The Real Effects of Payment Anonymity: Evidence from the Crypto Ecosystem

Jiayi Wei*

June 1, 2026

Abstract

Does the design of payment systems cause real welfare consequences? Despite the rapid growth of cryptocurrency, the welfare effects of anonymous digital payments remain empirically unknown. I study the staggered rollout of over 30,000 Bitcoin ATMs across U.S. counties, kiosks that convert cash to cryptocurrency with minimal identification, as a laboratory for estimating these consequences. BTM entry increases total crime by 6.9 percent, with effects concentrated in offenses that exploit anonymity, irreversibility, and cash-to-crypto conversion, while crimes unrelated to payment design barely respond. Tracing the causal chain further, BTM entry fuels local drug markets and raises overdose mortality. Crime-induced losses erode household balance sheets, raising credit delinquency in affected neighborhoods. These harms are regressive. Minority and Hispanic communities, young adults, and areas with weaker financial infrastructure absorb the largest increases in victimization and the steepest declines in financial health, simultaneously widening pre-existing inequality in safety, public health, and credit access. To scale these local effects to the aggregate level, I first use on-chain forensic analysis to trace BTM-originated funds along their downstream paths, identifying the routing intermediaries each transaction passes through and the share that

*INSEAD, Boulevard de Constance, 77305 Fontainebleau, France. Email: jiayi.wei@insead.edu.

terminates at illicit endpoints. I then build an informational model in which anonymity at each stage weakens transaction traceability and raises social harm. Calibrated to these on-chain flow shares, the model translates the local causal estimates into aggregate welfare statements, yielding a national-level measure of the cost generated by anonymous payment infrastructure.

JEL Classification: K42, G23, O33, R23, I18

Keywords: payment system, anonymity, crypto, Bitcoin ATM, crime, credit, inequality

1 Introduction

The digitization of payments has been one of the major successes of modern finance. Each generation of payment technology has made transacting faster, cheaper, and more inclusive. Cards displaced checks, electronic transfers displaced paper, and mobile money brought hundreds of millions of unbanked households into the formal financial system. The records that digital payments generate have become valuable in their own right, expanding access to credit for borrowers who lack traditional histories [Ghosh et al., 2026]. Cryptocurrency appears to extend this frontier. Over 650 million people now own crypto assets, and stablecoin networks settle upward of $10 trillion per year.[1] New access points carry digital finance into neighborhoods that banks have left behind, raising local financial participation and new business formation [Cong et al., 2026], an effect that mirrors the wealth gains traditional bank branch expansions delivered to underserved households [Célérier and Matray, 2019]. Viewed through the standard lens of payment innovation, this expansion should be good news.

Yet one attribute separates this innovation from every previous step in that sequence. A payment system is defined not only by how it moves value but by what it remembers. Money itself is, at bottom, a substitute for a societal record of transactions [Kocherlakota, 1998]. Cash remembers nothing, and its physical nature kept that forgetfulness local and small in scale. The instruments that could move value globally and instantly remembered everything, recording the identities of the transacting parties along with the full transaction details. That record is what deters crime by exposing illicit payments to enforcement [Wright et al., 2017]. It is also what underpins credit scoring, tax compliance, and the recovery of stolen funds. Cryptocurrency unbundles the record. Its public ledger remembers every transaction perfectly, but it remembers no identities. The link between a person and an address is created or severed at the edges of the system, where cash turns into crypto.

[1] Global crypto ownership reached 659 million at the end of 2024 [Crypto.com, 2025]. Stablecoin settlement volume, adjusted by Visa to exclude bot and other inorganic activity, was approximately $10 trillion over the trailing twelve months [Visa, 2026].

The services that operate at these edges, together with mixers and unregulated exchanges downstream, supply the anonymity of cash at the scale of digital finance.

Whether this should worry a regulator depends on who demands the anonymity. The payments literature identifies two motives. Anonymity protects honest users, because a transaction record can be used against the person who created it, through data harvesting, price discrimination, or theft, and markets undersupply this protection [Kahn et al., 2005, Garratt and van Oordt, 2021, Ahnert et al., 2025]. Anonymity also serves evasion, because the record is what makes enforcement possible [Rogoff, 2017]. The difficulty is that anonymity cannot be sold selectively. A technology that shields the honest user's history necessarily shields the criminal's, since the supplier cannot observe the motive. One product faces two demand curves, and the policy question reduces to their composition. How much of the demand for payment anonymity is protection, how much is evasion, and what happens to real outcomes when its price falls?

Existing evidence cannot answer this question. Forensic measurement of illicit prevalence disagrees by an order of magnitude, with early work attributing nearly half of Bitcoin transaction volume to illegal activity [Foley et al., 2019] and subsequent work putting the share near three percent [Makarov and Schoar, 2021]. Even within market-microstructure work, prevalence of manipulative activity is itself disputed [Griffin and Shams, 2020]. The deeper problem is that prevalence is silent about the equilibrium. Suppose every illicit transaction that once settled in cash now settles in crypto. Measured prevalence would be enormous, yet the total amount of crime would be exactly what it was before. The policy question is whether anonymous payment infrastructure creates illicit activity or merely re-denominates activity that would have occurred anyway. A stock of illicit volume, however precisely measured, cannot separate the two. Only the causal response of real outcomes can, and this paper provides it.

The paper proceeds in three steps. First, I exploit the staggered rollout of more than 30,000 Bitcoin ATMs (BTMs) across U.S. counties to estimate the causal effect of anony-

mous payment access on crime, drug markets, overdose mortality, and household credit. Second, I use transaction flows on the blockchain to measure how the crypto system produces anonymity, stage by stage. Starting from the on-chain records of Bitcoin Depot, the largest U.S. operator, I trace where 1.4 million BTM customers send their funds, measure how much of the total anonymity each stage of the system supplies, and test whether the stages reinforce or replace one another. Third, I build a model in which regulation at each stage of the crypto transaction chain determines how traceable payments are, and traceability determines crime. The model combines the causal estimates with the on-chain measurements to ask where identity requirements earn the highest return. Throughout, the object of study is the anonymity margin within the payment technology, not crypto access itself. The headline findings are stark. BTM entry raises local crime by 6.9 percent and overdose mortality by 45 percent, and on chain the heaviest concentration of purchases sits just below the dollar threshold at which the machine starts asking for identification.

The BTM is a physical kiosk that converts cash into cryptocurrency. A user walks into a gas station, feeds in cash, and walks out with Bitcoin, often without presenting identification for amounts below state-specific thresholds. Each installation is therefore a geographically precise, discrete shock to the local availability of anonymous payment technology. Three features make the setting unusually well suited to measurement. The precise footprint of each machine allows its effects to be traced in county-level and tract-level administrative records on crime, mortality, and consumer credit. Operators set dollar thresholds below which no identification is required, and these thresholds differ across operators and states, so the degree of anonymity varies within the same technology. This variation separates the effect of anonymity from the effect of the machine. And because the largest operator is a publicly traded company filing audited financial disclosures with the SEC, the on-chain transaction record can be validated against an external ground truth.

Why would a cash-to-crypto kiosk cause crime? The logic follows from the two demand motives. If the demand a BTM meets is mostly protection, its arrival shifts payment habits

and raises user surplus, and real outcomes connected to crime should not move. If a substantial share is evasion, the machine lowers an input cost. A drug purchase, a laundering operation, or a fraud payout requires a payment that cannot be traced back to its participants, and cheaper anonymity raises the expected return to these activities. Four testable predictions follow. Crimes that depend on untraceable payment should rise more than crimes that do not. The effect should be stronger where identification requirements are more permissive. If the machine merely changes how existing crime is paid, physical quantities such as drug volumes and overdose deaths should stay flat, and if it creates new illicit activity, they should rise. And because the identity-verification threshold acts as a screen, users whose motive is to avoid identification should reveal themselves by transacting just below it. The data confirm all four.

BTM entry increases total crime by 6.9 percent. The pattern across offenses pins down anonymity as the operative channel and confirms the first prediction. Classifying all 51 NIBRS offense codes by their connection to anonymous payments yields a monotonic gradient, with high-relevance crimes rising 0.095 log points, medium-relevance crimes 0.052, and low-relevance crimes 0.028. A complementary decomposition, grounded in the payment system vulnerability framework of Kahn and Roberds [2009], maps each crime to the technological attribute it exploits. The channels tied to the financial function of the machine, namely irreversibility, pseudonymity, and cash-to-crypto conversion, all produce significant effects. The channel that reflects only the machine's physical presence does not. The degree of anonymity matters within the technology as well, confirming the second prediction. Dividing treated counties by the permissiveness of the identification requirements they face, the crime response is concentrated in the most permissive tercile and is essentially zero in the strictest. What drives the crime increase is what BTMs do, not where they are.

The central identification concern is that operators may place machines where crime is already rising. Entry models address this concern directly. Gas station density is the single strongest predictor of BTM placement, followed by the Hispanic population share

and population density, while local crime conditions have no independent predictive power once economic and demographic characteristics are controlled for. The baseline design is a staggered difference-in-differences with a not-yet-treated comparison group, and the pre-trends are flat. To push further, I construct a Bartik instrument that interacts each county's pre-period share of the national stock of gas stations with the national growth of the BTM industry, so that exposure is driven by the commercial infrastructure through which the expansion could be channeled rather than by local conditions. The first stage is strong, and the instrumental-variable estimate is somewhat larger than the uninstrumented one, which suggests the baseline results are if anything conservative.

These estimates could in principle reflect population sorting, policing attention, or the spatial reallocation of existing crime rather than new criminal activity. Geocoded incident data from the City of Chicago, 2012 to 2024, rule out each alternative. After BTM entry, tract demographics are unchanged and investigatory stops do not rise. And crime increases in tracts both near and far from the machine, with the larger increase closer to it, a pattern that reallocation cannot generate. The Chicago data also replicate the county-level results at the census-tract and police-beat level, with increases of 4 to 9 percent, confirming that the effects are not artifacts of aggregation.

The third prediction separates new crime from a change in how existing crime is paid. A drug deal that switches from cash to crypto leaves the quantity of drugs and the number of deaths unchanged, so these quantities rise only if real activity expands. They rise. Drug seizure weight increases 75 percent after BTM entry, with cocaine up 189 percent and heroin up 73 percent, a ranking that mirrors the substances most commonly traded on darknet markets. Overdose mortality increases by 45 percent, roughly five additional deaths per year in the median treated county. BTM entry creates illicit activity rather than merely moving it onto a new payment instrument. At the EPA value of a statistical life, the implied mortality cost is about \$60 million per county per year, against the \$20,000 to \$50,000 in revenue a typical machine generates. The incidence is regressive. Victimization rises

disproportionately among Hispanic residents and young adults, and the victim-to-offender relationship shifts toward stranger-perpetrated crime. Consumer credit delinquency rises with a lag, concentrated in low-income neighborhoods and fully mediated by the crime increase. The communities that financial innovation was supposed to include are the ones that absorb the harm.

The county-level estimates establish that anonymous payment access causes real harm, but the BTM is only one node in a larger system. Anonymity in a crypto payment is produced in three stages. Entry services convert cash into crypto and determine whether a person's identity is attached to an address. Transfer services such as mixers break the on-chain link between sender and receiver. Exit services such as unregulated exchanges convert crypto back into spendable value without restoring the identity trail. The second part of the paper measures this supply chain with transaction flows, asking how much of the total anonymity each stage contributes and whether the stages operate as complements or substitutes. The entry stage I measure directly at the transaction level. Starting from a proprietary cluster of seed addresses belonging to Bitcoin Depot, the largest operator in the BTM industry, I retrieve every Bitcoin transaction those addresses touch and recover 2.4 million cash-funded purchases by 1.4 million customers between November 2018 and May 2026. The reconstruction is externally validated, as the median transaction size computed on chain tracks the median disclosed in the operator's audited SEC filings with a correlation of 0.97. The median purchase is $179, and 86 percent of customer wallets receive exactly one payout and never appear again. Following the funds forward, I classify each customer's destinations into regulated exchanges, unregulated exchanges, peer-to-peer platforms, remittance services, mixers, gambling services, and darknet markets. The resulting flow decomposition shows what share of the anonymity users obtain comes free at the machine and what share is purchased downstream, and whether users who enter anonymously also buy more obfuscation later.

The same data reveal the demand for anonymity at the level of the individual transaction

and confirm the fourth prediction. Bitcoin Depot begins verifying identity at $250. The distribution of purchase sizes piles up just below that number, with two and a half times the mass that a smooth distribution would place there and a bunching estimate of $\hat{b} = 0.58$. The excess mass sits at $250 in dollar terms throughout a period in which the Bitcoin price moves from $30,000 to $67,000 to $19,000 and back, so it reflects the identity threshold rather than round-number habits. Structuring tests identify nearly 13,000 person-days on which a user splits one purchase into several sub-threshold visits, and most of these users spread the visits across multiple wallets to avoid leaving a pattern. A regulatory shift confirms that the threshold is what these behaviors target. In January 2024 California's Digital Financial Assets Law imposed a $1,000 daily cap that displaced the operator's $250 trigger as the binding margin, and the structuring rate fell by 43 percent at exactly that quarter. These users are paying real costs, in foregone purchase size, in fees, and in repeated trips, for nothing other than not showing identification. The individual-level bunching on chain and the county-level crime gradient off chain are two sides of the same margin.

The third part of the paper turns these measurements into a price for anonymity. An accounting identity divides the dollar harm a BTM county suffers by the dollar flow its kiosks process, both measured at the same county-quarter level, and scales the resulting damage rate by system-wide volume, the logic used to compute the social cost of carbon. Counted alone, those costs would imply that the optimal quantity of anonymity is zero. So the second step embeds the accounting in a model where anonymity is one product facing the two demand curves of the opening question, and no stage of the supply chain can screen one from the other. Tightening any identification rule deters crime and prices honest users out of privacy they value, so the optimal strength of each rule is interior. The measurements supply both sides of that trade-off, the harm per anonymous dollar from the causal estimates and the willingness to pay from the bunching evidence. The entry-stage answer is that free anonymity below a fixed threshold sits on the permissive side of the optimum.

The paper makes three contributions. First, it provides the first causal estimates of the

real effects of anonymity in digital payment systems. Theory establishes that anonymity has genuine value and that markets may undersupply it [Kahn et al., 2005, Garratt and van Oordt, 2021, Ahnert et al., 2025]. The cost side has rested on prevalence measurement that remains contested [Foley et al., 2019, Makarov and Schoar, 2021, Cong et al., 2025], or on evidence from cash, which is anonymous but cannot move value at scale [Wright et al., 2017, Rogoff, 2017]. The estimates here trace the full chain from payment design through crime, drug markets, mortality, and household credit. They mirror the finding that removing cash from the streets reduced crime [Wright et al., 2017]. They also sharpen recent evidence that crypto kiosks raise both financial participation and fraud [Cong et al., 2026], by isolating anonymity itself rather than cash or crypto adoption as the operative margin. Second, the paper links off-chain causal identification with on-chain forensic measurement at the level of the individual transaction, an approach not previously applied to the study of payment system design. It provides the first flow-based measurement of how the crypto payment system produces anonymity. The reconstruction of a major operator's full transaction record, validated against audited SEC filings, anchors the entry stage at the level of the individual transaction. Following the flows forward measures how much of the total anonymity each stage of the system supplies and whether the stages operate as complements or substitutes. This turns the blockchain from a source of prevalence statistics into a measurement instrument for the supply of and demand for anonymity. Third, the model connects the empirical results to policy. It prices payment anonymity from both sides, the harm an anonymous dollar carries and the price its buyers pay, and locates current regulation relative to the optimum stage by stage. Existing models of payment privacy price anonymity but do not embed the multi-stage structure through which it is actually produced.

The remainder of the paper is organized as follows. Section 2 provides institutional background on Bitcoin ATMs and the broader payment-system landscape. Section 3 describes the data sources. Section 4 estimates the causal effect of BTM entry on local crime. Section 5 addresses endogeneity in BTM placement with entry models and a Bartik instrument.

Section 6 uses Chicago as a within-city mechanism laboratory. Section 7 measures the real consequences in drug markets, overdose mortality, victim composition, and household credit. Section 8 measures the anonymity supply chain on chain. Section 9 scales the local estimates to a national cost and embeds the measurements in a model of optimal anonymity. Section 10 concludes.

2 Institutional Background

2.1 The Payment System Landscape

A payment instrument can be characterized along four attributes that jointly determine its economic and regulatory role. These are (i) whether the user's real-world *identity* is recorded, (ii) whether the system's *scale* is local-only or global, (iii) whether the medium is *digital* (electronically transferable) or physical, and (iv) whether each transaction is *traceable* on an auditable ledger.

Two endpoints have historically dominated. Physical cash is anonymous, local, non-digital, and not traceable. Bank deposits, payment cards, and electronic transfers are identity-recorded, global, digital, and fully traceable through regulated intermediaries. Table 1 summarizes the attribute bundles. The post-2001 anti-money-laundering framework and the post-2008 expansion of card and mobile networks accelerated substitution from cash toward identity-recorded electronic instruments. A substantial public-finance and central-bank literature argues that this substitution reduces criminal demand for currency [Rogoff, 2017, Sands, 2016], even as it raises distinct concerns about financial privacy, surveillance, and the political economy of intermediation [Kahn, 2018, Brunnermeier et al., 2019]. The key feature of this two-instrument world is that anonymity and scale could not be obtained together. A user who wanted anonymity was confined to cash and its local, physical limits, and a user who wanted global digital reach had to surrender identity. No payment instrument offered both, and it is precisely this missing combination that cryptocurrency supplies.

2.2 The Cryptocurrency Payment Landscape

Cryptocurrency reopened the trade-off by introducing a third combination of attributes. Identity is pseudonymous because addresses carry no legal names. Reach is global because a single network is usable from any internet connection. The medium is digital by construction. And transactions are fully traceable on the public ledger, but only to addresses, not to legal identities. In the years after Bitcoin's 2009 launch, an ecosystem of intermediaries developed that re-imposed identity at specific entry and exit points. Centralized exchanges (Coinbase, Binance, Kraken) operate under federal Bank Secrecy Act registration, require customer identification, and monitor transactions for reporting to the Financial Crimes Enforcement Network. Around this regulated core, a parallel ecosystem of decentralized exchanges, peer-to-peer markets, cross-chain bridges, mixing services, and physical kiosks provides routes that combine cryptocurrency's pseudonymity with reduced identity disclosure. The aggregate is best described as an *anonymity supply chain*. Each component either preserves an existing identity link or severs one of three links. Stage 1 (entry) severs the link between real-world identity and on-chain address. Stage 2 (mid-chain obfuscation) severs the link between sender and receiver in the on-chain transaction graph. Stage 3 (exit) severs the link between on-chain address and purpose of use. Figure A.4 maps the components against the three stages, identifies the named operators or services in each, and shows that Stage 1 is the unique gateway through which all anonymity-preserving flows must pass. Whoever controls this gateway observes the demand that feeds the entire chain, and identity rules imposed there constrain everything downstream.

2.3 Bitcoin ATMs as the On-Ramp Gateway

Bitcoin ATMs are physical kiosks that exchange cash for cryptocurrency with minimal identification requirements. Within the anonymity supply chain in Figure A.4, BTMs are one of the Stage 1 entry routes, where the link between real-world identity and on-chain address is severed, a point I return to in Section 8.

A BTM looks and operates like an ordinary ATM, typically standing in a gas station or convenience store, converting cash into cryptocurrency within minutes. Using one is simple. The customer scans a wallet QR code from a smartphone, verifies identity with a phone number (or a government ID above the operator's KYC threshold), inserts cash, and receives Bitcoin at the spot rate plus the operator's markup. Bitcoin dominates the menu, accepted at nearly every machine, with other cryptocurrencies offered at a minority of installations (Figure 2). The flow is almost entirely one-way. Bitcoin Depot's annual filings describe its kiosks as offering "one-way exchange of cash-to-Bitcoin," with sell-side machines below one percent of the fleet.

The first U.S. BTM was installed in Austin, Texas in 2014. The industry grew slowly through 2017, accelerated during the 2017–2018 cryptocurrency boom, contracted during the subsequent downturn, and then expanded explosively from 2020 onward, reaching more than 30,000 active machines by 2023 (Figure 3). Three industry features are relevant. First, BTM pricing far exceeds the 0.5 to 2 percent fees charged by centralized online exchanges such as Coinbase and Binance. Bitcoin Depot's 10-K filings disclose markups of around 15 percent over the spot exchange rate, plus a flat $3 fee per transaction. Users who pay this premium are buying immediate settlement, cash-based access without a bank account, physical proximity, and a degree of anonymity that online exchanges do not provide. Second, the identity-verification regime is set primarily at the operator level. All operators register with the Financial Crimes Enforcement Network as money services businesses and choose their own KYC thresholds, the dollar amounts below which a phone number suffices and above which a government ID is required. Operators must also comply with state law, and state requirements range from full identification of every customer to no BTM-specific rules during much of the sample period. The combination of operator thresholds and state law generates the cross-state variation in effective Stage 1 anonymity that the identification strategy below exploits. Third, the largest U.S. operator, Bitcoin Depot, is publicly traded (NYSE: BTM) and files quarterly audited disclosures. The on-chain reconstruction in Section 8 is built from

proprietary seed addresses and the public blockchain, and the audited disclosures provide an independent benchmark against which its accuracy can be checked.

2.4 Geographic Distribution and Host Venue Mix

The U.S. Bitcoin ATM industry combines broad geographic penetration with a concentrated host-venue mix. Penetration is the striking fact. By end-2023 approximately 31,000 active machines operated in every state and in 1,880 of 3,143 U.S. counties, and about 23 percent of ZIP codes hosted at least one kiosk (Figures 3 and A.1). Roughly one in four ZIP codes thus offers a walk-up conversion point from cash to Bitcoin. Cross-state density (machines per 100,000 residents) is led by Nevada, Texas, and Florida and is lowest in Hawaii and several Northeast states (Figure 4), so states differ in the intensity of coverage rather than its presence. At the venue level, the dominant host categories are gas stations and convenience stores, followed by smoke shops, check-cashing outlets, and laundromats (Figure 5). This co-location pattern reflects two simultaneous forces. On the supply side, operators seek venues with 24-hour foot traffic and existing cash-handling infrastructure, which gas stations and convenience stores already provide. On the demand side, customers who prefer anonymous transactions are more likely to frequent cash-oriented retail. This co-location is the basis of the identification strategy. Section 5 builds a Bartik-style instrument from the pre-2014 stock of host venues.

2.5 Bitcoin ATMs and Criminal Activity

Law enforcement agencies, the Federal Trade Commission, and state attorneys general have documented extensive criminal use of BTMs. The FBI's Internet Crime Complaint Center (IC3) identifies BTMs as a primary payment vector for romance scams, government impersonation fraud, and tech support schemes. The FTC reports that consumer losses from cryptocurrency-related fraud exceeded $1 billion in 2022, with BTMs as the most commonly cited payment method for in-person fraud encounters. State-level enforcement actions have

shut down BTM operations in multiple jurisdictions for Bank Secrecy Act violations, including failure to file suspicious activity reports and willful blindness to structuring (breaking large transactions into amounts below reporting thresholds).

The theoretical link between anonymous payment infrastructure and crime operates through several channels. Payment instrument design theory [Kahn and Roberds, 2009] identifies anonymity, irreversibility, and transaction speed as the key attributes that determine a payment system's vulnerability to criminal exploitation. BTMs score high on all three dimensions. Transactions are pseudonymous and often fully anonymous below KYC thresholds, settlements are irreversible once confirmed on the blockchain [Cong and He, 2019], and the conversion from cash to cryptocurrency occurs in minutes. The FATF's risk-based framework [FATF, 2015] assigns the highest vulnerability scores to financial products that combine anonymity with rapid, irrevocable settlement, precisely the attribute bundle that BTMs provide. In this respect the BTM extends to the cryptocurrency domain the role that large-denomination banknotes have historically played in the underground economy [Rogoff, 1998].

3 Data and Descriptive Statistics

This paper assembles nineteen data sources to trace the causal chain from BTM entry to local crime, drug markets, overdose mortality, and household finance, and to follow the resulting funds on the blockchain. This section describes each source in turn.

3.1 Bitcoin ATM Data

The Bitcoin ATM data are a kiosk-level record of the universe of U.S. Bitcoin ATMs, web-crawled from Coin ATM Radar, which has tracked operator listings continuously since 2013. The data cover 31,192 kiosks installed by 209 distinct operators in 1,880 of 3,143 U.S. counties. For each kiosk I observe the exact location (street address, latitude, and longitude,

geocoded to the census tract and county), the installation date, the operator's identity, the host venue (for example Circle K, Shell, or H-E-B, classified into venue categories), the supported cryptocurrencies, the transaction direction (buy-only or two-way), the opening hours, and the operator's KYC tiers, that is, the dollar limits at which phone verification or government identification become required. Operator disclosures in the 10-K filings of the publicly traded operators Bitcoin Depot and CoinFlip corroborate the aggregate machine counts, and FinCEN money services business registries corroborate the operator identities.

For the causal analysis I collapse the kiosk records into a county-year panel covering 2014 through 2023, and into tract-year and police-beat-year panels for the Chicago analysis. Each panel records, for every geographic unit and year, the count of active kiosks, the cumulative number of distinct operators, and a binary first-entry indicator that triggers the staggered difference-in-differences design.

3.2 Off-Chain Outcome and Covariate Data

Crime outcomes are sourced from the FBI National Incident-Based Reporting System (NIBRS), which records all Group A offenses at the incident level. I work with incident-level records for 2013 through 2023 and clean them following Dahl and DellaVigna [2009] and Huck [2024], who also use NIBRS as the outcome variable. Agencies with inconsistent reporting are dropped, specifically agency-years with gaps of more than ten consecutive days without a reported offense, agencies that do not report in all twelve months, and agencies with fewer than fifty reporting days in a year. I retain only agencies with county- or place-level jurisdiction, since incidents reported by state or federal agencies cannot be attributed to a specific county. The remaining incidents are aggregated to county-year counts by offense category, restricted to counties that report in at least six of the eleven sample years, with only observed county-years included rather than imputed zeros, so the panel is unbalanced. *Drug seizures* are measured from the property segments of the same NIBRS incident records, which report the quantity of drugs seized by drug type, aggregated to county-year

totals. *Overdose mortality* combines the NCHS model-based county-level drug-poisoning death rates for 2013 through 2021 [Rossen et al., 2020] with CDC provisional counts for 2022 and 2023. *Household credit* conditions are measured with the county-level Credit Insecurity Index published by the Federal Reserve Bank of New York, which is constructed from the New York Fed Consumer Credit Panel (Equifax) and available for 2018 through 2023. *Demographic and economic covariates* come from the Census ACS five-year tables, the BEA Regional Economic Accounts, the FDIC Summary of Deposits (bank branch density), the BLS QCEW (employment by NAICS), the FHFA county house price index, HUD fair market rents, and the Vera Institute Incarceration Trends dataset (county jail population). *Local market characteristics* relevant for BTM siting include gas station density (NAICS 447, from QCEW 2005–2023), which enters the entry selection models as a covariate. The Bartik-style instrument instead uses the 2005 county stock of gasoline stations with convenience stores (NAICS 447110, from the same QCEW source). Both are presented in Section 5.

Chicago microdata support the within-city mechanism analysis in Section 6. Geocoded incident records from the Chicago Data Portal (2001–2025) and Investigatory Stop Reports from the Chicago Police Department (2016–2024) are aggregated to tract-quarter and beat-quarter panels.

3.3 On-Chain Data

The on-chain analysis (Section 8) draws on three sources. *Operator seed addresses* are clusters of seed addresses attributed to several BTM operators, mainly Bitcoin Depot, from a proprietary source. *The Bitcoin blockchain* is sourced from Google BigQuery's `crypto_bitcoin` dataset, which mirrors the full public ledger as block, transaction, input, and output tables, so that for every transaction I observe the sending and receiving addresses, the BTC amounts, and the timestamp. Daily BTC/USD prices come from CoinMetrics. *Destination labels* are assembled from nine public and proprietary sources (labeled A–I in the analysis), covering 3.1 million unique Bitcoin addresses after deduplication. I extend this labeled universe with

the standard clustering heuristics from the blockchain-forensics literature, namely common-input-ownership clustering, change-address detection, and peeling-chain tracing [Meiklejohn et al., 2013].[2] For external validation of the on-chain reconstruction, I extract two disclosed series from Bitcoin Depot's quarterly 10-K and 10-Q filings on SEC EDGAR, the rolling 12-month median kiosk transaction size (2020Q1–2025Q4) and the audited cost-of-crypto-purchased line (2022Q2–2025Q4). Section 8.3 compares the reconstructed series against these disclosures.

3.4 Summary statistics

Table 2 reports summary statistics for the kiosk data and the county-year crime panel. Summary statistics for the on-chain dataset, including the address universe by category and the size distribution of customer payouts, are reported in Section 8.3 as part of the on-chain analysis.

4 The Effect of BTM Entry on Crime

4.1 Empirical Strategy

The enforcement record in Section 2.5 documents extensive criminal use of Bitcoin ATMs, but documented use does not establish that kiosks raise local crime. Offenses routed through a kiosk may be displaced from other payment channels, or they may have occurred regardless of the kiosk. I therefore ask whether total reported crime in a county changes when its first kiosk arrives, using a staggered difference-in-differences design. Treatment is the first installation of a kiosk in a county, an event that is staggered across counties between 2014 and 2023 and, once it occurs, is never reversed in the sample. The estimating equation is

[2] Common-input-ownership clustering assigns addresses to the same entity when they appear as inputs to the same transaction. Change-address detection attributes the one-time output that returns leftover funds to the sender. Peeling-chain tracing follows the characteristic sequence in which a large balance is paid out one transaction at a time.

the event-study specification

$$Y_{it} \;=\; \alpha_i + \lambda_t + \sum_{e=-5}^{5} \beta_e \, D_{it}^{e} + \varepsilon_{it}, \tag{1}$$

where Y_{it} is the crime outcome of county i in year t, α_i and λ_t are county and year fixed effects, and D_{it}^{e} is an indicator equal to one when year t falls e years after the county's first kiosk installation. The indicator for $e = -1$ is omitted, so each coefficient β_e measures the crime response e years from entry relative to the year before entry. Estimating Equation (1) by least squares is unreliable under staggered adoption, because already-treated counties enter the implicit comparison group and the resulting coefficients average heterogeneous cohort effects with weights that can be negative [Callaway and Sant'Anna, 2021]. I therefore estimate the dynamic coefficients and their post-entry average with the Callaway and Sant'Anna [2021] estimator, which computes an average treatment effect for each entry cohort and year by comparing that cohort with counties not yet treated, and then aggregates these group–time effects across cohorts. The event window is $[-5, +5]$ years around entry. Inference uses the multiplier bootstrap with standard errors clustered at the county level.

Identification rests on two assumptions. Parallel trends requires that, absent BTM entry, crime in treated counties would have evolved like crime in not-yet-treated counties. No anticipation requires that crime does not respond before the kiosk physically arrives, which is plausible because operators do not announce installations in advance and the entry analysis in Section 5.1 shows that siting follows slow-moving county fundamentals rather than short-run crime shocks. Both assumptions are testable in their pre-treatment implications. I report the full path of pre-treatment event-study coefficients and a joint Wald test of their significance with every estimate.

The estimation sample is the NIBRS county–year panel described in Section 3.2, with 10,488 observations covering 2013 through 2023. The primary dependent variable is log(count+1), which accommodates the skewed distribution of crime counts and yields estimates inter-

pretable as approximate percentage changes. The tables also report counts, rates per 10,000 residents, and log rates.

4.2 A Crime Taxonomy and Testable Predictions

An aggregate crime effect, on its own, cannot distinguish a causal mechanism from confounding. If BTM entry merely proxies for unobserved neighborhood change, the crime response should be diffuse and spread across offense types roughly in proportion to their baseline frequency. If the kiosk itself drives the response, the effect should concentrate in crimes that can exploit its technological attributes, the anonymity, irreversibility, and speed with which it converts cash into digital value [Kahn and Roberds, 2009, FATF, 2015]. I formalize this logic by classifying all 51 NIBRS offense codes along two dimensions. The classification uses only the statutory definition of each offense and these technological attributes, so the predictions below are fixed before any treatment effect is estimated.

The first dimension, payment relevance, sorts offenses by how much the crime relies on the kiosk's payment function, either to commit the offense or to move its proceeds. High-relevance offenses (14 codes) require the payment function itself, for example fraud schemes in which the victim deposits cash at a kiosk, ransomware demanding cryptocurrency, or drug transactions settled through anonymous conversion. Moderate-relevance offenses (17 codes) have an indirect connection, generating cash proceeds that could be laundered through a kiosk or occurring in the commercial environment around it, as with burglary or robbery near a host store. No-relevance offenses (20 codes) are driven by interpersonal violence, sexual motives, or regulatory violations with no connection to the payment function, such as homicide, sexual assault, and vandalism, and serve as a placebo. Figure 6 displays the full assignment. If the crime response operates through the payment function, treatment effects should decline monotonically with payment relevance, $ATT_{\text{High}} > ATT_{\text{Moderate}} > ATT_{\text{None}} \approx 0$.

The second dimension asks which specific technological trait a crime exploits, and yields

predictions about timing as well as magnitude. Irreversibility, the impossibility of chargebacks, is what makes wire fraud, extortion, and confidence schemes viable. These schemes need only a kiosk and a victim, so their response should emerge almost immediately after entry. Pseudonymity, the severed audit trail between cash and digital value, serves drug violations, ransomware, and identity theft. This is the theoretically central group, because pseudonymous conversion is the kiosk's defining innovation over wire transfers, prepaid cards, and money orders [Foley et al., 2019]. Its response should build with a delay as anonymity-dependent enterprises adapt to the new infrastructure. Cash-to-crypto conversion provides a post-crime laundering step for burglary, theft, and fencing. Its response should arrive last, after the property crime itself and the conversion of loot to cash. Cash attraction, finally, captures the possibility that the kiosk raises crime through its physical presence alone, by concentrating cash-carrying individuals at predictable locations. If the payment function rather than the physical presence drives the aggregate effect, this group and the 20-code placebo, which is shared with the first dimension, should both show no response. Figure 7 displays the assignment.

The two dimensions are complementary. The relevance gradient tests for a dose–response pattern but cannot identify which technological trait is operative, while the trait decomposition identifies operative channels but implies no intensity ordering, because its groups represent different mechanisms rather than decreasing dosage. Because the groups differ in size, from 6 codes to 20, I construct the group-level outcome as the average of logs,

$$\bar{Y}_{it}^{(k)} \;=\; \frac{1}{J_k} \sum_{j \in \mathcal{G}_k} \log\Big(\text{crime}_{ijt} + 1\Big), \tag{2}$$

where $\mathcal{G}_k$ is the set of offense codes in group k and J_k its size. This measure weights offense types equally, so that low-frequency but high-relevance crimes such as wire fraud are not drowned out by high-volume offenses such as larceny. I then estimate Equation (1) separately for each group, with log(total count + 1) as a secondary check.

4.3 Baseline Effect on Total Crime

Table 3 reports the aggregate effect and Figure 8 the event-study dynamics. According to the log specification, total reported crime rises by about 7 percent after BTM entry. Relative to the sample mean of 477 incidents per county and year, this corresponds to roughly 33 additional reported incidents annually in the average treated county. The identifying assumptions find support in the pre-treatment window. The five pre-entry coefficients fluctuate tightly around zero and the joint Wald test does not reject parallel trends ($p = 0.97$ in the log specification and $p \geq 0.86$ across the other dependent variables). The effect builds gradually, from near zero in the entry year to about 11 percent three years later. The growing profile is consistent with gradual criminal adoption of the kiosk rather than a one-time level shock, a pattern Section 4.5 examines directly.

4.4 The Payment-Relevance Gradient

Table 4 and Figure 9 report estimates of Equation (1) for the three relevance tiers. Crimes that depend directly on the kiosk's payment function, such as fraud schemes, extortion, and drug offenses, rise by about 10 percent after entry. Crimes with only an indirect connection to the payment function, such as burglary that generates loot to launder or robbery near a host store, rise by about 5 percent. Crimes unrelated to the payment function, such as homicide and sexual assault, change by less than 3 percent, an estimate indistinguishable from zero.

Two features of this pattern matter. First, the effect doubles at each step up the gradient. This is what one expects if the kiosk's payment function is the treatment. General neighborhood decline cannot produce it, because decline would raise all crime, not crime in proportion to its payment relevance. Second, the placebo is flat. Counties that received kiosks did not simply become more criminal. The dynamics tell the same story. The high-relevance group rises steadily for five years after entry, the moderate group follows the same path at half the height, and the placebo never moves. No tier moves before entry.

4.5 Decomposition by Technological Trait

Table 5 and Figure 10 report the same estimation for the five trait groups. Read side by side, the event studies show the three technological-trait groups switching on in sequence. Offenses that exploit irreversibility, such as wire fraud and extortion, jump within the first year of entry. Offenses that exploit pseudonymity, such as drug violations, ransomware, and identity theft, start later but grow fastest and end as the largest of the five groups. Offenses that use the kiosk to launder stolen proceeds, such as burglary and theft, respond with the longest delay. The two remaining groups never respond. Robbery and pickpocketing around the kiosk stay flat, and so does the placebo.

The sequence matches how long each criminal use takes to set up. A fraud scheme needs only a kiosk and a victim, so it exploits the kiosk at once. Drug sellers and other anonymity-dependent users adopt the kiosk gradually, so the pseudonymity response builds over time. Laundering comes at the end of a chain. The offender first steals, then fences the goods for cash, and only then converts the cash, so this response arrives last. The two null groups carry the identifying message. If kiosks raised crime by drawing cash-carrying customers to a store, robbery near the kiosk should have risen. It did not. What matters is what the kiosk does, not where it is.

