Online Betting and Financial Distress: Evidence from Brazil

Rômullo Carvalho[1], Sérgio Firpo[2], Guilherme Freitas*[1], and Pedro Scatimburgo[1]

[1]StoneCo Economic Research [2]Insper Institute of Education and Research, São Paulo, Brazil

This version: August 2026

Abstract

Online betting became a mass phenomenon in Brazil. About 25 million individuals bet on authorized platforms in 2025, which are only a part of the total betting market. We investigate the impact of betting on several credit and financial outcomes of individuals who are clients of a financial institution. We match data at the transaction level, which allows us to date each individual's first transfer to a betting operator, authorized or not, to the Central Bank's credit registry. Using a staggered difference-in-differences design, conditional on pre-treatment characteristics, we compare those whose first bet falls in or after January 2025 with individuals never observed betting. We find that starting to bet raises the probability of delinquency by 19.7% and the credit-card delinquency rate by 38.7% from their pre-entry levels. Also, lenders cut the total credit limit by 16%, and the monthly change in the account balance falls by 43%. The cut is 2.6 times larger for those who had less credit before their first bet. Part of the rise in delinquency is a consequence of pre-existing distress: the cumulative probability of starting to bet is 2.6 pp higher after an individual enters delinquency. Bettors nationally are younger and likely of lower income than the individuals we observe, so our estimates are, if anything, attenuated.

Keywords: online betting; credit supply; credit risk; financial distress.
JEL codes: G21; G28; G51; D14; L83.

*Corresponding author: guilherme.marcal@stone.com.br. We thank Vinicius Carrasco, Gabriel de Almeida, Matheus Calvelli, Guilherme Cota, Nina Cunha and João Pedro Vieira for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the institution. All errors are our own.

1 Introduction

Fixed-odds online betting became a mass phenomenon in Brazil before it became a legal one. Over 20 million individuals were already sending money to a betting platform in 2024, when the authorization regime was not yet in force and the flow reached operators mostly through payment intermediaries registered under unrelated activity codes (Banco Central do Brasil, 2024). Law No. 14.790/2023 and its implementing regulations brought the activity inside a federal authorization regime, in force from January 2025 (Section 2). It did not open the activity; it formalized one that was already widespread. Brazil is by now among the largest betting markets in the world: on authorized platforms alone, gross gaming revenue – what bettors wagered less what was paid back to them, and so, by construction, their aggregate net loss – reached about R$37 billion in 2025 (Secretaria de Prêmios e Apostas, 2026).

Formalization changed which operators the payment record can attribute to betting, and Figure 1 shows the two resulting views of the expansion side by side. Panel (a) is the official statistic from the Central Bank, and its jump in January 2025 dates the arrival of the register rather than the arrival of the market. Panel (b) is the measure we use in this paper, computed for the individuals in our sample rather than for the country and built by identifying betting payees from the transaction record itself, so that it reaches authorized and unauthorized operators alike over the full window (Section 3.2); on it, the share of these individuals sending Pix to a betting operator roughly doubles over two years and rises steadily, without a break at the entry into force. What changed in January 2025 was the visibility of the market, not its trajectory.

Participation is intermittent: the median bettor deposits in two of their first six months, so a monthly share below one in ten cumulates over the window to almost one in five of the individuals in our sample having bet at some point, of whom those who started from January 2025 onward are the larger half (Section 3.3). The money sent does not come back. Nine in ten of those who bet withdraw less from operators than they deposit: what goes to betting is, for almost all of them, money lost.

What betting does to the financial position of the person who starts is a public question in Brazil. One side is the bettor's own behavior: whether money sent to operators comes at the cost of paying bills, of saving, or of servicing debt. The other is the lender's: whether a betting habit, once visible, changes the credit the person is offered. Both are hard to separate from a third possibility, that the distress came first and the betting followed. Bettors differ from non-bettors along many unobserved dimensions, and financial distress can be a consequence of betting or its cause – someone already under strain may turn to betting as an

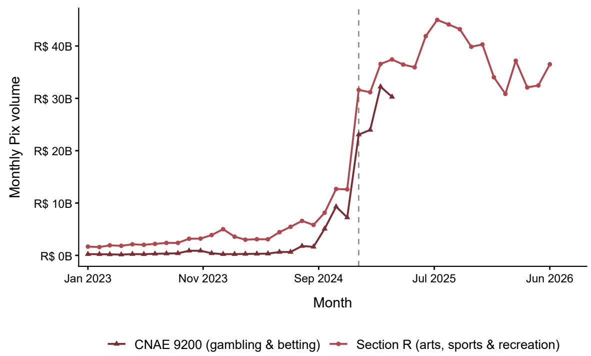

(a) The market: official Pix statistics

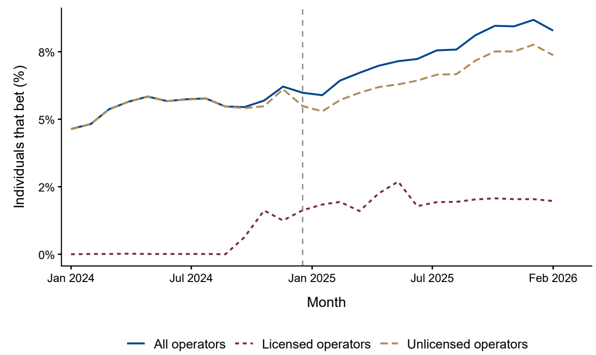

(b) Our data: individuals at one institution

Figure 1: Market expansion around the regulatory transition.

Notes: Panel (a), from published statistics of the Central Bank of Brazil, is the monthly Pix volume for the whole country to CNAE Section R (arts, culture, sports & recreation) and to CNAE 9200-3/99 (gambling & betting), and so reaches only operators filed under those codes (public LAI response, protocol 18810.012487/2025-61). Panel (b) is computed from our own data: among individuals with an active account at the institution whose records we use, the share that sent Pix to a betting operator in the month, with betting payees identified from the transaction record (Section 3.2); three series, to any operator, to an authorized one, and to one outside the regime. The two are not comparable in level or in unit: volume in reais against a share of individuals, and the gambling code against authorized and unauthorized operators alike. The dashed line marks the entry into force of the regulated regime (January 2025).

attempted way out – so observing betting and financial outcomes for the same individual through time does not tell the two directions apart.

We bring two data sources together for the same individuals, both from a Brazilian payments and credit institution. The first is the transaction-level record of the banking account, which identifies the payee of each transfer and so dates the month in which a person first sends money to a betting operator – what we call the first bet. The second is the credit registry of the Central Bank of Brazil (SCR). Matched by taxpayer number, the two yield a panel running from January 2024 to April 2026 that carries, for the same individual and month, the amount deposited with betting operators, two measures of delinquency, the total credit limit, and the monthly change in the account balance – a joint observation of exposure, credit and liquidity that neither administrative records nor household surveys currently provide. A large share of these individuals also carry out some entrepreneurial activity, which is one reason the balance we observe is a measure of liquidity rather than of household saving alone. We also classify each betting payee by product type, which is how we observe that the market these individuals face is mostly casino rather than sports (Appendix A.3).

The 2025 wave of first bets supplies both the sample the design needs and cohorts entering in different months, all under a single regime, so that the first bet means the same thing for every cohort – which matters where cohorts serve as each other's controls. Those who were already betting before 2025 faced a market whose counterparty risk plausibly selected on a tolerance for risk we neither observe nor can condition on, so we set them aside and compare the individuals whose first bet falls in or after January 2025 (*new bettors*) with those who never bet (*never-bettors*), in a staggered difference-in-differences design estimated with the doubly robust estimator of Callaway and Sant'Anna (2021). Neither group is betting in the months that precede entry, so the pre-period contains no betting by construction. The regulation itself is not the treatment: it shifts the environment of both groups alike and is absorbed by the calendar-month effects, leaving the timing of the individual's own first bet as the variation the design exploits.

Starting to bet worsens the individual's financial position, on both sides of the credit relationship and beyond it. The delinquency flag rises by 19.7% and the credit-card delinquency rate by 38.7%, while the total credit limit, a quantity set by lenders rather than chosen by the borrower, is cut by 16%. On the real side, the monthly change in the account balance falls by at least 43%. Part of the rise in delinquency is distress that came before the first bet, and the design measures it rather than assuming it away: the two groups already differ before treatment, new bettors holding a smaller credit portfolio and limit and a larger share of high-cost credit, and a placebo on a false event signs that selection, so netting it out leaves the increase at about 14% in the flag and 35% in the card rate. Reversing the roles of the two events, so that entry into delinquency is the treatment and betting the outcome, dates the same selection instead of inferring it: the cumulative probability of starting to bet is 2.6 pp higher after an individual enters delinquency ($p < 0.01$).

The two effects do not have the same source. The rise in delinquency needs nothing of the lender: money sent to a betting operator is money not available for debt service. The limit cut appears instead to be a screening response, with the start of betting entering the lender's assessment as new information about the borrower, on a route we do not observe (Section 5.1). The cut is 2.6 times larger for individuals whose credit limit before the first bet was in the lower half of the distribution, and it does not vary with age; the delinquency flag rises by the same amount whatever the prior credit position and concentrates among the youngest. The regressivity is therefore one of credit supply rather than of realized risk. The drawdown in the account balance does not follow that gradient: it is if anything larger among those who had more credit before starting, and it reaches beyond the credit relationship.

Our contributions are four. The first is the estimand: the effect on the individuals who actually start betting, rather than an intention-to-treat effect read off legalization, in a population most of whom never bet. Because who starts is not random, we read the placebo on a false event as a signed bound rather than as a test to be passed, and we date distress relative to the first bet rather than assuming its order. The second is that one of the outcomes is set by the lender: the credit limit is a supply-side quantity, so the credit response to betting is measured rather than inferred from what the borrower does. The third is the joint observation of the payment flow, the credit registry and the account balance for the same individuals over the same months, which lets exposure, realized default and liquidity be read against one another. The fourth is coverage: we capture betting through unauthorized as well as authorized operators, in a market that is mostly casino rather than sports, and so not the market on which evidence identified off sportsbook legalization rests.

Related literature

The direction from betting to distress is where the recent applied work concentrates. In bank and card microdata the evidence is associative: gambling spending as a share of income goes with financial distress and, at the extreme, with unemployment and elevated mortality (Muggleton et al., 2021). The wave of U.S. sports-betting legalization supplies the causal counterparts: betting crowds out saving and investment rather than consumption (Baker et al., 2024), and credit scores fall and bankruptcies rise among the financially weaker (Hollenbeck et al., 2025; Goss and Mangrum, 2026). What closes the loop back onto distress is a credit-supply reaction: observable borrower behavior can trigger repricing and rationing under asymmetric information (Stiglitz and Weiss, 1981; Agarwal et al., 2009), and Gao et al. (2023) show, with fintech-lender data, that of all consumption categories gambling is the one that most reduces a lender's willingness to extend credit. Our design makes the opposite trade to those studies: legalization settles selection by construction, at the cost of an intention-to-treat effect on a population most of whom never bet, while dating treatment to the individual's own first bet delivers the effect on those who actually bet and makes selection a threat to be handled rather than assumed away.

On the direction from distress to betting, a long tradition sets an "entertainment" motive for gambling against a "desperation" one. Blalock et al. (2007) give the canonical statement, with lottery sales rising with the poverty rate across thirty-nine U.S. states, and Herskowitz (2021) the sharper mechanism: betting is a second-best, high-variance liquidity technology under financial constraints, and a randomized savings treatment that relaxes the constraint lowers demand for it. On betting sustaining itself, Allcott et al. (2022) identify, from a randomized experiment, both habit formation and partially unrecognized self-control problems in a product delivered through the same always-available smartphone interface. We do not decompose the two for betting, but that literature is why betting, once started, need not be self-correcting, and why effects estimated over a one-year horizon should be read as the beginning of the response rather than the whole of it.

For Brazil, the evidence is recent and still thin. Lauro and Sacramento (2026), using betting data for 1,046 bettors on a Brazilian platform, report a median return near −72%, deepening for those who persist, a pattern that echoes the negative expected value and near-absent learning of Brazilian retail day trading (Chague and Giovannetti, 2020); Abreu and Bressan (2026), in turn, characterize the determinants of betting among Brazilian investors.

The rest of the paper is organized as follows. Section 2 describes the regulatory context, and Section 3 the data, the identification of betting operators and the bettor population. Section 4 sets out the design and defines the estimands. Section 5.1 reports the consequences of starting to bet for credit and liquidity, Section 5.2 where those consequences concentrate, and Section 5.3 the reverse arm, in which entry into delinquency itself precedes the first bet. Section 6 concludes.

2 Institutional and regulatory context

Brazil banned casinos and most forms of gambling in 1946, yet an informal betting market thrived for decades – from *jogo do bicho* (an illegal numbers lottery) to the brief legalization of bingo halls between 1993 and 2004. Law No. 13.756/2018 created fixed-odds betting as an exclusive public service of the Federal Government and gave the Ministry of Finance two years to regulate it (Brazil, 2018). The deadline lapsed in December 2022 without regulation, and between 2018 and 2023 operators based mostly abroad served Brazilian bettors without a domestic legal entity or a public registry (CNPJ).

Law No. 14.790/2023 closed the regulatory gap (Brazil, 2023). It imposed taxation on operators' gross revenue, required federal authorization to operate, and mandated that only legal entities incorporated under Brazilian law, with headquarters and administration in the national territory, were eligible for authorization. To implement the new regime, the federal government created the Secretariat of Prizes and Betting (SPA/MF), housed within the Ministry of Finance (Decree No. 11.907/2024), with the mandate to regulate, authorize, and supervise fixed-odds betting operators. Since 2024 the SPA/MF has issued the regulations (*portarias*) that operationalized the law; Table 1 lists them.

Two regulations matter for the payment infrastructure and, therefore, the data we observe. Regulation No. 615/2024 restricted deposits and withdrawals to transfers between the bettor's own registered account and the operator's account, accepting only Pix (Brazil's instant payment system, launched in November 2020 and now near-universal), debit or prepaid cards, and bank wires (Secretaria de Prêmios e Apostas, 2024); in practice, Pix carries roughly 97% of all betting deposits (Pay4Fun, 2025). Regulation No. 722/2024 set technical requirements for the platforms, including the `.bet.br` domain, and imposed

Table 1: Chronology of the regulation of fixed-odds betting in Brazil

Date	Instrument	Effect
1946	Decree-Law 9.215	Bans casinos and most forms of gambling.
Dec 2018	Law 13.756/2018	Creates the fixed-odds betting modality (Art. 29).
Dec 2022	—	Two-year deadline to regulate expires without action.
Jul 2023	Prov. Measure 1.182	Reopens the topic; provides for commercial exploitation.
Dec 2023	Law 14.790/2023	Regulatory framework: taxation, federal authorization, domestic legal entity required.
Jan 2024	Decree 11.907/2024	Creates the SPA/MF within the Ministry of Finance.
Apr 2024	Reg. 615/2024	Restricts deposits/withdrawals to Pix, debit/prepaid cards, and bank wires.
May 2024	Reg. 722/2024	`.bet.br` domain; CPF and facial-biometric registration.
May 2024	Reg. 827/2024	Online casino rules: certified algorithms, 85% min. RTP, probability disclosure.
Sep 2024	Reg. 1.475/2024	Only firms that filed a request may operate during adjustment period.
Oct–Nov 2024	Blocking orders	Anatel/SPA begin blocking unauthorized domains.
Jan 2025	Entry into force	Only authorized operators, with domestic tax registry and `.bet.br`.
Mar 2025	Reg. 566/2025	Bars financial institutions from holding illegal operators' accounts.
Nov 2025	Reg. 2.579/2025	Centralized self-exclusion platform; mandatory prudential self-limits.

Sources: laws and decrees from the Presidency; SPA/MF regulations (*portarias*) from the Federal Official Gazette.

bettor-registration rules: full CPF validation and facial-recognition biometrics at sign-up, with re-authentication every seven days. Before regulation, most betting payments were routed through payment intermediaries registered under unrelated activity codes, making the betting destination difficult to identify from the payment record alone. The new rules tied each deposit to a single, biometrically verified individual and routed transfers through Pix to the operator's own tax registry (CNPJ), making betting flows directly observable in the payments data for authorized operators. Unauthorized operators and intermediaries that route their payments remain outside this observable perimeter; identifying them requires the procedure described in Section 3.2 and Appendix A.1.

The regulated regime took full effect on January 1, 2025, from which date only authorized operators could operate nationally. Enforcement relied on a partnership between the SPA/MF and Anatel (the telecommunications regulator), which began administratively blocking the domains of unauthorized operators in October–November 2024. A separate front targeted the payments channel: Regulation No. 566 of March 2025 barred financial institutions from holding accounts or processing deposits of unauthorized operators, closing the remaining payment route.

Despite these measures, the illegal market has proven resilient. The Senate's Parliamentary Commission of Inquiry into online betting (*CPI das Bets*), installed in November 2024, documented widespread money-laundering schemes and tax evasion through unauthorized operators. Its final report was voted down in June 2025, but the evidence it gathered informed subsequent enforcement actions. The Federal Court of Accounts (TCU), in Ruling No. 1,296/2026, estimated that 41–51% of the market remained with illegal operators even after the regime's entry into force (Tribunal de Contas da União, 2026). By mid-2026, Anatel had blocked nearly 50,000 unauthorized domains at the SPA/MF's request, yet operators frequently migrate to new ones (Secretaria de Prêmios e Apostas, 2026).

Although the regulatory framework uses the umbrella term "fixed-odds betting," the platforms it covers span a wide range of products. Figure 2 illustrates the spectrum: from sports-first operators offering pre-match and live odds on football leagues, through mixed platforms that combine a sportsbook with online casino games to casino-first operators whose landing pages default to slot games, crash games, and virtual table games.

