGER_{10} are substantively consistent – as with IGER_1, we observe a significantly higher ATT overall, and a lower ATT for the high-income group relative to the low-to-mid income groups.

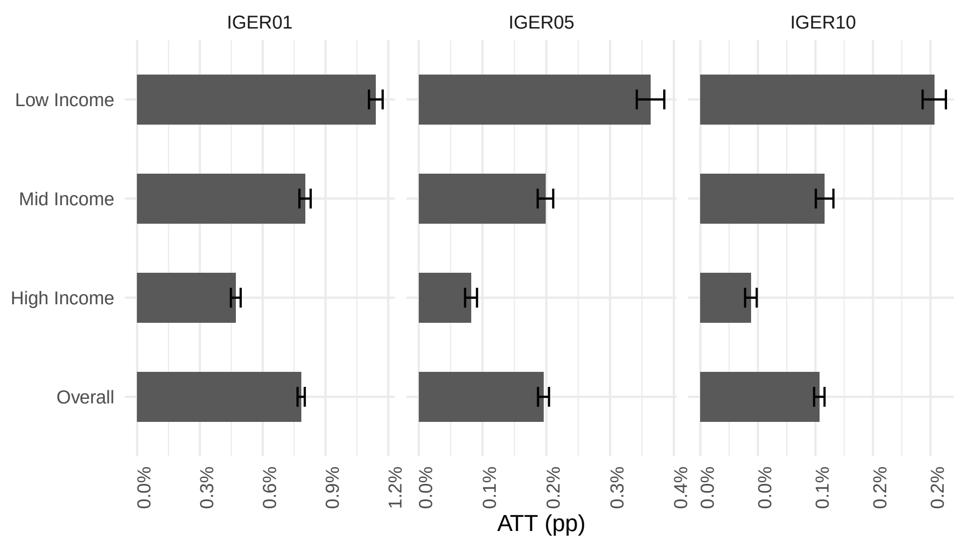

Figure W 6: **$\mathbf{IGER_1}$, $\mathbf{IGER_5}$, and $\mathbf{IGER_{10}}$ ATT Estimate Comparison**

Note. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Income group affiliation uses state-specific terciles on income distribution.

Web Appendix E5. Results Using Simultaneous Legalization of Online and Retail Sports Betting As Treatment States

In this section, we summarize our main results when we only include states that launched RSB and OSB simultaneously (i.e., in the same month): Arizona, Kansas, Ohio, Virginia, and Wyoming. This is in contrast to our main analysis, which includes all 11 treatment states in Table 1, encompassing all three market-launch patterns. The results are shown in Figures 7 and 8. We observe the same substantive conclusions – similar increases in observed OSG spending across income groups, and significantly increasing $IGER_1$ with most of the increase coming from low-to-middle income individuals.

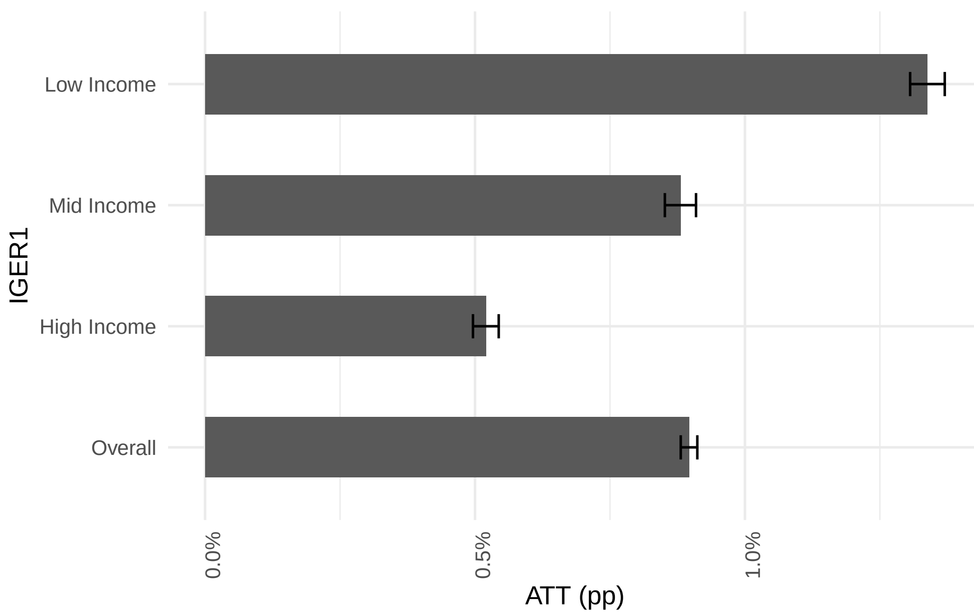

Figure W 7: **$IGER_1$ ATT estimates from states that launched RSB and OSB simultaneously**

Note. Results obtained using Arizona, Kansas, Ohio, Virginia, and Wyoming as treatment states, as these states launched RSB and OSB in the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Income group affiliation uses state-specific terciles on income distribution.

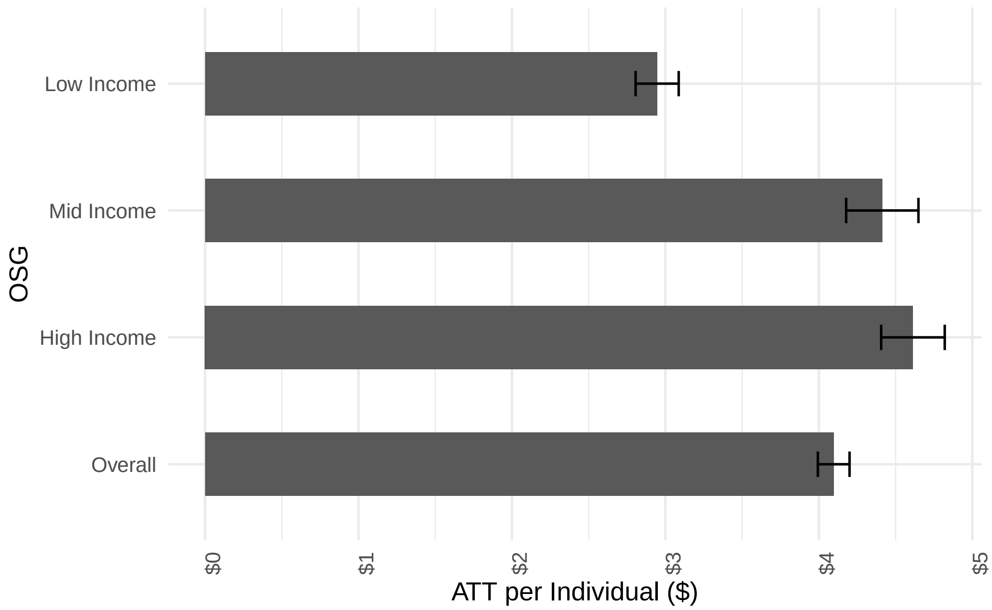

Figure W 8: **OSG ATT estimates from states that launched RSB and OSB simultaneously**

Note. Results obtained using Arizona, Kansas, Ohio, Virginia, and Wyoming as treatment states, as these states launched RSB and OSB in the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Income group affiliation uses state-specific terciles on income distribution.