The two decompositions also discipline each other. The moderate tier of the gradient mixes laundering-related property crime with proximity crime, and the trait decomposition shows that its entire response comes from the laundering part. Taken together, a confounder would have to mimic a very specific signature, an immediate rise in fraud, a delayed and accelerating rise in drug and ransomware offenses, a late rise in burglary, and no change in robbery, assault, homicide, or sexual offenses. No single county-level shock plausibly generates that pattern, which is why I read the decomposition as evidence that the aggregate effect is causal and runs through the kiosk's payment function.

4.6 Heterogeneity by Identification Requirements

The taxonomy infers the anonymity mechanism from which crimes respond. A sharper test uses variation in how much anonymity the kiosks themselves provide. Operators set a first-tier threshold below which a customer can buy without presenting government identification, and these thresholds differ across operators and states. I measure county-level anonymity as the average first-tier threshold across the kiosks a county receives and split treated counties into terciles of this measure.

Table 6 and Figure 11 report the result. In the most permissive tercile, where the average kiosk lets a customer convert more than $1,126 without identification, crime rises steadily after entry and exceeds 20 percent above the pre-entry level by the fourth year. Averaged over the post-entry window, the response is 0.187 log points. The middle tercile averages 0.071, roughly a third of the permissive response. The strictest tercile averages 0.003, essentially no response. Pre-treatment coefficients are flat in the strict and middle terciles. In the permissive tercile they sit slightly below zero, so if anything the estimate understates the post-entry increase. The dose-response panel shows the same ordering. Within every kiosk-count bin the effect rises with permissiveness, and counties holding more than ten kiosks in the most permissive tercile report crime about 27 log points above counties with none.

The machines are physically identical across terciles. What differs is the amount of cash a customer can convert without leaving a name. If the kiosk raised crime through its physical presence or its cash flows, the identification threshold should not matter. It matters decisively, which confirms the prediction that the effect should be stronger where identification requirements are more permissive.

As a final check, I weight the 51 offense types by their empirical association with BTMs rather than by the discrete taxonomy. Following the text-based methodology of Baker et al. [2016], I measure each offense type's co-mention frequency with Bitcoin ATMs in a corpus of news articles and use the normalized frequencies as weights in a crime index. The ATT on the log weighted index is comparable in magnitude to the unweighted total but estimated

with notably tighter precision, as concentrating power on BTM-relevant offenses strips out noise from unrelated crime. Appendix A details the construction of the weights and reports the estimates, with the weight distribution in Figure A.2 and the event-study dynamics in Figure A.3.

In sum, BTM entry raises reported crime by about 7 percent, and the increase decomposes along two pre-specified dimensions into a pattern specific to the kiosk's technological attributes. Crimes respond in proportion to their dependence on the payment function and on timescales that match the trait each exploits, while physical-proximity crimes and placebo offenses do not respond at all. The response also scales with the identification thresholds the kiosks apply, which ties the effect to the degree of anonymity within the technology itself. The next section addresses the remaining concern that entry timing itself may be endogenous.

5 Instrumental Variable Analysis

5.1 County-Level Determinants of BTM Entry

The crime estimates in Section 4 treat the arrival of the first kiosk as the treatment. A natural concern is that operators choose where to place kiosks, so I begin by asking which county characteristics predict where and when BTMs appear. The answer serves two purposes. It shows the dimensions along which treated and untreated counties differ, which disciplines the comparisons behind the difference-in-differences design, and it tests directly whether operators select on crime. It also identifies the supply-side force, the pre-existing stock of host venues, on which the instrument in Section 5.2 is built.

I estimate two complementary models. The cross-sectional logit takes the form

$$\Pr(\text{BTM}_i = 1 \mid \mathbf{X}_{i,2014}) = \Lambda\Big(\mathbf{X}'_{i,2014}\boldsymbol{\beta} + \gamma_s\Big), \tag{3}$$

where BTM_i equals one if county i installs at least one BTM during 2015–2023, $\mathbf{X}_{i,2014}$ is a vector of standardized covariates measured in 2014, the year before the sample opens, γ_s is a state fixed effect, and $\Lambda(\cdot)$ is the logistic CDF. I report average marginal effects computed by the delta method. The second model is a discrete-time hazard that asks not whether but when a county receives its first kiosk. The hazard of first entry is

$$h_{it} = \Pr(T_i = t \mid T_i \geq t) = 1 - \exp\Big[-\exp\Big(\alpha_t + \mathbf{X}'_{i,t-1}\boldsymbol{\beta} + \gamma_s\Big)\Big], \tag{4}$$

where T_i is the year in which county i installs its first BTM, so each county contributes observations until entry and exits the risk set thereafter. The complementary log-log link makes the model the discrete-time counterpart of a continuous-time proportional hazards model [Jenkins, 1995]. The year effects α_t trace the baseline hazard of the national adoption wave and $\mathbf{X}_{i,t-1}$ contains the same covariates as the logit, now time-varying and lagged one year. Agreement between the two models confirms that the cross-sectional patterns are not artifacts of collapsing a nine-year adoption process into a single binary outcome. In both models all continuous covariates are standardized to mean zero and unit variance and standard errors are clustered at the state level.

The covariates enter in five blocks. Demand-side composition includes population density and the Black and Hispanic population shares. Economic conditions include median household income, the poverty rate, unemployment, the Gini coefficient, and GDP growth. Financial infrastructure includes bank branch density and average deposits per branch. Local market characteristics include house price growth, fair market rent as a proxy for operating costs, and gas station density, the dominant host venue category for BTM kiosks. The final block adds a local crime proxy, the county jail population rate from the Vera Institute's Incarceration Trends data. The jail rate is available for 96 percent of the counties in the estimation sample. A crime rate built from NIBRS would be available for only 16 percent, so the jail rate avoids that severe selection.

Table 7 reports the logit and Figure 12 plots the average marginal effects. Three variables dominate. A one-standard-deviation increase in gas station density raises the probability of ever receiving a BTM by 22 percentage points, the largest effect in the model. The same increase in the Hispanic population share raises the probability by 13 percentage points. One interpretation, suggested by industry coverage, is demand for cross-border remittances, for which BTMs offer lower fees than traditional money transmitters.[3] Population density adds 9 percentage points. Among the remaining covariates, income inequality and median income enter positively, fair market rent enters negatively, so operators favor unequal, denser markets but avoid expensive locations, and poverty and unemployment carry no independent weight. The crime proxy is the central null result. Conditional on the other covariates, the jail population rate does not predict entry, and all other coefficients are unchanged by its inclusion.

The Black population share illustrates why the gas station control matters. In specifications with demographics alone, counties with larger Black populations appear less likely to receive kiosks. Once gas station density enters the model, the coefficient shrinks sharply and loses significance in the logit, though a smaller negative estimate persists in some hazard specifications. The reason is that the two variables move together. Counties with larger Black populations have systematically fewer gas stations and convenience stores, with a partial correlation of -0.23 conditional on the other covariates. Operators are not avoiding these counties. They are following host venues, and host venues are scarcer there.

The hazard estimates in Table 8 tell the same story about timing. Gas station density, the Hispanic share, and population density remain the dominant predictors of when a county receives its first kiosk. Bank branch density is negative and significant, so kiosks arrive earlier where traditional banking access is thinner, consistent with BTMs partially substituting for bank infrastructure. The jail population rate is again insignificant once state fixed effects

[3]Industry reports indicate that a substantial share of BTM transactions involve remittances to Latin America, where recipients convert Bitcoin to local currency at lower fees than traditional money transmitters [e.g., Coin ATM Radar, 2023].

are included, and Figure 13 plots the full coefficient path of the hazard specification with the crime proxy. Crime does not predict the timing of entry any more than it predicts its location.

The conclusion is unambiguous. Operators follow gas stations and demographics, not crime. Selection on demographics does not by itself threaten the crime estimates. The difference-in-differences design never requires treated and untreated counties to look alike, only to trend alike, and county fixed effects absorb every level difference in population, income, or retail structure. A problem arises only if the fundamentals that attract operators also predict changes in crime around the time of entry. Two results speak against that. The fundamentals that predict entry are slow-moving stocks rather than shocks, which supports the no-anticipation assumption in Section 4.1, and the pre-treatment coefficients in Section 4.3 show no differential crime movement in the years before the kiosk arrives. The entry analysis also hands the next subsection its instrument. The strongest predictor of entry is a stock of host venues that was already in place in 2005, nine years before the first U.S. kiosk, and that operators take as given when they expand.

5.2 Instrumental Variables

Section 4.3 shows that reported crime rises by about 7 percent after a county's first kiosk arrives, with flat pre-trends and an increase that concentrates in offenses exploiting the kiosk's technology. That design already rules out the leading threats to a staggered comparison. The Callaway–Sant'Anna estimator avoids the negative-weight bias of two-way fixed effects, county fixed effects absorb every level difference between treated and untreated counties, the flat pre-treatment coefficients rule out pre-existing differential trends, and the unresponsive placebo offenses rule out a diffuse county shock that raises all crime. Two concerns survive. The first is selection on future shocks. Operators choose where and when to enter, and pre-treatment tests have limited power against violations of parallel trends that begin only at entry [Roth, 2022]. If operators entered counties just before crime would have risen for other

reasons, the difference-in-differences estimate would attribute that rise to the kiosk. The second concern is that the staggered design captures only the first installation. Expansion after entry is a sequence of further operator choices that the entry analysis cannot discipline, because counties exit the hazard risk set at first installation, and the average treated county ends the sample with 11.8 kiosks. If the unobserved forces that attract those additional machines also raise crime, a least-squares regression of crime on the kiosk count is biased upward. If operators instead stop adding machines where vandalism risk or regulatory scrutiny is high, the bias runs downward.

Both concerns disappear if the number of kiosks a county receives can be moved by a force unrelated to the county's own crime path. The entry analysis identifies such a force. Operators follow host venues, and the stock of host venues was in place long before Bitcoin existed. I therefore construct a shift-share instrument [Bartik, 1991],

$$Z_{it} \;=\; s_i \;\times\; g_{-i,t}, \tag{5}$$

where s_i is county i's share of the national stock of gas station and convenience store establishments in 2005 (NAICS 447110, from the Quarterly Census of Employment and Wages) and $g_{-i,t}$ is the leave-one-out national cumulative BTM count in year t. I compute $g_{-i,t}$ from the full universe of 31,192 installations across 1,880 treated counties rather than from the crime sample alone, which would understate the national expansion and reintroduce sample-driven correlation with the error term.

The economic logic of the instrument is the contracting structure of the industry. BTM operators are franchise-like networks that contract with retail chains rather than with individual stores, as when Bitcoin Depot agreed in 2021 to place kiosks across roughly nine hundred Circle K locations.[4] Deals of this form channel the national expansion into counties in proportion to their pre-existing stock of suitable venues. The instrument thus replaces

[4]Bitcoin Depot, press release "Bitcoin Depot Announces Long-Term Partnership with Circle K," available at https://ir.bitcoindepot.com/news-events/press-releases/detail/11/bitcoin-depot-announces-long-term-partnership-with-circle.

both operator choices with forces outside the county. Where the predicted kiosks land is set by the 2005 retail geography, measured four years before Bitcoin was invented and nine years before the first U.S. kiosk, so it cannot reflect any anticipation of cryptocurrency infrastructure. How many arrive, and when, is set by the national expansion of the industry, which no single county's crime shocks can move.

Identification follows the share-exogeneity framework of Goldsmith-Pinkham et al. [2020] and requires that the 2005 share is orthogonal to subsequent county-level crime shocks, conditional on fixed effects. The share itself is a slow-moving feature of retail geography shaped by highway networks and commuting patterns, not by crime fluctuations a decade later. It is worth stating exactly what could violate the exclusion restriction in this design. County fixed effects absorb every level difference between counties, and year effects absorb the national expansion itself. The instrument therefore fails only if counties with many gas stations in 2005 followed different crime trends during the rollout years for reasons unrelated to BTMs. Two concrete stories fit that description. Gas stations cluster along highways, and drug markets spread along highways during the sample period. Gas-station-heavy retail economies may also have declined as commerce moved online. The specifications absorb the most likely sources of such trends, with log population interacted with year capturing trends that follow county size and state-by-year fixed effects capturing state drug policy, policing, and economic cycles. More importantly, both stories predict that high-share counties should show crime changes even before kiosks arrive. They do not. In untabulated balance tests the instrument does not predict pre-treatment changes in crime or demographics, and the pre-treatment coefficients of the 2SLS event study below are flat.

Table 9 reports the first stage and Figure 14 plots the coefficient across specifications. The instrument is a strong predictor of the cumulative BTM count, with a Kleibergen–Paap F statistic of 98.5, roughly ten times the conventional weak-instrument threshold [Stock and Yogo, 2005]. The coefficient is also stable. If the instrument worked through correlated local trends rather than through hosting capacity, adding state-by-year fixed effects would

attenuate the first stage, yet the coefficient is essentially unchanged.

The reduced form in Table 10 regresses crime directly on the instrument and requires only instrument exogeneity, not the full exclusion restriction. The instrument significantly predicts crime in counts, rates, and logs. Counties with more gas stations in 2005 experience larger crime increases whenever the national BTM network expands. This is the simplest causal statement the design delivers, and it does not depend on the functional form of the first stage.

Table 11 presents OLS and 2SLS estimates of the structural equation

$$Y_{it} \;=\; \beta \cdot \mathrm{BTM}_{it} + \alpha_i + \gamma_t + X_{it}'\delta + \varepsilon_{it}, \tag{6}$$

where BTM_{it} is the cumulative number of kiosks in county i in year t and β is the effect of one additional kiosk. In the log specification one additional kiosk raises crime by about 0.5 percent. The per-kiosk effect is small, but counties accumulate kiosks, and at the average treated county's 11.8 machines the implied total effect is about 6 percent. The coefficient is positive in every specification and shrinks when crime is allowed to trend with county size, a demanding control given that the instrument assigns more kiosks to larger retail markets.

The 2SLS estimates are similar to or moderately larger than their OLS counterparts. Two standard forces explain why instrumenting does not reduce the estimate. Recorded kiosk counts overstate true operational machine-years whenever machines go offline or are relocated without being de-registered, which attenuates OLS. And 2SLS recovers the effect for complier counties, those whose kiosk count responds most to hosting capacity. These are counties with extensive highway retail, where a kiosk serves a large cash-intensive catchment area, so a local average treatment effect above the population average is plausible.

Figure 15 provides the sharpest test of the identifying assumption by comparing OLS and 2SLS event studies. The specification interacts event-time indicators with the county's eventual treatment intensity $d_i = \log(1+\max_t \mathrm{BTM}_{it})$ and instruments each interaction with

the corresponding instrument interaction,

$$Y_{it} = \sum_{\tau=-5}^{5} \beta_\tau \cdot \mathbf{1}[\text{rel_time}_{it} = \tau] \cdot d_i + \alpha_i + \gamma_t + \varepsilon_{it}, \tag{7}$$

with $\tau = -1$ as the reference period. Using the eventual dose rather than the contemporaneous count avoids conflating cumulative kiosk growth, which rises mechanically with event time, with the dynamic treatment path. The contrast between the two estimators speaks directly to the selection concern that opened this section. The OLS pre-treatment coefficients are negative and significant, so selection is real. Counties that eventually accumulate many kiosks were on declining crime paths before entry, and least squares confounds that selection with the treatment. The 2SLS pre-treatment coefficients are flat at zero, which is exactly what the design requires, since a confounded instrument would inherit the differential pre-trends. After entry the 2SLS effect appears immediately and then plateaus, the same early-onset signature as the irreversibility channel in Section 4.5, and consistent with a level shift in crime rather than a diverging trend.

Table 12 subjects the baseline estimate to state-level clustering, dropping the largest five percent of counties, restricting to counties above the median population, and restricting the sample to 2015 through 2023. The estimate remains positive and significant in every variation. It is essentially unchanged by state-level clustering and by the later sample window, and it nearly quadruples when the largest counties are dropped, so the per-kiosk effect is substantially larger in smaller counties. The heterogeneity is itself informative about reverse causality. If operators targeted high-crime locations, the effect should be largest in big-city counties where they have the most choice. The data point the opposite way, toward a single kiosk mattering more in a thin local market.

The section closes where it began, with the two concerns that motivated the instrument. Selection on future shocks would have left pre-trends in the instrumented event study, and there are none. Endogenous intensity is replaced by variation from 2005 retail geography

and national industry growth, and the dose–response estimate survives every control set and sample cut. What remains is the comparison across designs. The staggered estimator rests on parallel trends across entry cohorts. The instrument rests on the exogeneity of a retail share fixed four years before Bitcoin existed. Neither assumption implies the other, so a single confounder would have to defeat both designs through unrelated channels. Yet the two designs return the same answer, a cumulative effect of about 6 percent from the instrument next to the 7 percent average treatment effect of first entry. When variation from operator timing and variation from pre-Bitcoin retail geography deliver the same estimate, the explanation left standing is the treatment itself. I conclude that BTM entry causes the increase in local crime.

6 Mechanism Analysis

The county-level difference-in-differences in Section 4 and the Bartik instrument in Section 5 both identify a positive causal effect of Bitcoin ATM entry on local crime. This section asks how that effect operates. Three alternative mechanisms could generate a county-level treatment effect without any net increase in offenses. Higher-crime populations could move into BTM neighborhoods. Police could direct more attention to treated areas and detect crimes that would otherwise go unrecorded. Existing crime could relocate toward the kiosks from nearby blocks. I use Chicago as a single-city laboratory in which the geographic resolution of the data permits a sharper test of the causal claim and a direct test of each mechanism. The evidence rejects all three.

6.1 Higher-Granularity Test: Within-City Evidence

Chicago suits this purpose for two reasons. The Chicago Police Department publishes every reported incident with exact coordinates, which permits analysis at the census tract–quarter and police beat–quarter level, units two to three orders of magnitude smaller than counties.

And the city's BTM deployment is dense and staggered. Kiosks reach 252 of 857 census tracts over 2015–2024, creating rich variation in treatment timing within a single policing regime, judicial system, and labor market. I extract Chicago installations from the master BTM panel (Section 3), geocode each machine to its tract and beat, and date treatment from the quarter of first entry. The 605 never-treated tracts serve as the comparison group. Estimation follows the same heterogeneity-robust difference-in-differences framework as the county analysis, with unit and quarter fixed effects and standard errors clustered at the unit level. For the quarterly within-city panels I report the imputation estimator of Borusyak et al. [2024], which remains efficient when the panel is long and many small units carry few incidents per period, alongside log and Poisson pseudo-maximum-likelihood (PPML) specifications.

Table 13 reports the tract-level estimates. The log specification yields a coefficient of 0.042 and the PPML semi-elasticity is 0.055, increases of about 4 to 6 percent, both significant at the 1 percent level and stable when I exclude tracts with zero pre-treatment crime. Figure 16 shows the dynamics. The increase is concentrated in battery, narcotics, assault, and deceptive practice, the same payment-relevant categories that drive the national result, while theft and motor vehicle theft decline. Pre-treatment coefficients center on zero with no trend, and post-treatment effects emerge within two quarters and reach roughly three to five additional incidents per tract–quarter by six quarters after entry.

The police beat results sharpen the picture further. Beats are the CPD's smallest patrol assignment unit, 277 areas averaging about 3.5 km^2, of which 156 receive a machine. Table 14 shows a log coefficient of 0.163 and a PPML semi-elasticity of 0.083, both significant at the 1 percent level, with the level specification pointing to about 8.5 additional incidents per beat–quarter at marginal significance. The beat estimates exceed the tract estimates, as expected when the effect is concentrated in small areas around the kiosk and dilutes as the spatial unit grows. The event study in Figure 17 shows flat pre-trends and post-treatment coefficients rising to roughly ten to fifteen additional incidents per beat–quarter by the sixth

quarter.

That the result replicates across two independent geographic partitions of the same city substantially narrows the set of viable confounders. Any omitted-variable story would need to operate at the sub-neighborhood level, move tracts and beats in the same way, and produce a monotonically growing post-treatment path with no pre-trend. The remaining candidates are the three mechanisms above, which I now test in turn.

6.2 Mechanism Tests

The within-city replication establishes that recorded crime rises where kiosks enter, but it does not by itself establish that the machines cause new criminal activity. This subsection develops the three alternative stories that could produce the same pattern without any net increase in offenses, derives the testable prediction of each, and confronts each prediction with data that only the Chicago setting provides. The evidence rules out all three.

6.2.1 Population Sorting

The first alternative story is that BTM entry changes who lives nearby rather than what anyone does. A kiosk that converts cash into crypto with little identification is a convenient amenity for people already engaged in illicit activity, so entry could draw such individuals to relocate into the neighborhood. Recorded crime would then rise because the resident pool shifted, not because the machine enabled new offenses. The testable implication is direct. The composition of the local population must change after entry, and it must change early enough to drive the crime response.

I test this prediction with tract-level demographics from the American Community Survey five-year estimates, 2015–2023, estimating the same difference-in-differences design on six composition outcomes. These are total population, log population, and the shares of white, Black, Hispanic, and male residents. The event window spans four years before to three years after entry to match the annual frequency of the ACS.

The data reject this story. Table 15 shows that none of the six outcomes is significant at the 5 percent level. Total population moves by about 30 residents against a tract mean of roughly 4,000, under one percent of the base. The white share shows a marginal decline of about one percentage point, and the remaining coefficients are economically negligible. The event studies in Figure 18 hover near zero before and after entry with no systematic drift. Timing rejects the story independently. The crime response appears within two quarters of entry, while residential turnover in urban housing markets operates over years [Sampson, 2012]. A sorting story would also raise crime across all categories as the population changes, not produce the concentration in payment-relevant offenses that the data show. Compositional change cannot account for the crime response.

6.2.2 Police Detection

The second alternative story is that BTM entry makes police look harder, so more of the crime that was always there gets recorded. Regulators have repeatedly flagged kiosks as a money-laundering vector, and officers could respond by patrolling BTM neighborhoods more intensively. The extra scrutiny would convert previously undetected offenses into recorded ones, and the treatment effect would measure detection rather than crime. The testable implication is again direct. Proactive policing activity must rise in treated areas after entry.

I measure proactive policing with Chicago's Investigatory Stop Reports, which log every investigatory stop by CPD officers from 2016 to 2024 with date, beat, and the demographics of the stopped person. I aggregate stops to the beat–quarter level and estimate the same event-study design as for crime.

The data reject this story as well. Figure 19 shows the event study for total stops. Every post-treatment coefficient is statistically indistinguishable from zero, and the point estimates dip slightly below zero in the first quarters after entry before returning toward zero. Because the aggregate null could mask offsetting group-level shifts, Figure 20 decomposes stops by the race and sex of the stopped person. No demographic group experiences a rise in stops

after entry. Stops of Black subjects, the group most often invoked in detection narratives, decline modestly in the quarters immediately after entry and then recover toward zero, and stops of white, Hispanic, and male subjects stay flat. If anything, proactive policing falls slightly while recorded crime rises, the opposite of what a detection mechanism requires. The composition of the crime response points the same way. The increase is concentrated in battery, narcotics, and deceptive practice, offense types that enter the data predominantly through victim and witness reports rather than street stops.

6.2.3 Spatial Reallocation

The third alternative story is that BTM entry moves crime around rather than creating it. A kiosk that converts cash on the spot makes its corner a more attractive place to do illicit business than the next one, so existing offenses could simply relocate from nearby blocks toward the machine. Crime in treated tracts rises, crime next door falls, and the net citywide effect is zero. The distinction matters for welfare. Displacement redistributes harm rather than creating it, and enforcement around kiosk sites could in principle offset the problem entirely. The story implies two testable predictions. Crime in areas adjacent to treated tracts must fall after entry, and the decline must be steeper closer to the kiosks.

I test the first prediction with a paired difference-in-differences that estimates the treatment effect against two separate control groups. Near controls are untreated tracts that share a boundary with a treated tract. Far controls are untreated tracts more than 3 km from any treated tract, matched to treated tracts by coarsened exact matching on pre-treatment crime quintiles. Both control groups inherit the treatment timing of their paired treated tract. If treated tracts pull crime from their neighbors, near controls have artificially depressed crime, so the treatment effect measured against near controls should exceed the effect measured against far controls. Figure 21 shows the opposite. Both control groups exhibit parallel pre-trends, and post-treatment the far-control ATT (0.060^{***}) exceeds the near-control ATT (0.031^{*}). The compressed gap against near controls indicates that crime

in adjacent tracts rises along with crime in treated tracts, a pattern of mild outward spillover rather than siphoning.

The second prediction concerns the distance gradient, which I test with continuous variation in proximity. For each tract I construct a time-varying inverse-distance-weighted exposure index, $\text{Exposure}_{it} = \sum_j d_{ij}^{-2}$, summing over all machines j active in quarter t, where d_{ij} is the distance from tract i to machine j. I interact the standardized index with event-time indicators,

$$\ln(\text{crime}_{it} + 1) = \sum_{\tau} \beta_\tau \cdot \mathbf{1}[t - g_i = \tau] \cdot \text{Exposure}_i + \alpha_i + \gamma_t + \varepsilon_{it},$$

where g_i is a tract's own treatment quarter or, for untreated tracts, the treatment quarter of the nearest treated neighbor. Figure 22 maps the exposure index across the city and shows the cross-sectional variation that identifies the β_τ. Under displacement, post-treatment coefficients should be negative because high-exposure untreated tracts lose crime to nearby kiosks. Positive coefficients indicate spillover. The evidence again rules out displacement. Figure 23 plots the estimates. Pre-treatment coefficients cluster around zero. Post-treatment, the marginal effect of exposure drifts upward and reaches significance at $t + 7$ (0.036, 95 percent confidence interval [0.010, 0.063]). One additional standard deviation of proximity to active machines raises crime by about 4 percent seven quarters after entry. The direction again contradicts displacement and matches the paired design.

The Chicago evidence delivers two conclusions. First, the county-level crime effect survives, and in some specifications strengthens, at the tract and beat level. The increase of roughly 4 to 9 percent across units and functional forms is not an artifact of geographic aggregation, county-level confounders, or NIBRS reporting conventions. Second, the three mechanism tests narrow the interpretation to new criminal activity. Demographic composition does not move. Investigatory stops do not rise, and for some groups they fall. Both displacement tests point in the wrong direction for reallocation and instead reveal mild out-

ward spillover, which implies that the county estimates, if anything, understate the total effect.

7 Downstream Welfare Consequences

The preceding sections establish that Bitcoin ATM entry causes an economically meaningful increase in criminal activity and that the increase reflects new offenses rather than sorting, detection, or displacement. This section asks what those offenses cost. I trace the consequences through three measurable channels. Drug seizure weight captures whether the crime response carries a real expansion of drug supply. Overdose mortality captures whether the supply expansion harms public health. Victim composition captures who bears these costs. A final subsection develops a simple model that links crime losses to household credit, the channel through which the harm propagates into financial outcomes. Together these analyses move from incident counts to the human costs that determine the social value of anonymous payment infrastructure.

7.1 Beyond Incident Counts: Drug Seizures and Market Severity

Incident counts are an incomplete measure of harm. A county that records the same number of drug offenses before and after BTM entry may still experience a substantial deterioration in public safety if the volume of drugs flowing through those offenses grows. The DEA evaluates local drug threats by seizure weight rather than arrest counts for precisely this reason.[5] I therefore ask whether the increase in drug incidents translates into a measurable expansion of drug supply, the necessary condition for the public health consequences examined next. Table 16 reports summary statistics for the seizure and overdose outcomes used throughout this section.

A positive seizure response admits three interpretations. Law enforcement may learn

[5]DEA, *National Drug Threat Assessment* (annual). See also Caulkins et al. [2005].

to surveil BTM locations and detect more drugs without any change in underlying supply. Dealers may migrate to BTM-based laundering and concentrate pre-existing flows at newly observable points. Or BTMs may lower the friction of anonymous value transfer and let suppliers expand operations, a genuine increase in total supply. Only the third interpretation implies real market expansion. The three make different predictions. Detection predicts a jump on both the extensive margin (any seizure) and the intensive margin (seizure weight). Displacement predicts offsetting declines in neighboring areas. Scale expansion predicts a sustained, accelerating increase in weight accompanied by downstream harm.

The decomposition by substance sharpens these tests into a specific prediction. If BTMs raise crime by lowering the cost of moving criminal proceeds anonymously, the supply response should be largest for the drugs whose distribution depends most on anonymous payment. Cocaine and heroin carry the highest legal penalties and dominate darknet trade [Foley et al., 2019, Böhme et al., 2015], so their distribution networks gain the most from a cheap cash-to-crypto channel. Hard drugs should therefore show the strongest supply response, including expansion into counties they did not previously reach.

Table 17 reports Callaway–Sant'Anna estimates on both margins, and the results align with the prediction. On the intensive margin (Panel B), the ATT on total seizure weight per capita is 0.562 in inverse-hyperbolic-sine units, roughly 75 percent more weight seized in treated counties, with pre-trend p-values above 0.96 (Figure 24). Cocaine shows the largest response (1.060, about 189 percent). Heroin rises by about 73 percent and methamphetamine by about 47 percent. Cannabis also rises strongly (about 153 percent), and the next subsection traces that response to its own mechanism, the cash economy created by the cannabis banking exclusion.

The extensive margin (Panel A) completes the alignment and rules out the detection interpretation. The probability of recording any drug seizure does not rise after entry. It falls slightly, by 4.1 percentage points, significant at the 5 percent level. If BTM entry merely drew enforcement attention, both margins should rise together. Against this aggregate de-

cline, only cocaine and heroin, the two darknet-linked hard drugs, expand on the extensive margin (13.4 and 15.9 percentage points, both $p < 0.01$), while cannabis and methamphetamine show no change. BTM entry opens hard-drug distribution into counties that previously recorded no seizures, exactly where the anonymity-dependence prediction says the new infrastructure should matter most. For the other substances, the effect operates purely through intensification of existing markets. Figure 25 summarizes the decomposition.

As a functional-form check, I re-estimate the weight effects by Poisson QMLE following Chen and Roth [2024] (Table 18). The proportional effects are similar in sign and ordering, with overall seizure weight rising by about 57 percent, and the event-study path of the Poisson estimates in Figure 26 mirrors the inverse-hyperbolic-sine dynamics.

The evidence is most consistent with scale expansion. Detection is ruled out by the declining aggregate extensive margin. Displacement is hard to reconcile with the magnitudes, since a 75 percent overall increase and a near-tripling for cocaine would require implausibly large spatial substitution if existing supply were merely reallocating, and the Chicago tests in Section 6.2.3 reject displacement directly at fine geography. The temporal profile, with effects emerging gradually and accelerating, matches infrastructure buildup rather than an enforcement shock.

The geocoded Chicago data permit one further decomposition that county seizure data cannot. Seizure weight measures what law enforcement recovers but does not separate supply-chain offenses from end-user possession. I classify each Chicago narcotics incident by its description field, coding possession as demand-side and manufacture, delivery, and trafficking as supply-side, and estimate tract-level regressions separately for each. Table 19 shows a sharp asymmetry. Supply-side offenses rise strongly, with a log coefficient of 0.101 and a Poisson semi-elasticity of 0.429, while demand-side offenses show no significant change in any specification. If BTMs serve as financial infrastructure for distribution networks, the technology should shift the supply curve outward and load the response onto supply-chain offenses, which is what the data show. The within-city evidence thus corroborates the

county-level conclusion that BTM entry expands drug market scale.

7.2 Cannabis Banking Exclusion

The results so far pool all counties together. State cannabis policy creates variation that lets me test the drug-payment channel more directly. The argument runs in three steps. First, cannabis remains a Schedule I substance under federal law even in states that have legalized recreational sales, so banks, which are federally regulated, refuse deposits from cannabis businesses. Second, dispensaries in legal states therefore operate almost entirely in cash. Industry estimates put 70–75 percent of cannabis transactions in cash, roughly $20 billion per year that cannot enter a bank account. Third, a BTM is built for exactly this problem. It converts cash into cryptocurrency on the spot, with no bank involved and little identification. So if BTMs raise drug activity by serving as the payment rail for cash drug markets, the effect of BTM entry should be larger in recreational-cannabis states, where the pool of cash searching for such a rail is largest.

The prediction is a difference in treatment effects across states, which calls for a triple difference-in-differences specification,

$$Y_{\ell st} \;=\; \alpha_\ell + \gamma_{st} + \beta_1 \,\mathrm{BTM}_{\ell t} + \beta_2 \,\mathrm{BTM}_{\ell t} \times \mathrm{RecCannabis}_{st} + \varepsilon_{\ell st}, \tag{8}$$

where $\mathrm{RecCannabis}_{st}$ equals one if state s has legalized recreational cannabis sales by year t. County fixed effects α_ℓ absorb time-invariant county characteristics. State-by-year fixed effects γ_{st} absorb all statewide shocks in each year, such as legalization itself, state tax changes, and state drug policy, so the standalone $\mathrm{RecCannabis}_{st}$ term drops out and β_2 is identified by comparing BTM and non-BTM counties within the same state and year. The coefficient β_1 captures the BTM effect in states without legal recreational cannabis, and β_2 captures the additional amplification in recreational-cannabis states.

Table 20 reports the results for seizure weight. The baseline BTM effect ($\beta_1 = 0.481$)

confirms the drug channel in non-cannabis states. The interaction is large, with $\beta_2 = 0.795$ for total drugs, significant at the 5 percent level. The BTM–drug effect is about 165 percent larger in recreational-cannabis states. The amplification is strongest and most precisely estimated for cannabis seizures themselves ($\beta_2 = 1.043$, $p < 0.01$), exactly where the banking-exclusion mechanism operates. Cocaine, heroin, and methamphetamine show positive but only marginally significant amplification, suggesting that the cannabis cash economy spills over into broader drug supply chains.

Total crime provides a placebo test that validates this interpretation. If the amplification merely reflected the kinds of places that legalize cannabis, say states whose BTM counties were trending toward more illicit activity of every sort, then the interaction should appear for crime in general, not just for drug seizures. Table 21 shows it does not. The interaction for total crime is a precise zero ($\beta_2 = 0.006$, $p > 0.80$), so BTM entry raises overall crime by a similar amount in cannabis and non-cannabis states. The amplification is specific to drug supply, exactly where the mechanism operates. Abundant unbanked cannabis cash, paired with a machine that moves that cash into crypto cheaply, fosters the drug business in these counties rather than criminal activity across the board.

The result carries a direct policy implication. The SAFE Banking Act, a proposed federal law that would allow banks to serve state-legal cannabis businesses without federal penalty, would shrink the pool of unbanked cannabis cash. With less cash searching for a way into crypto, the amplification documented here should weaken. Giving legal cannabis businesses access to banking is therefore not only an industry convenience but also a tool for cutting off the cash that feeds the BTM drug channel.

7.3 From Drug Supply to Human Lives: Overdose Mortality

The seizure results establish that BTM entry expands the quantity of controlled substances circulating in local communities. The critical welfare question is whether the expansion costs lives. Overdose mortality is the most severe consequence of drug availability, and each death

carries foregone lifetime earnings, medical expenditures, and long-run community disruption. The CDC estimates the annual economic burden of overdose deaths at roughly $1 trillion, concentrated in the kinds of communities where BTMs disproportionately locate.[6]

Three channels could connect BTM entry to overdose deaths. Expanded supply enlarges the exposed user pool, so deaths rise with quantity. The drug mix shifts toward cocaine and heroin, which carry higher overdose risk per unit than cannabis, so mortality can rise even without a quantity increase. And cryptocurrency-facilitated distribution may carry high-potency synthetic opioids, whose value-to-weight ratio suits darknet logistics. The quantity channel implies gradual acceleration that mirrors the supply expansion, while the potency channel implies sharper onset.

I use NCHS model-based county-level overdose death rates (2013–2021), supplemented by CDC provisional data for 2022–2023. The model-based estimates smooth over suppressed cells in rural counties and so provide broader coverage than raw vital statistics [Rossen et al., 2020]. The panel covers 1,045 counties and 9,869 county-year observations. One limitation is that the data report all-cause overdose deaths without substance detail, which precludes a direct test of the composition and potency channels.

Table 22 reports five specifications. The preferred specification (column 1) applies the IHS transformation to the death rate per 100,000. The ATT is 0.371 (SE 0.078, $p < 0.001$), with a pre-trend p-value of 0.67. The magnitude implies an increase in the overdose death rate of roughly 45 percent. Against the sample mean of 23 deaths per 100,000, that is about 10 additional deaths per 100,000 residents, or roughly five additional deaths per year in a county of median treated size (≈50,000 residents). Valued at the EPA's value of a statistical life of $11.6 million, the implied mortality cost is on the order of $60 million per median treated county per year, three orders of magnitude above the annual revenue of a typical BTM installation ($20,000–$50,000).[7]

The remaining columns of Table 22 report level, count, extensive-margin, and Poisson

[6] See Florence et al. [2021] on the economic burden of the opioid crisis.
[7] VSL from Viscusi [2018]. BTM revenue from Coin ATM Radar (2023).

specifications, with diagnostics detailed in the table notes, and Figure 27 plots the event studies for these additional specifications.

The event-study trajectory of the preferred specification (Figure 28, left panel) shows effects near zero at entry that strengthen through three years out, mirroring the gradual drug supply expansion. This profile fits the quantity channel and fits poorly with a pure potency story, which would predict sharper onset. The seizure and mortality results together establish the chain from BTM entry through supply expansion to human deaths. The remaining question is who bears these costs.

