Competition in the Cryptocurrency Exchange Market*

Junyi Hu[†], Anthony Lee Zhang[‡]

June 2023

Abstract

How do cryptocurrency exchanges compete with each other? We show that small and large crypto exchanges appear to be complements, rather than substitutes, as traditional oligopoly theory would predict. When large exchanges list new tokens, trade volumes on small exchanges increase, and small exchanges become more likely to list. We rationalize these facts in a model where small exchanges have captive customer bases, and rely on arbitrage trade with large exchanges for liquidity provision. Our results imply that large exchanges' listing decisions play a systemically important "leader" role in determining trade volumes and listings on other exchanges.

*We thank Simon Trimborn, Yizhou Xiao, the wagmi chat, and seminar participants at CUHK and the Hong Kong Conference for Fintech, AI, and Big Data for helpful comments. We are especially grateful to Jinfei Sheng for comments that greatly improved the paper. We are grateful to Yichun He, Haichuan Wang, Wanyi Wang, and Chupei Zhang for excellent research assistance.

[†]Fudan University, School of Economics; junyihu19@fudan.edu.cn.

[‡]University of Chicago, Booth School of Business; anthony.zhang@chicagobooth.edu.

1 Introduction

Cryptocurrency exchanges are financial intermediaries which allow customers to trade cryptoassets, against either fiat currencies or other cryptoassets. The crypto exchange market is surprisingly fragmented. There are over 1000 different crypto exchanges, offering essentially the same few assets to trade.[1] There are over 100 active cryptocurrency exchanges in the United States alone, compared to only 16 exchanges for equity trading.[2] There are a number of very large crypto exchanges, but their market share is modest: the top 2 crypto exchanges are responsible for only around 15% of total BTC trading volume, as of 2022. This paper analyzes the structure of strategic interactions between crypto exchanges. How do crypto exchanges compete with each other? In particular, if cryptocurrency exchanges compete for trade volume of the same coin within a fixed customer base, why does exchange market structure not consolidate into a monopoly or oligopoly, where all customers trade on a small number of large and liquid exchanges? From a normative perspective, how large of a role do the largest few exchanges play in the ecosystem, given that the market for crypto exchanges appears fairly competitive?

We begin by demonstrating a number of surprising facts about competition between crypto exchanges. Suppose a large exchange lists a new token. If exchanges were competing over a fixed customer base, trade volumes of the token on smaller exchanges should *decrease*, and small exchanges who have not already listed the token should be *less* likely to list, due to the entry of a large competitor. We find exactly the opposite patterns empirically. When a large exchange lists a new token, trade volumes of the token on smaller exchanges *increase*, and small exchanges become *more* likely to list the token. In other words, large and small exchanges appear to behave like economic complements, rather than economic substitutes.

We rationalize these results in a simple model in which a "periphery" of small crypto exchanges have captive customer bases, and rely partially on arbitrage flows with a large and deep "core" exchange for liquidity provision to their customers. When a core exchange lists a new token, trade volume on peripheral exchanges increases, as arbitragers bridge inventory shocks to peripheral exchanges' captive customers into the deep liquidity on the core exchange. Thus, peripheral exchanges anticipate increased profits on a token after it is available for trading on core exchange, giving peripheral exchanges incentives to follow core exchanges' listing decisions. The model makes predictions about the structure of price

[1] For example, a partial list of exchanges can be found on Blockspot.io, where the general exchanges are classified as either a buy/sell platform, exchange, derivatives exchange, futures exchange, P2P exchange, or NFT marketplace.

[2] The list of exchanges can be found on the SEC website. Note that 12 of these exchanges are run by three groups: Intercontinental Exchange Inc NYSE, Nasdaq Inc, and Cboe Global Markets.

correlations across exchanges, and how price correlations relate to the volume-increase and listing-following effects, which we verify empirically.

Positively, our results indicate that crypto exchanges do not compete like firms in oligopolistic product markets. Instead, small peripheral exchanges essentially behave as "costly windows" for their captive customer bases to access the deep liquidity available on a small number of large "core" exchanges. Our results imply that, while the crypto exchange market appears unconcentrated, the large players play a systemically important "leader" role which is understated by their market share of trading volume: their listings decisions, which at present are essentially unregulated, play a substantial role in driving listing decisions and trade volumes on the long tail of smaller exchanges.

We use data on 500 large crypto tokens' prices and trades volumes across 262 exchanges from January 2017 to July 2022. Using this data, we demonstrate three stylized facts. First, when a large exchange lists a new token, trade volumes of the token on small exchanges tend to *increase*, by around 22%-160% across different specifications. Second, small exchanges tend to follow large exchanges' listing decisions: large exchange listings associate with an increase in the number of peripheral exchanges which list a token. Both of these results are surprising in light of standard theories of oligopolistic competition: they suggest that large exchanges are economic complements rather than substitutes to small exchanges. Third, we find that the entry of a large exchange tends to also decrease the *dispersion* of token prices across small exchanges by around 6%-33%, suggesting that arbitrage flows across large and small exchanges may play a role in explaining these findings.

We construct a simple model to rationalize these results. We model the strategic interactions between a single "core" exchange and a number of "peripheral" exchanges, which have captive customer bases and partially rely on imperfect arbitrage with the core exchange for liquidity provision to their customers. There is a single risky asset, or "token", which can be traded. The core exchange has infinite market depth. Each peripheral exchange has a set of captive customers, who receive inventory shocks for the risky asset, and who can only trade on the peripheral exchange. Inventory shocks have an aggregate and idiosyncratic component, so customers of a given peripheral exchange may on net want to buy or sell a token. Customers have holding costs for the asset, so aggregate inventory shocks generate pressure on peripheral exchange prices. Each peripheral exchange also has a set of arbitrageurs, who can trade on the peripheral exchange and the core exchange to partially close price gaps for the risky asset. Arbitrageurs have inventory costs, implying that they cannot fully close price gaps induced by inventory shocks. Peripheral exchanges collect fees depending on trade volume, and list the coin if anticipated fees are greater than an exogeneous cost of listing.

In the absence of the core exchange, trade on peripheral exchanges is generated only by the idiosyncratic component of customers' inventory shocks. If customers have positive inventory positions on average, they cannot sell these positions to others, so the token price must decrease significantly to clear the market. Inventory shocks thus have relatively large effects on prices, and trade volumes are relatively low. When the core exchange lists the coin, arbitrageurs trade to partially close the price gaps between the peripheral exchange and the core exchange. This effectively gives peripheral exchange customers partial access to core exchange liquidity, decreasing the price impact of aggregate inventory shocks. Moreover, arbitrage activity generates increased trade volume on the peripheral exchange, which also increases the expected profits of the peripheral exchange.

The model explains our three stylized facts. Peripheral exchanges' customers are fully captive, so the core exchange's entry does not directly cannibalize the peripheral exchange's customers; however, the entry of the core exchange allows arbitrage trade with the peripheral exchange, causing trade volumes to increase. Since peripheral exchanges' profits from listing a coin are higher when they anticipate higher trading volumes, peripheral exchanges thus have an incentive to follow core exchanges' token listing decisions. The model also predicts that core exchange listings should decrease the cross-sectional dispersion of token prices across peripheral exchanges should both decrease, since core-periphery arbitrage trade cause peripheral exchange prices to cluster close to core exchange prices.

The model makes two additional predictions, which we bring to the data. First, price correlations between exchanges should have a core-periphery structure. Peripheral exchange prices consist of the core exchange's price, plus noise generated by inventory shocks of the peripheral exchange's customers; thus, the correlation between a peripheral exchange's price and the core exchange's price should be greater than the correlation between two peripheral exchanges' prices. Second, the three phenomena we have documented should be associated with each other across exchanges: peripheral exchanges which rely more on arbitrage with the core exchange should have stronger price correlations with the core exchange, larger volume increases when the core exchange lists, and a larger tendency to follow the core exchange's listing decisions.

From a positive perspective, our results shed light on the nature of competition between crypto exchanges. Small and large crypto exchanges do not compete in the manner of classic models of oligopolistic competition; rather, peripheral exchanges appear to act like "costly windows" for captive customer bases to deep liquidity on core exchanges. We do not have direct evidence on the nature of the frictions that cause small exchanges customer bases to be captive. One possibility is that there are regulatory or jurisdictional barriers, preventing customers from accessing core exchanges directly; small peripheral exchanges thus act like financial intermediaries, who navigate the institutional and regulatory barriers to offering crypto trading services in a given local environment, but largely outsource liquidity provision to large international core exchanges through arbitrage. Another possibility is that consumers are unsophisticated, are attracted to peripheral exchanges through advertising, and face search frictions for directly trading on core exchanges. Peripheral exchanges are thus able to collect spreads from the fact that their customers are not sophisticated enough to switch to the deeper core exchange. Both hypotheses suggest that consumers may be better off if these frictions were eliminated, allowing customers to trade directly in the large core exchanges. This would disintermediate both peripheral exchanges and the arbitrageurs which profit from the spreads between core and peripheral exchanges.

Our results also imply that a small set of core exchanges play a systemically important "leader" role in crypto markets. Large exchanges' market shares of total trading volume is modest – the top 2 exchanges account for only around 15% of total trading volume in BTC as of 2022. But these market shares likely understate the importance of large exchanges, since their listing decisions have substantial power to affect token trading volumes, liquidity, and the decisions of other exchanges whether to list tokens for trading. Quantitatively, in a simple back-of-envelope calculation combining the forces we analyze here, we find that Binance's decision to list a coin increases daily total trade volumes by around 1046pp, and Coinbase's listings increase volumes by around 348pp, within the 10-day window after listing. A large part of this effect is indirect: the volume increase consists of a 295pp (60pp) increase in volume directly on Binance (Coinbase), and a 199pp (213pp) increase on incumbent exchanges, and new exchanges which list the token following Binance or Coinbase's listing decisions. Despite the large power core exchanges have to shape market outcomes, core exchanges currently have a large degree of freedom to decide which assets to list.[3] Thus, regulators may wish to monitor the listing decisions of large crypto exchanges, for example requesting that exchanges provide data on tokens they plan to list, and the reasoning for listing these tokens.

This paper relates most closely to a few other papers that study cryptocurrency exchanges. Augustin, Rubtsov and Shin (2020) analyzes how the introduction of BTC futures contracts affects price discovery and price dispersion for BTC prices across exchanges. Makarov and Schoar (2020) show that there are often large and persistent deviations in BTC prices across exchanges, which are smaller within countries, and appear to be related to capital controls. Choi, Lehar and Stauffer (2022) analyze the "Kimchi premium", the phenomenon that BTC tends to trade at a premium in Korea relative to the USA. Makarov and Schoar (2019)

[3]Exchanges' ability to list tokens varies by jurisdiction, however; for example, exchanges serving US customers tend not to list tokens which the exchange believes are likely to violate US securities regulations.

characterizes the exchanges which drive BTC price discovery. Relative to this literature, our contribution is to analyze the strategic interactions between exchanges, by studying how the listing decisions of large exchanges affect smaller exchanges' trade volumes and listing decisions. We also relate to a number of papers that have analyzed the effect of large exchange listings on returns (Ante, 2019; Lemmen, 2022), though we do not focus on returns in this paper.

We also fit into the broader literature on cryptocurrencies and decentralized finance, which is surveyed in Harvey, Ramachandran and Santoro (2021), John, Kogan and Saleh (2022), and Makarov and Schoar (2022). A number of other papers, such as Chan et al. (2020) and Kogan et al. (2023), analyze retail investors' crypto trading strategies using investor-level data from cryptocurrency exchanges. Yu and Zhang (2022) show that demand for BTC increases with local economic policy uncertainty. Liu and Tsyvinski (2018) and Liu, Tsyvinski and Wu (2022) analyze the risk factors underlying cryptocurrency returns. Liu, Sheng and Wang (2021) contruct a tech index from ICO whitepapers which predicts crypto returns. Cong et al. (2020) discuss wash trading in crypto, and Amiram, Lyandres and Rabetti (2022) analyze which exchanges engage in wash trading to a larger extent. Li, Shin and Wang (2021) analyzes cryptocurrency pump-and-dump schemes. Augustin, Chen-Zhang and Shin (2022) analyze returns from "yield farming" liquidity provision strategies. Cong and He (2019) discuss smart contracts, and Cong, Li and Wang (2019) discuss tokenomics. Outside of the literature on cryptocurrencies and blockchains, Budish, Lee and Shim (2019) analyze how stock exchanges compete on the dimension of trading speed, and how this affects exchanges' innovation incentives.

The paper proceeds as follows. Section 2 describes institutional background around cryptocurrency exchanges and the data we use. Section 3 describes our stylized facts. Section 4 describes our model. Section 5 tests the predictions of our model empirically. We conclude in Section 6.

2 Data and Institutional Background

2.1 Institutional Background

Cryptocurrency exchanges, analogous to exchanges for stocks, bonds, and other financial assets, allow customers to exchange fiat currencies for cryptocurrencies. Crypto exchanges function in a custodial manner: they allow users to "deposit" fiat or cryptocurrencies, hold fiat currencies and cryptocurrencies on behalf of users, and allow users to trade their custodied fiat and crypto with other users of the exchange. For the vast majority of exchanges, trading in each assets is governed through limit-order books.

Like regular financial asset exchanges, users can deposit and withdraw fiat from the exchange through bank transfers or other means. A unique feature of cryptocurrency exchanges, relative to exchanges for stocks or other traditional financial assets, is that users can also deposit or withdraw cryptocurrencies from the exchange. Users can "withdraw" custodied assets, instructing the exchange to send funds held on her behalf to her own private "wallet" address. Users can also deposit cryptocurrencies, sending it to a designated "deposit" address, and receiving on-exchange custodially-owned crypto in exchange. The ability to deposit and withdraw implies that an important function of crypto exchanges is also to serve as "on/off-ramps" for crypto: allowing users to deposit fiat, exchange fiat for crypto, and withdraw crypto, or vice versa.