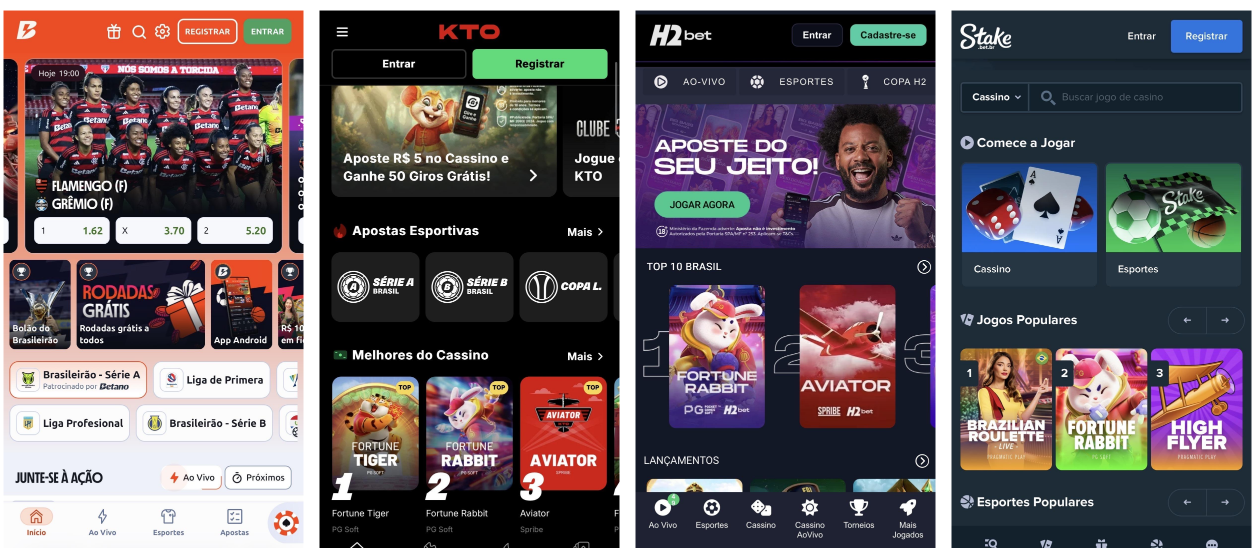

Figure 2: Illustrative examples of authorized betting platforms

Notes: Screenshots of four authorized operators under `.bet.br` domains, selected to illustrate the product spectrum, captured on July 2026. From left to right: a sports-first operator displaying pre-match odds for Brazilian soccer; a mixed operator offering both league betting and online casino games, with *Fortune Tiger*, *Fortune Rabbit*, and *Aviator* as its top three casino products; a second mixed operator whose landing page highlights *Fortune Rabbit* and *Aviator* as the two most popular games nationwide; and a casino-first operator where the home screen defaults to slot and table games.

The slot game *Fortune Tiger* – commonly known as *jogo do tigrinho* – became the most visible symbol of online casino proliferation in Brazil. Digital influencers promoted it using manipulated demonstration accounts that simulated winnings, triggering widespread public backlash; the SPA/MF opened over 400 enforcement proceedings against influencer advertising and ordered the removal of hundreds of profiles and posts (Secretaria de Prêmios e Apostas, 2026). Casino-only platforms are particularly prevalent among unauthorized operators. The regulatory framework treats these casino games as a subcategory of fixed-odds betting: Regulation No. 827/2024 requires operators to certify their algorithms, maintain a minimum return-to-player (RTP) rate of 85%, and disclose winning probabilities before each wager.

The channel through which bettors accessed platforms also changed after regulation. Until mid-2025, access was almost exclusively through mobile browsers since major app stores did not list real-money gambling apps for Brazil; by that point, betting sites under `.bet.br` had already become the second-largest internet destination in the country after Google (Poder360, 2025). Google Play reversed its policy in June 2025 and Apple followed in May 2026, and Figure 3 shows the immediate effect: Brazil went from a negligible share of global sportsbook and casino downloads to roughly 30% of the worldwide total in the second half of 2025. Native apps lower the friction of repeated use and offer capabilities that mobile browsers cannot fully replicate – such as push notifications with live odds, in-app betting during matches, and persistent home-screen access – which may deepen engagement among existing users.[1]

The scale of the industry is visible in its marketing as well. In the ranking of Brazil's 300 largest advertisers, fifteen betting companies bought a combined 2.9 billion reais in media in 2025, up from 2.0 billion by eleven companies the year before, which placed the sector second only to large retailers in share of total advertising spending (Kantar Ibope Media, 2026).[2] Advertising reaches authorized and unauthorized operators alike: the SPA/MF restricts promotional content for the former, while the latter advertise outside the perimeter of the rule.

The supply side of the market moved with the regime, and it did so almost entirely through the authorized layer. Counting the distinct betting destinations that receive Pix from the individuals we observe, the number of authorized operators active in a month goes from one over most of 2024 to twelve in December 2024, and then to 58 in January 2025, settling near 77 from mid-2025 onward, which is close to the size of the register itself. The number of destinations outside the regime does not fall with it: it runs between roughly 260 and 400 across the whole window, peaks in the first half of 2025, and only drifts down from

[1]Early adopters among Brazilian operators reported markedly higher user engagement after launching on the official stores (iGaming Business, 2026).

[2]A separate estimate by Tunad, covering television, radio and streaming but excluding digital-only channels, puts the figure at 1.4 billion reais for 2025, with open television as the central channel (Tunad, 2025; Lance!, 2025).

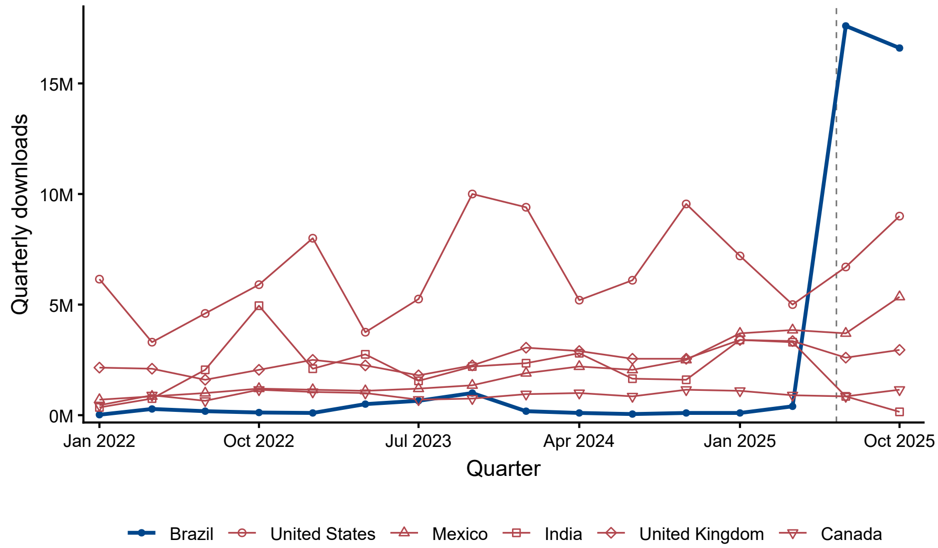

Figure 3: Sportsbook and casino app downloads by country

Notes: Quarterly downloads of apps classified as "Sportsbook & Casino" by Sensor Tower, for the six largest markets by 2025 volume. The dashed line marks the opening of Google Play to authorized gambling apps in Brazil (June 2025). Apple followed in May 2026. Brazil went from near zero to over 15 million downloads per quarter, accounting for roughly 30% of the global total in H2 2025. Source: Sensor Tower, *State of Mobile 2026*.

2026 as enforcement accumulates. What the regime added to the market, on this margin, was an authorized segment alongside the existing one rather than a replacement of it, which is the same conclusion the TCU reaches from the revenue side.

3 Data and sample

3.1 Sources

Both of our data sources come from a Brazilian payments and credit institution and are observed monthly for the same set of individuals. The first is the transaction-level record of each individual's account, which decomposes inflows, outflows, and the end-of-month balance; within the outflows, we identify the Pix transfers whose recipient is a betting operator (Section 3.2).

Betting platforms operate through a wallet held on the operator's side: to place bets, an individual funds that wallet by transferring money from their own bank account to the operator, and to cash out the operator transfers funds back to the same account.[3] What the account record captures are exactly these two legs – the deposit (funds sent from the individual's account to the operator) and the withdrawal (funds returning from the

[3] Pix is the dominant funding channel for betting deposits (Section 2). The fraction we miss – a funding via debit/pre-paid card or bank wire – is residual.

operator). We do not see individual bets, odds, or the wallet balance held on the platform; the operator-side ledger lies outside our data.

The second source is the credit registry of the Central Bank of Brazil's Credit Information System (SCR), a snapshot of every credit operation an individual holds in the national financial system, by modality. Some demographic data (age, gender, and geocoded address) complete our dataset, and the taxpayer identifier (CPF) links the sources.

The units we analyze are individuals, identified in the credit registry by their taxpayer number; a large share of them also carry out some entrepreneurial activity, and part of this group holds the account we observe under the tax registry of that activity rather than as a natural person. The balance we observe is therefore a transactional balance that may combine personal income with the cash generated by that entrepreneurial activity, a distinction that we return to in Section 5.1 and whose weight we test in Appendix B.4 by restricting the sample to natural-person accounts.

From these sources we build two panels on the same individuals. From the SCR come the credit outcomes: a delinquency flag (a balance overdue by ninety days or more), the credit-card delinquency rate (the delinquent balance over the portfolio in the card modalities), and the log of the total credit limit; the log of the total credit portfolio enters as a pre-treatment covariate rather than as a reported outcome. From the account flow comes the end-of-month balance and its monthly change.

This population accounts for a small fraction of the national betting market: roughly 1–3% of bettors by number of individuals.[4] The demographic composition is nevertheless broadly consistent with the national profile reported by the regulator: among individuals transacting with authorized operators, approximately 70% are under 40 (versus 73% nationally), with a comparable concentration in the 25–40 age range. The gap is concentrated among the youngest: only 11% are under 25, against 22% nationally, which is expected given that the sample over-represents working-age individuals with some entrepreneurial activity. The female share is higher (43% versus 32%), consistent with the gender composition of this population. The analytical value of the sample rests not on its scale but on the joint observation of betting flows, credit registry outcomes, and account balances at the individual level – a combination that neither administrative records nor household surveys currently provide.

The panels are monthly. The transaction data used to identify betting operators span 2020 to mid-2026, and the estimation panels run from January 2024 to mid-2026.

[4]The SPA/MF reports approximately 25 million distinct bettors on authorized platforms over 2025 (Secretaria de Prêmios e Apostas, 2026). Over calendar 2025, restricting our count to authorized operators yields 267,000 distinct individuals (1.1%); including all identified bettor operators raises it to 703,000 (2.8%).

3.2 Identification of betting operators

Identifying who bet requires identifying which Pix recipients (payees) belong to the betting ecosystem – a nontrivial task, for the reasons discussed in Section 2. We do so through four complementary approaches, set out in full in Appendix A: to the SPA/MF register of authorized operators, we add the payees flagged by their own transactional pattern, and then the payment intermediaries whose pattern is too diluted to flag, recovered instead from the confirmed bettors we can identify among their payers, before and after regulation. Because an authorized operator receives Pix under its own tax registry (CNPJ) only once the regime is in force, the register becomes visible to us from January 2025, whereas the transactional pattern, being a property of the payee's own Pix, spans the full sample.

We conduct external checks on each part of the list, reported in Appendix A.2. Two are worth stating here. The transactional score alone flags 91.5% of the authorized operators with enough volume. And of the 67 payment gateways that the SPA/MF identified as serving unauthorized operators – a list made public during the Senate's 2025 inquiry into betting (*CPI das Bets*) (Secretaria de Prêmios e Apostas, 2025) – our list contains 64 (96%), 21 of them recovered through the pre-regulation intermediary layer, because that is the layer a direct reading of the 2024 data cannot see.

This procedure yields 1,260 betting destinations: 87 authorized by the SPA/MF and 1,173 unauthorized payees – unauthorized operators and payment intermediaries, which our data do not let us tell apart. Because the two intermediary approaches are the ones identified indirectly, Appendix B.2 re-estimates every result of Section 5 on the universe that keeps only the register and the transactional score.

Classifying operators: casino versus sports We also try to classify each operator by the kind of betting it offers – casino games, continuous and high-frequency, versus sports betting, event-based and lower-frequency – since the two may carry different risks. This is not straightforward: we observe a Pix to a payee, not the game it financed, so the type can only be read from the operator behind it. This became possible only from January 2025, once the payee identifies the brand and its predominant product.

The label is built from the brands behind each operator, whose platform content we score on a casino-to-sports scale using a large-language model, sorting it into sports-type, casino-type, hybrid, or inconclusive when the content is too thin to score. The procedure classifies 255 operators with a determinable profile, covering about 66% of post-regulation bettors, and three checks support the labels (Appendix A.3).

One structural fact conditions everything built on this classification. Sports betting is almost exclusively a regulated phenomenon: 27 of the 28 sports-type operators are authorized, while the unauthorized operators concentrate casino platforms and inconclusive intermediaries (Appendix Table A2). At the bettor level, among those 255 operators, bettors classified as sports route about 97% of their flow through authorized operators, while those classified as casino route about 65% through unauthorized ones.

3.3 The bettor population

Betting is widespread in this population: almost one in five individuals whose account at the institution was active over the window sent Pix to a betting operator at some point, split between the 883,920 who started doing so from January 2025 onward and the 712,778 who were already betting before. Against the 7,631,337 individuals eligible to start – the never-bettors plus the 2025 entrants – the 2025 initiation rate is about 12%. Persistence is short: over the individual's own first six months of betting, 44.5% deposit in a single month and 36.0% in three or more. The number of operators splits the population in a similar way: 41.5% stay with one operator over those six months, while 28.8% reach five or more. The median bettor thus has two betting months and two operators. Amounts are small at the median and concentrated in the upper tail – R$79 deposited over the six months, against R$354 at the 75th percentile. The same dispersion appears when deposits are compared to the total outflows from the account: 0.6% at the median against 22.8% at the 90th percentile. Over those six months, 82.5% of bettors never withdraw from an operator and 90.8% deposit more than they withdraw (Table 2).[5] That 90.8% is the payments-side counterpart of what the wager-level evidence for Brazil finds where the operator's own ledger is observed: Lauro and Sacramento (2026), with the bet-by-bet records of 1,046 bettors on a Brazilian platform, report that 92.5% end with a negative cumulative result, a near-universal loss of the same order as the one Chague and Giovannetti (2020) document for retail day trading.

The classification of Section 3.2 assigns a dominant operator type to 46.6% of bettors, those whose deposits reach an operator we classify. They deposit more than the population as a whole, R$218 or R$195 at the median against R$79, so the two type columns are informative against each other rather than against the first. Casino bettors deposit more often and at more platforms: a median of three betting months and six distinct operators over the first six months, against two and two for sports bettors, with 56.1% of casino bettors depositing in three or more months and 56.2% reaching five or more operators, against 47.2% and 29.4% of sports bettors. Median deposits are of similar magnitude, while the upper tail is heavier

[5] Wagers and payouts settle inside the operator's wallet, which we do not observe (Section 3.1). Two limits follow. Deposits understate how much is wagered, because a balance won inside the wallet can be re-wagered without a new deposit. And the net result overstates losses, because winnings left with the operator never return to the account.

Table 2: Descriptive statistics: bettor profiles over the first six months of betting

		By operator type		*By age group*			
	All bettors	Casino	Sports	18–24	25–34	35–44	45+
Panel A: Persistence, deposits, and net result							
Betting months (median)	2	3	2	2	2	2	2
Three or more betting months (%)	36.0	56.1	47.2	35.9	37.5	37.9	35.9
Distinct operators (median)	2	6	2	2	2	2	2
Five or more operators (%)	28.8	56.2	29.4	33.5	32.9	31.1	25.9
Deposits (R$, median)	79	218	195	93	90	80	68
75th percentile	354	990	1,354	487	460	390	284
Never withdrew from an operator (%)	82.5	72.3	59.6	70.2	77.4	82.5	86.9
Net losers (%)	90.8	85.6	78.8	83.2	87.9	90.9	93.3
Panel B: Deposits as a share of account outflows							
Share of outflows (%, median)	0.6	2.3	2.4	1.8	1.2	0.7	0.4
90th percentile	22.8	30.4	44.7	38.5	32.0	22.8	13.3
per betting month	1.2	2.9	3.1	2.7	2.0	1.3	0.8
# of Bettors	1,241,364	496,148	82,213	99,434	386,056	327,242	254,989

Notes: All statistics are measured over each individual's first six months of betting, so that people who started at different dates are comparable. This requires the first deposit to precede the end of the data by at least six months, which corresponds to 82.9% of the 1,497,749 bettors observed in the January 2024–May 2026 flow window, against the 1,785,367 of Appendix Table A1, which spans the full identification window and 79 to 84% of each column. Deposits are Pix sent from the individual's account to an operator, withdrawals are Pix received from one; neither is a bet or a payout, since the operator's own ledger is outside our data (Section 3.1). A betting month is a month with at least one deposit. Net losers are individuals whose deposits exceed their withdrawals, and include by construction those who never withdrew. Panel B divides deposits to operators by total outflows over the same period. Each individual is assigned to the operator type taking the largest share of their deposits, considering the classification of Section 3.2; the two type columns cover 46.6% of the sample. Age is taken from the registration record, available for 86% of it; the rest appear only in the first column. Statistics are computed across individuals, one observation each; for the per-betting-month line, each individual is first summarised by the median across their own betting months.

on the sports side, R$1,354 at the 75th percentile versus R$990. As a share of total outflows from the account, deposits are again close at the median, 2.3% versus 2.4%, and the tails again diverge: 30.4% at the 90th percentile for casino bettors against 44.7% for sports bettors. The groups differ more in what comes back: 72.3% of casino bettors never withdraw from an operator and 85.6% deposit more than they withdraw, against 59.6% and 78.8% of sports bettors. Two features of these products may bear on this difference. Sports betting has a skill component, in that the bettor can act on information about the event and on the choice of odds, so a small group of informed bettors can sustain larger stakes and recover part of them, whereas casino outcomes are drawn from a fixed distribution the player cannot influence. And a casino round settles in seconds and is typically repeated many times within a session, so the operator margin accumulates over far more wagers per unit of deposit than in sports betting. Neither feature is measurable in our data, but both are consistent with the observed pattern: the casino profile is closer to repeated play with little realized return, the sports profile to occasional betting with a longer tail of amounts.

By age, persistence is flat and the weight of betting in the account is not. Between 35.9% and 37.9% of every age group deposits in three or more of their first six months, and the number of platforms falls only mildly, from 33.5% of the 18–24 group reaching five or more operators to 25.9% of the 45+ group. The share of outflows falls steeply: the median from 1.8% in the 18–24 group to 0.4% in the 45+ group, and the 90th percentile from 38.5% to 13.3%. Deposits vary much less, R$93 against R$68 over the same six months. Younger bettors also get money back more often: 70.2% of the 18–24 group never withdraw from an operator and 83.2% deposit more than they withdraw, against 86.9% and 93.3% of the 45+ group. What varies across age groups, then, is less how long people bet, or across how many platforms, than how large their deposits are relative to their account outflows. The withdrawal pattern fits that reading: the youngest bettors, who commit the largest share of their outflows, are also the ones who most often take money back from operators.