7.4 Who Bears the Cost? Victim Composition

A given increase in aggregate crime carries very different welfare implications depending on whose victimization rises. If additional offenses hit all groups proportionally, the loss is spread evenly. If the burden concentrates on groups that are already vulnerable, BTM entry not only reduces aggregate welfare but widens existing inequality in crime exposure, much as a regressive tax widens inequality in income. A substantial literature documents that crime costs already fall disproportionately on disadvantaged communities [Sampson, 2012, Sharkey, 2018, Anders, 2023]. I test whether anonymous payment infrastructure widens that gap.

I construct victim-level measures from NIBRS records, which report each victim's race and ethnicity, age, sex, and relationship to the offender. The analysis proceeds in two stages. I first estimate effects on victim shares, the percentage of victims in each category, which identifies which groups are disproportionately targeted. I then estimate per-capita victimization rates for each group separately and scale each ATT by the group's own baseline, which reveals how large each group's increase is in proportional terms. The two views are complements, since shares are zero-sum while rates are not.

Table 23 and Figure 29 report the share results, with all pre-trend p-values above 0.90. Among demographic groups, the minority victim share rises by 7.8 percentage points, driven

primarily by Hispanic victims (7.2 points), a shift of roughly one third of the baseline minority share of 22.4 percent. The concentration among Hispanic victims is consistent with the disproportionate exposure of Hispanic communities to cocaine and methamphetamine distribution networks and reinforces the drug-market mechanism. The age profile shifts toward younger adults. The share of victims aged 18–25 rises by 2.7 points while the share aged 65 and over falls by 2.1 points, which indicates market-generated offenses rather than crimes of opportunity against the physically vulnerable. The most diagnostic dimension is the victim–offender relationship. The stranger crime share rises by 3.7 points while the domestic violence share falls by 3.3 points, with no change in intimate partner or acquaintance crime. Payment-related offenses, whether drug transactions, fraud, or robbery around cash conversion points, involve parties without prior relationships, so a shift toward stranger-perpetrated offenses is the signature of the payment-driven criminal activity that BTMs enable.

Because shares are zero-sum, a rising minority share could in principle coexist with falling victimization for everyone. The per-capita rates rule this out. Figure 30 plots each group's ATT relative to its own baseline, with delta-method standard errors. Every group experiences a statistically significant increase. Hispanic victims show the largest proportional rise (about 250 percent), followed by stranger crime (227 percent), young adults (217 percent), and minority victims overall (215 percent). Even the groups at the bottom of the ranking, elderly victims (174 percent) and Black victims (179 percent), roughly double relative to baseline. The ranking also resolves an apparent paradox in the share results. The domestic violence share falls even though the domestic violence rate doubles, because stranger and drug-related offenses grow faster still. A policymaker reading only shares might conclude that BTM entry improves domestic violence outcomes. The rate analysis shows that no group is better off.

Three distributional conclusions follow. BTM entry widens existing disparities in crime exposure, with Hispanic communities absorbing both the largest compositional shift and

the largest proportional increase. The shift toward younger victims carries long-run costs beyond immediate victimization, since disruption at that life stage damages human capital accumulation and lifetime earnings [Sharkey, 2018]. And the shift toward stranger crime confirms that the new offenses are market-generated, which directs policy toward the financial infrastructure enabling the market rather than toward relationship-based violence prevention.

7.5 From Crime to Household Credit: A Mediation Model

The harms documented so far need not stop at public safety. If crime erodes household wealth and financially marginal households cannot absorb the losses, the crime channel should propagate into credit markets. This subsection formalizes that chain in a simple mediation model and derives the predictions that connect the crime estimates to household credit outcomes.

7.5.1 Environment

Consider county ℓ in period t. Three types of agents interact. Criminals choose effort given the payment technology. Households earn income, face crime risk, and service debt. Banks set lending terms given observed default risk. BTM entry is a binary treatment $D_{\ell t} = \mathbb{1}[t \geq g_\ell]$, where g_ℓ is the year of first BTM installation in county ℓ.

7.5.2 The Criminal's Problem

Criminal j in county ℓ chooses effort e_j to maximize the expected net payoff

$$\max_{e_j \geq 0} \ \pi_j \ = \ p(e_j, \kappa_\ell) \cdot R(e_j) - \Big[1 - p(e_j, \kappa_\ell)\Big] \cdot P - c(e_j), \tag{9}$$

where $p(e, \kappa)$ is the probability of completing the crime without apprehension, increasing in effort e and decreasing in κ, the cost of monetizing proceeds anonymously. $R(e)$ is the

monetary return from crime, P is the expected punishment cost, and $c(e)$ is the direct cost of effort.

BTM entry reduces κ. Before BTM availability, monetizing crime proceeds anonymously requires costly alternatives such as money mules, physical cash transport, or layered transactions through multiple accounts. After entry, the criminal walks into a gas station and converts cash to pseudonymous cryptocurrency in minutes. Formally, κ_ℓ drops from $\bar{\kappa}$ to $\underline{\kappa}$ when $D_{\ell t} = 1$.

The first-order condition, $\frac{\partial p}{\partial e}[R(e) + P] + p \cdot R'(e) = c'(e)$, yields an optimal effort $e_j^*(\kappa)$ that is decreasing in κ. Lower anonymity costs raise crime. Aggregating over the distribution of criminal types G_ℓ,

$$C_{\ell t} \;=\; \int e_j^*(\kappa_{\ell t})\, dG_\ell(j) + \varepsilon_{\ell t}. \tag{10}$$

This maps directly to the observed crime rate, with the treatment effect $\partial C/\partial D$ corresponding to the 6.9 percent ATT from the staggered DiD.

7.5.3 The Household's Problem

Household i in county ℓ lives for two periods. In period 1 the household earns income y_i, consumes c_{1i}, saves s_i, and carries existing debt d_i at interest rate r_d. In period 2 the household receives a stochastic income draw $y_{2i} \sim F_y$ and suffers a crime loss

$$L_i = \lambda_i \cdot C_{\ell t}, \tag{11}$$

where $\lambda_i > 0$ is household i's vulnerability to crime. The parameter captures direct losses from theft, fraud, and property damage, indirect costs such as medical expenses, defensive expenditure, and insurance premia, and behavioral costs such as reduced labor supply from fear. Vulnerability varies across households. Low-income households, renters, the uninsured, and residents of high-exposure neighborhoods face higher λ, mapping to the victim composition results above.

Default occurs when period-2 resources fall short of debt obligations,

$$\text{Default}_i \;=\; \mathbb{1}\Big[\lambda_i\, C_{\ell t} > \underbrace{s_i(1+r_f) + y_{2i} - (1+r_d)d_i}_{\equiv\ \text{buffer}_i}\Big]. \tag{12}$$

The buffer is the household's financial cushion, savings plus income minus debt service. Households with small buffers default at lower crime levels. The default probability is

$$\delta_i(C_{\ell t}) \;=\; F_y\Big((1+r_d)d_i - s_i(1+r_f) + \lambda_i\, C_{\ell t}\Big), \tag{13}$$

with comparative statics

$$\frac{\partial \delta_i}{\partial C_{\ell t}} = \lambda_i \cdot f_y(\cdot) > 0 \qquad \text{and} \qquad \frac{\partial^2 \delta_i}{\partial C_{\ell t}\,\partial \lambda_i} > 0 \text{ (for log-concave } f_y). \tag{14}$$

Crime raises default probability in proportion to vulnerability, and the cross-derivative implies that the crime-to-credit effect is stronger for more vulnerable households. Recent evidence shows how far liquidity- constrained households go to avoid high-interest borrowing, including selling plasma when conventional credit is unavailable [Dooley and Gallagher, 2024].

7.5.4 The Mediation Chain

Aggregating over the household type distribution H_ℓ, the county-level delinquency rate is $\Delta_{\ell t} = \int \delta_i(C_{\ell t})\, dH_\ell(i)$, and the total effect of BTM entry on delinquency decomposes as

$$\frac{\partial \Delta_{\ell t}}{\partial D_{\ell t}} \;=\; \underbrace{\frac{\partial \Delta_{\ell t}}{\partial C_{\ell t}}}_{\text{crime}\to\text{credit}} \;\times\; \underbrace{\frac{\partial C_{\ell t}}{\partial D_{\ell t}}}_{\text{BTM}\to\text{crime}}. \tag{15}$$

Both components are estimable. The second is the 6.9 percent ATT already in hand, and the first can be recovered by instrumenting crime with BTM entry.

The banking side amplifies the chain. Competitive banks set lending rates to break even,

so $(1-\bar{\delta}_\ell)(1+r_d) = (1+r_f) + \text{risk premium}$. When BTM-induced crime raises $\bar{\delta}_\ell$, banks raise rates or ration credit. Higher rates push additional households past the default margin, and rationing excludes marginal borrowers from formal finance altogether. The model thus delivers two predictions. Delinquency should rise in treated counties in proportion to the local crime response, and the rise should concentrate among high-vulnerability, low-buffer households, the same populations the victim composition results identify as bearing the heaviest victimization costs.

8 On-Chain Forensic Analysis

The preceding sections establish the real consequences of anonymous payment access. What they leave open is what that anonymity actually is. The crypto system runs on a public ledger that records every transaction, so anonymity there is not an absence of records. It is a produced good, supplied by identifiable services that sever specific informational links for users willing to pay in fees, in effort, or in forgone purchase size. This section turns to the ledger to study that production directly. The object of study is no longer what BTM entry does to a county but how the crypto system produces and sells anonymity, who buys it, at which stage, through which services, and in what combinations, the machinery that ultimately generates the consequences measured off chain.

The agenda runs at two levels. At the macro level, the goal is an anatomy of anonymity in the crypto payment system. As Section 2.2 describes, anonymity is not a single attribute but the product of a supply chain in which each stage severs one informational link, identity to address at entry, sender to receiver mid-chain, address to purpose at exit. The anatomy asks how much value flows through the anonymity-supplying services of each stage, what each layer costs, who buys which combination of stages, and how much each stage contributes to the anonymity that survives end to end. At the micro level, I work with one traceable cluster of flow, BTM customers whose funds enter the chain with clean stage-one

anonymity, cash converted to coin below an identification threshold. The cluster's origin is fully observed, its reconstruction can be validated against audited filings, and its behavior prices anonymity one transaction at a time when identification rules bind, which is what makes anonymity traceable through the rest of the chain. Subsection 8.1 lays out the measurement design that connects the two levels. Subsection 8.2 describes the supply side of the system, the services that produce anonymity at each stage and the volumes that flow through them. Subsection 8.3 builds and validates the entry-stage measurement, Subsection 8.4 estimates the demand for anonymity it reveals, and Subsection 8.5 follows the funds to their destinations. Subsection 8.6 assembles the anatomy, and Section 9 then scales the local estimates to a national cost and embeds the measurements in an informational model of optimal anonymity.

8.1 Measurement Design

Figure 1 maps the services operating at each stage of the supply chain to the informational links they sever. Each stage is an industry with producers, a technology, and a price, and the anatomy measures it as one.

Two measurement problems stand in the way. First, anonymity leaves no aggregate footprint. It is purchased one transaction at a time by users whose motives are invisible in totals, so prevalence statistics cannot reveal who buys it, at which stage, or in what combinations. The system must be measured at the level of individual flows. Second, an individual flow can only be interpreted if its origin is known. A mixer deposit means one thing when it arrives from an exchange theft and another when it arrives from a retail cash purchase, so tracing must start from a population whose origin is observed, not sampled from the middle of the graph.

The design solves both problems with a tracer population. I reconstruct the complete customer record of Bitcoin Depot, the largest publicly traded BTM operator in the United States, 1.4 million customers whose flows all begin with cash entering a kiosk. For each flow

I observe the purchase size, the timing, and the customer's entry-stage anonymity choice relative to the operator's $250 know-your-customer (KYC) threshold, and the reconstruction is validated against the operator's audited SEC filings. Following these flows forward through the chain then reveals which stages users engage, in which combinations, and at what revealed price. Figure 31 illustrates the tracing logic. The questions that turn on individual flows, the price of anonymity and the composition of demand, are answered with the tracer. The system-wide quantities require only labeled service addresses rather than known origins, and they come from joining a consolidated label database against the full transaction record, of which BTM-originated flow is one small, fully observed slice.

Identification of the stage contributions comes from two sources of variation. In the cross section, the destination decomposition shows which stage combinations users purchase. In the time series, two enforcement shocks described in the next subsection each shut down a single stage while leaving the others open. The section describes the supply side of the system (Subsection 8.2), reports the entry-stage measurement in full (Subsections 8.3 and 8.4), reports the tracing infrastructure and its label coverage (Subsection 8.5), and assembles the remaining objects of the anatomy together with the measurement design that identifies them (Subsection 8.6).

8.2 The Supply of Anonymity

Anonymity production is concentrated in identifiable services, and each stage has a distinct industrial organization. At entry, the producers are the BTM operators whose physical network the preceding sections study, together with the other cash on-ramps. Their technology is the cash purchase below an identification threshold, their capacity is the kiosk network itself, and their price is the markup, roughly 15 percent for the largest operator against the 0.5 to 2 percent that regulated exchanges charge. Entry is the only stage where anonymity comes bundled with currency conversion and is sold over a physical counter, which is why it reaches users who never touch the rest of the crypto system.

Mid-chain, mixers pool deposits from many senders and pay out to fresh addresses, mechanically severing the link between transaction inputs and outputs, and cross-chain bridges move value onto ledgers where tracing tools lose coverage. Both charge explicit fees, so mid-chain anonymity is the only layer sold as a standalone product. At exit, unregulated exchanges and peer-to-peer platforms convert crypto into spendable value with weak or no identity verification, and darknet markets accept it directly as payment. Their margin is a spread rather than a posted fee, the gap between the execution users accept there and what a regulated venue would offer. The supply side is therefore observable service by service, which is what makes an industry-level measurement possible.

The size of each segment is measured from the ledger without any reference to BTM flow. The consolidated label database described in Subsection 8.5 identifies 3.1 million addresses belonging to mixers, unregulated exchanges, peer-to-peer platforms, gambling services, darknet markets, and regulated exchanges, and joining those labels against the full transaction record yields the dollar volume passing through the anonymity-supplying services of each stage over time. These series are the industry sizes of the three stages. They show how large the market for anonymity is, how it has grown, and how its mix across stages has shifted as regulation and enforcement moved, and the BTM flow traced in the rest of this section enters them as the one slice whose origin is fully observed.

The industry also has a regulatory history that the volume series record. The seizure of Hydra Market in April 2022 removed the largest darknet exit, and the Treasury sanctions on Tornado Cash in August 2022 closed the dominant mixing service, the two largest enforcement events in the sample period. Each removed one stage's leading producer while leaving the other stages open, which is precisely the variation that identifies how the stages interact. What the supply side cannot reveal on its own is the demand side, who buys which combination of stages and what each stage adds to the protection the chain delivers. Those two objects require following individual flows of known origin, and the next three subsections build that capability.

8.3 Reconstructing the Entry Stage

Reconstruction proceeds in three steps, summarized in Figure 32. I first identify operator-controlled addresses, then extract customer payouts from operator-initiated transactions, and finally validate the result against the operator's audited disclosures.

The starting point is the operator seed cluster attributed to Bitcoin Depot in a proprietary blockchain intelligence dataset and verified against the company's own disclosures. From BigQuery's public Bitcoin dataset I retrieve every transaction in which at least one input or output belongs to this seed set, and I extend the cluster with two heuristics. The first is the common-input-ownership (CIO) heuristic of Meiklejohn et al. [2013]. Any address that appears as an input alongside a known operator address shares cryptographic control, because constructing such a transaction requires the private key of every input. The second is a change-detection rule. The largest non-customer output of an operator-initiated transaction is treated as change returning to the operator when it exceeds the second-largest output by a factor of ten, carries a dollar value above $1,000 at the day's BTC price, uses an address format consistent with the operator's wallet generation, and later appears as a spending input in another operator-initiated transaction. The combined procedure yields a depot cluster of 105,530 operator-controlled addresses delivering 2,397,315 customer payouts to 1,405,435 unique customer wallets between November 2018 and May 2026. Dropping the change output of each operator-initiated transaction leaves the customer payouts, and each payout is converted to dollars at the daily BTC closing price. The median payout is $179. Table 24 reports the on-chain address universe by category, and Table 25 reports the size distribution of customer payouts.

The change-detection rule does real work because the operator delivers most payouts in batches. Decomposing the 247,967 operator-initiated transactions by the number of inputs and customer outputs, transactions with several inputs and several customer outputs account for 29.2 percent of transactions but 67.6 percent of customer payouts, while one-input one-customer peeling chains account for 29.4 percent of transactions and only 3.2 percent

of payouts. Batching spreads the per-byte transaction fee across recipients and lets the operator consume accumulated change alongside customer deliveries. Because two thirds of payouts arrive inside transactions that mix customer deliveries with operator change, naive transaction counting would badly misstate customer volume, and the change rule is what separates the two. Appendix Figure A.5 documents the full delivery architecture.

Customer wallets are then grouped into persons. A single individual using one BTM may generate several distinct receive addresses, either because the receiving wallet rotates addresses automatically or because the user deliberately spreads activity across wallets to limit cross-transaction linkability, and both motives are well established in the on-chain literature [Androulaki et al., 2013]. Applying the CIO heuristic to the customer side groups the 1,405,435 wallets into 1,008,458 persons, sets of addresses cryptographically shown to share one owner. I exclude 34 mega-clusters of at least 1,000 wallets each, 189,803 wallets in total, as services or aggregators rather than retail customers. All numbers below refer to the retail base after this exclusion, 1,215,632 wallets held by 1,008,424 persons.

The reconstruction can be checked against audited data. Bitcoin Depot's quarterly SEC filings (ticker BTM) disclose two series I observe independently on chain. The first is the rolling twelve-month median kiosk transaction size. The second is the quarterly cost of cryptocurrency purchased, the operator's audited BTC outflow. Figure 33 displays the alignment. The markup-adjusted on-chain median tracks the disclosed series with a correlation of 0.967 and a median absolute deviation of 2.6 percent over the 19 quarters from 2021Q1 to 2025Q3. The on-chain retail quarterly volume tracks the audited crypto cost with a correlation of 0.958 and a median deviation of 9.0 percent over 11 quarters, with mean coverage of 96 percent. The cluster I reconstruct is therefore the same customer-facing flow the operator reports to investors. The validation also surfaces one limitation. In 2023Q2 a single mega-cluster of 1,264 wallets with an over-the-counter desk signature contributes $200 million of the $395 million total flow. That cluster is among the 34 excluded from the retail base.

Who uses these machines, and what do they do with the coins? Figure 34 profiles the

retail base. The customer population is overwhelmingly one-shot and small. 86 percent of retail wallets receive exactly one payout and never appear again, and 81 percent of retail persons are also one-time, so the one-time wallets are mostly one-time individuals rather than one person rotating addresses. The median customer puts $179 to $230 of cash through the operator over their entire observed history. Volume is nonetheless concentrated, with the top 10 percent of customers accounting for about 75 percent of all dollar volume, and among the 14 to 19 percent of customers who return, 70 to 80 percent fit all of their activity within a 30-day window. Most importantly for what follows, the coins do not sit still. Tracing each wallet to its first spending transaction, the median time between receipt and first outflow is 0.9 days, 59 percent of wallets spend within 24 hours, 97 percent spend within 30 days, and fewer than 0.5 percent hold for more than one year (Figure 35). The holding period is far too short for a store-of-value motive. BTM flow is payment in motion, not investment, which means the destination of the first outflow carries the economic content of the transaction and following the funds forward is the right measurement strategy.

8.4 Demand for Anonymity

The entry stage supplies anonymity through one concrete margin. Below the operator's $250 threshold a customer buys with cash and a phone number, and at $250 the kiosk begins verifying identity. If customers value the anonymity, they should reveal it by their behavior around that line. They do, on two margins, and the behavior responds when regulation moves the line.

Bunching at the threshold. The single-transaction margin is bunching. Figure 36 displays the dollar distribution of payouts together with a polynomial counterfactual. In the window before California's regulatory cap took effect (2021Q1 to 2023Q4, see below), 58.9 percent of all customer payouts and 68.5 percent of payouts below $1,500 fall under $250. The bin just below the threshold contains 2.52 times as many payouts as the bin just above

it, computed on 1.07 million payouts with mega-clusters excluded. If customers did not respond to the threshold, that ratio would be one. The polynomial counterfactual estimator of Kleven and Waseem [2013] returns a normalized excess mass of $\hat{b} = +0.58$ on the same sample, using a fifth-order polynomial fit outside a [$220, $270) exclusion zone and masking round-number bins that otherwise bias the counterfactual upward. For comparison, estimates in the literature range from about 0.05 for the EITC kink [Saez, 2010], through 0.1 to 0.5 for Pakistani tax notches [Kleven and Waseem, 2013], to about 1 for the U.S. conforming-loan limit [DeFusco and Paciorek, 2017] and 3 for U.K. stamp-duty notches [Best and Kleven, 2018]. An estimate of 0.58 sits squarely among well-documented strong behavioral responses to a binding regulatory margin.

The dollar-denomination falsification. A natural alternative story is that the bunching is a round-number habit in Bitcoin units rather than a response to the dollar threshold. The data reject this story. Over the sample the BTC price moves from $30,000 to $67,000 to $19,000 and back, yet the excess mass stays fixed at $250 in dollar terms and never moves with the price. Whatever customers are responding to is denominated in dollars, which is the currency of the identification rule, not the currency of the asset.

Structuring across the threshold. The multi-transaction margin is structuring, splitting a larger intended purchase into several sub-$250 payouts within a short window. For each wallet-day I record the maximum single payout and the same-day sum, and a structuring event is a day on which the maximum stays below $250 while the sum reaches $250 or more across at least two payouts. Figure 37 characterizes this population. In the pre-cap window, 5,267 wallet-days meet the definition. Aggregating to the person level, so that several wallets controlled by one CIO-clustered individual on the same day count as one event, raises the count to 12,943, which means 59 percent of person-level structuring spreads across multiple wallets in a single day's session. The multi-wallet share is a sophistication signal, because structurers actively diversify across receive addresses, which defeats any monitoring

the operator runs on its own customer records. It is also a methodological point, because each wallet is a separate customer from the operator's perspective, so multi-wallet structuring is invisible in operator data and detectable only through cryptographic clustering. Structuring is behaviorally specific rather than a mechanical feature of busy days. Only 9.8 percent of multi-payout wallet-days and 12.2 percent of multi-payout person-days fall in the structuring region, the median same-day structured sum is $351, and the mean number of sub-payouts is 3.5.

A regulatory regime shift. A regulatory shift confirms that the $250 trigger is the margin these behaviors target. On January 1, 2024, California's Digital Financial Assets Law (DFAL) took effect and imposed a flat $1,000-per-day cap on BTM transactions, enforced by the operator at the kiosk. After that date, a customer in a covered location can no longer assemble a large purchase from repeated sub-$250 visits, so the incentive to structure around $250 disappears. If the excess mass and the splitting behavior were driven by anonymity-seeking at the threshold, both should fall sharply beginning in 2024Q1. If they instead reflected round-number habits or small-budget customers, the cap should leave them unchanged. Figure A.6 plots the quarterly structuring rate, the share of retail person-days containing at least two sub-$250 payouts that sum to $250 or more. The pre-DFAL mean over 2021Q1 to 2023Q4 is 1.97 percent, with a peak of 2.83 percent in 2022Q4. The post-DFAL mean over 2024Q1 to 2025Q4 is 1.13 percent, a 43 percent relative decline that begins in exactly the quarter the law takes effect. The decline is nationwide rather than confined to California, because Vermont, Nebraska, and Louisiana adopted similar caps over 2024 and because the largest operator applies portions of its compliance infrastructure across states, so I read the evidence as a sharp regime-shift fact rather than a difference-in-differences estimate. Even under this conservative reading, the conclusion is the same one the bunching and the dollar-denomination falsification deliver. The behavior at $250 exists because of the identification threshold, and it responds when regulation moves the threshold.

The price of anonymity. These margins price the anonymity. At $\hat{b} = +0.58$, roughly 20,000 payouts in the pre-cap window are displaced from where a smooth distribution would put them into the zone just below $250, customers who intended to buy more and cut the amount to avoid showing identification. Structurers go further and pay in repeated trips and transaction fees. Customers paying the operator's roughly 15 percent markup are paying on the order of $25 per median transaction, and those bunching and structuring at the threshold are paying real costs in foregone purchase size and effort, for nothing other than not showing identification. The threshold therefore isolates a population whose revealed valuation of anonymity exceeds what they give up to keep it.

8.5 Where the Funds Go

The entry stage tells only the first part of the story. Whether the anonymity acquired at the machine survives depends on where the funds go next, so the second measurement task is to classify each customer's destinations by type. I trace each validated customer wallet to its spending transactions and classify the receiving addresses into the categories that define the supply chain. Regulated exchanges require identity verification at deposit, so flow that ends there re-enters the documented financial system. Unregulated exchanges and peer-to-peer platforms convert crypto to spendable value with weak or no identification. Remittance and payment services represent ordinary commerce. Mixers are purchased obfuscation. Gambling services and darknet markets are terminal uses, with the latter directly illicit. Flow that remains in personal wallets stays inside the pseudonymous system. The resulting decomposition delivers $\hat{\phi}$, the share of BTM-originated volume reaching illicit endpoints, and the stage-level shares used in the next subsection.

The tracing problem is tractable because customer behavior is focused. 36 percent of customers spend exactly once and the median customer makes two spending transactions. 39 percent of customer wallets and 35 percent of persons send to exactly one unique destination address, and the median is two destinations. Most customers arrive at the moment of spend-

ing with a single intent rather than a portfolio to build. Destination diversity concentrates in a small tail, with the top 10 percent of customers accounting for 87 to 89 percent of all unique customer-destination links (Figure 35). Two filters precede classification. First, 39,174 customers with more than 100 unique destinations each carry a service-like signature, are set aside for separate identity verification, and are excluded from the retail destination analysis. Their exclusion shrinks the unique-destination universe from 16.4 million to 3.82 million addresses while retaining 1.34 million customers, 97.2 percent of those with observed spending. Second, I assemble a consolidated label database from nine public and proprietary sources, anonymized here as sources A through I, spanning entity-category, ransomware, sanctions, and scam-address registries. Deduplicated across sources, the database covers 3.1 million unique Bitcoin addresses under a twelve-category taxonomy with source-priority resolution.

Joining the label database against the 3.82 million retail destinations shows how far direct labels reach. Only 0.28 percent of destination addresses carry a third-party label, but those labeled destinations capture 4.0 percent of customer reach and 63.4 percent of dollar volume, where the volume figure is inflated by co-spending and should be read as an upper bound on labeled coverage. The labeled slice identifies the major centralized exchange deposit clusters, including Coinbase and Binance, along with 473 destinations classified as darknet markets reaching 0.38 percent of customers and 124 destinations classified as mixing services. The unlabeled remainder is dominated by addresses that receive from exactly one customer, 84.6 percent of them, a pattern consistent with personal wallets that third-party databases do not catalog. The full decomposition therefore combines three inputs, direct labels where they exist, propagation of labels through short spending paths for the flow that transits intermediate hops before reaching a labeled service, and a stratified verification sample that bounds the composition of the unlabeled one-customer tail.

8.6 Assembling the Anatomy

The pieces are now in place to assemble the anatomy. For the entry stage the account is complete. Its technology is a regulatory threshold, its price is the roughly 15 percent markup plus the purchase size customers forgo at the line, its demand is the excess mass of $\hat{b} = +0.58$ and the structuring population behind it, and its response to regulation is the 43 percent decline after DFAL. The stage produces anonymity for free, and a measurable population consumes it at a real cost. The supply-side measurement of Subsection 8.2 sizes the stages the funds move through next. The two objects that remain are the ones only origin-known flow can deliver, the composition of demand across stages and the contribution of each stage to the anonymity that survives end to end.

Composition is measured by sorting traced flow across routes. A sub-threshold BTM purchase breaks the identity-to-address link at entry. Whether the anonymity persists depends on the route the funds take afterward. A customer whose funds end at a regulated exchange surrenders the anonymity at exit, because the deposit re-attaches an identity, so for this customer the entry stage supplied all the anonymity there was and the system took it back. A customer who routes through a mixer purchases additional sender-receiver unlinkability mid-chain. A customer who exits through an unregulated exchange, a peer-to-peer platform, or a darknet market preserves the broken identity trail through the off-ramp. Sorting customer flow across these routes measures what share of the anonymity users obtain comes free at the machine and what share is purchased downstream. Formally, the decomposition yields the fractions $\hat{\pi}_1$ for flow whose anonymity exists only at entry, $\hat{\pi}_{12}$ for flow that adds mid-chain mixing, $\hat{\pi}_{13}$ for flow that adds an anonymity-preserving exit without mixing, and $\hat{\pi}_{123}$ for flow that uses all three stages, the empirical weights on the stage combinations $S \subseteq \{1, 2, 3\}$ that the model of Section 9 takes as inputs. The labeled join of Subsection 8.5 already locates every route in the data, the major regulated exchange clusters on one side and the 124 mixing-service and 473 darknet destinations on the other, so completing the decomposition is a matter of extending label coverage rather than of finding the routes.

The same tracing logic also applies to a second origin-known population, flows that begin at addresses the label database flags as ransomware operators, sanctioned entities, or dark-net markets, so BTM customers are one tracer among two rather than the universe of the measurement. Comparing the retail-cash bundle with the illicit-origin bundle separates the anonymity demanded by ordinary cash users from the anonymity demanded by crime, and the two bundles bracket the composition of the system as a whole.

The same decomposition measures the dependence between stages at the individual level. The customers I observe bunching and structuring at the entry threshold are observed again at their destinations, so the joint distribution of entry-stage avoidance and downstream obfuscation is directly computable. If the customers who avoid identification at the kiosk are disproportionately the ones who route through mixers and anonymity-preserving exits, the stages are complements in demand. Anonymity-seekers layer them, and tightening any single stage degrades the whole chain for exactly the users who value it most. If instead the behaviors are independent or negatively related, the stages are substitutes. Squeezing one stage pushes demand toward the others, and single-stage regulation invites a waterbed effect. The two cases carry opposite implications for whether entry-stage rules such as KYC thresholds and daily caps reduce system-wide anonymity or merely relocate it, which is the question the coordination result in the model turns on.

Contribution is identified by the two enforcement shocks, the Tornado Cash sanctions of August 2022 and the Hydra Market seizure of April 2022, each read at both levels. At the system level, the response of the volume series of Subsection 8.2 around each event measures how total anonymity production absorbs the loss of one link. At the customer level, how BTM-originated flow reallocates across the remaining routes identifies the substitution elasticities between stages for the origin-known retail population. If the stages are complements, flow into the entire anonymity-preserving chain should contract when one link disappears. If they are substitutes, the same flow should reroute toward the surviving obfuscation options. Combining the two readings, the elasticities scaled by the system-wide stage volumes deliver

each stage's marginal contribution to total anonymity, the object that determines where a marginal dollar of enforcement buys the most. Together, the quantity series, the composition shares, the dependence test, and the shock estimates complete the anatomy.

9 An Informational Model of Payment Anonymity

The measurements of Section 8 serve two quantitative purposes. The first is an accounting identity that scales the local causal estimates of the preceding sections to a national cost of anonymous payment flow. The second is an informational model in which each measurement is the empirical analogue of a parameter, and which converts the accounting into a policy trade-off. I present them in that order.

9.1 From a Local Estimate to a National Cost

The scaling rests on three objects linked by one identity,

$$\underbrace{\tau}_{\text{local damage}} \xrightarrow{\div \bar{F}} \underbrace{c = \tau/\bar{F}}_{\text{damage per dollar}} \xrightarrow{\times Q} \underbrace{C = cQ}_{\text{national cost}} .$$

The numerator τ is the dollar harm per BTM-county-quarter, the causal crime and mortality effects of Sections 4 and 7 monetized at standard victim costs and the value of a statistical life. The denominator $\bar{F}$ is the average dollar flow per BTM-county-quarter, the operator's customer outflow reconstructed on chain in Section 8 and validated against Bitcoin Depot's audited \$2.4 billion of cumulative kiosk revenue, divided by the number of kiosk-quarters the operator runs. Their ratio c is dimensionless, dollars of social harm per dollar of anonymous flow. The scale factor Q is the total anonymous dollar flow through the system, the entry-stage volume series of Subsection 8.2.

The division of identification labor is deliberate. The damage τ is causal, estimated from the staggered difference-in-differences and Bartik instrument designs. The exposure $\bar{F}$

and the scale Q are descriptive accounting, observed directly on the ledger. Because τ and $\bar{F}$ are both measured at the county-quarter level of the same operating network, the ratio requires no attribution of harm to any individual machine. The construction is the empirical damage-function approach used to compute the social cost of carbon, a causal local damage estimate divided by local exposure and scaled by the aggregate quantity [Hsiang et al., 2017, Carleton et al., 2022], with anonymous payment volume in the role of emissions.

Two disciplines apply to the extrapolation. Applying c to Q assumes the harm content of a BTM dollar extends to anonymous dollars entering elsewhere, and the comparison of the retail-cash bundle with the illicit-origin bundle in Subsection 8.6 bounds how far the compositions differ. More fundamentally, the identity only counts costs. It treats every anonymous dollar as carrying harm and none as carrying value, so on its own it implies that the optimal quantity of anonymity is zero. That conclusion is an artifact of the accounting, not a result, and correcting it requires a model with both sides of the market.

9.2 An Informational Model of Payment-System Anonymity

The model is organized around a single trade-off. Anonymity is one product facing two demand curves. It is bought by users with a legitimate demand for privacy and by users who need it to profit from crime, and no stage of the supply chain can separate the two, because any screening device would require exactly the identification the product exists to remove. A policy that degrades anonymity for one population therefore degrades it for both, and the optimal strength of each informational link is interior rather than maximal.

The model treats a financial transaction as producing information along three dimensions, indexed by the three stages of the supply chain. The parameter $\theta_1 \in [0, 1]$ measures the strength of the identity-to-address link at the on-ramp, $\theta_2 \in [0, 1]$ the strength of the sender-receiver link mid-chain, and $\theta_3 \in [0, 1]$ the strength of the address-to-purpose link at the off-ramp. On the legitimate side, a user draws a privacy value v from a distribution F and transacts anonymously when v exceeds the cost of assembling the anonymity bundle she

wants, a cost that rises with every link her bundle must defeat through forgone purchase size, repeated trips, fees, and markups. Aggregating over these users gives the privacy surplus $V(\boldsymbol{\theta})$, decreasing in each θ_j. On the illicit side, each crime type τ requires breaking a set $S(\tau) \subseteq \{1, 2, 3\}$ of links to remain undetected, and the detection probability takes the Cobb-Douglas form

$$p(\boldsymbol{\theta}, \bar{\theta}_\ell \mid \tau) = \Big(\prod_{j \in S(\tau)} \theta_j^{\alpha_j} \Big)\, h(\bar{\theta}_\ell),$$

with α_j the detection elasticities and $h(\cdot)$ the county-level information externality. A crime is committed when its gain exceeds the entry cost plus the expected punishment, so the volume of type τ is $q_\tau(\boldsymbol{\theta}) = 1 - G_\tau\Big(c(\theta_1) + p(\boldsymbol{\theta} \mid \tau)\, \Phi\Big)$, with G_τ the distribution of criminal gains and Φ the penalty, and each unit of crime imposes the dollar-monetized harm $H(\tau)$.

Welfare nets the two sides, and its derivative at any stage displays the trade-off directly,

$$W(\boldsymbol{\theta}) = V(\boldsymbol{\theta}) - \sum_\tau H(\tau)\, q_\tau(\boldsymbol{\theta}), \qquad \frac{\partial W}{\partial \theta_j} = \underbrace{\frac{\partial V}{\partial \theta_j}}_{\text{privacy surplus lost}} - \sum_\tau H(\tau) \underbrace{\frac{\partial q_\tau}{\partial \theta_j}}_{\text{crime deterred}}.$$

The first term is negative because tightening any link prices marginal legitimate users out of anonymity they value. The second term is positive because the same tightening raises detection and deters crime. The optimal θ_j^* equates the two margins. Full traceability is optimal only if legitimate privacy value is zero, and full anonymity only if crime harm is zero, so neither corner survives the model. The policy question is then not whether anonymity is good or bad but on which side of θ_j^* current regulation sits, stage by stage.

Four measurements feed this structure. First, the bunching and structuring margins trace the demand curve for anonymity at entry, the $\partial V / \partial \theta_1$ side of the trade-off. The excess mass identifies the purchase value the marginal user forgoes rather than cross the identification line, the structuring events identify the extra transaction friction accepted, and the BTM fee premium of roughly \$25 per transaction over regulated exchanges provides a lower bound on revealed willingness to pay. Second, the illicit share $\hat{\phi}$ from the destination decomposition

splits the marginal anonymous dollar between the two demand curves. The fraction $1-\hat{\phi}$ of the demand priced at the threshold weights the surplus side, the fraction $\hat{\phi}$ weights the harm side, and the same share separates the illicit content of the flow on which the damage rate c of Subsection 9.1 is estimated. Third, the supply-chain weights $\hat{\pi}$ defined in Subsection 8.6 give the empirical distribution over the stage combinations S, and the comparison of the retail-cash bundle with the illicit-origin bundle performs empirically the separation of the two demand curves that the supply chain itself cannot perform. Without these weights the cross-stage interaction term in the model cannot be evaluated. With them it reduces to a closed-form weighted sum of the calibrated detection elasticities. Fourth, the stage-dependence test disciplines the detection technology. The Cobb-Douglas form assumes the stages are complements rather than additive or freely substitutable. If the same customers exhibit entry-stage avoidance and downstream obfuscation, the multiplicative structure is supported. If the behaviors are independent or substitutes, single-stage tightening relocates anonymity rather than reducing it, and the model's coordination result must be discounted accordingly.

