Does Market Fragmentation Improve Price Efficiency?[*]

Jonathan A. Brogaard[†] **Anne Haubo Dyhrberg**[‡] **Paige Rowberry**[§]

December 3, 2025

First Version: February, 2025

Chen and Duffie (2021) argue that market fragmentation can improve price efficiency by enabling traders to split their orders across venues, reducing their overall price impact, making their trading more aggressive. Testing this theory in equity markets is difficult due to regulatory requirements that distort market independence. The cryptocurrency marketplace provides an ideal setting given its lack of regulatory interference. Taking advantage of sudden and unexpected changes in coin listings, we use a difference-in-difference identification strategy to show that fragmentation improves price efficiency, consistent with Chen and Duffie (2021). The implications for the role of fragmentation are essential for improving market quality and investor welfare.

Key words: Market Fragmentation, Price Efficiency, Cryptocurrencies

JEL: D40, G14, G23

---

[*] We thank seminar participants at Macquarie University, the Philly Five Conference, and the Women in Microstructure Annual Meeting 2025 for valuable comments.

[†] University of Utah, United States of America, e-mail: brogaardj@eccles.utah.edu

[‡] University of Delaware, United States of America, e-mail: dyhrberg@udel.edu

[§] University of Utah, United States of America, e-mail: Paige.Rowberry@eccles.utah.edu

1. Introduction

The trading of a single security across multiple platforms, known as market fragmentation, is a defining feature of modern financial markets. Before 2000, the New York Stock Exchange and Nasdaq handled 95% of US equity trading volume (Haslag and Ringgenberg, 2023). Today, their combined volume share is less than 35%.[1] The same phenomenon exists in Europe, where the listing exchange only accounts for 40% of the volume share.[2] Despite this persistent and widespread trend, the effect of fragmentation on price efficiency remains unclear. Fragmentation is linked to both improved efficiency (O'Hara and Ye, 2011) and to wider spreads, reduced depth, and higher adverse selection, which reduces price efficiency (Chung and Chuwonganant, 2022, Hendershott and Jones, 2005, Gresse, 2017, Bennett and Wei, 2006, Foucault and Menkveld, 2008). Even studies of natural experiments, like a sudden market merger in Switzerland, find weak effects on efficiency (Dzieliński, Hagströmer, and Qu, 2025).

Chen and Duffie (2021) argue that market fragmentation improves price efficiency. The theory relies on the mechanics of order splitting across markets. In a single centralized market, assuming the trader does not split the orders over time, a large order immediately reveals high demand and drives up the price against the trader. This price impact acts as a cost and discourages trading. Fragmentation allows the trader to split this large order across multiple venues. If prices on these venues do not update immediately, a trader can buy on one exchange without immediately pushing up the price on others. As such, the price impact of the first trade does not instantly spill over to the second market. This separation allows an investor to execute the total order at a lower average cost. The reduction in trading costs induces investors with private information to trade more aggressively. As they trade more, prices reflect their information more quickly and efficiency improves. Our objective is to empirically test this theory and explore whether fragmen-

---

[1] See `https://www.cboe.com/us/equities/market_share/`.

[2] See (2023)European Securities and Markets Authority (ESMA) (ESMA).

tation improves price efficiency.

Testing these theoretical predictions in equity markets is challenging. Under regulatory frameworks such as Regulation National Market System (Reg NMS) and the Markets in Financial Instruments Directive (MiFID II), trading venues in the US and Europe are highly integrated. Order routing rules and best execution requirements link venues so closely that prices update nearly instantaneously across platforms. This limits the scope for traders to benefit from order-splitting across venues as the channel requires prices to update slowly. In addition, finding shocks to fragmentation is challenging. Natural experiments due to regulatory shocks like Reg NMS roll out slowly over time. Venue closures are rare. When they do occur, stock delistings typically coincide with mergers or financial distress. This confounds the effect of fragmentation with changes in fundamentals.

We therefore look outside equities to cryptocurrencies, a market without regulatory interference that offers quasi-exogenous shocks to fragmentation. This market allows us to more directly identify the order-splitting channel and its effect on price efficiency compared to previous work. We make two main contributions with our unique setting. First, we identify a causal link between fragmentation and price efficiency by exploiting shocks to fragmentation. We find that fragmentation monotonically improves price efficiency as predicted by Chen and Duffie (2021). Second, we deliver one of the first empirical tests of the order-splitting channel. We find suggestive evidence that traders split their orders when fragmentation increases and persistent evidence that traders manage to reduce their price impact across orders, consistent with the channel by Chen and Duffie (2021).

The cryptocurrency market provides a valuable natural laboratory for four reasons. First, cryptocurrencies trade across many different venues[3] that operate independently and without cen-

---

[3] Since an exchange has a specific regulatory definition, we describe these cryptocurrency markets as trading venues, although they are colloquially known as exchanges. These venues are known as centralized cryptocur-

tralized regulation. Second, in the absence of best-execution requirements, trade-through rules, or order-routing mandates, prices can diverge across venues (Makarov and Schoar, 2020). This imperfect price correlation across venues is necessary for order-splitting to reduce total price impact: a trader can execute on one venue before the price impact spills over to others.

Third, the cryptocurrency market experiences shocks to fragmentation through venue closures and coin delistings. We exploit two such events to test whether reducing fragmentation harms price efficiency and whether the order-splitting mechanism drives this effect. First, we use the unexpected collapse of FTX in November 2022, which provides a setting to test the Chen and Duffie (2021) predictions during a crisis period. Second, we use routine coin delistings by Binance to test the effect in different market conditions. Together, these events allow us to isolate the effect of fragmentation (e.g. the loss of a venue) from changes in venue competition (e.g. fees or technology).

Fourth, cryptocurrencies offer transparent and granular data across many exchanges. We observe trade and quote data at the millisecond level for 1,768 unique coin pairs across 16 major trading venues from March 2021 to May 2023. For each trade, we observe the execution time, price, volume, and whether the trade is buyer- or seller-initiated. We observe best bid and ask quotes updated in real-time. This granularity allows us to construct precise measures of price efficiency and price impact at the coin level, aggregating across all venues where a coin trades. Collectively, these features allow us to provide a direct empirical test of fragmentation on price efficiency and the order-splitting channel proposed by Chen and Duffie (2021), in a way that is not feasible in traditional financial markets.

Our first identification strategy exploits the unexpected collapse of FTX to test how a sud-

---

rency exchanges as opposed to decentralized exchanges and automated market makers that operate directly on the blockchain. The centralization refers to the venue being the custodian of the market participant's funds, not the centralization of the market.

den, market-wide reduction in fragmentation affects price efficiency. Prior to its collapse, FTX (including FTX-US) was the second-largest venue, representing 22% of trading volume. Its sudden shutdown was largely unexpected. Even sophisticated investors were caught off guard. The speed and severity of the collapse meant that traders could not reallocate funds in advance, making the shock as close to exogenous as possible in financial markets. This eliminated a major trading venue almost overnight and caused an immediate and significant decrease in market fragmentation for many cryptocurrencies.

This event simultaneously affected 139 coins, creating variation along two key dimensions. First, it affected both highly liquid coins like Bitcoin and Ethereum, which traded on many venues, and illiquid coins that traded on only two or three venues. Second, the shock varied in intensity. Some coins had only 5% of their volume on FTX, while others had over 40%. This variation allows us to test the marginal effect predicted by Chen and Duffie (2021): losing a venue should harm efficiency most when there are few venues to begin with.

The FTX shock also addresses the endogeneity problem identified by Babus and Parlatore (2022). Their model suggests traders might strategically fragment markets when investor disagreement is low. If traders have the same valuation of the asset and trade on the same side, they can look to other venues to reduce competition. This implies that a correlation between fragmentation and efficiency might simply reflect underlying disagreement levels rather than the causal effect of fragmentation itself. The FTX collapse mitigates this problem because it forces a change in fragmentation for all coins on its platform, regardless of their disagreement levels or any strategic considerations by traders.

We use a synthetic difference-in-differences (SDID) design (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021) to draw causal inference. Since coins on FTX differ from other coins in liquidity and trading patterns, we need a control group that looks like FTX coins before the collapse. The SDID algorithm builds this control group by reweighting non-FTX coins (and pre-collapse dates) so that the weighted control group closely matches the treated coins' pre-collapse patterns in volatility, volume, fragmentation, and the number of venues. This ensures that any post-collapse differences can be attributed to losing FTX rather than pre-existing differences. The SDID design also helps net out market-wide shocks during the collapse. To the extent that general uncertainty leads to broad changes in trading activity, those effects should be shared by both treated coins and the weighted control group through the common time effects.

The FTX collapse provides a test of both the central price efficiency prediction of Chen and Duffie (2021) and the marginal effect. Their model predicts that losing a venue hurts efficiency most when there are few venues to begin with. Going from three venues to two is much worse than going from ten venues to nine. Our results support this prediction. We find a small average effect for all FTX coins: price efficiency falls by 15% relative to its pre-period mean. The effect is larger for coins that faced a larger shock. For coins with at least 20% of their volume on FTX, the shock is like losing one of five (or fewer) venues. For this group, price efficiency falls by 125% of its pre-period mean. The loss in efficiency grows with the shock. For coins with over 40% exposure, which is similar to losing one of roughly 2.5 venues, the loss of efficiency is 162%.

Our second identification strategy examines routine coin delistings from Binance and complements our FTX analysis in two ways. First, the FTX collapse occurred during a market-wide crisis. While our SDID design controls for market-wide effects, testing the mechanism during normal market conditions provides additional evidence of generalizability. Second, Binance delistings are coin-specific rather than venue-wide shocks.

Because Binance holds a 66% market share, its decision to delist a coin creates a substantial decrease in fragmentation. However, coin delistings are not random. Delistings are triggered by regulatory pressure on privacy coins, project migrations to new blockchains, low liquidity, or prolonged project inactivity. To identify how reduced fragmentation affects price efficiency, we need a counterfactual that accounts for the non-random nature of delistings. We use a propensity score matching design to address this challenge. We run a propensity score based on the pre-period average volume, price in USD, volatility, number of venues, and fragmentation to match each delisted coin with a control coin. We find that after a delisting, price efficiency falls by a remarkable 417% relative to its pre-period mean. The persistence of the efficiency decline, even after controlling for liquidity and depth, indicates that the deterioration in efficiency is not merely a product of reduced market participation.

We then test the order-splitting channel driving our efficiency result. Chen and Duffie (2021) predict that with fewer venues, traders cannot split orders as much. This leads to higher price impact and larger order sizes. We find evidence of this channel in the FTX test. For the average coin, the price impact rises by 32% relative to its pre-period mean. For coins with at least 20% of their volume on FTX, order-splitting becomes much harder, and price impact rises by 162% relative to its pre-period mean. Price impact also rises in the Binance sample, which suggests order-splitting becomes harder, although the price impact results are not as strong once we control for the change in liquidity.

In an ideal setting, we would observe a random shock to the distribution of trading volume without changing the total volume in the market. However, because our shocks coincide with market-wide stress or declining coin fundamentals, we include controls to isolate the effect of fragmentation on price efficiency from the effects of market contraction and changes in coin-specific risk. In our regressions we control for trading volume in USD to account for fluctuations in total market activity. We include volatility to capture the riskiness of individual coins and the heightened uncertainty. To ensure that our results are driven by the distribution of volume (fragmentation) across venues we also control for the number of trading venues. For the Binance delistings, we add a control for the price level, as Binance often removes coins with deteriorating values.

Since our empirical setting is natural experiments, our analysis faces limitations. First, the FTX shock may have altered the composition of the market if risk-averse traders exited entirely. However, trading volumes for our control group remained stable, which suggests that the collapse of FTX did not trigger a universal exit from the asset class. Furthermore, most market participants trade on more than one exchange, such that when FTX collapsed they continued to trade on non-FTX exchanges. Second, the FTX collapse occurred during a period of heightened uncertainty. This raises the concern that our results reflect general market uncertainty rather than a reduction in fragmentation. We address this through our SDID design. Our SDID uses a control group that faced the same market-wide stress but did not lose the FTX venue, and we include time fixed effects. Third, even with an appropriate control group, the FTX shock represents a crisis event, which limits generalizability. We address this by analyzing routine delistings on Binance. The Binance setting introduces a fourth limitation because delisted coins often face deteriorating fundamentals and traders may abandon the coin entirely, which is why we match on price and trading volume.

Our results challenge the conventional wisdom that fragmentation is detrimental and show that it can improve price efficiency and investor welfare. While we do not have a formal model for welfare implications, both the direction and magnitude of our results imply that fragmentation is welfare improving for market participants. This is important given that regulators often view fragmentation as harmful to the welfare of market participants. As financial markets continue to evolve, policymakers should consider the benefits to fragmentation alongside its costs.

The remainder of the paper is structured as follows. Section 2 details the theoretical framework of Chen and Duffie (2021) and maps it to our empirical setting. In Section 3, we present our data alongside our measures for fragmentation and efficiency. Section 4 tests the impact of fragmentation on price efficiency using the collapse of FTX and Binance delistings as identification strategies. We explore the underlying order-splitting mechanism in Section 5. Section 6 concludes.

2. Theoretical Background and Mapping

This section explains how market fragmentation can improve price efficiency, drawing on the model of Chen and Duffie (2021). When traders can split orders across multiple venues, they pay lower trading costs overall, which encourages them to trade more. More trading means more prices reflect more information. Consider a trader who wants to sell a large position. If she must sell everything on one exchange, her own selling pressure pushes the price down against her. The larger her order, the worse the price she receives. This discourages her from trading as much as she would like and incorporating her private information into prices. Now suppose she can access two exchanges and splits her order evenly. Price impact grows faster than order size because it is convex. For example, selling 100 units on one exchange might move the price by 10 cents, but selling 50 units moves it by only 2.5 cents. By splitting her order across two exchanges, the trader faces a 2.5-cent impact on each venue instead of a 10-cent impact on one. Her total trading cost falls.

This requires that prices across venues do not move in perfect tandem. If selling on exchange A immediately pushed down the price on exchange B, splitting would offer no benefit and the trader would face the same total price impact regardless of how she divided her order. But when prices are imperfectly correlated, a trader can execute on the second exchange before it fully reflects her trade on the first. Fragmentation does make each individual exchange thinner, which raises the price impact of any single trade on that venue. However, what matters for the trader is not the price impact on one exchange but the total cost of executing her entire order across all exchanges. Chen and Duffie (2021) show that when prices are imperfectly correlated, the per-exchange price impact rises more slowly than the number of venues. The trader splits her order into more pieces, each piece is smaller, and the convexity of price impact means the savings from smaller orders outweigh the cost of thinner venues. Total trading costs fall, traders respond by trading more aggressively, prices incorporate information faster, and efficiency improves. Because fragmentation reduces the total price impact a trader faces across all venues, even though each individual trade may face steeper impact locally, we measure outcomes at the coin level, aggregating across venues. The benefits of fragmentation are largest when venues are few. Going from one exchange to two produces large efficiency gains, while going from nine to ten produces smaller gains.

2.1 Testing the Theory

To test the predictions of Chen and Duffie (2021) we use the cryptocurrency market. This market provides a useful setting for several reasons. First, the Chen and Duffie (2021) model assumes multiple independent exchanges trade the same asset. Relative to equities, cryptocurrency venues offer a closer match to the independence assumption. In the US, multiple exchanges are linked through the National Best Bid and Offer (NBBO) system, and brokers are required under Reg NMS to route orders to the best price across markets. This effectively integrates equity trading into a single consolidated market. European exchanges are not as tightly linked, but MiFID II's best execution requires brokers to optimize routing across trading venues, which aligns prices across exchanges. In addition, there are pan-European multilateral trading facilities that integrate trading in all listed stocks with the same trading rules. By contrast, cryptocurrency venues operate without a central routing or price aggregation mechanism. Each order book functions in isolation.

Second, Chen and Duffie (2021) assume exchanges are identical. This assumption allows them to abstract from competition between venues. While cryptocurrency venues are not identical, they nonetheless share similar structures. They all operate electronic limit order books and unlike equity exchanges, they all use maker-taker fees to incentivize liquidity provision; none of them are inverted. Some venues have flat fees, while others have tiered fees and volume discounts. Beyond fee structures, cryptocurrency trading venues differ in regulations dictated by the jurisdiction of the venue (e.g., whether the venue is US-based such as Coinbase, or international such as Upbit). These differences represent the competitive landscape of the market. Our paper abstracts from this form of competition, just as Chen and Duffie (2021) do. Our identification strategy makes this possible because the collapse of FTX or a delisting from Binance are shocks to the level of fragmentation (i.e., the number of venues), not to the competitive behavior (e.g., fee schedules, technology) of the surviving venues. Therefore we isolate the causal effect of a change in the number of competitors, directly mapping to the definition of fragmentation in Chen and Duffie (2021).

Third, Chen and Duffie (2021) assume traders can submit orders across exchanges without frictions. Once the accounts are opened and funded, submitting orders across multiple cryptocurrency venues is simple by using their online platform or a trading algorithm connected to the venue's API. Opening an account is similar to setting up a brokerage account in equities. Traders first create an account, verify their identity, and fund it before they can submit orders. The main friction traders face is funding the accounts. Funding takes time, especially for large deposits that require additional verification, similar to delays when transferring money into a retail brokerage account such as Fidelity or Robinhood. However, in equities, placing an order through one broker provides access to all exchanges. A retail investor therefore only needs one login, for one broker, to access all exchanges. In contrast, trading on multiple cryptocurrency venues, such as Binance and Coinbase, requires separate accounts and funding at each, much like needing approval and capital at both Fidelity and Robinhood to trade on two separate exchanges. Some venues, such as Gemini and Coinbase, offer dedicated institutional platforms that help reduce these barriers, but other venues provide fewer institutional services.