Web Appendix E6. Alternative Estimators

In this section, we evaluate the robustness of our main results by changing the estimator. Table 9 compares the main cross-validated GSC specification to a zero-factor GSC baseline and two popular staggered difference-in-differences estimators by Sun and Abraham ("SA") Sun and Abraham (2021) and Callaway and Sant'Anna ("CS") Callaway and Sant'Anna (2021). Because the SA and CS estimators do not incorporate latent factors, we use GSC_0 (i.e., two-way fixed effects) as the common baseline for statistical comparisons. The "μ Same" column reports whether GSC_0/SA and GSC_0/CS pairs fail to reject the null hypothesis of equal treatment effects at the 95% confidence level.

Of the 13 outcome variables, 10 fail to reject equal treatment effects for both estimator pairs. The remaining three reject equality in one of the two pairs, but in each case the point estimates are close in magnitude and all estimates are significantly different from zero in the same direction. We therefore interpret these as differences in precision rather than qualitative disagreements.

Comparing across the first two columns, the main cross-validated GSC estimates are generally close to their zero-factor counterparts, though the latent-factor specification produces notably larger point estimates for a few outcomes, such as NCPG helpline calls. In all cases, the direction and statistical significance of the effects are preserved. Overall, we conclude that the findings are robust across estimator families and that the latent-factor specification refines magnitudes for certain outcomes without altering the substantive conclusions.

Table W 9: Comparison to Staggered Diff-in-Diff Estimators

Outcome Variable	**GSC (main)**	**GSC$_0$**	**SA**	**CS**	**μ Same**
IGER$_1$	0.0078 (0.0001)	0.0073 (0.0001)	0.0071 (0.0005)	0.0074 (0.0002)	Both
IGER$_5$	0.0020 (0.0000)	0.0021 (0.0000)	0.0020 (0.0002)	0.0020 (0.0000)	Both
IGER$_{10}$	0.0010 (0.0000)	0.0011 (0.0000)	0.0011 (0.0001)	0.0011 (0.0000)	Both
IGER$_{15}$	0.0007 (0.0000)	0.0007 (0.0000)	0.0007 (0.0001)	0.0007 (0.0000)	Both
IGER$_1$ - Low Inc.	0.0114 (0.0002)	0.0105 (0.0002)	0.0100 (0.0007)	0.0106 (0.0002)	Both
IGER$_1$ - Mid Inc.	0.0080 (0.0001)	0.0077 (0.0001)	0.0073 (0.0005)	0.0076 (0.0002)	Both
IGER$_1$ - High Inc.	0.0047 (0.0001)	0.0047 (0.0001)	0.0047 (0.0003)	0.0048 (0.0002)	Both
OSG Spend	3.6412 (0.0650)	3.8689 (0.0414)	3.6897 (0.3837)	3.6260 (0.0619)	SA
OSG Spend - Low Inc.	2.6363 (0.0780)	2.6817 (0.0361)	2.6949 (0.3273)	2.6738 (0.0544)	Both
OSG Spend - Mid Inc.	3.8180 (0.1175)	4.2508 (0.0851)	3.7498 (0.3003)	3.5036 (0.1866)	SA
OSG Spend - High Inc.	4.0212 (0.1521)	4.4783 (0.0725)	4.4539 (0.5256)	4.4923 (0.0806)	Both
New Gamblers	0.0032 (0.0001)	0.0029 (0.0001)	0.0023 (0.0002)	0.0025 (0.0002)	CS
NCPG Calls per 1M people	7.0000 (1.0861)	4.2973 (0.8302)	5.1723 (1.4447)	5.1742 (1.7074)	Both

Note. "GSC (main)" reports point estimates from the cross-validated latent-factor specification used in the main text; the number of latent factors is selected by cross-validation following Xu (2017). "GSC$_0$" reports GSC with zero latent factors (i.e., a two-way fixed-effects specification), which serves as the common baseline for the statistical comparisons in the "μ Same" column. "SA" and "CS" represent the corresponding results of the staggered diff-in-diff estimators by Sun and Abraham Sun and Abraham (2021) and Callaway and Sant'Anna Callaway and Sant'Anna (2021), respectively. "μ Same" denotes whether alternative estimators (either "SA", "CS", or "Both") fail to reject the null hypothesis of equal treatment effect sizes relative to GSC$_0$ at the 95% confidence level. Standard errors are shown in parentheses for all estimators. The formal μ Same tests are anchored to the GSC$_0$ baseline.

Web Appendix E7. Results Excluding States that launched Sports Betting Around the COVID-19 Pandemic

In this section, we summarize our main results when we exclude states that launched sports betting around the onset of the COVID-19 pandemic (March 2020 through January 2021): Colorado, Illinois, Tennessee, and Virginia. As above, we observe the same substantive conclusions in this scenario.

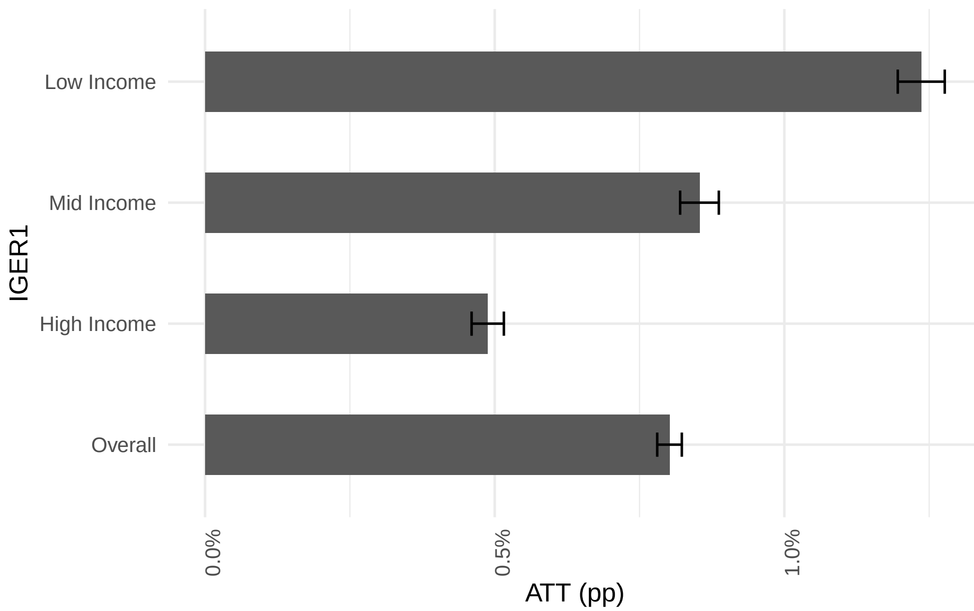

Figure W 9: **$\mathbf{IGER}_1$ ATT estimates without COVID-19 states**

Note. Estimates exclude states that launched OSB around the onset of the COVID-19 pandemic (March 2020 through January 2021): Colorado, Illinois, Tennessee, and Virginia. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