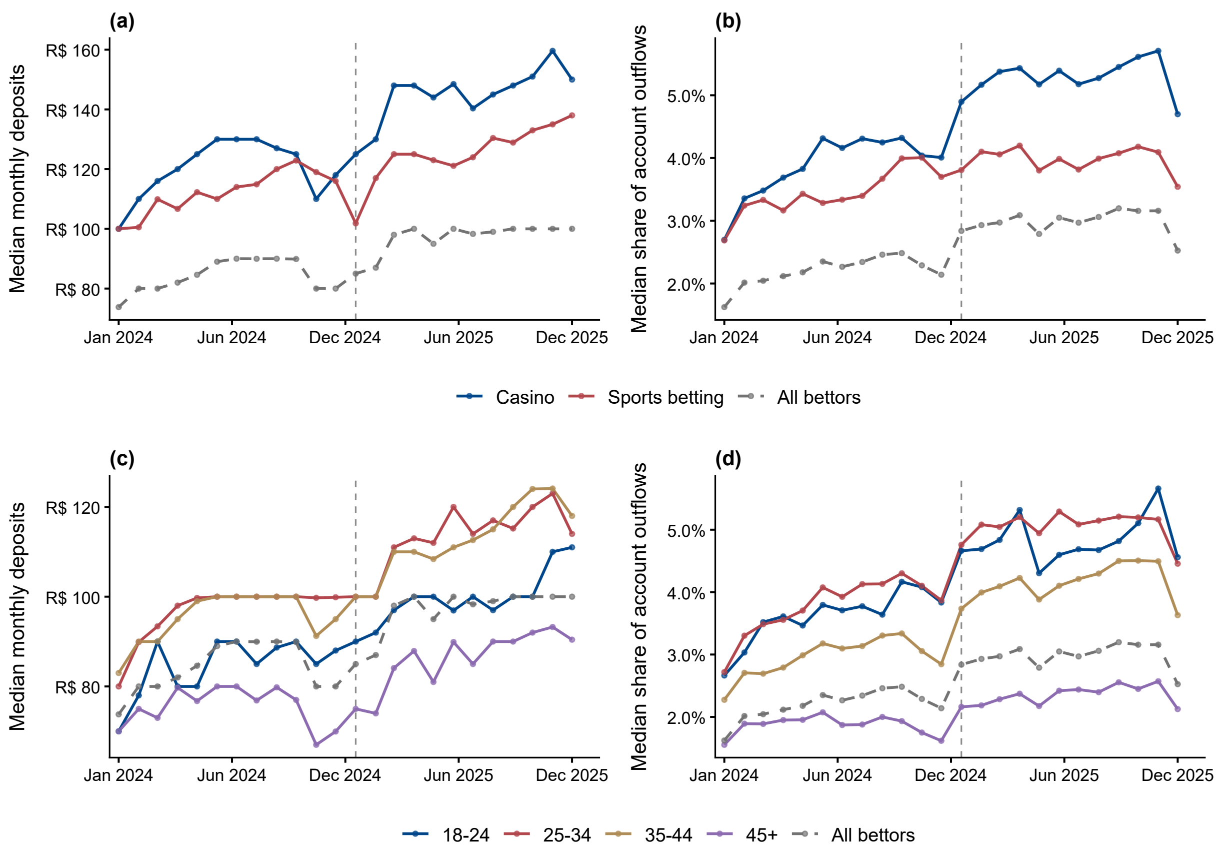

Figure 4: Betting intensity among continuing bettors.

Notes: Monthly medians of deposits (panels (a) and (c)) and of deposits as a share of the individual's total account outflows (panels (b) and (d)), by operator type in the top row and by age group in the bottom, with all bettors as the dashed reference series. Balanced panel of continuing bettors: individuals who deposited with an operator in both 2024 and 2025. Each median is taken over those who deposited in the month; points resting on fewer than 500 individuals are suppressed. The vertical line marks the beginning of the regulated regime in January 2025.

The statistics in Table 2 compare people at the same point in their own betting history. Figure 4 adds the time dimension, following the intensive margin among those who keep betting. It tracks the amount deposited and its share of account outflows month by month on a balanced panel of individuals who deposited with an operator in both 2024 and 2025, so movement in the series comes from behavior rather than from the arrival of new bettors. Deposits rise over the whole period for both operator types, including before January 2025: the median monthly deposit of casino bettors goes from R$118 in December 2024 to R$150 twelve months later, and of sports bettors from R$116 to R$138. On the share of account outflows the two types separate at the start of the regulated regime, the casino series stepping from 4.0% in December 2024 to 4.9% in January 2025 and climbing above 5% through the year, while the sports series holds near 4%. The two types are close at the median share of Table 2, and the gap appears among those who continue betting, the margin on which the two types most differ. The age panels reproduce the ordering of the table: the 45+ group sits below every other in both, and the youngest group commits the largest share of its account outflows, close to the 25–34 group, even though it deposits less each month than the 25–34 and 35–44 groups. Whether continuing bettors escalate or whether continuing already selects the more intense among them is not something the figure separates.

4 Empirical strategy

4.1 Treatment definition and estimation sample

The treatment is the individual's first bet. New bettors are those whose first Pix to a betting operator occurs in or after January 2025[6]; never-bettors never appear in the betting flow over the whole period; and individuals already betting before January 2025 are excluded, for the reasons given next. The entry month varies across new bettors, so adoption is staggered; and because the excluded group is precisely the one already betting in 2024, both the treated and the control group are non-bettors throughout the pre-event window by construction. Because the first months of the regime are also where the change in payment observability is concentrated, Appendix B.1 re-estimates the design dropping the January to March 2025 entry cohorts.

Individuals who started betting before January 2025 entered a different market from those who started after. Before regulation, reaching an operator carried the search costs and the counterparty risk of a gray market: no domestic legal entity stood behind the platform,

[6]What we observe is the deposit to the operator, not the wager: the wallet sits on the operator's side (Section 3.1), so a deposit is a necessary step to betting rather than the same event. We write *first bet* for the first deposit throughout and keep *deposit* for magnitudes.

and nothing guaranteed that a balance could be cashed out. Those who bet under those conditions are unlikely to be the same type as those who started after them, since accepting an unauthorized counterparty with no recourse plausibly selects on a tolerance for risk that we neither observe nor can condition on. From January 2025 the same act takes place inside a common institutional frame: as we described in Section 2, the operator holds an SPA licence and is incorporated locally under its own tax registry, the bettor is identified and verified at registration (through biometrics), funding and cash-out run through an account in the bettor's own name, and casino games run certified algorithms with a minimum return to player and disclose winning probabilities before the wager. Access and exposure changed as well, with betting apps reaching the official stores from mid-2025 and a nationwide advertising push on open television (Tunad, 2025; Lance!, 2025).

Betting spread accordingly – the share of individuals who bet in a given month roughly doubles between January 2024 and December 2025 (Figure 1, panel (b)) – so restricting the treated group to 2025 entrants is also where the mass of entry lies. What it takes to place a first bet is therefore close to uniform across the treated individuals we compare, which matters in a staggered design where cohorts serve as each other's controls: it is the same treatment entered at different dates, not a different activity for each cohort. The frame is common to the comparison group as well – never-bettors face the same licensing, the same advertising and the same availability – so the regulation shifts the environment of both groups and is absorbed by the calendar-month effects, leaving the individual's own onset as the only variation the design exploits. Still, enforcement is incomplete (Section 2) and part of our treated group deposits with unauthorized operators, which is why we do not restrict the operator list to the SPA register (Section 3.2). What the regime changed is the environment the modal new bettor met, not every transaction in it.

Sample. Table 3 describes the population the design speaks to: 338,263 new bettors and 3,689,513 never-bettors, the individuals from whom the estimation samples of Section 4.2 are drawn. What the table shows is what motivates the selection reading. Before treatment, new bettors were already smaller and in a weaker credit position than never-bettors, with a credit portfolio and limit well below and a higher share of high-cost credit; they are also somewhat younger and less often women. The difference between the groups thus begins before treatment, which is why Section 4.3.1 reads the estimates through a signed bound rather than through a naive comparison of levels.

Estimating on the whole universe is not computationally feasible, so the estimates run on a random draw from it: a deterministic hash of the taxpayer number targeting 100,000 new bettors and never-bettors at a five-to-one ratio, which makes the draw reproducible and

Table 3: Pre-treatment characteristics of the eligible universe

	Never	*New bettor*	*Norm. diff.*
Pre-treatment means (2024)			
Credit portfolio (R$)	70,490	26,741	−0.19
Credit limit (R$)	84,295	31,598	−0.28
Share of high-cost credit	0.268	0.362	+0.27
No. of credit operations	9.3	8.7	−0.01
No. of financial institutions	3.4	3.2	−0.06
Age (years)	41.0	38.9	—
Female share (%)	55.2	50.4	—
Individuals	3,689,513	338,263	

Notes: The eligible universe is the set of individuals who could have started betting through the institution and who appear in the credit registry over the window. It does not impose the covariate requirement that the conditional specification adds on top of it, which is what removes 23.4% of the treated and 6.3% of the controls (Section 4.1). New bettors are those whose first bet falls in or after January 2025 and never-bettors those who never appear in the betting flow; individuals already betting before January 2025 (518,823 of those visible here, against the 712,778 of Section 3.3, which counts on the wider active-account population) belong to neither group and are excluded from the design. Means are taken over 2024, before any new bettor has entered. The last column is the standardized mean difference between new bettors and never-bettors, on the five pre-treatment credit covariates, with portfolio and limit in logs; Imbens and Rubin (2015) treat values above 0.25 in absolute value as the range in which regression adjustment becomes unreliable, and the limit and the high-cost share sit at or beyond it. Age and gender are not covariates, so they carry no standardized difference. For computational cost the estimates are run on a random draw from this universe rather than on all of it, described in Section 4.2.

independent of any outcome. The draw imposes no condition on account activity, so it also reaches 763,328 never-bettors whose account had closed before the window; Appendix B.3 re-estimates the credit effects requiring an active account before entry and finds them somewhat larger. The conditional specification then requires the five pre-treatment covariates. This filter removes 23.4% of the treated against 6.3% of the controls, and therefore the conditional estimates rest on a population somewhat more established than the treated group as a whole, and the counts by specification are reported with each set of estimates. Appendix B collects the exercises that test these sample choices one at a time.

A caveat on the comparison group follows from the same measurement limitation. A never-bettor is someone we never observe sending Pix to an identified operator; such an individual could still bet through a channel we do not capture: a card or bank wire, an account held at another institution, or a recipient our procedure misses. This would leave some latent bettors in the comparison group and, like the other sources of measurement error discussed above, attenuate the estimated contrast toward zero rather than inflate it.

4.2 Estimation

The main design. Our goal is to estimate how an individual's financial position evolves around the month of their first bet, relative to the trajectory they would have followed had they not started. Let G_i be the month of i's first bet – the entry cohort, with $G_i = \infty$ for never-bettors – and let $k = t - G_i$ count the months elapsed since entry. Consider the following event-study model:

$$Y_{it} \;=\; \alpha_i \;+\; \lambda_t \;+\; \sum_{k \neq -1} \beta_k \, \mathbf{1}\{\, t - G_i = k \,\} \;+\; \varepsilon_{it}, \tag{1}$$

in which α_i and λ_t are individual and calendar-month fixed effects, and β_k measures the gap that opens at horizon k, normalized to the month before entry.

Four financial outcomes enter Y_{it}: i) a delinquency flag, equal to one in months with a balance overdue by ninety days or more in any modality in the financial system; ii) the credit-card delinquency rate, the delinquent balance over the portfolio in the card modalities; iii) the log of the total credit limit in the system, a measure of credit made available to the individual; and iv) the monthly change in the end-of-month account balance, a proxy for liquidity or savings, on the asinh scale, which accommodates the zeros and negative values a log cannot.

The estimand is therefore the effect of starting to bet on these outcomes for those who start and conditional on starting in or after January 2025. Importantly, the regulation is not the treatment. We do not estimate a "no regulation" counterfactual, and every number we report is an effect of starting to bet rather than an average effect of the regulation on the population.

Estimator. Entry is staggered and the effect may differ across cohorts and horizons, so the two-way fixed-effects estimator of (1) does not recover the estimand: already-treated individuals serve as controls for later entrants and enter with negative weights (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). We therefore target the group-time average treatment effect directly. Writing $Y_{it}(g)$ for the outcome individual i would record in month t having started betting in month g, $Y_{it}(\infty)$ for the outcome under never betting, and T for the last month of the panel, the parameter is

$$\mathrm{ATT}(g,k) \;=\; \mathbb{E}\Big[\, Y_{i,g+k}(g) - Y_{i,g+k}(\infty) \,\Big|\, G_i = g \,\Big], \tag{2}$$

the effect at horizon k on the cohort entering in g, the disaggregated counterpart of β_k. The event studies we report are the profile of $\mathrm{ATT}(k)$ over k, and the summary effect aggregates it in two layers, over cohorts within a horizon of K periods and then over horizons:

$$\theta^{\text{att}} = \frac{1}{K+1}\sum_{k=0}^{K} \text{ATT}(k), \qquad \text{ATT}(k) = \sum_{g} \Pr\Big(G_i = g \mid G_i + k \leq T\Big)\,\text{ATT}(g,k). \tag{3}$$

Each $\text{ATT}(g,k)$ is estimated by the doubly robust estimator of Callaway and Sant'Anna (2021), which applies the doubly robust estimating equation of Sant'Anna and Zhao (2020) cohort by cohort, so that consistency requires either the outcome regression or the propensity score to be correctly specified, but not both (for a survey of the recent literature with heterogeneous timing, see Roth et al., 2023).

Our preferred specification uses the not-yet-treated as the comparison group and conditions on five pre-treatment covariates over the pre-treatment window: credit portfolio, credit limit, share of high-cost credit, number of credit operations, and number of financial institutions. We also report the unconditional specification, and the never-treated comparison group as robustness. The event-study profiles we plot and the average effects we tabulate are the plug-in aggregations of (3), with the cohort weights taken from their sample frequencies.

An auxiliary design, from delinquency to betting. The design set out above carries the results of the paper, all of which run from betting to the individual's financial position. A second and subordinate estimation runs the opposite way, on the same panel, exchanging the roles of treatment and outcome: if deteriorating finances push individuals toward gambling, realized distress should itself precede and predict the onset of betting. The treatment becomes entry into delinquency, staggered by individual, and the outcome becomes betting itself. Writing G_i^d for the month of i's first delinquency and B_{it} for a betting outcome – an indicator of depositing with an operator in the month, the cumulative indicator of having started, or the asinh of the amount deposited – the estimand is

$$\text{ATT}^d(g,k) \;=\; \mathbb{E}\Big[\, B_{i,g+k}(g) - B_{i,g+k}(\infty) \,\Big|\, G_i^d = g \,\Big]. \tag{4}$$

This is the design behind Section 5.3, and it is estimated with the same CS-DR framework, the same clustering and the same aggregation as (2). Because a common third factor can push at once toward delinquency and toward betting, what it identifies is that realized distress precedes the bet, not that delinquency causes betting. Its role is to date the selection that the main design nets out through the placebo of Section 4.3.1 rather than to qualify the effects themselves.

Simultaneous confidence bands are obtained by multiplier bootstrap, with clusters at the individual level. Because each results table reports a family of outcomes estimated on the same individuals, we also adjust for multiple testing within family, and report the Bonferroni correction.

4.3 Identification

Two assumptions identify $\mathrm{ATT}(g,k)$. The first is no anticipation: outcomes do not move in advance of a bet not yet placed, so that $Y_{it}(g) = Y_{it}(\infty)$ for $t < g$. It is credible for the three credit outcomes, on distinct grounds. The limit is set by lenders, who cannot react to a bet that has not happened. The delinquency flag and the card rate record a balance already overdue by ninety days, so they can only register a shortfall that has occurred. It is less so for the account balance, because individuals might move money in to fund the first bet, producing a visible spike in the month before entry; for that outcome we allow one month of anticipation and take $k = -2$ as the reference period (Section 5.1).

The second is parallel trends: absent the first bet, new bettors would have followed the same trajectory as individuals not yet betting. Our preferred specification does not impose it unconditionally. A difference in levels does not by itself violate it, since the individual fixed effects absorb it; what does is a level that predicts a trajectory. New bettors start from a weaker credit position than never-bettors (Table 3), so we hold the pre-treatment credit position fixed and read the assumption conditional on it.

Conditioning removes selection on the pre-treatment level, but not selection on the path: if the individual whose finances are already deteriorating is the one more likely to start betting, the conditional assumption fails as well. We therefore measure how far it fails and bound what that failure implies for the estimates. Estimating (4) requires the same pair of assumptions with the roles of treatment and outcome exchanged. No anticipation there means that betting initiation must not run ahead of delinquency onset, which is precisely the movement (2) estimates. Each direction is therefore read while treating the other's timing as given, and neither is evidence on the other's dating.

4.3.1 Bounding violations of parallel trends

Group placebo. We first run a placebo on a false event. We restrict the panel to the pre-treatment window (January to December 2024), assign to every future new bettor a false entry in June 2024 – a date by which, by construction, neither group had yet placed a bet – and re-estimate the same CS-DR specification against the never-bettors. A placebo away from zero is the divergence already under way, which the estimate of interest inherits as a bias. We report it under both the unconditional and the conditional specification, the latter on the same covariate set as the corresponding effect.

Writing $\hat{\pi}$ for the placebo, the debiased estimate simply removes it,

$$\hat{\theta}^{\text{db}} = \hat{\theta}^{\text{att}} - \hat{\pi}, \tag{5}$$

carrying the divergence measured before entry into the post-entry window at the same average magnitude. What the placebo carries is the *sign* of $\hat{\pi}$ relative to the effect rather than its exact magnitude, and it carries it as a diagnostic on a fabricated event date rather than as a statement about the pre-treatment window of the design itself. When the placebo runs opposite to the estimated effect, the bias pushes against the finding: the observed effect understates the true one in magnitude. When the placebo runs in the same direction, the bias works in favor of the finding, which must accordingly be discounted and reported as a net or bounded quantity rather than as a clean point estimate. This classification is the *signed reading*, and it is what the last columns of Table 4 report.

Sensitivity analysis. We then follow the approach of Rambachan and Roth (2023), which replaces parallel trends with a bound on how large the post-treatment violation may be relative to the largest one visible before entry, and returns a robust confidence set for the average post-treatment effect at each admissible size $\bar{M}$. We summarize the exercise by the breakdown $\bar{M}$: the largest multiple at which that set still excludes zero, so that a breakdown of 0.5 means the finding survives a violation half the size of the worst one seen before entry.

The restriction as originally stated is symmetric – it admits a bias of either sign – whereas the placebo reading is signed, and the two part ways for the outcomes whose placebo runs opposite to the effect. We therefore also report the variant that intersects the relative-magnitudes set with the direction of the observed pre-trend. Because that variant fixes the direction of the bias while the restriction anchors its admissible size on the design's own pre-treatment window, the two are coherent only where that window points the same way, and the symmetric restriction governs otherwise: a direction the window does not display cannot be imposed on it. Of the outcomes we report, the change in the account balance meets the condition and the credit limit does not. Appendix C states both restrictions formally and reports the exercise in full; the breakdowns appear with each set of estimates.

5 Results

The main result of the paper is what starting to bet does to the individual's financial position, and what follows is ordered by it. Section 5.1 reports the average effect on delinquency, on the credit limit and on the account balance, and Section 5.2 where that effect concentrates. Section 5.3 then reverses the design and estimates the arm that runs the other way, from realized distress to the onset of betting. That arm is complementary rather than a headline: the placebo of Section 4.3.1 already removes the average selection from the delinquency estimates, and what the reversal adds is its timing.