In terms of access hours and geography, traders of cryptocurrencies face fewer frictions than in equity markets. Unlike equities, which trade only during set hours[4], cryptocurrency venues operate twenty-four hours a day, seven days a week, and are open to traders worldwide.[5] There are more than 250 active venues globally, making it easy for most individuals to gain access to multiple venues and trade.[6] Furthermore, most investors trade on more than one venue.[7] While traders manage separate balances across platforms, the combination of continuous trading, global accessibility, and direct control over order routing offers a setting that is largely consistent with the frictionless multi-market framework described in Chen and Duffie (2021).

Fourth, Chen and Duffie (2021) assume that prices are imperfectly correlated across venues. This assumption is important for the order-splitting channel to be effective. While anyone can trade cryptocurrencies globally, many venues cater to local fiat-cryptocurrency demand, which creates heterogeneous order flow. This leads to price divergences across venues, sometimes reflecting local demand imbalances or capital controls (Makarov and Schoar, 2020). In equities, by contrast, prices are highly correlated across exchanges because of order routing and cross-venue arbitrage, leaving little independent variation across venues (Aquilina, Budish, and O'Neill, 2022).

We use negative shocks (venue removals) to test a theory modeled on positive shocks (venue additions) and assume that the fragmentation-efficiency relationship is symmetric. We use neg-

---

[4] Several exchanges are proposing extending trading hours pending SEC approval. For example Nasdaq will allow trading 24 hour during weekdays in the second half of 2026

[5] Venues may block users in specific jurisdictions for legal reasons.

[6] According to data from `https://www.coingecko.com/en/exchanges` as of October 2023.

[7] As of 2025, Kraken reports that 69% of cryptocurrency investors maintain accounts on multiple exchanges.

ative shocks because they provide a cleaner identification. A positive shock, such as a new exchange opening, is less exogenous. Traders must choose to fragment their trades and migrate. Such a change would not be immediate, as volume would shift slowly over time. This makes it difficult to define a clear event window.

The exogenous negative shocks also account for what drives fragmentation. While the Chen and Duffie (2021) model largely takes fragmentation as given, other theories examine why markets might fragment endogenously. Babus and Parlatore (2022) show a centralized market can be disadvantageous when traders have similar beliefs because they all try to take the same side of the trade. In such cases, traders have a strategic incentive to fragment the market. This presents a potential endogeneity challenge for empirical work as some assets may be more or less fragmented based on the level of disagreement among the traders, not because of structural changes. In our setting, cryptocurrency valuations are generally heavily debated. This suggests the strategic incentive to fragment may be less prevalent as there are more traders on the opposite side with other beliefs. Nevertheless, traders of well-known coins, like Bitcoin, may still have correlated valuations. These traders might benefit more from fragmentation, according to Babus and Parlatore (2022). Therefore, there is room for endogeneity in market selection across coins. We account for this issue by using the FTX collapse and Binance delistings as negative shocks so that a change in fragmentation is not the decision of the trader, but external.

Although the cryptocurrency market fits the Chen and Duffie (2021) model framework well in terms of many independent venues with imperfectly correlated prices and low friction order submissions, we do face data limitations. As we discuss in more detail in Section 3, we have trade and quote data, not order level data.

3. Data and Metrics

We use high-frequency trade and quote data from Tardis.dev. For each trade, we have data on the execution time to the millisecond, the price, volume, and an indicator of whether the trade is buyer- or seller-initiated. Tardis.dev provides the best bid and ask quotes, updated every time there is a change to the quote price or size. Our dataset spans from March 2021 to May 2023 and includes 1,768 unique coin pairs across 16 trading venues: Binance, Coinbase, FTX, Gate-io, Kraken, KuCoin, Binance US, Bybit, Crypto.com, Bitstamp, FTX US, Gemini, AscendEX, Upbit, Bitflyer, and Coinflex. We focus exclusively on spot trading pairs to ensure comparability across trading platforms. We restrict the sample to coins traded on at least two venues to ensure continuity of trading after the collapse of FTX and the delisting events on Binance. We process the raw trade and quote data through several steps to remove outliers and standardize naming conventions across trading venues. For more detail on this process, see Appendix A.

We measure price efficiency using variance ratios and autocorrelation coefficients, based on the random walk hypothesis proposed by Lo and MacKinlay (1988). In an efficient market, price changes follow a random walk and returns are uncorrelated over time. We calculate the autocorrelation of returns for each coin at the one-minute interval. The autocorrelation coefficient at lag k is:

$$\rho_k = \left| \frac{\mathrm{Cov}(\mathrm{Ret}_t, \mathrm{Ret}_{t-k})}{\mathrm{Var}(\mathrm{Ret}_t)} \right| \tag{1}$$

where Ret_t is the return at time t. In an efficient market, ρ_k should be close to zero for $k > 0$.

The variance ratio compares the variance of returns over different time horizons. We compute it as:

$$\text{Variance Ratio} = \left| \frac{\text{Var}(\text{Ret}_{c\text{-minute}})}{b \times \text{Var}(\text{Ret}_{a\text{-minute}})} - 1 \right| \tag{2}$$

where $c = a \times b$, and a and c represent different time intervals (e.g., 1-minute and 5-minute). A variance ratio close to zero indicates returns follow a random walk. To make sure that we include both short-term and longer-term variations we compute different horizons: 15-second to 1-minute returns, 1-minute to 5-minute returns, 5-minute to 15-minute returns, and 15-minute to 30-minute returns. We take the absolute value and multiply the variance ratio and autocorrelation measures by -1 so that higher values correspond to higher efficiency.

We compute the daily volume-weighted price impact at different future horizons (h) in basis points:

$$\text{Price Impact}_t(h) = \sum_{i \in \text{trades on day } t} \frac{2 \times D_{it} \times (M_{t+h} - M_t)}{M_t} \times 10,000 \tag{3}$$

where for each trade i, D_{it} is the trade direction (+1 for buys, -1 for sells), M_t is the mid-quote before the trade, and M_{t+h} is the mid-quote at h minutes (or seconds) after the trade at time t. Since our sample consists of both highly liquid and less liquid coins, we include a variety of time horizons : 30-second, 1-minute, 5-minute, and 30-minute intervals. We volume-weight across trades.

We use principal component analysis (PCA) rather than reporting individual measures across horizons or simple averages for several reasons. First, PCA accounts for the correlation structure among our measures. Autocorrelation at various lags and variance ratios across different time horizons are correlated by construction, as they all capture deviations from the random walk hypothesis. Simple averaging would give equal weight to what are essentially overlapping signals, whereas PCA extracts the common variation. Second, PCA provides an optimal weighting scheme. The first principal component captures the maximum variance in the data, creating a more powerful test than arbitrary equal weighting. Third, our measures have different scales and variances. PCA automatically standardizes these differences, whereas simple averaging would inappropriately give more weight to measures with larger numerical ranges. By using PCA, we create a summary measure that reflects the shared informational content across our multiple efficiency proxies while addressing scaling and correlation issues. Finally, using principal components allows us to present results parsimoniously rather than reporting results across 16 separate regressions (four efficiency measures and four price impact measures across two natural experiments).

Our main conclusions are robust to this aggregation choice. In the internet appendix, we replicate our core results using each individual efficiency and price impact measure as the dependent variable. The signs, magnitudes, and statistical significance remain consistent across individual components of efficiency and price impact measures. We therefore present the principal component results (*Price Efficiency PC* and *Price Impact PC*) in our main tables for parsimony.

3.1 Measuring Fragmentation

Although the cryptocurrency market is highly fragmented in terms of the number of available venues, trading volume is largely concentrated among a small number of large platforms. Table 1 reports daily trading volumes (in USD), the market share, the number of symbols listed on each venue, and the median trade size (in USD).[8] The market share is defined as a venue's median daily volume relative to the total median daily volume across all venues. Binance alone accounts for almost 66% of total volume, while the three largest venues, Binance, Coinbase, and FTX, together capture 83% of total volume.[9] Trading activity is also concentrated across coins. The

---

[8] All volumes are converted to USD using Binance's daily average midpoint price. For coins not listed on Binance, we instead use the venue with the longest trading history, breaking ties by highest trading volume.

[9] Prior to its collapse, FTX and FTX-US had a market share of approximately 22%.

ten largest coins represent 47% of total daily volume, and four of these coins are only traded on Binance and Binance US.

While the venues are similarly structured, they nonetheless cater to different traders. This is reflected in the breadth of coins offered. Table 1 shows that Binance, despite being the largest venue by volume, does not list the most coins. Gate-io, with only 3.4% market share, lists 1,071 spot coins during the sample period. This pattern suggests that some venues emphasize a small set of highly liquid coins, while others focus on offering access to a wide variety of smaller, less liquid coins. The clientele differences are further reflected in trade size. Venues with relatively few listed coins, such as Bitstamp (168) and FTX US (49), exhibit the largest median trade sizes ($2,224 and $2,893, respectively). By contrast, venues with extensive listings, such as Gate-io (1071), have much smaller trades (median size $259). These differences reflect the variation in local traders as discussed by Chen and Duffie (2021), which is necessary for prices to be imperfectly correlated across venues.

[**Table 1** here]

Chen and Duffie (2021) measure fragmentation as the number of exchanges. In their model, this definition is natural because exchanges are assumed to be identical, with the same share of trading volume. Under that assumption, fragmentation and the number of exchanges are equivalent.

In practice, trading activity is uneven across venues. Simply counting venues therefore overstates fragmentation when one or two venues dominate volume. For example, two venues with a 50/50 split of volume are more fragmented than two venues with an 80/20 split, even though the number of venues is the same. As shown in Table 1, overall volume in cryptocurrencies is concentrated among a few large trading venues. To capture fragmentation empirically, we measure fragmentation using the Herfindahl–Hirschman Index (HHI), which quantifies how concentrated trading volume is across venues. Specifically, for each coin-day we compute:

$$\mathrm{HHI}_{it} = \sum_{e=1}^{E} \left(\frac{\mathrm{Volume}_{ite}}{\sum_{e=1}^{E} \mathrm{Volume}_{ite}} \right)^2 \tag{4}$$

where Volume_{ite} is the trading volume of coin i on trading venue e at time t, where E is the total number of trading venues trading coin i at time t.

Following Haslag and Ringgenberg (2023), we define fragmentation as one minus the HHI:

$$\mathrm{Fragmentation}_{it} = 1 - \mathrm{HHI}_{it} \tag{5}$$

A higher fragmentation value means trading volume is more evenly distributed across trading platforms, while a lower value indicates volume concentration on fewer trading venues. Across all coins the average level of fragmentation is 0.17 and the average number of venues is 2.4 as shown in Panel A of Table 2. Since interpreting the level of fragmentation can be difficult, consider this example. Suppose ETH-USDT is traded on two venues: 90% on Binance and 10% on Kraken. This results in a HHI and fragmentation level of:

$$\mathrm{HHI} = 0.9^2 + 0.1^2 = 0.82, \quad \mathrm{Fragmentation} = 1 - 0.82 = 0.18.$$

This scenario reflects the average coin in our sample and shows that trading volume is relatively concentrated. The standard deviation of 0.2 implies that a one-standard-deviation move from the mean corresponds to an 80/20 volume split between Binance and Kraken in our example. When we look at the extremes, we observe coins with very high fragmentation. For instance, ETH–BTC trades on all but one venue, and BTC–USD exhibits the highest average fragmentation across days in our sample at 0.69. The maximum observed fragmentation is 0.82 for DAI–USD.

If we compare the most (Panel B) and least (Panel C) liquid coin terciles, based on USD trading volume, we see that the most liquid coins are traded across the highest number of venues (3.3 on average) and exhibit the greatest fragmentation. By contrast, the least liquid coins trade on only 1.6 venues on average, with a median of just one venue. Recall that to ensure the continuity of trading after the FTX collapse or Binance delisting, we restrict our sample to coins that are traded on at least two venues at some point during the sample period. However, even the least liquid coins have instances of highly fragmented volume with maximum values close to that of the most liquid coins.

[**Table** 2 here]

This variation of fragmentation based on volume is important to understanding of the marginal effect of fragmentation. In our sample we have coins where FTX or Binance represent a substantial or small proportion of the volume. When FTX collapses or Binance delists a coin, the shock therefore varies in size. As previously described, Chen and Duffie (2021) measure fragmentation as the number of venues while assuming that volume is equally distributed across venues. To map our volume shocks to fragmentation to the predictions in Chen and Duffie (2021), we consider two scenarios. In the first scenario, one coin is traded on two venues, each with 50% of volume. In the second scenario, one coin is traded on 10 venues each with 10% of volume according to the model. When one venue closes, the number of venues is only reduced by one, but the volume shock is 50% in scenario one and 10% in scenario two. In our setting this would be reflected in FTX having 50% or 10% of volume in a coin prior to its collapse. In Chen and Duffie (2021), adding venues improves price efficiency, but at a diminishing rate. As such, going from 9 to 10 venues has a smaller impact on price informativeness than going from 1 to 2 venues. We test this marginal effect by segmenting coins into buckets based on how much volume was traded on FTX prior to its collapse. We describe this analysis in the next section.

The following analysis uses our measures in a regression format, so we first present the pairwise correlations in Table 3, which highlights the relationships between our variables. Both the number of venues and fragmentation are positively correlated with our efficiency measures and negatively correlated with price impact. The number of venues, for instance, has a correlation of 0.27 with Efficiency PC and -0.23 with Price Impact PC. Fragmentation shows a similar but weaker pattern, with a correlation of 0.06 with Efficiency PC and -0.11 with Price Impact PC. This suggests that as trading activity spreads across more platforms, prices tend to be more efficient and the price impact across venues is lower.

[**Table 3** here]

4. Testing the Impact of Fragmentation on Efficiency

This section tests the hypothesis that an exogenous reduction in market fragmentation reduces price efficiency. To implement this test, we exploit two distinct shocks: a market-wide venue closure and localized asset-level delistings.

4.1 Venue-Level Shock to Fragmentation

We begin by examining the collapse of FTX and FTX US in November 2022 as a large, plausibly exogenous shock to fragmentation. This event is useful for four reasons. First, the collapse was a genuine shock. While there were warning signs, the speed and severity of FTX's downfall were not widely anticipated. Even sophisticated investors appear to have been caught off guard.[10] This limits the scope for traders to reallocate funds in advance and reduces concerns that anticipatory behavior drives our results.

Second, while FTX was the second-largest trading venue, representing 22% of volume in the month before the collapse, the volume share of FTX varied across coins. This is important for understanding the marginal effect of fragmentation. Recall that Chen and Duffie (2021) predict that fragmentation improves price efficiency, but at a decreasing rate. For a coin like Bitcoin that trades on many exchanges with more evenly distributed volume, the removal of FTX is a smaller shock than for a coin where FTX handles the majority of global volume. As such, we expect coins with a smaller FTX volume share to experience a smaller deterioration of price efficiency and vice versa.

Third, the coins on FTX also vary in terms of liquidity. The shock impacted both the most and least liquid assets simultaneously. This is a key advantage, as it reduces concerns that our results merely reflect a pre-existing correlation between liquidity and fragmentation. Specifically, Haslag and Ringgenberg (2023) show that fragmentation improves market quality for large stocks while small stocks experience worse market quality.

Finally, the shock allows us to address the endogeneity problem identified by Babus and Parlatore (2022). Recall that their model suggests traders might strategically fragment markets when investor disagreement on an asset is low. This creates a significant challenge for empirical work. A simple correlation between fragmentation and efficiency might just capture this underlying disagreement level as traders intentionally fragment the market looking for liquidity, not the causal effect of fragmentation itself. The surprise of the FTX collapse mitigates this problem. It is as close as we can get to an exogenous event that forces a change in fragmentation for all coins on

---

[10] Financial News (Nov. 14, 2022): “Hedge funds ‘extremely surprised’ by FTX collapse”

its platform, regardless of their investor disagreement levels.

The events that led to the collapse of FTX are multifaceted. The trading venue was founded in 2019 by Gary Wang and Sam Bankman-Fried, who had started a proprietary-trading and market making firm years before called Alameda Research. The two entities were supposed to be independent, but in reality they were heavily integrated. Alameda was able to borrow unlimited funds from FTX without the risk of margin calls or liquidations. Most of the debt was collateralized by FTX's native and illiquid coin FTT, making Alameda vulnerable to volatility in the price of FTT.

In the span of 10 days, FTX went from being the second largest cryptocurrency trading venue to going bankrupt. On November 2nd, 2022, a leaked Alameda balance sheet [11] revealed their reliance on FTT. Market participants started to withdraw funds from FTX and Binance announced a large sell-off of FTT. As a result, the price dropped. On November 7th, FTX announced that they faced a liquidity shortage and sought a bailout from venture capital funds and Binance. Initially Binance proposed to purchase the non US part of FTX, but walked away from the deal on November 9th. On November 11, 2022, FTX filed for bankruptcy. Sam Bankman-Fried and the CEO of Alameda Research, Caroline Ellison, were arrested in December and have since been sentenced to 25 and 2 years, respectively, for fraud and money laundering.