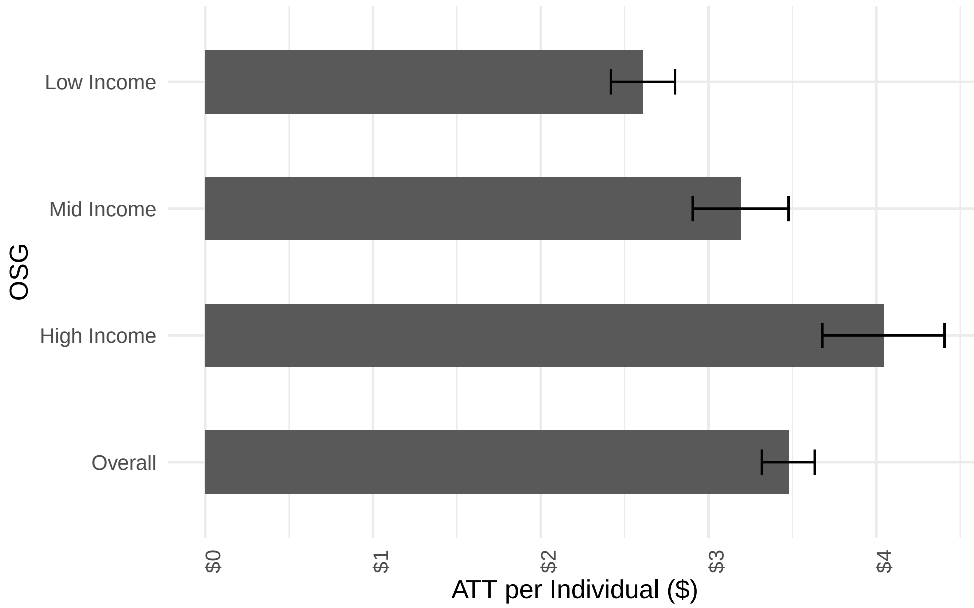

Figure W 10: **OSG ATT estimates without COVID-19 states**

Note. Estimates exclude states that launched OSB around the onset of the COVID-19 pandemic (March 2020 through January 2021): Colorado, Illinois, Tennessee, and Virginia. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Estimates exhibit similar levels and patterns as those presented in the main analysis.

Web Appendix E8. Excluding COVID Time Periods from Analysis

In this robustness check, we remove all data from March to December 2020 from the observation period. By doing so, we use pre-COVID data when setting up the pretreatment weights, and only use post-COVID data to evaluate the ATTs.

For the three treatment states with online launches during the COVID-19 period (March–December 2020)—Illinois (June 2020), Colorado (May 2020), and Tennessee (November 2020)—we shift their treatment dates forward to January 2021. The remaining eight treatment states, whose online launches occurred after December 2020, are unaffected by this adjustment. The results remain substantively consistent with our main model specification.

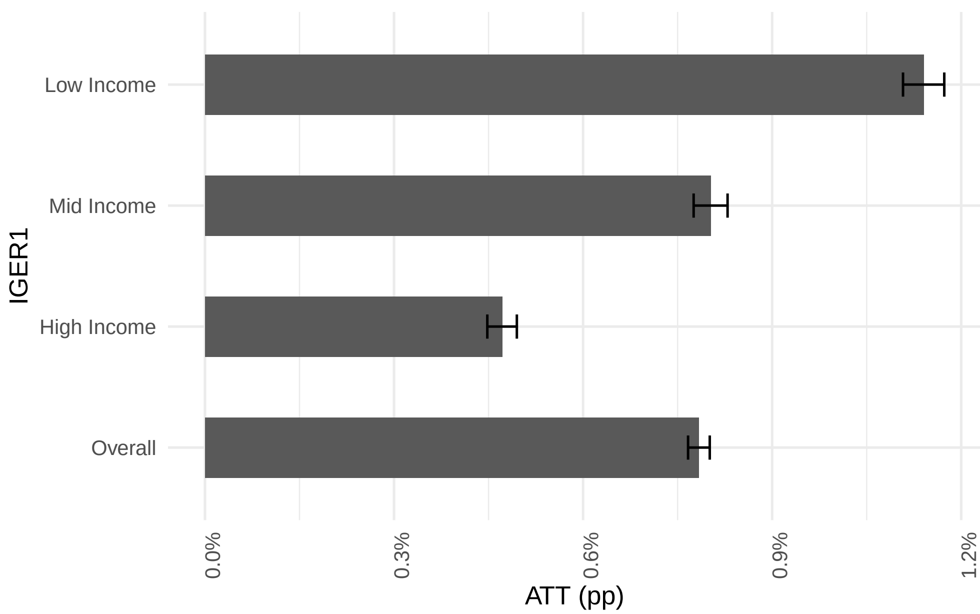

Figure W 11: **IGER$_1$ ATT estimates excluding COVID-19 time periods**.

Note. For states that launched OSB during this time period (Colorado, Illinois, and Tennessee) we move the treatment date forward to January 2021. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Estimates exhibit similar levels and patterns as those presented in the main analysis.

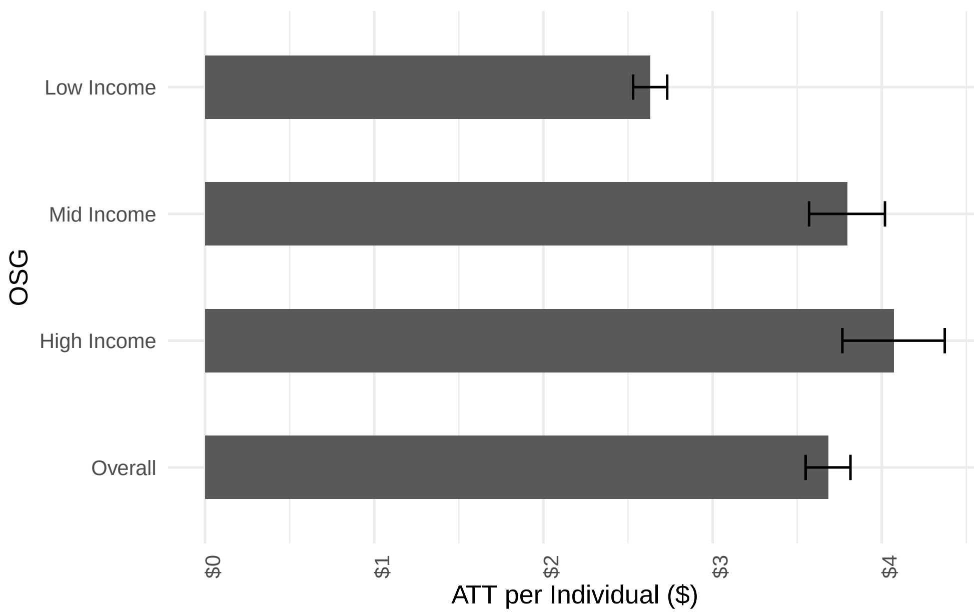

Figure W 12: **OSG ATT estimates excluding COVID-19 time periods**

Note. We exclude March 2020 through December 2020. For states that launched OSB during this time period (Colorado, Illinois, and Tennessee) we move the treatment date forward to January 2021. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023). Estimates exhibit similar levels and patterns as those presented in the main analysis.

Web Appendix E9. Placebo Treatment Timing

In this section, we evaluate the robustness of our main results by changing the observed times of treatment, a practice known as "placebo treatments." For each treated state, we created an artificial time of treatment by reducing that state's pretreatment period by 25% of the time between sample start and actual time of treatment. Intuitively, a 95% confidence interval for the estimated treatment effect should be indistinguishable from zero 95% of the time since no actual treatment occurred. Empirically, we find this to be true. Confidence intervals for the overall treatment effect estimates contained zero 93% of the time.

Web Appendix E10. Placebo Categories

To assess the internal validity of our analysis, we estimate treatment effects on large spending categories where we would not a priori expect a systematic relationship to online market launch. These placebo categories were specifically selected to capture a substantial share of the consumer wallet, including essential services (health insurance, home insurance), major discretionary spending (air travel, lodging), and high-frequency purchases (big-box retailers). If our identification strategy were flawed, we would likely detect spurious effects in at least some of these high-volume categories.