As an example of the role cryptocurrency exchanges play in the process of using cryptocurrencies, in Appendix A.1, we describe in detail how a customer would use crypto exchanges and cryptocurrency on-chain transfers to perform an international funds transfer, exchanging, for example, fiat currency in the USA for fiat in the Philippines. In short, a customer would exchange USD fiat for cryptocurrencies using a US crypto exchange, and send the crypto to the funds receipient, who would then exchange the crypto for Phillipine fiat currency. Using cryptocurrencies to perform such transfers is convenient because it allows consumers to partially circumvent capital controls and other restrictions imposed by policymakers, as well as fees charged by intermediary financial institutions who facilitate traditional international remittances. Appendix A.1 also briefly discusses the regulation of crypto exchanges. Crypto exchanges have nontrivial difficulty in expanding across jurisdictions for a number of reasons. First, since crypto exchanges must allow both crypto and fiat deposits and withdrawals, exchanges must be able to integrate with local banking systems for fiat funds transfers. Secondly, due to the necessity of integrating with local banking systems, crypto exchanges logistically must work with local financial regulators, and are subject to varying regulations depending on the jurisdictions they operate within.

There are a number of other uses of cryptocurrencies besides remittances: users in countries with high inflation or low confidence in financial institutions might buy and self-custody cryptocurrencies as a store of value.[4] Cryptocurrencies can also be used to perform a number of simple financial transactions, within the space of "decentralized finance."[5] In addition,

[4]See CNBC and Rest Of World for a discussion of the use of cryptocurrencies as a store of value in Lebanon.

[5]For example, market participants can use stablecoin tokens to purchase other blockchain tokens, such as ETH, MKR, or UNI, using an automated market maker protocol such as Uniswap. Market participants can also lend stablecoin tokens on lending and borrowing protocols, such as Aave and Maker, allowing them

many market participants purchase cryptocurrencies on centralized exchanges to speculate on crypto price appreciation.

While we focus on the role of crypto exchanges in facilitating spot trading, crypto exchanges also offer consumers other products. Some exchanges also offer individuals the ability to short cryptocurrencies or use leverage, as well as various derivative products such as futures and options; see, for example, Bybit and Binance. Some exchanges also pay consumers interest rates on users' custodied cryptocurrencies, in a manner similar to bank deposits or CDs. Two such examples are Binance and Coinbase. Like traditional exchanges, crypto exchanges also offer market data and analytics products; see, for example, Binance.

There are a very large number of crypto exchanges: according to Blockspot.io, as of 2023 there are over 1,000 different exchanges offering fairly similar assets to trade. A small number of very large exchanges account for a nontrivial, but modest, fraction of total market share. Figure 1 shows the market share of the top 2 exchanges in our data for BTC volume. The market share of the top 2 exchanges is fairly large, reaching 15% in 2022. Moreover, this is likely an underestimate of the top exchanges' market shares, since small exchanges are anecdotally known to falsely overreport or manipulate trade volumes (Cong et al., 2020).

2.2 Data

The primary dataset we use in this paper is from cryptotick.com, which collects coin trade price and quantity data from a broad set of cryptocurrency exchanges. Cryptotick obtains this data by querying APIs provided by the exchanges, and timestamping data using the same synchronized clock (UTC time) for all exchanges. The dataset contains hourly OHLCV data, that is, open, high, low and close prices, as well as total trade volume, each hour on each exchange. One data series in the cryptotick data represents one trading pair on a given exchange; that is, a pair involving trading of one cryptocurrency for either a fiat currency, a stablecoin (that is, a cryptocurrency designed to be worth the same as a fiat currency), or another major cryptocurrency such as BTC or ETH. We aggregate the data to daily data for each trading pair-exchange ID, taking the average of the prices for each open hour within a day, and adding the trade volumes across all hours within a day.

to receive positive interest rates, and also to use these assets as collateral to borrow other assets. Market participants can speculate on the prices of assets using derivatives-like contracts, on platforms such as dydx. These functionalities are enabled not by trusted centralized financial intermediaries, or legal contracts, but by pieces of code embedded in the blockchain, which programmatically perform transactions, like exchanging one crypto token for another, in a fully automated manner which does not involve any human discretion. In this way, the blockchain ecosystem allows individuals to engage in a number of simple financial transactions, such as trading assets, borrowing and lending, and speculating, in a way that does not require trust in any legal system, financial institution, or other human entity.

Our data spans January 2017 to July 2022, and there are 264 exchanges and 12,417 coins in the raw dataset. We restrict our sample in two ways: by the cryptocurrency in the trading pair, and by the denominator that the cryptocurrency is traded against. Since many coins are not actively traded, we first restrict our sample to trading pairs involving the top 500 cryptocurrency coins ranked by coinmarketcap.com on September 3, 2022. We also restrict our sample to three kinds of trading pairs: pairs involving one of our 27 major fiat currencies;[6] pairs involving BTC or ETH, which are the two largest cryptocurrencies by market cap; and pairs involving one of the three major stablecoins (USDT, USDC, BUSD).

For all trading pairs, we convert all coin prices and volumes to USD terms; for fiat pairs, we convert using same-day USD-fiat exchange rates, and for crypto pairs, we convert using daily prices of cryptocurrencies and stablecoins from Yahoo finance. We aggregate data to coin-level data by taking the average of the prices for each trading pair involving same coin within a day weighted by its USD trading volume, and adding volumes across all trading pairs of the same coin within a day. The final sample consists of 482 coins across 262 exchanges. Appendix B contains more details about the data cleaning process.

We identify the listing date of a coin on an exchange by observing the first date it appears on an exchange in our price and volume data. We round to the nearest day: if the first trade we observe occurs before 12:00PM, we identify the listing date as the previous day, and if the first trade occurs after 12:00PM, we identify the listing date as the current day. For the central exchange listings, our analysis requires taking a stance on which exchanges are central. We adopt two different specifications: we either treat the largest two exchanges at the present date, Binance and Coinbase, as central. Summary statistics of the data are shown in Table 1.

3 Stylized Facts

We proceed to document three stylized facts about how large exchange listings affect crypto market outcomes, which together suggest that large exchanges are complements rather than substitutes to small exchanges.

3.1 Listings and Trade Volume

Fact 1. *When a large exchange lists a token, incumbent small exchanges which have previously listed the token experience increases in token trading volume.*

[6]These 27 major fiat currencies include: NZD USD KRW JPY CNY IDR SGD VND TWD AUD PKR ZAR TRY MXN BRL CHF ILS PLN GBP RUB EUR CAD HKD INR SAR AED SEK.

To demonstrate Fact 1, we estimate exchange-coin level difference-in-differences (DID) specifications, analyzing how coin trading volumes change when a large exchange lists a new token. We begin with the following flexible DID specification:

$$\log(Volume_{c,e,t}) = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_{c,e} + \eta_t + \epsilon_{c,e,t} \quad (1)$$

where c indexes coins, e indexes exchanges, and t indexes days. $\log(Volume_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t. $treat_{c,k,t}$, with k from -10 to 10, is a series of dummy variables that equal 1 if there are exactly k days from the large exchange listing date to time t, and 0 otherwise. $treat_{c,-11,t}$ and $treat_{c,11,t}$ are dummy variables that equal 1 when time t is more than 10 days before and after the large exchange listing date, respectively. Coins that have never been listed on large exchanges are considered as the control group. If coin c has not experienced a large exchange listing, $treat_{c,-11,t}$ will always equal 1 and other dummy varibales always equal 0. To avoid collinearity problem, we set observations that are exactly 10 days before large exchange listings as the referncence group, which means that $treat_{c,-10,t}$ is always 0.

We plot estimated coefficients of $treat_{c,k,t}$ with k from -10 to 11. The results are shown in Figure 2. When a large exchange lists a token, there is a large increase in trade volumes on small exchanges that have previously listed the token. The difference between the treated and control groups is small in magnitude prior to listings, though there is a slight pre-trend prior to 3 days before listings, and a significant increase in volumes for Coinbase in the 3 days before listings. We believe this is potentially due to a gap between the announcement of listings and their implementation.[7] Directly after the large exchange lists, we observe a sharp increase in trading volume: volumes increase by around 125% and 91%, respectively, for coin-exchange pairs following a coin listing by Binance and Coinbase, respectively, relative to coin-exchange pairs that did not experience a listing event. The coefficients decrease over time, but are still positive and significant 10 days after listings. Though we do not model this effect, we believe the long-run decrease in volumes may be due to entry and cannibalization: we will show that the large exchange's listing tends to induce many other small exchanges to list the token, which may tend to cannibalize market share from incumbent small exchanges which have already listed the token.

[7] Binance usually announces coin listings one day before they are available for trading, while Coinbase typically announces them one or two days in advance.

We then estimate the following differences-in-differences specification:

$$\begin{aligned}\log(Volume_{c,e,t}) =& \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> 30 \text{ days})_{c,t} + \\ & \beta_3 PreThreedayListing_{c,t} + \delta_{c,e} + \eta_t + \epsilon_{c,e,t}\end{aligned} \tag{2}$$

$Listing(\text{0-30 days})_{c,t}$ and $Listing(> 30 \text{ days})_{c,t}$ are dummy variables; $Listing(\text{0-30 days})_{c,t}$ is equal to one for coin c on date t if a large exchange has listed coin c prior to date t but later than date $t-30$, and analogously $Listing(> 30 \text{ days})_{c,t}$ is one if a large exchange has listed coin c prior to date $t-30$. $PreThreedayListing_{c,t}$ is a dummy variable which is equal to one for coin i on date t if a large exchange decides to list coin i between date $t+1$ and date $t+3$. We include this to absorb the effect, from Figure 2, that there is a slight pre-trend in volumes a few days before the large exchange lists the token. The effect of including $PreThreedayListing_{c,t}$ in the regressions is that the coefficients β_1 and β_2 are identified based on outcomes more than 3 days before listings. We cluster standard errors at the coin-exchange pair level.

Essentially, (2) is a difference-in-differences specification. For any coin c, the identification of β_1, β_2, and β_3, are driven by "incumbent" exchanges e which have listed coin c before the large exchange lists the coin. Other coin-exchange pairs serve as a control group, contributing to identifying the time fixed effects η_t. Our coefficients of interest are β_1 and β_2,which measure the extent to which small exchanges which have listed coin c before the large exchange experience unusual increases in volume after the large exchange lists, relative to other coins besides the listed coin. β_1 measures the short-run effect from 0 to 30 days, and β_2 measures the effect after 30 days.

Estimates from (2) are shown in Table 2. The eight columns reflect different combinations of large exchange listings and fixed effects. The first four columns examine the effects of Binance listings, while the last four columns focus on Coinbase listings. For Columns (1) and (5), we control for day and coin fixed effects. Columns (2) and (6) include fixed effects for the country that the exchange operates in, and columns (3) and (7) include exchange fixed effects. Columns (4) and (8) include day and coin-exchange pair fixed effects, controlling for any sources of unobserved heterogeneity which affects all coins traded on a given exchange on a given day. The results are broadly in line with the graphical evidence in Figure 2. The $Listing$(0-30 days) coefficients obtained across all specifications are positive and significant. Quantitatively, we find that Binance listings increase volumes within 30 days by 22% to 160%, and Coinbase listings increase 30-day volumes by around 68% to 94%. Consistent with Figure 2, the estimated effects are smaller after 30 days, though in most cases they are still positive and significant.[8]

Fact 1 is counterintuitive because, if large and small exchanges compete over a fixed pool of customers, large and small exchanges should be *substitute* goods: the entry of large exchanges should tend to cannibalize small exchange trade volumes, as customers migrate to the larger and more liquid large exchanges. Instead, this stylized fact suggests that the large exchange is a *complementary* good to small exchanges: when the large exchange enters, customers on small exchanges have some reason to increase trading, rather than switch to the large exchange.

3.2 Listing Following

Fact 2. *Small exchanges tend to follow large exchanges' listing decisions: after a large exchange lists a new token, many small exchanges that previously did not list the token quickly list it.*

As a descriptive evidence for this prediction, we plot the listing times of small exchanges relative to large exchange listings in Figure 3. For each coin listed by a large exchange, the histogram shows the fraction of small exchanges listing within 100 days before or after the large exchange's listing. There is a sharp increase to the right of 0, illustrating that small exchanges tend to list coins just after large exchanges do. To demonstrate this fact more formally, we estimate the following flexible coin-level specifications:

$$\Delta\#Exchanges_{c,t} = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \tag{3}$$

The dependent variable, $\Delta\#Exchanges_{c,t}$, is the net change in the number of exchanges which list coin c in time t. All other variables are defined above. [9]δ_c and η_t are respectively coin and time fixed effects. The coefficients of interest are β_k, which measure how many listings by small exchanges for a given coin increases after large exchange listing events, compared to coins that did not experience large exchange listings. We cluster standard errors

[8]A caveat in interpreting this result is that volume numbers reported by exchanges are known to be subject to manipulation and wash trading (Cong et al., 2020). One particular concern is if, in response to a large exchange's entry decision, smaller exchanges increasingly engage in wash trading, as suggested by Amiram, Lyandres and Rabetti (2022). While this hypothesis can partly account for our volume result, it does not explain the later stylized facts we will show: that small exchanges follow large exchanges' listing decisions, that price dispersion would decrease across small exchanges, and that these tendencies are associated with the extent of price correlation across exchanges.

[9]Thus, $\Delta\#Exchanges_{c,t}$ is equal to total new listings minus delistings; however, delistings are relatively rare in our data. There are 13180 listings of these top 500 coins by all exchanges in our sample period, and 237 delisting in contrast.

at the coin level. The results are shown in Figure 4. The coefficient estimates are quite flat prior to large exchange listing. Consistent with Figure 3, we find that a large number of small exchanges list tokens within 2 days after they are listed by a large exchange. Binance and Coinbase listings lead to at most 1.25 and 1.40 more net listings per day, respectively, for coins that are recently listed by these large exchanges, relative to coins that have not experienced the listing events.