Figure 1 shows the total daily trading volume (in USD) for coins on FTX near the collapse date. Despite the news stories, volume remained relatively steady before November 8, 2022, and then spiked as traders ran on the venue. Volume decreases sharply when the shutdown becomes imminent, reaching zero on November 11. Fragmentation decreases along with the November 8th run on the venue as evident from the sharp decline in Figure 2. Fragmentation drops from 0.37 to 0.24 and remains low 30 days after for coins traded on FTX. In contrast, coins not traded on FTX show stable fragmentation levels throughout this period. Thus the collapse of FTX acted as

---

[11] CoinDesk: Alameda balance sheet story (Nov. 2, 2022)

a concentrated shock to its listed coins, while the broader market did not experience a meaningful change to fragmentation. Given this clear reaction in volume and shock to fragmentation, we use November 8 as our event day in the subsequent analysis. Alternatively, we use an event window, where we block out the days from November 2 to November 10 and compare 30 days before and after. The results for price efficiency are reported in Internet Appendix Table A.1, and the results for price impact are reported in Internet Appendix Table A.2. Both are similar in significance, sign, and magnitude to the main results.

[**Figure 1** here]

[**Figure 2** here]

To make causal inference about the effect of fragmentation on price efficiency, we need a suitable counterfactual for the coins traded on FTX. The set of coins listed on FTX is not random. As we discussed in Table 1, FTX, like other venues, lists a mix of both highly liquid and illiquid coins but lists a relatively small selection of 175 coins. We therefore construct a control group that resembles the treated coins in the pre-period. To do so, we employ a Synthetic Difference-in-Differences (SDID) approach (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). SDID combines elements of Difference-in-Differences and Synthetic Control by choosing weights by coin (and by pre-period date) so that the weighted control group closely matches the treated group's pre-treatment outcome trajectory, while also balancing selected pre-period characteristics. Unlike propensity-score matching, SDID does not require selecting a single matched control for each treated coin. Instead, it considers all coins not listed on FTX and assigns higher weight to control coins whose pre-period behavior is most similar to the treated group, as well as downweighting pre-period days that would otherwise exert undue influence. In our implementation, we construct the synthetic control by weighting control coins to align with the treated coins in the pre-period on volatility, logged trading volume, fragmentation, and the number of trading venues. We use logged trading volume to capture liquidity differences and volatility to capture risk. We also include fragmentation and the number of venues to ensure that treated coins are compared to controls with similar pre-period market structure. We require that treated and potential control coins trade on at least two venues and trade every day in the 30 days before and after the collapse of FTX. This leaves us with 139 treated coins and up to 1,404 potential control coins.

The SDID approach also helps address potential reverse causality between market fragmentation and liquidity, as discussed by Dzieliński, Hagströmer, and Qu (2025). Because more liquid coins tend to trade across a greater number of venues, it can be difficult to determine whether fragmentation increases liquidity or whether liquidity attracts fragmentation. By constructing a synthetic control group that matches treated and untreated coins on pre-period liquidity measures, specifically volatility and trading volume, the SDID design helps ensure that any post-treatment differences can more plausibly be attributed to the loss of the FTX venue rather than to pre-existing liquidity differences.

Table 4 shows the average pre and post-period levels of the variables we use to compare the treated coins and their synthetic controls: the number of venues, fragmentation, volume, and volatility.[12] The table shows that the coins in the synthetic control group do not experience meaningful changes to volume or volatility. They are traded on exactly the same number of venues throughout the period and the volume distribution across venues barely changes, as evident by the 0.01 decrease in fragmentation. In contrast, the treated coins on FTX experience significant reductions to the number of venues, fragmentation, and volume, as well as an increase in volatility after the collapse of FTX. The significant volume reduction for the treated group is expected, as

---

[12] Table 4 shows results summary statistics for matches based on the price impact outcome, not the efficiency outcome variable. The summary statistics are slightly different in the pre-period for volume and number of exchanges when we instead match on efficiency. Since they are so similar, we only show the results for price impact.

many traders were unable to exit and withdraw funds from the exchange before it shut down. The stable volume in the synthetic control group, however, demonstrates that this was not a market-wide exit from the asset class. While the pre-period number of venues, volume and volatility are not exactly the same for the treated and control groups, they are as close as possible given the sample of control coins available. Turning to the outcome variables, both treated and control coins experience decreases in price efficiency in the post-period, with the treated group declining by 0.20 compared to 0.15 for the synthetic controls. Price impact increases for both groups, but the increase is larger for treated coins (0.62) than controls (0.29), consistent with the predicted effect of losing a trading venue.

[**Table** 4 here]

Turning to the outcome variable price efficiency, Table 4 shows that the treated and control coins experience a decrease in efficiency in the post-period. Both groups follow parallel trends in the pre-period and then start to diverge as evident in Figure 4, which compares the pre and post-period for the coins with at least 10% volume on FTX and their synthetic control group. Prior to the collapse the coins on FTX and their controls follow a similar pattern. However after the bankruptcy, coins on FTX experience an immediate drop in efficiency, which remains low throughout the post-period. On the contrary, the control group keeps a somewhat stable level, slightly decreasing in the post-period. This divergence between treated and control groups in the post-period suggests that the efficiency losses among FTX coins cannot be explained by market uncertainty alone.

[**Figure** 4 here]

To implement the SDID, we estimate unit and time weights using pre-treatment data following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021) so that the weighted control group matches the treated group's pre-treatment outcome trajectory and selected covariates. We then estimate the standard DiD specification below on the resulting SDID-weighted sample:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}, \tag{6}$$

where Outcome_{it} is the principal component of efficiency for coin i on date t. Post_t is a dummy variable which is equal to 1 on November 8, 2022 and 30 days after, while 0 before. The dummy variable Treat_i equals 1 if coin i is traded on FTX and 0 otherwise. X_{it} includes control variables (log volume, number of trading venues, fragmentation, and volatility), α_i is a coin fixed effect, and Θ_t is a date fixed effect. SDID estimates unit weights for control coins and time weights for pre-period dates to minimize pre-treatment imbalance in outcomes (and included covariates). We implement SDID by applying these weights and estimating the DiD coefficient on $\text{Post}_t \times \text{Treat}_i$. Because identification comes from this reweighting, we do not interpret or report coefficients on X_{it} or summary fit statistics like R^2 for the SDID specifications. Instead, we assess fit by verifying close pre-treatment alignment, shown in Figure 4 and Internet Appendix Figures A.1 and A.2.

We first test the main hypothesis that a reduction in fragmentation harms price efficiency. As shown in Table 5, Column 1, the FTX collapse led to a modest average reduction in price efficiency of -0.13, or 15% of the pre-period mean, but it is only significant at 10%. This modest average effect is expected and consistent with the marginal effect of fragmentation discussed by Chen and Duffie (2021). Out of 139 treated coins, 109 had less than 10% of their volume on FTX, which translates to a small shock (similar to losing one of at least 10 venues). The effect on price efficiency becomes much stronger when we analyze coins that experienced a larger centralization shock, as predicted by the theory. We segment the treated group into coins that had at least 10%, 20%, 30%, or 40% of their trading volume on FTX, which maps to the model's setup of losing one of 10, 5, 3.3, or 2.5 venues, respectively. As shown in Columns 2-5, the negative effect on price efficiency is monotonic and highly significant. When FTX has at least 20% volume share (losing one of five venues), price efficiency decreases by 125% (-1.09) of the pre-period mean. The effect grows to a 162% decrease for coins with over 40% volume share.

[**Table** 5 here]

In addition to the cutoff analysis, we model the treatment effect as a continuous variable in Table 6. This allows us to test the marginal impact of the shock while controlling for differences across treatment intensities. As shown in Column 1, the coefficient on price efficiency is -1.48 and significant at the 1% level. Using the average pre-collapse FTX volume share of 0.12 and a pre-period mean of 0.87, this implies that efficiency decreases by 20% for a typical treated coin ($\frac{0.12\times-1.48}{0.87} \approx -0.20$).

A potential concern is that our results reflect a temporary exit of liquidity providers rather than a structural change in price efficiency. If the FTX collapse simply forced liquidity providers out of the market, the drop in efficiency might be a short-term crisis that reverses once these providers migrate to surviving venues. To test for this, Column 2 of Table 6 splits the post-period into an early phase (days 0–22) and a final week (days 23–29). If the effect were merely a migration delay, price efficiency should mean-revert as liquidity pools on surviving exchanges. Instead, the coefficients for price efficiency are statistically indistinguishable across both periods. The decline in efficiency is persistent, remaining at -1.57 in the final week. This persistence suggests that the efficiency losses are not driven by temporary market confusion or trader exit, but by the structural loss of the trading venue.

[**Table** 6 here]

4.2 Asset-Level Shocks to Fragmentation

Our FTX results demonstrate that sudden, market-wide reductions in fragmentation harm price efficiency. While our SDID design controls for common time effects during the FTX collapse a question remains about generalizability. Does the fragmentation-efficiency relationship hold outside of extreme market conditions? The FTX collapse occurred during a period of massive liquidations and heightened uncertainty. Even with our control group, the mechanism might operate differently when traders are calm versus when they are in crisis mode. To test whether our findings generalize beyond uncertainty periods, we examine routine coin delistings from Binance that occur across different market conditions, including periods of relative market stability.

Binance is the largest trading venue in the cryptocurrency market, with nearly a 70% market share in our sample, making delistings a substantial shock to fragmentation. Unlike the collapse of FTX, which was a rare venue-level shock affecting a broad set of coins simultaneously during a crisis, delistings are coin-level shocks that occur routinely across varying market conditions. These delistings are announced in advance and occur for reasons unrelated to recent changes in price efficiency or broader market stress. Binance announces delistings on its website seven days before they take effect, giving users time to sell or withdraw funds. On the delisting day, trading is suspended, deposits are stopped within 24 hours, and users are able to withdraw the delisted coin from the platform. The withdrawal window lasts up to a few months after the delisting. Traders are not at risk of losing their funds. Binance ultimately converts any remaining balance into stablecoin after withdrawal deadlines, ensuring users can recover their capital even if they fail to withdraw the delisted coin. This contrasts sharply with the FTX collapse, where traders faced recovery uncertainty. Since the delisting is announced before trading ends, we use the announcement date as the event date and then remove the seven days between the announcement and delisting.

During our sample period from March 2021 to May 2023, we identify 34 delisting events for coins that continued trading on at least one other venue. As detailed in Internet Appendix Section C, these events span diverse market conditions and stem from various causes, including low liquidity, regulatory compliance, and project instability. The full list of delisted assets and event dates is reported in Internet Appendix Table A.3. Although these events generate fragmentation shocks, several also coincide with coin-specific fundamentals that can influence efficiency. It is therefore important that we build an appropriate control group.

Each delisting type generates the same fundamental shock: traders lose access to Binance as a trading venue. Figure 5 plots daily trading volume in USD 30 days before to 30 days after a delisting announcement. Before the announcement the average daily trading volume across the delisted coins hovers around two million USD. Volume then rises to roughly five million USD two days before the announcement. On the announcement day, trading volume spikes above 20 million USD, and the following day it climbs further to about 23 million USD. Afterward, volumes decline sharply and fall below pre-announcement levels, as these coins continue to trade only on smaller, non-Binance venues. The spike indicates abnormal trading volume as traders exit their positions prior to the delisting.

[**Figure 5** here]

Figure 6 illustrates the coinciding change in fragmentation surrounding Binance delistings. Prior to the delisting announcement, fragmentation is 0.25 on average. The day before the announcement, as trading volume spikes in Figure 5, fragmentation rises slightly. After the announcement, however, fragmentation falls back to 0.25 within two days, and declines further to 0.20 by day 8, reaching 0.15 by day 10. By day 30, fragmentation drops to a low of roughly 0.10. Since the market is reacting to the delisting in the days after the announcement, we remove one week after the announcement and then compare 30 days before and after for our main tests.

[**Figure** 6 here]

To identify how reduced fragmentation affects price efficiency, we need a counterfactual that accounts for the non-random nature of delistings. To do so, we take each delisted coin and run a propensity score match based on pre-period averages.[13] We match on the same variables as in the FTX analysis (logged volume, volatility, number of venues, and fragmentation) to control for liquidity, risk, and fragmentation levels. Additionally, we match on the price level (USD). This addition is necessary because, unlike the broad FTX shock, Binance delistings often target assets with specific price characteristics. Extremely low nominal prices can signal abandoned tokens or stalled development. We use these criteria to match one-to-one and create a control group of 34 coins to our 34 delisted coins.

Table 7 shows the pre and post averages of the matching variables and outcome variables for the delisted coins and the control coins. The match of delisted and controls is close when considering the pre-period number of venues, fragmentation level, volume, and volatility. In addition, we only observe significant changes to the matching variables for the delisted coins, not the controls: the number of venues and fragmentation decrease by 38% and 48% of the pre-period mean, respectively, while the control group sees no changes to either variable. Volume decreases for both groups but the change is only significant at 10% for the control group. Volatility increases significantly for delisted coins (0.02), while controls show no significant change. The match on price is more difficult. The delisted coins have a much higher average price than the controls $(2,102 versus 81)$, though neither group experiences a significant price change after the delisting. Turning to the outcome variables, the patterns align with our predictions. Price impact increases significantly for delisted coins (0.80), while controls experience a slight and insignificant decline (-0.06). Price efficiency drops sharply for delisted coins (-0.98), compared to a modest -0.04 for

---

[13] We are unable to use the SDID approach as it requires each coin to trade every day. Since some of these coins are delisted due to low liquidity, we use the alternative method of propensity scores which matches based on the pre-period average.

controls. The results are presented graphically in Internet Appendix Figure A.3.

[**Table** 7 here]

[**Figure** 7 here]

To estimate the impact of delisting on price impact and price efficiency, we use the following coin-level difference-in-difference specification:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \Theta_t + \alpha_i + \varepsilon_{it} \tag{7}$$

where Outcome_{it} is efficiency or price impact for coin i on event day t. The event day ranges from 30 days before and after the delisting announcement, with the 7 days immediately after the announcement removed. The interaction term, $\text{Post}_t \times \text{Treat}_i$, equals 1 for the 30 days after the announcement and 0 for the 30 days before, for the delisted coins. Since we include coin (α_i) and event day (Θ_t) fixed effects the individual Treat_i and Post_t variables are omitted. We include coin fixed effects (α_i) to absorb time-invariant heterogeneity across assets. The inclusion of event-day fixed effects (Θ_t) controls for common dynamics relative to the announcement (e.g., typical trading patterns on the day of a delisting news release). It is worth noting that this staggered structure, where delistings occur on different calendar dates, allows us to distinguish event-time from calendar-time dynamics. This contrasts with our FTX analysis, where the shock occurred simultaneously for all assets. In that single-event setting, event-time (days relative to the collapse) and calendar-time (actual dates) are perfectly collinear, making it impossible to include separate fixed effects for both. We control for time-varying characteristics X_{it}: *Number of Venues*, logged USD *Volume*, *Volatility*, *Price*, and *Depth*. The price control is important in the Binance setting because many delisted coins exhibit substantial declines in value before delisting. Including price ensures that the effects are not simply reflecting differential price trajectories between treated and control coins. In contrast, the FTX collapse affected both high- and low-priced assets simultaneously, so price movements are largely absorbed by the event-day fixed effects and correlated volume changes.

Interpreting the Binance results requires us to address the challenge that delistings generate coin-specific market contractions rather than a simple redistribution of trading activity. Binance accounts for 73% of trading volume for the average delisted coin and over 90% for 11 delisted coins. These coins trade on few other venues. The average is 2.8 venues. Consequently, traders may find it unprofitable to migrate to another trading venue and abandon the coin. Thus, the market for that coin may contract as well as centralize. We see evidence of this in Table 7 and Internet Appendix Figure A.3. Trading volume drops by 16% and the available depth decreases by 88% relative to the pre-period mean. This indicates that the redistribution of volume is incomplete within 30 days, or that traders exit the coin. We therefore control for volume and depth to hold observable market participation constant. We further control for the number of venues to isolate the change in volume distribution from the reduction in trading outlets. However, this introduces a bad control problem. The delisting causes the contraction in liquidity and the drop in venue count. Controlling for these variables effectively absorbs part of the treatment effect. We therefore interpret the specification with controls as a conservative lower bound. It tests if price efficiency degrades due to fragmentation itself rather than from the loss of market participation.

Table 8 reports the price efficiency regression coefficients. Columns 1 and 3, which estimate the regression without controls, confirm the univariate results and show a significant drop in efficiency (-0.95 and -0.91 depending on the fixed effect specification). In Columns 2 and 4, we add the control variables. The coefficient for price efficiency is smaller relative to the base specification, but it remains statistically significant and economically large, showing a remarkable decrease of -0.64, or 417% relative to the pre-period mean. Since the effect remains strong after controlling for market participation, trader exit cannot fully explain the decline in efficiency. We thus attribute this result to the decline in fragmentation.[14]

[**Table 8** here]

5. Testing the Order Splitting Mechanism

The results in the previous section establish that venue closures and delistings lead to a decline in price efficiency. We now turn to the mechanism behind this decline. Chen and Duffie (2021) propose that fragmentation improves efficiency by enabling informed traders to split parent orders across multiple venues. This behavior mitigates the price impact of any single trade and encourages more aggressive, informative trading. When venues are removed, this channel is restricted. Consequently, traders must execute orders on fewer venues, leading to higher price impact.