5.1 Effects of betting on credit and liquidity

Table 4 reports the average effect on the three credit outcomes and on the change in the account balance. The conditional column carries our main results, and the four point to the same deterioration: starting to bet raises the probability of being delinquent by 3.7 pp and the credit card delinquency rate by 4.4 pp, cuts the credit limit by 17.4 log points, and draws the monthly change in the account balance down. In proportional terms, and measuring the two delinquency effects against the mean of the treated in the months before their own entry, that is +19.7% on the flag (on a base of 18.5%), +38.7% on the card rate (on a base of 11.4%), about −16% on the limit and about −43% on the change in the balance. All four are significant at 0.1% and survive a Bonferroni correction within the family. The placebo column and the signed reading of Section 4.3.1 separate how much of each magnitude can be read as an effect of betting, and for the two delinquency measures they revise the figures down.

In the unconditional specification the limit rises, by about +22%. That is not an effect of betting. The unconditional group placebo differs from zero in every outcome, so the design does not hold without the covariates: before starting to bet, new bettors are individuals whose credit is expanding, and the unconditional specification attributes that expansion to entry. Conditioning on the pre-treatment covariates inverts the sign of the limit to a cut, because the covariates absorb the expansion, but the conditional placebo does not vanish either.[7]

The limit has a placebo of opposite sign to the estimated effect, which taken alone would make the conditional estimate a conservative lower bound and would deepen the cut from about −16% to about −22%. The pre-treatment window of the design runs the other way, so we do not read the cut as a floor and report −16% as the estimate. The conditional event study is consistent with it, though the pre window is not flat: the limit drifts down over the four months before entry, reaching −0.030 at $k = -1$, and the profile then steepens, to

[7]The choice of comparison group does not drive these estimates. Replacing the not-yet-treated with the never-bettors in the unconditional specification returns 0.011 against 0.010 for the delinquency flag, 0.024 against 0.022 for the card rate and 0.134 against 0.196 for the limit; all three remain significant at 0.1%, and the pre-entry expansion of the limit is somewhat smaller once the future-treated are removed from the comparison group. The same exercise on the change in the account balance is reported with that outcome below.

Table 4: Average effect on financial outcomes: unconditional, conditional, and signed bound

	CS-DR			Debiased	Breakdown $\bar{M}$	
Outcome	Uncond.	Cond.	Cond. placebo	estimate	Sym.	Signed
Delinquency (flag)	0.010***	0.037***	0.011***	0.026	0.25	0.25
	(0.002)	(0.003)	(0.002)			
Card delinquency rate	0.022***	0.044***	0.004**	0.040	0.75	0.75
	(0.002)	(0.002)	(0.002)			
Log limit	0.196***	−0.174***	0.069***	−0.243	0.50	> 2
	(0.011)	(0.013)	(0.011)			
Change in balance†	−0.605***	−0.569***	0.091***	−0.661	0.25	> 2
	(0.034)	(0.034)	(0.019)			
Clusters (individuals)	599,530	268,233				
New bettors	97,287	37,619				
Never-bettors	502,243	230,614				
Observations	7,899,529	5,669,676				

Note: Average effect on the treated θ^{att} of (3), dynamically aggregated (Section 4), from the CS-DR estimator of Callaway and Sant'Anna (2021) with a not-yet-treated comparison, over 28 months (January 2024–April 2026). Standard errors clustered at the individual level in parentheses; every raw p is below 0.001 and survives Bonferroni within the family. The panel is unbalanced – an individual enters the credit registry only in months with a reported operation – so an individual contributes on average 13.2 of the 28 months in the unconditional sample and 21.1 in the conditional one. The unconditional column uses the whole panel; the conditional column, the placebo and the debiased estimate require the five pre-treatment covariates and rest on the 268,233 individuals that have all five, a selective restriction discussed in Section 4.1. For the card delinquency rate, defined only for cardholders, the clusters are 490,663 and 242,163 and the observations 5,709,636 and 4,496,905. The placebo is the false-event (June 2024) group placebo and the debiased estimate is the conditional effect minus it, equation (5). The last two columns give the breakdown $\bar{M}$ of the Rambachan and Roth (2023) sensitivity analysis under the symmetric (8) and sign (9) restrictions (Section 4.3.1); the grid runs to 2, so > 2 denotes survival of the whole grid, and the underlying post effects with original and robust intervals are in Appendix Table C1. † The change in the account balance is on the asinh scale, so the proportional reading holds where the monthly change is positive and describes a deepening of the same order where it is already negative; it is estimated on the 29-month active-account panel with one month of anticipation (described with that outcome in the text), on a random subsample of 30,000 new bettors and 150,000 never-bettors. Its two breakdown columns are computed on the *unconditional* event study of that panel rather than on the conditional column of this row, which is what was estimated for the exercise (Appendix Table C1). * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

−0.050 by the first month after entry and −0.308 by the twelfth (Figure 5). The drift is of the same sign as the effect and is the violation the sensitivity analysis of Section 4.3.1 is anchored on. Part of it may not be a violation at all: an individual whose first bet was funded from an account we do not observe has an entry date recorded later than the true one, so some of the months we treat as pre-treatment are already post-treatment. We anchor the sensitivity analysis on the drift as measured, which is the conservative reading of it.

Both delinquency measures rise, but with a placebo of the same sign, so we report them net of it: +2.6 pp in the flag and +4.0 pp in the rate, about +14% and +35% against the pre-entry rates of the treated, 18.5% and 11.4%.[8] They are the only two outcomes whose

[8]The event study corroborates the size of the placebo by an independent route. A missed payment is recorded only after ninety days, so the first three horizons cannot register a default caused by a bet placed at $k = 0$, and what they measure is deterioration already under way; they average +1.22 pp in the flag, against the

placebo works in favor of the finding, and part of what remains is selection of an individual who was already in distress, the reverse channel we date in Section 5.3.

The sensitivity analysis formalizes the signed reading. For the same-sign outcomes the symmetric restriction is the appropriate approach: the card delinquency rate is the more robust of the two, surviving a post-treatment violation three-quarters the size of the largest pre-treatment one, while the delinquency flag is the less robust and is the finding most dependent on a well-behaved pre-trend. For the limit the two diagnostics disagree in sign: the false-event placebo runs against the effect, while the design's own pre-treatment window runs with it, its largest violation being -0.013 in first differences. The sensitivity analysis is anchored on that window, so the symmetric restriction is the reading we publish, and the cut survives a post-treatment violation half the size of the largest pre-treatment one. The sign-restricted variant returns the whole grid for this outcome, but only by imposing the direction the window contradicts; we report it in Appendix C without reading it.

Our delinquency effect is smaller than what the recent U.S. estimates imply for bettors. Those studies are identified off state legalization, so they recover an intent-to-treat effect on the whole population and reach the per-bettor object by dividing by take-up: Goss and Mangrum (2026) scale a rise of 0.31 pp on a base of 10.71% by a quarterly take-up of 3.1 pp and report an implied 10 pp among induced bettors, and Hollenbeck et al. (2025) argue that with adoption of 13–20% the effect on actual bettors is five to ten times their estimates. We estimate the effect on the individuals who actually start betting, without passing through take-up: +3.7 pp, or +2.6 pp after the debiasing, on a pre-entry base of 18.5%. The two are not one biased version of the other; they describe different populations, and Section 5.3 below says how they differ. Because entering delinquency itself raises the probability of starting to bet, whoever a change in access induces is drawn disproportionately from individuals already deteriorating, whereas our estimate averages over everyone who entered. Two features run the other way. We observe betting only as Pix sent from the observed account to identified operators, so individuals who bet through a channel we do not capture sit in the comparison group (Section 4.1), which attenuates the contrast; and our window ends twelve months after entry, whereas the U.S. profiles are still rising three years out.

The credit card is where our results and that evidence part. Hollenbeck et al. (2025) find no effect on card delinquency and read it as the flexibility of the instrument: the card is what a household draws on in order not to miss a payment elsewhere, so distress surfaces in other products instead. That reading presumes the credit limit stays where it was, and in our data

+1.1 pp of the false-event placebo. Restricted to $k \geq 3$, the average effect is +4.4 pp in the flag and +5.3 pp in the card rate, against the +3.7 and +4.4 over all thirteen horizons.

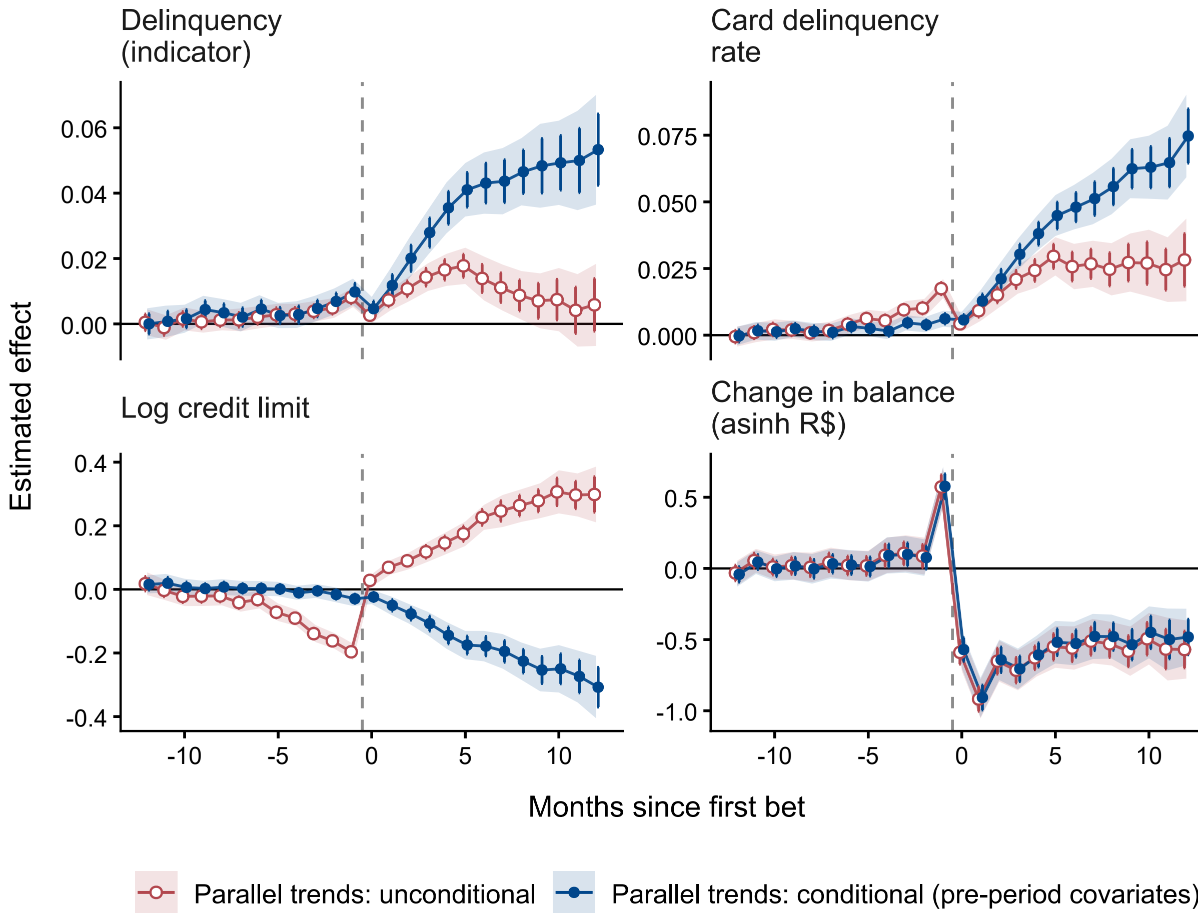

Figure 5: Event study of financial outcomes.

Notes: Event-study profile (3) by month relative to the first bet, for the three credit outcomes and for the monthly change in the account balance, under unconditional and conditional parallel trends; samples, estimator and standard errors as in Table 4. Thin bars: pointwise 95% confidence intervals; wide, light bars: simultaneous bands. The fourth panel rests on a different panel and a different specification, and is placed alongside the others for the direction of the trajectory rather than for a comparison of levels; the positive spike it shows in the month before entry is the funding of the first deposit, a month that the one-month anticipation places inside the treated window and therefore outside the average reported for that outcome. The conditional pre window is what anchors the sensitivity analysis of Section 4.3.1: the breakdown $\bar{M}$ scales the admissible post-treatment violation by the largest violation visible here, and is reported by outcome in Table 4 and in full in Appendix Table C1 and Figure C1.

it does not – the limit is cut as soon as the betting starts, and the card is where the break then shows up.

The rise in delinquency has a reading that requires nothing of the lender: money sent to an operator is money not available for debt service. For the tenth of bettors who commit more than a fifth of their account outflows to operators (Table 2, Panel B), a diversion of that order can turn a payment that would have been made into one that is missed.

We interpret the limit reduction as the banking system's response to an individual who has started betting, and a potential mechanism for it starts from what the lender can see. The limit we measure is the total across the financial system, and the institution that cuts it need not be the one whose account the deposit left, so the transfer itself is not generally visible to it. However, this behavior might become visible to the lender through other means, such as the behavioral markers commercial bureaus sell or because the individual is also betting from an account held at the lender. Whichever route it takes, the fact that the individual has started betting enters the lender's screening models as new information about the borrower's type, and the limit falls in the month it does – the response asymmetric-information models predict of a lender that observes a new behavioral marker (Stiglitz and Weiss, 1981; Agarwal et al., 2009; Gao et al., 2023).

The delinquency effect can then compound the cut. Once recorded, a delinquency is visible to every lender in the system, so the deterioration that follows entry may feed further reduction over the horizon: the two effects need not be independent, and the limit profile deepens as the months pass (Figure 5).

In addition to credit, the account balance itself gives a real-side reading of financial distress. The deposits come out of the account (Section 4), and what happens to the balance that remains is the fourth outcome of Table 4. For the share of this population that carries out an entrepreneurial activity, the end-of-month balance is not only a household cushion, it is also that activity's accumulated saving and working-capital buffer, so a persistent drawdown erodes the margin that absorbs revenue shocks.

We estimate the same CS-DR design on the monthly change in the balance, on the asinh scale, which accommodates the zeros and negative values that a log cannot. Two features of the specification follow from what the outcome measures. First, the change in the balance is informative only for accounts that move money, so we restrict the sample to the individuals whose account was already active in the pre-treatment window; a dormant account has no buffer to draw down, and including it adds zeros that dilute the contrast rather than measuring it. Second, the month before entry carries a spike in funding – the deposit leg of the mechanism, as the individual moves money in before betting – so we allow one month of anticipation, which keeps that spike out of the pre-treatment block.

Starting to bet draws the balance down. In reais the monthly change falls by about R$81 (Appendix B.4), which cumulates to roughly R$1,050 over the thirteen months we observe; on the asinh scale of Table 4 the same estimate is -0.57 in the conditional specification, a proportional reading of about -43% that holds where the monthly change is positive and describes a deepening of the same order where it is already negative. Its group placebo is of opposite sign to the effect, and here the design's own pre-treatment window agrees with the placebo on direction, which it does not for the limit; the estimate is therefore a conservative floor: removing it deepens the drawdown to about -48%, and under the sign-restricted sensitivity analysis the estimate survives the entire $\bar{M}$ grid. The comparison group does not drive it either: in the unconditional specification, taking the never-bettors instead of the not-yet-treated returns -0.594 against -0.605. The drawdown moves in the same direction as the credit contraction (last panel of Figure 5), and it is where the shock reaches liquidity and accumulated saving rather than the credit relationship alone.

The drawdown has a counterpart in this literature. Baker et al. (2024) find that betting does not reduce total household spending but crowds out saving almost one for one, and Muggleton et al. (2021) report the same direction in the United Kingdom, with heavier gamblers holding smaller savings and making smaller pension contributions. Marionneau et al. (2024) find the association on the population closest to ours – individuals already carrying debt, observed through their bank transactions – though cross-sectionally, so part of what it measures is the distress that precedes the bet (Section 5.3). In our setting the buffer being displaced is the transactional balance itself, so the displacement runs on a different margin: what is given up is not a long-horizon return, as in a brokerage balance, but the liquidity that absorbs a bad month of revenue.

5.2 Heterogeneity

The average effects can be partitioned in many ways, and we report three that the descriptive evidence and the existing literature point to. The prior credit position is the split on which that literature has most to say, since the deterioration it documents is concentrated among borrowers who were already constrained. Age is where our own statistics separate bettors most: persistence hardly varies across age groups, while deposits as a share of account outflows fall steeply with it (Section 3.3), so the same betting behavior weighs unequally on the account. And the product type separates two profiles that those statistics already distinguish, the casino one closer to repeated play with little realized return. The first two are predetermined; the third is fixed at entry ($t = 0$) rather than before it, and does not inherit the identification of the pooled design in the same way.

The three are built in two different ways, and the standard error of the contrast depends on which. The credit-position split partitions the estimation sample at a median, comparison group included, so that each half is identified against controls from the same half and the two halves need not bracket the pooled estimate; they share no units, and the variance of the difference is the sum of the two variances,

$$\Delta = \hat{\theta}^{\text{att}}_{\text{low}} - \hat{\theta}^{\text{att}}_{\text{high}}, \qquad \widehat{\text{se}}(\Delta) = \sqrt{\widehat{\text{se}}^2_{\text{low}} + \widehat{\text{se}}^2_{\text{high}}}. \tag{6}$$

The age and product-type splits instead identify each group against the same pool of never-bettors. Their estimates are therefore correlated, and the standard error of the difference between two of them,

$$\Gamma = \hat{\theta}^{\text{att}}_a - \hat{\theta}^{\text{att}}_b, \tag{7}$$

comes from a joint bootstrap over the shared controls.

5.2.1 Prior credit position

We partition the sample at the median of the pre-treatment credit limit, a proxy for the borrower's prior credit position, and re-estimate the conditional CS-DR effect θ^{att} of (3) within each subgroup. The gradient is the difference Δ of (6).

We find that the shock is regressive, and the gradient is in the amount of credit supplied (Table 5). The thin-credit individual takes about 2.6 times more limit cut in log points (-0.193 against -0.075, a difference significant at $p < 0.001$), or -17.6% against -7.2% in proportional terms (Figure 6).[9] The probability of becoming delinquent rises equally in the two subgroups (a null difference, $p = 0.92$), and the card delinquency rate differs by 0.9 pp against levels of 4.6 and 3.7, significant only at the 5% level. The balance drawdown runs the other way: the thick-credit half draws the balance down more than the thin-credit one (-0.592 against -0.356, $p = 0.011$), and the ordering is not an artifact of the base being drawn down, which is similar in the two halves before entry (R\$50 a month in the thick-credit half against R\$58 in the thin-credit one) and puts the implied drawdown at about R\$22 a month against R\$18. What the lender does is graded by the prior credit position; the erosion of the buffer is not, and if anything falls more heavily on the half that holds more of it.