We then estimate the following regression specification:

$$\begin{aligned}\Delta\#Exchanges_{c,t} =& \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> \text{30 days})_{c,t} + \\ & \beta_3 PreThreedayListing_{c,t} + \delta_c + \eta_t + \epsilon_{c,t}\end{aligned} \tag{4}$$

All other variables are defined above. The results are shown in Table 3. Within 30 days of listing by Binance and Coinbase, the average net listings per day increases by 0.11 and 0.12, respectively. The effects are also persistent. Average net listings per day still increase by 0.01, 30 days after the listing events for coins that have been listed by both Binance and Coinbase, relative to coins that have not experienced the listings. We observe little anticipation effects for both Binance and Coinbase listings, with a significance level at 90%.

If large and small exchanges were substitute goods, small exchanges would expect lower profits from listing a coin if a large exchange has already listed. The fact that we observe small exchanges following large exchanges suggests that small exchange perceive *higher* profits from listing coins after a large exchange has entered. Fact 2 thus provides further evidence that small exchanges view large exchanges as complements to themselves, rather than substitutes.

3.3 Listings and Price Dispersion

Fact 3. *When a large exchange lists a token, the dispersion of token prices across small exchanges decreases.*

Figure 5 shows how large exchange listings affect price dispersion across small exchanges. For each coin affected by a listing, we measure the median, 25th, and 75th percentile prices across exchanges, and normalize all percentiles by the median. We then plot the average of the normalized quantiles across coins affected by listings, for a 10-day window before and after large exchange listings for Binance and Coinbase. There is a nontrivial amount of price dispersion across small exchanges: the p75-p25 spread is around 70bps for coins Binance lists, and around 40bps for coins Coinbase lists. Visually, dispersion declines substantially after listings, by around 20bps and 15bps respectively for Binance and Coinbase.

We then test this prediction in regression form. We first estimate the following flexible specification:

$$Dispersion_{c,t} = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \tag{5}$$

$Dispersion_{c,t}$ is the standard deviation of log prices across all exchanges for coin c at date t. δ_c and η_t are respectively coin and time fixed effects. All other variables are defined above. The coefficients of interest are β_k, which measure how much price dispersion for a given coin decreases after large exchange listing events, compared to coins that did not experience large exchange listings. We cluster standard errors at the coin level. Figure 6 presents the results. Immediately after large exchange listings, dispersion increases dramatically in the first or second day of listings, and shows a gradual decreasing trend thereafter. The effect is statistically significant for Binance, but not for Coinbase. Quantitatively, five days after listing, dispersion decreases by approximately 0.03 and 0.005, or 33% and 6% in percentage terms, for Binance and Coinbase listings, respectively. The general pattern is similar for Binance and Coinbase listings. The coefficients before large exchange listings are small and insignificant in magnitude, so there is little evidence of differential pre-trends in dispersion for coins affected by listings.

We then estimate the following regression specification:

$$\begin{aligned} Dispersion_{c,t} =& \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> 30 \text{ days})_{c,t} + \\ & \beta_3 PreThreedayListing_{c,t} + \delta_c + \eta_t + \epsilon_{c,t} \end{aligned} \tag{6}$$

The results are presented in Table 4. The estimates in Column (1) show that Binance listings lead to a 0.03 decrease in price dispersion, relative to coins not affected by listings. This is equivalent to a 33% reduction in dispersion. Unlike our volume results, the effects on price dispersion persist after 30 days. The average daily normalized price dispersion decreases by 0.04, or 44%. For Coinbase listings in Column (2), we still observe a 0.005 and 0.01 decrease in dispersion, or 6% and 11% in percentage terms, within 30 days and 30 days after Coinbase listings, respectively, although the coefficients are not significant. The estimates of β_3 are insignificant, implying that we do not observe anticipation effects on price dispersion for either Binance or Coinbase listings.

The fact that large exchange listings decrease small exchange price dispersion suggests that a part of the mechanism behind large and small exchange complementarity is arbitrage flows: the entry of large exchanges may cause the overall token market to become more efficient through increased arbitrage trade, which may increase small exchanges' trade volumes and thus their profits from listing the token. In the following section, we will build a model to formalize this intuition, and show how it can explain the stylized facts that we have demonstrated.

4 Model

We construct a model where a single token is traded on an infinitely deep "core" exchange, and a number of "peripheral" exchanges with lower depth. The model allows us to analyze how exchanges' prices are related to each other, and how the core exchange's listing decisions affect the peripheral exchanges' trade volumes and listing decisions. Technically, the model builds on the literature on double-auction models with inventory costs (Vayanos, 1999; Vives, 2010; Du and Zhu, 2017; Chen and Duffie, 2021; Zhang, 2022).

There is a single risky asset, which we will call a "token". There is a core exchange, on which the price of the asset is ψ; the core exchange is infinitely deep, in the sense that there are market makers with infinite capacity, offering to buy or sell arbitrary amounts of the asset at price ψ. We assume ψ has mean μ_ψ and standard deviation σ_ψ. There are also N peripheral exchanges, indexed by j. There are two kinds of market participants on the peripheral exchanges: users, who demand liquidity, and arbitrageurs, who trade against price deviations between the peripheral exchange and the core exchange, subject to inventory costs.

Each peripheral exchange has a unit measure of users with some demand to trade the risky asset, in order to reduce inventory costs. Users are constrained to only trade on exchange j. User i has utility ψ per unit of the risky asset, and suffers inventory costs $\frac{\gamma_j}{2}x^2$ if she holds a net position x in the risky asset. User i begins with $x_{i,0}$ units of the risky asset. This position could be thought of as either a literal inventory position, or more generally as a demand shock for the risky asset; for example, i may receive information that induces her to want to take a long or short position in the risky asset. Hence, i's monetary utility for receiving z net units of the risky asset, thus ending with $x = x_{i,0} + z$ units of the risky asset, is:

$$u_i(z) = \psi(z + x_{i,0}) - \frac{\gamma_j}{2}(z + x_{i,0})^2 \tag{7}$$

Users' inventory position has a systematic and an idiosyncratic component. The standard deviation of $x_{i,0}$ across users on exchange j is $\sigma_{I,j}$. We assume:

$$x_{i,0} = \eta_j + \xi_{ij}$$

where η_j has mean μ_j and standard deviation $\sigma_{A,j}$. We assume η_j is uncorrelated with ψ, and $\eta_j, \eta_{j'}$ are uncorrelated for all peripheral exchanges j, j'. η_j can thus be thought of as an aggregate inventory shock which affects all users on exchange j. We assume exchange j charges a quadratic trading fee to users; if the user trades z units of the asset, she pays a fee $\frac{\tau_j}{2} z^2$ to the exchange. The assumption that trading fees are quadratic simplifies the analysis, but can be relaxed without changing the qualitative results. Since users are atomistic, each user's trades have a negligible effect on overall exchange prices, so users ignore their price impact. If a user purchases z units of the asset at price p_j with position $x_{i,0}$, her total value is thus:

$$V_i = \underbrace{\psi \left(z_i + x_{i,0}\right)}_{\text{Fundamental Value}} - \underbrace{p_j z_i}_{\text{Net Payment}} - \underbrace{\frac{\gamma_j}{2} \left(z_i + x_{i,0}\right)^2}_{\text{Inventory Costs}} - \underbrace{\frac{\tau_j}{2} z_i^2}_{\text{Exchange Fees}} \tag{8}$$

where we have ignored the agent's initial wealth for simplicity, since it only additively shifts V_i and does not affect any decisions agents make. Differentiating, agents i's marginal utility for purchasing an additional unit of the asset is:

$$\frac{\partial V_i}{\partial z_i} = \psi - p_j - \gamma_j \left(z_i + x_{i,0}\right) - \tau_j z_i$$

Setting to 0, agent i's demand for the asset, as a function of the price p_j, is:

$$z_i = \frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - p_j}{\gamma_j + \tau_j} \tag{9}$$

Integrating over all users, aggregate demand from users on exchange j at price p is:

$$Z_{user,j} \left(p_j\right) = \int_{-\infty}^{\infty} z_i \left(x\right) dF_{x_{i,0}} \left(x\right) = \frac{-\gamma_j}{\gamma_j + \tau_j} \eta_j + \frac{\psi - p_j}{\gamma_j + \tau_j} \tag{10}$$

Each peripheral exchange j has a unit measure of atomistic arbitrageurs, who can trade the risky asset on j as well as the core exchange. We will assume arbitrageurs for exchange j cannot trade on other peripheral exchange. Let k index arbitrageurs. Arbitrageurs have utility linear in money. They cannot hold net inventory, so they must buy on the core exchange as much as they sell on the peripheral exchange. Let z_k be the net amount arbitrageur k buys on j and sell on the core exchange. Arbitrageurs face quadratic inventory costs for arbitrage: they incur a cost $\frac{\zeta_j}{2} z_k^2$ for arbitraging k net units of the asset. Arbitrageurs pay the same fees as users: if they trade a quantity z_k, they pay fee $\frac{\tau_j}{2} z_k^2$. Hence, arbitrageurs' value for buying z_k units at price p_j on exchange j, and selling at price ψ on the core exchange, is:

$$V_k \left(z_k\right) = z_k \left(\psi - p_j\right) - \frac{\zeta_j}{2} z_k^2 - \frac{\tau_j}{2} z_k^2$$

Differentiating, arbitrageurs' marginal utility for purchasing an additional unit of the asset is:

$$\frac{\partial V_k}{\partial z_k} = \psi - p_j - \zeta_j z_k - \tau_j z_k$$

Arbitrageur k's demand for the risky asset at price p_j is thus:

$$z_k = \frac{\psi - p_j}{\zeta_j + \tau_j}$$

Integrating demand over the unit measure of arbitrageurs on exchange j, we have:

$$Z_{arb,j}(p_j) = \frac{\psi - p_j}{\zeta_j + \tau_j} \tag{11}$$

Peripheral exchange j's profits, if trade volume is z_i for each user, are: $\int_{-\infty}^{\infty} z_i^2(x)\, dF_{x_{i,0}}(x)$. We assume exchange j has some cost C_j of listing tokens. Hence, exchange j will list the risky asset if it anticipates profits greater than C_j from listing. We will think of the core exchange's listing decisions as exogeneous.

4.1 Equilibrium

In equilibrium, aggregate demand from users and arbitrageurs sums to 0 on each exchange. Thus, adding (10) and (11), market clearing on exchange j requires:

$$Z_{user,j}(p_j) + Z_{arb,j}(p_j) = \left(\frac{-\gamma_j}{\gamma_j + \tau_j}\eta_j + \frac{\psi - p_j}{\gamma_j + \tau_j}\right) + \frac{\psi - p_j}{\zeta_j + \tau_j} = 0$$

The following proposition solves for prices, volumes, and exchange profits when the core exchange does not list the token.

Proposition 1. *When the core exchange does not list the token, the equilibrium price on exchange j is:*

$$p_{j,0}^* = \psi - \gamma_j \eta_j \tag{12}$$

Expected squared trade quantity is:

$$\mathbb{E}\left[z_{i,0}^{*2}\right] = \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \sigma_{I,j}^2 \tag{13}$$

Exchange j's profit from listing the token is:

$$\pi^*_{j,0} = \frac{\tau_j}{2}\left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2 \sigma^2_{I,j} \tag{14}$$

Exchange j lists the token if its cost of listing is lower than (14).

The following proposition solves prices, volumes, and exchange profits when the core exchange does list the token.

Proposition 2. *When the core exchange does list the token, the equilibrium price on exchange j is:*

$$p^*_{j,1} = \psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j \tag{15}$$

Expected squared trade quantity is:

$$\mathbb{E}\left[z^{*2}_{i,1}\right] = \left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2\left[\sigma^2_{I,j} + \left(\frac{\gamma_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2\eta^2_j\right] \tag{16}$$

Exchange j's profit from listing the token is:

$$\pi^*_{j,1} = \frac{\tau_j}{2}\left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2\left[\sigma^2_{I,j} + \left(\frac{\gamma_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2\eta^2_j\right] \tag{17}$$

Exchange j lists the token if its cost of listing is lower than (17).

The intuitions behind Propositions 1 and 2 are as follows.

Prices. Expression (12) states that the price on exchange j, in the absence of the core exchange, is the "efficient price" ψ, minus the aggregate inventory shock η_j times users' inventory cost γ_j. If the aggregate component of inventory shocks η_j is positive, there is no other exchange for users to sell their endowments to; exchange j's price must then be higher than ψ in order to clear the market. The gap between exchange j's price and ψ depends on η_j, and users' cost of holding inventory, γ_j.

When the core exchange lists the token, arbitrageurs trade against this price gap, buy from the CEX at price ψ and selling to the peripheral exchange. Arbitrage cannot perfectly close the gap, because arbitrageurs also face transaction fees and inventory costs. Comparing (12) and (15), arbitrageurs decrease the effect of inventory shocks on prices by a factor:

$$\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j} \tag{18}$$

When trading costs τ_j are low, and arbitrageurs' inventory costs ζ_j are low relative to users' costs γ_j, prices will tend to converge towards ψ significantly when the core exchange lists.