If the order-splitting mechanism is the primary driver of our results, we should observe an increase in price impacts when fragmentation declines. Turning to SDID estimation in Table 9, we see that the group of all coins traded on FTX experiences a significant increase in price impact of 0.20 (Column 1), or 32% relative to the pre-period average. This suggests that even for the average coin, the loss of a single venue immediately made it harder to trade without moving prices. Just as with efficiency, this effect on price impact is monotonic in the shock's intensity. Columns 2–5 show that as the shock to fragmentation increases, so does the price impact. For the 20% volume group, price impact increases by 162% (1.02) relative to the pre-period average.

---

[14] To ensure our main findings are not solely driven by the principal component (PC) analysis, we replicate the tests in this section using individual measures of efficiency as dependent variables. Internet Appendix Table A.4 and Table A.5 estimate the SDID and DID specifications on autocorrelation and variance ratios. Both suggest that efficiency declines with less fragmentation.

At higher cutoffs, the effect on price impact continues to increase (1.31 and 1.35), although the coefficients become less statistically significant due to the small number of treated coins in these high-intensity groups.

To test the marginal impact of the shock while controlling for differences across treatment intensities, we examine the results from our continuous event study in Column 3 of Table 6. The coefficient on Price Impact is 3.34 and is significant at the 1% level. This suggests that the adverse effects on price impact scale with the intensity of the fragmentation shock. Using the average pre-collapse FTX volume share of 0.12 and a pre-period mean of -0.63, this implies a 64% increase in price impact ($\frac{0.12\times 3.34}{-0.63} \approx 0.64$). Furthermore, Column 4 shows that this increase in price impact is persistent. The coefficients for the early post-period and the final week are identical, indicating that liquidity conditions did not recover within the month following the collapse.

[**Table 9** here]

We test the order-splitting channel in the Binance delisting sample by examining price impact in Table 10. Columns 1 and 3 show that without controls, the price impact principal component increases significantly. However, when we add controls in Columns 2 and 4, the magnitude and significance of the price impact coefficient decrease. Although the price impact is still positive, it is not significant at conventional levels. The weaker results in the Binance sample likely reflect the reduced statistical power of the smaller sample with 34 treated coins and the fact that controlling for volume absorbs much of the liquidity shock that drives price impact.

[**Table 10** here]

To ensure our findings are not solely driven by the price impact principal component (PC), we replicate the specific tests from each identification strategy using the individual measures of efficiency and price impact. We examine the FTX shock by repeating the SDID specification for price impact at various horizons in Internet Appendix Table A.6. Turning to the Binance delistings, we repeat the difference-in-differences tests for the individual price impact measures in Internet Appendix Table A.7.

Our results are largely consistent with Chen and Duffie (2021). After delistings and venue collapses, price impact increases and price efficiency declines. The strong and statistically significant decline in price efficiency (417%) demonstrates that reducing fragmentation harms efficiency even in non-crisis periods. Taken together, the evidence suggests that when traders lose venues for order-splitting, their ability to execute trades without moving the price deteriorates, leading to less informative prices.

6. Conclusion

Financial markets in the US and Europe have grown increasingly fragmented over the past 25 years. The NYSE and Nasdaq handled almost all trade volume prior to the year 2000 but are now reduced to executing closer to 35% of volume combined. The same is true in Europe where the listing exchange only handles 40% of volume. Despite this widespread trend, we do not have clear theoretical guidance or empirical support for how fragmentation affects price efficiency. Chen and Duffie (2021) provide a new theoretical framework that links fragmentation and price efficiency through the order-splitting channel.

We test the theoretical predictions by Chen and Duffie (2021) and overcome the empirical challenges faced by previous work by using high-frequency data from 1,768 coin pairs across 16 cryptocurrency trading venues. The cryptocurrency setting provides unique advantages for studying fragmentation. Unlike traditional exchanges, cryptocurrency venues operate independently without best-execution rules or centralized routing. This allows us to observe fragmentation effects without regulatory interference. Further, we exploit two natural experiments in this market: the sudden collapse of the venue FTX and coin delistings on Binance. These events allow us to obtain causal inference of fragmentation's effect on price efficiency.

We find that fragmentation improves price efficiency and the effects are also economically large. For the FTX collapse, coins with over 20% of their volume on FTX see their price efficiency drop by 125% of the pre-period mean. This effect grows to 162% for coins with at least 40% of volume on FTX. Furthermore, when coins are delisted from Binance and lose access to the world's largest trading venue, their price efficiency falls by a remarkable 417% relative to the pre-period mean. Our results are consistent with the diminishing marginal effect of fragmentation as described by Chen and Duffie (2021). We find evidence that this result is driven by the order-splitting channel. When traders lose a venue, their ability to split orders is impaired, which increases their trading costs. Specifically, the FTX collapse leads to a 32% increase in price impact for the average coin, growing to 162% for the 20% volume group.

Our findings matter for market design and regulation. Regulators often view market fragmentation as harmful and seek to consolidate trading. Our evidence suggests the opposite: fragmentation can improve efficiency by giving traders more tools to manage their orders and reduce price impact. Although we do not estimate a formal model, these gains in efficiency strongly imply that fragmentation is welfare-improving for market participants. By ensuring prices reflect their true value, fragmentation allows participants to allocate capital effectively, trade at lower costs, and operate with greater confidence. This challenges the conventional wisdom that centralized markets are always better. As financial markets continue to evolve, policymakers should consider the benefits of fragmentation alongside its costs. The ability of traders to split orders across venues appears to be a key channel through which fragmentation improves the efficiency of prices.

We acknowledge, however, that the unique characteristics of our empirical setting also introduce limitations. While the lack of regulation in cryptocurrency markets provides an environment to test economic theory, our findings may not be fully generalizable to the highly regulated and institutionally complex US and European equity markets. The mechanisms that connect fragmented venues in traditional markets, such as smart order routers and inter-market sweep orders, are absent in our setting. Furthermore, the investor base and the nature of information in cryptocurrency markets may differ from those in traditional finance. Therefore, while our paper provides clear causal evidence of a fundamental economic channel, further research is needed to understand how this channel operates in the presence of the specific regulations and institutional features that govern modern equity markets.

References

Aitken, M., H. Chen, and S. Foley, 2017, “The Impact of Fragmentation, Exchange Fees and Liquidity Provision on Market Quality,” *Journal of Empirical Finance*, 41, 140–160.

Allen, J., and M. Wittwer, 2023, “Centralizing Over-the-Counter Markets,” *Journal of Political Economy*, 131(12), 3310–3351.

Aquilina, M., E. Budish, and P. O’Neill, 2022, “Quantifying the High-Frequency Trading “Arms Race”,” *Quarterly Journal of Economics*, 137(1), 493–564.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager, 2021, “Synthetic Difference-in-Differences,” *American Economic Review*, 111(12), 4088–4118.

Babus, A., and C. Parlatore, 2022, “Strategic Fragmented Markets,” *Journal of Financial Economics*, 145(3), 876–908.

Baldauf, M., and J. Mollner, 2021, “Trading in Fragmented Markets,” *Journal of Financial and Quantitative Analysis*, 56(1), 93–121.

Battalio, R., S. A. Corwin, and R. Jennings, 2016, “Can Brokers Have it All? On the Relation Between Make-take Fees and Limit Order Execution Quality,” *Journal of Finance*, 71(5), 2193–2238.

Bennett, P., and L. Wei, 2006, “Market Structure, Fragmentation, and Market Quality,” *Journal of Financial Markets*, 9(1), 49–78.

Bessembinder, H., J. Hao, and K. Zheng, 2020, “Liquidity Provision Contracts and Market Quality: Evidence from the New York Stock Exchange,” *Review of Financial Studies*, 33(1), 44–74.

Brolley, M., and D. A. Cimon, 2020, “Order-Flow Segmentation, Liquidity, and Price Discovery:

The Role of Latency Delays," *Journal of Financial and Quantitative Analysis*, 55(8), 2555–2587.

Budish, E., R. S. Lee, and J. J. Shim, 2019, "Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation," Working Paper.

Cespa, G., and X. Vives, 2022, "Exchange Competition, Entry, and Welfare," *Review of Financial Studies*, 35(5), 2570–2624.

Chao, Y., C. Yao, and M. Ye, 2019, "Why Discrete Price Fragments US Stock Exchanges and Disperses Their Fee Structures," *Review of Financial Studies*, 32(3), 1068–1101.

Chen, D., and D. Duffie, 2021, "Market Fragmentation," *American Economic Review*, 111(7), 2247–2274.

Chowdhry, B., and V. Nanda, 1991, "Multimarket Trading and Market Liquidity," *Review of Financial Studies*, 4(3), 483–511.

Chung, K. H., and C. Chuwonganant, 2022, "Regulation NMS and Market Quality," *Financial Management*, 41(2), 285–317.

Colliard, J.-E., and T. Foucault, 2012, "Trading Fees and Efficiency in Limit Order Markets," *Review of Financial Studies*, 25(11), 3389–3421.

Comerton-Forde, C., V. Grégoire, and Z. Zhong, 2019, "Inverted Fee Structures, Tick Size, and Market Quality," *Journal of Financial Economics*, 134(1), 141–164.

Comerton-Forde, C., and T. J. Putniņš, 2015, "Dark Trading and Price Discovery," *Journal of Financial Economics*, 118(1), 70–92.

Degryse, H., F. De Jong, and V. v. Kervel, 2015, "The Impact of Dark Trading and Visible Fragmentation on Market Quality," *Review of Finance*, 19(4), 1587–1622.

Dzieliński, M., B. Hagströmer, and C. Qu, 2025, "Exchange Competition, Fragmentation, and Market Quality," Working Paper.

European Securities and Markets Authority (ESMA), 2023, "Evolution of EEA Share Market Structure Since MiFID II," *ESMA TRV Risk Analysis*.

Fischer, T. G., C. Krauss, and A. Deinert, 2019, "Statistical Arbitrage in Cryptocurrency Markets," *Journal of Risk and Financial Management*, 12(1), 31.

Foley, S., and T. J. Putniņš, 2016, "Should We Be Afraid of the Dark? Dark Trading and Market Quality," *Journal of Financial Economics*, 122(3), 456–481.

Foucault, T., O. Kadan, and E. Kandel, 2013, "Liquidity Cycles and Make/Take Fees in Electronic Markets," *Journal of Finance*, 68(1), 299–341.

Foucault, T., and A. J. Menkveld, 2008, "Competition for Order Flow and Smart Order Routing Systems," *Journal of Finance*, 63(1), 119–158.

Gomber, P., S. Sagade, E. Theissen, M. C. Weber, and C. Westheide, 2016, "Spoilt for Choice: Order Routing Decisions in Fragmented Equity Markets," *CFR Working Paper*.

Gresse, C., 2017, "Effects of Lit and Dark Market Fragmentation on Liquidity," *Journal of Financial Markets*, 35, 1–20.

Haslag, P., and M. C. Ringgenberg, 2023, "The Demise of the NYSE and NASDAQ: Market Quality in the Age of Market Fragmentation," *Journal of Financial and Quantitative Analysis*, 58(7), 2753–2782.

Hatheway, F., A. Kwan, and H. Zheng, 2017, "An Empirical Analysis of Market Segmentation on US Equity Markets," *Journal of Financial and Quantitative Analysis*, 52(6), 2399–2427.

Hendershott, T., and C. M. Jones, 2005, “Island Goes Dark: Transparency, Fragmentation, and Regulation,” *Review of Financial Studies*, 18(3), 743–793.

Hollifield, B., P. Sandås, and A. Todd, 2017, “Latency Arbitrage When Markets Become Faster,” *Working Paper*.

Liu, X., W. Chen, and K. Zhu, 2023, “Market Fragmentation in Cryptocurrency Exchange,” Working Paper.

Liu, Y., A. Tsyvinski, and X. Wu, 2022, “Common Risk Factors in Cryptocurrency,” *Journal of Finance*, 77(2), 1133–1177.

Lo, A. W., and A. C. MacKinlay, 1988, “Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test,” *Review of Financial Studies*, 1(1), 41–66.

Madhavan, A., 2000, “Market Microstructure: A Survey,” *Journal of Financial Markets*, 3(3), 205–258.

Makarov, I., and A. Schoar, 2020, “Trading and Arbitrage in Cryptocurrency Markets,” *Journal of Financial Economics*, 135(2), 293–319.

Malinova, K., and A. Park, 2015, “Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality,” *Journal of Finance*, 70(2), 509–536.

——— , 2020, “’Sniping’ in Fragmented Markets,” *Working Paper*.

Menkveld, A. J., 2008, “Splitting Orders in Overlapping Markets: A Study of Cross-Listed Stocks,” *Journal of Financial Intermediation*, 17(2), 145–174.

Menkveld, A. J., and M. A. Zoican, 2017, “Need for Speed? Exchange Latency and Liquidity,” *Review of Financial Studies*, 30(4), 1188–1228.

O'Hara, M., and M. Ye, 2011, "Is Market Fragmentation Harming Market Quality," *Journal of Financial Economics*, 100(3), 459–474.

Pagano, M., 1989, "Trading Volume and Asset Liquidity," *Quarterly Journal of Economics*, 104(2), 255–274.

Pagnotta, E. S., and T. Philippon, 2018, "Competing on Speed," *Econometrica*, 86(3), 1067–1115.

Petrella, G., 2009, "MiFID, Reg NMS and Market Quality: Lessons for the Future," *Journal of Financial Regulation and Compliance*, 17(3), 230–245.

van Kervel, V., A. Kwan, and P. J. Westerholm, 2023, "Order Splitting and Interacting with a Counterparty," *Journal of Financial Markets*, 66, 100850.

Zargar, F. N., and D. Kumar, 2019, "Informational Inefficiency of Bitcoin: A Study Based on High-Frequency Data," *Research in International Business and Finance*, 47, 344–353.

Table 1
Venue Market Share and Trading Statistics

This table presents the daily median and mean trading *Volume* in USD by cryptocurrency trading venue, alongside the number of listed *Coins*, the median *Number of Trades* per day, the median *Trade Size* in USD, and the trading venues' *Market Share*. The *Market Share* is calculated as the trading venue's median daily volume relative to the total median volume across all trading venues.

	Median Volume [1]	Mean Volume [2]	Coins [3]	Number of Trades [4]	Trade Size [5]	Market Share (%) [6]
Binance	14,904,648,755	17,368,709,014	839	24,446,453	610	65.71
Coinbase	2,376,713,923	2,981,112,092	468	3,447,094	689	10.48
FTX	1,514,433,454	1,646,886,025	175	1,120,991	1,351	6.68
Gate-io	779,719,089	865,689,668	1071	3,009,900	259	3.44
Kraken	776,043,323	936,348,202	417	482,924	1,607	3.42
Kucoin	667,610,514	739,812,649	883	5,668,588	118	2.94
Binance US	431,111,758	563,448,850	330	612,033	704	1.90
Bybit-spot	328,511,275	435,078,666	285	661,491	497	1.45
Crypto.com	312,570,675	343,720,526	536	4,621,070	68	1.38
Bitstamp	200,537,072	293,483,067	168	90,184	2,224	0.88
FTX US	157,383,104	176,145,432	49	56,149	2,803	0.69
Gemini	111,053,701	142,391,490	124	149,522	743	0.49
Ascendex	91,670,034	110,344,617	319	329,231	278	0.40
Upbit	21,757,634	40,580,877	139	38,469	566	0.10
Bitflyer	7,044,453	10,307,355	6	3,621	1,945	0.03
Coinflex	339,334	610,429	13	865	392	0.00

Table 2
Fragmentation and Venue Count by Coin Tercile

This table presents the average, median, standard deviation, and maximum statistics of *Fragmentation* and the number of *Venues* for coin terciles. *Fragmentation* is 1 - Herfindahl–Hirschman Index. We group the coins into three terciles based on logged average trading volume (USD) for the coin across our sample. Panel A shows statistics for the whole sample. Panels B–D show statistics by volume tercile.

	Mean [1]	Median [2]	St.Dev. [3]	Max [4]
Panel A: Full Sample				
Fragmentation	0.17	0.08	0.20	0.82
Venues	2.39	2.00	1.72	15.00
Panel B: Highest Tercile				
Fragmentation	0.20	0.13	0.21	0.82
Venues	3.34	3.00	2.18	15.00
Panel C: Middle Tercile				
Fragmentation	0.16	0.05	0.20	0.79
Venues	1.91	2.00	0.95	7.00
Panel D: Lowest Tercile				
Fragmentation	0.15	0.00	0.19	0.75
Venues	1.55	1.00	0.63	5.00

Table 3
Correlation Matrix

This table reports pairwise correlations among the main market variables. *Auto Corr* is defined previously, *Var Ratio PC*, *Efficiency PC*, and *Price Impact PC* are first principal components of our variance ratios, efficiency measures (autocorrelation and variance ratios), and price impact measures. *Fragmentation* is as previously defined. *Venues* is the number of venues a coin is traded on, *Volume* is the natural log of the volume in USD, and *Volatility* is measured as the difference between the maximum and minimum midquote, divided by the average midquote.