The split at the median of the prior limit exposes the magnitude gradient to some degree of reversion to the mean. The prior limit is also correlated with income, since income is one of the inputs to how it is set, but it embeds the credit relationship itself – tenure with lenders, product mix, and the lender's own risk model – and these data cannot separate the two. We read the gradient as one of the borrower's standing with lenders, of which income is a component.

The concentration among borrowers with the weakest prior credit position is a robust regularity in this literature, and our estimate of its size is close to the existing ones. Hollenbeck et al. (2025) find the deterioration "overwhelmingly concentrated" among pre-legalization subprime consumers (an 11-point credit-score fall, +12% collections, +12% bankruptcy likelihood, against small improvements for the super-prime), and Baker et al. (2024) find

[9]The gradient is estimated holding the five pre-treatment covariates of the main design fixed inside each half, the prior limit among them. Dropping the prior limit from the covariate set within the halves, on the grounds that the split already conditions on it, attenuates the gradient to 1.5 times (-0.165 against -0.110, $p = 0.07$) and leaves the two delinquency gradients indistinguishable from zero. We report the specification whose covariate set matches the pooled estimate of Table 4, and read the size of the gradient as specification-dependent while its direction is not.

Table 5: Heterogeneity by prior credit position

Outcome	Thin credit	Thick credit	Difference	$\bar{M}$ (Thin / Thick) Sym.	Signed
Log limit	-0.193^{***}	-0.075^{***}	-0.118^{***}	0.25 / 0.00	> 2 / > 2
	(0.028)	(0.015)	(0.031)		
Card delinquency rate	0.046^{***}	0.037^{***}	0.009^{*}	0.25 / 0.50	0.25 / 0.50
	(0.003)	(0.003)	(0.004)		
Delinquency (flag)	0.036^{***}	0.036^{***}	0.001	0.25 / 0.25	0.25 / 0.25
	(0.004)	(0.004)	(0.005)		
Change in balance	-0.356^{***}	-0.592^{***}	0.237^{*}	—	—
	(0.048)	(0.080)	(0.093)		

Note: Subgroup effects θ^{att} of (3), estimated within each half of the conditional credit panel of Table 4 split at the median pre-treatment limit, with *thin credit* the half below the median and *thick credit* the half above; estimator, sample, aggregation and standard errors as in that table. The change in balance is on the asinh scale and on the 29-month active-account panel described in the note to it. The *Difference* column is the gap Δ of (6), whose standard error follows from the subgroups being disjoint. The last two columns give the breakdown $\bar{M}$ of the Rambachan and Roth (2023) sensitivity analysis under the symmetric (8) and sign (9) restrictions (Section 4.3.1), run within each subgroup (Thin / Thick) and anchored on its own pre-treatment window; the two coincide for the delinquency outcomes, and the symmetric one is what we read for the limit, whose pre-treatment window runs with the effect and therefore does not warrant the direction the sign restriction imposes. Under it the thin-credit leg survives a post-treatment violation a quarter the size of the largest pre-treatment one, while the thick-credit leg does not exclude zero even at $\bar{M} = 0$; the gradient between the two rests on the z-test of (6) and not on this exercise. The exercise is not run for the change in balance. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

low-savings households cutting net investment by 41% of the mean against 14% overall, roughly a factor of three. Our thin-to-thick ratio in the limit cut sits in the same range. What differs is where the gradient lives. There it is in realized distress – the constrained borrower is the one who defaults, is collected on, or files. Here it is on the lender's side: the limit cut is graded by the prior position while the delinquency flag rises by the same amount in both halves. The regressivity is thus mainly one of credit supply rather than of realized risk, and it is the screening response of Section 5.1 applied most sharply where the prior position leaves least margin.

5.2.2 Age group

Age is recorded at registration of the banking account, and is fixed before the first bet, so partitioning by it raises none of the concerns that attach to a split on a post-treatment variable. We form four groups (18–24, 25–34, 35–44, and 45+) and estimate the conditional CS-DR effect of each against the same pool of never-bettors, so the gap between the extreme groups is the difference Γ of (7). Because four groups trace a profile rather than a contrast, we also test its monotonicity with an inverse-variance weighted regression of the group effects on the midpoint age of each, and report the p-value of the slope.

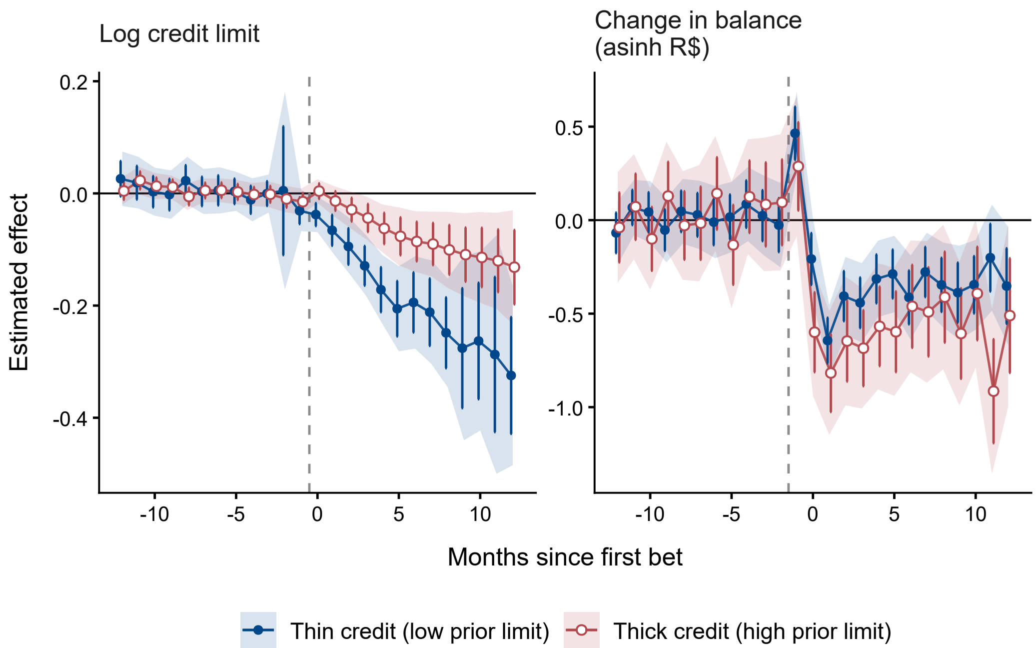

Figure 6: Event study by prior credit position: the credit limit and the account balance.

Notes: Event-study profile (3) (conditional) by month relative to the first bet, for the log credit limit (left) and the monthly change in the account balance (right), with one series per half of the split at the median pre-period credit limit (thin credit: below; thick credit: above). Standard errors clustered at the individual level; thin bars are pointwise 95% confidence intervals, wide bars are simultaneous bands. The two panels rest on different samples and specifications, as in Table 5: the limit on the 28-month credit panel (January 2024–April 2026), the balance on the 29-month active-account panel with one month of anticipation, which is why the dashed line marking the start of the treated period sits one month earlier in the right panel. The month immediately to the right of that line is therefore already treated, and it carries the funding of the first deposit: the balance rises in that month in both halves and falls thereafter, and that month is outside the average reported in Table 5. The comparison is of the trajectory within each panel and of the ordering between the two halves, not of levels across panels.

The age gradient is sharp, and among the credit outcomes it is confined to delinquency (Table 6). The youngest group breaks about 2.6 times as often as the oldest: the delinquency flag rises 6.0 pp at 18–24 against 2.3 pp at 45 or older (a gap of 3.7 pp, $p < 0.001$), and the card delinquency rate 6.8 pp against 2.9 pp (a gap of 4.0 pp on the unrounded estimates, $p < 0.001$); both profiles decline monotonically across the four groups, and the slope of each is significant at the same level (Figure 7). The gradient is one of response and not of base: the pre-entry delinquency rate of the treated is almost the same in the two extreme groups (18.2% at 18–24 against 17.2% at 45 or older), so in relative terms the flag rises +33.0% in the youngest group against +13.4% in the oldest. The limit cut, by contrast, is flat.

Importantly, the group where the effects concentrate is under-represented in our sample: only 11% of the individuals we observe transacting with authorized operators are under 25, against 22% of bettors nationally (Section 3), a gap that follows from a population weighted toward working-age individuals with some entrepreneurial activity. Our pooled estimates are therefore averages over an older-than-national mix of bettors; at the national composition the delinquency effects would be larger, while the limit cut, which does not vary with age, would be unchanged.

Table 6: Heterogeneity by age group

Outcome	18–24	25–34	35–44	45+	Gap (young−old)	Trend p
Delinquency (flag), pp	6.0***	4.5***	3.9***	2.3***	3.7***	< 0.001
	(0.91)	(0.43)	(0.48)	(0.44)	(1.01)	
Card delinquency rate, pp	6.8***	5.5***	4.9***	2.9***	4.0***	< 0.001
	(0.91)	(0.40)	(0.40)	(0.36)	(0.94)	
Log limit	−0.188***	−0.168***	−0.183***	−0.144***	−0.044	0.15
	(0.051)	(0.025)	(0.023)	(0.023)	(0.055)	
Change in balance	−0.464***	−0.516***	−0.599***	−0.628***	0.164	< 0.001
	(0.099)	(0.056)	(0.059)	(0.071)	(0.121)	
Clusters, credit	233,151	241,992	241,642	240,579		
Clusters, balance	129,532	135,619	134,793	134,095		

Note: Group effects θ^{att} of (3), each group identified against the same pool of never-bettors; estimator, aggregation, panels and standard errors as in Table 4. The two delinquency rows and their gap are in percentage points, the limit in log points, the change in balance on the asinh scale. *Gap* is the difference Γ of (7) between the 18–24 and the 45+ group. *Trend p* is for the slope of an inverse-variance weighted regression of the four group effects on the midpoint age of each; it tests the monotonicity of the profile, which the gap between the extremes alone does not. The *Clusters* rows count the units entering each group's estimation, that is, the group itself plus the shared control pool. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

The pattern is the mirror image of the credit-position split of Section 5.2.1. What the lender does is graded by the prior credit position and not by age; what happens to the borrower afterward is graded by age and not by the prior limit. Exposure to the credit cut is thus uniform across the life cycle while the realized delinquency it feeds into is concentrated among the young, and the group placebos are small and of similar size in every group, so the gradient is not the residue of a differential pre-trend. The balance drawdown does run monotonically with age, and in the opposite direction, which we read as compositional: an older individual holds a larger balance and therefore has more of it to draw down.

The profile lines up with the only comparable causal evidence. Goss and Mangrum (2026) find the delinquency response to legalization concentrated under 40, weaker at 40–64 and absent above 65, with the under-40 effect roughly three and a half times the 40–64 one; that ratio does not depend on take-up, so it is comparable to our directly measured 2.6, though their cut is under-40 against 40–64 and ours the youngest group against the oldest. Hollenbeck et al. (2025) report the same direction more weakly (effects "directionally strongest for low-income younger men"), and Abreu and Bressan (2026) find age the most stable determinant of Brazilian betting participation and frequency. Designs identified off different shocks, in different credit systems, reproduce the gradient. What our setting adds is the lender's side, which none of them observes: the limit cut does not vary with age, so credit

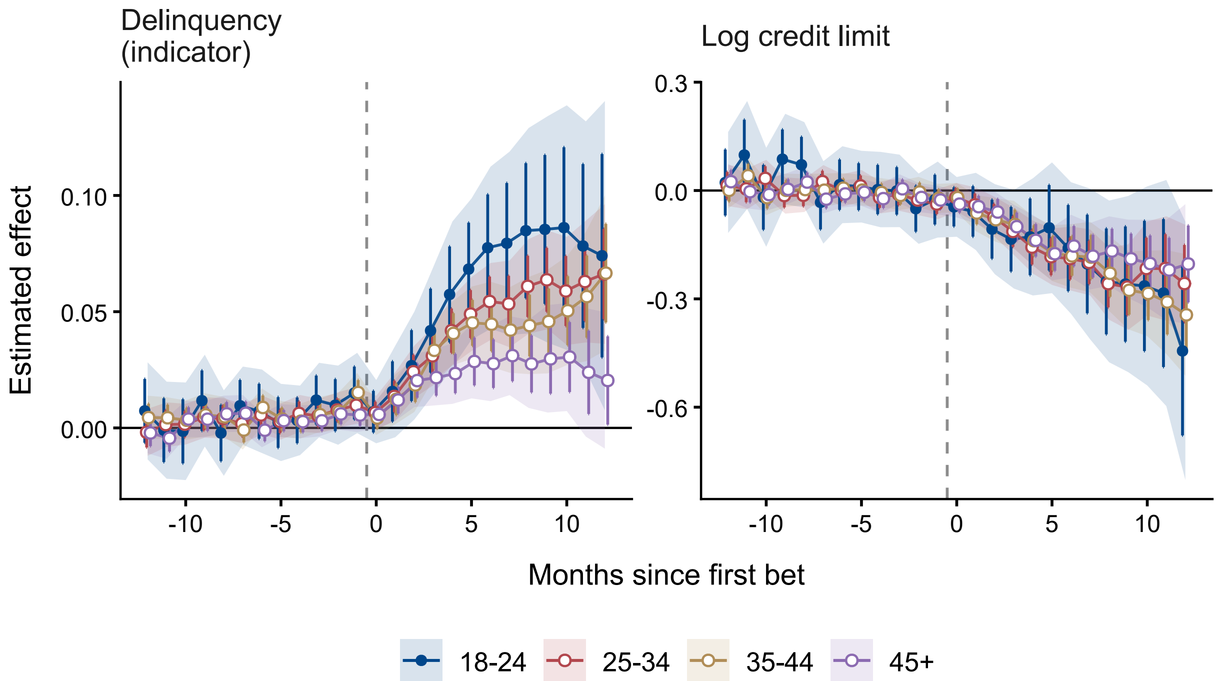

Figure 7: Event study by age group: delinquency and the log credit limit.

Notes: Event-study profile (3) (conditional) by month relative to the first bet, for the delinquency flag (left) and the log credit limit (right), one series per age group; samples, estimator and standard errors as in Table 6. Shaded bands are pointwise 95% confidence intervals. The delinquency profiles fan out by age after entry, while the limit paths lie on top of one another.

supply is graded on the prior credit position (Section 5.2.1) rather than on the dimension where realized risk concentrates.

5.2.3 Product Type: Casino versus Sports Betting

The third split is by product type. Casino games are continuous and high-frequency, sports betting is event-based and lower-frequency, and Section 3.3 shows the two profiles differing in persistence, in the number of platforms used, and in how often money comes back from the operator. We take the classification of operators by product from Section 3.2.

We assign a type to each bettor from the predominant product type of the Pix destinations in the month of entry ($t = 0$). Restricting the window to $t = 0$ is deliberate: because the credit shock materializes from entry onward, a longer window would risk contaminating the type label with the outcome's own reaction. Let s_i be the share of an individual's classified Pix directed to sports operators. To limit attenuation from individual measurement error, we form two arms at the extremes of the s_i distribution, a sports arm ($s_i \geq 0.55$) and a casino arm ($s_i \leq 0.30$), leaving the intermediate range out of focus. The type is stable enough to serve as a fixed characteristic: s_i correlates 0.90 across consecutive quarters and 97.6% of bettors keep the same classification between them.

This creates two disjoint arms, each identified against the same pool of never-bettors by a conditional CS-DR estimator, and the value of interest is the gap Γ of (7) between them.

The only significant difference is in the limit cut (Table 7). In the casino arm the limit falls by about 22%; in the sports arm by about 9%, an estimate at the margin of significance whose interval already includes zero under the sensitivity exercise, which is what the last two columns record as n.s. On the delinquency margins the two types are statistically indistinguishable, and the balance drawdown is of the same magnitude in both arms, so the erosion of saving is common to the two product types and what distinguishes casino is the credit cut alone. The difference grows with event time, so the averages understate it (Figure 8): in the casino arm the cut drops abruptly already at $t = 0$ and deepens steadily without reverting, to about -37% at twelve months, while in the sports arm it is shallow, non-monotone, and does not exclude zero by the end of the horizon.

Two features keep this product type gap suggestive. The game type and the operator's authorization status are strongly correlated – sports operators are almost all authorized, casino concentrates among the unauthorized – so the raw gap mixes type and legality, and isolating the type within the authorized operators, the gap is no longer significant. And the group placebo of the limit has a gap of the same sign, so part of the difference may be a differential pre-trend. What survives is the direction and its robustness: run separately by arm, the sensitivity analysis of Rambachan and Roth (2023) keeps the casino limit cut away from zero at a post-treatment violation half the size of the largest pre-treatment one, while the sports cut never excludes zero. The delinquency gap is null between the types.

This split matters for comparing our findings to the literature, because that evidence is almost entirely about sports betting. Baker et al. (2024), Hollenbeck et al. (2025) and Goss and Mangrum (2026) all date treatment to sportsbook legalization, and Hollenbeck et al. (2025) treat online casino as a confound, re-estimating without the six *iGaming* states to confirm that sports betting drives their results. Our comparison runs the other way, and to whatever extent product type rather than legal status carries the difference, those estimates would be a lower bound for a market like Brazil's, where casino platforms account for the larger share of both operators and bettors (Tables A2 and 2). The shape points the same way: the casino cut deepens steadily over the horizon without reverting, which is what Baker et al. (2024) describe on the betting side when they report net investment falling further with each quarter elapsed and conclude that the substitution is not self-correcting.