Trade volumes and exchange profits. In the absence of a core exchange, trade on peripheral exchanges is generated only by the idiosyncratic component of inventory shocks: (13) states that volume depends on the variance of users' endowments $\sigma_{I,j}^2$, as well as a factor which reflects how large transaction fees τ_j are relative to users' inventory costs γ_j. When a core exchange enters, trade is generated by both the idiosyncratic and aggregate components of inventory shocks, since arbitrageurs can buy on the core exchange and sell to the peripheral exchange. (16) shows that expected squared trade volume can be cleanly decomposed into an the sum of (13), and an extra term reflecting the aggregate shock η_j, and the multiplier (18) capturing how active arbitrageurs are. Thus, expected squared trade volume of peripheral exchanges is strictly higher when the core exchange lists the token. Since profits are proportional to squared trade volume, peripheral exchanges' profits are also higher when the core exchange lists.

Next, using these propositions, we derive a number of predictions to bring to the data.

4.2 Comparative Statics and Predictions

We now describe how our model rationalizes the three stylized facts we have shown.

Prediction 1. *Consider all peripheral exchanges j which list a given token before the core exchange does. These exchanges will experience increases in trading volume for the token, when the core exchange lists the token.*

Prediction 1 corresponds to Fact 1. This prediction follows directly from comparing (13) and (16), and the intuition that the aggregate component of inventory shocks also contributes to trade volume after the CEX lists the token.

Prediction 2. *Consider all peripheral exchanges j which list a given token before the core exchange does. The volatility of token prices on these exchanges will fall after the core exchange lists the token. The cross-sectional dispersion of prices across these exchanges will also fall after the core exchange lists the token.*

Prediction 2 corresponds to Fact 3. This prediction follows because, from (12), the variance of j's prices when the core exchange does not list is

$$\sigma_\psi^2 + \gamma_j^2 \sigma_{A,j}^2$$

Whereas the variance when the core exchange lists is, from (15), the smaller quantity:

$$\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2$$

Intuitively, arbitrage with the core exchange decreases the effect of inventory shocks on peripheral exchanges' prices, limiting volatility. The prediction about dispersion follows similarly. Suppose for simplicity that exchanges are symmetric, so $\sigma_{A,j}^2 = \sigma_A^2$ for all exchanges. The dispersion of peripheral exchange prices around ψ is $\gamma_j^2 \sigma_A^2$ without the core exchange, and the lower quantity

$$\left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_A^2$$

with the core exchange. Again, arbitrage with the core exchange causes peripheral exchange prices to cluster more closely around ψ.

The next prediction is about the "listing following" effect: peripheral exchanges who have not yet listed a given token will have a stronger incentive to list, after the core exchange has listed the token.

Prediction 3. *Listings will tend to follow the core exchange: when the core exchange lists the token, some peripheral exchanges which previously did not list the token will choose to list the token. Formally, peripheral exchanges' profit with the core exchange, (17), is greater than peripheral exchanges' profit without the core exchange, (14), so the set of peripheral exchanges which lists the token is strictly larger after the core exchange enters.*

Prediction 3 corresponds to Fact 2. This prediction follows from (14) and (17). When the core exchange enters, expected profits on all peripheral exchanges increase. Thus, once the core exchange lists the token, all peripheral exchanges which have already listed have no incentive to unlist, even if the listing decision is fully reversible and the cost can be recovered. Moreover, some exchanges which previously did not list the token will find it profitable to list the token. This prediction essentially implies that the core exchange is a complement to peripheral exchanges; in particular, this prediction contrasts with standard models of imperfect competition, in which the entry of a large competitor should cannibalize smaller competitors, and decrease incentives for entry.

Next, we use the model to derive two additional predictions, which should hold in the data if our model described the mechanism at work in the data. The first prediction concerns the structure of prices correlations across exchanges; we first derive expressions for these correlations.

Proposition 3. *The coefficient of determination R^2 between the core exchange's price, and peripheral exchange j's price, is:*

$$R^2_{j,CE} = \frac{Cov^2\left(p^*_j, \psi\right)}{Var\left(p^*_j\right)Var\left(\psi\right)} = \frac{\sigma^2_\psi}{\sigma^2_\psi + \left(\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2 \gamma^2_j \sigma^2_{A,j}} \tag{19}$$

The R^2 between the prices of exchanges j and j' is:

$$\begin{aligned} R^2_{j,j'} &= \frac{Cov^2\left(p^*_j, p^*_{j'}\right)}{Var\left(p^*_j\right)Var\left(p^*_{j'}\right)} \\ &= \frac{\sigma^2_\psi}{\left[\sigma^2_\psi + \left(\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2 \gamma^2_j \sigma^2_{A,j}\right]} \frac{\sigma^2_\psi}{\left[\sigma^2_\psi + \left(\frac{\zeta_{j'}+\tau_{j'}}{\gamma_{j'}+\zeta_{j'}+2\tau_{j'}}\right)^2 \gamma^2_j \sigma^2_{A,j'}\right]} \end{aligned} \tag{20}$$

Prediction 4. *We always have:*

$$R^2_{j,CE} \geq R^2_{j,j'}$$

That is, the correlation between the core exchange price and the price on any peripheral exchange j is stronger than the correlation between the prices on peripheral exchanges j and j'.

In words, Prediction 4 states that the structure of price correlations between exchanges inherits the core-periphery structure of the exchange network: peripheral exchanges' prices are more correlated with the core exchange than they are with each other. This is because, from expression (15), each peripheral exchange's price is equal to the core exchange's price, plus an error term reflecting aggregate inventory shocks on the peripheral exchange which are imperfectly eliminated by arbitrageurs. Thus, $R^2_{j,CE}$ reflects the correlation of the core exchange price ψ, with a price which is ψ plus a noise term, whereas $R^2_{j,j'}$ reflects the correlation between two prices which are each equal to ψ plus a noise term.

An additional prediction is that, if the effects we observe are truly driven by arbitrage flows, then price correlations, volume effects of listings, and the "listing following" effect should be associated across exchanges.

Prediction 5. *When peripheral exchanges differ mainly in their arbitrage costs ζ_j, peripheral exchanges whose prices are more correlated with the core exchange should tend to a larger increase in trade volumes when core exchanges list, and should have a stronger tendency to* *follow the core exchange's listing decisions. Formally, we have:*

$$\frac{\partial R^2_{j,CE}}{\partial \zeta_j} = \frac{-2\sigma_\psi^2 \sigma_{A,j}^2 \gamma_j^2 \left(\zeta_j + \tau_j\right)\left(\gamma_j + \tau_j\right)}{\left[\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2\right]^2 \left(\gamma_j + \zeta_j + 2\tau_j\right)^3} \leq 0$$

$$\frac{\partial \Delta \mathbb{E}\left[z_{i1}^{*2}\right]}{\partial \zeta_j} = -\frac{2(\gamma_j + \tau_j)^2}{(\gamma_j + \zeta_j + 2\tau_j)^3}\left(\frac{\eta_j}{\sigma_{I,j}}\right)^2 \leq 0$$

$$\frac{\partial \Delta \pi_j^*}{\partial \zeta_j} = -\frac{2(\gamma_j + \tau_j)^2}{(\gamma_j + \zeta_j + 2\tau_j)^3}\left(\frac{\eta_j}{\sigma_{I,j}}\right)^2 \leq 0$$

Prediction 5 follows if there are differences in how "connected" peripheral exchanges are to the core exchange, which in our model corresponds to the arbitrageur inventory cost parameter ζ_j. For a peripheral exchange with lower ζ_j, prices will tend to be more correlated with the core exchange; the core exchange's listings will tend to increase volumes more; and the core exchange's listing decisions will increase the peripheral exchange's profits more, implying that the peripheral exchange has a stronger incentive to "follow" the core exchange's listing decisions. If Prediction 5 holds in the data, this indicates empirically that three separate phenomena – price correlations, volume increases, and listing following – are statistically associated, increasing our confidence that they are driven by the same underlying economic phenomenon.

4.3 Discussion of Assumptions

Our baseline model assumes a simple model in which users are tied to a single peripheral exchange, and cannot move across exchanges. If users were able to move across peripheral exchanges and the core exchange, perhaps at some cost, this would cause exchanges to become partially substitutes for each other; one exchange's listing decision could potentially cannibalize volume from other exchanges, as users move to the exchange which has newly listed the token. This force would tend to push against the effects that we find, causing exchanges to tend to be substitutes instead of complements. If the user substitution force were strong enough, listings could tend to decrease trade volumes, and the core exchange's decision to list may cannibalize enough volume that it induces peripheral exchanges to unlist. This runs counter to the evidence we find empirically. We thus assume away this effect for expositional simplicity.

Our baseline model also does not feature the "listing pump" effect, that core exchange listings tend to associate with increased token prices, which is emphasized in a number of academic and industry studies(Ante, 2019; Lemmen, 2022; Talamas, 2021). Since the main focus of our paper is to analyze the structure of competition between exchanges, we do not discuss the listing pump effect in detail. However, there are a number of ways to derive the listing pump effect in the context of our model. One approach would be to assume that aggregate inventory shocks η_j have positive means; that is, users on peripheral exchanges have a net positive endowment of the asset. Inventory costs then tend to depress prices on peripheral exchanges, and the entry of the core exchange will tend to alleviate this negative price pressure and raise token prices. The listing pump effect could also be microfounded from a richer multi-period model, in which the entry of the core exchange increases market depth and decreases volatility of the token, thus pushing prices upwards through a "liquidity premium" effect.

A number of other assumptions are made largely for expositional simplicity. We assume the core exchange has infinite depth; it is sufficient for our effects that the core exchange's depth is finite, but much greater than peripheral exchanges' depth. We assume arbitrageurs can only trade the peripheral exchange against the core exchange; it is sufficient that the cost of doing this is lower than the cost of arbitraging two peripheral exchanges against each other. In our conversations with practitioners, most market makers in practice appear to trade smaller exchanges with larger core exchanges. One reason for this is that being a market maker on an exchange often involves direct negotiations with the exchange for special access, and it is potentially difficult to enter into multiple of these agreements at once. We assume there is a single core exchange; in practice, there are a number of bigger exchanges which likely behave more like core exchanges, and smaller exchanges which behave more like peripheral exchanges. In our empirical analysis, we treat Binance and Coinbase as core exchanges, and the long tail of smaller exchanges as peripheral.

5 Empirical Tests

Our model shows that the listing effects by large exchanges are essentially caused by their core position. We proceed to test Predictions 4 and 5 empirically.

5.1 Core-Periphery Structure of Price Correlations

Prediction 4 of our model states that price correlations should have a core-periphery structure: shocks to prices on any peripheral exchange are only transmitted to other exchanges through the core exchange, so the correlations between peripheral exchange prices and core exchange prices should be greater than the correlations between different peripheral exchange prices.

To test this prediction, we calculate the return correlations of Bitcoin between all pairs of exchanges. For each exchange pair, we calculate return correlations using the entire time period where we have coverage for both exchanges in the pair.

Figure 7 plots the probability density function (PDF) of return correlations, for all exchange pairs excluding pairs involving Binance or Coinbase, and exchange pairs between peripheral exchanges and either Binance or Coinbase. As the model predicts, price correlations between a peripheral exchange and a core exchange are in general stronger than price correlations between two peripheral exchanges. Quantitatively, the 25th, 50th, and 75th percentiles of price correlations among all exchanges are 0.51, 0.72, and 0.88. The percentiles are 0.70, 0.84, and 0.93 for pairs involving Binance, and 0.71, 0.85, and 0.94 for pairs involving Coinbase.

5.2 Interaction between Correlation Structure and Volume Increases/Listing Following

Prediction 5 posits that peripheral exchanges with higher price correlations to the core exchange will see greater volume increases and exhibit a stronger tendency to follow the core exchange's listing decisions. We assess each of these sub-predictions separately.

5.2.1 Correlation Structure and Volume Increases

To examine whether peripheral exchanges with stronger price correlations with the core exchange experience larger volume increases when the core exchange lists, we estimate the following variant of Specification (2):

$$\begin{aligned} Log(Volume)_{c,e,t} = & \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(\text{0-30 days})_{c,t} \times Correlation_e + \\ & \beta_3 Listing(> \text{30 days})_{c,t} + \beta_4 Listing(> \text{30 days})_{c,t} \times Correlation_e + \\ & \beta_5 PreThreedayListing_{c,t} + \beta_6 PreThreedayListing_{c,t} \times Correlation_e + \\ & \beta_7 Correlation_e + \delta_{c,e} + \eta_t + \epsilon_{c,e,t} \end{aligned} \tag{21}$$

All variables are defined as in previous specifications. $Correlation_e$ is the correlation between Bitcoin returns between exchange e and a given core exchange using the entire time period where we have coverage for both exchanges in the pair. The coefficients of interest are β_2 and β_4: when these are coefficients are positive, exchanges with greater correlations with the core exchange will tend to have larger volume increases within 30 days and 30 days after core exchange listings.

The results are shown in Table 5. We find that β_2 and β_4 are positive and significant, confirming that exchanges with stronger price correlations with the core exchange experience larger volume increases when the core exchange lists. In Columns (4) and (8), our preferred specifications, exchanges with an additional 0.01 return correlation with Binance and Coinbase experience an additional 2.6% and 2.7% increase in daily average volume within 30 days of their listings. The effects are 3.3% and 3.6% after 30 days of Binance and Coinbase listing, respectively.

5.2.2 Correlation Structure and Listing Following

Next, we show that peripheral exchanges which have stronger price correlations with core exchanges also tend to follow core exchanges' listing decisions more closely. To measure the propensity for a given peripheral exchange to follow a core exchange's listing decisions, we define Listing Following Probability$_e$, the listing following probability to a core exchange for peripheral exchange e, as the number of listings within 30 days of core exchange listings over the number of listings within the [-100, 100] days window:

$$\text{Listing Following Probability} = \frac{\#\text{ Listings within [0,30] days window}_e}{\#\text{ Listings within [-100,100] days window}_e} \tag{22}$$

Intuitively, Listing Following Probability$_e$ measures the tendency for exchange e to list very quickly (within 30 days) after a core exchange lists. When Listing Following Probability$_e$ is 1, in all cases where e lists a token in a 100 day window of a core exchange's listing, the listing occurs from 0 to 30 days after the core exchange's listing.