Table W 10: Placebo Checks on Large, Unrelated Spending Categories (Alphabetical Order)

Category	Pre–Treat Mean	ATT	SE	p–value
Air Travel	34.47	0.26	0.43	0.553
Big–Box Retailers	342.71	2.57	1.82	0.204
Cable / Internet Providers	56.05	0.85	0.58	0.101
Health Insurance Premiums	1.92	0.90	1.88	0.600
Home & Auto Insurance	42.93	0.12	0.75	0.900
Lodging / Accommodation	49.82	0.47	0.71	0.501
Online Travel Agencies	5.88	–0.08	0.20	0.701

All estimated treatment effects are statistically insignificant and economically small relative to pre-treatment spending levels. The absence of any statistically significant effects across these diverse, theoretically unrelated spending areas reinforces the plausibility of our identification strategy and suggests our gambling-specific results reflect genuine causal effects and not artifacts of the dataset.

Web Appendix E11. IGER Using Net Spend

As a robustness check, we re-compute IGER_1 using *net* OSG spend (deposits minus withdrawals) in place of gross deposits when constructing the rolling three-month gamble-to-income measure. The results are directionally similar to our baseline specification: the overall IGER_1 ATT is 0.57 percentage points, compared to 0.78 percentage points when using gross deposits. The income gradient also persists, with the largest increases among lower-income individuals. These findings indicate that the qualitative conclusions do not depend on whether the numerator is gross deposits or net spend, even though the level is smaller when net spend is used (0.57 versus 0.78 percentage points overall).

Web Appendix E12. Panelist-Weighted Analyses

In this robustness check, we use the weighted average of state-level ATTs, in which the weight of each state is proportionate to the number of panelists in that state. In the main paper, the overall average effect is equal to an unweighted average of the treatment effects associated with each state. The results remain substantively consistent with our main model specification.

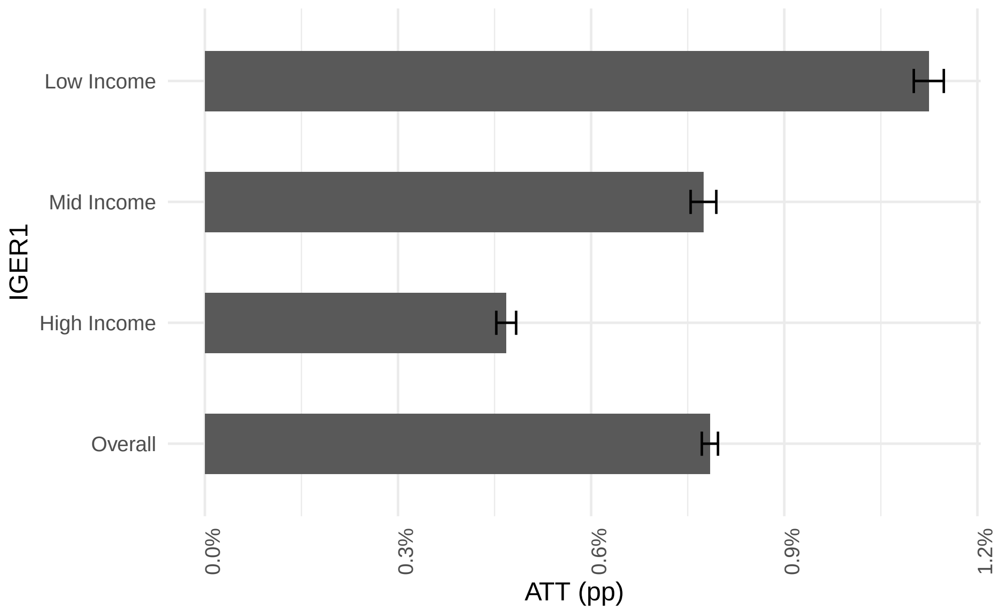

Figure W 13: **IGER_1 ATT Weighted by Panelist**

Note. We use a weighted average of state-level ATTs based on panelist count. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

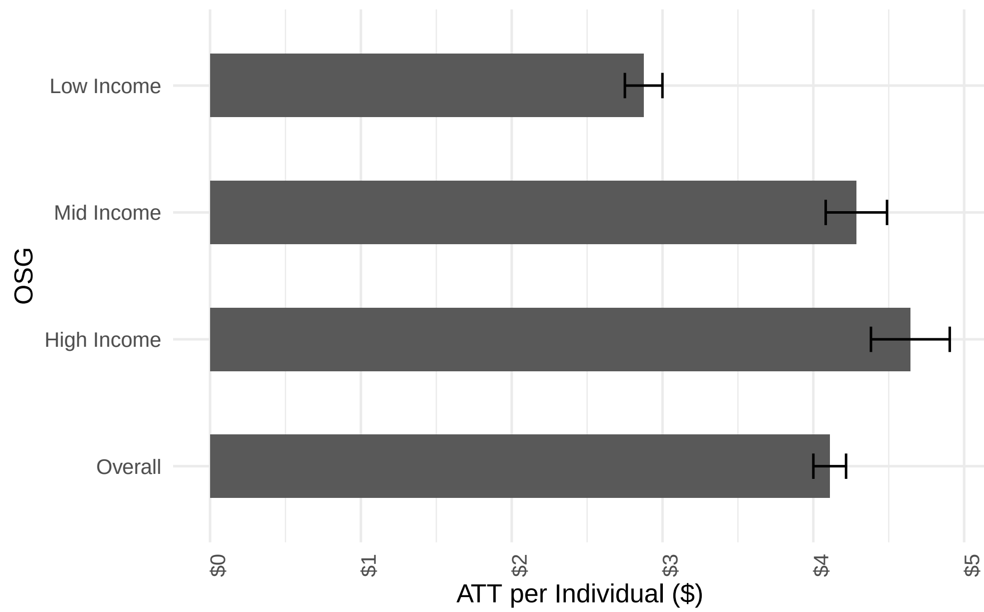

Figure W 14: **OSG ATT Weighted by Panelist**

Note. We use a weighted average of state-level ATTs based on panelist count. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

Web Appendix E13. Persistent $IGER_1$ Subgroup Analysis

Figure 15 reports the persistent $IGER_1$ treatment effect by income subgroup. When expressed per full-panel member, the persistent component contributes 0.15 pp for the low-income tercile, 0.12 pp for the middle-income tercile, and 0.08 pp for the high-income tercile. All estimates are statistically significant at the 5% level. These subgroup patterns mirror the overall $IGER_1$ income gradient reported in the main text, confirming that the regressive pattern holds specifically for the persistent component of elevated deposit-to-income exposure. In a separate estimation restricted to pre-launch OSG users—individuals with observed OSG deposits before their state's online market launch—the persistent $IGER_1$ increase is 3.94 pp (from a base of approximately 1.5%).