Table 7: Heterogeneity by product type

					$\bar{M}$ (C / S)	
Outcome	Casino	Sports	Gap (C−S)	p (gap)	Sym.	Signed
Delinquency (flag), pp	4.9***	4.0***	1.0	0.42	0.25 / 0.00	0.25 / 0.00
	(0.50)	(1.05)	(1.21)			
Card delinquency rate, pp	5.9***	6.3***	−0.4	0.68	0.50 / 0.25	0.50 / 0.25
	(0.46)	(1.00)	(1.09)			
Log limit	−0.248***	−0.097*	−0.151	0.007	0.50 / n.s.	> 2 / n.s.
	(0.025)	(0.049)	(0.056)			
Change in balance	−0.453***	−0.556***	0.103	0.40	—	—
	(0.048)	(0.114)	(0.122)			
Bettors	10,814	2,216				
Never-bettors			230,614			
Clusters (individuals)			243,644			

Note: Arm effects θ^{att} of (3), each arm identified against the same pool of never-bettors; estimator, aggregation, panels and standard errors as in Table 4. The two delinquency rows and their gap are in percentage points, the limit in log points, the change in balance on the asinh scale. *Gap* is the difference Γ of (7) between the casino and the sports arm. The last two columns give the breakdown $\bar{M}$ of the Rambachan and Roth (2023) sensitivity analysis, run for each arm separately (Casino / Sports) under the symmetric (8) and sign (9) restrictions (Section 4.3.1); "n.s." marks an arm whose average post effect does not exclude zero even at $\bar{M} = 0$, which the sign restriction cannot rescue. The post effects and intervals behind these columns are in Appendix Table C2; the exercise is not run for the change in balance. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

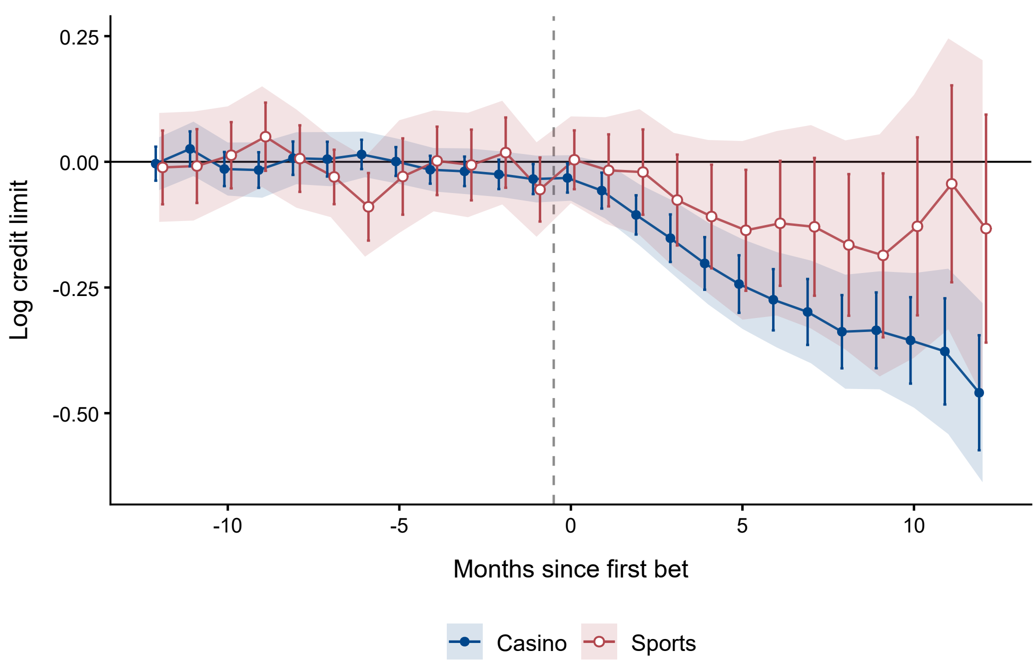

Figure 8: Event study of the log credit limit by product type.

Notes: Event-study profile (3) (conditional) by month relative to the first bet, for the log credit limit, one series per arm; samples, estimator and standard errors as in Table 7. Thin bars are pointwise 95% confidence intervals, wide bars are simultaneous bands.

5.3 Financial distress and starting to bet

The effects reported so far run from the first bet to the individual's financial position. Betting and financial distress can run in both directions, however, and part of the delinquency effect of Section 5.1 is distress that was already there: starting to bet is, in part, a marker of preexisting financial distress. The placebo nets that selection out of the estimates; this subsection dates it. In the pre-treatment period, new bettors were already smaller and in a weaker credit position than never-bettors, with a much lower credit portfolio and limit and a higher share of high-cost credit (Table 3). That is a difference in levels, and it does not date the distress relative to the first bet. The estimates of (4) do: the cumulative probability of starting to bet is 2.6 pp higher after an individual enters delinquency (Table 8). Over the window, 8.4% of the eligible individuals visible in the credit registry started betting, which is the sample this estimate comes from, against 11.6% in the wider universe of active accounts (Table 3, Panel A, and Section 3.3); so 2.6 pp is close to a third of the initiation rate where it is estimated.

The three measures of B_{it} do not behave alike. The monthly probability of betting and the amount deposited with operators move the same way and lose significance under conditioning; the cumulative measure is the one that persists. The event study behind the table shows it rising over the months that follow the first delinquency, so the 2.6 pp averages over the year after entry rather than describing its endpoint: it is the equally weighted mean of thirteen horizons of a cumulative outcome, which matches the level reached at about six months, while the twelfth month reaches 5.3 pp (Figure 9). The conditional specification here uses six pre-treatment covariates, the five of the main design plus age. What this arm does not do is displace the effects of Section 5.1. The selection it dates is the same one the placebo already removes from those estimates, which is why we report the delinquency effects net of it and read this result as the order of the two events rather than as a rival magnitude.

This result relates to the desperation motive of Blalock et al. (2007), where lottery sales across U.S. states rise with the poverty rate, and gives it an individual-level counterpart. Herskowitz (2021) supplies the mechanism: betting is an attempt to raise a sum of money when there is no other route to it, so people bet more when they face an expense they cannot cover, and bet less once a randomized savings device gives them a safe way to put that money together. Our estimate comes at the same mechanism from the opposite and complementary side: rather than a safe route being opened, betting begins at a moment of realized distress, when a debt has already gone unpaid. Survey evidence from Brazil points the same way: in the two most recent waves of a nationally representative survey of Brazilian savers and investors, the strongest predictor of how often an individual bets is the belief that betting

Table 8: Effect of entering delinquency on betting

Outcome	Unconditional	Conditional
Monthly probability of betting (pp)	0.69***	0.43
	(0.06)	(0.74)
Started betting, cumulative (pp)	1.8***†	2.6**
	(0.11)	(0.92)
Amount deposited (asinh)	0.036***	0.022
	(0.003)	(0.046)
Clusters (individuals)	160,000	45,034
Treated (entered delinquency)	40,000	21,403

Note: Estimates of (4), dynamically aggregated, from the CS-DR estimator of Callaway and Sant'Anna (2021) over 28 months (January 2024–April 2026); standard errors clustered at the individual level in parentheses. The pre window of this design does not test parallel trends, and the reason is one of construction: the treated group is restricted to individuals whose first bet falls in or after the month of their first delinquency, so their cumulative initiation is identically zero before entry and the pre coefficient reports the initiation rate of the comparison group rather than a differential trend. Its magnitude is fixed by the base-period convention (Appendix D) and therefore bounds nothing, and for the same reason the Rambachan and Roth (2023) sensitivity analysis of Section 4.3.1 is not run here; what identifies the effect is the timing of initiation relative to the delinquency, and the post-treatment coefficients are invariant to that convention. The conditional specification requires the six pre-treatment covariates and rests on the 45,034 individuals that have all six, which is what widens its standard errors: the requirement binds harder on the comparison group than on the treated, so the three-to-one ratio of the panel falls to about one-to-one.
† The comparison group is the not-yet-delinquent; restricting it to individuals who never become delinquent over the window returns 1.5 pp against 1.8, a difference of 0.24 pp and about an eighth of the estimate. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

supplements income, and betting is more frequent among younger respondents, those with lower incomes, and those reporting financial stress (Abreu and Bressan, 2026).

6 Conclusion

This paper examined the effects of starting to bet on credit and financial distress in Brazil. The 2025 regulation of fixed-odds betting supplies the setting rather than the treatment: it concentrates a large wave of first bets, which gives the design its staggered cohorts and a pre-period in which neither group is betting, and it places the individuals we compare inside a common institutional frame. What identifies the effects is the timing of each individual's own first bet. Starting to bet carries consequences on both sides of the credit relationship and on the account itself. Delinquency rises, the credit limit held across the financial system is cut, and the end-of-month account balance is drawn down. The rise in delinquency needs nothing of the lender: money sent to an operator is money not available for debt service. The limit reduction appears instead to be a screening response, in which having started to bet enters the lender's assessment as new information about the borrower. The two sides

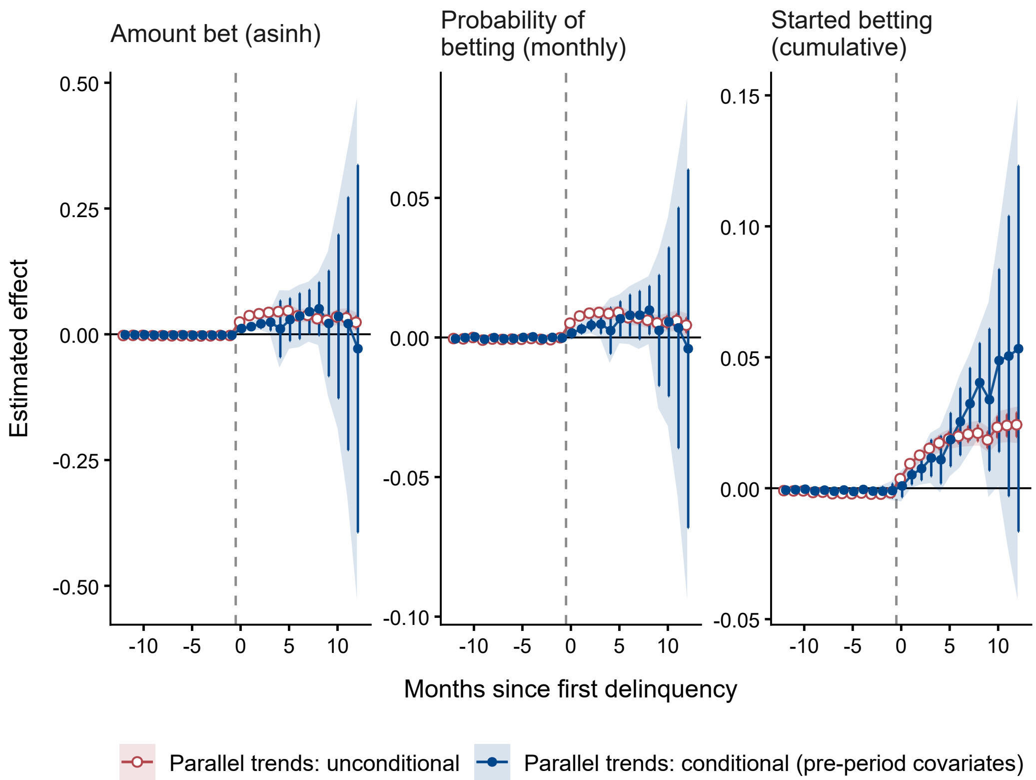

Figure 9: Event study of the delinquency-to-betting design.

Notes: CS-DR estimates of (4) (unconditional and conditional); treatment is entry into delinquency and the outcome is starting to bet, by month relative to the first delinquency. Sample: 160,000 individuals (40,000 treated) over 28 months (January 2024–April 2026) in the unconditional specification and 45,034 (21,403 treated) in the conditional one; standard errors clustered at the individual level. Thin bars are pointwise 95% confidence intervals, wide bars are simultaneous bands.

of the credit relationship are graded differently: the reduction falls hardest on borrowers with less credit before they start and does not vary with age, while delinquency rises by the same amount whatever the prior credit position and concentrates among the youngest. The regressivity is one of credit supply rather than of realized risk. The balance drawdown does not follow that gradient: it is if anything larger among those who had more credit before starting, and it reaches beyond the credit relationship. Part of the rise in delinquency is not a consequence of the bet but distress that came before it: those who start were already in a weaker credit position, and reversing the design shows that entry into delinquency itself precedes and predicts the first bet, which is why we report the delinquency effects net of the selection the placebo signs.

Our contributions are four. The first is the estimand: the effect on individuals who actually start betting, rather than an intention-to-treat effect read off legalization. Since who starts is not random, we read the placebo on a false event as a signed bound rather than as a specification test to be passed, and we date distress relative to the bet by reversing the roles of the two events. The second is that one of our outcomes is set by the lender: the credit limit is a supply-side quantity, so the response to entry is measured rather than inferred from what the borrower does. The third is the joint observation of three records for the same individuals over the same months: the payment flow that measures the exposure, the credit registry that measures the response in credit and in realized default, and the account balance, which measures liquidity and accumulated saving. The fourth is coverage: we capture betting through unauthorized as well as authorized operators, in a market that is mostly casino rather than sports and that evidence identified off sportsbook legalization does not reach.

External validity. The individuals we observe are older than the population of bettors the regulator reports (Secretaria de Prêmios e Apostas, 2026), and that difference has a known sign: the delinquency response concentrates in the youngest group while the limit cut does not vary with age (Section 5.2.2), so at the national composition the delinquency effects would be larger and the limit cut unchanged. They are also likely of higher income, since a large share of them carry out some entrepreneurial activity, whereas the national universe of bettors includes a sizable participation of social-program beneficiaries (Banco Central do Brasil, 2024); if the response is stronger where there is less margin, as the gradient by prior credit position suggests, that difference points the same way.

Policy and welfare implications. Brazil's current instruments act on the betting operator and on the bettor. Operators must hold a federal authorization, are confined to a dedicated domain and to transfers from the bettor's own registered account, and are blocked when unauthorized; bettors face a centralized self-exclusion platform and mandatory self-limits on time and amount since November 2025, and social-program beneficiaries are barred from registering. Where the financial system appears, it is as the operator's counterparty: institutions may not hold the accounts of unauthorized operators (Table 1). A provisional measure of May 2026, after our sample window closes, reaches the credit decision itself, barring institutions from extending credit tied to a client's transfers for betting and blocking for twelve months the registration of individuals who enroll in the federal debt-renegotiation program *Desenrola Brasil* (Brazil, 2026). What we measure is on the other margin: not credit extended to fund a bet, but credit already held and withdrawn once betting starts, and withdrawn most where the borrower's prior standing with lenders was weakest. The rise in delinquency is smaller than the raw estimate suggests, since part of that estimate is distress that preceded the first bet.

Part of the entry into betting is downstream of financial distress, so an instrument that acts at the moment of the first bet reaches the individual after the deterioration is under way, whereas the renegotiation program acts on individuals already in distress. The instruments addressed to the bettor also operate through the register of authorized operators, and authorized operators account for about a quarter of the flow we observe. The segment they leave out is nonetheless measurable, since payment data identify betting payees the register does not list.

The costs we measure also belong in the fiscal arithmetic of legalization. Goss and Mangrum (2026) express a state's case for legalizing as tax revenue generated per additional delinquency, and Baker et al. (2024) note that part of the revenue a state collects is offset by the public cost of managing overextended borrowers. In Brazil the two sides of that arithmetic are not independent. The tax base is the operator's gross gaming revenue, which is by construction the aggregate net loss of bettors, so the losses exceed the revenue they generate by the inverse of the rate: revenue is a fraction of the losses, not a quantity to be set against them. The other leg of the tax, the levy on the individual's winnings, collects close to nothing: over three months on a single authorized platform, Lauro and Sacramento (2026) find no bettor with a monthly profit above the minimum wage, and none crossing the threshold at which winnings are taxed. Our estimates add an item no revenue calculation carries: the contraction of credit among the borrowers who had least of it, a cost with no offsetting revenue.

The natural next step of this paper is consumption. The balance drawdown says that liquidity falls without saying what it displaces; classifying account outflows by destination would show whether betting is funded out of accumulated saving or substitutes for consumption.

A Construction of the betting-operator list

This appendix states in full the procedure summarized in Section 3.2: the four approaches through which a Pix payee is identified as belonging to the betting ecosystem, the external checks on each of them, and the classification of operators by product type. The list this procedure yields is what defines who counts as a bettor, so the criteria are reported here rather than left to a replication file.

A.1 The four approaches

- **(i) Authorized operators:** The public SPA/MF list of authorized operators, with the corresponding brand names, taken as ground truth. Because an authorized operator receives Pix under its own tax registry only once the regime is in force, this becomes visible to us from January 2025.

- **(ii) Payees identified by their transactional behavior:** Betting deposits leave a distinctive pattern: small, round-valued, recurrent transfers concentrated at night. We score every payee that receives Pix from the individuals in our data on this pattern and flag those that carry it. The score reads six transactional features at the payee level: a low median ticket, a high share of round amounts (both integer values and multiples of five), a high share of micro-deposits, a concentration in night and early-morning hours, and high recurrence per payer.[10] We require most of these features rather than all, so as not to drop a genuine operator that misses one, such as sports platforms whose activity peaks at match times rather than late at night. Reading only the payee's own pattern, this approach spans the full sample.

- **(iii) Payment intermediaries, before regulation:** Some payees mix betting money with large volumes of unrelated payments, diluting their own transactional pattern so that approach (ii) fails to flag them. We recover them from their payer base instead. We first identify confirmed bettors in the regulated period, then look back at where those same individuals were sending Pix in 2024: a payee whose 2024 transactions came

[10]This approach follows and extends the methodology the Central Bank used to size the market from Pix data (Banco Central do Brasil, 2024). Formally, six binary criteria: a median ticket of at most R$50; at least 85% integer amounts; at least 60% multiples of five; at least 60% of transfers up to R$50; at least 35% of transfers between 6 p.m. and 5 a.m.; and at least three transfers per payer. A payee meeting at least five of the six is flagged. Two further guards apply to all three unauthorized approaches (ii)–(iv): a minimum of 100 distinct payers, and a blocklist of economic-activity codes whose deposit profile mimics betting but is unrelated to it (mainly fuel stations, plus other retail and public-service codes), with a few known betting intermediaries re-included by hand.

mostly from these confirmed bettors was itself a betting intermediary, even though nothing in its own pattern revealed it at the time.

- **(iv) Payment intermediaries, after regulation:** The same recovery applied within the regulated period, capturing intermediaries that continued routing bets after it.[11]

Table A1: Construction of the list of betting payees, by approach

Approach	Construction	Identifiable to us	Payees	Individuals
(i)	Authorized operators (SPA/MF register)	From Jan 2025	87	428,568
(ii)	Payees flagged by their own transactional pattern	Full sample	660	1,109,758
(iii)	Payment intermediaries recovered from their payers, before regulation	Retrospective	289	910,310
(iv)	Payment intermediaries recovered from their payers, after regulation	From Jan 2025	224	558,079
	Total (distinct)		1,260	1,785,367

Note: The list is built through four complementary approaches spanning the pre- and post-regulation phases. *Identifiable to us* records when each approach first becomes visible in our data: authorized operators are read from the SPA/MF register and appear under their own tax registry only from the January 2025 regime; the transactional pattern is a property of the payee's own Pix and so operates across the full sample; and the two intermediary approaches are recovered from the share of a payee's payers who are confirmed bettors. *Payees* counts distinct betting payees (CNPJ) assigned to each approach; the 1,260 total splits into 87 authorized by the SPA/MF and 1,173 unauthorized. *Individuals* counts distinct individuals who sent Pix to payees in each approach over the identification window (2020 to mid-2026); because an individual may transact with payees in more than one approach, the per-approach counts overlap and do not sum to the total, which is the total distinct count.