With these inputs the system-wide welfare cost decomposes as

$$\Delta W_{\text{system}} = \underbrace{\Delta W_1}_{\substack{\text{from BTM DiD/IV}}} + \underbrace{\Delta W_2}_{\substack{\text{calibrated from 2022}\\ \text{Tornado Cash sanctions}}} + \underbrace{\Delta W_3}_{\substack{\text{calibrated from 2022}\\ \text{Hydra Market seizure}}} + \underbrace{\Delta W_{\text{interact}}(\hat{\pi})}_{\substack{\text{cross-stage, from}\\ \text{on-chain } \hat{\pi} \text{ above}}}.$$

Here $\Delta W = W(1,1,1) - W(\boldsymbol{\theta}^{\text{now}})$ values the move from current policy to full traceability, an accounting benchmark at the corner of the policy space rather than the optimum, and the same inputs evaluate the marginal condition above to locate which side of θ_j^* each stage currently sits on. The off-chain causal sections deliver ΔW_1 directly. The natural-experiment calibrations deliver ΔW_2 and ΔW_3 with sensitivity bands. The on-chain reconstruction delivers the $\hat{\pi}$ vector that closes the system by quantifying the cross-stage interaction term, and the bunching margins deliver the privacy-surplus side that turns the calculation into a trade-off rather than a one-sided cost count. The welfare cost per dollar of anonymous

payment volume, with confidence intervals propagated through statistical, calibration, and functional-form uncertainty, is the object this calibration delivers.

10 Conclusion

This paper traces a causal chain from the physical expansion of anonymous payment infrastructure through crime, drug markets, and public health to household financial outcomes, and then prices the anonymity that drives the chain. The staggered deployment of more than 30,000 Bitcoin ATMs across U.S. counties provides the identifying variation, and the public blockchain provides the measurement.

The headline causal result is an increase in total reported crime of about 7 percent after a county's first kiosk arrives, with a monotonic gradient across a pre-specified payment-relevance classification. Offenses that exploit the kiosk's payment function rise by about 10 percent, indirectly connected offenses by about 5 percent, and placebo offenses not at all. The Chicago case study replicates the finding within a single city, with effects of roughly 4 to 9 percent across census tracts and police beats, and rules out population sorting, enhanced policing, and spatial displacement by direct test. Bartik instrument estimates built from pre-Bitcoin retail geography deliver the same magnitude and point, if anything, to the difference-in-differences estimates being conservative.

The downstream costs are large and unevenly borne. Drug seizure weight rises by about 75 percent, concentrated in cocaine and heroin, the two darknet-linked hard drugs, and the effect is amplified where the cannabis banking exclusion leaves dispensary cash searching for a payment rail. Overdose mortality rises by about 45 percent, roughly five additional deaths per year in a median treated county. At the EPA's value of a statistical life that is a mortality cost on the order of $60 million per county and year, three orders of magnitude above the revenue of a typical kiosk. Supply-side narcotics offenses rise while demand-side offenses do not, and the victimization burden concentrates on Hispanic communities, young

adults, and the targets of stranger-perpetrated crime.

The on-chain reconstruction turns these local estimates into a measurement of the anonymity market itself. From the raw blockchain I recover 2.4 million customer payouts to 1.4 million wallets of the largest U.S. operator and validate the series against its audited disclosures, with correlations above 0.95. Demand for anonymity is directly visible at the $250 identification threshold. Purchase sizes bunch below the line, structurers split larger purchases across visits and across wallets in patterns only cryptographic clustering can detect, and when California's DFAL cap removed the value of staying below the line the structuring rate fell by 43 percent in the exact quarter the law took effect. The destination decomposition shows where the anonymized funds go, and the county-level crime response rises monotonically in the identification permissiveness of the kiosks a county receives. The bunching on chain and the crime gradient off chain are mirror images of the same margin. The identification threshold is where BTM design translates into real-world harm.

The final section assembles these measurements into an answer to the question the paper set out to ask, what payment anonymity costs and what it is worth. The cost side is an accounting identity. The causal estimates put a dollar harm on each BTM county and quarter, the reconstruction measures the dollar flow through the same counties, and the ratio prices the harm per anonymous dollar, the same logic that prices a ton of carbon. Counted alone, those costs would imply that the optimal quantity of anonymity is zero. The value side rejects that answer. The typical customer converts $179 in cash and pays roughly $25 more than a regulated exchange would charge, a price only someone who values not being identified would pay. Part of that demand is criminal, the structurers and the darknet-bound dollars visible in the data. Part of it is remittances, unbanked customers, and ordinary privacy. The ledger shows both populations standing in the same queue, and no identification rule can tell them apart at the point of sale.

The model combines the two sides into a single rule. Tighten identification until the crime deterred stops outweighing the privacy priced out. Three regulatory shocks, one per

stage of the supply chain, locate current rules relative to that point. At entry, California's 2024 daily cap removed the value of staying below $250 and structuring fell by 43 percent in the quarter the law took effect. Mid-chain, the 2022 sanctions on Tornado Cash struck purchased obfuscation, and at exit the seizure of Hydra Market removed the largest illicit endpoint, the two natural experiments that calibrate the stage-level terms of the model. The cross section delivers the same message. Counties whose kiosks demand identification at low amounts show no crime response at all, while the most permissive tercile drives the entire effect and exceeds 20 percent by the fourth year. Anonymity rules bind, behavior moves when they move, and the measured harm concentrates exactly where the rules are loosest. That is what the permissive side of the optimum looks like. That the communities bearing the heaviest costs have the weakest voice in regulatory proceedings adds urgency to moving the rules.

The question is larger than Bitcoin ATMs. Cash is receding, central banks are designing digital currencies, and every payment system must decide how much anonymity to build in. Until now that choice was made without measurement, because anonymous flow was invisible by definition. Cryptocurrency broke the constraint. Its anonymity is sold at physical machines with known locations and recorded on a ledger anyone can read. This paper uses that window to put numbers on both sides of the choice, the crime, drug supply, and deaths that anonymous dollars carry, and the price legitimate users pay to keep their transactions private. Payment anonymity is a design choice with real effects, and it can now be set by measurement rather than by assumption.

References

Toni Ahnert, Peter Hoffmann, and Cyril Monnet. Payments and privacy in the digital economy. *Journal of Financial Economics*, 163:103960, 2025.

John Anders. The long run effects of de jure discrimination in the credit market: How redlining increased crime. *Journal of Public Economics*, 222:104882, 2023.

Elli Androulaki, Ghassan O. Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. Evaluating user privacy in Bitcoin. In *Financial Cryptography and Data Security (FC 2013)*, volume 7859 of *Lecture Notes in Computer Science*, pages 34–51. Springer, 2013.

Scott R Baker, Nicholas Bloom, and Steven J Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

Timothy J Bartik. Who benefits from state and local economic development policies? Technical report, W.E. Upjohn Institute for Employment Research, 1991.

Michael Carlos Best and Henrik Jacobsen Kleven. Housing market responses to transaction taxes: Evidence from notches and stimulus in the u.k. *Review of Economic Studies*, 85 (1):157–193, 2018.

Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–238, 2015.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285, 2024.

Markus K. Brunnermeier, Harold James, and Jean-Pierre Landau. The digitalization of money. *NBER Working Paper*, (26300), 2019.

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

Tamma Carleton, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viscusi, Jiacan Yuan, and Alice T Zhang. Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *Quarterly Journal of Economics*, 137(4):2037–2105, 2022.

Jonathan P Caulkins, Peter Reuter, Martin Y Iguchi, and James Chiesa. *How Goes the "War on Drugs"? An Assessment of U.S. Drug Problems and Policy*. RAND Corporation, 2005.

Claire Célérier and Adrien Matray. Bank-branch supply, financial inclusion, and wealth accumulation. *Review of Financial Studies*, 32(12):4767–4809, 2019.

Jiafeng Chen and Jonathan Roth. Logs with zeros? some problems and solutions. *Quarterly Journal of Economics*, 139(2):891–936, 2024.

Coin ATM Radar. Bitcoin atm industry statistics and use cases. https://coinatmradar.com/charts/, 2023. Accessed: June 2026.

Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754–1797, 2019.

Lin William Cong, Campbell R Harvey, Daniel Rabetti, and Zong-Yu Wu. An anatomy of crypto-enabled cybercrimes. *Management Science*, 71(4):3622–3633, 2025.

Lin William Cong, Tingyu Feng, Yican Liu, and Fangzhou Lu. Crypto ATMs: Material effects of virtual currencies. Working paper, SSRN 5617630, 2026.

Crypto.com. Crypto market sizing report 2024. https://crypto.com/en/research/2024-crypto-market-sizing-report, 2025. Accessed: June 2026.

Gordon Dahl and Stefano DellaVigna. Does movie violence increase violent crime? *Quarterly Journal of Economics*, 124(2):677–734, 2009.

Anthony A. DeFusco and Andrew Paciorek. The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy*, 9(1):210–240, 2017.

John M Dooley and Emily A Gallagher. Blood money: Selling plasma to avoid high-interest loans. *Working Paper*, 2024.

FATF. Risk-based approach: Virtual currencies. Technical report, Financial Action Task Force, Paris, 2015.

Curtis Florence, Feijun Luo, and Ketra Rice. The economic burden of opioid use disorder and fatal opioid overdose in the United States, 2017. *Drug and Alcohol Dependence*, 218: 108350, 2021.

Sean Foley, Jonathan R Karlsen, and Tālis J Putniņš. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies*, 32(5): 1798–1853, 2019.

Rodney Garratt and Maarten R C van Oordt. Privacy as a public good: A case for electronic cash. *Journal of Political Economy*, 129(7):2085–2127, 2021.

Pulak Ghosh, Boris Vallée, and Yao Zeng. Fintech lending and cashless payments. *Journal of Finance*, 81(2):1053–1101, 2026.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

John M Griffin and Amin Shams. Is Bitcoin really untethered? *Journal of Finance*, 75(4): 1913–1964, 2020.

Solomon Hsiang, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D J Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser. Estimating economic damage from climate change in the united states. *Science*, 356(6345):1362–1369, 2017.

John R. Huck. Psychological externalities of investing: Evidence from stock returns and crime. *Review of Financial Studies*, 37(7):2273–2314, 2024.

Stephen P. Jenkins. Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics*, 57(1):129–138, 1995.

Charles M. Kahn. Payment systems and privacy. *Federal Reserve Bank of St. Louis Review*, 100(4):337–344, 2018.

Charles M Kahn and William Roberds. Why pay? an introduction to payments economics. *Journal of Financial Intermediation*, 18(1):1–23, 2009.

Charles M Kahn, James McAndrews, and William Roberds. Money is privacy. *International Economic Review*, 46(2):377–399, 2005.

Henrik Jacobsen Kleven and Mazhar Waseem. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *Quarterly Journal of Economics*, 128(2):669–723, 2013.

Narayana R Kocherlakota. Money is memory. *Journal of Economic Theory*, 81(2):232–251, 1998.

Igor Makarov and Antoinette Schoar. Blockchain analysis of the bitcoin market. NBER Working Paper 29396, National Bureau of Economic Research, 2021.

Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the ACM Internet Measurement Conference*, pages 127–140, 2013.

Kenneth S Rogoff. Blessing or curse? foreign and underground demand for euro notes. *Economic Policy*, 13(26):261–303, 1998.

Kenneth S. Rogoff. *The Curse of Cash: How Large-Denomination Bills Aid Crime and Tax Evasion and Constrain Monetary Policy*. Princeton University Press, Princeton, NJ, 2017.

Lauren M Rossen, Brigham Bastian, Margaret Warner, Diba Khan, and Yolanda Chong. County-level drug overdose mortality: United States, 2006–2019. *NCHS Data Brief*, (394), 2020.

Jonathan Roth. Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322, 2022.

Emmanuel Saez. Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2(3):180–212, 2010.

Robert J Sampson. *Great American City: Chicago and the Enduring Neighborhood Effect*. University of Chicago Press, 2012.

Peter Sands. *Making it Harder for the Bad Guys: The Case for Eliminating High Denomination Notes*. Number 52 in M-RCBG Associate Working Paper Series. Harvard Kennedy School, Cambridge, MA, 2016.

Patrick Sharkey. *Uneasy Peace: The Great Crime Decline, the Renewal of City Life, and the Next War on Violence*. W. W. Norton, 2018.

James H Stock and Motohiro Yogo. Testing for weak instruments in linear IV regression. *Identification and Inference for Econometric Models*, pages 80–108, 2005.

Visa. Visa onchain analytics dashboard. https://visaonchainanalytics.com/transactions, 2026. Accessed: June 2026.

W Kip Viscusi. Pricing lives: International guideposts for safety. *Economic Record*, 94: S1–S10, 2018.

Richard Wright, Erdal Tekin, Volkan Topalli, Chandler McClellan, Timothy Dickinson, and Richard Rosenfeld. Less cash, less crime: Evidence from the electronic benefit transfer program. *Journal of Law and Economics*, 60(2):361–383, 2017.

Figure 1: Anonymity Tools by Supply Chain Stage and Identity Links Severed

Notes: Top panel shows the identity chain: a fully traceable cryptocurrency transaction connects real-world identity to counterparty through five information links (L1–L5). Bottom panel maps each provider to its supply chain stage and indicates which identity links it disrupts (filled circles). Stage 1 (entry) tools sever links 1–3, converting fiat to crypto without identity disclosure. Stage 2 (obfuscation) tools sever links 4–5, disrupting the on-chain transaction graph. Stage 3 (exit) tools sever link 5, enabling anonymous conversion to real-world value. †Subject to OFAC sanctions or law enforcement seizure since 2022.

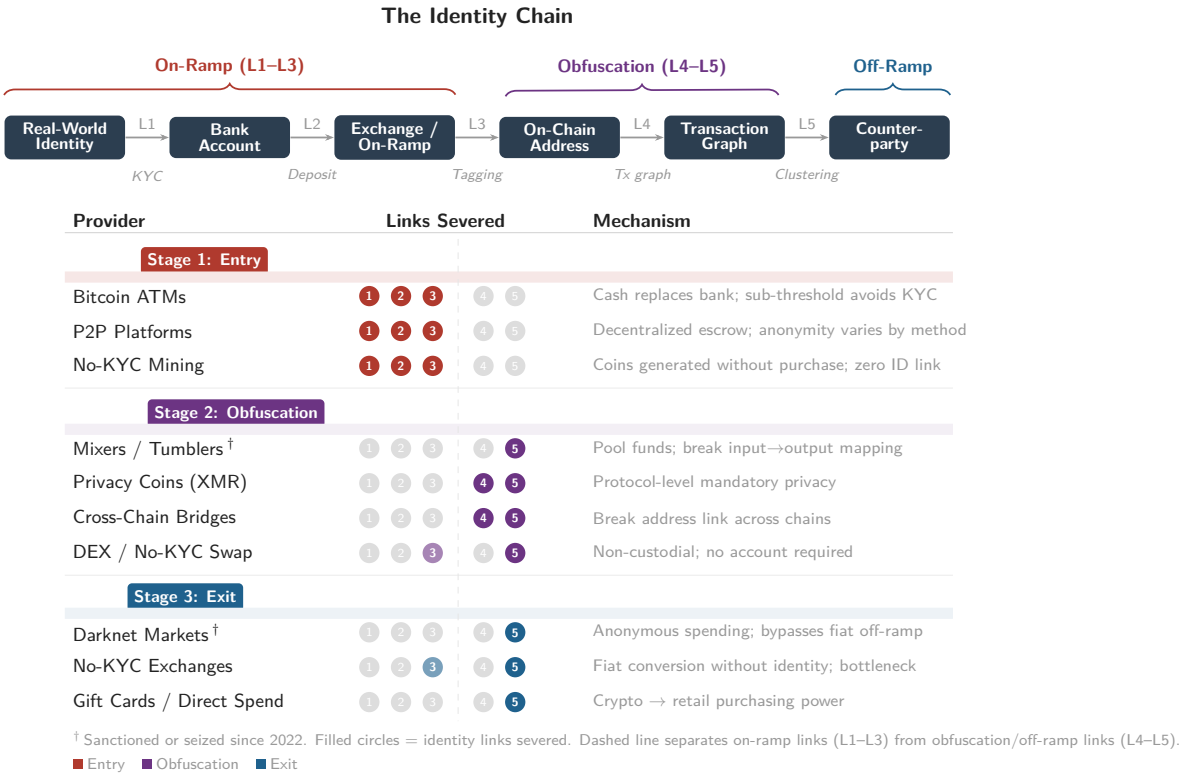


Figure 2: Cryptocurrency Types Accepted at Bitcoin ATMs

Notes: Share of active US Bitcoin ATM installations by accepted cryptocurrency, as of December 2023. Bitcoin (BTC) is universally supported. A substantial share of machines also accepts Litecoin (LTC) and Ethereum (ETH), which offer faster settlement times and lower fees for small transactions. Monero (XMR) and Dash acceptance, while present, is limited. The near-universal support for Bitcoin motivates the focus on Bitcoin ATMs as the unit of analysis; results are robust to restricting the sample to BTC-only machines. Source: Coin ATM Radar.

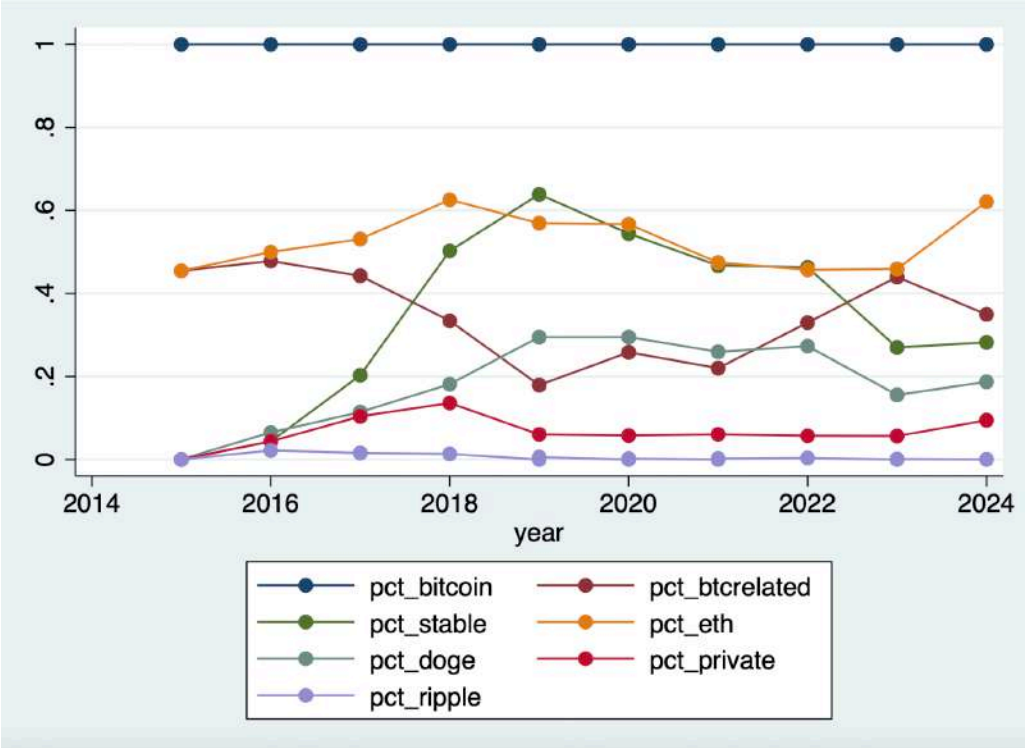


Figure 3: Geographic Diffusion of Bitcoin ATMs, 2015–2024

Notes: Quarterly share of U.S. states, counties, and ZIP codes with at least one active Bitcoin ATM, 2015Q3 through 2024Q3. Diffusion was negligible before 2016, spread across states first, and accelerated at the county and ZIP level during 2020–2021 in tandem with the Bitcoin price run-up. By the end of the sample every state, about 60 percent of counties (1,880 counties), and about 23 percent of ZIP codes host at least one kiosk. Source: Coin ATM Radar kiosk records.

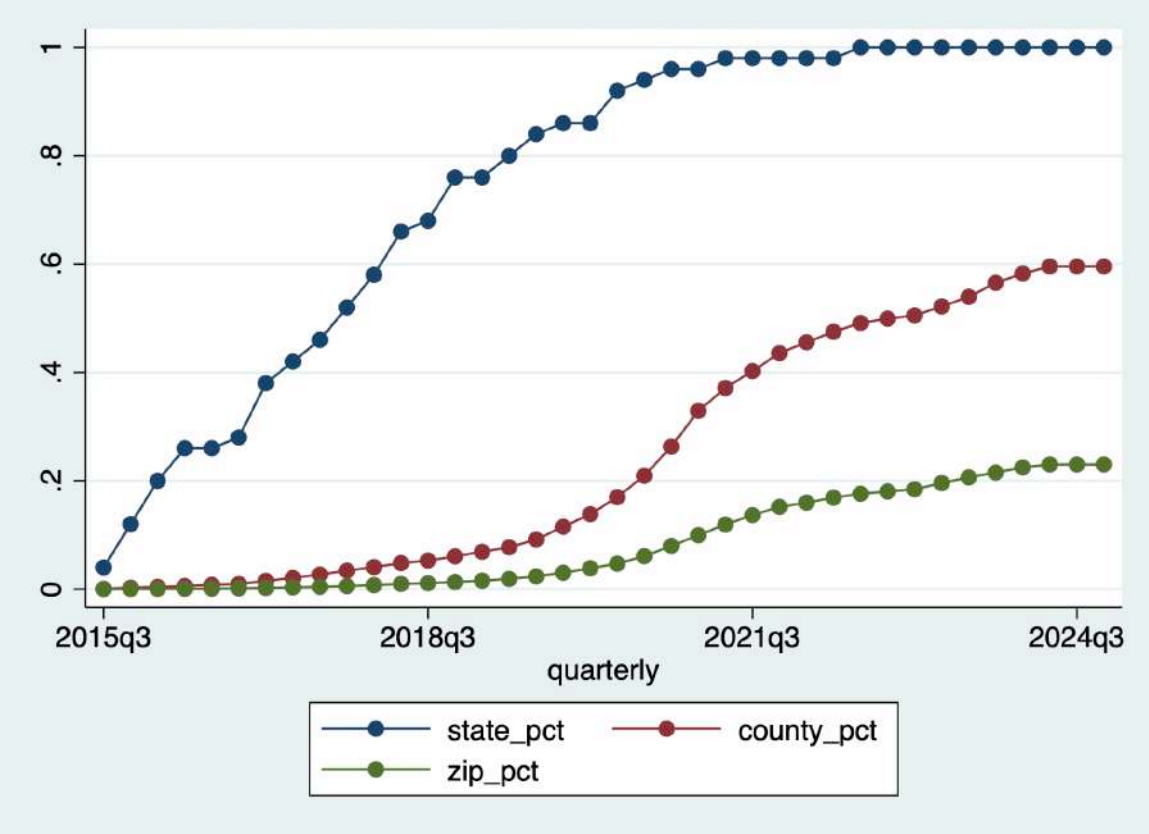


Figure 4: Geographic Distribution of Bitcoin ATMs: State-Level Density

Notes: State-level Bitcoin ATM density (machines per 100,000 residents), averaged over the full sample period 2013–2023. Nevada, Texas, and Florida rank among the highest-density states, partly reflecting favorable state money-transmitter licensing environments (captured by the BDI) and large populations of remittance-sending Hispanic residents. Hawaii and some Northeast states rank near the bottom, consistent with their hostile regulatory postures (see Appendix A). Source: Coin ATM Radar; population from the 2020 Census.

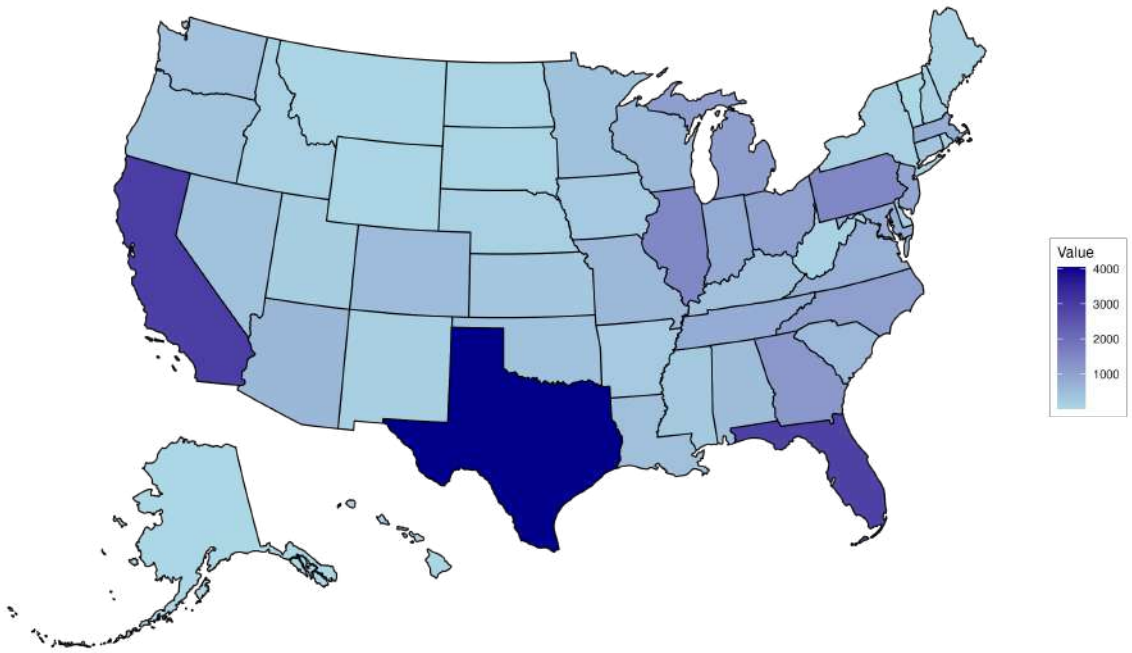


Figure 5: Bitcoin ATM Host Venue Types

Notes: Distribution of Bitcoin ATM installations by host venue category, based on the full Coin ATM Radar dataset (2013–2023, $N \approx 31{,}000$ machines). Gas stations and convenience stores together account for the majority of installations, consistent with the operator strategy of embedding machines in high-foot-traffic, 24-hour locations that serve unbanked and cash-intensive consumers. The dominance of gas stations motivates including gas station density as a supply-side control variable in the entry models (Section 5.1).

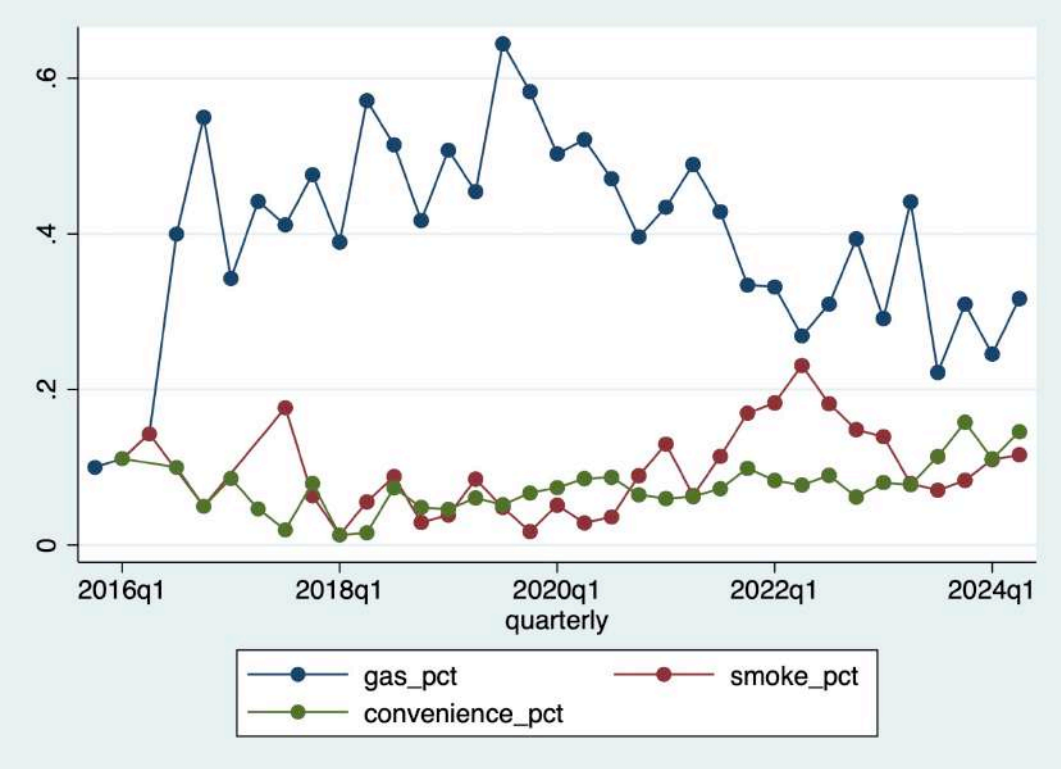


Figure 6: Crime Classification by Payment Relevance (3-Tier)

Notes: Each bar represents one of the 51 NIBRS offense codes, grouped by payment relevance to BTM infrastructure. High (14 codes): crimes whose financial logistics directly exploit BTM payment functions. Moderate (17 codes): crimes facilitated by BTM physical presence or whose proceeds can be laundered through BTMs. None (20 codes): crimes with no plausible BTM mechanism (placebo).



Figure 7: Crime Classification by BTM Technological Trait Exploited (5-Group)

Notes: Each bar represents one of the 51 NIBRS offense codes, grouped by the specific BTM technological trait exploited. Irreversibility (6 codes): non-reversible transactions enable fraud and extortion. Pseudonymity (7 codes): anonymous value transfer enables drug markets, ransomware, and trafficking. Cash-to-Crypto (10 codes): stolen proceeds laundered via BTM conversion. Cash Attraction (8 codes): BTM locations concentrate cash-carrying targets. Not Relevant (20 codes): placebo group identical to the 3-tier None category.

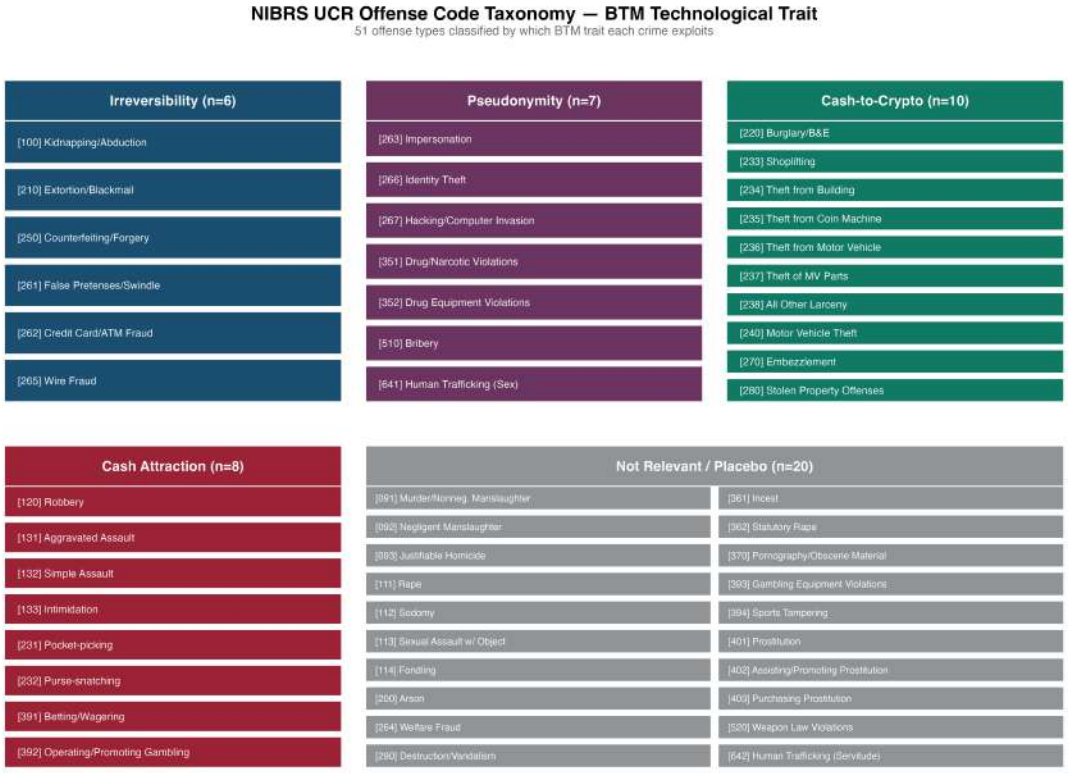


Figure 8: Event Study: Effect of BTM Entry on Total Crime

Notes: Callaway-Sant'Anna (2021) difference-in-differences with doubly-robust inverse probability weighting. Three panels show effects on crime count, rate per 10,000 county population, and log-transformed count. Comparison group is not-yet-treated counties. Estimates represent ATT at each event time. Vertical dashed line at $t = -1$ indicates reference period (normalized to zero). Shaded bands show 95% simultaneous confidence intervals constructed via multiplier bootstrap. Sample: 10,488 county-year observations, NIBRS 2013–2023. County-level clustering. BTM = Bitcoin automated teller machine.

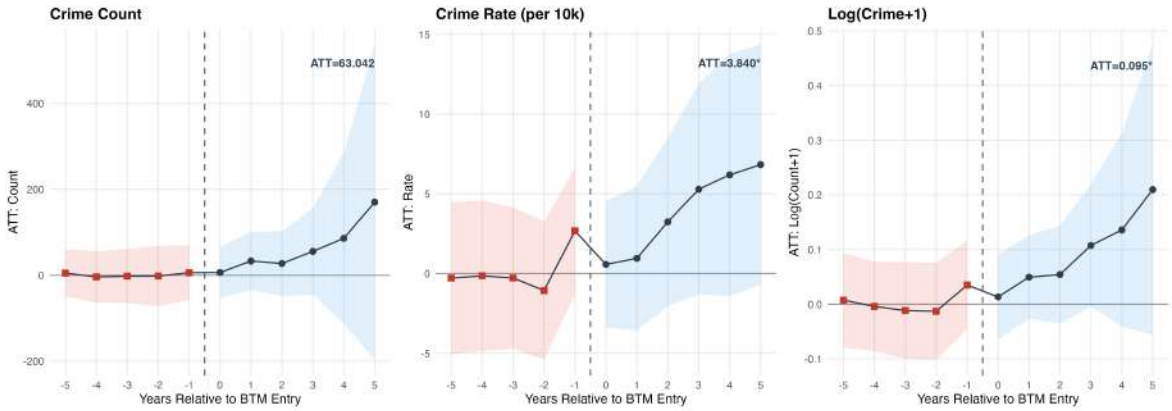


Figure 9: Event Study: Effect of BTM Entry on Crime by Payment Relevance

Notes: Three panels showing event-study estimates for the average-of-logs dependent variable ($\bar{Y}_{it} = J^{-1} \sum_j \log(\text{crime}_j + 1)$) by payment relevance tier. Callaway–Sant'Anna (2021) with not-yet-treated comparison group, doubly-robust IPW, unbalanced panel. Shaded bands show 95% confidence intervals. Sample: 10,488 county–year observations, NIBRS 2013–2023. County-level clustering.

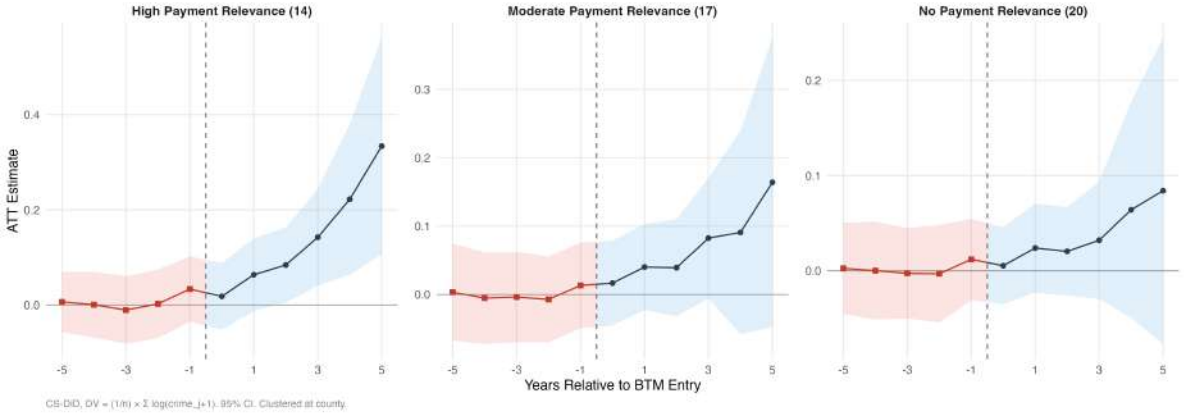


Figure 10: Event Study: Effect of BTM Entry on Crime by Technological Trait

Notes: Five panels showing event-study estimates for the average-of-logs dependent variable ($\bar{Y}_{it} = J^{-1}\sum_j \log(\text{crime}_j + 1)$) by BTM technological trait exploited. Callaway–Sant'Anna (2021) with not-yet-treated comparison group, doubly-robust IPW, unbalanced panel. Shaded bands show 95% confidence intervals. Sample: 10,488 county–year observations, NIBRS 2013–2023. County-level clustering.