	Auto Corr [1]	Var Ratio PC [2]	Efficiency PC [3]	Price Impact PC [4]	Fragmentation [5]	Venues [6]	Volume [7]	Volatility [8]
Auto Corr	1.00							
Var Ratio PC	0.63	1.00						
Efficiency PC	0.80	0.97	1.00					
Price Impact PC	-0.13	-0.23	-0.22	1.00				
Fragmentation	0.03	0.06	0.06	-0.11	1.00			
Venues	0.19	0.28	0.27	-0.23	0.58	1.00		
Volume	0.40	0.58	0.57	-0.26	0.13	0.46	1.00	
Volatility	0.04	0.06	0.06	0.37	-0.07	-0.09	0.20	1.00

Table 4
Comparing Treated and Synthetic Control Groups

The table shows the average *Price Impact*, *Price Efficiency*, *Number of Venues*, *Fragmentation*, logged USD *Volume*, and *Volatility* for each coin traded on FTX (Treated) and the synthetic control group made up of coins not traded on FTX. The coins in the synthetic control group are assigned weights to best match the pre-period trajectory of the treated coins so that any change in the price impact or price efficiency is not due to differences or changes in the *Number of Venues*, *Fragmentation*, *Log(Volume)*, and *Volatility*. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Treated			Synthetic Control		
	Pre [1]	Post [2]	Difference [3]	Pre [4]	Post [5]	Difference [6]
Venues	5.68	4.87	-0.81	2.74	2.74	0.00
Fragmentation	0.39	0.27	-0.12	0.26	0.25	-0.01
Volume (USD)	14.47	13.87	-0.59	12.55	12.41	-0.15
Volatility	0.04	0.07	0.03	0.06	0.07	0.02
Price Impact	-0.63	-0.01	0.62	-0.19	0.10	0.29
Price Efficiency	0.87	0.67	-0.20	0.23	0.09	-0.15

Table 5
Price Efficiency and the FTX Collapse

This table shows the estimation of the equation:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}$$

where Outcome_{it} is the specific measure of price efficiency for coin i on date t. Post_t equals 1 on November 08, 2022 onward and 0 in the 30 days prior. Treat_i equals 1 if coin i was traded on FTX and 0 otherwise. The control group is constructed using SDID by choosing unit and time weights so that the weighted control group matches the treated coins' pre-period outcomes (and covariates) in logged USD volume, volatility, number of venues, and fragmentation, following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). We estimate the DiD coefficient on $\text{Post}_t \times \text{Treat}_i$ using these SDID weights. X_{it} includes logged USD volume, number of trading venues, fragmentation, and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. We estimate the model using all coins and then segment the treated sample based on the pre-collapse FTX volume share being at least 10%, 20%, 30%, or 40% during the month prior to the collapse. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Full Sample [1]	10% [2]	20% [3]	30% [4]	40% [5]
Treat × Post	-0.13*	-0.70***	-1.09***	-1.19***	-1.42***
	(0.071)	(0.14)	(0.23)	(0.24)	(0.28)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	92,580	86,040	85,140	84,900	84,660

Table 6
FTX Collapse: Continuous Treatment and Temporal Effects

This table examines the effect of the FTX collapse on Price Efficiency and Price Impact using a continuous treatment intensity (Columns 1 and 3) and tests for effect attenuation over time (Columns 2 and 4). We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Pre-Volume Share FTX}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}$$

In Columns 2 and 4, we split Post_t into EarlyPost_t (days 0–22) and FinalWeek_t (days 23–29). Pre-Volume Share FTX_i is the average share of trading volume on FTX for coin i in the month before the collapse. X_{it} includes control variables: *volatility*, *number of trading venues*, and *logged trading volume*. α_i represents coin fixed effects, and Θ_t represents date fixed effects. Standard errors are clustered by coin and date. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Price Efficiency		Price Impact	
	[1]	[2]	[3]	[4]
Post × Pre-Volume Share FTX	-1.48***		3.34***	
	(0.40)		(0.93)	
Early Post (Days 0–22) × Pre-Volume Share FTX		-1.62***		3.34***
		(0.49)		(0.90)
Final Week (Days 23–29) × Pre-Volume Share FTX		-1.57***		3.34***
		(0.55)		(1.17)
Log(Volume)	0.16***	0.18***	0.10*	0.10*
	(0.037)	(0.046)	(0.056)	(0.056)
Volatility	0.29	0.94	8.21***	8.21***
	(0.76)	(0.97)	(1.21)	(1.21)
Venues	-0.12***	-0.14***	-0.013	-0.013
	(0.023)	(0.027)	(0.036)	(0.036)
Coin FE	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes
Observations	9,445	9,445	9,445	9,445
R^2	0.669	0.662	0.572	0.572

Table 7
Delisted and Control Coins Around Delisting Event

This table reports average *Number of Venues*, *Fragmentation*, logged USD *Volume*, *Volatility*, *Price* in USD, *Price Impact*, and *Price Efficiency* for the delisted coins and their matched control group. The control group is defined using propensity score matching on *Number of Venues*, *Fragmentation*, logged USD *Volume*, *Price*, and *Volatility*. We report averages for the pre and post-periods—30 days before and after the delisting announcement—and their differences. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Delisted Coins			Controls		
	Pre [1]	Post [2]	Difference [3]	Pre [4]	Post [5]	Difference [6]
Number of Venues	2.80	1.74	-1.06^{***}	2.78	2.85	0.06
Fragmentation	0.26	0.13	-0.12^{***}	0.24	0.25	0.01
Volume	12.28	10.32	-1.95^{***}	12.37	12.16	-0.21^{*}
Volatility	0.08	0.10	0.02^{**}	0.09	0.08	−0.01
Price	2102.46	1142.56	−959.91	80.81	89.09	8.28
Price Impact	-0.22	0.58	0.80^{**}	-0.12	-0.18	−0.06
Price Efficiency	0.15	-0.83	-0.98^{***}	0.23	0.20	−0.04

Table 8
Delisting on Price Efficiency

This table presents the estimation of the regression:

$$\text{Efficiency}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \Theta_t + \alpha_i + \varepsilon_{it}$$

where Efficiency_{it} is the price efficiency principal component. The event days range from 30 days before and after the delisting announcement, with the 7 days immediately after the delisting removed. Controls include *number of venues*, *log(Volume)*, *volatility*, and *logged price*. Standard errors are clustered by coin and event day. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Price Efficiency			
	[1]	[2]	[3]	[4]
Post × Treat	-0.94***	-0.64***	-0.95***	-0.69***
	(0.19)	(0.20)	(0.16)	(0.18)
Venues		-0.15		-0.13
		(0.12)		(0.13)
Log(Volume)		0.19***		0.19***
		(0.055)		(0.055)
Volatility		1.89**		1.89**
		(0.75)		(0.75)
Log(Price)		0.046		0.046
		(0.13)		(0.13)
Coin FE	Yes	Yes	Yes	Yes
Event FE	Yes	Yes	No	No
Date FE	No	No	Yes	Yes
Observations	3,835	3,848	3,810	3,823
R^2	0.581	0.585	0.633	0.619

Table 9
Mechanisms of the FTX Shock: Price Impact

This table reports estimates from:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}$$

where Outcome_{it} is the specific measure of price impact for coin i on date t. Post_t equals 1 on November 08, 2022 onward and 0 in the 30 days prior. Treat_i equals 1 if coin i was traded on FTX and 0 otherwise. The control group is constructed using SDID by choosing unit and time weights so that the weighted control group matches the treated coins' pre-period outcomes (and covariates) in logged USD volume, volatility, number of venues, and fragmentation, following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). We estimate the DiD coefficient on $\text{Post}_t \times \text{Treat}_i$ using these SDID weights. X_{it} includes logged USD volume, number of trading venues, fragmentation, and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. We estimate the model using all coins and then segment the treated sample based on the pre-collapse FTX volume share being at least 10%, 20%, 30%, or 40% during the month prior to the collapse. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Full Sample	10%	20%	30%	40%
	[1]	[2]	[3]	[4]	[5]
Treat × Post	0.20**	0.79**	1.02**	1.31*	1.35
	(0.090)	(0.33)	(0.50)	(0.67)	(0.88)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	92,580	86,040	85,140	84,900	84,660

Table 10
Mechanisms of Delisting: Price Impact

This table tests the order-splitting channel by estimating the effect of delisting on price impact. We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \Theta_t + \alpha_i + \varepsilon_{it}$$

where Outcome_{it} is the price impact principal component. The event days range from 30 days before and after the delisting announcement, with the 7 days immediately after the delisting removed. Controls include *number of venues*, *log(Volume)*, *volatility*, and *logged price*. α_i denotes coin fixed effects and Θ_t denotes date fixed effects. Standard errors are clustered by coin and event day. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Price Impact			
	[1]	[2]	[3]	[4]
Post × Treat	0.99***	0.18	0.99***	0.31
	(0.33)	(0.28)	(0.28)	(0.25)
Venues		-0.13		-0.13
		(0.17)		(0.14)
Log(Volume)		-0.24***		-0.21***
		(0.088)		(0.070)
Volatility		9.80***		10.1***
		(1.13)		(1.19)
Log(Price)		0.13		0.23
		(0.22)		(0.22)
Coin FE	Yes	Yes	Yes	Yes
Event FE	Yes	Yes	No	No
Date FE	No	No	Yes	Yes
Observations	3,848	3,848	3,823	3,823
R^2	0.476	0.563	0.543	0.616

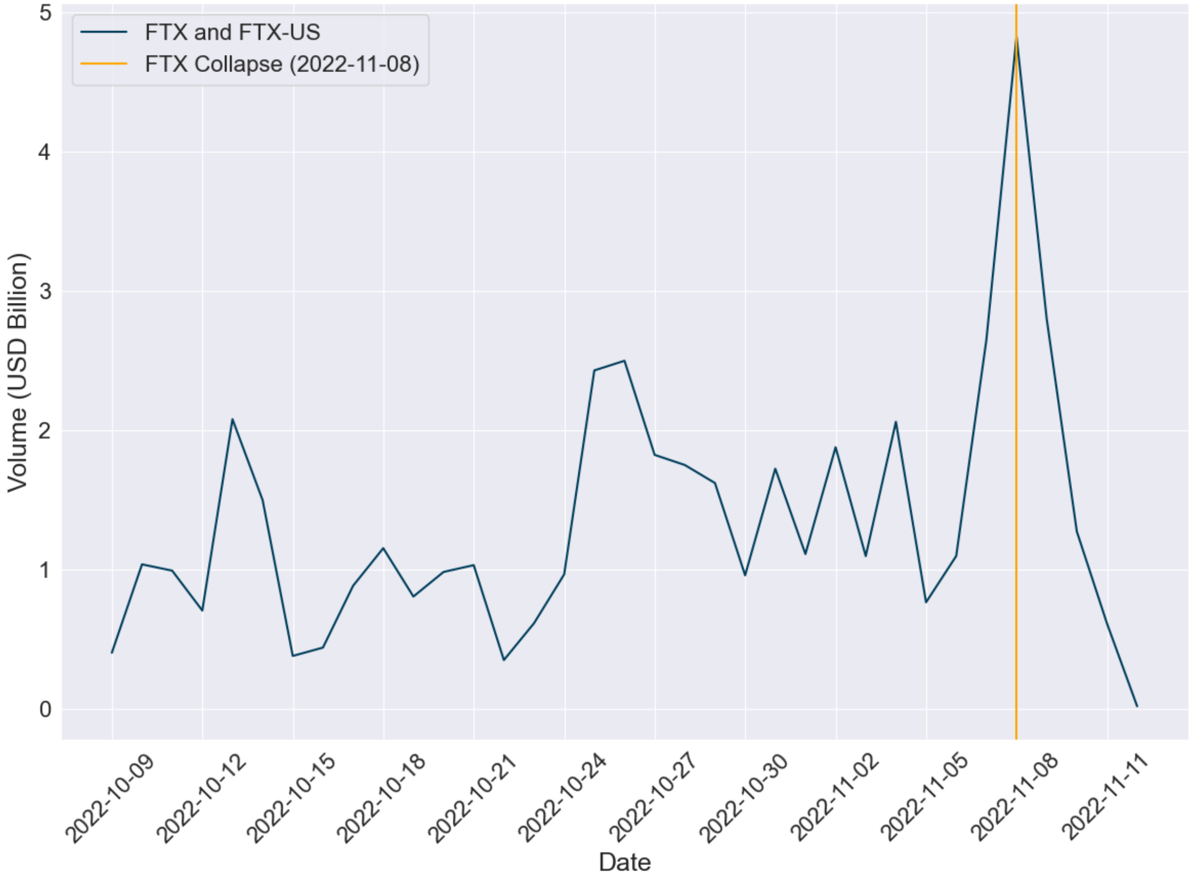

Figure 1. Trading Volume on FTX

This figure shows the total daily trading volume (in USD) for coins traded on FTX and FTX-US during the period around the FTX collapse on November 08, 2022. The vertical line indicates the collapse date.

Figure 2. Fragmentation on FTX and Other Venues

This figure compares the two-day moving average market fragmentation for coins traded on FTX and FTX-US with those not traded on FTX around the FTX collapse on November 08, 2022. The y-axis displays fragmentation, calculated as 1 minus the Herfindahl-Hirschman Index (1-HHI). The vertical line indicates the collapse date.

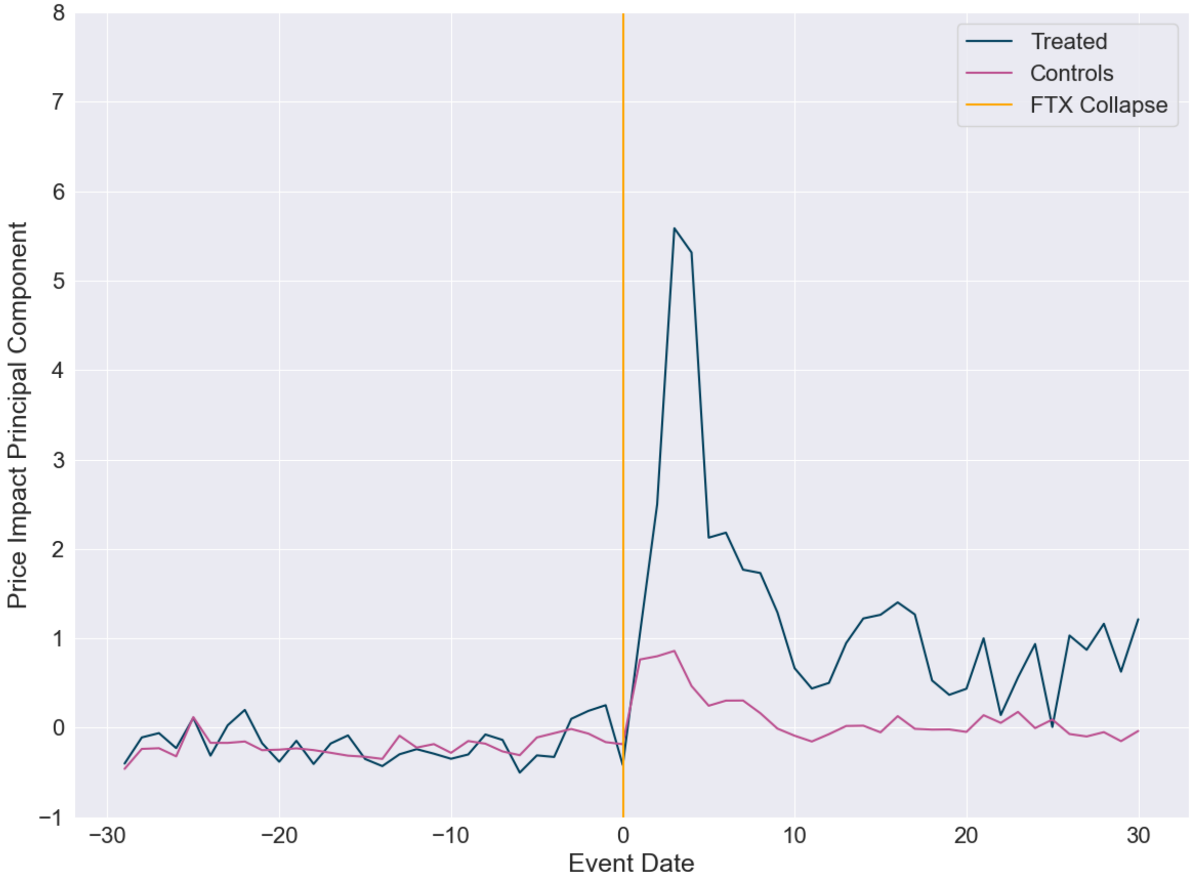

Figure 3. Price Impact Trends
This figure shows trends from a synthetic difference-in-differences analysis comparing price impact for coins in the treated group, defined as those with at least 10% of trading volume on FTX, and coins in the control group, which were not traded on FTX. The controls are matched based on the pre-period trends of number of venues, fragmentation, volume, and volatility. The x-axis shows days relative to the FTX collapse on November 08, 2022, with day 0 marking the event.