Using these two measures, we then estimate whether exchanges with higher price correlations with a core exchange also have higher listing following probabilities. We estimate the specification:

$$ListingFollowingProbability_e = \alpha + \beta Correlation_e + \epsilon_e \tag{23}$$

where e indexes exchanges. As above, $Correlation_e$ is the correlation in BTC prices between exchange e and the core exchange. We expect a positive β, suggesting that peripheral exchanges with higher price correlations with a given core exchange also tends to follow core exchanges' listing decisions.

As shown in Figure 8, $Correlation_e$ and $ListingFollowingProbability_e$ are positively correlated, whether we treat Binance or Coinbase as the core exchange: that is, exchanges with stronger positive price correlations with a given core exchange tend to follow the core exchange's listing decisions more closely. In regressions, both β coefficients are both significant at the 95% level. Quantitatively, exchanges with an 0.01 higher BTC return correlation with a core exchange has a 0.19% (0.22%) higher probability of following the core exchange's listing decisions, respectively, for Binance (Coinbase).

5.3 Volume Decomposition

Finally, we conduct a decomposition analysis to quantitatively compare the different channels through which core exchange listings affect trade volume. We estimate three coin-level DID regressions:

$$\begin{aligned}
\log(Volume_{c,t}[incumbent]) &= \sum_{k=-11}^{11} \left(\beta_k^{inc}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \\
\log(Volume_{c,t}[incumbent + entrant]) &= \sum_{k=-11}^{11} \left(\beta_k^{inc} + \beta_k^{ent}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \\
\log(Volume_{c,t}[incumbent + entrant + central]) &= \sum_{k=-11}^{11} \left(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}
\end{aligned} \tag{24}$$

Where, as in (3), $treat_{c,k,t}$ are dummies for time since a core exchange has listed a token. We construct $\log(Volume_{c,t})$ in three successively broader ways: using only volume on incumbent exchanges, which list token c at least 10 days before core exchange listings; using incumbents as well as any new exchanges which enter following the core exchange's entry, excluding volume on the core exchange itself; and using volume on incumbents, entrants, and the core exchange itself. We then estimate the coefficients β_k^{inc}, $(\beta_k^{inc} + \beta_k^{ent})$, and $(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen})$ through separate DID estimates of regression equations using these three different definitions of volume. Using these successively broader volume definitions, we can decompose the extent to which volume increases associated with core exchange entry are generated by volume increases on incumbents, as in Fact 1; the entry of new exchanges which follow the core exchange, as in Fact 2; and the direct effect of trade volumes generated by the core exchange. The results are shown in Figure 9. The estimated coefficients are larger for broader definitions of volume.

We can then recover β_k^{inc}, β_k^{ent}, and β_k^{cen} simply by taking differences between the estimated coefficients. These estimates allow us to do a simple accounting decomposition of each $(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen})$, the total volume increase induced by core exchange listing after k days, into three components: β_k^{cen}, the direct effect of trade volume on the core exchange; β_k^{inc}, the effect on incumbent volume; and β_k^{ent}, the effect on entrant volume. For example, for Coinbase 0 days after a listing event, we estimate β_k^{inc}, β_k^{ent}, and β_k^{cen} to be of 1.38, 0.31, and

0.34 respectively. Exponentiating the sum and substract one, $\exp\left(\beta_k^{inc} + \beta_k^{cen} + \beta_k^{ent}\right) - 1$, Coinbase's listing associates with a total volume increase of 661pp, relative to pre-listing incumbent volume. We infer that this effect can be thought of as the net result of a 297pp volume increase among incumbents, a 36pp increase in volume from entrants, and a 40pp increase directly on the core exchange, where the latter two numbers are percentages as a fraction of pre-listing incumbent volume.[10]

We show our estimates of β_k^{inc}, β_k^{cen}, and β_k^{ent} in the bottom row of Figure 9. Taking the exponent of the blue line on the top row, Binance listings are associated with an average 1046pp increase in total volume within first 10 days, and a 266pp increase after ten days. The numbers for Coinbase are 388pp and 64pp after 10 days. From the second row, the short-run effect for Binance decomposes into a 87pp increase of incumbent volume, a 60pp increase associated with entrants, and a 295pp increase on Binance itself. 10 days after listing, however, the incumbent volume effect decreases to -68pp, the entrant effect increases to 538pp, and the effect on Binance volume decreases slightly to 125pp. Coinbase listings are associated with an average 388pp increase in total volumes within first 10 days, and a 64pp increase after 10 days. Within first 10 days of listing, incumbent volumes increase by 103pp, compared to 54pp for entrants, and 60pp for Coinbase itself. After 10 days, the incumbent effect declines to -56pp, the entrant effect increases to 190pp, whereas the effects on Coinbase are stable at roughly 27pp.

As another way to view these results, we calculate the ratio:

$$LR_k = \frac{\exp\left(\beta_k^{inc} + \beta_k^{ent}\right) - 1}{\exp\left(\beta_k^{cen}\right) - 1} \tag{25}$$

The ratio (25) can be interpreted as saying, for each dollar of trade volume a core exchange creates by listing a token, how many dollars of trade volume are generated on other (incumbent and entrant) exchanges. Within first 10 days of listing, we estimate that average LR_k is equal to 0.67 for Binance, and 3.55 for Coinbase: thus, every dollar captured by Binance listing generates 0.67 cents of trade on other exchanges, and every dollar captured by Coinbase generates \$3.55 dollars of trade volume on other exchanges. The coefficient for Binance is smaller largely because the denominator is larger: both exchanges generate similar spillover effects from listing, but Binance generates a larger direct increase in volume. Thus, for both core exchanges, a large fraction of the volume increases associated with listings are due to the indirect effects on other exchanges.

[10] To calculate these percentages, we exponentiate each β and subtract one.

6 Conclusion

In this paper, we argued that crypto exchanges do not compete like firm in oligopolistic product markets. Instead, small "peripheral" exchanges appear to be complementary to large "core" exchanges. Peripheral exchanges experience increased trading volume after core exchanges list, and tend to follow the listing predictions of core exchanges. Our results suggest that the large core crypto exchanges are potentially systemically important players in crypto markets, in a way that is understated by their modest shares of overall trading volume. The decision of a core exchange to list a token induces large indirect trading volume increases, both on incumbent exchanges which have previously listed the token, and on entrant exchanges which follow the core exchange's listing decisions. Exchanges currently make their listing decisions with substantial discretion and little oversight; policymakers may wish to monitor or regulate core exchanges' listing decisions, given the large effects these decisions have on crypto market outcomes.

References

Amiram, Dan, Evgeny Lyandres, and Daniel Rabetti. 2022. "Cooking the Order Books: Information Manipulation and Competition among Crypto Exchanges." *Available at SSRN.*

Ante, Lennart. 2019. "Market reaction to exchange listings of cryptocurrencies." *Blockchain Research Lab Working Paper Series.*

Augustin, Patrick, Alexey Rubtsov, and Donghwa Shin. 2020. "The impact of derivatives on spot markets: Evidence from the introduction of bitcoin futures contracts."

Augustin, Patrick, Roy Chen-Zhang, and Donghwa Shin. 2022. "Reaching for Yield in Decentralized Financial Markets." *Available at SSRN 4063228.*

Budish, Eric, Robin S Lee, and John J Shim. 2019. "A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?" National Bureau of Economic Research.

Chan, Qing, Wenzhi Ding, Chen Lin, and Alberto G Rossi. 2020. "An inside look into cryptocurrency exchanges." *Available at SSRN 3759062.*

Chen, Daniel, and Darrell Duffie. 2021. "Market fragmentation." *American Economic Review*, 111(7): 2247–74.

Choi, Kyoung Jin, Alfred Lehar, and Ryan Stauffer. 2022. "Bitcoin microstructure and the kimchi premium." *Available at SSRN 3189051.*

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

Cong, Lin William, Xi Li, Ke Tang, and Yang Yang. 2020. "Crypto wash trading." *Available at SSRN 3530220.*

Cong, Lin William, Ye Li, and Neng Wang. 2019. "Tokenomics: Dynamic adoption and valuation." *Working paper, University of Chicago.*

Du, Songzi, and Haoxiang Zhu. 2017. "What is the optimal trading frequency in financial markets?" *The Review of Economic Studies*, 84(4): 1606–1651.

Harvey, Campbell R, Ashwin Ramachandran, and Joey Santoro. 2021. *DeFi and the Future of Finance.* John Wiley & Sons.

John, Kose, Leonid Kogan, and Fahad Saleh. 2022. "Smart Contracts and Decentralized Finance." *Available at SSRN.*

Kogan, Shimon, Igor Makarov, Marina Niessner, and Antoinette Schoar. 2023. "Are cryptos different? evidence from retail trading." *Evidence from Retail Trading (March 20, 2023).*

Lemmen, Jan. 2022. "Cross-listings: Cryptocurrency Characteristics and Abnormal Returns." *Master Thesis.*

Li, Tao, Donghwa Shin, and Baolian Wang. 2021. "Cryptocurrency pump-and-dump schemes." *Available at SSRN 3267041.*

Liu, Yukun, Aleh Tsyvinski, and Xi Wu. 2022. "Common risk factors in cryptocurrency." *The Journal of Finance*, 77(2): 1133–1177.

Liu, Yukun, and Aleh Tsyvinski. 2018. "Risks and returns of cryptocurrency." *Working paper, Yale University and NBER.*

Liu, Yukun, Jinfei Sheng, and Wanyi Wang. 2021. "Technology and cryptocurrency valuation: Evidence from machine learning." *Available at SSRN 3577208.*

Makarov, Igor, and Antoinette Schoar. 2019. "Price discovery in cryptocurrency markets." Vol. 109, 97–99.

Makarov, Igor, and Antoinette Schoar. 2020. "Trading and arbitrage in cryptocurrency markets." *Journal of Financial Economics*, 135(2): 293–319.

Makarov, Igor, and Antoinette Schoar. 2022. "Cryptocurrencies and decentralized finance (DeFi)." National Bureau of Economic Research.

Talamas, Roberto. 2021. "Analyzing the Crypto Exchange Pump Phenomenon." Messari.

Vayanos, Dimitri. 1999. "Strategic trading and welfare in a dynamic market." *The Review of Economic Studies*, 66(2): 219–254.

Vives, Xavier. 2010. *Information and learning in markets: the impact of market microstructure.* Princeton University Press.

Yu, Gloria Yang, and Jinyuan Zhang. 2022. "Flight to bitcoin." *Available at SSRN 3278469.*

Zhang, Anthony Lee. 2022. "Competition and manipulation in derivative contract markets." *Journal of Financial Economics*, 144(2): 396–413.

Figure 1: Market Share of the Top 2 Exchanges for BTC Rrading over time

This figure shows the market share of total BTC trading volume, for the 2 largest exchanges as of 2022. These exchanges are Binance and Coinbase. Data source: Cryptotick.

Figure 2: Large Exchange Listings and Trade Volumes: DID Estimates

The figure depicts estimates from Specification (1):

$$\log(Volume_{c,e,t}) = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_{c,e} + \eta_t + \epsilon_{c,e,t}$$

along with 95% confidence intervals. The outcome variable is log coin-exchange-day level logarithmic dollarized volume. $\delta_{c,e}$ represents coin-exchange fixed effects, and η_t represents day fixed effects. Observations exactly 10 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin-exchange pair level. Data source: Cryptotick.

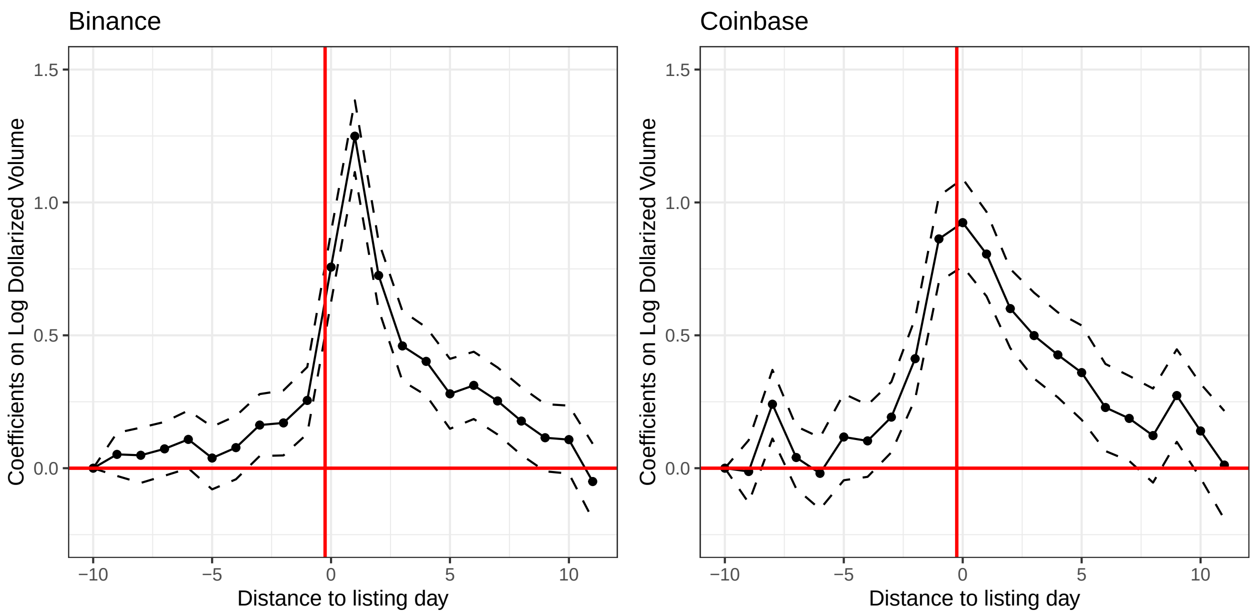

Figure 3: Listing Times of Small Exchanges Relative to Large Exchanges

This figure shows the listing following pattern of all exchanges relative to the large exchange's listing. The x-axis denotes the time interval between an exchange's listing date and a large exchange's listing date for the same coin. The red vertical line indicates zero time interval with the large exchange listing. The y-axis is the mass of each time interval bar. Data source: Cryptotick.