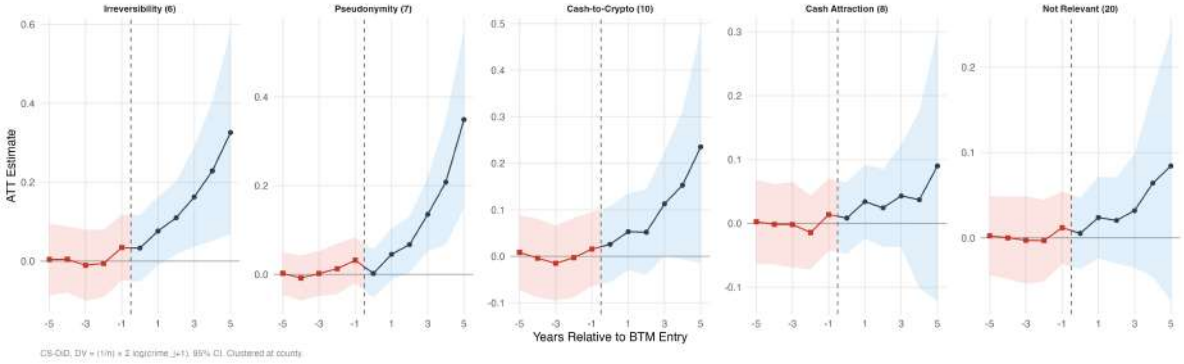


Figure 11: Crime Response by County KYC Permissiveness

Notes: County anonymity is the average first-tier KYC threshold across the kiosks in the county. Treated counties are split into terciles, cuts at $787 and $1,126. Panel A plots event-study coefficients of the log crime rate relative to the year before first BTM entry, estimated separately within each tercile with county and year fixed effects. Panel B plots coefficients on bins of the county's kiosk count relative to no-kiosk county-years, by tercile. Standard errors clustered at the county level. Sample is the NIBRS county-year panel, 2013–2023. Shaded bands are 95% CIs.

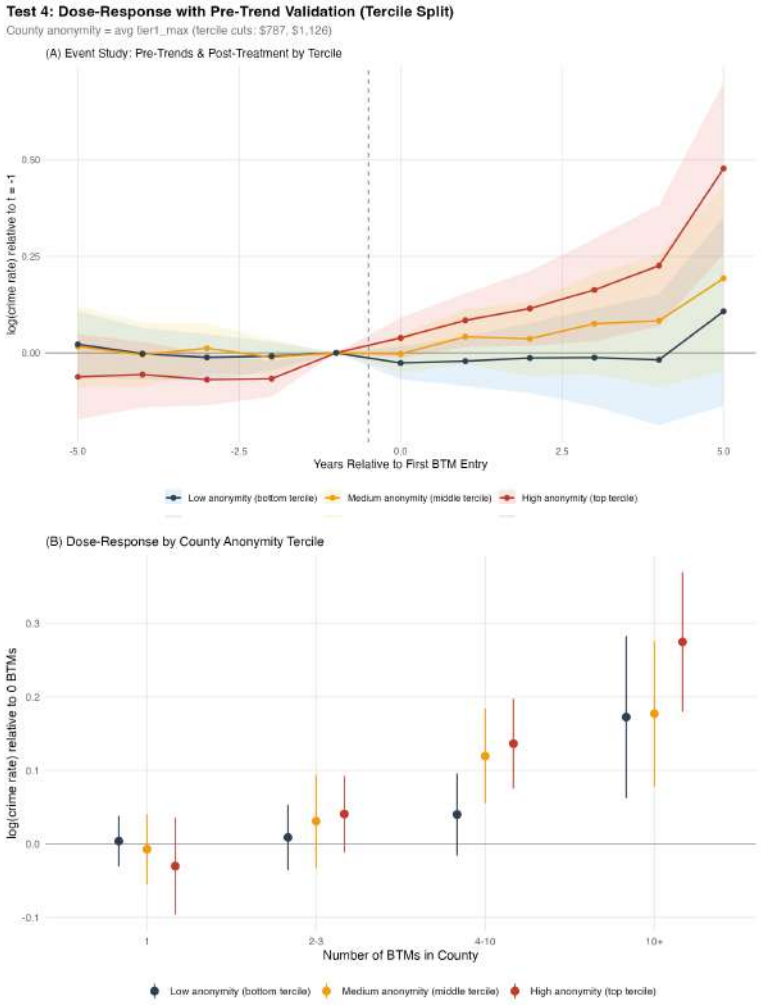


Figure 12: Average Marginal Effects on the Probability of Bitcoin ATM Entry (Cross-Sectional Logit, With Crime Proxy)

Notes: This figure plots average marginal effects from the cross-sectional logit model of Bitcoin ATM entry with demographic, economic, financial infrastructure, and local market controls, state fixed effects, and the log jail population rate (per 100,000) as a crime proxy (Column 4 of Table 7). The figure displays a subset of the covariates reported in the table. Jail population data are from the Vera Institute of Justice Incarceration Trends dataset, covering approximately 96% of U.S. counties. The dependent variable equals one if the county installs at least one BTM during 2015–2023. Covariates are standardized and measured in 2014. Bars show 95% confidence intervals from standard errors clustered at the state level. $N = 3{,}893$ counties; 2,807 with BTM entry.

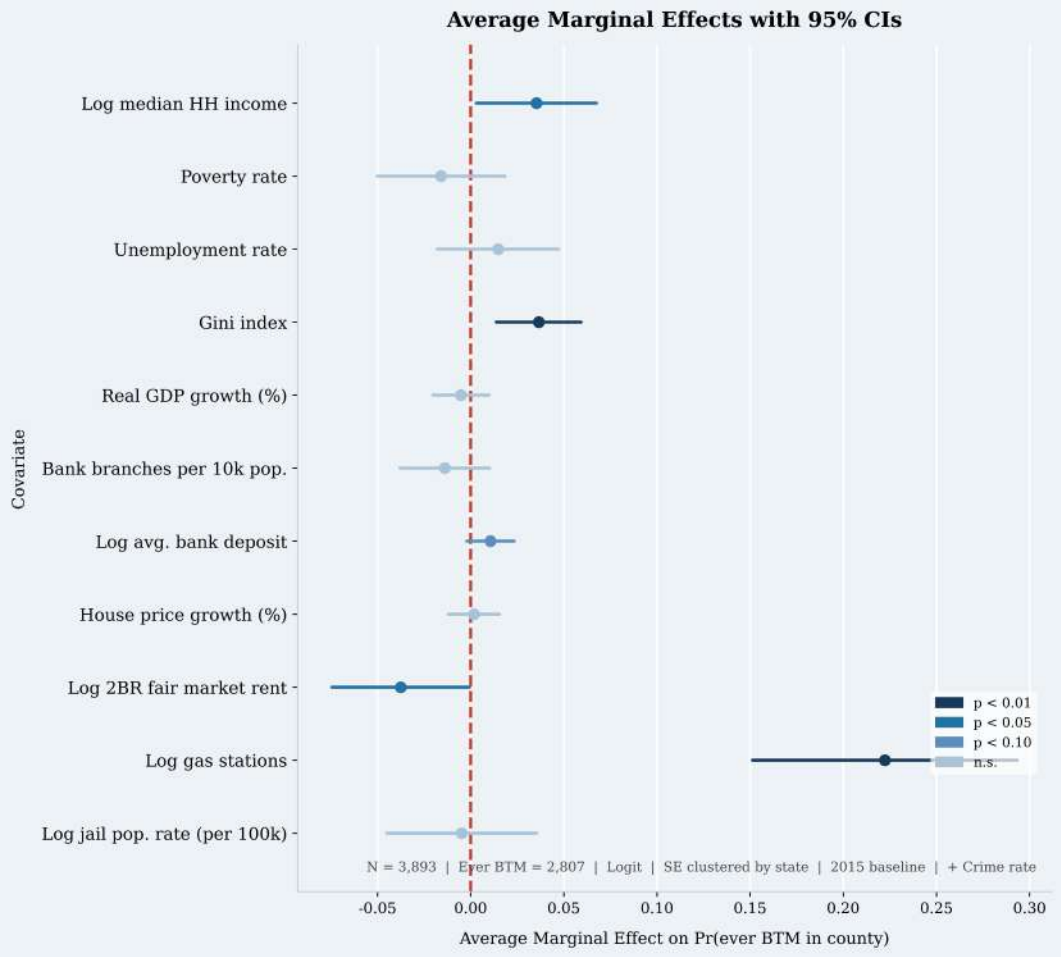


Figure 13: Standardized Coefficients for the Timing of First Bitcoin ATM Entry (Discrete-Time Hazard, With Crime Proxy)

Notes: This figure plots standardized coefficients from a complementary log-log discrete-time hazard model of first BTM entry with demographic, economic, financial infrastructure, and local market controls, state fixed effects, and the log jail population rate (per 100,000) as a crime proxy (Column 8 of Table 8). The figure displays a subset of the covariates reported in the table. Jail population data are from the Vera Institute of Justice Incarceration Trends dataset, covering approximately 96% of U.S. counties. Covariates are time-varying and lagged one year. Bars show 95% confidence intervals from standard errors clustered at the state level. N = 25,180 county–years; 2,342 entry events.

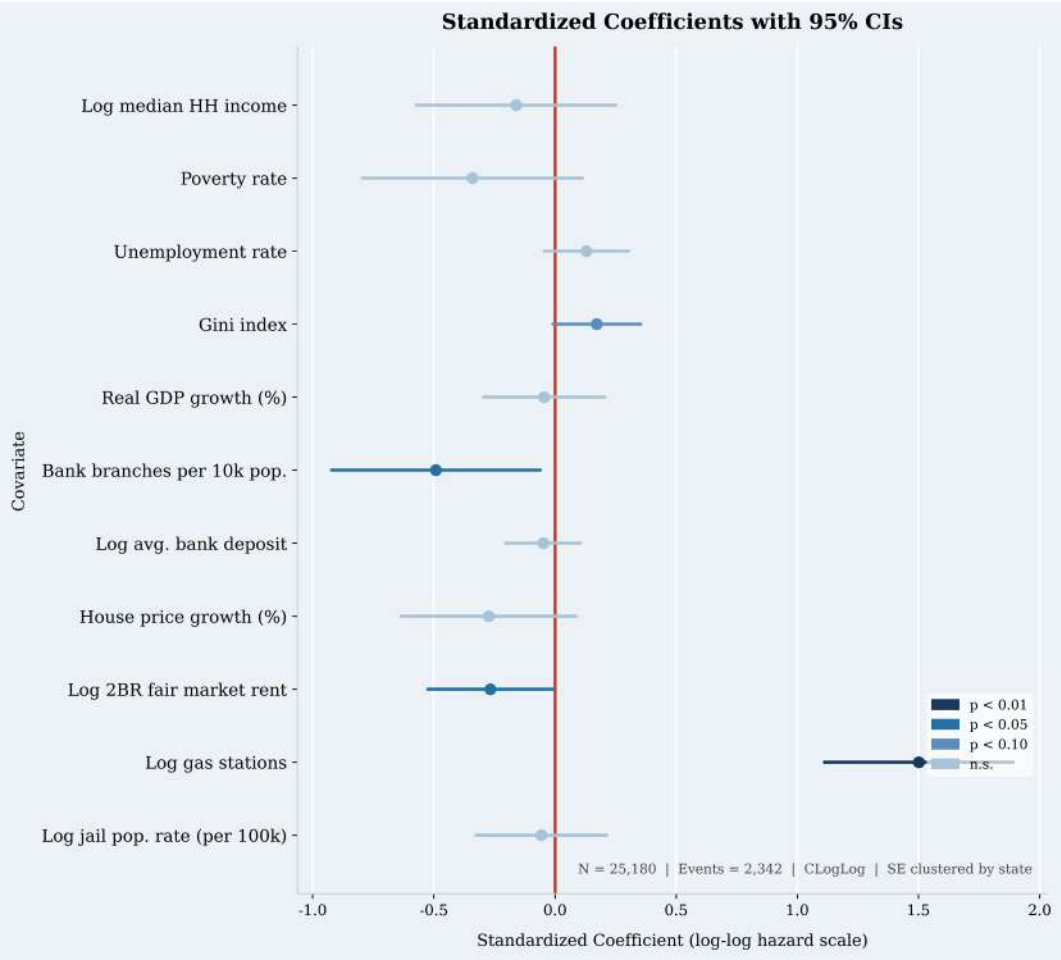


Figure 14: First-Stage Relationship: Bartik Instrument and BTM Count

Notes: This figure displays the first-stage coefficient of the Bartik instrument $Z_{it}^{\text{Bartik}} = s_i \times g_{-i,t}$ on county-level cumulative BTM count across three specifications: (1) county + year FE, (2) adding log population × year interactions, and (3) adding state × year FE. Point estimates with 95% confidence intervals are shown. The instrument is constructed from each county's 2005 share of the national stock of gas station and convenience store establishments (NAICS 447110, QCEW) interacted with the leave-one-out national BTM count from the full universe of 1,880 treated counties. Kleibergen-Paap F-statistics range from 59.5 to 98.5 across specifications. Standard errors clustered at the county level.

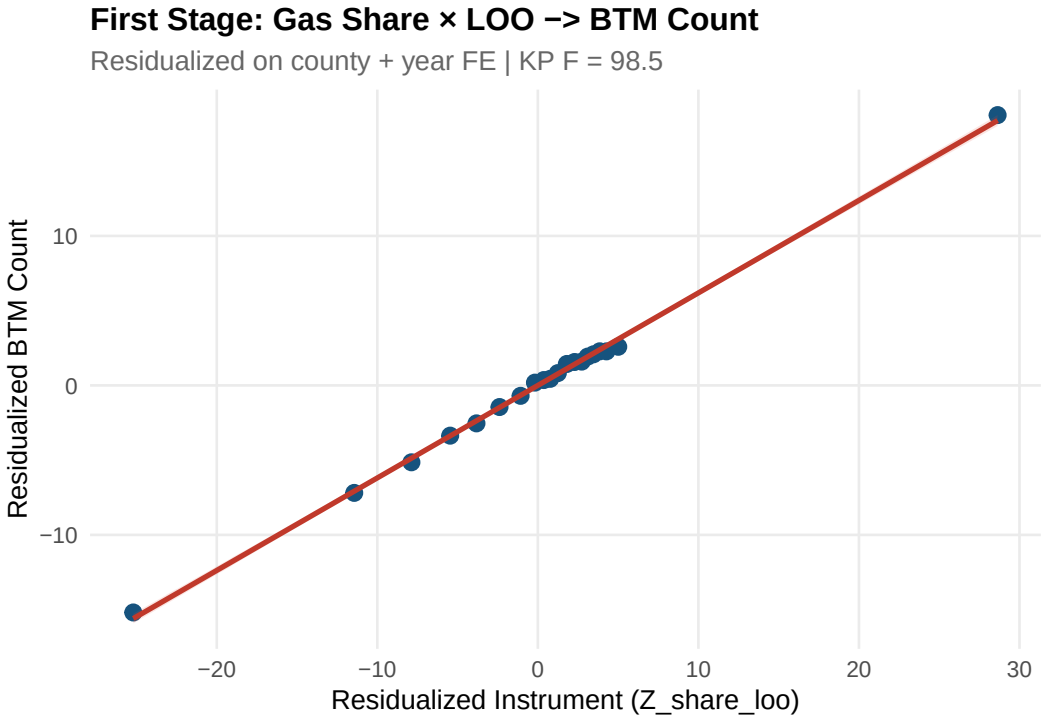


Figure 15: Event Study: OLS vs. 2SLS Effect of BTM Count on Crime Rate

Notes: This figure plots event-study coefficients from the dose-response specification $\text{Crime}_{it} = \sum_{\tau=-5}^{5} \beta_\tau \cdot \mathbf{1}[\text{rel_time} = \tau] \cdot \text{dose}_i + \alpha_i + \gamma_t + \varepsilon_{it}$, where $\text{dose}_i = \log(1 + \max \text{BTM}_i)$ is the time-invariant eventual treatment intensity. OLS (red circles) uses log dose directly; 2SLS (blue triangles) instruments the dose interactions with $Z_{it} = \text{gas_conv}_{i,2005} \times g_{-i,t}/1000$, where $g_{-i,t}$ is the leave-one-out national BTM count computed from all 1,880 treated counties. The reference period is $t = -1$ (normalized to zero). Shaded ribbons show 95% confidence intervals. Error bars show ± 1.96 standard errors. Vertical dashed line at $t = -0.5$ separates pre- and post-treatment periods. Kleibergen-Paap first-stage $F = 98.5$. Sample: 10,488 county-year observations covering 838 treated and 207 never-treated counties, NIBRS 2013–2023. Standard errors clustered at the county level.

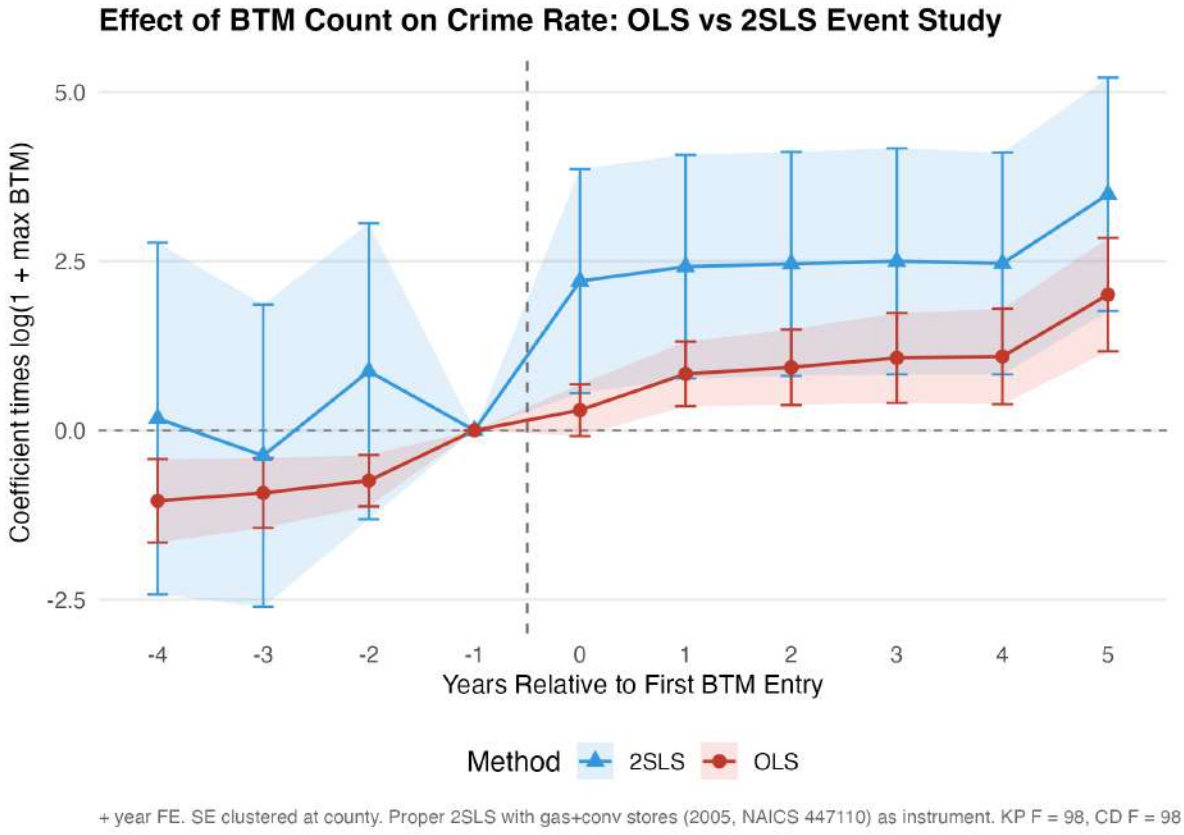
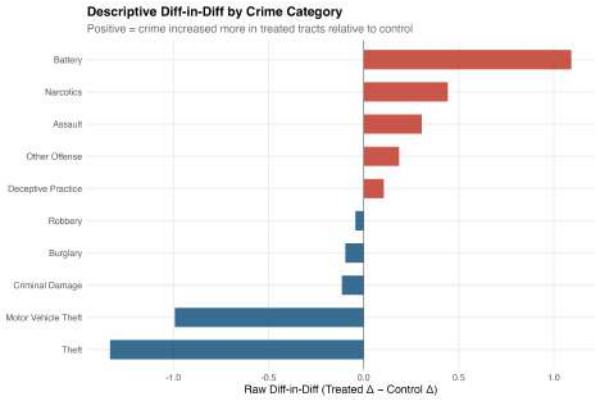
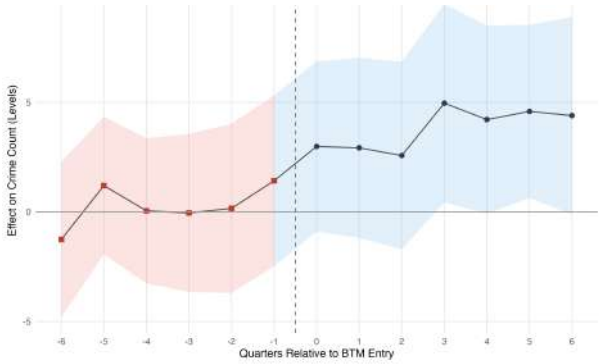


Figure 16: BTM Entry and Crime: Category Decomposition and Tract-Level Event Study

Notes: Left panel: descriptive difference-in-differences by crime category, showing the change in crime counts in treated tracts relative to control tracts. Right panel: event-study estimates for total crime at the Census tract–quarter level. Specification includes tract and quarter fixed effects with standard errors clustered at the tract level. Sample comprises 858 Chicago Census tracts observed 2012–2024. Confidence intervals at 95% level.



(a) DiD by Crime Category



(b) Event Study: Tract Level

Figure 17: Event Study: BTM Entry and Crime at Police Beat–Quarter Level

Notes: Event-study estimates for total crime at the police beat–quarter level. Specification includes beat and quarter fixed effects with standard errors clustered at the beat level. Sample comprises 277 Chicago police beats observed 2012–2024. Confidence intervals at 95% level.

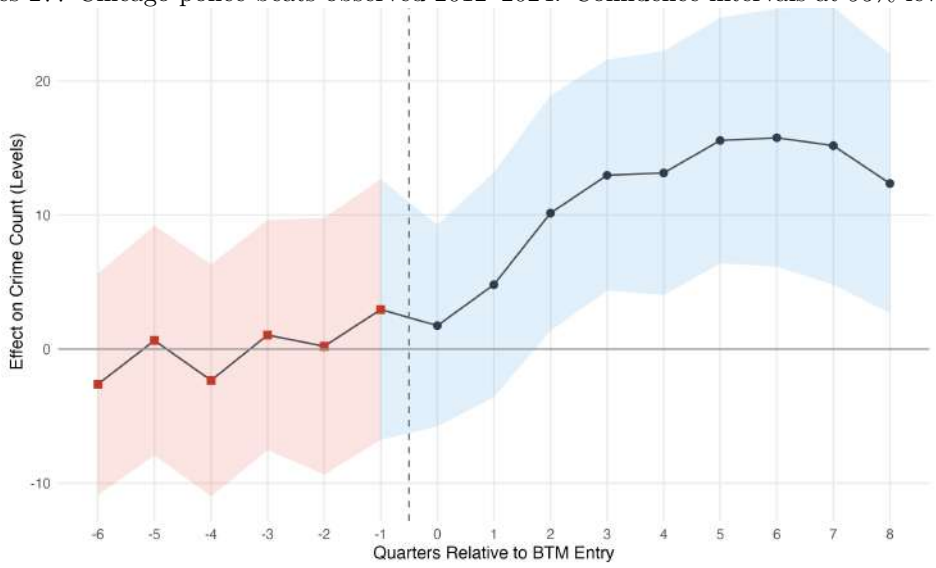


Figure 18: Population Sorting: Demographic Composition Event Studies

Notes: Four-panel event study showing effects of Bitcoin ATM entry on population demographic composition by Census tract and year. Panels display coefficients on share of population that is White, Black, Hispanic/Latino, and Male. Specifications use either Callaway–Sant'Anna (2021) or BJS imputation method. Data from American Community Survey 2015–2023, conducted annually. Specifications include Census tract fixed effects and year fixed effects. Standard errors clustered at tract level. None of the demographic effects are statistically significant at conventional levels.

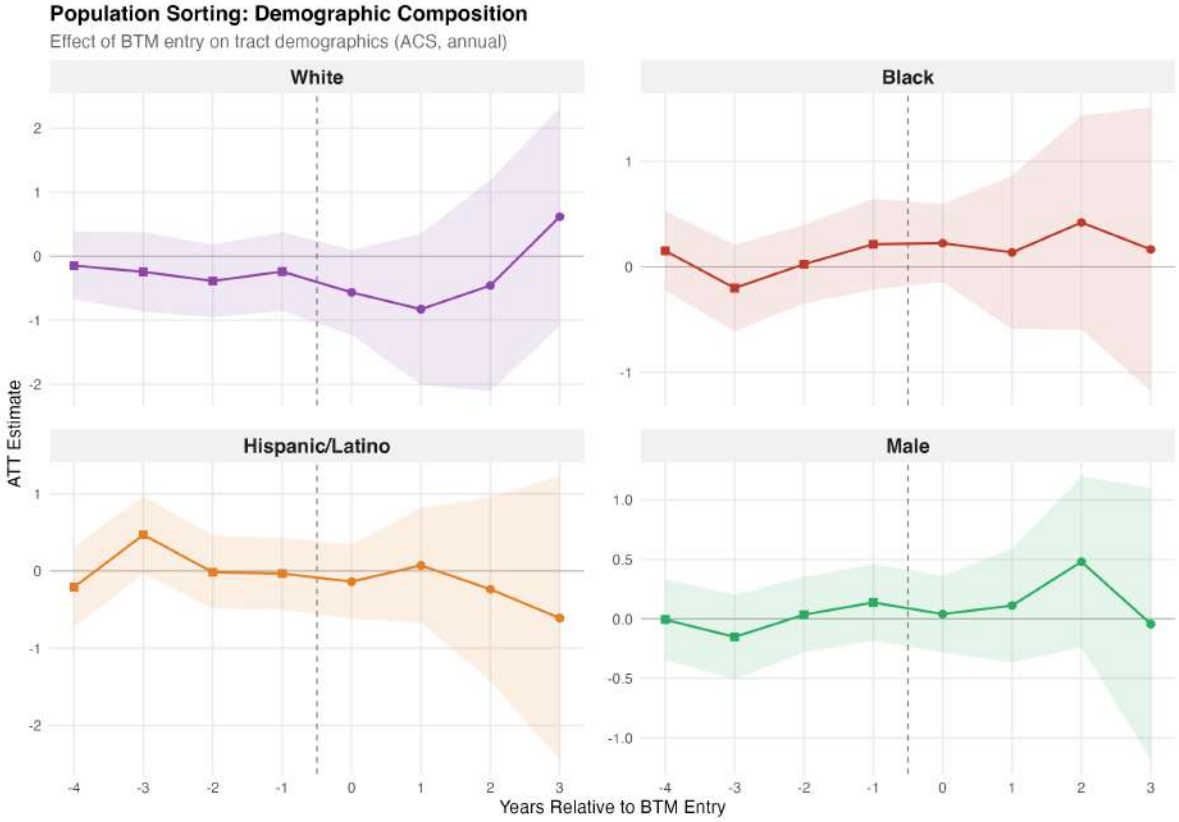


Figure 19: Event Study: Investigatory Stops Following BTM Entry

Notes: Event study showing effects of Bitcoin ATM entry on total investigatory stops and frisks (ISR) at the police beat–quarter level. Specification uses Callaway–Sant'Anna (2021) or TWFE event study estimator. Includes police beat fixed effects and quarter fixed effects with standard errors clustered at the beat level. Data from Chicago Police Department ISR records 2016–2024 (Q1–Q4). Confidence intervals at 95% level. Sample covers police beats with observed ISR activity in sample period.

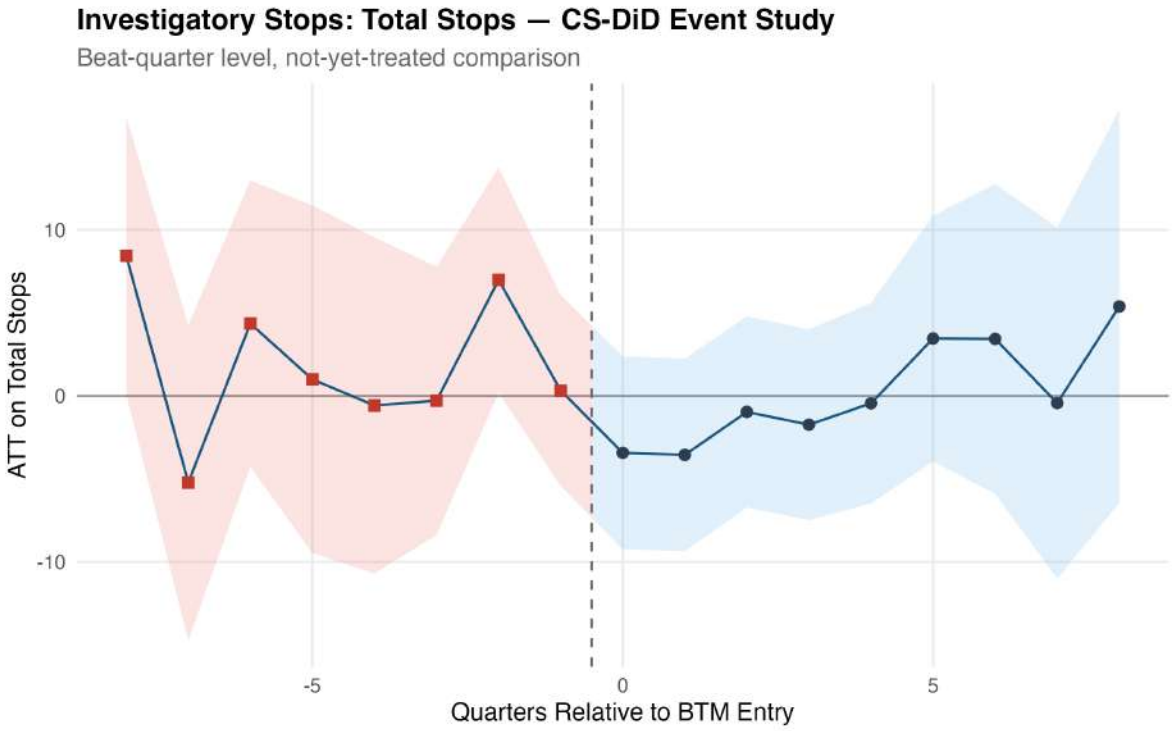


Figure 20: Investigatory Stops by Demographic Group

Notes: Four-panel event study showing effects of Bitcoin ATM entry on investigatory stops and frisks disaggregated by demographic group. Panels display separate event study coefficients for stops of Black, White, Hispanic/Latino, and Male subjects. Specifications use Callaway–Sant'Anna (2021) or TWFE estimator. Data at police beat–quarter level from Chicago Police Department ISR records 2016–2024. Includes police beat fixed effects and quarter fixed effects with standard errors clustered at the beat level. Confidence intervals at 95% level.

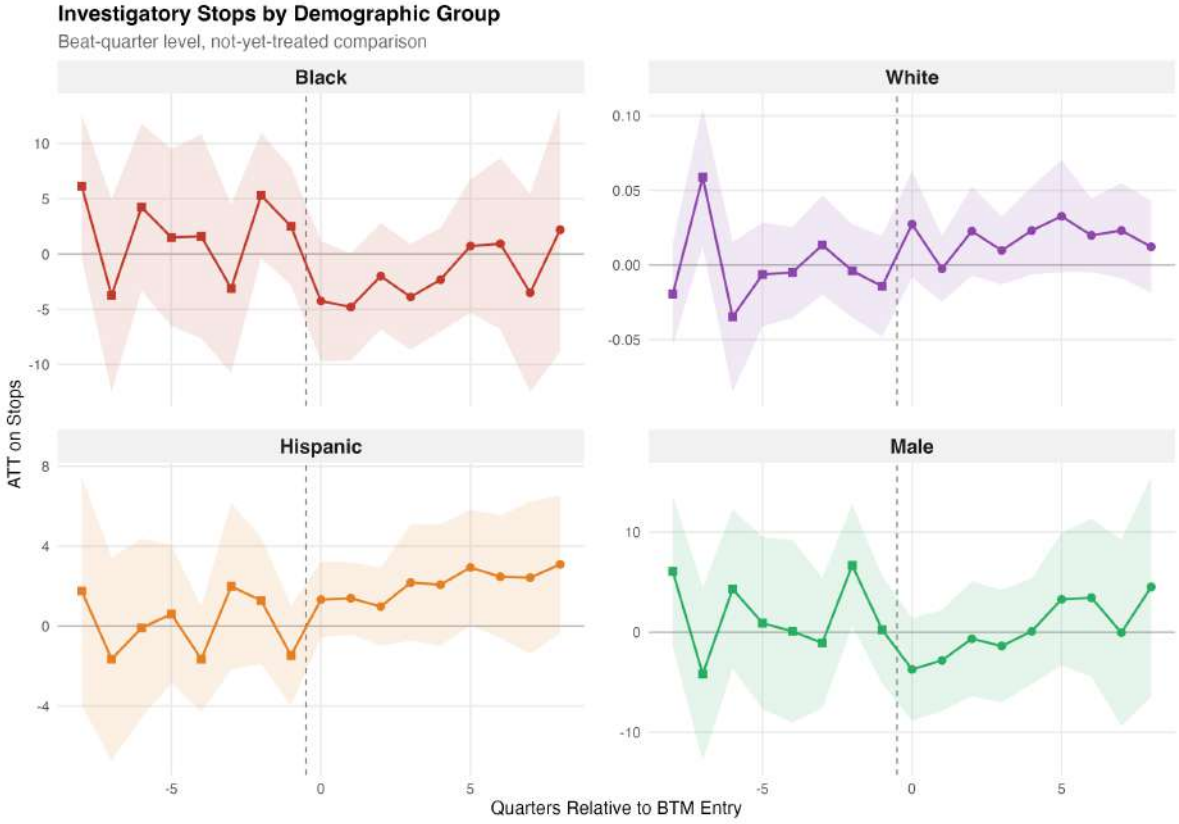


Figure 21: Spatial Reallocation Test: Near vs. Far Control Group Event Studies

Notes: Event study estimates of the effect of Bitcoin ATM entry on ln(crime + 1) in treated tracts, using two distinct control groups. "Near Control" (red triangles) consists of untreated tracts that share a boundary with a treated tract. "Far Control" (blue circles) consists of untreated tracts located more than 3 km from any treated tract, matched to treated tracts via coarsened exact matching (CEM) on pre-treatment crime quintiles (2013–2015). Each control tract inherits the treatment timing of its paired treated tract. The specification includes Census tract and quarter fixed effects; standard errors are clustered at the tract level. The reference period is $t = -1$. Under a siphon (displacement) hypothesis, the near-control effect should exceed the far-control effect post-treatment because crime is drawn away from nearby areas.

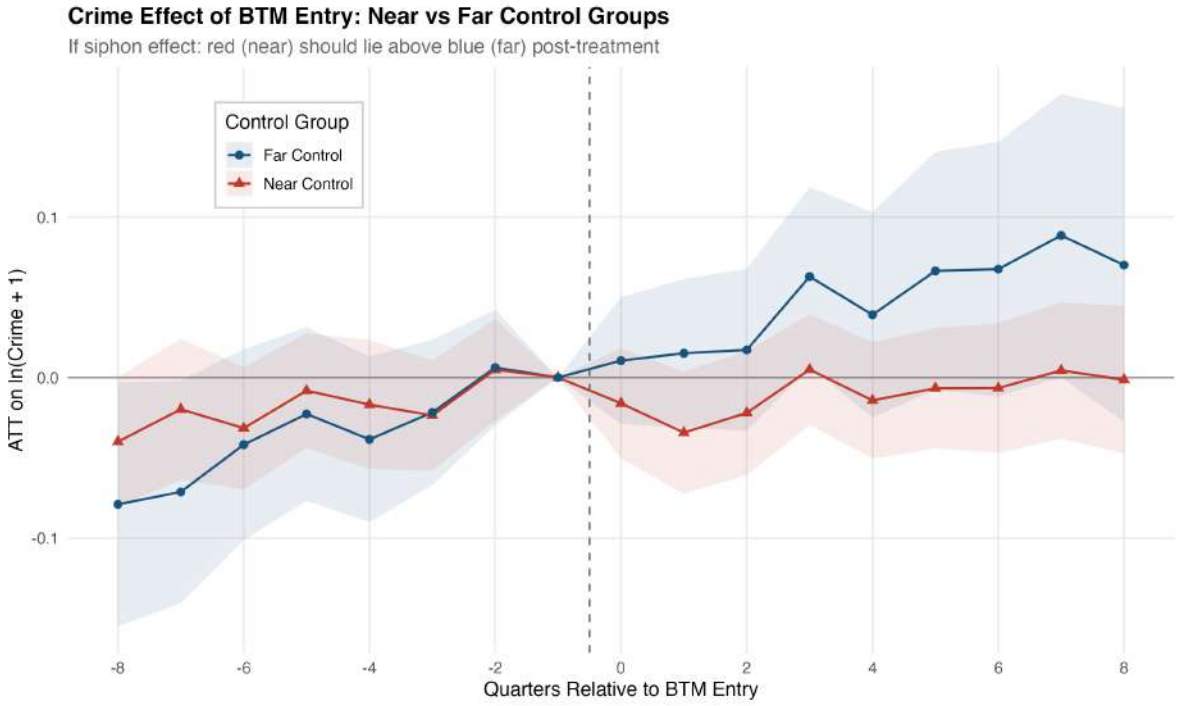


Figure 22: Spatial Distribution of BTM Exposure Across Chicago Census Tracts

Notes: Map of Chicago census tracts showing the average inverse-distance-weighted (IDW) exposure index for untreated tracts over the sample period (2012–2024). For each quarter q, exposure is $\text{Expo}_{i,q} = \sum_{j:\,\text{entry}_j \le q} 1/d_{ij}^2$, summing only over BTMs active by quarter q; the plotted value is the time-average $\bar{\text{Expo}}_i = T^{-1} \sum_q \text{Expo}_{i,q}$. Color intensity (plasma scale, log-transformed) indicates average spatial proximity to active BTMs. Red-shaded tracts are treated (received at least one BTM). Grey tracts have no exposure data. Projection: NAD83 / Illinois East (EPSG:26971).