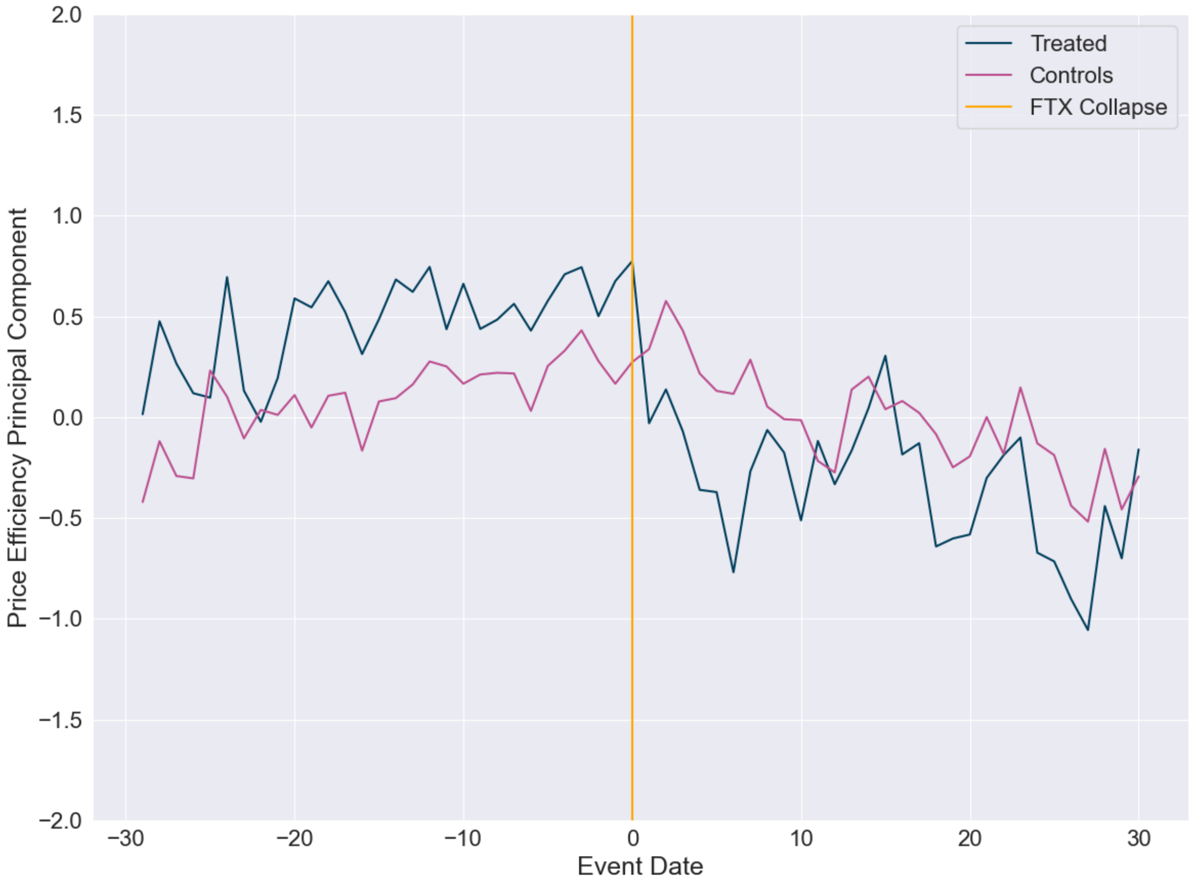

Figure 4. Price Efficiency Trends
This figure shows trends from a synthetic difference-in-differences analysis comparing price efficiency for coins in the treated group, defined as those with at least 10% of trading volume on FTX, and coins in the control group, which were not traded on FTX. The controls are matched based on the pre-period trends of number of venues, fragmentation, volume, and volatility. The x-axis shows days relative to the FTX collapse on November 08, 2022, with day 0 marking the event.

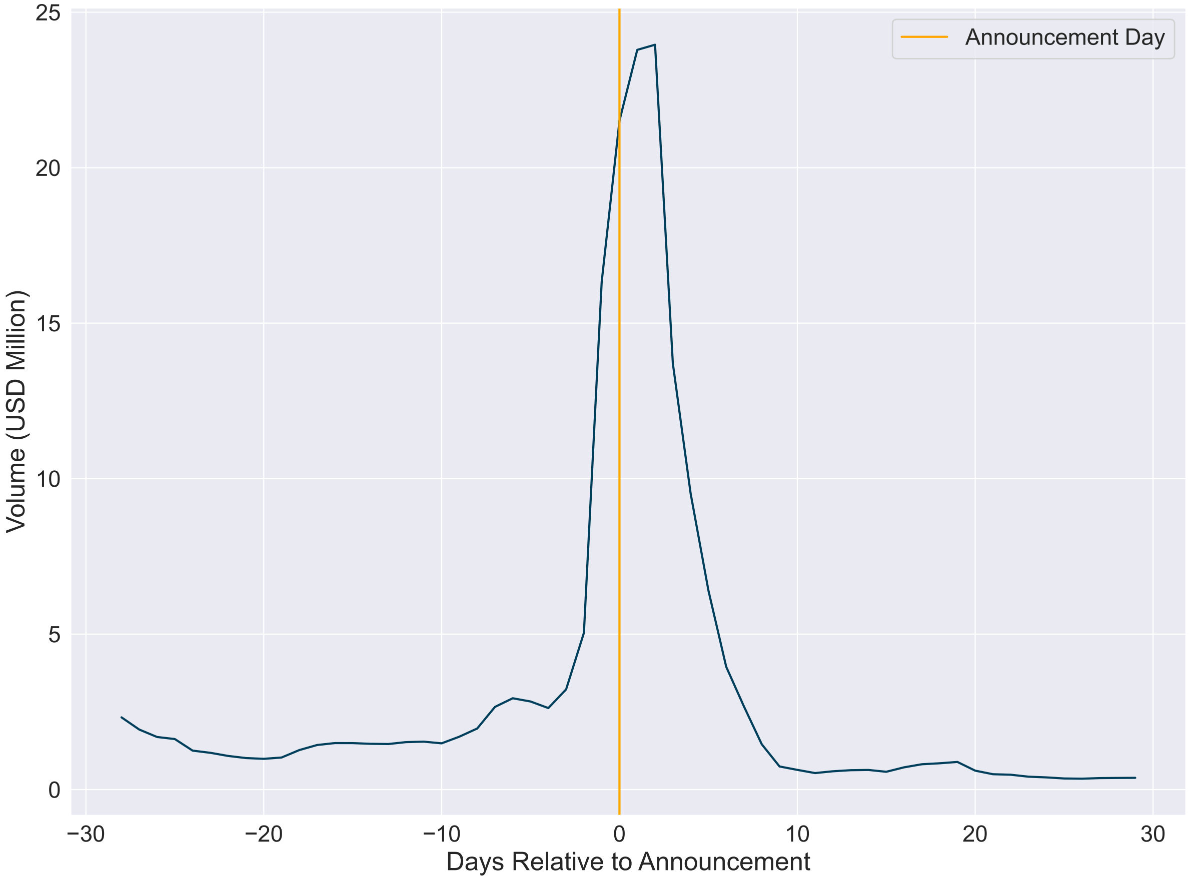

Figure 5. Volumes Surrounding Delisting Events

This figure shows the 4-day moving average trading volume (in USD) across trading venues for coins that were traded on Binance, delisted from Binance, and continued trading on at least one other trading venue. The y-axis shows the average daily trading volume in million USD. The x-axis shows the days relative to the delisting event, with day 0 marking the delisting date. The figure covers a 30-day window, spanning 30 days before and 30 days after the delisting.

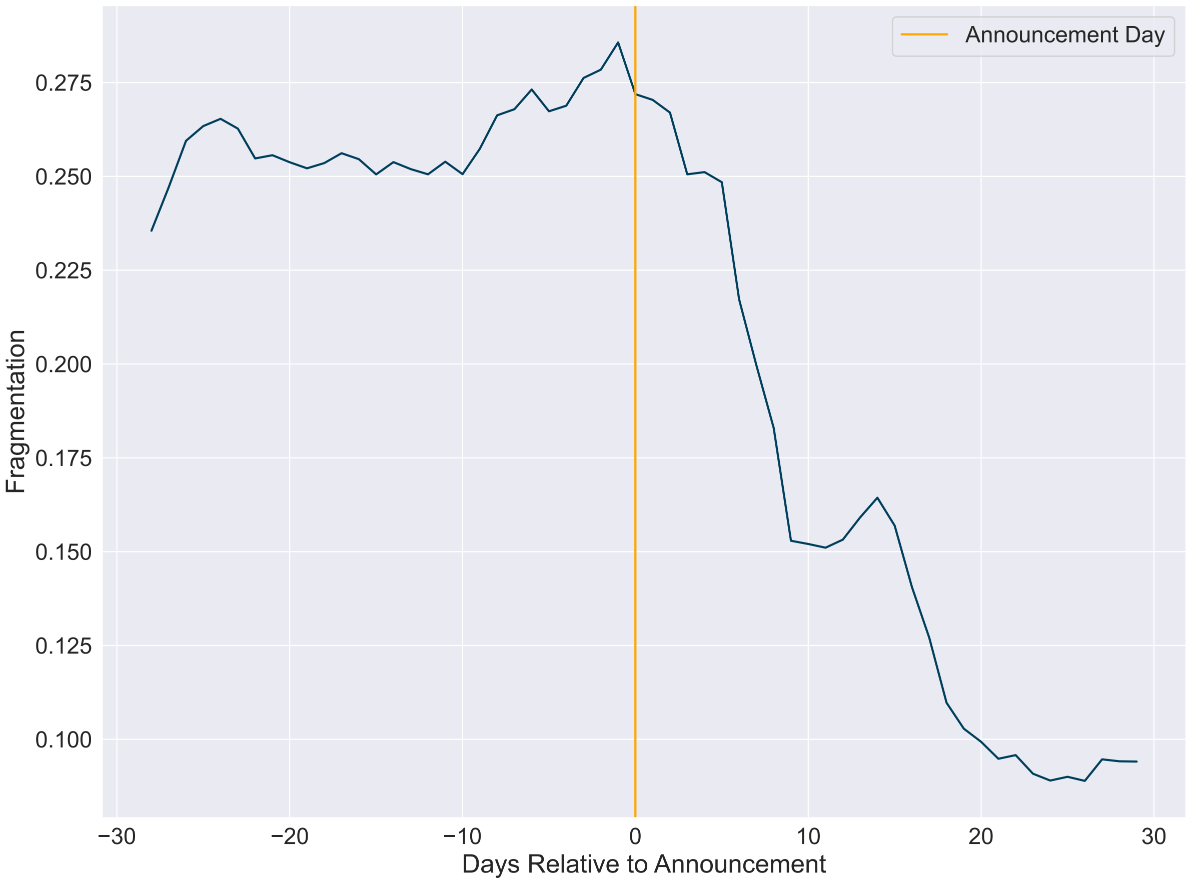

Figure 6. Fragmentation Surrounding Delisting Events

This figure shows the 4-day moving average level of fragmentation for coins that were traded on Binance, were delisted from Binance, and continued trading on at least one other trading venue. Fragmentation is 1-Herfindahl-Hirschman Index. The x-axis shows the days relative to the delisting event, with day 0 marking the day the coin was removed from Binance. The figure covers 30 days before and 30 days after the delisting.

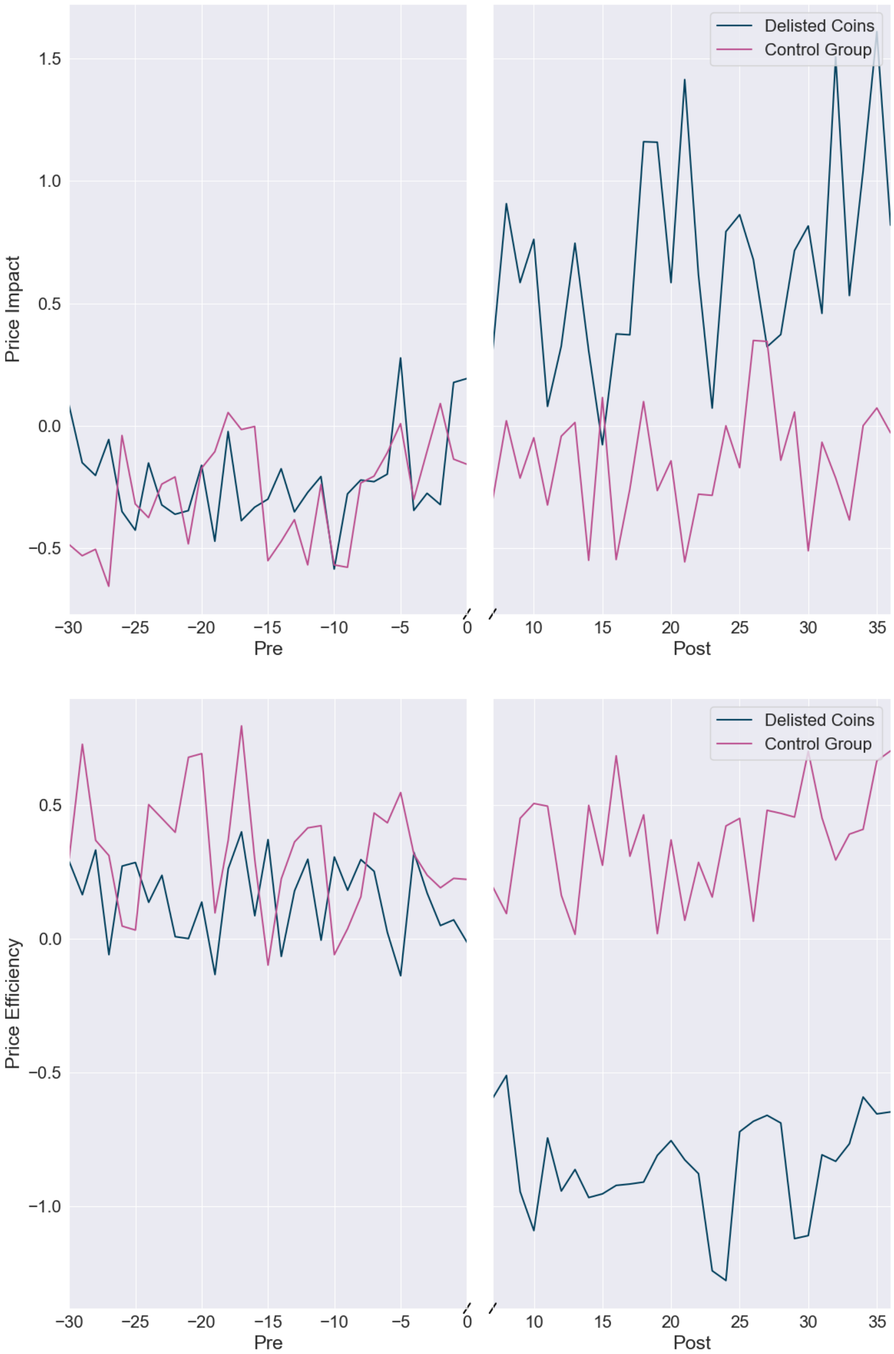

Figure 7. Price Impact and Efficiency for Delisted and Control Coins
This figure shows the average level of the principal component price impact (top panel) and price efficiency (bottom panel) for coins that were delisted from Binance and the control group. The control group is identified using the propensity scores on pre-period number of venues, fragmentation, logged volume, USD price, and volatility. The x-axis shows 30 days before and after the delisting announcement, with a gap 7 days after the announcement.

Internet Appendix

A. Data Processing Detail

We pre-process the data from Tardis.dev in several steps. We first drop observations where the bid or ask price is zero, trade volume is negative, the bid price exceeds the ask price, the bid-ask spread exceeds 5%, or returns between consecutive 5-millisecond timestamps exceed 5%. At the daily level, we include only trading venue–coin–day combinations with at least 10 trade observations.

We resample the data to 5-millisecond intervals to capture high-frequency movements while reducing reporting errors. For quotes, we retain the last quote within each 5-millisecond interval. For trades, we aggregate all trades within each interval by summing trade amounts in the same direction (buy or sell) and computing a volume-weighted average price.

After resampling, we remove extreme outliers and data inconsistencies. We correct coin pairs reported in reverse order (e.g., changing *USDT-BTC* to *BTC-USDT*). We exclude coins that experience price changes exceeding 10,000% that remain elevated for at least seven days, and coin-days with prices surging over 1,000% in a single day (e.g., Terra (LUNA) due to its extreme volatility). We remove observations where a coin's price on a trading venue deviates by more than 100% from the median price across other trading venues on the same day; if a coin is flagged as an outlier for more than 30 days on any trading venues, we remove it entirely from the dataset. We also remove coin-day observations where the average midpoint-price exceeds $75,000; for non-Bitcoin (BTC) coins, we exclude prices above $60,000. After applying these filters, we winsorize the efficiency measures and volatility at the 1% and 99% levels monthly to mitigate the impact of remaining outliers.

We then standardize coin naming conventions across trading venues and exclude remaining futures contracts. We rename 'XBT' to 'BTC', reorder coin pairs for uniformity (e.g., changing

PAXETH to ETH-PAX, and remove symbols such as underscores and hyphens from coin names. Since Tardis.dev data does not classify security types into futures and spot trading pairs, we remove futures contracts using specific criteria. We exclude coins with names containing terms like “futures,” “PERP,” or “perpetual”; symbols ending with digits or containing date patterns; and leveraged coins and tokens. We also manually verify each coin’s security type on each trading venue’s website to ensure accuracy. After identifying the set of spot currencies, we convert all trade volumes into U.S. dollars using daily average trading venues rates.