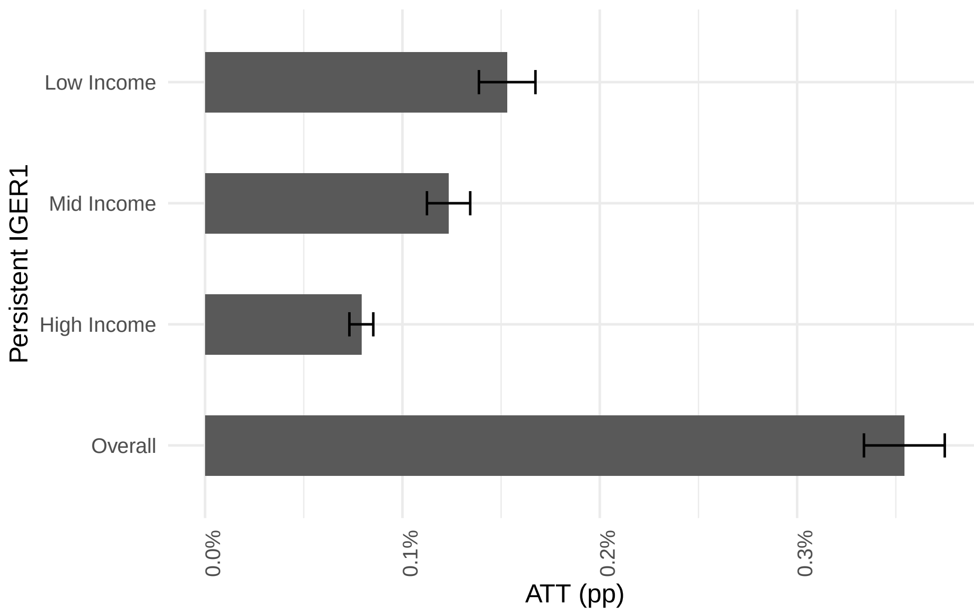

Figure W 15: **Persistent $IGER_1$ Treatment Effect by Subgroup**

Note. Average treatment effects on persistent $IGER_1$ (individuals above the 1% deposit-to-income threshold in both the current and prior non-overlapping three-month windows). Bars show the ATT estimated via Generalized Synthetic Control; error bars are 95% confidence intervals.

Web Appendix E14. Persistence Decomposition at Higher IGER Thresholds

In the main text, we decompose the IGER_1 effect into persistent and non-persistent components, where persistent IGER_k denotes individuals above the $k\%$ deposit-to-income threshold in both the current and prior non-overlapping three-month windows. Here we replicate this decomposition for IGER_5 and IGER_{10}.

Figure 16 reports the dynamic persistent and non-persistent IGER_5 effects. The persistent component (bottom panel) exhibits a clean post-launch increase of 0.09 pp, with tight confidence intervals and a gradual buildup consistent with the IGER_1 pattern. The non-persistent component (top panel) shows positive point estimates throughout, though with notably wider confidence intervals in the early post-treatment months. This early-period imprecision reflects high cross-state heterogeneity in the initial surge of new 5%-threshold crossers—a pattern that stabilizes as the flow of new crossers normalizes across states.

Figure 17 shows the corresponding decomposition for IGER_{10}. Again, the persistent component displays a gradual buildup and the non-persistent component shows a sharper initial spike that moderates over time. The persistent IGER_{10} effect is 0.05 pp and the non-persistent effect is 0.07 pp, both statistically significant.

Across all three thresholds, the share of the total effect attributable to persistent versus non-persistent gambling is strikingly stable: approximately 42–44% persistent and 56–58% non-persistent. This consistency suggests that the relative balance between newly elevated and chronically elevated deposit-to-income ratios is a robust feature of the treatment effect, not an artifact of the 1% threshold chosen for the main analysis.

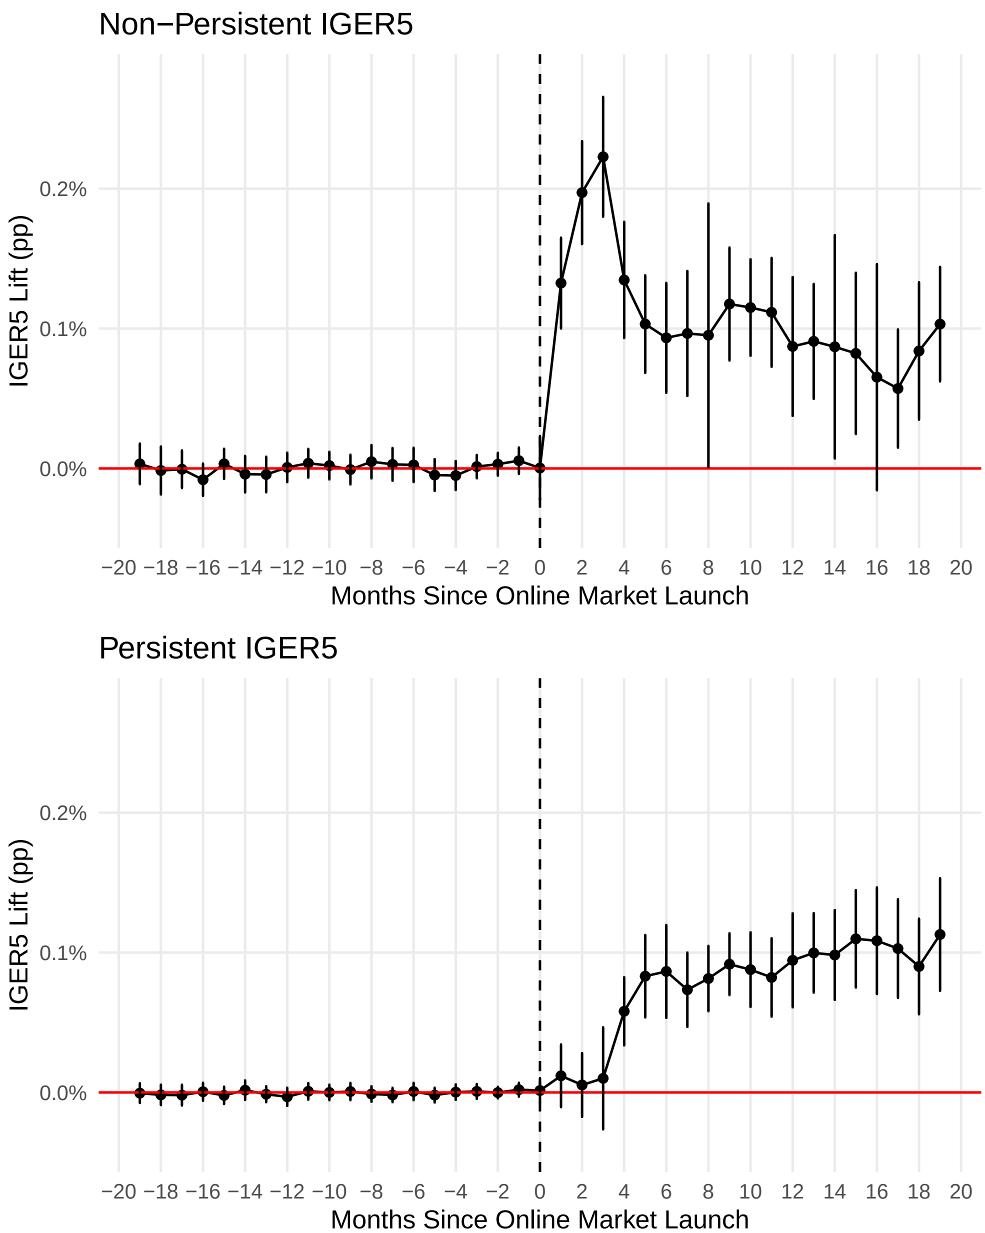

Figure W 16: **Dynamic Persistence Decomposition: IGER$_5$**

Note. Top panel: non-persistent IGER$_5$ rate lift (individuals above the 5% threshold now but not three months ago). Bottom panel: persistent IGER$_5$ rate lift (above 5% in both windows). Both panels share the same y-axis scale. Error bars are 95% bootstrap confidence intervals. The dashed vertical line marks the online market launch date.