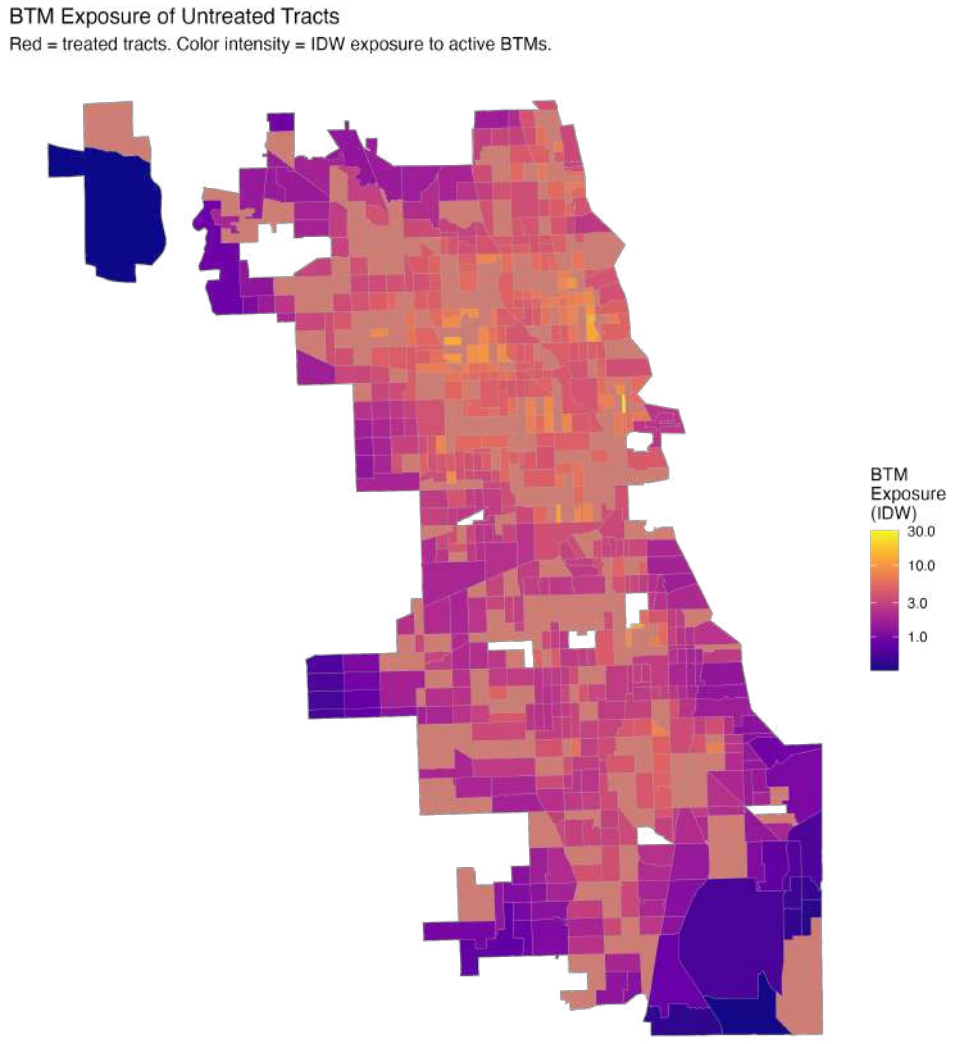


Figure 23: IDW Exposure Event Study: Marginal Effect of BTM Proximity on Crime

Notes: Event study of the marginal effect of BTM exposure on ln(crime + 1). Exposure is measured as an inverse-distance-weighted (IDW) index of active Bitcoin ATMs surrounding each Census tract, standardized to mean zero and unit standard deviation. The specification interacts event-time indicators with the continuous exposure measure: $\ln(\text{crime}_{it}+1) = \sum_\tau \beta_\tau \cdot \mathbf{1}[t - g_i = \tau] \cdot \text{Exposure}_i + \alpha_i + \gamma_t + \varepsilon_{it}$, where g_i is the treatment quarter for treated tracts or the nearest treated neighbor's treatment quarter for untreated tracts. Sample includes all Chicago Census tracts; Census tract and quarter fixed effects are absorbed; standard errors are clustered at the tract level. Coefficients represent the effect of one additional standard deviation of BTM exposure at each event time. Under a siphon hypothesis, post-treatment coefficients should be negative (high-exposure untreated tracts lose crime to nearby BTMs).

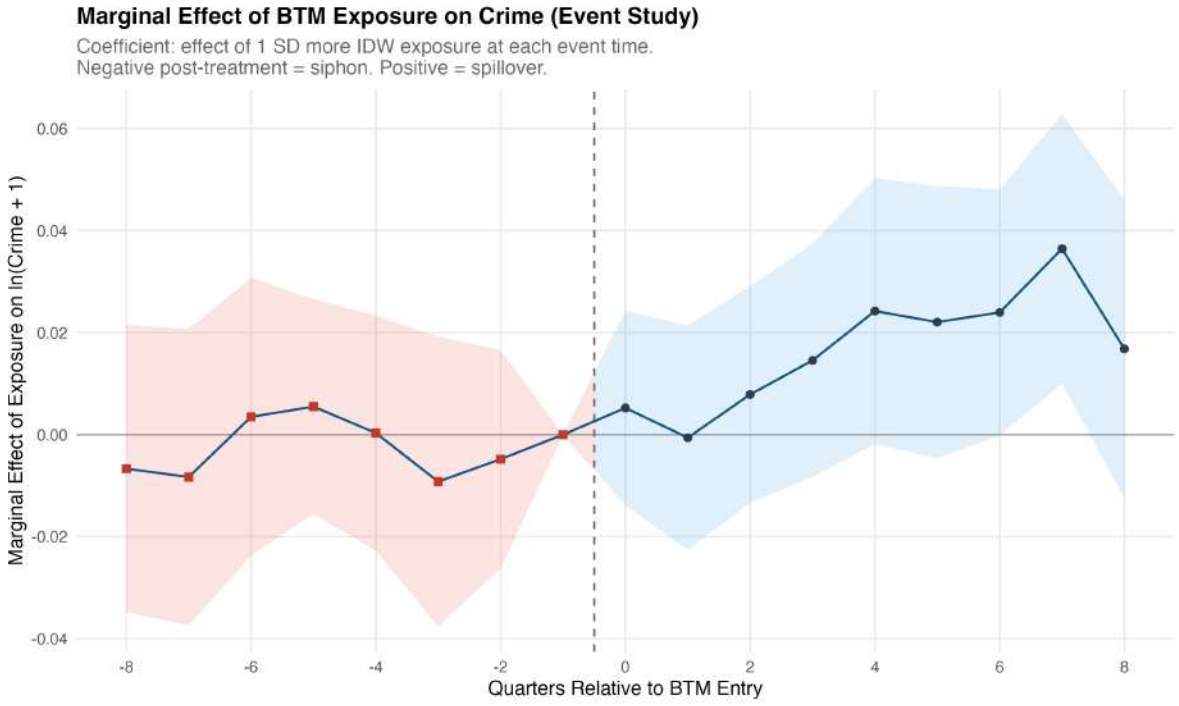


Figure 24: Event Study: Effect of BTM Entry on Drug Seizure Weight by Drug Type

Notes: Callaway-Sant'Anna (2021) staggered difference-in-differences estimates of BTM entry on drug seizure outcomes. Top row: extensive margin (probability of any seizure; ATT in percentage points). Bottom row: intensive margin (IHS-transformed grams per 10,000 population). "Hard" drugs include cocaine (incl. crack) and heroin; "soft" drugs include cannabis. NIBRS seizure data with unit conversion to grams (excluding dosage units and plants). Comparison group: not-yet-treated counties. Shaded bands show 95% confidence intervals. Reference period $t = -1$ normalized to zero. Sample: ≈10,400 county-year observations, NIBRS 2013–2023. Standard errors clustered at the county level.

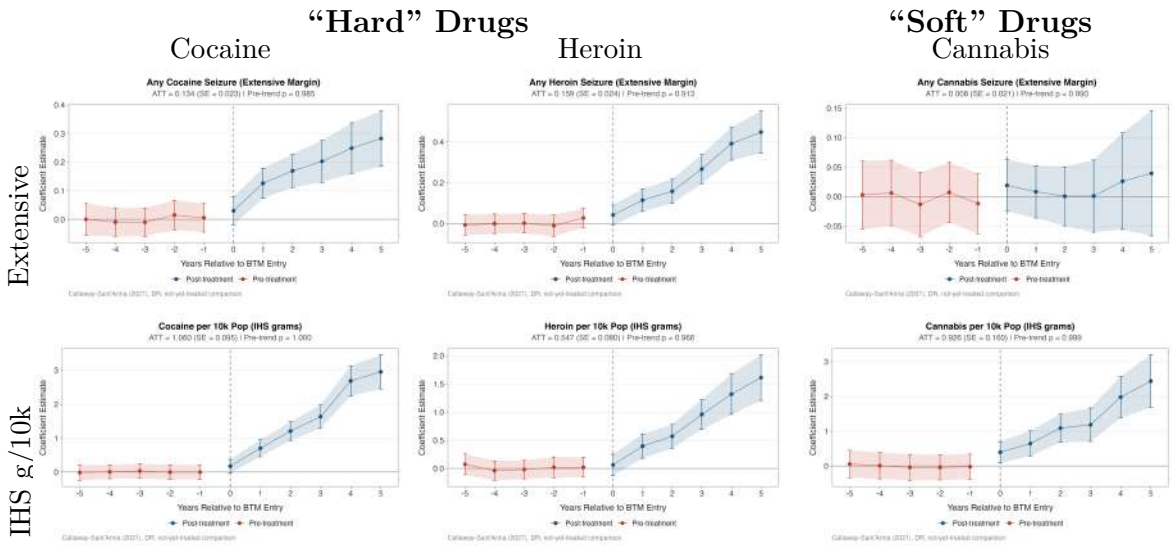


Figure 25: Effect of BTM Entry on Drug Seizure by Type: Occurrence and Weight

Notes: Coefficient plots showing Callaway-Sant'Anna (2021) ATT estimates with 95% confidence intervals. Left panel: extensive margin (any seizure > 0, probability in percentage points). Right panel: IHS-transformed grams per 10,000 population. Hard drugs: heroin, cocaine (incl. crack), methamphetamine/amphetamines. Soft drugs: cannabis. Comparison group: not-yet-treated counties. Standard errors clustered at the county level. Data: NIBRS drug seizure records, 2013–2023.

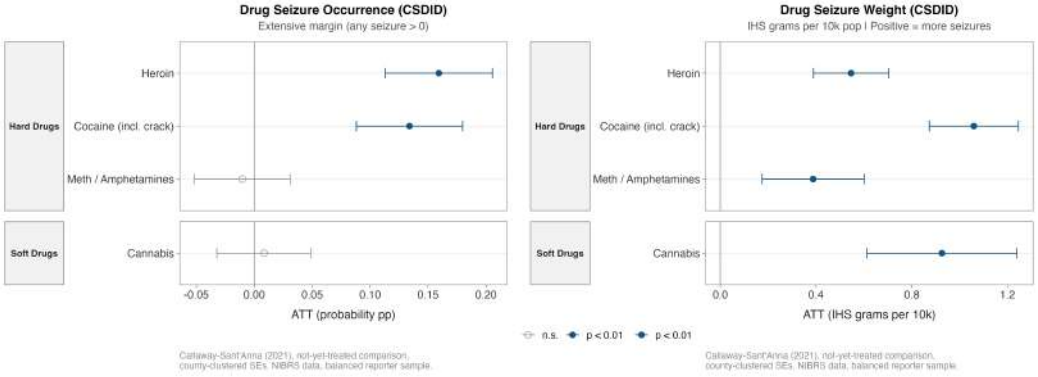


Figure 26: Event Study: Poisson QMLE Estimates of BTM Entry on Drug Seizure Weight

Notes: Poisson quasi-maximum likelihood estimates of BTM entry on drug seizure weight (grams per 10,000 population), following Chen & Roth (2024, *QJE*). Poisson QMLE handles zero-inflated outcomes natively without log or IHS transformation. County and year fixed effects; standard errors clustered at the county level. Coefficients on the log scale: $\exp(\hat{\beta}) - 1$ gives the proportional treatment effect on the mean. Outcomes winsorized at the 99th percentile to tame extreme outliers from large single busts. Never-treated counties included as comparison group with reference period $t = -1$ normalized to zero. Vertical dashed line at $t = -0.5$. Shaded bands show 95% confidence intervals. Sample: ≈10,400 county-year observations, NIBRS 2013–2023.

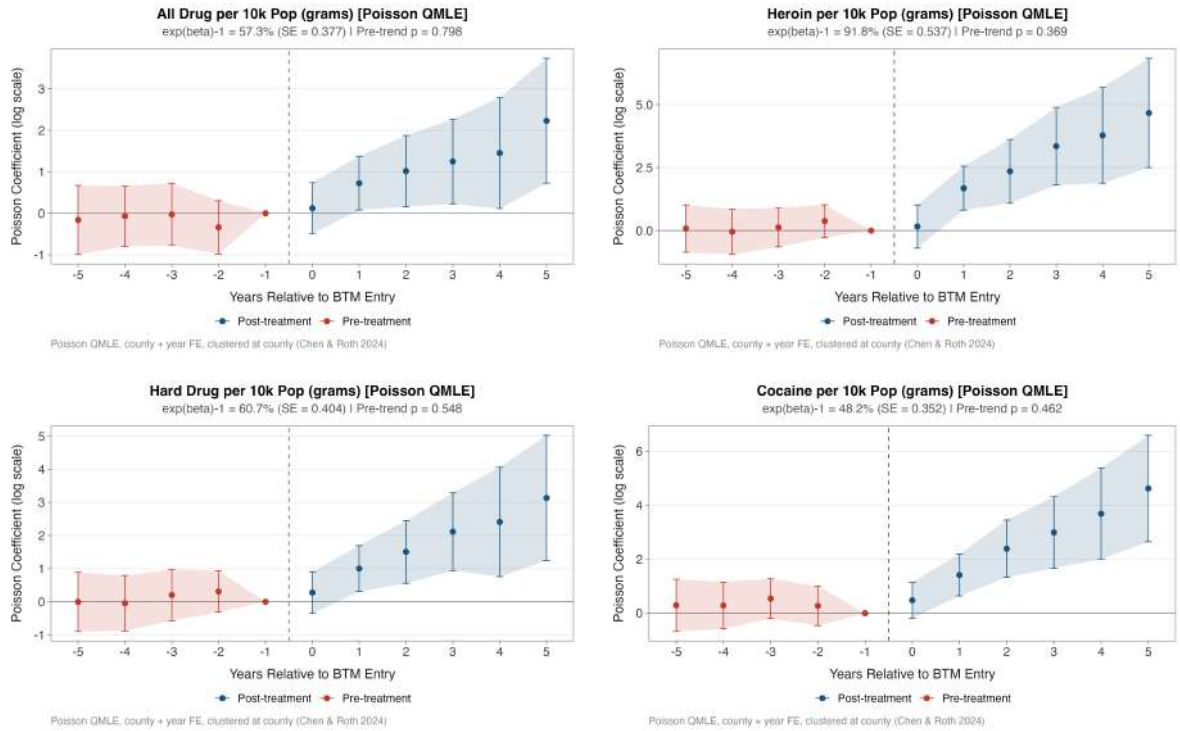


Figure 27: Event Study: Overdose Mortality — Additional Specifications

Notes: Callaway-Sant'Anna (2021) staggered DiD estimates. Top left: overdose death rate per 100,000 in levels. Top right: IHS-transformed back-calculated death counts. Bottom left: extensive margin (any overdose death in county-year). Bottom right: Poisson QMLE estimate following Chen and Roth (2024, QJE). NCHS model-based county estimates (2013–2021) supplemented by CDC provisional data (2022–2023). Comparison group: not-yet-treated counties. Shaded bands show 95% confidence intervals. Reference period $t = -1$ normalized to zero. Standard errors clustered at the county level.

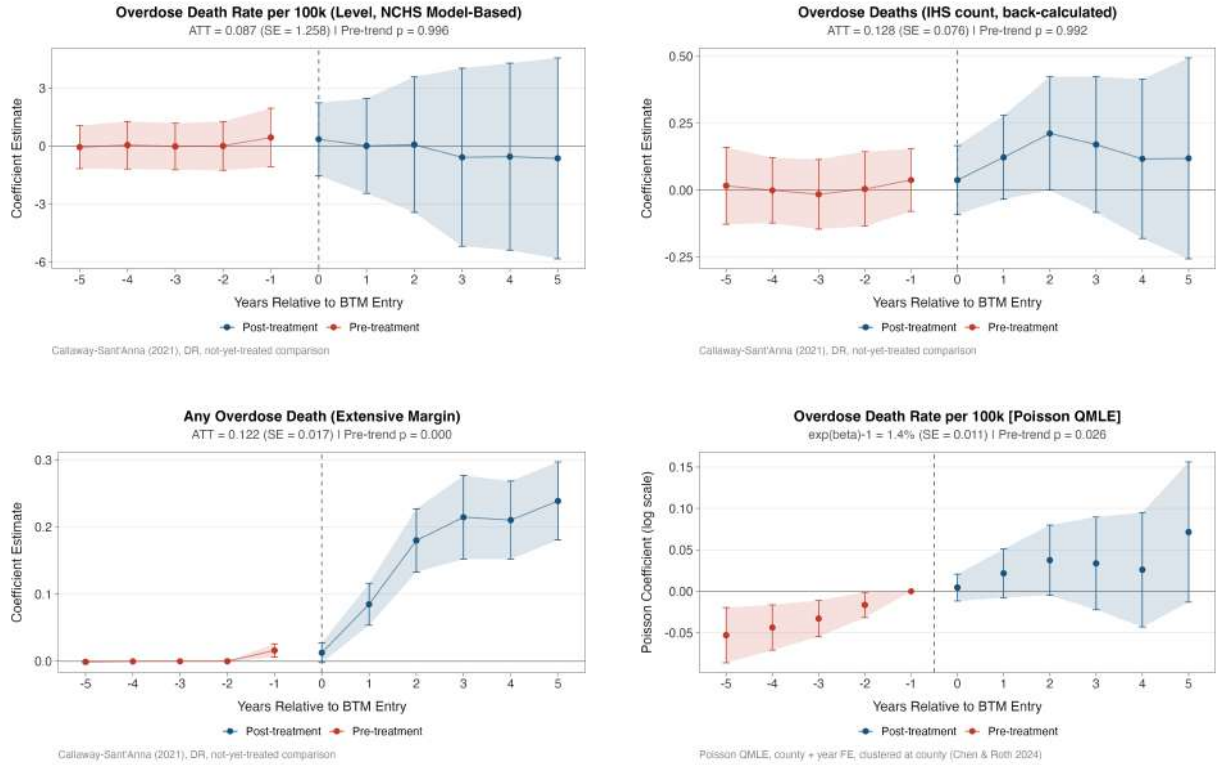


Figure 28: Event Study: Effect of BTM Entry on Overdose Death Rate

Notes: Left panel: Callaway-Sant'Anna (2021) staggered DiD estimates with IHS-transformed overdose death rate per 100,000 population as the dependent variable. Right panel: Poisson QMLE estimates following Chen and Roth (2024, QJE) with death rate in levels. NCHS model-based county estimates (2013–2021) supplemented by CDC provisional data (2022–2023). Comparison group: not-yet-treated counties. Shaded bands show 95% confidence intervals. Reference period $t = -1$ normalized to zero. Standard errors clustered at the county level.

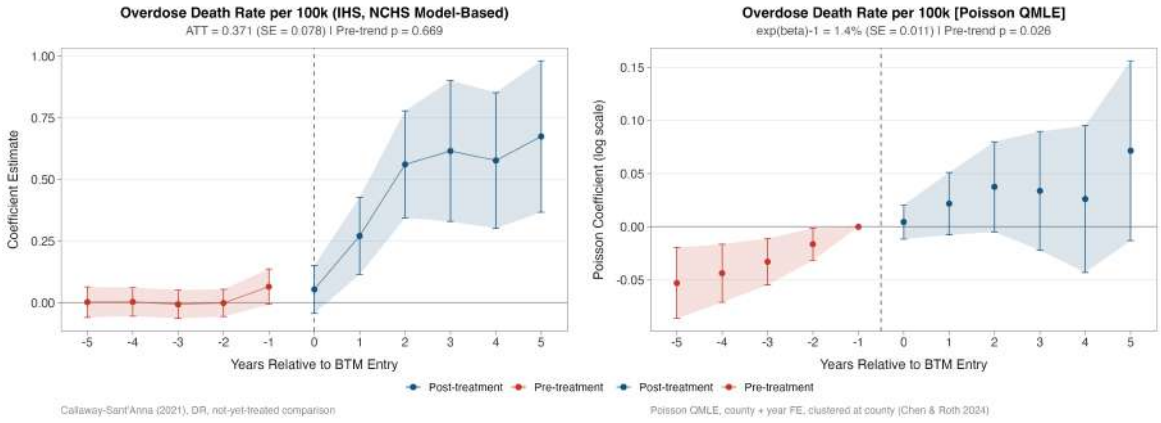


Figure 29: Effect of BTM Entry on Victim Composition (Share Outcomes)

Notes: Coefficient plot showing Callaway-Sant'Anna (2021) ATT estimates with 95% confidence intervals for victim composition share outcomes. Outcomes are percentage-point shares (not per-capita rates): positive values indicate the share of that victim category increases after BTM entry. Comparison group: not-yet-treated counties. Standard errors clustered at the county level. Filled circles indicate $p < 0.01$; triangles indicate $p < 0.10$; hollow circles indicate not significant. Data: NIBRS victim records, 2013–2023.

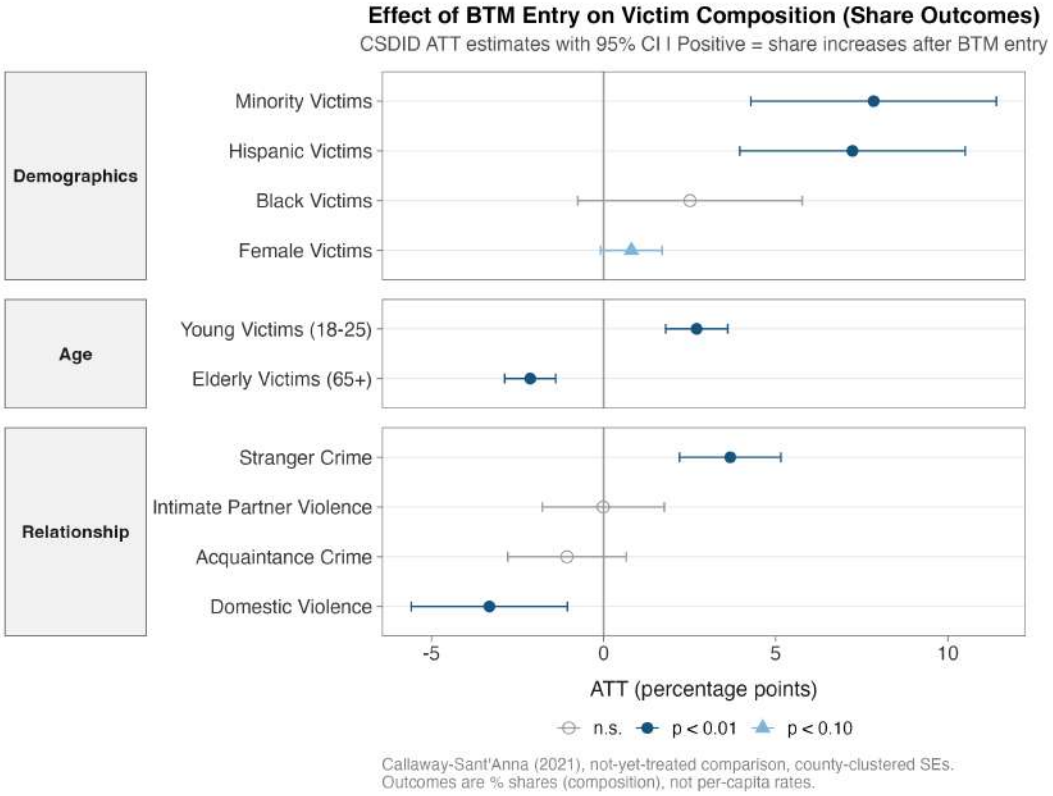


Figure 30: Proportional Effect of BTM Entry on Victim Rates by Group

Notes: Each point shows the proportional effect: ATT divided by baseline mean, multiplied by 100, representing the percentage increase in victimization rate relative to the pre-treatment mean. Because different victim groups have vastly different baseline rates, comparing level ATTs across groups is misleading. The proportional effect normalizes by the baseline mean, enabling comparable magnitudes across groups. Callaway-Sant'Anna (2021), not-yet-treated comparison, county-clustered standard errors. 95% CIs computed via delta method. Data: NIBRS victim records per 10,000 population, 2013–2023.

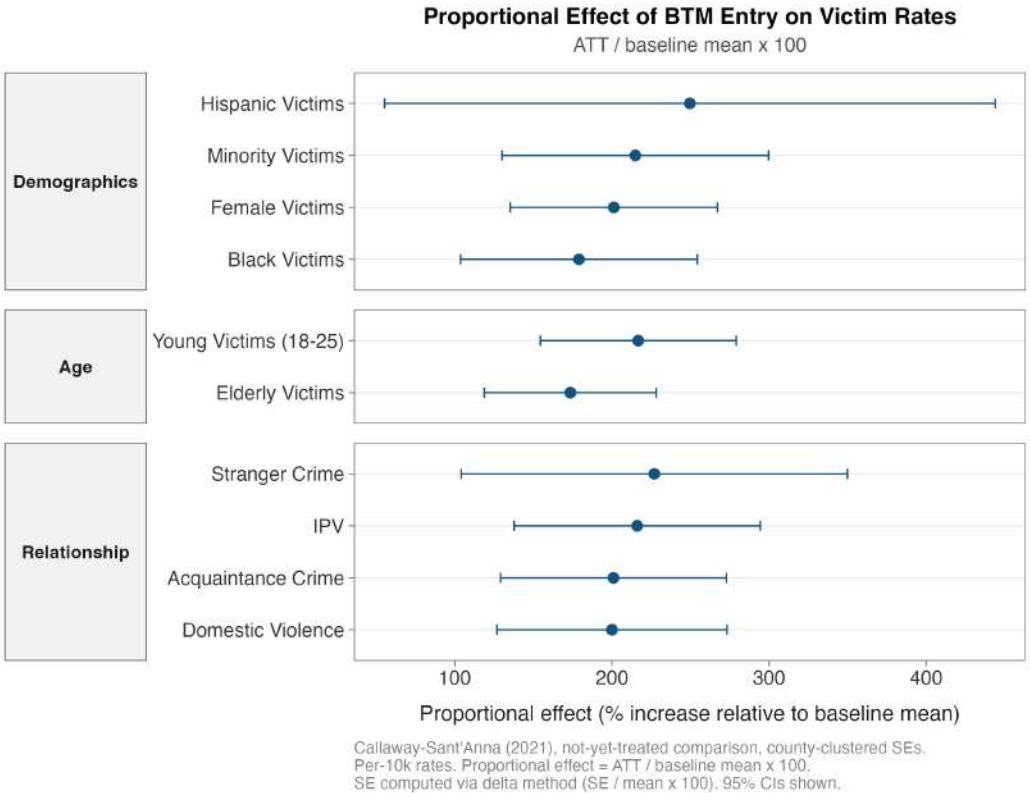


Figure 31: Stylized On-Chain Fund Flow from BTM Deposit to Exit

Notes: Stylized representation of on-chain fund flows originating from Bitcoin ATM deposits. Left: BTM operator cluster built from a proprietary set of seed addresses identified through common hot-wallet ownership (Bitcoin Depot data). Funds flow through intermediate hops ("peeling chains" that break address-level links), into obfuscation services (mixers, cross-chain bridges, or BTC-to-Monero conversion), and finally to exit-stage providers (non-compliant exchanges, darknet markets, or OTC desks). Node size represents approximate relative transaction volume. This figure illustrates the tracing strategy underlying the cross-stage elasticity estimation: shocks at any stage can be measured through volume changes at all other stages.

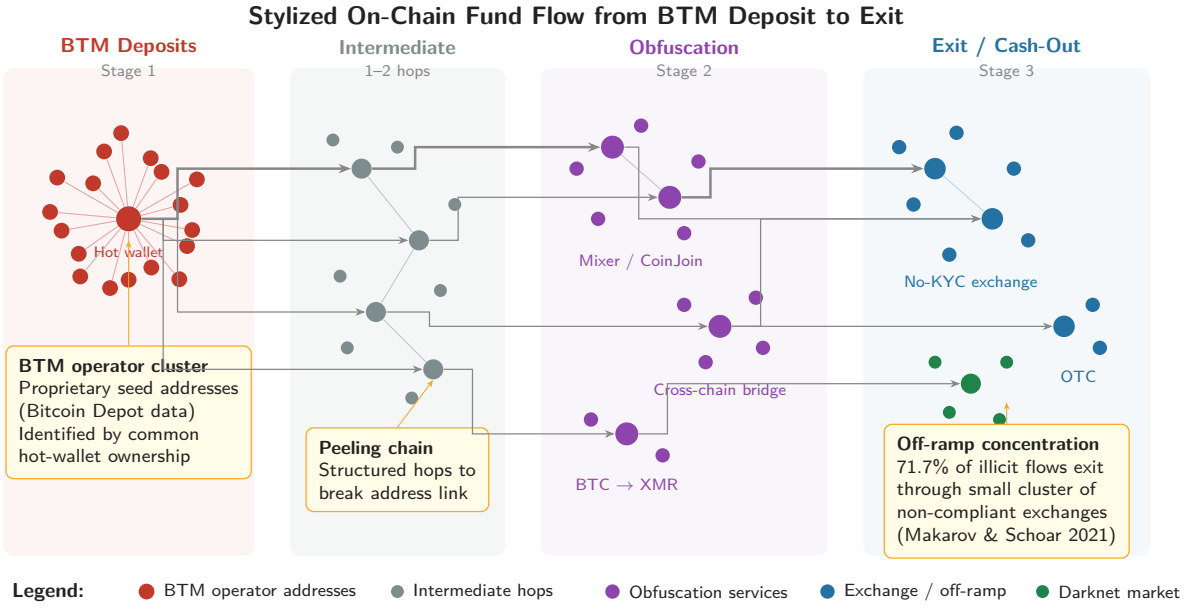


Figure 32: On-chain forensic analysis: identification, extraction, tracing, classification

Notes. Four-stage pipeline for the on-chain forensic reconstruction of Bitcoin Depot's customer activity (November 2018–May 2026). IDENTIFY: the operator seed cluster (proprietary blockchain-intelligence dataset) is extended via common-input-ownership clustering and a five-criteria change-detection rule to a depot cluster of 122,945 addresses. EXTRACT: 2.40 million customer payouts to 1.41 million validated wallets; CIO clustering aggregates these to 1.01 million persons; after excluding 34 mega-clusters of ≥1,000 wallets each (services and aggregators) the retail base is 1.17 million wallets / 970 thousand persons. External 10-K validation: Pearson $r = 0.967$ on the disclosed rolling 12-month median (2021Q1–2025Q3) and $r = 0.958$ on the disclosed quarterly crypto-cost (2023Q1–2025Q3). TRACE: 25.4 million hop-1 spending transactions recovered via BigQuery hash-JOIN. CLASSIFY: 3.82 million retail destinations joined against a 3.1-million-address consolidated label database assembled from nine public and proprietary sources (anonymized as A–I).

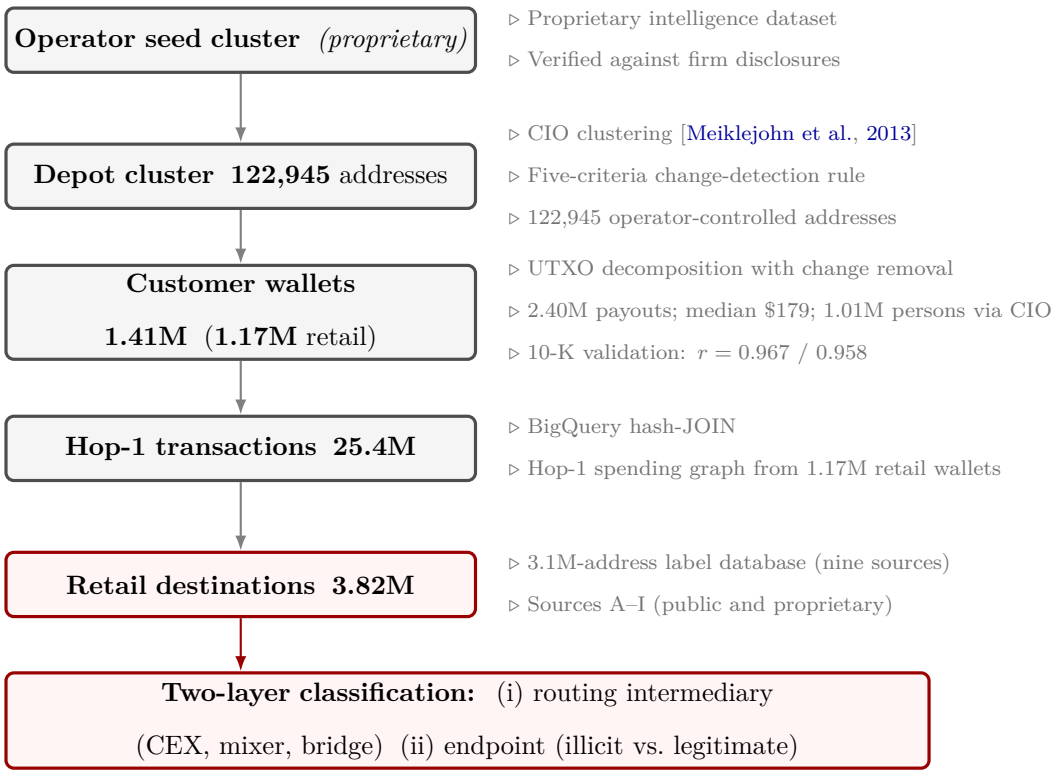


Figure 33: External validation: on-chain rolling 12-month median vs. Bitcoin Depot 10-K disclosed median

Notes. Quarterly median customer payout computed from the on-chain reconstruction (markup-adjusted to convert BTC cost into implied cash receipt at the daily BTC/USD price) is plotted against the rolling 12-month median kiosk transaction size that Bitcoin Depot discloses in each quarterly 10-K/10-Q filing. Window covers 19 quarters from 2021Q1 to 2025Q3. Pearson correlation is 0.967 and the median absolute deviation is 2.6% across the full window. Both series remain in the $140–$400 band throughout and move together in both level and trend. The alignment confirms that the on-chain cluster being reconstructed is the same customer-facing flow Bitcoin Depot reports to investors.

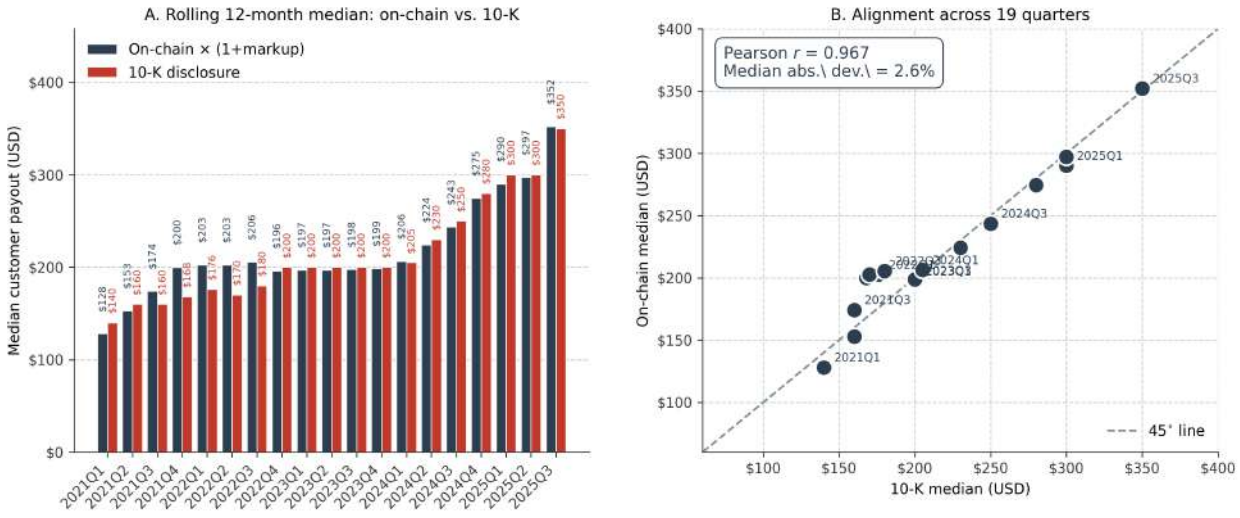


Figure 34: Customer behavior at the kiosk: depot-customer layer (Stage 1)

Notes. Four customer-level traits computed at both the wallet level (navy bars, N = 1,215,632 retail wallets) and the person level (red bars, N = 1,008,424 retail CIO clusters). Panel A is the bar histogram of the number of customer payouts received from Bitcoin Depot per customer, capped at 10+; 86% of wallets and 81% of persons receive exactly one payout. Panel B is a log-scale histogram of lifetime USD sum received per customer, with median \$179 / \$230 and 90th percentile \$3,775 / \$4,168. Panel C is the Lorenz curve of dollar-volume concentration, with customers ordered from highest to lowest by lifetime sum and the 45-degree line marking the equality reference; the top 10% of customers carry approximately 75% of all dollar volume. Panel D is a log-scale histogram of the span in days between first and last receipt, conditional on at least two payouts. Sample window 2018Q4–2026Q2; mega-clusters (34 persons with $\geq$ 1,000 wallets each) excluded as services or aggregators.