A.2 External validation

We conduct three external checks on the list. First, we compare approach (ii) with the SPA register: the transactional score alone flags 91.5% of the authorized operators with enough volume. Second, we look up the highest-volume intermediaries from approaches (iii) and (iv) on *Reclame Aqui*, a consumer-complaint website, and read their complaints for betting terms and brand names – of those with a page on the site, 77 of 99 (before regulation)

[11] Approaches (iii) and (iv) are applied within each period, with a bar that adapts to each. Because betting is more prevalent after regulation, a given concentration of known bettors among a payee's payers is less discriminating then, so we require a higher share afterward to hold the standard of evidence constant. Doing so also keeps an intermediary's betting role from being diluted, since that flow may have shifted between the two periods.

and 24 of 38 (after) carried betting-related complaints.[12] Third, we compare our list with the payment gateways that the SPA/MF identified as serving unauthorized operators – a list compiled jointly by the SPA/MF and the industry association (ANJL), made public during the Senate's 2025 inquiry into betting (*CPI das Bets*) (Secretaria de Prêmios e Apostas, 2025). Of the 67 gateways it names, our list contains 64 (96%): 41 through approach (ii), 21 through approach (iii), and 2 among the authorized operators of approach (i). The 21 recovered through approach (iii) matter most, because that is the layer a direct reading of the 2024 data cannot see.

A.3 Classifying operators: casino versus sports

The label is built from the brands behind each operator. What links a tax registry to its brands depends on regulatory status: for authorized operators, the SPA/MF register already lists the corresponding brand names; for unauthorized operators and intermediaries, we take the brand mentions in the complaints filed against each operator on *Reclame Aqui*. In both cases we then access each brand's site and score its homepage on a casino-to-sports scale using a large-language model, which reads the page text and a screenshot as two separate scores and averages them, falling back to the brand name when the page has no usable content; brand scores are aggregated up to the operator by the number of mentions. An operator is classified as sports-type when its sports ratio is at least 0.55, casino-type when it is at most 0.30, hybrid in between, and inconclusive when the content is too thin to score. What we score is the core business – where a brand concentrates product, marketing, and users – not the products it merely lists, since almost every platform offers both. The procedure classifies 255 operators with a determinable profile, covering about 66% of post-regulation bettors (Table A2).

Three checks support the labels. The textual and visual scores are highly correlated ($r = 0.79$, with categorical agreement of 82%); the largest sportsbooks (Bet365, Betano, Sportingbet) fall at the top of the sports axis and the well-known casino brands at the bottom; and pre-regulation transactional behavior corroborates the label through an independent channel, since sports bettors show a higher average ticket (R$80 against R$51), lower frequency, fewer platforms, and a spike in activity on football national-league match days that casino bettors do not display.

[12] We searched for complaints carrying betting terms (*aposta*, *saque*, *depósito*, *cassino*, *tigrinho*, *rodada*, and the like) and the brands named in them. Intermediaries with no page could not be checked this way; given the precision of the checks on those we could verify, we kept them in the list.

Table A2: Casino-versus-sports classification of operators, by regulatory status

Profile	Regulated	Unregulated	Total
Casino-type	16	87	103
Inconclusive	0	79	79
Hybrid	34	6	40
Sports-type	27	1	28
Lottery-type	4	0	4
Fantasy-type	1	0	1
Total	82	173	255

Note: Distribution of the 255 operators with a determinable profile, by type and by regulatory status (authorized by the SPA/MF or not). Type is assigned from the sports ratio of each operator's platform content; hybrid and inconclusive operators are not assigned to either type. Lottery-type and fantasy-type operators are residual categories excluded from the casino–sports analysis.

B Additional results

This appendix collects the exercises that test, one at a time, the choices the main specification makes: which entry cohorts enter the treated group, which payees count as betting operators, which individuals enter the sample, and on what scale the outcomes are measured. None of the four moves the reading of Section 5.1. The sensitivity analysis of Rambachan and Roth (2023), which disciplines the same estimates against violations of parallel trends, is collected separately in Appendix C.

B.1 Excluding the initial entry cohorts

The first months of the regime are where the change in payment observability is concentrated: an individual whose deposits had been routed through an intermediary in 2024 can surface as a first-time bettor in early 2025 because the destination became identifiable, and not because the betting began. The operator list of Section 3.2 holds this margin down by reaching back through the pre-regulation intermediary layer, and dropping the January to March 2025 entry cohorts tests what is left of it (Table B1). The limit cut is essentially unchanged, and the delinquency outcomes attenuate modestly in the pattern expected of a shorter post-treatment window rather than of contamination.

B.2 Restricting the operator list to the direct layers

The two intermediary layers of Appendix A.1 are the ones identified indirectly, and they are also the ones that decide whether a 2025 deposit is a first bet or the continuation of

Table B1: Robustness to excluding the initial entry cohorts (January–March 2025)

Outcome	Baseline	Excluding initial cohorts
Log limit (log points)	-0.174^{***}	-0.164^{***}
	(0.013)	(0.017)
Delinquency (flag), pp	3.66^{***}	3.08^{***}
	(0.26)	(0.32)
Card delinquency rate, pp	4.41^{***}	3.71^{***}
	(0.22)	(0.29)

Note: CS-DR estimates (conditional) of the average effect θ^{att} of (3) for the three credit outcomes, re-estimated dropping the January–March 2025 entry cohorts, which concentrate the payment reclassification of Section 4.1. The limit is reported in log points, the delinquency outcomes in percentage points. Standard errors, clustered at the individual level, in parentheses. The credit cut is essentially stable – the limit moves from -0.174 to -0.164 in log points, about -16.0% against -15.1%, a 6% fall in the coefficient – which is the pattern expected when the effect is not an artifact of reclassification, since removing the suspect cohorts does not knock the coefficient down. The delinquency outcomes attenuate modestly, an attenuation consistent with the shorter post-treatment window of the cohorts that remain rather than with contamination. Estimating the effect separately by entry cohort points the same way: the initial cohorts show no anomalously larger effect than the later ones, the limit cut registering -0.24 in the January 2025 cohort against -0.20 in the July 2025 cohort. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

a habit that ran through an intermediary in 2024. Restricting the list to the two direct layers therefore redefines the treated group at its root, in two directions at once: it drops the individual whose only identified destination was an intermediary, and it moves into the new-bettor group the individual whose 2024 intermediary history is no longer visible and whose first direct deposit falls in 2025.

The exchange is close to neutral in aggregate. The restricted universe retains 89% of the new bettors of the full one but 99.6% of the amount deposited, because what leaves is a mass of low-intensity bettors reached only through intermediaries and what enters is a smaller and heavier group reclassified out of the pre-regulation cohort. The estimates do not weaken (Table B2): the limit cut deepens from -0.174 to -0.202 log points, about -16% against -18%, while the two delinquency outcomes and the balance drawdown each move by roughly a tenth of their own magnitude and keep their sign and significance. What dependence there is runs in the conservative direction: the broader definition the paper uses is the one that yields the smaller cut.

B.3 Requiring an active account before entry

This exercise restricts both groups to individuals whose account was already active before entry, which is the restriction the balance specification imposes (Section 5.1), so that the credit and the balance results rest on one population. The restriction is demanding on the control pool: it retains 76.4% of the new bettors of the credit panel but only 25.3% of

Table B2: Robustness to restricting the operator list to the two direct layers

Outcome	Full list, (i)–(iv)	Direct layers, (i)–(ii)
Log limit (log points)	-0.174^{***}	-0.202^{***}
	(0.013)	(0.014)
Delinquency (flag), pp	3.66^{***}	3.22^{***}
	(0.26)	(0.22)
Card delinquency rate, pp	4.41^{***}	4.17^{***}
	(0.22)	(0.25)
Change in balance†	-0.605^{***}	-0.558^{***}
	(0.034)	(0.030)

Note: CS-DR estimates of the average effect θ^{att} of (3), re-estimated on the universe that results from keeping only layers (i) and (ii) of Table A1, the official register and the transactional score, and dropping the two intermediary layers. The restriction is applied to the Pix destinations before the first-bet date is computed, so it redefines who is a new bettor, when treatment begins, and who falls back into the never-bettor group; it is a different universe rather than a subsample of the same one. The limit is in log points, the delinquency outcomes in percentage points, and the change in balance on the asinh scale. Standard errors, clustered at the individual level, in parentheses. The three credit rows are the conditional specification, on the covariate set of Table 4, and the baseline column reproduces the estimates reported there and in Table B1. † The change in balance is the *unconditional* specification of the active-account panel in both columns, since that is what was estimated on the restricted universe, so its baseline is the unconditional -0.605 of Table 4 rather than the conditional headline of -0.569. $^{*}\,p < 0.05$; $^{**}\,p < 0.01$; $^{***}\,p < 0.001$.

the never-bettors, since a dormant account still carries credit in the registry and so meets every other sample requirement. On the conditional sample that identifies the effect the attrition is far milder, because the covariate requirement already selects individuals with a credit relationship: 33,064 of the 37,619 new bettors (87.9%) and 115,783 of the 230,614 never-bettors (50.2%) remain, for 148,847 individuals, or 136,877 for the card delinquency rate. The estimates do not weaken. The limit falls by 0.206 against 0.174 in the baseline, the delinquency flag rises by 3.91 against 3.66 pp and the card rate by 4.52 against 4.41 pp.

What the restriction buys is a cleaner falsification test. The unconditional group placebo, significantly different from zero in every outcome on the full panel, becomes indistinguishable from zero for both delinquency measures (0.000 and 0.001, standard error 0.002). The conditional placebo of the card rate also vanishes (-0.000, standard error 0.002), so that outcome no longer requires the bounded reading, and the remaining conditional placebos shrink while keeping their sign: 0.040 against 0.069 for the limit and 0.006 against 0.011 for the delinquency flag. The pre-treatment window tightens accordingly, the largest conditional violation falling from 0.029 to 0.020 for the limit. The debiased estimate of (5) is therefore almost unchanged, -0.246 for the limit against -0.243, while the delinquency readings rise to 3.33 and 4.56 pp. The point estimate moves toward a debiased estimate that stays where it was, the pattern expected if dormant controls attenuate the effect and inflate the placebo rather than generate it.

Two qualifications frame how much this exercise can carry. It is not an improvement in balance: what it removes is disproportionately the thin-portfolio never-bettors, who are the ones closest to the treated, and the standardized differences of the conditional sample deteriorate slightly, the limit from −0.230 to −0.249 and the high-cost share from 0.209 to 0.248, both still inside the conventional threshold. Nor is it a bias correction: it also drops 23.6% of the treated, those who bet while their account lies dormant, so the estimates above are the effect for bettors whose account is live, a different population rather than a cleaner estimate of the same one. Within the unconditional specification, whose levels the paper does not read causally, the never-bettor comparison becomes unstable at the resulting ratio (the limit at −0.005 against 0.134), which is a further reason to read the restriction through the conditional column alone. The sensitivity analysis of Section 4.3.1 is not re-estimated on this sample; the exercise is read through the point estimates and the placebos.

B.4 Measurement in reais and natural-person accounts

Two further exercises test the scale on which the outcomes are measured and the legal nature of the account. The first replaces the log and asinh transformations by the level in reais, on the same sample and with the same estimator, so that the effect is read per individual rather than proportionally (Figure B1). Winsorized at the 99th percentile, the limit falls by R$2,681 (s.e. 226), while the monthly change in the balance falls by R$81 (s.e. 9); both keep the sign and the significance of the transformed specification. Without winsorizing, the point estimates inflate by roughly an order of magnitude and the standard errors by about thirteen times (the limit at R$13,258 with a standard error of 2,962), which is the expected signature of a few extreme accounts dominating the mean and the reason the reported specification is the transformed one.

The second restricts the sample to natural-person accounts, dropping those held under the tax registry of the individual's entrepreneurial activity (Figure B2). The control pool survives the restriction, retaining 85.5% of the never-bettors in the credit panel and 79.5% in the account panel, and the estimates are indistinguishable from the full-sample ones: the limit at −0.181 against −0.174, the delinquency flag at 3.56 against 3.66 pp, the card rate at 4.42 against 4.41 pp, and the change in the balance at −0.575 against −0.591.[13] The effects are not driven by the entrepreneurial leg of this population.

[13]The balance comparison is run against the estimate on the subsample of the level exercise (−0.591) rather than against the headline (−0.569), because that exercise draws its own subsample; the two columns of the comparison are therefore internally consistent.

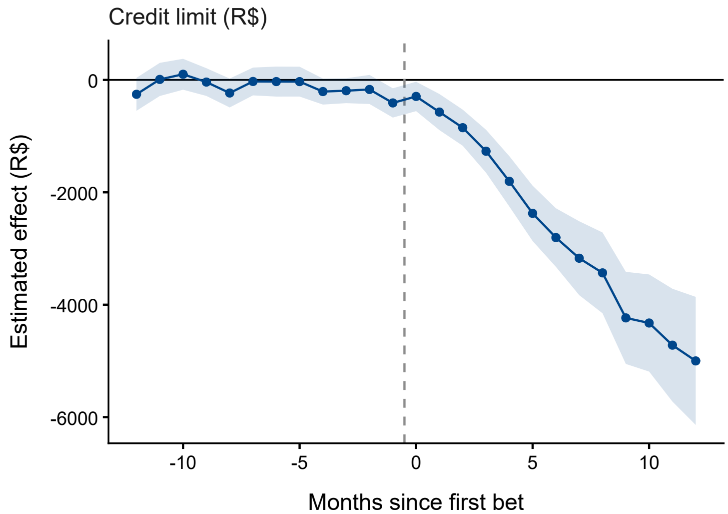

Figure B1: Robustness: credit measured in reais.

Notes: CS-DR estimates of the event-study profile (3) (conditional) for the credit limit measured in reais and winsorized at the 99th percentile, by month relative to the first bet. The shaded band is the pointwise 95% confidence interval, since this exercise stores the analytical standard errors rather than the multiplier-bootstrap bands. The vertical scale is in reais and is not comparable in magnitude with the log specifications of the main text; the corresponding averages, and the unwinsorized counterparts, are in the text of this appendix.

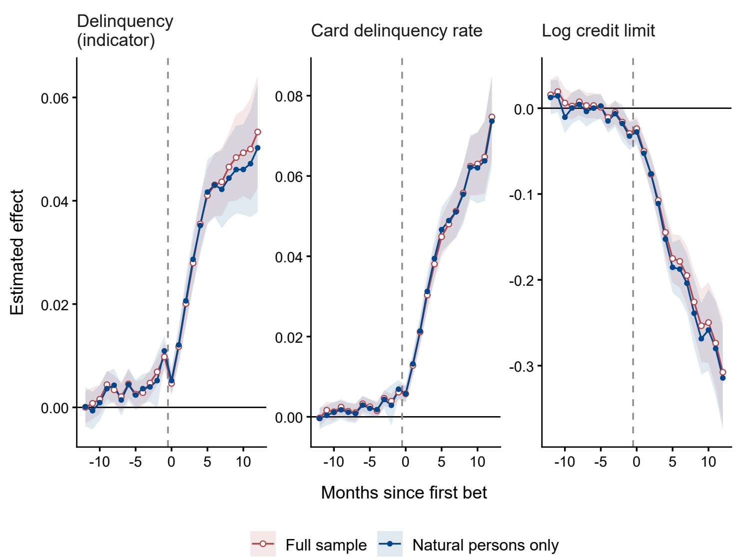

Figure B2: Robustness: sample restricted to natural-person accounts.

Notes: CS-DR estimates of the event-study profile (3) (conditional) for the three credit outcomes, comparing the full estimation sample with the sample restricted to natural-person accounts, by month relative to the first bet. Standard errors clustered at the individual level; the shaded bands are pointwise 95% confidence intervals, standardized across the two series so that they are comparable with each other. The restriction keeps 85.5% of the never-bettors of the credit panel.

C Sensitivity to violations of parallel trends

Section 4.3.1 describes the sensitivity analysis of Rambachan and Roth (2023) and the signed variant of it that we read, but states neither formally and reports only the breakdown $\bar{M}$ alongside each set of estimates. This appendix supplies the two restrictions imposed on the post-treatment violation, and the exercise in full: the average post-treatment effect each breakdown is computed on, the original and robust confidence intervals, and the curves that trace the robust set as the admissible violation grows.

The two restrictions. Decomposing the event-study profile of (3) into effect and bias, $\theta^{\text{es}}(k) = \tau_k + b_k$, the target is $\tau = \frac{1}{13}\sum_{k=0}^{12} \tau_k$ – the same object as the θ^{att} reported throughout – and the relative-magnitudes restriction bounds the post-treatment violation by the largest one visible before entry,

$$\Delta^{\text{RM}}(\bar{M}) = \Big\{ b : \ \Big| b_{k+1} - b_k \Big| \ \leq \ \bar{M} \max_{s<0} \Big| b_{s+1} - b_s \Big| \ \ \forall\, k \geq 0 \Big\}. \tag{8}$$

The breakdown $\bar{M}$ is the largest value for which the 95% robust confidence set for τ still excludes zero. Restriction (8) is symmetric, whereas the reading of Section 4.3.1 is signed; the signed bound proper intersects it with the direction of the observed pre-trend,

$$\Delta^{\text{RM,s}}(\bar{M}) = \Delta^{\text{RM}}(\bar{M}) \ \cap \ \Big\{ b : \ \text{sign}(\hat{\pi})\, b_k \geq 0 \ \ \forall\, k \geq 0 \Big\}, \tag{9}$$

where $\hat{\pi}$ is the false-event placebo of (5). It does not bind for the same-sign outcomes. For the opposite-sign ones it binds, and whether we read it turns on a second condition: the restriction fixes the direction of the post-treatment bias while (8) anchors its admissible size on the design's own pre-treatment window, so the two are coherent only where that window points the same way. It does for the change in the account balance, whose largest pre-treatment violation is +0.086 and whose pre-treatment profile stands at +0.087 at the reference month; it does not for the credit limit, whose largest violation is −0.013 in the pooled sample, −0.036 in the thin-credit subgroup and −0.040 in the casino arm, with the profile reaching −0.029, −0.031 and −0.034 at $k = -1$. We therefore read the sign-restricted interval for the balance and the symmetric one for the limit, and report both columns throughout. In every application $\hat{\pi} > 0$ – the conditional placebo is positive in all four outcomes of Table 4, and likewise in each subgroup of Table 5 and each arm of Table 7 – so the restriction actually imposed is $b_k \geq 0$. Neither restriction nests the debiased estimate of (5): at its tightest, $\bar{M} = 0$, (8) returns $\hat{\theta}^{\text{att}}$ itself rather than $\hat{\theta}^{\text{db}}$.

Table C1: Sensitivity to parallel-trends violations: breakdown $\bar{M}$ by credit outcome

Outcome	Average post effect	Original 95% CI	Breakdown $\bar{M}$ Sym.	Breakdown $\bar{M}$ Signed	Signed robust CI at $\bar{M}=1$	Placebo
Log limit	−0.174 (0.013)	[−0.200, −0.148]	0.50	> 2	[−0.263, −0.137]	opp.
Card delinquency rate	0.044 (0.002)	[0.040, 0.048]	0.75	0.75	[−0.012, 0.045]	same
Delinquency (flag)	0.037 (0.003)	[0.032, 0.041]	0.25	0.25	[−0.049, 0.044]	same
Change in balance†	−0.605 (0.034)	[−0.671, −0.539]	0.25	> 2	[−0.669, −0.514]	opp.