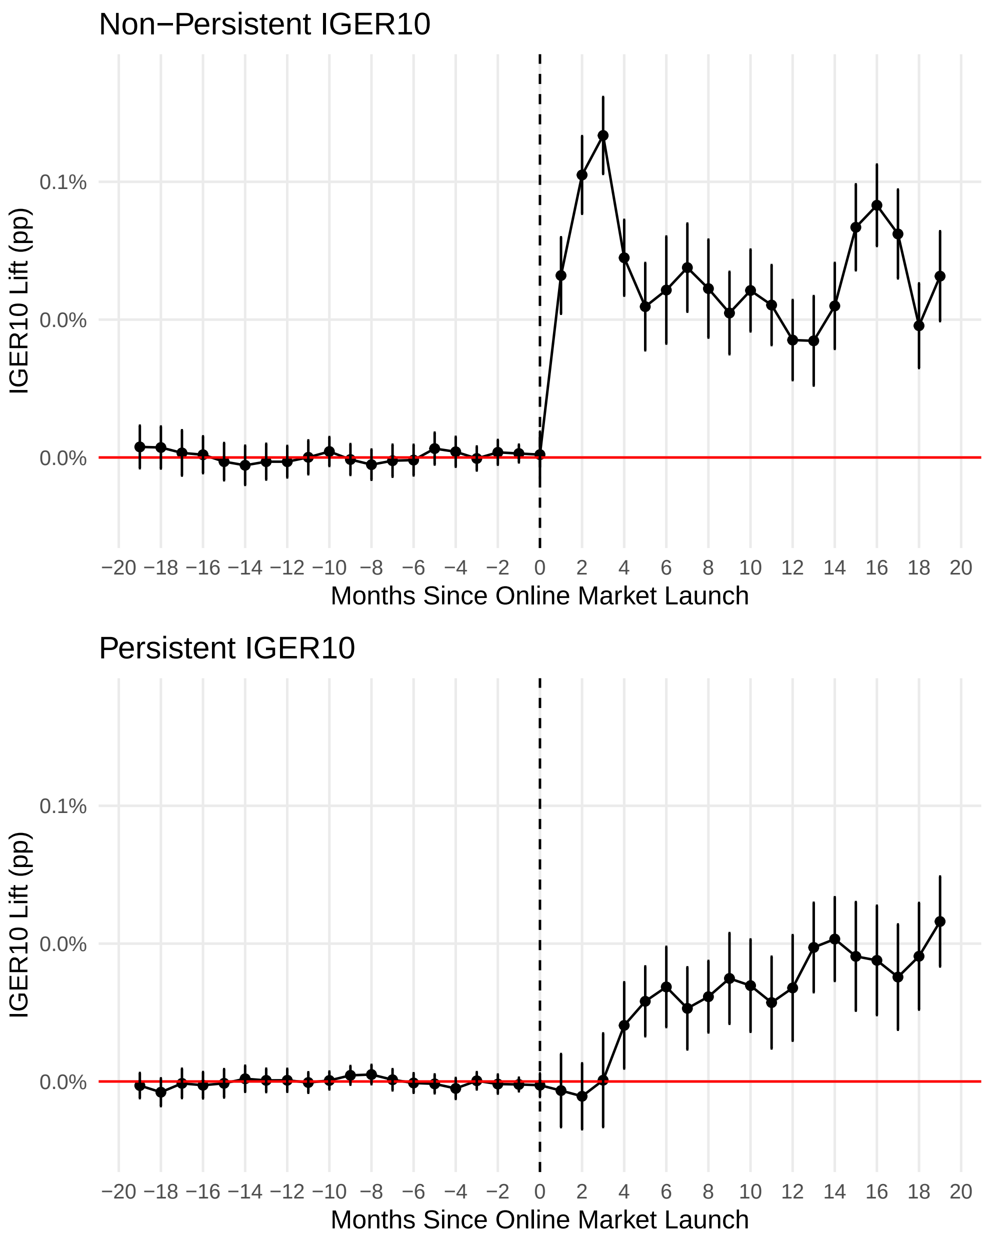

Figure W 17: **Dynamic Persistence Decomposition: IGER$_{10}$**

Note. Top panel: non-persistent IGER$_{10}$ rate lift (individuals above the 10% threshold now but not three months ago). Bottom panel: persistent IGER$_{10}$ rate lift (above 10% in both windows). Both panels share the same y-axis scale. Error bars are 95% bootstrap confidence intervals. The dashed vertical line marks the online market launch date.

Web Appendix E15. Narrow-Merchant Sensitivity

Our main analysis defines OSG spending using deposits to all 20 tracked operators. Some of these operators offer products beyond pure sportsbook activity, including horse racing (TVG Network), poker (PartyPoker), online casino gaming (Golden Nugget Online Gaming), and DFS-heavy platforms (Underdog Sports). To assess whether the main results depend on including these ambiguous operators, we re-estimate the pipeline after excluding these four merchants (which together account for approximately 6.2% of total deposits).

Table 11 compares the broad and narrow merchant-set ATTs for the primary outcomes. The narrow-merchant estimates are virtually identical to the broad estimates, confirming that the main results are not driven by the inclusion of ambiguous non-sportsbook operators.

Table W 11: Broad vs. Narrow Merchant Set ATT Comparison

	Broad (20 merchants)		Narrow (16 merchants)	
Outcome	ATT	SE	ATT	SE
OSG Spend (overall)	\$3.64	(0.065)	\$3.64	(0.065)
IGER_1 (overall)	0.78 pp	(0.088)	0.78 pp	(0.088)
IGER_1 (low-income)	1.14 pp	(0.167)	1.14 pp	(0.167)
IGER_1 (mid-income)	0.80 pp	(0.137)	0.80 pp	(0.137)
IGER_1 (high-income)	0.47 pp	(0.120)	0.47 pp	(0.120)

Note. Broad set includes all 20 tracked OSG operators. Narrow set excludes TVG Network, PartyPoker, Golden Nugget Online Gaming, and Underdog Sports. Both specifications use the same GSC methodology and control states.

Web Appendix E16. Income-Stability Sensitivity

Our main income measure is a Consumer Edge gross-income estimate derived from direct-deposit activity. To assess whether the regressive $IGER_1$ pattern depends on income measurement quality, we restrict the sample to members with stable income, defined as those whose monthly income coefficient of variation (CV) does not exceed 0.50. This filter retains 54.9% income-observed members with relatively predictable income trajectories and excludes those with the most variable or intermittent direct deposits. We additionally drop two control states (Vermont and Idaho) whose post-filter sample sizes fall below 200 members, leaving 10 control states; all 11 treatment states are retained. Table 12 compares the full-sample and stable-income ATTs for the primary outcomes.