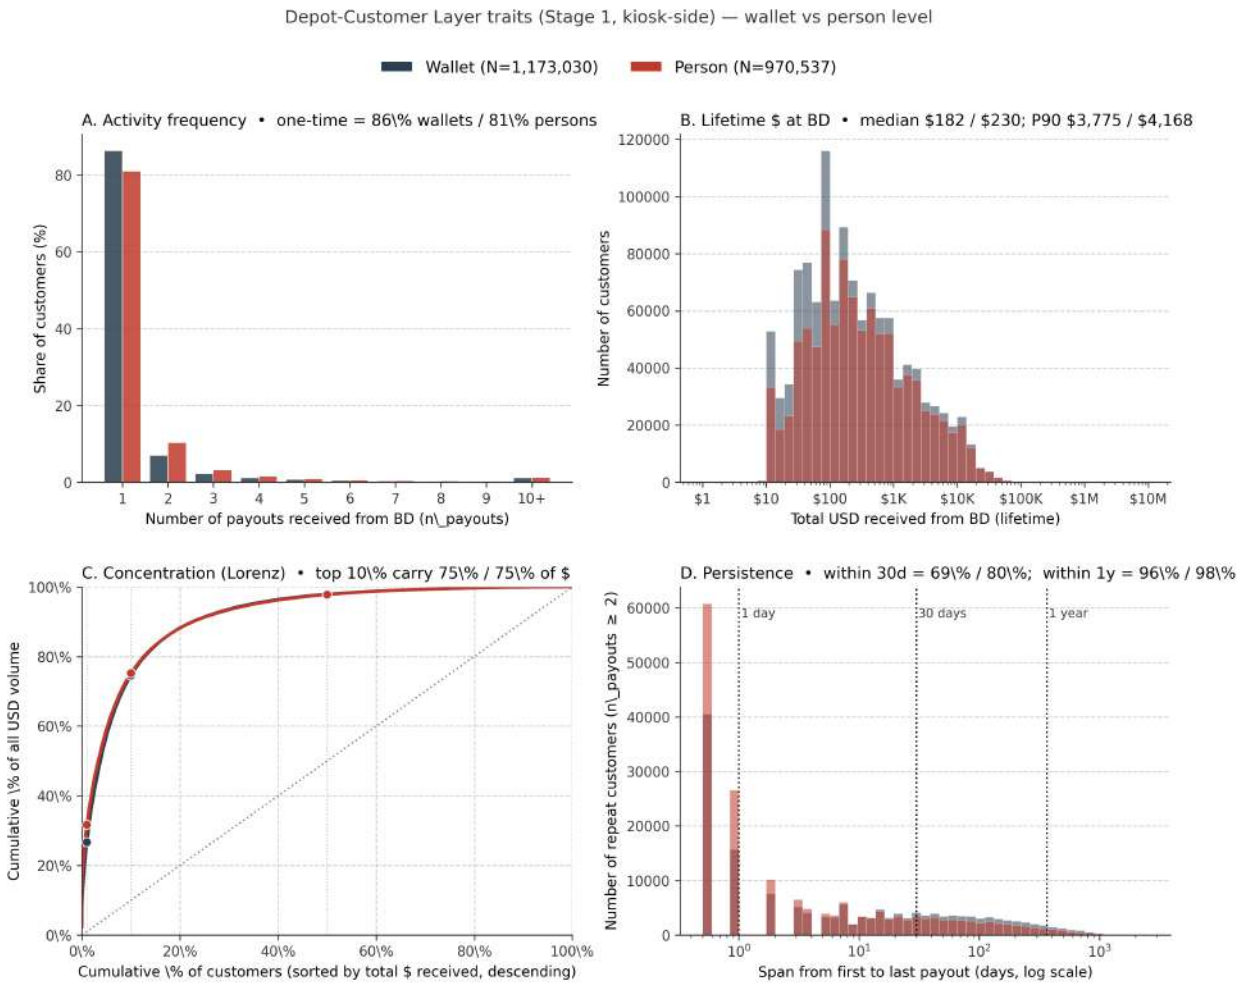


Figure 35: Customer behavior after the kiosk: customer-destination layer (Stage 2)

Notes. Four downstream-behavior traits computed for each retail customer at both the wallet and person level. Hop-1 is the first on-chain transaction in which the customer wallet appears as an input rather than as a receiving output. Panel A plots the frequency of hop-1 spending events per customer (cap at 10+). Panel B is a log-scale histogram of unique destination addresses reached per customer. Panel C is the headline: a log-scale histogram of the time in days from first receipt at Bitcoin Depot to first hop-1 outflow. Median time-to-spend is 0.9 days; 59% of wallets spend within 24 hours; 97% spend within 30 days. This is inconsistent with an investment holding-period interpretation of BTM use and consistent with flow-through cash-to-Bitcoin conversion in which the motivation for the kiosk transaction is the downstream transaction. Panel D is the Lorenz curve of destination concentration, ordering customers by number of unique destinations. Sample window 2018Q4–2026Q2, retail customers with at least one observed hop-1 spending transaction (1,192,398 wallets / 985,190 persons after mega-cluster exclusion).

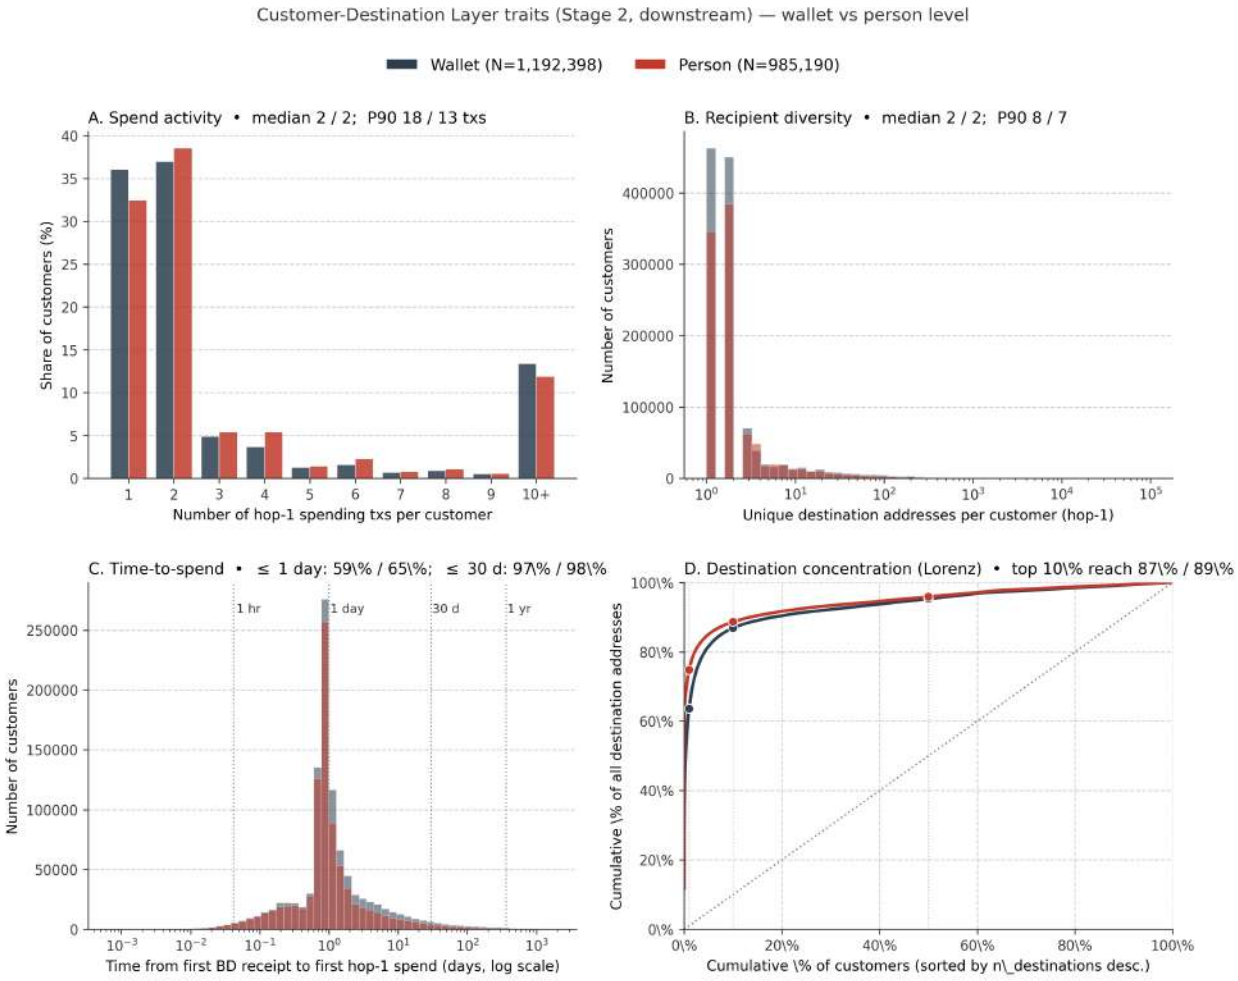


Figure 36: Bunching at the $250 KYC threshold

Notes. The figure displays the size distribution of customer payouts from Bitcoin Depot in the pre-DFAL window 2021Q1–2023Q4, restricted to the retail customer base (34 mega-clusters of 1,000 or more wallets excluded as services or aggregators). Panel A shows the full $0–$1,500 histogram with the $250 KYC threshold annotated by the vertical dashed line. Panel B zooms to $50–$600, applies the Kleven–Waseem (2013) excess-mass estimator with a fifth-order polynomial counterfactual fit on bins outside the [$220, $270) exclusion zone, and masks round-number bins ($150, $200, $300, $400, $500) from the fit to prevent them from biasing the counterfactual upward at the threshold. The Kleven–Waseem normalized excess mass over the exclusion zone is $\hat{b} = +0.58$ and the just-below to just-above bin ratio $N[\$225, \$250)/N[\$250, \$275)$ equals 2.52. Sample size in Panel B is N = 1,068,147 retail customer payouts.

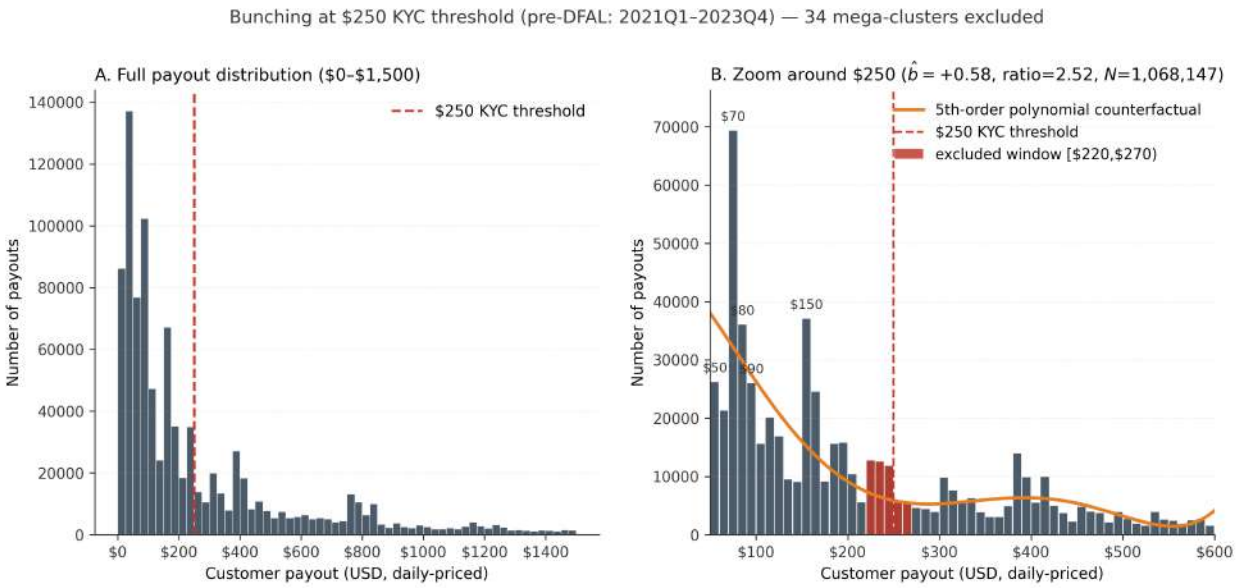


Figure 37: Structuring at the $250 KYC threshold — wallet level vs. person level

Notes. Each row characterizes structuring at a different aggregation unit: the top row uses the on-chain customer address × day (wallet level), the bottom row uses the v2 CIO cluster × day (person level). Panel A in each row plots the distribution of same-day total USD on multi-payout days, with the structuring sub-population (max single payout $< \$250$ *and* same-day sum $\geq \$250$) highlighted in red. Panel B in each row is the (sum, max) scatter with the bottom-right "structuring quadrant" shaded; points there satisfy the joint condition that identifies structuring. Sample is the pre-DFAL window 2021Q1–2023Q4, retail customers only, mega-clusters excluded. Wallet-level structuring covers 5,267 events (9.8% of multi-payout wallet-days). Person-level structuring covers 12,943 events (12.2% of multi-payout person-days). The 2.46× ratio implies that 59% of person-level structuring uses multiple wallets within a single day — a sophistication signal that operator-side KYC data cannot detect but on-chain CIO clustering surfaces.

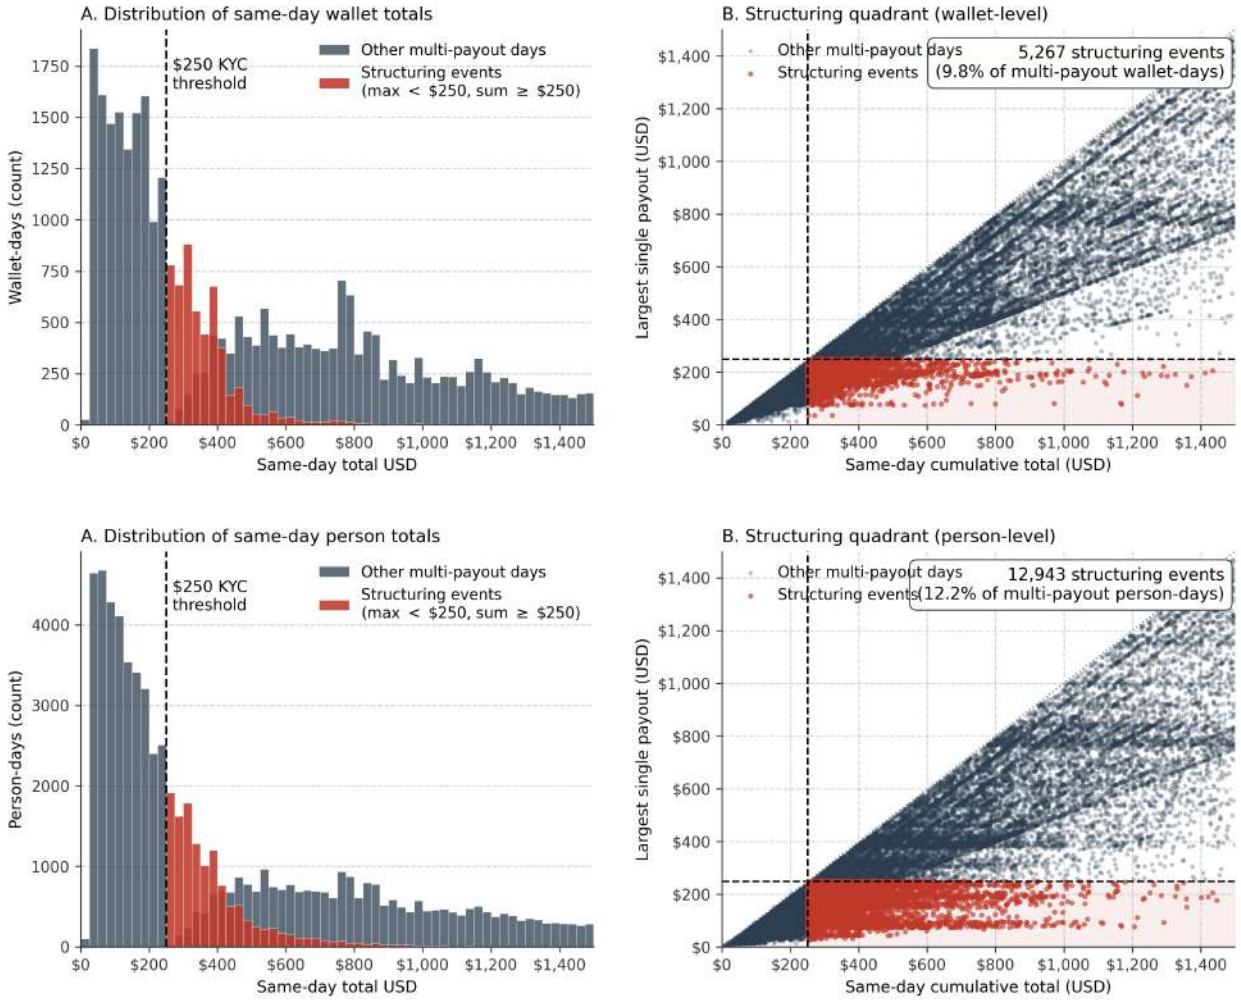


Table 1: Attribute bundles of payment instruments

Notes. The table classifies payment instruments along the four attributes introduced in Section 2.1. *Identity* records whether the user's real-world identity is recorded at the point of transaction. *Scale* records whether the system is local-only or usable globally. *Digital* records whether value is electronically transferable. *Traceable* records whether each transaction appears on an auditable ledger. Cryptocurrency is pseudonymous on both the identity and traceability dimensions because the public ledger records addresses rather than legal names.

	Identity	Scale	Digital	Traceable
Cash	anonymous	local	no	none
Bank / card	recorded	global	yes	full
Crypto	**pseudonymous**	**global**	**yes**	**pseudonymous**

Table 2: Summary statistics

Notes. Panel A reports kiosk-level statistics from the Coin ATM Radar data described in Section 3.1 (31,192 kiosks). Kiosks and operators per county are computed over the 1,880 counties with at least one kiosk. KYC limits are the operator-disclosed dollar thresholds. Panel B reports statistics for the NIBRS county-year crime panel, 2013–2023, restricted to counties reporting in at least six of the eleven years. Drug seizure weight is from NIBRS property records and is winsorized at the 99th percentile to limit the influence of recording errors. Overdose deaths are from CDC mortality files. Panel C reports 2015 baseline county covariates from the Census ACS, the FDIC Summary of Deposits, and the BLS QCEW.

	Mean	SD	P25	Median	P75	N
Panel A: Bitcoin ATM kiosks						
Kiosks per county	16.59	46.38	2.00	4.00	14.00	1,880
Operators per county	4.72	4.53	1.00	3.00	7.00	1,880
KYC tier-1 limit (USD)	1,372	2,965	300	960	2,000	30,715
KYC daily limit (USD)	21,781	12,351	15,000	19,000	25,000	31,028
Panel B: County-year panel, 2013–2023						
Crime incidents	477.0	599.1	171.0	309.0	504.0	10,488
Crime per 10,000 residents	68.2	40.5	39.4	61.8	89.4	10,463
Drug seizure weight (kg)	36.3	176.9	0.0	0.0	2.4	10,423
Overdose deaths	35.4	85.4	5.0	11.0	28.0	9,869
Population	136,829	306,978	24,992	46,178	115,508	10,463
Panel C: Baseline county covariates, 2015						
Median household income (USD)	46,830	12,251	38,826	45,111	52,250	3,141
Poverty rate (%)	16.7	6.6	12.0	16.0	20.3	3,142
Minority population share (%)	16.5	16.6	4.5	10.0	22.9	3,142
Unemployment rate (%)	7.8	3.6	5.5	7.5	9.7	3,142
Bank branches per 10,000 residents	4.64	3.00	2.81	3.78	5.38	3,112
Gas stations	35.8	75.7	8.0	16.0	34.0	3,117

Table 3: Effect of Bitcoin ATM Entry on Total Crime

Notes: Callaway and Sant'Anna (2021) ATT (simple aggregation) with doubly-robust inverse probability weighting and a not-yet-treated comparison group. Pre-trend p-values are from joint Wald tests of pre-treatment coefficients ($t = -5$ through $t = -2$). Standard errors clustered at the county level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)
	Count	Rate (per 10k)	Log(Count+1)	Log(Rate+1)
BTM entry	46.452	2.598	0.067*	0.055
	(34.834)	(1.952)	(0.035)	(0.038)
Pre-trend p	1.000	0.856	0.965	0.925
Outcome mean	477.013	68.231	5.791	4.047
Observations	10,488	10,488	10,488	10,488
Counties			1,045	
Treated counties			838	

Table 4: Effect of Bitcoin ATM Entry on Crime by Payment Relevance (3-Tier Taxonomy)

Notes: Each column reports the Callaway and Sant'Anna (2021) overall ATT ($\widehat{\theta}^S$, simple aggregation) for crimes grouped by payment relevance to BTM infrastructure. High = 14 offense codes directly exploiting BTM payment functions (fraud, extortion, identity theft, hacking, drugs, human trafficking). Moderate = 17 codes facilitated by BTM physical presence or proceeds laundering (robbery, assault, burglary, theft, bribery, gambling). None = 20-code placebo group with no plausible BTM mechanism. Panel A: $\bar{Y}_{it} = J^{-1} \sum_j \log(\text{count}_j + 1)$, weighting each offense type equally. Panel B: log(total group count + 1). Not-yet-treated comparison group, doubly-robust IPW. Pre-trend p-values from joint Wald test of pre-treatment coefficients. Full event study dynamics in Figure 9. Standard errors clustered at the county level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)
	High	Moderate	None
	(14 codes)	(17 codes)	(20 codes)
Panel A: Average-of-Logs			
ATT	0.095***	0.052*	0.028
	(0.031)	(0.031)	(0.019)
Pre-trend p	0.960	0.998	0.998
Outcome mean	1.548	1.857	0.863
Panel B: Log(Total Count+1)			
ATT	0.068*	0.070*	0.055
	(0.040)	(0.037)	(0.037)
Pre-trend p	0.958	0.954	0.989
Outcome mean	4.520	5.147	4.131
Observations		10,488	
Counties		1,045	
Treated counties		838	

Table 5: Effect of Bitcoin ATM Entry on Crime by Technological Trait (5-Group Taxonomy)

Notes: Each column reports the Callaway and Sant'Anna (2021) overall ATT ($\widehat{\theta}^S$, simple aggregation) for crimes grouped by the BTM technological trait exploited. Irreversibility = fraud and extortion exploiting non-reversible transactions (6 codes). Pseudonymity = offenses exploiting anonymous value transfer (7 codes). Cash-to-Crypto = property crimes whose proceeds are laundered via BTM conversion (10 codes). Cash Attraction = crimes drawn to BTM locations by cash concentration (8 codes). Not Relevant = 20-code placebo group identical to the 3-tier "None" category. Panel A: $\bar{Y}_{it} = J^{-1}\sum_j \log(\text{count}_j + 1)$. Panel B: log(total group count + 1). Not-yet-treated comparison group, doubly-robust IPW. Pre-trend p-values from joint Wald test of pre-treatment coefficients. Full event study dynamics in Figure 10. Standard errors clustered at the county level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1) Irreversibility (6 codes)	(2) Pseudonymity (7 codes)	(3) Cash-to-Crypto (10 codes)	(4) Cash Attraction (8 codes)	(5) Not Relevant (20 codes)
Panel A: Average-of-Logs					
ATT	0.110***	0.081***	0.073**	0.032	0.028
	(0.038)	(0.027)	(0.036)	(0.028)	(0.021)
Pre-trend p	0.980	0.893	0.997	0.994	0.998
Outcome mean	1.736	1.187	2.255	1.495	0.863
Panel B: Log(Total Count+1)					
ATT	0.092**	0.059	0.086**	0.040	0.055
	(0.042)	(0.038)	(0.037)	(0.039)	(0.038)
Pre-trend p	0.982	0.954	0.980	0.917	0.988
Outcome mean	3.684	3.726	4.783	4.140	4.131
Observations			10,488		
Counties			1,045		
Treated counties			838		

Table 6: KYC Threshold Heterogeneity and the BTM–Crime Effect

Notes: The dependent variable is the log crime rate. County anonymity is the average first-tier KYC threshold across the kiosks a county receives, the cash amount a customer can convert without presenting identification. Treated counties are split into terciles of this measure, with cuts at $787 and $1,126. A higher threshold means weaker identification requirements. Panel A reports the average of the post-entry event-study coefficients (event years 0 through 5) from an event study with county and year fixed effects, estimated separately within each tercile, matching Panel A of Figure 11. The pre-trend p-value is from a joint Wald test of the pre-entry coefficients. Panel B reports coefficients on bins of the county's kiosk count relative to county-years with no kiosks, estimated within each tercile plus all zero-kiosk county-years, matching Panel B of Figure 11. The sample is the NIBRS county-year panel, 2013–2023. Standard errors are clustered at the county level.

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

	(1) Strict (bottom tercile)	(2) Middle (middle tercile)	(3) Permissive (top tercile)
Panel A: Post-Entry ATT (event-study average)			
ATT	0.003	0.071	0.187***
	(0.057)	(0.058)	(0.053)
Pre-trend p	0.533	0.666	0.044
Observations	3,335	2,465	2,731
Treated counties	319	243	274
Panel B: Dose-Response (kiosk-count bins vs. zero)			
1 kiosk	0.004	-0.007	-0.029
	(0.017)	(0.024)	(0.034)
2-3 kiosks	0.009	0.031	0.041
	(0.023)	(0.032)	(0.027)
4-10 kiosks	0.040	0.120***	0.136***
	(0.028)	(0.033)	(0.031)
10+ kiosks	0.172***	0.177***	0.275***
	(0.056)	(0.051)	(0.048)

Table 7: County-Level Determinants of Bitcoin ATM Entry
(Cross-Sectional Logit, 2015 Baseline)

Notes: This table reports logit coefficients (columns 1–2) and the corresponding average marginal effects (columns 3–4) from cross-sectional logit regressions of Bitcoin ATM entry. The dependent variable equals one if the county ever installs at least one BTM during 2015–2023. All continuous covariates are standardized (mean 0, SD 1) and measured in 2014, the year before the sample opens. Columns (1) and (3) report the full specification with demographic, economic, financial infrastructure, and local market controls and state fixed effects; columns (2) and (4) add the log jail population rate from the Vera Institute Incarceration Trends dataset (∼96% county coverage). Standard errors clustered at the state level (50 clusters) are reported in parentheses. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Logit Coefficients		Avg. Marginal Effects	
	(1)	(2)	(3)	(4)
	Full + State FE	+ Crime	Full + State FE	+ Crime
Log pop. density (per sq mi)	1.106***	1.116**	0.0908**	0.0909**
	(0.424)	(0.449)	(0.0368)	(0.0387)
Pct. Black	-0.015	-0.148	-0.0012	-0.0120
	(0.223)	(0.193)	(0.0183)	(0.0157)
Pct. Hispanic	1.632***	1.550***	0.1340***	0.1262***
	(0.594)	(0.565)	(0.0478)	(0.0445)
Log median HH income	0.456**	0.434**	0.0375**	0.0353**
	(0.217)	(0.202)	(0.0173)	(0.0165)
Poverty rate	-0.328	-0.195	-0.0269	-0.0159
	(0.213)	(0.213)	(0.0176)	(0.0175)
Unemployment rate	0.155	0.181	0.0127	0.0147
	(0.196)	(0.207)	(0.0159)	(0.0166)
Gini index	0.471***	0.449***	0.0387***	0.0366***
	(0.131)	(0.137)	(0.0109)	(0.0116)
Real GDP growth (%)	-0.074	-0.064	-0.0060	-0.0052
	(0.096)	(0.093)	(0.0080)	(0.0077)
Bank branches per 10k pop.	-0.263	-0.171	-0.0216*	-0.0139
	(0.161)	(0.152)	(0.0130)	(0.0123)
Log avg. bank deposit	0.101	0.131*	0.0083	0.0107*
	(0.094)	(0.079)	(0.0078)	(0.0065)
House price growth (%)	-0.010	0.022	-0.0008	0.0018
	(0.084)	(0.086)	(0.0069)	(0.0070)
Log 2BR fair market rent	-0.578*	-0.461*	-0.0474**	-0.0375**
	(0.304)	(0.243)	(0.0237)	(0.0188)
Log gas stations	2.755***	2.731***	0.2262***	0.2223***
	(0.484)	(0.539)	(0.0324)	(0.0362)
Log jail pop. rate (per 100k)		-0.059		-0.0048
		(0.254)		(0.0205)
Observations	4,068	3,893	4,068	3,893
Ever BTM	2,907	2,807	2,907	2,807
State FE	Yes	Yes	Yes	Yes
SE cluster	State	State	State	State

Table 8: County-Level Determinants of the Timing of First Bitcoin ATM Entry (Complementary Log-Log Discrete-Time Hazard, 2015–2023)

Notes: This table reports coefficients from complementary log-log discrete-time hazard regressions of first Bitcoin ATM entry, 2015–2023. The dependent variable equals one in the year of first BTM installation; the county exits the risk set after entry. All continuous covariates are standardized (mean 0, SD 1) and lagged one year. Covariate blocks match the cross-sectional logit (Table 7) exactly. Columns (1) and (4) include demographics only; (2) and (5) add economic conditions (including poverty rate), financial infrastructure, and local market controls (including house price growth); (3) and (6) add the credit insecurity index (restricted to 2019–2023); (7) and (8) add the jail population rate from the Vera Institute Incarceration Trends dataset (∼96% county coverage). Odd columns exclude state fixed effects; even columns include them. t-statistics based on standard errors clustered at the state level (50 clusters) are reported in brackets. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1) Demo	(2) +Econ/Fin/Mkt	(3) +CI	(4) Demo	(5) +Econ/Fin/Mkt	(6) +CI	(7) +Jail	(8) +Jail
Log pop. density (per sq mi)	0.725***	0.272	0.664***	1.446***	0.843***	1.371***	0.287	0.751***
	[2.717]	[1.550]	[2.661]	[5.862]	[7.706]	[6.020]	[1.462]	[4.559]
Pct. Black	-0.597***	-0.224**	-0.687***	-0.612***	-0.257*	-0.911***	-0.341**	-0.371**
	[-2.787]	[-2.133]	[-2.842]	[-3.832]	[-1.839]	[-3.769]	[-2.269]	[-2.044]
Pct. Hispanic	-3.432	0.998	-1.345	-0.492	1.600***	-2.120	1.090	1.591**
	[-1.007]	[1.466]	[-0.838]	[-0.521]	[2.860]	[-1.227]	[1.620]	[2.446]
Log median HH income		0.277	0.115		-0.027	-0.146	0.252	-0.160
		[1.099]	[0.456]		[-0.144]	[-0.778]	[0.989]	[-0.763]
Poverty rate		-0.262	-0.258		-0.227	-0.220	-0.473*	-0.340
		[-1.190]	[-1.072]		[-1.388]	[-1.035]	[-1.849]	[-1.472]
Unemployment rate		0.212	0.198		0.253**	0.286*	0.141	0.129
		[1.228]	[1.386]		[2.255]	[1.717]	[0.857]	[1.454]
Gini index		0.169**	0.205***		0.110	0.170***	0.215**	0.172*
		[2.180]	[3.451]		[1.567]	[2.639]	[2.568]	[1.860]
Real GDP growth (%)		0.116*			-0.007		0.071	-0.045
		[1.830]			[-0.081]		[0.769]	[-0.352]
Bank branches per 10k pop.		-0.203**	-0.352***		-0.277***	-0.372**	-0.288*	-0.491**
		[-2.025]	[-3.646]		[-2.985]	[-2.204]	[-1.794]	[-2.236]
Log avg. bank deposit		-0.087	0.086		-0.042	0.026	-0.084	-0.049
		[-0.918]	[1.340]		[-0.478]	[0.334]	[-0.914]	[-0.620]
House price growth (%)		0.127	0.090		-0.110	-0.147	0.024	-0.274
		[0.993]	[0.932]		[-1.138]	[-1.464]	[0.142]	[-1.488]
Log 2BR fair market rent		-0.483**	-0.283		-0.270**	-0.225***	-0.405**	-0.267**
		[-2.293]	[-1.241]		[-2.193]	[-3.153]	[-2.214]	[-2.034]
Log gas stations		1.210***			1.477***		1.201***	1.501***
		[5.152]			[9.695]		[5.761]	[7.545]
Credit insecurity score			-0.007			-0.033		
			[-0.083]			[-0.242]		
Log jail pop. rate (per 100k)							0.296*	-0.057
							[1.754]	[-0.411]
Observations	32,994	28,133	12,583	32,994	28,133	12,583	25,180	25,180
Events	2,907	2,821	2,292	2,907	2,821	2,292	2,342	2,342
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	No	No	No	Yes	Yes	Yes	No	Yes
Jail rate (Vera)	No	No	No	No	No	No	Yes	Yes

Table 9: First Stage: Gas Station Bartik Instrument and BTM Count

Notes: This table reports first-stage regressions of county-level cumulative BTM count on the Bartik instrument $Z_{it} = s_i \times g_{-i,t}$, where s_i is the county's 2005 gas station establishment share (NAICS 447110, from QCEW) and $g_{-i,t}$ is the leave-one-out national cumulative BTM count. The leave-one-out national total is computed from the universe of BTM installations across all counties, not restricted to the crime analysis sample. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1)	(2)	(3)
	Baseline	+ Log Pop×Year	+ State×Year
Z_{it}^{Bartik}	0.718***	0.743***	0.717***
	(0.072)	(0.096)	(0.074)
Kleibergen-Paap F	98.5	59.5	93.2
Partial R^2	0.009	0.006	0.009
N	10,488	10,488	10,488
County FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Log Pop×Year	No	Yes	No
State×Year	No	No	Yes

Table 10: Reduced Form: Effect of Bartik Instrument on Crime Outcomes

Notes: This table reports reduced-form regressions of crime outcomes directly on the Bartik instrument $Z_{it} = s_i \times g_{-i,t}$, without instrumenting for BTM count. All specifications include county and year fixed effects. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Crime Rate per 10,000	(2) Crime Count	(3) Log(Crime+1)
Z_{it}^{Bartik}	0.051^{***}	4.527^{***}	0.004^{***}
	(0.014)	(0.870)	(0.001)
N	10,462	10,488	10,488
County FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Table 11: OLS and 2SLS Estimates of the Effect of BTM Count on Crime

Notes: This table reports OLS and 2SLS estimates of the effect of cumulative BTM count on county-level crime outcomes. The endogenous variable is the cumulative number of Bitcoin ATMs in county i in year t. The excluded instrument is the Bartik shift-share $Z_{it} = s_i \times g_{-i,t}$, where s_i is the county's 2005 gas station establishment share (NAICS 447110) and $g_{-i,t}$ is the leave-one-out national BTM count. Panel A uses crime rate per 10,000 population; Panel B uses the crime count level; Panel C uses log(crime + 1). All specifications include county and year fixed effects. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) Baseline		(2) + Log Pop×Year	
	OLS	2SLS	OLS	2SLS
Panel A: Crime rate per 10,000				
BTM count	0.053***	0.071***	0.034**	0.050**
	(0.020)	(0.025)	(0.016)	(0.024)
Panel B: Crime count				
BTM count	5.238***	6.309***	1.599***	2.093***
	(1.079)	(1.421)	(0.471)	(0.762)
Panel C: Log(crime + 1)				
BTM count	0.005***	0.005***	0.001	0.001*
	(0.001)	(0.001)	(0.000)	(0.001)
KP F-stat	98.5	98.5	59.5	59.5
N	10,488	10,488	10,488	10,488
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Table 12: Robustness of 2SLS Estimates: Crime Rate per 10,000

Notes: This table reports robustness checks for the baseline 2SLS specification (Table 11, Panel A, column 1). Column (1) clusters standard errors at the state level instead of the county level. Column (2) drops the largest 5% of counties by population. Column (3) restricts to counties above the median population. Column (4) restricts the sample to 2015–2023. The instrument is the Bartik shift-share $Z_{it} = s_i \times g_{-i,t}$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

	(1) State Clustering	(2) Drop Top 5% Pop.	(3) Urban Only	(4) 2015–2023 Only
BTM count	0.071**	0.279***	0.032**	0.072***
	(0.029)	(0.062)	(0.015)	(0.025)
N	10,488	—	—	—
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Table 13: BTM Entry and Crime: Census Tract–Quarter Specification

Notes: Regression coefficients showing the effect of Bitcoin ATM entry on total crime counts at the Census tract–quarter level. Models (1)–(2): OLS with raw and log crime outcomes; Models (3)–(4): OLS with demeaned outcomes; Models (5)–(6): Poisson specifications. All specifications include Census tract and quarter fixed effects. Standard errors clustered at the Census tract level in parentheses. Sample: unbalanced panel of 857 Chicago Census tracts observed quarterly from 2011 through 2025 (N = 49,942). The panel is unbalanced because tract-quarters with no recorded incidents are not filled with zeros. Columns (2), (4), and (6) use the one-quarter-ahead outcome, which shortens each tract's series by one quarter. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Dependent Variables:	crime_total	f_crime_total	ln_crime	f_ln_crime	crime_total	f_crime_total
Model:	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	Poisson	Poisson
Variables						
did	1.095	0.7143	0.0421***	0.0407**	0.0548***	0.0522***
	(1.743)	(1.786)	(0.0159)	(0.0158)	(0.0190)	(0.0186)
Fixed-effects						
tract_fips11	Yes	Yes	Yes	Yes	Yes	Yes
quarter	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	49,942	49,083	49,942	49,083	49,942	49,083
Squared Correlation	0.88730	0.88756	0.93625	0.93627	0.91738	0.91792
Pseudo R^2	0.19174	0.19200	0.98567	0.98752	0.83874	0.83886
BIC	469,486.5	461,267.0	11,908.0	11,559.0	467,506.5	458,216.6

Clustered (tract_fips11) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 14: BTM Entry and Crime: Police Beat–Quarter Specification

Notes: Regression coefficients showing the effect of Bitcoin ATM entry on total crime counts at the police beat–quarter level. Models (1)–(2): OLS with raw and log crime outcomes; Models (3)–(4): OLS with demeaned outcomes; Models (5)–(6): Poisson specifications. All specifications include police beat and quarter fixed effects. Standard errors clustered at the police beat level in parentheses. Sample: balanced panel of 305 Chicago police beat codes observed quarterly from 2011 through 2025 (N = 18,300). Beat-quarters with no recorded incidents are filled with zeros. Columns (2), (4), and (6) use the one-quarter-ahead outcome and lose each beat's last quarter. The Poisson estimator drops one beat with no variation in crime. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Dependent Variables:	total_crime	f_total_crime	ln_crime	f_ln_crime	total_crime	f_total_crime
Model:	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	Poisson	Poisson
Variables						
did	8.502*	8.491*	0.1626***	0.1393***	0.0828***	0.0819***
	(4.503)	(4.392)	(0.0302)	(0.0264)	(0.0198)	(0.0192)
Fixed-effects						
beat_num	Yes	Yes	Yes	Yes	Yes	Yes
quarter	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	18,300	17,995	18,300	17,995	18,240	17,936
Squared Correlation	0.83850	0.84409	0.88802	0.90392	0.85172	0.85805
Pseudo R^2	0.14580	0.14865	0.56876	0.60757	0.76807	0.77906
BIC	199,059.1	195,093.1	33,961.4	30,793.8	402,129.0	376,996.6

Clustered (beat_num) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 15: Population Sorting: BTM Entry and Demographic Composition

Notes: Regression coefficients testing for population sorting following Bitcoin ATM entry at the Census tract–year level. Dependent variables: (1) total population, (2) log population, (3) share White, (4) share Black, (5) share Hispanic/Latino, (6) share Male. All specifications include Census tract and year fixed effects. Standard errors clustered at the Census tract level in parentheses. Sample: 858 Chicago Census tracts, 2015–2023 ($N = 7{,}662$; American Community Survey 5-year estimates). $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Dependent Variables:	total_population	ln_pop	race_white	race_black	race_hislatin	male
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
did	30.44	0.0067	-1.025*	0.4119	0.0710	0.2478
	(32.87)	(0.0073)	(0.6063)	(0.3752)	(0.4647)	(0.2838)
Fixed-effects						
tract_fips11	Yes	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	7,662	7,662	7,635	7,635	7,635	7,635
R^2	0.97771	0.98774	0.96512	0.99118	0.97741	0.66901
Within R^2	0.00056	0.00029	0.00125	0.00053	1.15×10^{-5}	0.00030

Clustered (tract_fips11) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 16: Summary Statistics: Drug Seizure and Overdose Outcomes

Notes: Sample restricted to NIBRS-reporting counties, 2013–2023. "Treated" counties received at least one Bitcoin ATM during the sample period; "control" counties were never treated. Drug seizure weight in grams per 10,000 population (solid measurement units only: g/kg/oz/lb). Panel A winsorizes at the 99th percentile to cap extreme outliers from large individual seizures. IHS denotes the inverse hyperbolic sine transformation, $\text{IHS}(x) = \ln(x + \sqrt{x^2 + 1})$. Hard drugs: cocaine (incl. crack), heroin, morphine, opium, other narcotics, meth/amphetamines. Overdose death rates are NCHS model-based county estimates (2013–2021), supplemented by CDC provisional data for 2022–2023. Normalized differences reported following Imbens and Wooldridge (2009). † denotes $|\Delta_{\text{norm}}| > 0.25$.