We calculate a variety of metrics to measure the effect of fragmentation on price efficiency. We also calculate variables to control for the number of venues and volatility.

B. Venue-Level Shocks to Fragmentation

This section presents supplementary analysis on the venue-level shock discussed in Section 4.1, specifically the collapse of FTX. We include robustness checks for the event window and parallel trends visualizations.

B.1 Alternative Event Window Specification

To ensure our results are not driven by the immediate volatility of the bank run period preceding the bankruptcy filing of FTX, we replicate our analysis using an alternative event window. In this specification, we exclude the days from November 2, 2022, to November 10, 2022. We compare the 30-day period immediately following this blocked window to the 30-day period preceding it.

Table A.1 reports the results for price efficiency using both the cutoff SDID. Table A.2 reports the corresponding results for the price impact mechanism. The results demonstrate that the deterioration in price efficiency and the increase in price impact remain statistically significant and similar in magnitude to the baseline specification.

B.2 Parallel Trends by Exposure Level

To visualize the divergence between treated coins and synthetic controls across different treatment intensities, we provide Figure A.1 and Figure A.2. Figure A.1 plots the trends for the Price Impact Principal Component for coins with at least 20%, 30%, and 40% of their pre-collapse volume on FTX. Similarly, Figure A.2 plots the trends for the Efficiency Principal Component for the same volume thresholds. Prior to the collapse, trends are largely parallel. Additionally, the divergence in price impacts and efficiency between treated coins and control coins is monotonic with respect to the coin's reliance on FTX. The effects are most pronounced for coins that lost a venue representing a larger share of their trading volume.

C. Asset-Level Shocks to Fragmentation

This section presents robustness checks and visual diagnostics related to the asset-level shocks discussed in Section 4.2, specifically the Binance delisting events.

C.1 Delisting Sample Characteristics

Table A.3 lists all coins that were delisted from Binance or Binance US during our sample period, as confirmed on Binance's official website. As with the FTX analysis, we only consider coins that are traded on at least one other venue, resulting in 34 events.

These delistings span different market conditions and coins are delisted for various reasons. During the 13 months from March 2022 to March 2023, 34 trading pairs and 24 unique coins are delisted. The distribution shows clustering around the FTX collapse in November 2022 when 9 pairs delisted, with 7 occurring on November 28. Other clusters of delisting occur in October 2022 (5 pairs) and January 2023 (5 pairs). This clustering reflects Binance's tendency to conduct batch reviews. The remaining delistings are more sporadic.

The trading pairs span four base currencies: Bitcoin (BTC) with 14 pairs, Tether (USDT) with 10 pairs, Ethereum (ETH) with 9 pairs, and one USD pair. Several tokens were delisted across multiple base currencies simultaneously, indicating comprehensive removal rather than selective pair management. Examining the quote coins reveals that 24 unique tokens were affected, with some experiencing staggered delistings across different trading pairs. For instance, BEAM was first delisted from its BTC pair in November 2022, followed by its USDT pair in January 2023, suggesting a phased approach. Similarly, HNT-USDT appears twice in our data with different delisting dates on Binance and Binance US.

Prior to 2023, delisting a coin was often triggered by specific incidents or regulatory developments. After 2023, Binance implemented more comprehensive and routine evaluations of the coins on their platform considering four risk categories: legal and compliance risk (privacy projects and tokens under regulatory scrutiny), project risk (low user adoption, development abandonment, or team incidents), market risk (poor liquidity and low market capitalization), and ethical risks (fraudulent or market manipulation evidence).[15] We see examples of all categories in our delisting sample. For example, the privacy coin BEAM was removed amid regulatory pressure on privacy-focused assets[16], FTX coin (FTT) was delisted following the collapse of the FTX trad-

---

[15] See: Binance: Delisting criteria FAQ.

[16] Binance announcement: BEAM delisting (Jan 26, 2023)

ing venue[17], Helium (HNT) was removed during its migration to Solana[18], Groestlcoin (GRS) and Nebulas (NAS) were taken down for market-risk reasons (e.g., low liquidity)[19], and Mithril (MITH) was delisted due to prolonged project inactivity[20].

C.2 Matched Control Group Diagnostics

To validate the propensity score matching approach used in the main analysis, Figure A.3 visualizes the pre- and post-delisting trajectories for delisted coins versus control coins. The figure plots the trends for the matching variables: Logged USD Volume, Price, Volatility, Number of Venues, and Fragmentation. These graphs show that while delisted coins experience a sharp contraction in volume and price relative to controls, the matching procedure effectively aligns the pre-event trends of the control group.

D. Disaggregated Results

To ensure our main findings are not solely driven by the principal component (PC) aggregation, we replicate both event studies using individual measures of efficiency and price impact as dependent variables.

For price efficiency, Table A.4 reports the Synthetic Difference-in-Differences (SDID) treatment effects on individual efficiency measures for the FTX collapse, including autocorrelation of returns and variance ratios across multiple time horizons. Table A.5 reports the corresponding difference-in-differences results for autocorrelation coefficients and variance ratios for Binance

---

[17] Binance post: FTT delisting notice

[18] Binance post: HNT migration/delisting

[19] Binance announcement: GRS and NAS delisting

[20] Binance announcement: MITH delisting

delistings. Consistent with the main results, efficiency declines are more severe for coins that had a higher share of trading volume on FTX prior to the collapse.

For price impact, Table A.6 reports the SDID treatment effects on price impact measured at 30-second, 1-minute, 5-minute, and 30-minute horizons for the FTX collapse. Table A.7 reports the results for price impact measures at different time horizons for Binance delistings. The signs and significance levels reported in the main PCA results are consistent across the individual disaggregated metrics for both events.

Table A.1
Price Efficiency Excluding Immediate Collapse Window

This table reports the effect of the FTX collapse on price efficiency excluding the period from November 2 to November 10 using a Synthetic Difference-in-Differences (SDID) approach. This estimator reweights control coins to match the pre-treatment trends and levels of the treated coins. We estimate:

$$\text{Efficiency}_{it} = \beta(\text{Treat}_i \times \text{Post}_t) + \gamma X_{it} + \alpha_i + \delta_t + \varepsilon_{it}$$

Post_t is an indicator equal to 1 from November 10, 2022 onward. Treat_i is a binary indicator equal to 1 if the coin was traded on FTX, segmented by volume share thresholds. The control group is defined based on the pre-period trajectory of logged volume, volatility, number of venues and fragmentation following the SDID methodology of (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). X_{it} includes logged volume in USD, number of trading venues, fragmentation and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. Controls are optimized. We estimate the model using all coins and then segment the treated sample based on the volume share of FTX being at least 10%, 20%, 30% or 40% during the month prior to the collapse. We present significance at 1%, 5% and 10% as ***, **, and * respectively.

	Full Sample	10%	20%	30%	40%
	[1]	[2]	[3]	[4]	[5]
Treat × Post	-0.089	-0.64***	-1.11***	-1.20***	-1.41***
	(0.076)	(0.15)	(0.24)	(0.25)	(0.32)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	78,693	73,134	72,369	72,165	71,961

Table A.2
Mechanisms Excluding Immediate Collapse Window: Price Impact

This table reports the effect of the FTX collapse on price impact excluding the period from November 2 to November 10 using a Synthetic Difference-in-Differences (SDID) approach. This estimator reweights control coins to match the pre-treatment trends and levels of the treated coins. We estimate:

$$\text{Price Impact}_{it} = \beta(\text{Treat}_i \times \text{Post}_t) + \gamma X_{it} + \alpha_i + \delta_t + \varepsilon_{it}$$

Post_t is an indicator equal to 1 from November 10, 2022 onward. Treat_i is a binary indicator equal to 1 if the coin was traded on FTX, segmented by volume share thresholds. The control group is defined based on the pre-period trajectory of logged volume, volatility, number of venues and fragmentation following the SDID methodology of (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). X_{it} includes logged volume in USD, number of trading venues, fragmentation and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. Controls are optimized. We estimate the model using all coins and then segment the treated sample based on the volume share of FTX being at least 10%, 20%, 30% or 40% during the month prior to the collapse We present significance at 1%, 5% and 10% as ***, **, and * respectively.

	Full Sample [1]	10% [2]	20% [3]	30% [4]	40% [5]
Treat × Post	0.19**	0.63*	0.85	1.17	1.19
	(0.088)	(0.36)	(0.55)	(0.75)	(0.95)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	78,693	73,134	72,369	72,165	71,961

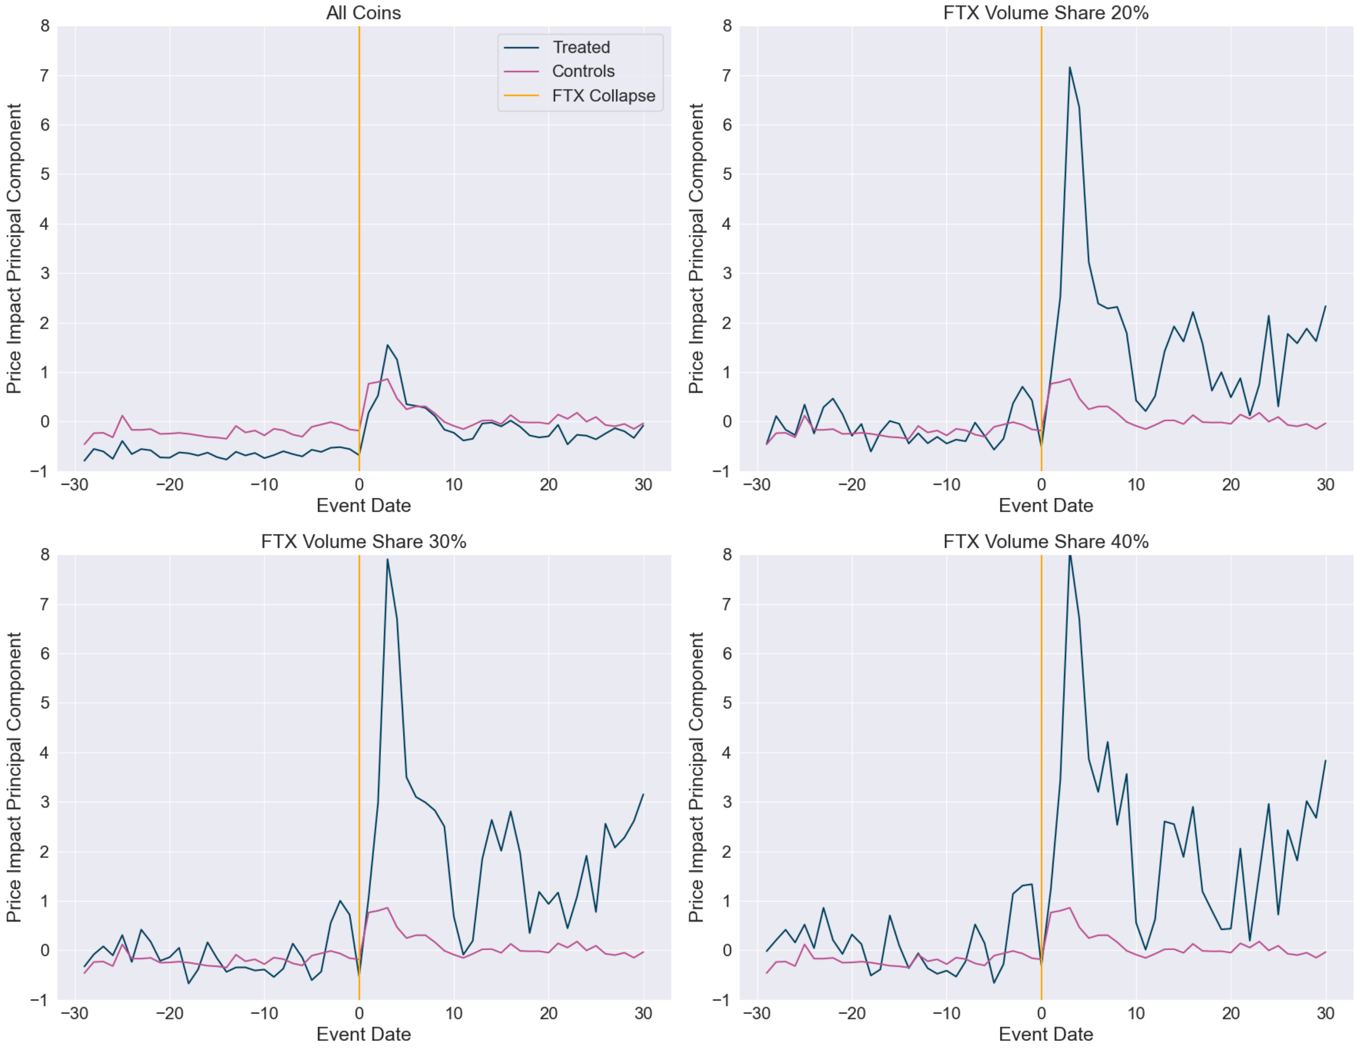

Figure A.1. Price Impact Trends by FTX Volume Share

This figure shows parallel trends from a synthetic difference-in-differences (SDID) analysis comparing price impact for all coins in the treated group, and for a subset where FTX represented at least 20%, 30%, or 40% of trading volume. The control group are coins not traded on FTX weighted to match the pre-period trends in number of venues, fragmentation, volume, and volatility. The vertical line indicates event day 0: the date FTX collapsed on November 08, 2022.

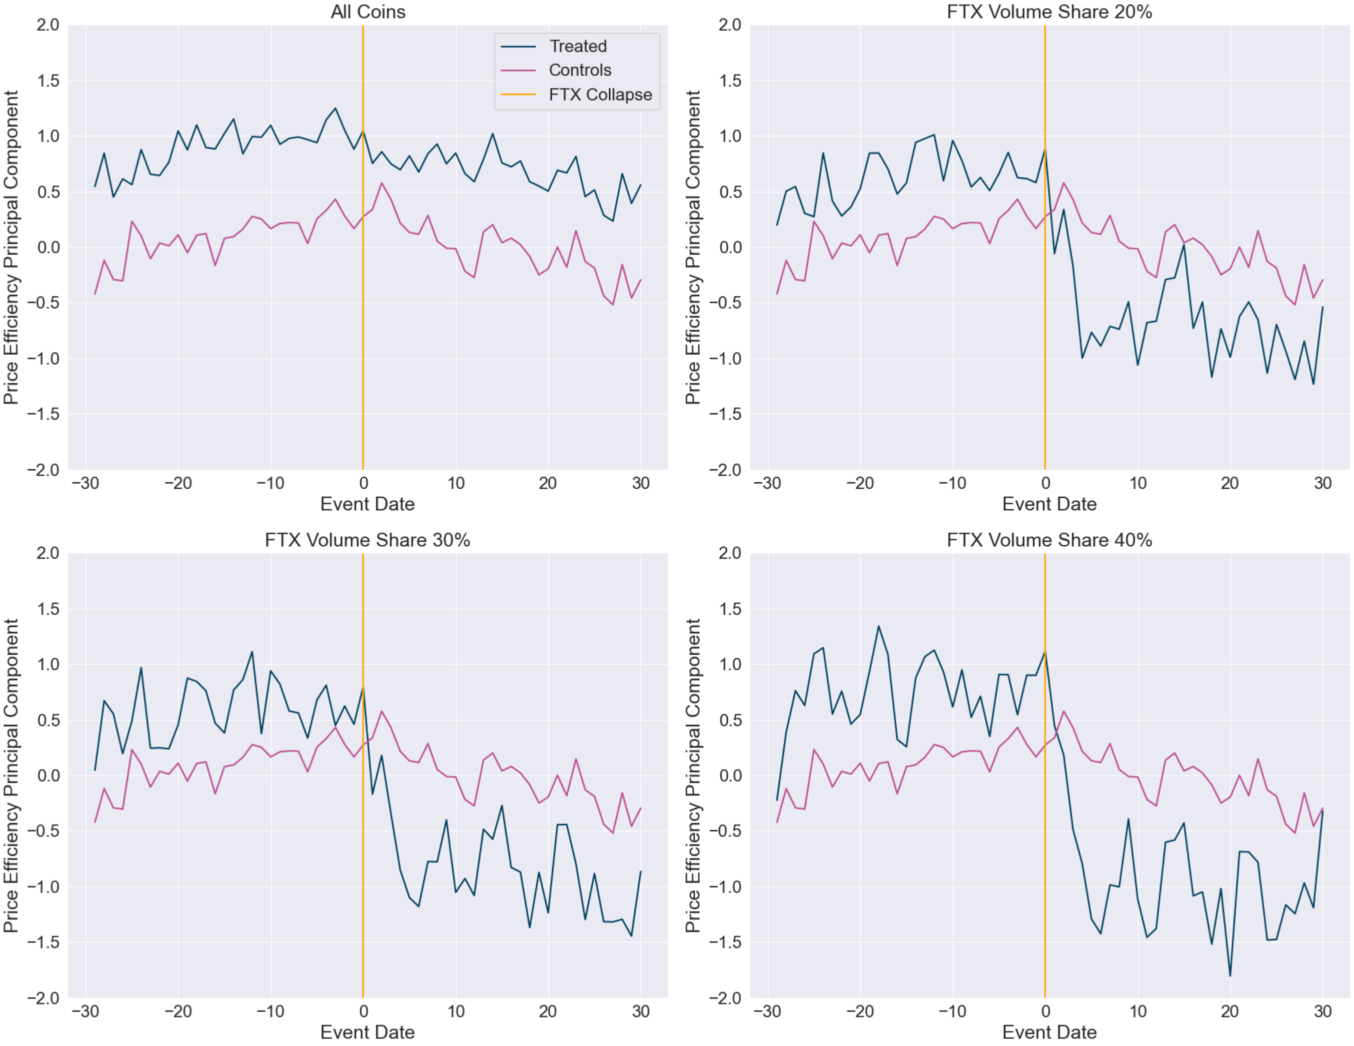

Figure A.2. Price Efficiency Trends by FTX Volume Share

This figure shows parallel trends from a synthetic difference-in-differences (SDID) analysis comparing price efficiency for all coins in the treated group, and for a subset where FTX represented at least 20%, 30%, or 40% of trading volume. The control group are coins not traded on FTX weighted to match the pre-period trends in number of venues, fragmentation, volume, and volatility. The vertical line indicates event day 0: the date FTX collapsed on November 08, 2022.