Table W 12: Full Sample vs. Stable-Income Subsample ATT Comparison

	Full Sample (12 controls)		Stable Income (10 controls)	
Outcome	ATT	SE	ATT	SE
OSG Spend (overall)	\$3.64	(0.065)	\$3.33	(0.054)
$IGER_1$ (overall)	0.78 pp	(0.088)	0.67 pp	(0.066)
$IGER_1$ (low-income)	1.14 pp	(0.167)	0.85 pp	(0.109)
$IGER_1$ (mid-income)	0.80 pp	(0.137)	0.77 pp	(0.118)
$IGER_1$ (high-income)	0.47 pp	(0.120)	0.44 pp	(0.083)

Note. Sample includes all 11 treatment states and 12 control states. Stable-income subsample restricts to members whose monthly income coefficient of variation does not exceed 0.50 and excludes two control states (Vermont and Idaho) with fewer than 200 post-filter members; all 11 treatment states are retained. Both specifications use the same GSC methodology. ATTs are moderately smaller in the stable-income subsample, consistent with the exclusion of members with more variable financial profiles, but the regressive $IGER_1$ pattern across income groups—low > middle > high—persists.

Web Appendix F. Dynamic Results by State

In this section we report dynamic OSG and $IGER_1$ effects by state for pretreatment parallel trends evaluation. Pretreatment effects near zero indicate that our synthetic controls effectively mimic the treated state's outcomes prior to the intervention. In most states and time periods this holds true, suggesting that the synthetic control would have followed a similar path posttreatment in the absence of an intervention.

Web Appendix F1. OSG: Online-with-Offline

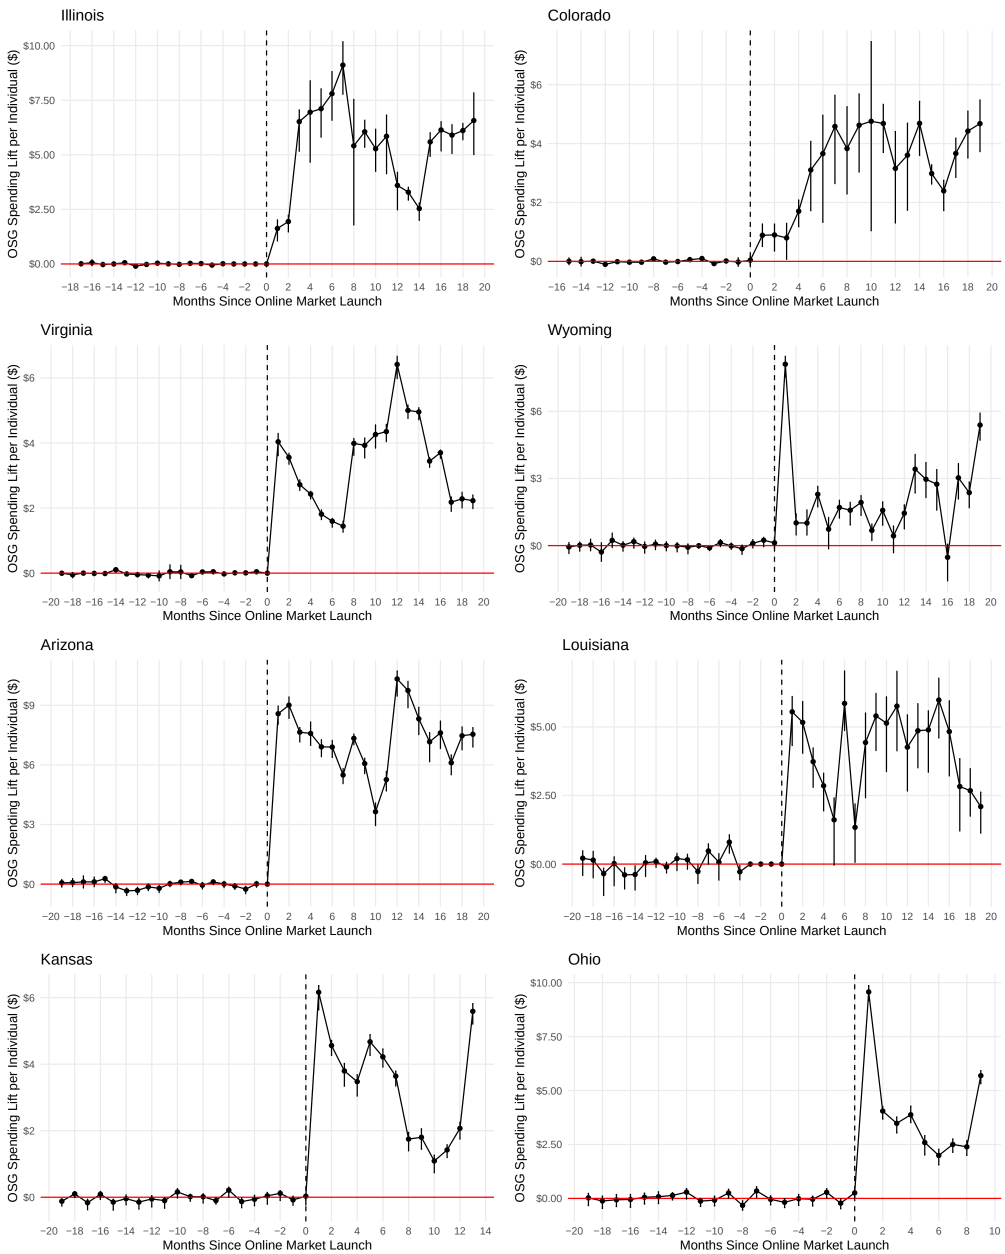

Figure W 18: **OSG Dynamic ATT by State (OWO)**

Note. Dynamic treatment effects over time for online sports gambling (OSG) spending per individual. Posttreatment observation spans 9 months for Ohio and 13 months for Kansas, all other states have at least 19 months posttreatment. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

Web Appendix F2. OSG: Online-After-Offline and Online-Only

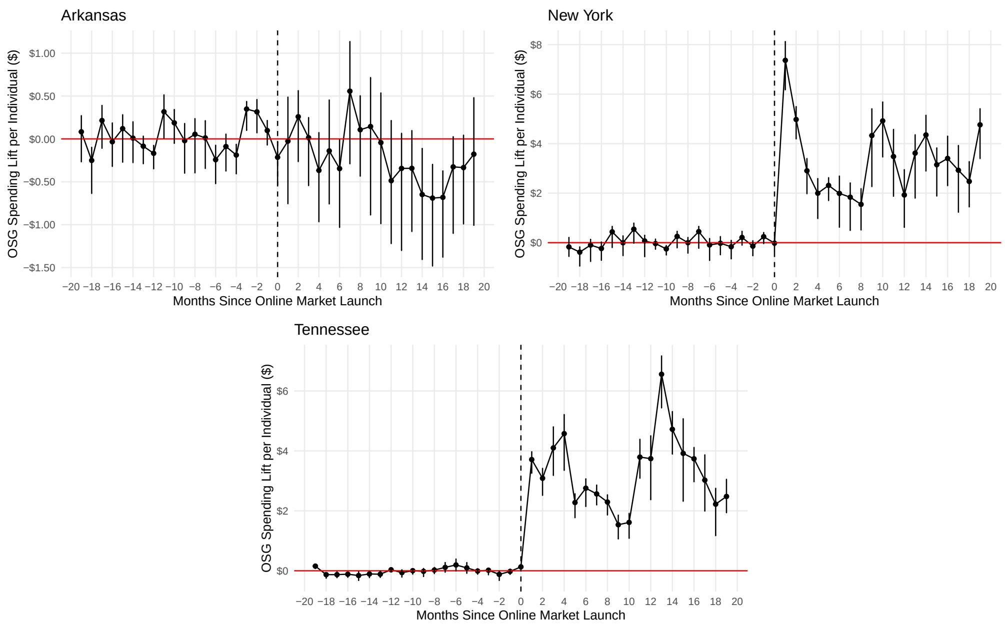

Figure W 19: **OSG Dynamic ATT by State (OAO and OO)**

Note. Dynamic treatment effects over time for online sports gambling (OSG) spending per individual. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