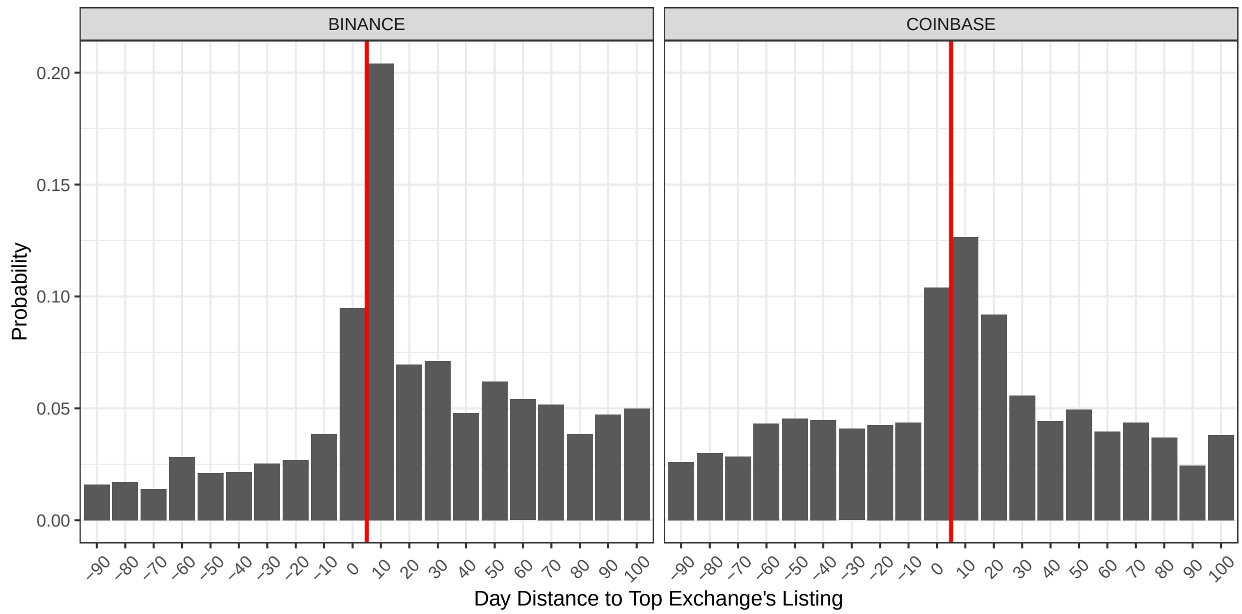

Figure 4: Large Exchange Listings and Small Exchange Listings

The figure depicts estimates from Specification (3):

$$\Delta\#Exchanges_{c,t} = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

along with 95% confidence intervals. The outcome variable is the number of exchanges that list coin c betweem day $t-1$ and t. δ_c represents coin fixed effects, and η_t represents day fixed effects. Observations exactly 10 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin level. Data source: Cryptotick.

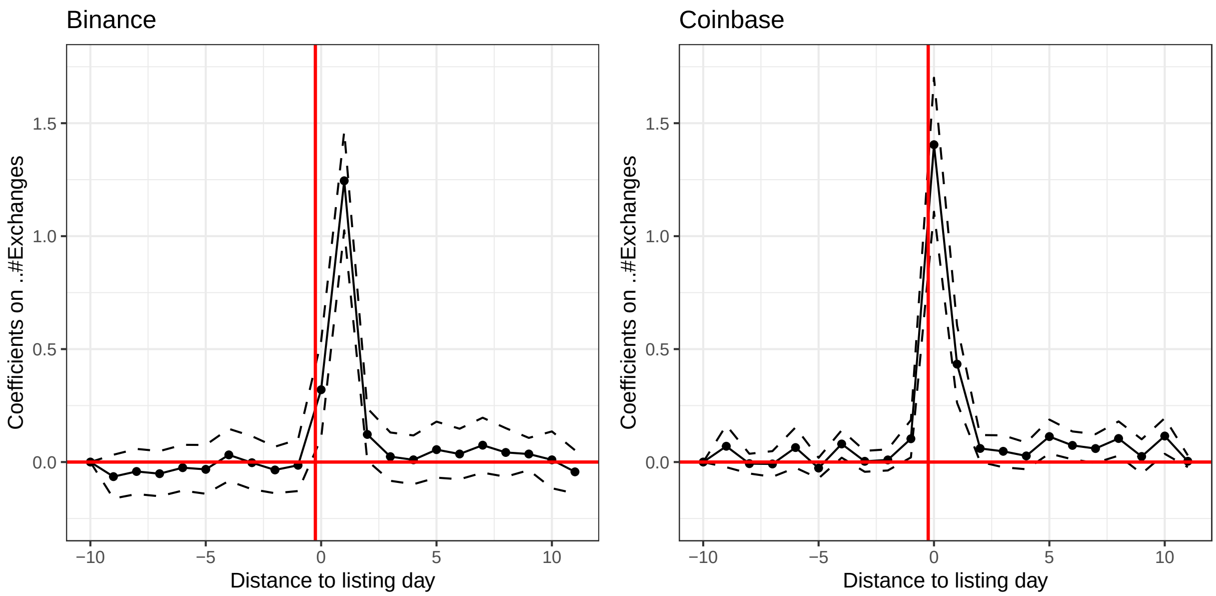

Figure 5: Large Exchange Listings and Price Dispersion

This figure illustrates how the interquartile range of token prices changes around listing events. For each coin affected by listings, we calculate the 25th and 75th percentiles of coin prices across exchanges, normalized by the median price. The figure shows the average of the normalized quantiles across coins affected by listings, for Binance and Coinbase respectively. In total, we observe 281 Listings for Binance, and 49 for Coinbase. The x-axis denotes the time interval between an exchange's listing date and a central exchange's listing date for the same coin. The red vertical line indicates the listing time. Data source: Cryptotick.

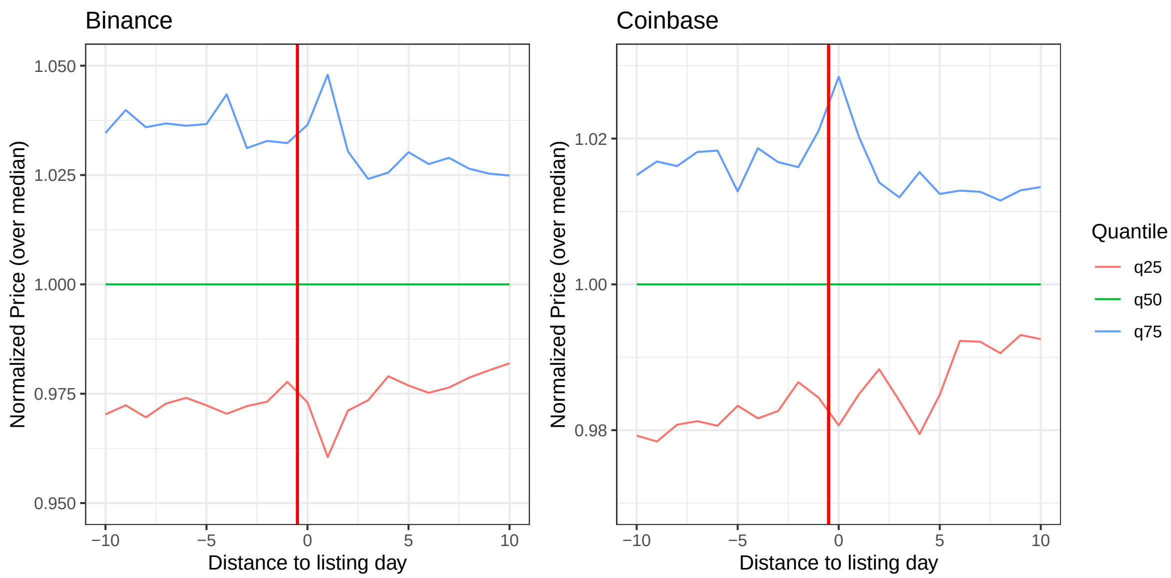

Figure 6: Large Exchange Listings and Price Dispersion: DID Estimates

The figure depicts estimates from Specification (5):

$$Dispersion_{c,t} = \sum_{k=-11}^{11} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

along with 95% confidence intervals. The outcome variable is coin-day level price dispersion, calculated as the standard deviation of log prices across exchanges for coin c at day t. δ_c represents coin fixed effects, and η_t represents day fixed effects. Observations exactly 10 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin level. Data source: Cryptotick.

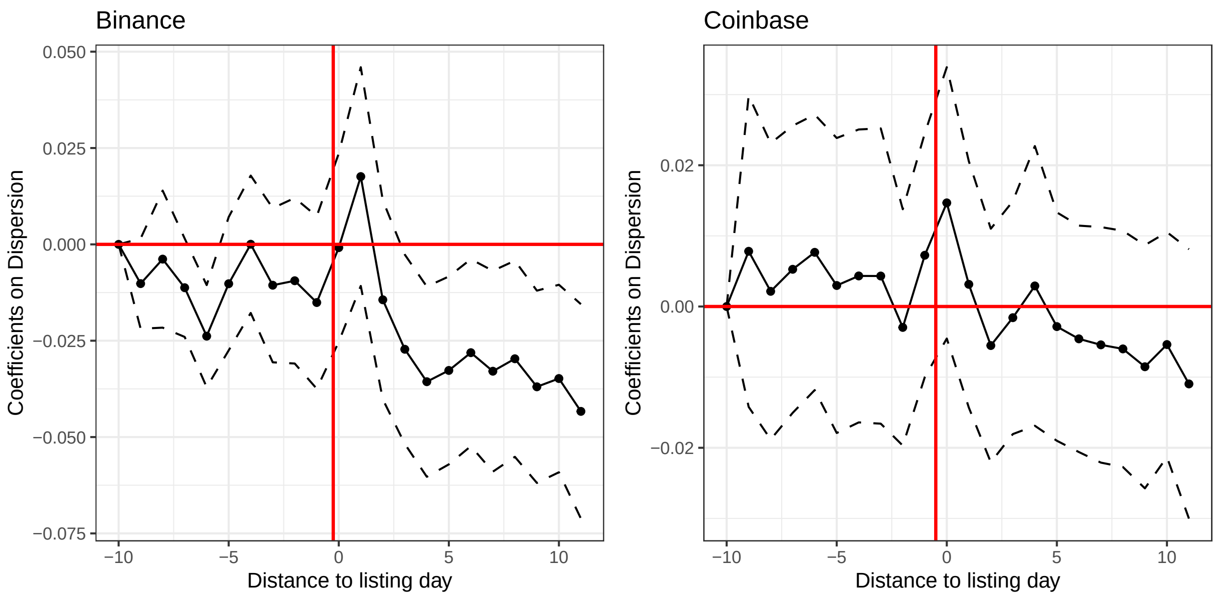

Figure 7: Return correlation of Bitcoin at the Exchange Pair Level

This figure shows the distribution of the pairwise correlations of BTC returns, for all exchange pairs excluding pairs involving Binance or Coinbase, and exchange pairs between peripheral exchanges and either Binance or Coinbase. For each exchange pair, we calculate return correlations using the entire time period where we have coverage for both exchanges in the pair. Data source: Cryptotick.