	Treated		Control		Norm.	
	Mean	SD	Mean	SD	Diff.	*N*
Panel A: Drug Seizure Weight (g per 10k pop, winsorized p99)						
All drugs (solid)	2,786.6	13,135.0	2,649.1	13,895.6	0.007	10,407
Hard drugs	287.1	1,331.8	337.4	1,617.6	−0.024	10,407
Cocaine (incl. crack)	64.0	322.9	51.4	317.2	0.028	10,407
Heroin	20.2	116.4	21.1	127.6	−0.005	10,407
Meth/amphetamines	110.1	519.2	127.9	594.8	−0.022	10,407
Cannabis	2,048.7	9,808.4	1,818.1	9,936.0	0.017	10,407
Panel B: IHS Drug Seizure Weight (g per 10k pop)						
All drugs (solid)	3.16	3.84	2.77	3.72	0.073	10,407
Hard drugs	2.16	2.88	1.99	2.90	0.042	10,407
Cocaine (incl. crack)	1.25	2.14	0.81	1.83	0.157	10,407
Heroin	0.79	1.65	0.60	1.61	0.080	10,407
Meth/amphetamines	1.67	2.49	1.66	2.62	0.003	10,407
Cannabis	2.90	3.67	2.41	3.50	0.097	10,407
Panel C: Extensive Margin						
Any drug seizure	0.456	0.498	0.405	0.491	0.073	10,423
Any hard drug seizure	0.450	0.498	0.398	0.490	0.074	10,423
Any cocaine seizure	0.389	0.487	0.271	0.445	0.178	10,423
Any heroin seizure	0.341	0.474	0.220	0.414	0.193	10,423
Any meth seizure	0.422	0.494	0.373	0.484	0.070	10,423
Any cannabis seizure	0.447	0.497	0.382	0.486	0.093	10,423
Panel D: Drug Incidents						
Count	245.0	777.5	80.4	270.2	0.200	10,423
IHS	2.78	3.16	2.08	2.63	0.172	10,423
Panel E: Overdose Mortality						
Death rate per 100k	23.3	14.6	21.8	15.5	0.071	9,869
Deaths (approx. count)	41.6	92.5	5.2	5.4	0.393†	9,869
Any overdose death	0.995	0.072	0.970	0.170	0.132	9,869
Panel F: Demographics						
Population	162,630	334,391	22,005	12,580	0.420†	10,407
County-years	8,554		1,869			10,423
Counties	838		207			1,045

Table 17: Effect of BTM Entry on Drug Seizure Weight: Callaway-Sant'Anna Estimates

Notes: Callaway-Sant'Anna (2021) staggered difference-in-differences estimates. Panel A reports extensive margin effects (probability of any seizure). Panel B reports IHS-transformed grams per 10,000 population. Comparison group: not-yet-treated counties. Standard errors clustered at the county level in parentheses. Sample restricted to NIBRS-reporting counties, 2013–2023. "Hard" drugs: cocaine (incl. crack), heroin, morphine, opium, other narcotics, meth/amphetamines. "Solid" weight uses only records in weight units (g/kg/oz/lb); excludes liquids, dosage units, and plant counts. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
	All Drugs	Hard	Cocaine	Heroin	Meth	Cannabis
Panel A: Extensive Margin (Any Seizure)						
ATT	-0.041^{**}	-0.030	0.134^{***}	0.159^{***}	-0.010	0.008
	(0.020)	(0.020)	(0.023)	(0.024)	(0.021)	(0.021)
Pre-trend p	0.994	0.971	0.985	0.913	0.967	0.990
Outcome mean	0.447	0.441	0.368	0.319	0.413	0.435
N	10,423	10,423	10,423	10,423	10,423	10,423
Panel B: IHS-Transformed Weight (g per 10k pop)						
ATT	0.562^{***}	0.544^{***}	1.060^{***}	0.547^{***}	0.388^{***}	0.926^{***}
	(0.149)	(0.118)	(0.095)	(0.080)	(0.109)	(0.160)
Pre-trend p	1.000	1.000	1.000	0.966	0.995	0.999
Outcome mean	3.217	2.127	1.171	0.753	1.671	2.816
N	10,407	10,407	10,407	10,407	10,407	10,407
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	County	County	County	County	County	County

Table 18: Effect of BTM Entry on Drug Seizure Weight: Poisson QMLE Estimates

Notes: Poisson quasi-maximum likelihood estimates following Chen & Roth (2024, *QJE*). Dependent variable: drug seizure weight in grams per 10,000 population. Poisson QMLE handles zero-inflated outcomes natively; $\exp(\hat{\beta})-1$ gives the proportional treatment effect on the mean. Outcomes winsorized at the 99th percentile. County and year fixed effects; standard errors clustered at the county level. Sample restricted to NIBRS-reporting counties, 2013–2023. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
	All Drug	Hard	Cocaine	Heroin	Meth	Cannabis
$\exp(\hat{\beta}) - 1$	0.573*	0.607*	0.482*	0.918**	0.430	0.871***
	(0.378)	(0.404)	(0.352)	(0.537)	(0.426)	(0.449)
Pre-trend p	0.798	0.548	0.463	0.369	0.496	0.372
Outcome mean	4,166.4	296.1	61.7	20.3	113.3	2,007.3
N	10,387	10,375	9,692	9,130	10,279	10,210
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	County	County	County	County	County	County
Winsorized	$p99$	$p99$	$p99$	$p99$	$p99$	$p99$

Table 19: Narcotics Crime: Demand vs. Supply Decomposition

Notes: Regression coefficients decomposing narcotics crime into demand-side (possession) and supply-side (manufacture, delivery, trafficking) offenses following Bitcoin ATM entry. Columns (1)–(2): OLS with raw counts; Columns (3)–(4): OLS with log outcomes; Column (5): OLS with supply share; Column (6): Poisson for supply offenses. All specifications include Census tract and quarter fixed effects. Standard errors clustered at the Census tract level in parentheses. Sample: unbalanced panel of 843 Chicago Census tracts observed quarterly from 2001 through 2025 ($N = 59{,}697$). Tract-quarters appear only when at least one narcotics incident is recorded. The Poisson estimator in column (6) drops tracts with no variation in supply offenses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Dependent Variables:	demand	supply	ln_demand	ln_supply	supply_share	supply
Model:	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	OLS	Poisson
Variables						
did	-1.553	0.6762*	-0.0067	0.1008***	0.0183	0.4286***
	(1.688)	(0.3537)	(0.0523)	(0.0387)	(0.0133)	(0.1064)
Fixed-effects						
tract_fips11	Yes	Yes	Yes	Yes	Yes	Yes
quarter	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	59,697	59,697	59,697	59,697	59,697	59,601
Squared Correlation	0.55517	0.58891	0.73034	0.58829	0.18511	0.70354
Pseudo R^2	0.09297	0.13933	0.43986	0.35714	0.99024	0.57582
BIC	482,087.7	338,064.2	109,907.2	105,633.1	10,391.7	204,582.6

Clustered (tract_fips11) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 20: Cannabis Banking Exclusion and Drug Seizures: Triple Difference-in-Differences

Notes: Each column reports a separate regression of drug seizure weight (inverse hyperbolic sine) on BTM entry, an interaction with recreational cannabis legalization status, county fixed effects, and state-by-year fixed effects. "btm" is a binary indicator equal to one after the first Bitcoin ATM is installed in the county. "btm × rec" interacts the BTM indicator with an indicator for states that have legalized recreational cannabis sales by year t. Standard errors are clustered at the county level. The sample covers NIBRS-reporting counties, 2013–2023.

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent Variables:	ihs_drug_total	ihs_cocaine	ihs_heroin	ihs_meth	ihs_cannabis	ihs_other_opioid
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
btm	0.4807***	0.6286***	0.4550***	0.4318***	0.5766***	0.5047***
	(0.1430)	(0.1157)	(0.1046)	(0.1174)	(0.1429)	(0.0983)
btm_x_rec	0.7951**	0.4600*	0.3928*	0.5607*	1.043***	0.3220
	(0.3383)	(0.2479)	(0.2295)	(0.2967)	(0.3522)	(0.2337)
Fixed-effects						
countyfips5	Yes	Yes	Yes	Yes	Yes	Yes
state_fips-year	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	10,364	10,364	10,364	10,364	10,364	10,364
R^2	0.82651	0.68431	0.64207	0.77090	0.79882	0.59978
Within R^2	0.00759	0.01187	0.00878	0.00788	0.01087	0.00929

Clustered (countyfips5) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 21: Cannabis Banking Exclusion and Crime: Triple Difference-in-Differences

Notes: Column (1) reports OLS with log crime rate as the dependent variable; column (2) reports a Poisson specification with total crime counts. "btm" is a binary indicator equal to one after the first Bitcoin ATM is installed in the county. "btm × rec" interacts the BTM indicator with an indicator for states that have legalized recreational cannabis sales by year t. Standard errors are clustered at the county level. The sample covers NIBRS-reporting counties, 2013–2023.

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent Variables:	ln_crime	crime_total
Model:	(1)	(2)
	OLS	Poisson
Variables		
btm	0.0531***	0.0185
	(0.0147)	(0.0178)
btm_x_rec	0.0064	0.0092
	(0.0318)	(0.0255)
Fixed-effects		
countyfips5	Yes	Yes
state_fips-year	Yes	Yes
Fit statistics		
Observations	10,449	10,449
Squared Correlation	0.91631	0.96475
Pseudo R^2	1.0461	0.94076
BIC	12,091.4	281,051.5

Clustered (countyfips5) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 22: Effect of BTM Entry on Overdose Mortality

Notes: Callaway-Sant'Anna (2021) staggered DiD estimates (columns 1–4) and Poisson QMLE following Chen and Roth (2024, QJE) (column 5) of the effect of first Bitcoin ATM entry on overdose mortality. Overdose death rates are NCHS model-based county estimates (2013–2021), supplemented by CDC provisional data for 2022–2023. IHS denotes the inverse hyperbolic sine transformation. "Deaths (IHS)" uses back-calculated death counts from rate × population. Column 5 reports $\exp(\hat{\beta}) - 1$ (proportional change) with delta-method standard errors; the dependent variable is the death rate in levels. Comparison group: not-yet-treated counties. Standard errors clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Callaway-Sant'Anna DiD				Poisson
	(1)	(2)	(3)	(4)	(5)
	Rate/100k (IHS)	Rate/100k (Level)	Deaths (IHS)	Any OD Death	Rate/100k
ATT	0.371***	0.087	0.128*	0.122***	0.014
	(0.078)	(1.258)	(0.076)	(0.017)	(0.011)
Pre-trend p	0.669	0.997	0.992	0.000	0.026
Outcome mean	3.640	23.036	3.188	0.991	22.860
N	9,869	9,869	9,869	9,869	9,869
Counties	1,045	1,045	1,045	1,045	1,045
County FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Table 23: Effect of BTM Entry on Victim Composition (Share Outcomes)

	ATT	SE	Pre-trend p	$\bar{Y}$
Panel A: Demographics				
% Minority Victims	7.843***	(1.820)	0.990	22.43
% Hispanic Victims	7.229***	(1.671)	0.964	11.35
% Black Victims	2.509	(1.663)	0.979	13.98
% Female Victims	0.807*	(0.456)	0.940	51.34
Panel B: Age				
% Young Victims (18–25)	2.703***	(0.460)	0.909	14.57
% Elderly Victims (65+)	−2.134***	(0.378)	0.970	11.97
Panel C: Relationship to Offender				
% Stranger Crime	3.677***	(0.752)	0.985	9.79
% Intimate Partner Violence	−0.011	(0.904)	0.931	30.16
% Acquaintance Crime	−1.064	(0.877)	0.911	34.06
% Domestic Violence	−3.319***	(1.157)	0.952	49.09
Counties	440			
County-years	4,836–4,840			

Notes: Callaway-Sant'Anna (2021) staggered DiD estimates. Comparison group: not-yet-treated counties. Standard errors (in parentheses) clustered at the county level. Outcomes are percentage-point shares (composition), not per-capita rates. $\bar{Y}$ = sample mean of the outcome variable. Pre-trend p = joint Wald test of pre-treatment coefficients. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Data: NIBRS victim records, 2013–2023.

Table 24: Address Taxonomy: Bitcoin Depot On-Chain Ecosystem

Notes: Classification of the 2,000,074 unique Bitcoin addresses appearing in the Bitcoin Depot transaction graph, 2018Q4 through 2026Q2. Seed addresses are identified through proprietary blockchain intelligence. Absolute counts for the seed cluster and the heuristic expansion are not displayed (ND) because, combined with the union, they would reveal the size of the proprietary seed list. Within the heuristic expansion, co-input expansion applies the common-input-ownership heuristic, the change-address heuristic identifies operator hot wallets from high-confidence change patterns, and the spending-pattern confirmation validates candidate operator addresses by their downstream spending behavior. Rows A2a through A2c report each method's share of the de-duplicated expansion, and the shares exceed 100 percent because a single address may satisfy multiple heuristics. Liquidity providers are the addresses attributed to Cumberland DRW and the Gemini exchange named in Bitcoin Depot's Form 10-K. Service-like counterparties are payout recipients flagged by the high-frequency filter (more than 100 distinct downstream destinations) and excluded from the customer sample. Clean customer wallets are recipients within each reconstruction batch's validated window, and the clean sample contains 2,397,315 payouts totaling $2.42 billion. Panels A, B, and C are mutually disjoint. Source: Bitcoin blockchain; classification validated against Bitcoin Depot's audited SEC disclosures.

	Category	Count	Share (%)
A.	**Bitcoin Depot-Controlled (union)**	**105,530**	**5.3**
A1.	Seed cluster (proprietary)	ND	
A2.	Heuristic expansion (union of three methods)	ND	
A2a.	Co-input expansion		3.1
A2b.	Change-address heuristic		54.5
A2c.	Spending-pattern confirmation		45.9
B.	**External Entities**	**39,372**	**2.0**
B1.	Liquidity providers (Cumberland DRW, Gemini)	198	
B2.	Service-like counterparties (high-frequency filter)	39,174	
C.	**Customer Wallets**	**1,855,172**	**92.8**
C1.	Clean customer sample (validated windows)	1,405,435	70.3
C2.	Outside validated windows	449,737	22.5
	Total	**2,000,074**	**100.0**

Table 25: Distribution of Bitcoin Depot Customer Payouts

(a) Summary Statistics

Statistic	Value
Payouts	2,397,315
Median	$179
25th percentile	$60
75th percentile	$605
90th percentile	$2,070

(b) Size Distribution

Size Bin	Count	Share (%)
$0–50	521,759	22.5
$50–100	392,299	16.9
$100–250	479,005	20.7
$250–500	293,940	12.7
$500–1K	246,443	10.6
$1K–5K	288,943	12.5
$5K+	97,105	4.2

Notes: Distribution of 2,397,315 customer payouts identified in the Bitcoin Depot on-chain reconstruction, 2018Q4–2026Q2, after removing liquidity-provider transfers, change outputs, and internal consolidation transactions. Dollar values are converted at the daily BTC/USD price, and percentile statistics use the 2,319,499 payouts with non-missing prices. The reconstructed quarterly median tracks the median kiosk transaction size disclosed in Bitcoin Depot's SEC filings with correlation 0.97, and the reconstructed quarterly volume tracks the audited cost-of-crypto-purchased line with correlation 0.96. Source: Bitcoin blockchain; Bitcoin Depot SEC filings.

A Construction of the BTM-Exposure Weighted Crime Index

This appendix details the BTM-exposure weighted crime index referenced in Section 4.5. The taxonomy analysis treats all offenses within a group equally, but the link between Bitcoin ATMs and crime varies across specific offense types even within a group. To exploit this heterogeneity, I construct continuous offense-level weights following the text-as-data methodology of Baker et al. [2016].

I compile a corpus of news articles from RavenPack News Analytics that mention Bitcoin ATMs, cryptocurrency ATMs, or related terms. For each of the 51 NIBRS offense types, I define a keyword dictionary of crime-specific terms covering both formal UCR descriptions and colloquial variants, and I count the number of BTM-related articles that co-mention each offense type. The weight assigned to offense j is the normalized co-mention frequency

$$w_j = \frac{n_j + \lambda}{\sum_{k=1}^{51}(n_k + \lambda)}, \tag{16}$$

where n_j is the number of BTM-related articles mentioning offense type j and λ is a Laplace smoothing parameter that ensures no offense receives zero weight. The weights sum to one across the 51 offense codes. Figure A.2 displays the distribution. The highest weights attach to false pretenses and swindle, wire fraud, hacking, extortion, and identity theft, the offense types most frequently discussed in media coverage of BTM-facilitated crime. Results are robust to raw count weighting and TF-IDF weighting, with high Spearman rank correlations of the weights across methods.

Figure A.3 plots the Callaway–Sant'Anna event-study dynamics of the log weighted index. The ATT is comparable in magnitude to the unweighted total but estimated with notably tighter precision. The improvement in precision is itself informative, since concentrating statistical power on the crime types most relevant to BTM activity reduces noise

from unrelated offenses, consistent with the evidence in Section 4.5 that BTMs affect crime selectively rather than uniformly. The event-study coefficients rise monotonically over the post-treatment window, pre-treatment estimates are flat and not jointly distinguishable from zero, and the largest pre-treatment estimate is roughly an order of magnitude smaller than the steady-state post-treatment effect.

B Appendix Figures

Figure A.1: Geographic Distribution of Bitcoin ATMs: County-Level Density

Notes: County-level Bitcoin ATM density (machines per 100,000 residents) as of December 2023. Counties shaded white have no active machines. The map reveals strong spatial concentration in Sun Belt metros (Texas, Florida, California, Georgia), the Great Plains corridor, and major Rust Belt cities. Rural counties in the Northeast and Mountain West are largely unserved. Source: Coin ATM Radar. Population denominator from the 2020 Census.

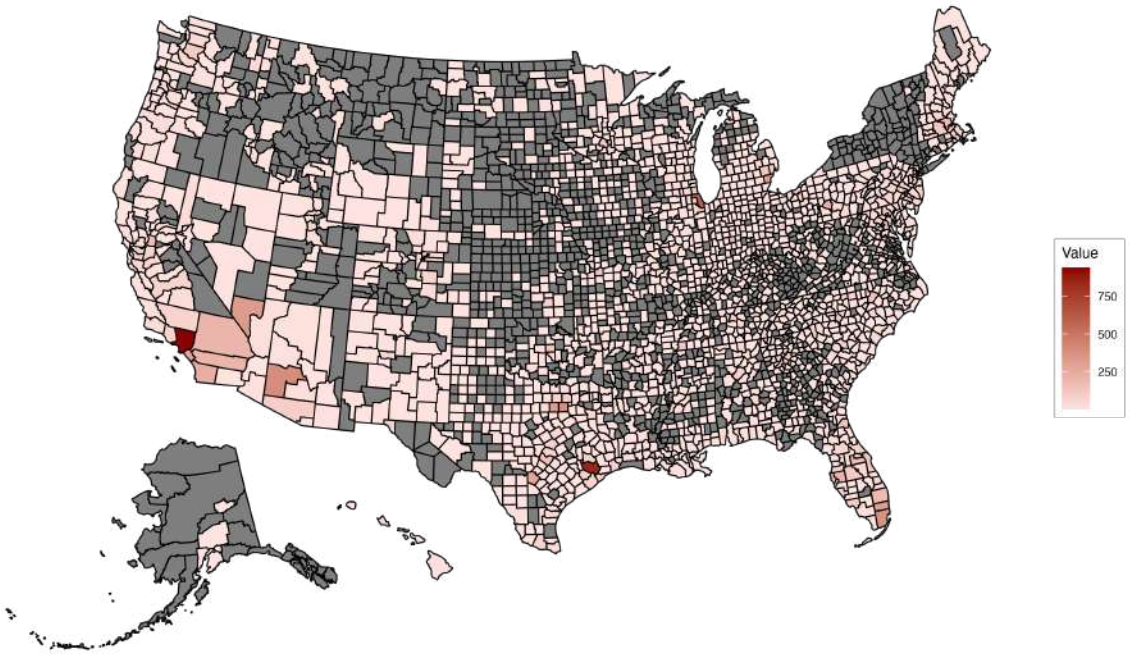


Figure A.2: BTM-Crime Exposure Weights by Offense Type

Notes: Radar chart displaying exposure weights for the top BTM-related offense types in the crime index. Weights are constructed from news-article co-mention frequencies in the RavenPack archive, following the text-as-data procedure described in Appendix A. Each NIBRS offense code receives a weight proportional to how often it is mentioned together with Bitcoin ATM terms, with Laplace smoothing. The largest weight goes to false pretenses and swindle, followed by extortion, hacking, and wire fraud. The full weighting scheme covers all 51 offense codes. This figure shows the twelve offenses with the highest weights.

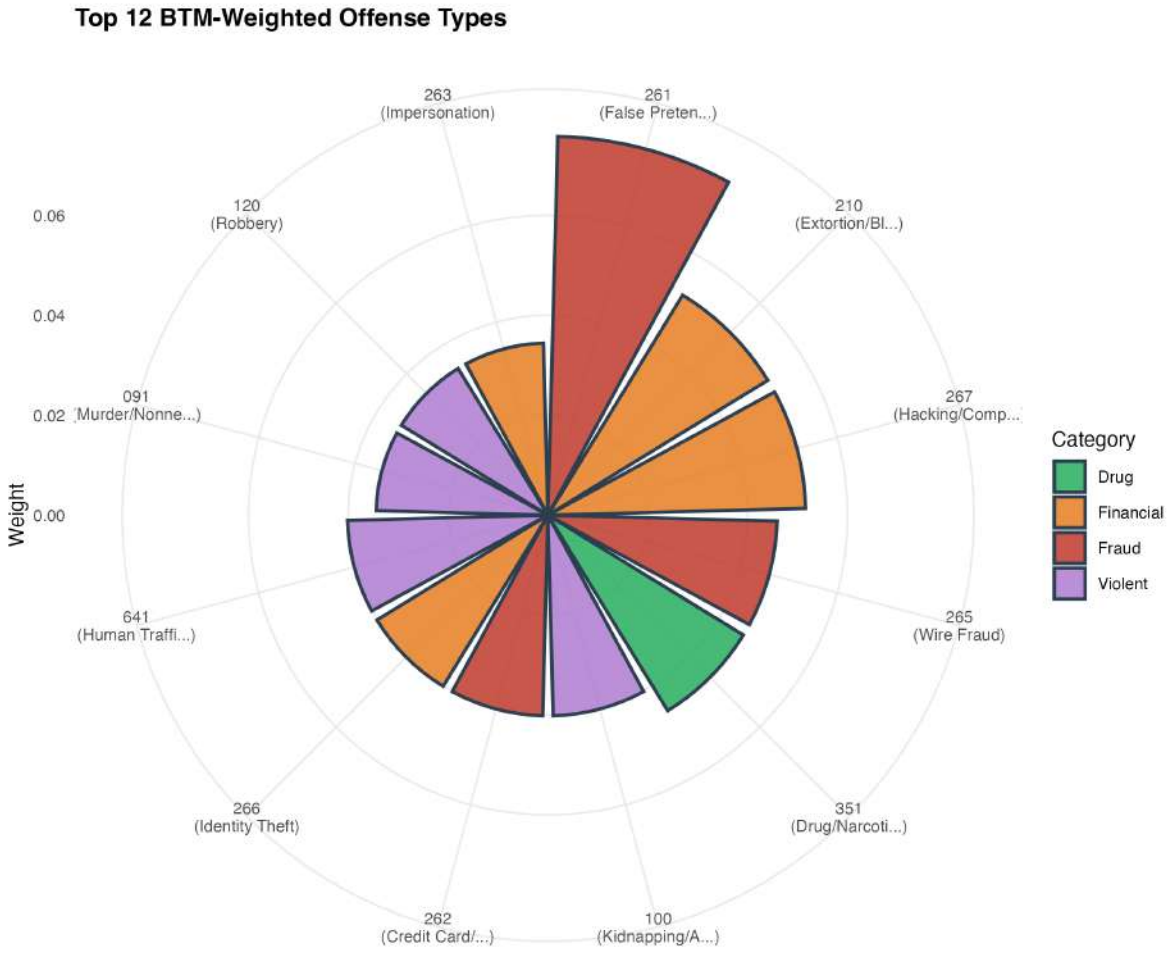


Figure A.3: Event Study: BTM-Exposure Weighted Crime Index

Notes: Callaway-Sant'Anna (2021) difference-in-differences on a weighted crime index. Weights reflect each NIBRS offense code's documented relevance to Bitcoin ATM-related criminal activity. Three panels show effects on weighted crime count, rate per 10,000 population, and log-transformed count. Pre-treatment estimates are flat. The joint pre-trend F-test across $t \in \{-5, \ldots, -2\}$ yields $p = 0.129$, so the pre-treatment coefficients are not jointly distinguishable from zero. Reference period normalized to zero at $t = -1$. Shaded bands show 95% simultaneous confidence intervals. Sample: 10,488 county-year observations, NIBRS 2013–2023. County-level clustering.

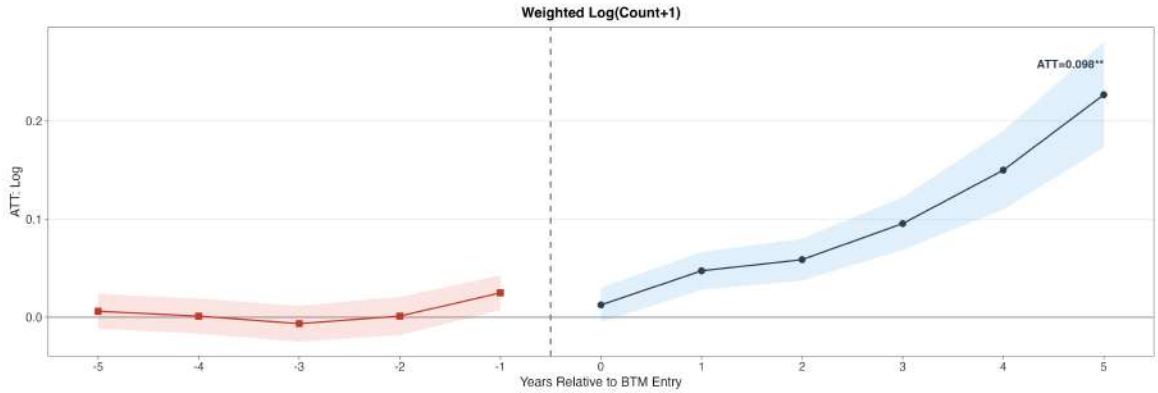


Figure A.4: The Anonymity Supply Chain: Flow of Value Through Entry, Obfuscation, and Exit

Notes: Each node represents a category of service providers within the anonymity supply chain. Providers are organized into three vertically specialized stages: entry (fiat-to-crypto conversion), obfuscation (on-chain transaction graph disruption), and exit (crypto-to-real-value conversion). Right column shows downstream welfare channels: costs (red) and benefits (green). Arrow thickness indicates approximate relative volume. Dashed borders indicate providers subject to OFAC sanctions, law enforcement seizure, or forced shutdown since 2022. Volume labels report annual estimates or post-shock changes. Sources: Chainalysis (2025), Elliptic (2025), TRM Labs (2025), Bitcoin Depot 10-K (2024).

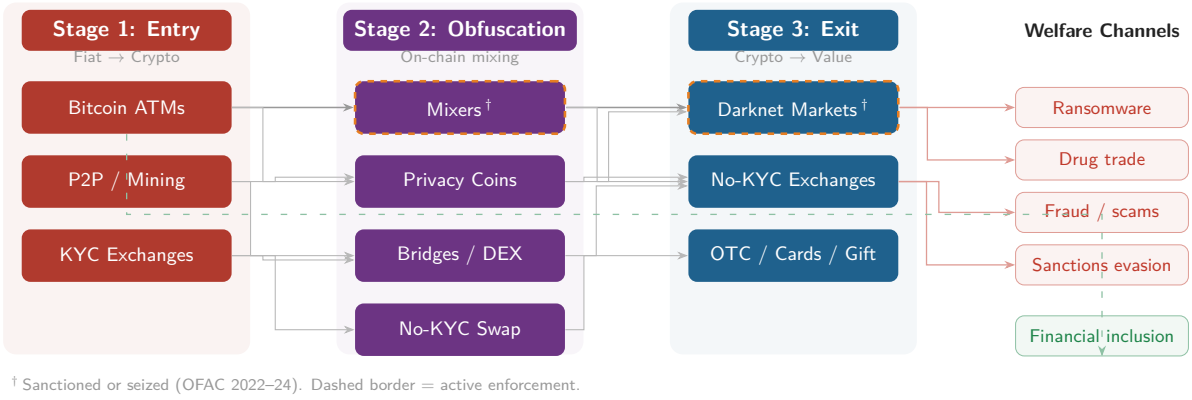


Figure A.5: Operator transaction architecture: four customer-delivery patterns

Notes. The figure decomposes the 247,967 operator-initiated transactions in the validated payout file along two dimensions: the number of inputs and the number of distinct customer outputs. Panel A schematizes the four patterns — peeling chain (one input, one customer plus one change), batch fan-out (one input to $k \geq 2$ customers), composite batched payment (multiple inputs to multiple customers), and sweep-and-pay (multiple operator UTXOs consolidated to fund one customer payment). Panel B compares each pattern's share of transactions to its share of customer payouts. The dominant fact is the asymmetry between the two columns: peeling chains and sweep-and-pay are infrastructure-style operations that account for 45.5% of operator transactions but deliver only 4.9% of customer payouts, whereas batched transactions deliver 95% of payouts from a minority of transactions. Pure internal UTXO management transactions with no customer output are not present in this panel because the clean payouts master is filtered to customer-delivering transactions.

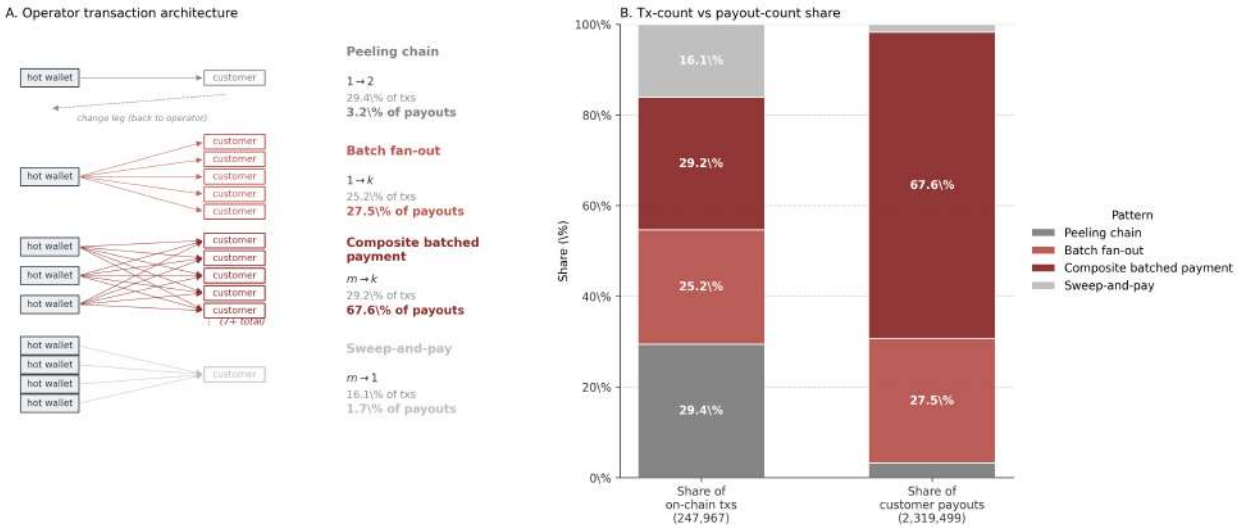


Figure A.6: Quarterly same-day structuring rate and the California DFAL cap

Notes. The figure plots the quarterly share of retail person-days that contain at least two sub-$250 customer payouts summing to $250 or more (structuring at the $250 threshold). The vertical dashed line marks January 1, 2024, when California's Digital Financial Assets Law took effect and imposed a flat $1,000-per-day cap on BTM transactions for California residents. Pre-DFAL (2021Q1–2023Q4) mean structuring rate is 1.97% with peak 2.83% in 2022Q4; post-DFAL (2024Q1–2025Q4) mean is 1.13%, a 43% relative decline. The mechanism is that the $1,000-per-day cap is strictly tighter than Bitcoin Depot's internal $250 KYC threshold and is operator-enforced at the kiosk, removing the need (and ability) for exposed customers to structure around $250. The drop is nationwide because DFAL was followed by similar caps in Vermont (mid-2024), Nebraska, Louisiana, and selected local ordinances, and because Bitcoin Depot's quarterly disclosures attribute approximately 20% of the 2024 revenue decline to the California limits. Sample: 1,173,030 retail wallets / 970,537 retail persons across 2021Q1–2025Q4, mega-clusters excluded.