Web Appendix F3. $IGER_1$: Online-with-Offline

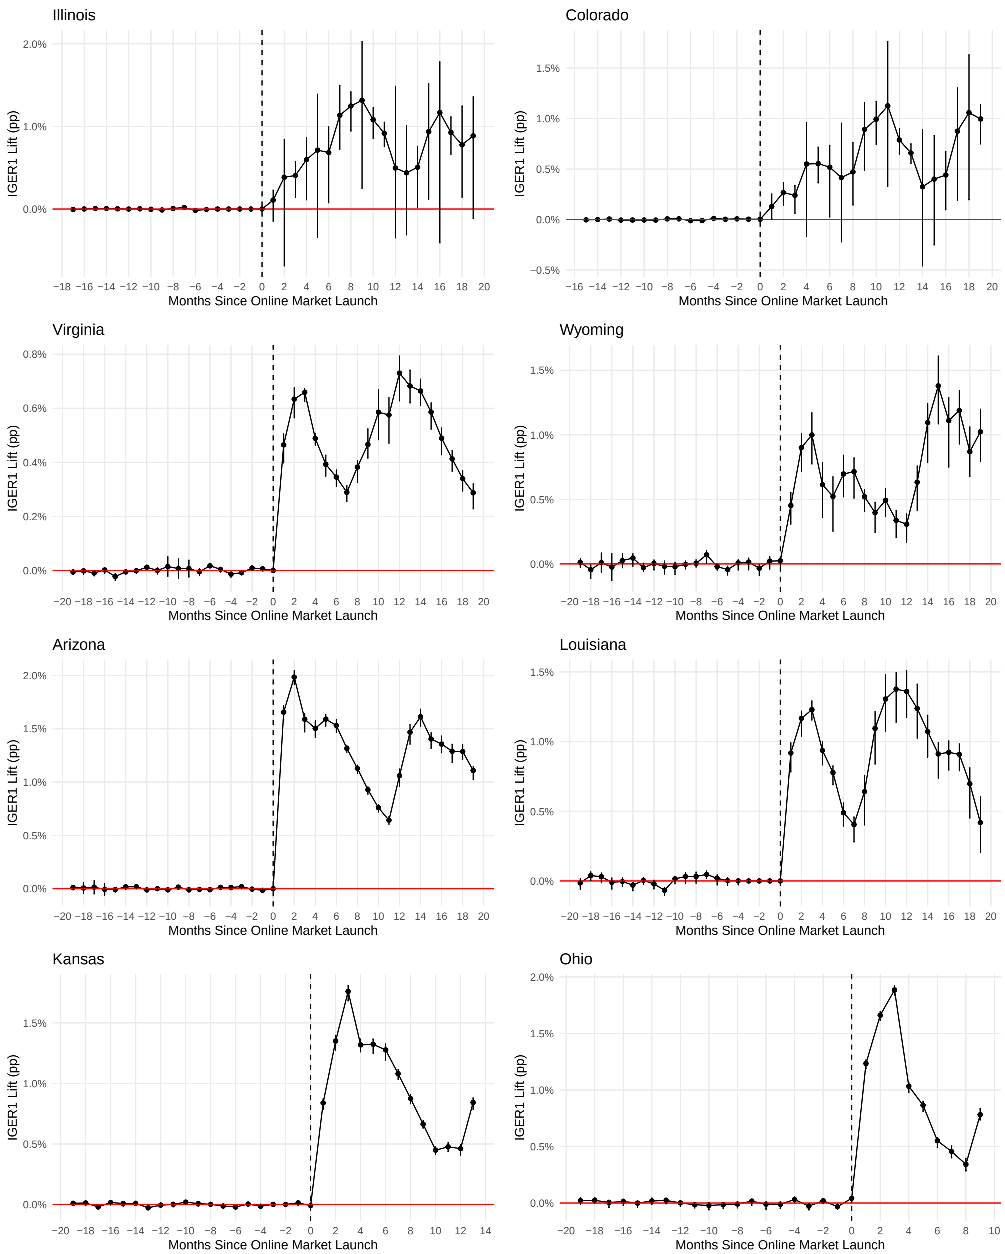

Figure W 20: **$IGER_1$ Dynamic ATT by State (OWO)**

Note. Dynamic treatment effects on $IGER_1$ over time. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Posttreatment observation spans 9 months for Ohio and 13 months for Kansas, all other states have at least 19 months posttreatment. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).

Web Appendix F4. $IGER_1$: Online-After-Offline and Online-Only

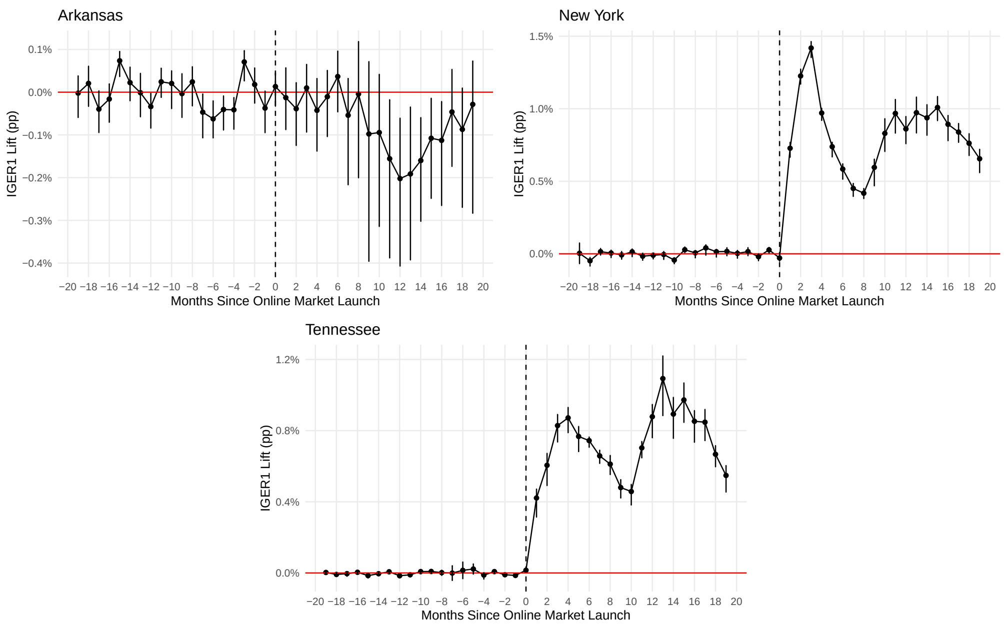

Figure W 21: **$IGER_1$ Dynamic ATT by State (OAO and OO)**

Note. Dynamic treatment effects on $IGER_1$ over time. Month 0 is the first calendar month in which legal online wagering was available; because some launches occurred mid-month, month 0 may reflect partial exposure. Web Appendix E5 shows similar results in the subset of states where online and offline betting launched within the same month. Error bars are 95% confidence intervals computed using the procedure from Li and Sonnier (2023).