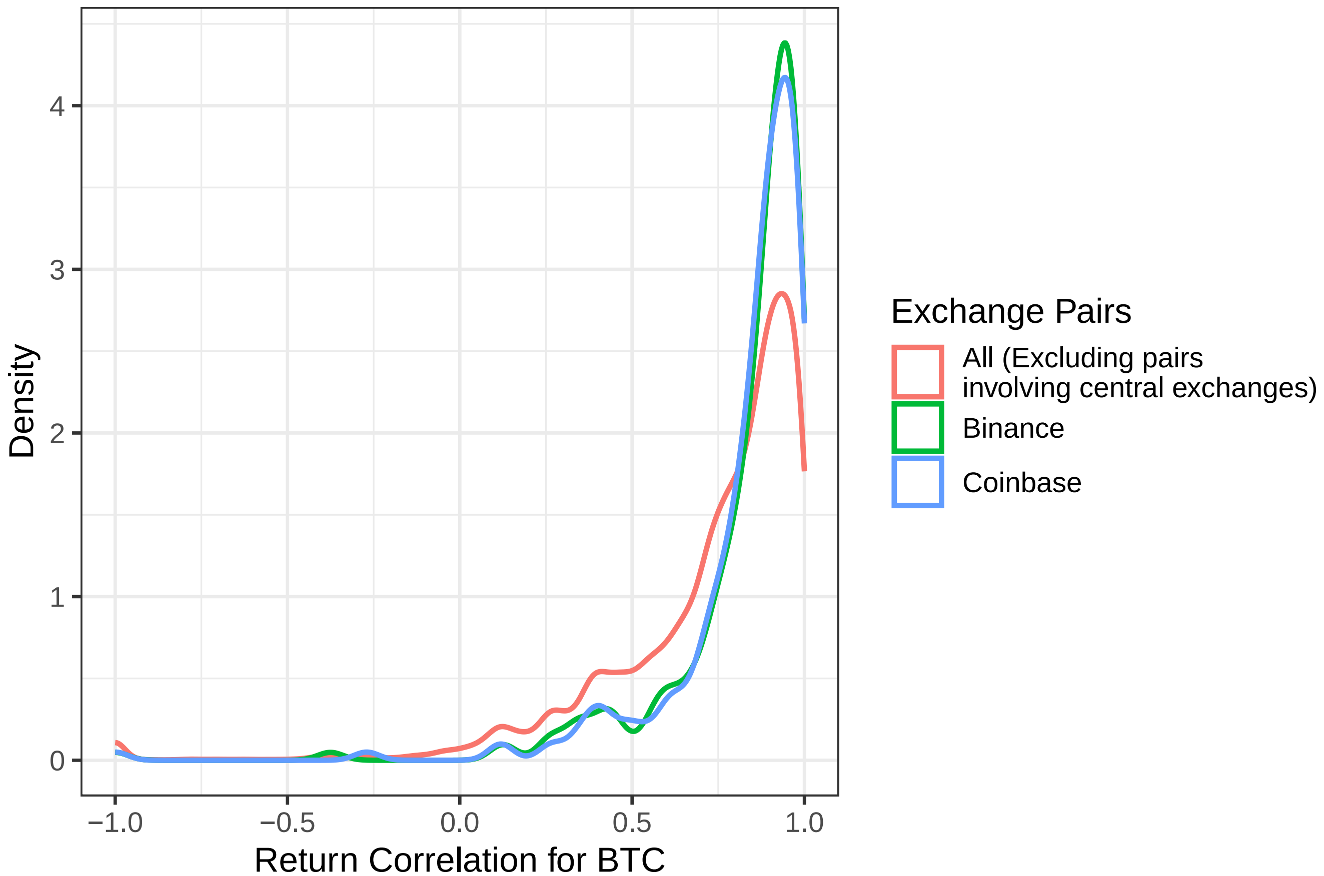

Figure 8: Correlation-listing Correlation for Central exchanges

This figure displays the relationship between BTC return correlation and listing following probability for all exchanges relative to central exchanges. Each data point represents an exchange. The x-axis denotes the return correlation of Bitcoin between an exchange and a central exchange using the entire time period where we have coverage for both exchanges in the pair. The y-axis denotes the listing following probability between an exchange and a central exchange, defined in Equation (22). The red line indicates the fitted linear regression curve. The correlation coefficient and its p-value are reported. For central exchanges, we focus on Binance and Coinbase. Data source: Cryptotick.

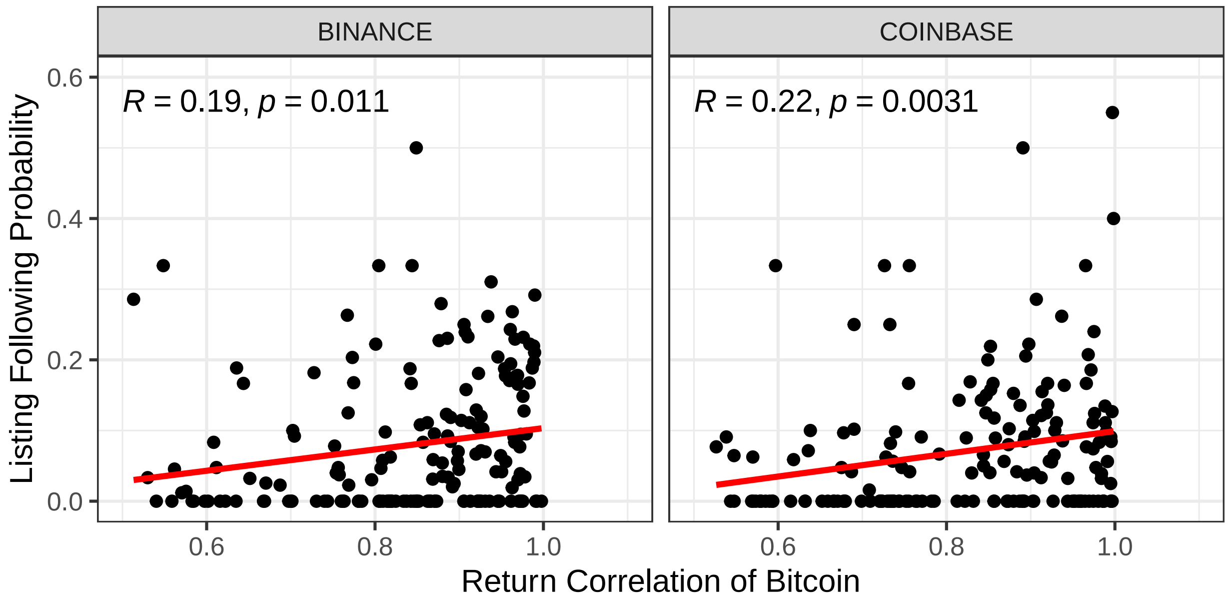

Figure 9: Correlation-listing Correlation for Central exchanges

The top row of the figure depicts estimates from Specification (24):

$$\log(Volume_{c,t}[incumbent]) = \sum_{k=-11}^{11} \left(\beta_k^{inc}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

$$\log(Volume_{c,t}[incumbent + entrant]) = \sum_{k=-11}^{11} \left(\beta_k^{inc} + \beta_k^{ent}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

$$\log(Volume_{c,t}[incumbent + entrant + central]) = \sum_{k=-11}^{11} \left(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen}\right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

along with 95% confidence intervals. We categorize exchange-coin pairs into three groups according to their listing time and central status: incumbent (exchanges that list at least 10 days before central exchange listings or coins that are not listed by central exchanges), central (central exchange), and entrant (all remaining exchanges). The outcome variable is coin-day level logarithmic dollarized volume, computed by filtering the relevant exchange-coin pairs, adding up the dollarized volume, and taking the logarithm. The top row shows the above estimation. The red line shows estimates of β_k^{inc}, using only incumbent volume as the dependent variable. The green line shows estimates of $(\beta_k^{inc}+\beta_k^{ent})$, using incumbent plus entrant volume as the dependent variable. The blue line shows estimates of $(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen})$, using incumbent, entrant, and central exchange volume as the dependent variable. δ_c represents coin fixed effects, and η_t represents day fixed effects. Observations exactly 10 days before central exchange listings are set as the reference group. Standard errors are clustered at the coin level. We then recover $\beta_k^{inc}, \beta_k^{ent}$, and β_k^{cen} simply by taking differences between the estimated coefficients. The bottom row shows our estimate. The red line shows estimates of β_k^{inc}, the green line shows estimates of β_k^{ent}, and the blue line shows estimates of β_k^{cen}. Data source: Cryptotick.

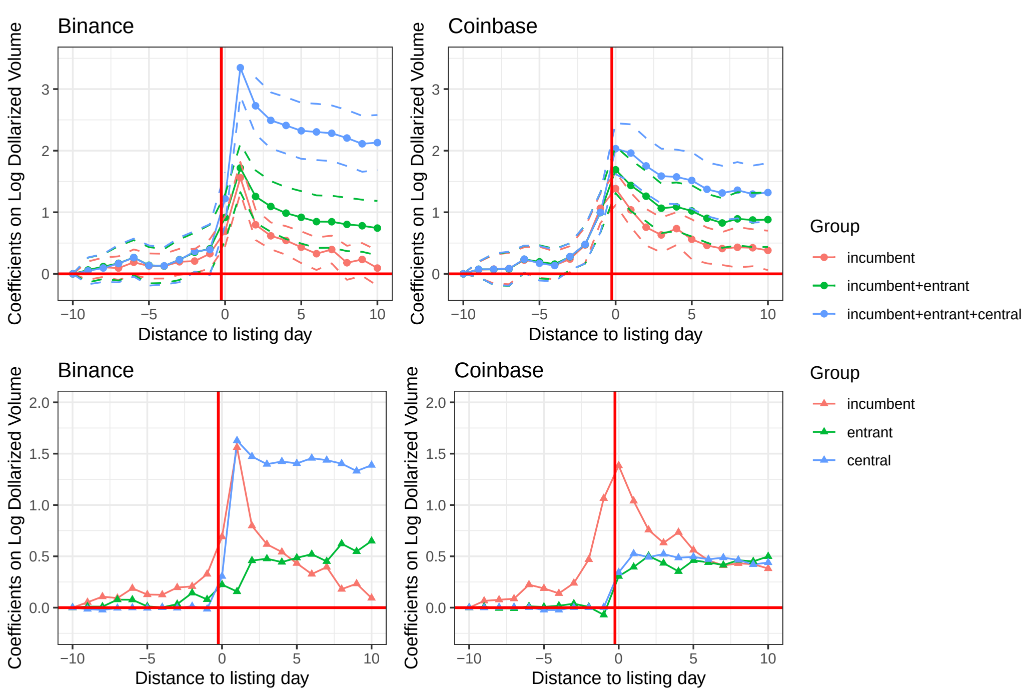

Table 1: Summary Statistics

This table presents summary statistics on variables related to coin outcomes, central exchange's listings, and other exchange level variable relative to the central exchange. Panel A shows descriptive statistics for exchange's return correlation of Bitcoin and listing following probability with regard to Binance and Coinbase. Panel B summarizes the coin-level coin outcomes and central exchange listings, Panel C shows the coin-exchange level variables. For each variable, we show the number of non-missing observations, the mean, the standard deviation and the 25th, 50th and 75th percentile values.

Panel A: Exchange Level

	Obs.	Mean	SD	p25	p50	p75	Obs.	Mean	SD	p25	p50	p75
	BINANCE						*COINBASE*					
BTC Return Correlation	245	0.77	0.21	0.45	0.84	0.97	246	0.76	0.24	0.44	0.83	0.98
Listing Following Prob	208	0.08	0.11	0	0.03	0.23	204	0.06	0.09	0	0.02	0.17

Panel B: Coin Level

	Obs.	Mean	SD	p25	p50	p75
Price Dispersion	403917	0.09	0.15	0	0.09	0.25
Net Listings	476513	0.01	0.23	0	0	0
Log (Volume)	476518	14.3	4.5	8.9	15.0	19.0
Listing (0-30 days) Binance	476518	0.02	0.13	0	0	0
Listing (> 30 days) Binance	476518	0.47	0.50	0	0	1
Pre Three-day Listing Binance	476518	0	0.04	0	0	0
Listing (0-30 days) Coinbase	476518	0.01	0.10	0	0	0
Listing (> 30 days) Coinbase	476518	0.14	0.35	0	0	1
Pre Three-day Listing Coinbase	476518	0	0.03	0	0	0

Panel C: Coin-Exchange Level

	Obs.	Mean	SD	p25	p50	p75
Log (Volume)	5762901	11.96	4.00	6.64	12.35	16.59
Listing (0-30 days) Binance	5763021	0.01	0.11	0	0	0
Listing (> 30 days) Binance	5763021	0.82	0.38	0	1	1
Pre Three-day Listing Binance	5763021	0	0.03	0	0	0
Listing (0-30 days) Coinbase	5763021	0.01	0.12	0	0	0
Listing (> 30 days) Coinbase	5763021	0.42	0.49	0	0	1
Pre Three-day Listing Coinbase	5763021	0	0.03	0	0	0

Table 2: Large Exchange Listings and Trade Volumes: DID Estimates

This table presents estimates from Specification (2):

$$\log(Volume_{c,e,t}) = \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> 30 \text{ days})_{c,t} + \beta_3 PreThreedayListing_{c,t} + \delta_{c,e} + \eta_t + \epsilon_{c,e,t}$$

$\log(Volume_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t. $Listing(\text{0-30 days})_{c,t}$ and $Listing(> 30 \text{ days})_{c,t}$ are dummy variables, equal to one for coin c on date t if a central exchange has listed coin c prior to date t but later than date $t-30$, and prior to date $t-30$, respectively. $PreThreedayListing_{c,t}$ is a dummy variable which is equal to one for coin i on date t if a central exchange decides to list coin i between date $t+1$ and date $t+3$. Columns (1) to (4) are result based on listings on Binance. Columns (5) to (8) are results based on listings on Coinbase. Standard errors are clustered at the coin-exchange pair level. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	Log(Volume)							
	Binance				Coinbase			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	1.6***	0.75***	0.49***	0.22**	0.94***	0.94***	0.76***	0.68***
	(0.12)	(0.09)	(0.08)	(0.07)	(0.07)	(0.06)	(0.05)	(0.05)
Listing (> 30 days)	0.88***	0.30**	0.16	0.01	0.43***	0.55***	0.55***	0.57***
	(0.13)	(0.10)	(0.09)	(0.08)	(0.09)	(0.07)	(0.06)	(0.06)
Pre Three-day Listing	0.96***	0.62***	0.48***	0.26***	0.99***	1.0***	0.99***	0.91***
	(0.11)	(0.09)	(0.08)	(0.07)	(0.07)	(0.06)	(0.05)	(0.05)
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	Yes	Yes	Yes	No	Yes	Yes	Yes	No
Country FE	No	Yes	No	No	No	Yes	No	No
Exchange FE	No	No	Yes	No	No	No	Yes	No
Coin-Exchange Pair FE	No	No	No	Yes	Yes	No	No	Yes
Adjusted R^2	0.20	0.40	0.62	0.79	0.20	0.40	0.62	0.79
Observations	5,762,901	5,602,902	5,762,901	5,762,901	5,762,901	5,602,902	5,762,901	5,762,901

Table 3: Exchange following and listing: DID estimates

This table shows estimates from specification (6):

$$\Delta\#Exchanges_{c,t} = \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> 30 \text{ days})_{c,t} + \beta_3 PreThreedayListing_{c,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