Note: Sensitivity analysis of Rambachan and Roth (2023) applied to the conditional event-study profile ATT(k) of (3), on the conditional credit sample of Table 4 (268,233 individuals, 28 months); implementation as described above. The average post effect is the equally weighted mean of the thirteen post-treatment coefficients; standard errors, clustered at the individual level, in parentheses. *Sym.* imposes (8) and *Signed* imposes (9); > 2 denotes survival of the whole grid, and *Placebo* records the sign of the conditional group placebo relative to the estimated effect. The signed robust interval is reported for every outcome; for the two same-sign outcomes, whose signed breakdown is below one, it already includes zero at $\bar{M}=1$ and conveys nothing the breakdown does not. † For the change in the account balance the exercise is run on the *unconditional* event study of the active-account panel, with one month of anticipation, so its average post effect reproduces the unconditional rather than the conditional column of Table 4. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Implementation. Both variants run on the conditional event-study coefficients over $k \in [-12, 12]$, with $\bar{M}$ on a grid from 0 to 2 in steps of 0.25 and equal weight on the thirteen post-treatment periods. Because the procedure builds the variance analytically from the influence function, the event study is re-estimated for every exercise in this appendix with a universal base period and without the multiplier bootstrap, so its coefficients are not numerically the ones plotted in Figure 5; the resulting average post effect reproduces the conditional estimates of Table 4 to three decimals.

Table C1 and Figure C1 cover the pooled sample and the balance drawdown, and Table C2 the two bet-type arms. The counterpart by prior credit position stays in the main text, in the last two columns of Table 5.

D The pre-treatment window of the delinquency-to-betting design

The sensitivity analysis of Appendix C is not run for the delinquency-to-betting design of Table 8, for the reason given in the note to that table: the pre-treatment coefficient of (4) is minus the initiation rate of the comparison group over the same months, a known quantity rather than a test.

Its magnitude is a matter of the base-period convention, and Figure D1 shows both. Under the varying base period, the convention of Figure 9, each pre coefficient is a month-on-month

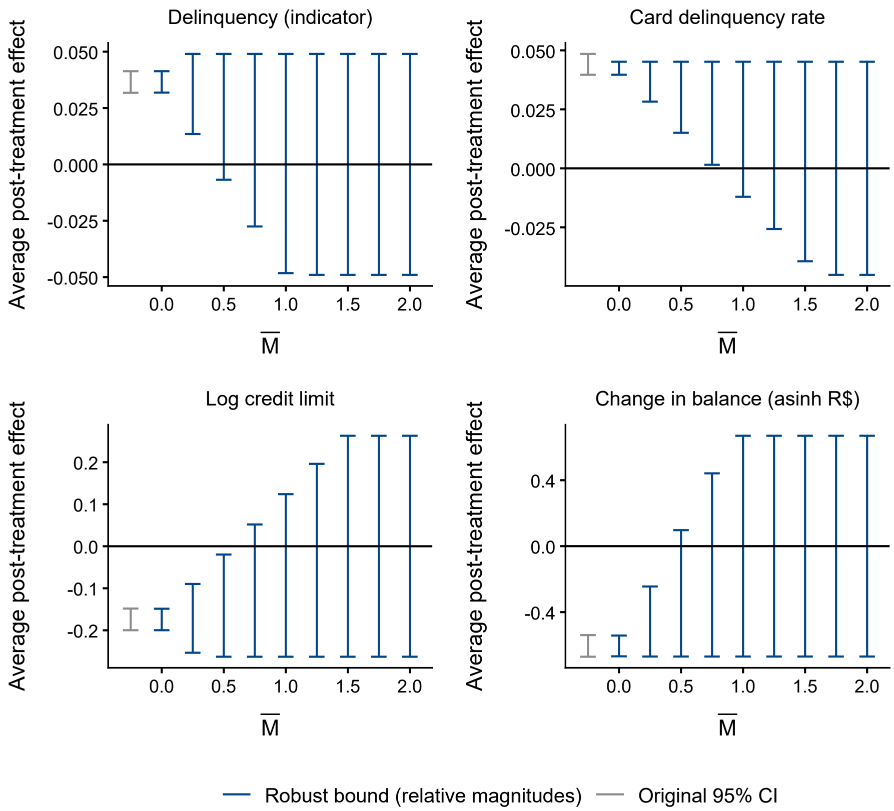

Figure C1: Sensitivity to parallel-trends violations, by credit outcome.

Notes: sensitivity analysis of Rambachan and Roth (2023) applied to the conditional event-study profile ATT(k) of (3), one panel per credit outcome and a fourth for the monthly change in the account balance, on the credit sample of Table 4 for the first three and on the active-account panel, with one month of anticipation and under the unconditional specification, for the fourth. Each panel plots the robust confidence set for the average post-treatment effect as the allowed post-violation size $\bar{M}$ grows (in multiples of the largest pre-treatment violation); the original CS-DR interval is at the left. The breakdown $\bar{M}$ of Table C1 is the point at which the robust set first includes zero. The curves shown are the *symmetric* restriction (8) only; which restriction we read for which outcome is set out above. The sign-restricted counterpart (9) is not plotted, and its breakdown and robust intervals are in the last two numeric columns of Table C1.

first difference and averages -0.17 pp in the unconditional specification and -0.09 pp in the conditional one. Under the universal base period, in which every horizon is measured against the month before entry, the same estimates give $+2.62$ and $+1.50$ pp twelve months before entry, declining monotonically to zero at $k = -1$. The two are the same fact at two scales, and the second is of the same sign as the effect and of its order of magnitude: applying the signed reading of Section 4.3.1 to it would return $+0.46$ pp in the unconditional specification against the $+1.76$ estimated. We therefore read neither as a bound. The thirteen post-treatment coefficients are identical under the two conventions, to three decimals, since for $k \geq 0$ both measure against $k = -1$; the estimate of Table 8 does not depend on the choice.

A further exercise supports the reading. Restricting the treated group to individuals who became delinquent from January 2025 onward – which aligns the treatment window with the

Table C2: Sensitivity to parallel-trends violations, by product type

Outcome	Average post effect	Robust 95% CI at $\bar{M} = 0$	Breakdown $\bar{M}$ Sym.	Breakdown $\bar{M}$ Signed	Signed robust CI at $\bar{M} = 1$
Panel A. Casino arm					
Delinquency (flag)	0.049	[0.040, 0.059]	0.25	0.25	—
Card delinquency rate	0.059	[0.050, 0.068]	0.50	0.50	—
Log limit	−0.248	[−0.298, −0.199]	0.50	> 2	[−0.509, −0.177]
Panel B. Sports arm					
Delinquency (flag)	0.040	[0.019, 0.060]	0.00	0.00	—
Card delinquency rate	0.063	[0.044, 0.083]	0.25	0.25	—
Log limit	−0.097	[−0.196, 0.002]	n.s.	n.s.	—

Note: Sensitivity analysis of Rambachan and Roth (2023) (Section 4.3.1) run separately for each arm of Table 7, anchored on that arm's own pre-treatment window; it is the material summarized by the $\bar{M}$ columns of that table. The sign restriction (9) is the valid reading for the limit, whose arm placebos run against the effect. "n.s." marks a cell whose robust interval already includes zero at $\bar{M} = 0$, so that no restriction on the direction of the bias can rescue it. The signed robust interval is reported where it is informative: for the four delinquency cells the signed breakdown is itself below one, so the interval includes zero at $\bar{M} = 1$ and adds nothing to the breakdown. The average post effect reproduces the arm θ^{att} of Table 7, since the universal base period leaves the post-treatment coefficients unchanged.

window in which the outcome was available, at the cost of 29% of the treated – raises the cumulative effect to +2.79 pp unconditional and +4.89 pp conditional rather than reducing it, so the earlier cohorts were diluting the estimate: their post-treatment window falls largely in 2024, when betting was still rare in this population. We report the restricted-cohort estimates as a direction rather than as a magnitude: with 11,869 treated in the conditional specification the longest horizons are thinly identified.

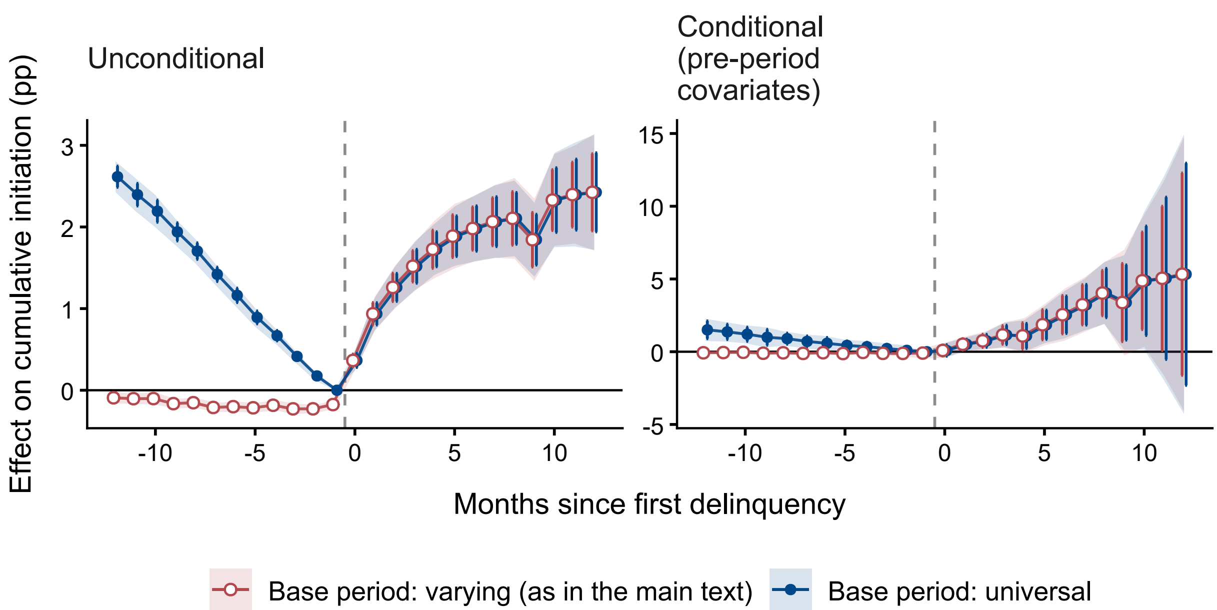

Figure D1: Base-period convention in the delinquency-to-betting design.

Notes: CS-DR estimates of (4) for the cumulative indicator of having started betting, by month relative to the first delinquency, under the two base-period conventions and in both specifications. Under the varying convention, which is what Figure 9 plots, each pre-treatment coefficient is a month-on-month first difference; under the universal convention every horizon is measured against $k = -1$, which is therefore zero by definition. Bars are pointwise 95% confidence intervals and the shaded band is the simultaneous band; the vertical scales differ across panels. The post-treatment coefficients coincide under the two conventions because both measure against $k = -1$ for $k \geq 0$, so the divergence is confined to the pre-treatment window and is the initiation rate of the comparison group, cumulated or not.

References

Daniel Pereira Alves de Abreu and Aureliano Angel Bressan. What makes Brazilians gamble in bets? empirical evidence from ANBIMA's investor x-ray for 2024 and 2025. *Revista Brasileira de Finanças*, 24(1):e202602, 2026.

Sumit Agarwal, Paige Marta Skiba, and Jeremy Tobacman. Payday loans and credit cards: New liquidity and credit scoring puzzles? *American Economic Review*, 99(2):412–417, 2009. doi: 10.1257/aer.99.2.412.

Hunt Allcott, Matthew Gentzkow, and Lena Song. Digital addiction. *American Economic Review*, 112(7):2424–2463, 7 2022. doi: 10.1257/aer.20210867.

Scott R. Baker, Justin Balthrop, Mark J. Johnson, Jason D. Kotter, and Kevin Pisciotta. Gambling away stability: Sports betting's impact on vulnerable households. NBER Working Paper 33108, National Bureau of Economic Research, 2024.

Banco Central do Brasil. Análise técnica sobre o mercado de apostas online no Brasil e o perfil dos apostadores [technical analysis of the online betting market in Brazil and the profile of bettors]. Estudos Especiais do Banco Central 119, Banco Central do Brasil, Brasília, 2024. Technical Note 513/2024-BCB/SECRE, 23 September 2024.

Garrick Blalock, David R. Just, and Daniel H. Simon. Hitting the jackpot or hitting the skids: Entertainment, poverty, and the demand for state lotteries. *American Journal of Economics and Sociology*, 66(3):545–570, 2007. doi: 10.1111/j.1536-7150.2007.00526.x.

Brazil. Lei nº 13.756, de 12 de dezembro de 2018 [law no. 13.756 of 12 december 2018], 2018. Creates the fixed-odds betting modality (Art. 29).

Brazil. Lei nº 14.790, de 29 de dezembro de 2023 [law no. 14.790 of 29 december 2023], 2023. Regulatory framework for fixed-odds betting.

Brazil. Provisional measure no. 1,355 of 4 may 2026, 2026. Establishes the *Novo Desenrola Brasil* household debt-renegotiation program; Art. 16 bars financial institutions from granting credit operations directly linked to a client's transfers for fixed-odds betting. In force at the time of writing; not converted into law.

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

Fernando Chague and Bruno Giovannetti. Day-trading stocks for a living? *Revista Brasileira de Finanças*, 18(3):1–4, 2020. doi: 10.12660/rbfin.v18n3.2020.81949.

Clément de Chaisemartin and Xavier D'Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.

Mingze Gao, Henry Leung, Linhui Liu, and Buhui Qiu. Consumer behaviour and credit supply: Evidence from an australian fintech lender. *Finance Research Letters*, 57:104205, 2023. doi: 10.1016/j.frl.2023.104205.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.

Jacob Goss and Daniel Mangrum. Sports betting across borders: Spatial spillovers, credit distress, and fiscal externalities. Staff Reports 1184, Federal Reserve Bank of New York, 2026.

Sylvan Herskowitz. Gambling, saving, and lumpy liquidity needs. *American Economic Journal: Applied Economics*, 13(1):72–104, 2021. doi: 10.1257/app.20180177.

Brett Hollenbeck, Poet Larsen, and Davide Proserpio. The financial consequences of legalized sports gambling. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC '25)*. Association for Computing Machinery, 2025. doi: 10.1145/3736252.3742536. Forthcoming, *Management Science*; SSRN working paper 4903302.

iGaming Business. Keeping players engaged: how features enhance the sportsbook experience in LatAm. https://igamingbusiness.com/content-os/latin-america-sportsbook-engagement-features-brazil-market/, 2026. Accessed July 2026.

Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015.

Kantar Ibope Media. Ranking dos 300 maiores anunciantes do Brasil [ranking of the 300 largest advertisers in Brazil]. Published in *Meio & Mensagem*, "Agências & Anunciantes" special edition, 2026. Data for 2025. Betting firms (15 companies) invested R$ 2.9 billion in media purchases, up from R$ 2.0 billion (11 companies) in 2024.

Lance! Saiba quanto as empresas gastaram em publicidade com apostas esportivas [how much firms spent on sports-betting advertising]. News report, 2025. URL https://www.lance.com.br/lance-negocios/saiba-quanto-as-empresas-gastaram-em-publicidade-com-apostas-esportivas.html.

Leonardo M. Lauro and Luiz C. Sacramento. Sports gambling in Brazil: Evidence from real bets. Working paper, PUC-Rio and Instituto Pensi, 2026. March 2026.

Virve K. Marionneau, Aino E. Lahtinen, and Janne T. Nikkinen. Gambling among indebted individuals: an analysis of bank transaction data. *European Journal of Public Health*, 34 (2):342–346, 2024. doi: 10.1093/eurpub/ckad117.

Naomi Muggleton, Paula Parpart, Philip Newall, David Leake, John Gathergood, and Neil Stewart. The association between gambling and financial, social and health outcomes in big financial data. *Nature Human Behaviour*, 5(3):319–326, 2021. doi: 10.1038/s41562-020-01045-w.

Pay4Fun. Levantamento da Pay4Fun: 97% das transações para casas de apostas foram feitas por Pix. Reported by Yogonet Brasil, 24 Feb. 2025. https://www.yogonet.com/brasil/noticias/2025/02/24/4546-levantamento-da-pay4fun-97-das-transacoes-para-casas-de-apostas-foram-feitas-por-pix, 2025. Accessed July 2026.

Poder360. Sites de apostas só perdem para o Google em acessos no Brasil. https://www.poder360.com.br/poder-economia/sites-de-apostas-so-perdem-para-o-google-em-acessos-no-brasil/, 2025. Based on SimilarWeb data via Aposta Legal. Accessed July 2026.

Ashesh Rambachan and Jonathan Roth. A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591, 2023.

Jonathan Roth, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244, 2023.

Pedro H. C. Sant'Anna and Jun Zhao. Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122, 2020.

Secretaria de Prêmios e Apostas. Portaria spa/mf nº 615, de 17 de abril de 2024 [ordinance spa/mf no. 615 of 17 april 2024], 2024. Regulates payment methods for fixed-odds betting: transfers limited to Pix, debit or prepaid cards, and bank wires, moving between the bettor's own registered account and the operator's account.

Secretaria de Prêmios e Apostas. Nota informativa sei nº 567/2025/mf: gateways de pagamento e instituições financeiras associados a operadores ilegais de apostas [payment gateways and financial institutions linked to unlicensed betting operators]. Nota Informativa 567/2025, Secretaria de Prêmios e Apostas, Ministério da Fazenda, Brasília, 2025. Response to Requerimento nº 393/2025 of the Senate's parliamentary inquiry into online betting (CPI das Bets); compiled from SPA/MF monitoring begun in 2024 and a report shared by the Associação Nacional de Jogos e Loterias (ANJL); accompanying spreadsheet SEI 49370002; signed 19 March 2025; Processo nº 19995.002439/2025-14.

Secretaria de Prêmios e Apostas. Panorama periódico do mercado regulado de apostas de quota fixa [periodic overview of the regulated fixed-odds betting market]. Relatório periódico, Secretaria de Prêmios e Apostas, Ministério da Fazenda, Brasília, 2026. Data from 1 January to 30 December 2025.

Joseph E. Stiglitz and Andrew Weiss. Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3):393–410, 1981.

Tribunal de Contas da União. Acórdão 1.296/2026 – tcu – plenário [ruling no. 1,296/2026 – tcu – plenary]. Acórdão 1296/2026, Tribunal de Contas da União, Brasília, 2026. TC 007.885/2025-8. Audit of the regulation and supervision of fixed-odds betting; reports 41–51% illegal market share (R$ 26–40 billion per year).

Tunad. Bets investem mais de r$ 1,4 bilhão em mídia em 2025 e reforçam a centralidade da tv aberta [betting firms spent over r$ 1.4 billion on media in 2025, reinforcing the centrality of open television]. Industry report, 2025. URL https://tunad.io/blog/bets-investem-mais-de-r-14-bilhao-em-midia-em-2025-e-reforcam-a-centralidade-da-tv-aberta/.