Table A.3
Coin Delistings from Binance

This table lists coins that were delisted from Binance or Binance US, as confirmed on Binance's website. HNT-USDT appears twice in our data since it was delisted on Binance and Binance-US on separate dates.

Coin Pair	Delisting Date	Coin Pair	Delisting Date
CVC-ETH	2022-03-03	SRM-BTC	2022-11-28
XEM-ETH	2022-03-03	SRM-USDT	2022-11-28
RENBTC-BTC	2022-04-11	XEM-BTC	2022-11-28
GRS-BTC	2022-07-08	CVC-BTC	2022-12-09
NAS-BTC	2022-07-08	MITH-USDT	2022-12-14
NAS-ETH	2022-07-08	REP-USD	2022-12-14
DNT-BTC	2022-10-12	BLZ-ETH	2023-01-06
NBS-USDT	2022-10-12	ICX-ETH	2023-01-06
TCT-USDT	2022-10-12	BEAM-USDT	2023-01-26
HNT-BTC	2022-10-14	BTS-BTC	2023-01-27
HNT-USDT	2022-10-14	KEY-ETH	2023-01-27
FTT-BTC	2022-11-15	JUV-BTC	2023-02-10
FTT-USDT	2022-11-15	KEY-USDT	2023-02-10
BEAM-BTC	2022-11-28	CVP-ETH	2023-03-02
BTCST-USDT	2022-11-28	ONT-ETH	2023-03-16
PNT-BTC	2022-11-28	HNT-USDT	2023-03-24
REP-BTC	2022-11-28	KNC-ETH	2023-03-24

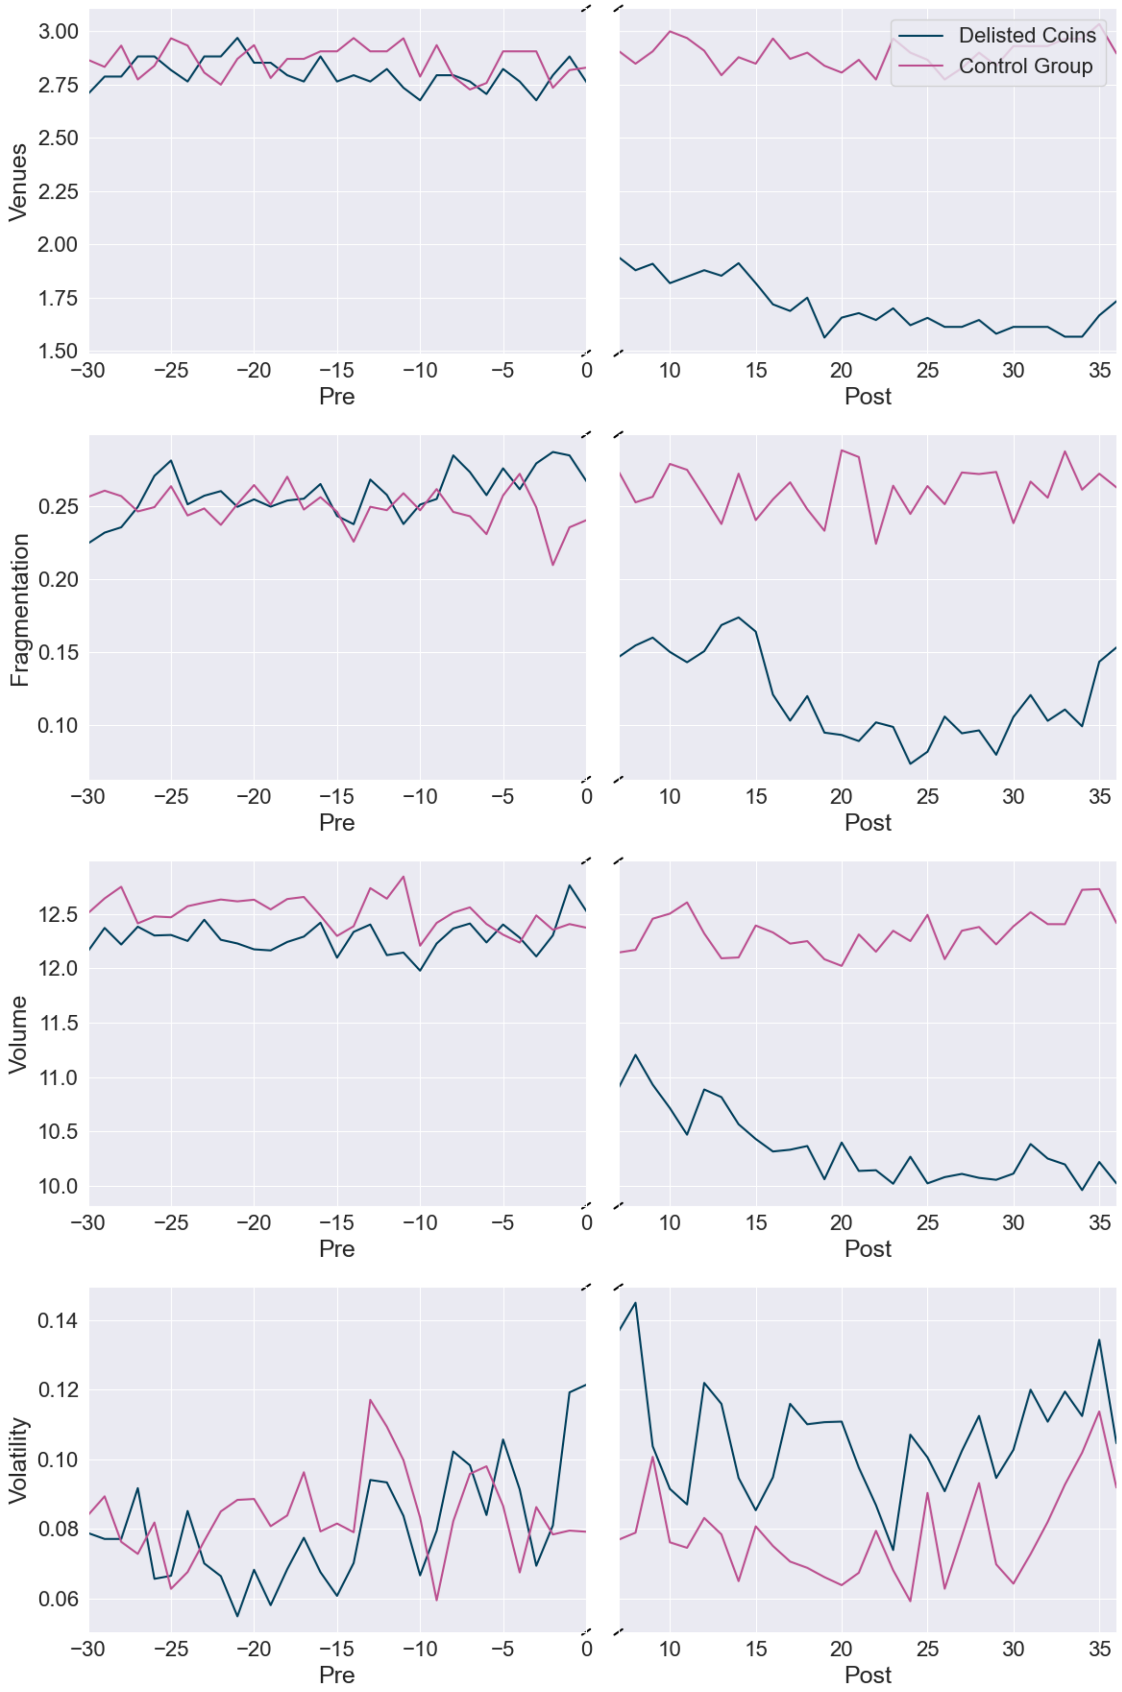

Figure A.3. Matched Variables for Delisted and Control Coins

This figure shows the trends of the *Number of Venues*, *Fragmentation*, logged USD *Volume*, and *Volatility* for the delisted coins and the control group 30 days before and after the delisting announcement. We take out the 7 days immediately after the announcement to exclude the noise as market participants react to the news.

Table A.4
FTX SDID on Individual Efficiency Measures

This table reports the effect of the FTX collapse on individual efficiency measures using a Synthetic Difference-in-Differences (SDID) approach. This estimator reweights control coins to match the pre-treatment trends and levels of the treated coins. We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}$$

Post_t is an indicator equal to 1 from November 08, 2022 onward. Treat_i is a binary indicator equal to 1 if the coin was traded on FTX, segmented by volume share thresholds. The control group is defined based on the pre-period trajectory of logged volume, volatility, number of venues and fragmentation following the SDID methodology of (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). X_{it} includes logged volume in USD, number of trading venues, fragmentation and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. Controls are optimized. We estimate the model using all coins and then segment the treated sample based on the volume share of FTX being at least 10%, 20%, 30% or 40% during the month prior to the collapse We present significance at 1%, 5% and 10% as ***, **, and * respectively.

	Full Sample [1]	10% [2]	20% [3]	30% [4]	40% [5]
Panel A: Autocorrelation					
Treat×Post	-0.0028	-0.024***	-0.040***	-0.039***	-0.042**
	(0.0042)	(0.0071)	(0.012)	(0.012)	(0.018)
Panel B: Variance Ratio 1 Min					
Treat×Post	-0.028***	-0.083***	-0.12***	-0.13**	-0.13*
	(0.0094)	(0.026)	(0.041)	(0.052)	(0.066)
Panel C: Variance Ratio 5 Min					
Treat×Post	-0.018**	-0.076***	-0.13***	-0.14***	-0.19***
	(0.0088)	(0.021)	(0.031)	(0.036)	(0.035)
Panel D: Variance Ratio 15 Min					
Treat×Post	-0.013**	-0.062***	-0.11***	-0.13***	-0.15***
	(0.0055)	(0.014)	(0.020)	(0.019)	(0.030)
Panel E: Variance Ratio 30 Min					
Treat×Post	-0.013***	-0.038***	-0.048***	-0.056***	-0.064***
	(0.0039)	(0.0087)	(0.015)	(0.015)	(0.023)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	92,580	86,040	85,140	84,900	84,660

Table A.5
Delisting DID on Individual Efficiency Measures

This table presents the individual efficiency measures for the Delisting Difference-in-Differences analysis. We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \Theta_t + \alpha_i + \varepsilon_{it}$$

where Outcome_{it} is the specific price efficiency measure (Autocorrelation or Variance Ratio) for coin i on event day t. The event days range from 30 days before and after the delisting announcement, with the 7 days immediately after the delisting removed. Post_t is a dummy variable equal to 1 for the 30 days after the announcement and 0 for the 30 days before. Treat_i equals 1 for the delisted coins and 0 for the control coins. X_{it} includes time-varying control variables: *Number of Venues*, logged USD *Volume*, *Volatility*, and logged USD *Price*. α_i represents coin fixed effects, and Θ_t represents event-day fixed effects. Standard errors are clustered by coin and event day. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Auto Corr (1m)	Var Ratio (1m)	Var Ratio (5m)	Var Ratio (15m)	Var Ratio (30m)
	[1]	[2]	[3]	[4]	[5]
Panel A: No Controls					
Post × Treat	-0.045***	-0.10***	-0.095***	-0.074***	-0.043***
	(0.014)	(0.027)	(0.027)	(0.018)	(0.012)
R-squared	0.315	0.560	0.398	0.254	0.106
Panel B: Controls					
Post × Treat	-0.044***	-0.043	-0.053*	-0.053**	-0.035**
	(0.015)	(0.026)	(0.028)	(0.020)	(0.014)
Venues	-0.0070	-0.013	-0.024	-0.0035	-0.0056
	(0.0062)	(0.016)	(0.018)	(0.014)	(0.0096)
Log(Volume)	0.0072*	0.018**	0.028***	0.016***	0.0044
	(0.0039)	(0.0077)	(0.0066)	(0.0053)	(0.0045)
Volatility	0.12**	0.48***	0.34***	0.0014	-0.096
	(0.056)	(0.085)	(0.098)	(0.067)	(0.065)
Log(Price)	-0.0020	0.053***	0.023	-0.0093	0.0022
	(0.0086)	(0.018)	(0.018)	(0.015)	(0.0092)
Coin FE	Yes	Yes	Yes	Yes	Yes
Event FE	Yes	Yes	Yes	Yes	Yes
Observations	3,848	3,848	3,848	3,848	3,848
R-squared	0.334	0.618	0.450	0.264	0.107

Table A.6
FTX SDID on Individual Price Impact Measures

This table reports the effect of the FTX collapse on individual price impact measures using a Synthetic Difference-in-Differences (SDID) approach. This estimator reweights control coins to match the pre-treatment trends and levels of the treated coins. We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \alpha_i + \Theta_t + \varepsilon_{it}$$

Post_t is an indicator equal to 1 from November 08, 2022 onward. Treat_i is a binary indicator equal to 1 if the coin was traded on FTX, segmented by volume share thresholds. The control group is defined based on the pre-period trajectory of logged volume, volatility, number of venues and fragmentation following the SDID methodology of (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). X_{it} includes logged volume in USD, number of trading venues, fragmentation and volatility. Coin and date fixed effects are included. Standard errors are bootstrapped with 100 replications. Controls are optimized. We estimate the model using all coins and then segment the treated sample based on the volume share of FTX being at least 10%, 20%, 30% or 40% during the month prior to the collapse We present significance at 1%, 5% and 10% as ***, **, and * respectively.

	Full Sample	10%	20%	30%	40%
	[1]	[2]	[3]	[4]	[5]
Panel A: Price Impact 30 sec					
Treat×Post	3.81*	16.7**	22.1**	25.6*	26.2
	(2.02)	(7.08)	(10.2)	(14.1)	(19.1)
Panel B: Price Impact 1 min					
Treat×Post	4.04*	17.4**	22.6**	27.3*	28.5
	(2.07)	(7.72)	(11.2)	(15.4)	(20.9)
Panel C: Price Impact 5 min					
Treat×Post	4.86**	19.5**	25.7**	33.6**	34.9
	(2.11)	(8.07)	(12.1)	(16.0)	(21.9)
Panel D: Price Impact 30 min					
Treat×Post	4.93**	16.8**	21.0	31.8*	32.6*
	(2.34)	(8.33)	(13.3)	(16.2)	(19.3)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	92,580	86,040	85,140	84,900	84,660

Table A.7
Delisting DID on Individual Price Impact Measures

This table presents the individual price impact measures for the Delisting Difference-in-Differences analysis. We estimate:

$$\text{Outcome}_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \gamma X_{it} + \Theta_t + \alpha_i + \varepsilon_{it}$$

where Outcome_{it} is the specific price impact measure at different horizons for coin i on event day t. The event days range from 30 days before and after the delisting announcement, with the 7 days immediately after the delisting removed. Post_t is a dummy variable equal to 1 for the 30 days after the announcement and 0 for the 30 days before. Treat_i equals 1 for the delisted coins and 0 for the control coins. X_{it} includes time-varying control variables: *Number of Venues*, logged USD *Volume*, *Volatility*, and logged USD *Price*. α_i represents coin fixed effects, and Θ_t represents event-day fixed effects. Standard errors are clustered by coin and event day. Significance at 1%, 5%, and 10% is denoted by ***, **, and * respectively.

	Price Impact (30s) [1]	Price Impact (1m) [2]	Price Impact (5m) [3]	Price Impact (30m) [4]
Panel A: No Controls				
Post × Treat	21.7**	23.6***	30.6***	29.4***
	(8.39)	(8.56)	(9.41)	(9.43)
R-squared	0.473	0.468	0.432	0.343
Panel B: Controls				
Post × Treat	1.92	2.91	6.41	9.60
	(7.10)	(7.28)	(8.42)	(7.98)
Venues	-3.06	-3.66	-4.25	-2.76
	(4.08)	(4.15)	(5.07)	(5.47)
Log(Volume)	-5.29**	-5.88**	-7.21***	-7.04***
	(2.34)	(2.39)	(2.33)	(2.41)
Volatility	217***	239***	291***	295***
	(27.0)	(28.9)	(33.2)	(38.8)
Log(Price)	1.03	2.76	4.30	7.12
	(5.37)	(5.61)	(6.64)	(6.96)
Coin FE	Yes	Yes	Yes	Yes
Event FE	Yes	Yes	Yes	Yes
Observations	3,848	3,848	3,848	3,848
R-squared	0.546	0.547	0.521	0.412