$\Delta\#Exchanges_{c,t}$ is the net change in the number of exchanges which list coin c in time t. $Listing(\text{0-30 days})_{c,t}$ and $Listing(> 30 \text{ days})_{c,t}$ are dummy variables, equal to one for coin c on date t if a central exchange has listed coin c prior to date t but later than date $t-30$, and prior to date $t-30$, respectively. $PreThreedayListing_{c,t}$ is a dummy variable which is equal to one for coin i on date t if a central exchange decides to list coin i between date $t+1$ and date $t+3$. Columns (1) and (2) are result based on listings on Binance and Coinbase, respectively. Standard errors are clustered at the coin level. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	Net Listings	
	Binance	Coinbase
	(1)	(2)
Listing (0-30 days)	0.11***	0.12***
	(0.009)	(0.01)
Listing (> 30 days)	0.01***	0.01**
	(0.003)	(0.004)
Pre Three-day Listing	0.04*	0.04*
	(0.02)	(0.062)
Day FE	Yes	Yes
Coin FE	Yes	Yes
Adjusted R^2	0.11	0.12
Observations	515,157	514,596

Table 4: Coin dispersion and listing: DID estimates

This table shows estimates from specification (6):

$$Dispersion_{c,t} = \beta_1 Listing(\text{0-30 days})_{c,t} + \beta_2 Listing(> 30 \text{ days})_{c,t} + \beta_3 PreThreedayListing_{c,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

$Dispersion_{c,t}$ is the standard deviation of log price across exchanges of coin c at time t, across exchanges. $Listing(\text{0-30 days})_{c,t}$ and $Listing(> 30 \text{ days})_{c,t}$ are dummy variables, equal to one for coin c on date t if a central exchange has listed coin c prior to date t but later than date $t-30$, and prior to date $t-30$, respectively. $PreThreedayListing_{c,t}$ is a dummy variable which is equal to one for coin i on date t if a central exchange decides to list coin i between date $t+1$ and date $t+3$. Columns (1) and (2) are result based on listings on Binance and Coinbase, respectively. Standard errors are clustered at the coin level. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	Dispersion	
	Binance	Coinbase
	(1)	(2)
Listing (0-30 days)	-0.03*	-0.005
	(0.01)	(0.007)
Listing (> 30 days)	-0.04**	-0.01
	(0.01)	(0.01)
Pre Three-day Listing	-0.007	0.003
	(0.01)	(0.008)
Day FE	Yes	Yes
Coin FE	Yes	Yes
Adjusted R^2	0.47	0.47
Observations	403,917	403,917

Table 5: Coin volume, exchange correlation, and listing

This table presents estimates of Specification (21). The dependent variables are log trading volume. $Listing(\text{0-30 days})_{c,t}$, $Listing(> 30 \text{ days})_{c,t}$, and $PreThreedayListing_{c,t}$ are all indicators that equal to one for coin c on date t if a central exchange has listed coin c prior to date t but later than date $t-30$, prior to date $t-30$, and between date $t+1$ and date $t+3$. $Correlation_e$ is the return correlation of Bitcoin between the central exchange and the peripheral exchange $e$using the entire time period where we have coverage for both exchanges in the pair. Columns (1) to (4) are result based on listings on Binance. Columns (5) to (8) are results based on listings on Coinbase. Standard errors are clustered at the coin-exchange pair level. Standard errors in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	Log (Volume)							
	Binance				Coinbase			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	-7.0***	-5.1***	-2.3***	-2.2**	-1.0	-0.86	-0.85*	-1.8***
	(0.96)	(1.2)	(0.69)	(0.80)	(0.64)	(0.59)	(0.42)	(0.47)
Listing (0-30 days) × Correlation	9.2***	6.4***	3.0***	2.6**	2.2**	2.0**	1.8***	2.7***
	(1.0)	(1.3)	(0.74)	(0.87)	(0.68)	(0.64)	(0.45)	(0.51)
Listing (> 30 days)	-3.9***	-4.6**	-1.5*	-3.1**	-0.85	-0.93*	-1.1***	-2.7***
	(1.1)	(1.5)	(0.59)	(0.98)	(0.54)	(0.54)	(0.30)	(0.53)
Listing (> 30 days) × Correlation	5.3***	5.3**	1.8**	3.3**	1.5**	1.7**	1.9***	3.6***
	(1.2)	(1.6)	(0.63)	(1.1)	(0.58)	(0.58)	(0.32)	(0.58)
Pre Three-day Listing	-3.9***	-4.7***	-3.6***	-3.5***	-1.1	-1.0	-0.85*	-1.7***
	(1.0)	(1.0)	(0.86)	(0.82)	(0.75)	(0.73)	(0.50)	(0.47)
Pre Three-day Listing × Correlation	5.2***	5.9***	4.4***	4.0***	2.3**	2.2**	2.0***	2.8***
	(1.1)	(1.1)	(0.92)	(0.89)	(0.80)	(0.79)	(0.53)	(0.50)
Correlation	3.9**	1.4			7.7***	5.0***		
	(1.2)	(1.6)			(0.49)	(0.64)		
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	Yes	Yes	Yes	No	Yes	Yes	Yes	No
Country FE	No	Yes	No	No	No	Yes	No	No
Exchange FE	No	No	Yes	No	No	No	Yes	No
Coin-Exchange Pair FE	No	No	No	Yes	Yes	No	No	Yes
Adjusted R^2	0.26	0.42	0.62	0.79	0.26	0.42	0.62	0.79
Observations	5,750,665	5,590,666	5,750,665	5,750,665	5,750,665	5,590,666	5,750,665	5,750,665

Appendix

A Supplementary Material for Section A

A.1 Logistics of International Remittances using Cryptocurrencies

As an extended example which illustrates the role of cryptocurrency exchanges in the usage of crypto, we describe the process of transferring funds internationally using cryptocurrencies. Suppose, for example, an individual in the USA who wished to transfer funds to a individual in the Philippines using cryptocurrencies. Such a transfer would follow the following steps:

1. The US-based individual would deposit fiat, using a bank transfer or other means, into a crypto exchange operating in the USA, and use these funds to purchase cryptocurrencies custodied on the exchange.
2. The US-based individual would "withdraw" her crypto to her private blockchain wallet.
3. The US-based individual could then send her cryptocurrencies to the wallet address of the individual in the Philippines.
4. The Philippines-based individual would "deposit" her crypto into a crypto exchange.
5. The Philippines-based individual would sell her crypto on the exchange for Philippines fiat currency, and then withdraw this, using a bank transfer or other means, to regular Philippines fiat currency.

The total fees charged in the course of this transaction include fees charged by exchanges for depositing, trading, and withdrawing in steps 1, 2, 4, and 5, as well as transaction fees charged for the blockchain transfer in step 3. The fees charged by exchanges vary. For the largest exchange, Binance, deposits and withdrawals are free, and purchases are charged around 0.1%, with discounts for very large trades and traders. Some smaller exchanges charge higher fees. The crypto transfer in step 3 has fees ranging from fractions of a cent to a few US dollars. Fees vary based on the degree of blockchain network congestion, but fees are generally independent of the value of the transaction. These transfers thus have competitive pricing, relative to some countries with inefficient traditional financial infrastructure.

An important benefit of crypto transfers is that they allow users to circumvent various regulations, such as capital controls as well as know-your-customer and anti-money-laundering provisions, imposed by national financial regulators. Crypto wallets are pieces of software or hardware, in which the security of funds is guaranteed through private-key cryptography. Self-custodied cryptocurrencies are not stored with any trusted intermediary: rather, a "private key" – a long numeric code, kept only on the user's hardware device – is used to prove to the blockchain network that the user owns her tokens, and to direct the network to take actions such as transfer tokens to other wallets. Crypto "miners", which build the blockchain by inserting proposed transactions in new "blocks", are incentivized to mine by newly minted crypto tokens they are given, and transaction fees which are paid by users for each transaction that they "mine". Since miners have no access to individuals' private keys, they have no ability to take funds from individuals' wallets.

It is logistically very difficult for regulators to enforce capital controls and other transfer restrictions directly on crypto transfers at the blockchain level, that is, step 3. of the process above. Firstly, there is no public mapping from addresses to individuals, so regulators cannot easily tell who owns a wallet, or even what country a wallet's owner resides in. Secondly, even if regulators were able to identify a set of wallets to impose potential transfer limitations on, enforcing transfer restrictions is difficult to to the structure of blockchain mining, because transactions are processed by geographically dispersed miners in an essentially discretion-free manner. Hypothetically, for example, if US-based Ethereum miners were instructed by US regulators to stop processing transactions from certain wallets, these transactions would only have to wait in the "mempool" of proposed transactions until a non-US miner not subject to the restriction mined a block and included the transaction.[11]

Crypto exchanges play a critical role in the process of sending funds due to their role in steps 1, 2, 4, and 5 of the funds transfer process. They serving as "on/off-ramps", by allowing deposits and withdrawals of crypto or fiat, and the trading of fiat for crypto. Since on-blockchain crypto transfers cannot easily be restricted, regulators have instead focused on imposing financial regulations through exchanges. For example, in the USA, a 2019 joint statement by the CFTC, FinCEN, and the SEC announced that crypto exchanges were classified as money services businesses, and thus are subject to KYC and AML rules under the Bank Secrecy Act of 1970. US-based crypto exchanges thus must gather identifying information about their customers to comply with these requirements. Crypto exchanges in

[11] One class of exceptions to this rule is that the administrators of certain tokens, such as the Circle (USDC) and Tether (USDT) USD stablecoins, include code in the "smart contracts" governing their tokens which allows them to freeze the funds of certain "blacklisted" wallets. These token administrators cooperate with regulators to freeze the funds of wallet addresses identified as being involved in hacks or other criminal activity. See, for example, Coindesk and Cointelegraph. However, freezing funds is only possible if, at the creation of the token, administrators include the capability to blacklist tokens, and the majority of crypto tokens do not have built-in blacklist functionality.

many other countries with strict financial regulations are subject to similar requirements.

There are other ways to exchange fiat for cryptocurrencies besides custodial crypto exchanges. Users can simply exchange cryptocurrencies for fiat informally in social networks. Peer-to-peer exchanges, such as LocalBitcoins, also exist, which pair buyers and sellers of crypto in a manner that does not involve exchange custody of assets. Various institutions existing in legal gray areas also offer to exchange fiat for crypto across countries; for example, black market exchanges in Argentina allow individuals to exchange Argentinian pesos for USD, as well as various cryptocurrencies.[12]

B Data Cleaning

This part introduces our data cleaning process. The raw dataset contains hourly data for each trading pair on each exchange, including price and volume variables. Our main goal is to create a daily coin-exchange pair level dataset, with price and volume variables for each coin on each exchange for each day. We followed five steps:

1. **Aggregating data at the daily level.** For each trading pair on each exchange within a day, we aggregate the hourly data by taking the average price of each open hour and the total volume of all hours.

2. **Focusing on top 500 coins and common trading pair.** We restrict our sample in two ways: by the cryptocurrency in the trading pair, and by the denominator that the cryptocurrency is traded against. Since many coins are not actively traded, we first restrict our sample to trading pairs involving the top 500 cryptocurrency coins ranked by coinmarketcap.com on September 3, 2022. We also restrict our sample to three kinds of trading pairs: pairs involving one of our 27 major fiat currencies [13] ; pairs involving BTC or ETH, which are the two largest cryptocurrencies by market cap; and pairs involving one of the three major stablecoins (USDT, USDC, BUSD).

3. **Converting prices and volumes to USD terms.** For fiat pairs, we convert using same-day USD-fiat exchange rates, and for crypto pairs, we convert using daily prices of cryptocurrencies and stablecoins from Yahoo finance.

4. **Aggregating data at the coin level.** For each top 500 coin on each exchange on each day, we aggregate data across all trading pairs involving the same coin, taking the

[12] See Devon Zuegel.

[13] These 27 major fiat currencies include: NZD USD KRW JPY CNY IDR SGD VND TWD AUD PKR ZAR TRY MXN BRL CHF ILS PLN GBP RUB EUR CAD HKD INR SAR AED SEK.

average of the prices for each trading pair involving same coin within a day weighted by its USD trading volume, and adding volumes across all trading pairs of the same coin within a day.

5. **Winsorizing and imputating data.** Due to price outliers, we also winsorize data by a maximum and minimum of: 2 and 0.5 times the median price of each coin on each day, respectively. We do that in order to measure dispersion properly, or there will be some explosive numbers for SD of log prices. Moreover, we impute the missing data for some coins that are listed on the exchange but have no trading for a few days. [14]We assign the price as the missing value and the trading volume as 0.

Finally, we obtained daily prices and volumes in USD terms for each coin listed on exchanges. We used the coin-exchange level data from step 5 for most of our analysis. For some of our analysis, we further aggregated this data to the coin level or the exchange level by taking the average price and the total volume.

C Proofs

C.1 Proof of Proposition 1

Prices. When the CEX does not list the token, arbitrageurs have no activity. Market clearing requires aggregate demand from all users on exchange j to equal 0. Hence, from (10), we need:

$$Z_{user,j}\left(p_j\right) = \int_{-\infty}^{\infty} z_i\left(x\right) dF_{x_{i,0}}\left(x\right) = \frac{-\gamma_j}{\gamma_j+\tau_j}\eta_j + \frac{\psi - p_j}{\gamma_j+\tau_j} = 0$$

Solving for p_j, we have:

$$p^*_{j,0} = \psi - \gamma_j \eta_j \tag{26}$$

This is (12).

Trade quantities. To solve for expected squared trade quantity, note that user i's trade quantity is (9). Plugging in for $\psi - p_j$ using (26), we have:

$$z^*_{i,0} = \frac{-\gamma_j}{\gamma_j+\tau_j} x_{i,0} + \frac{\gamma_j \eta_j}{\gamma_j+\tau_j}$$

[14]Specifically, if there are missing days for each coin listed on any exchanges that lie between the first date and last date that appear in the data, we impute these observations. Number of observations increase by 19% from 5,763,021 to 6,835,510.

Thus, we have:

$$\mathbb{E}\left[z_{i,0}^{*2}\right] = \int_{-\infty}^{\infty} \left(\frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - p_{j0}^*}{\gamma_j + \tau_j} \right)^2 dF_{x_{i,0}}(x)$$

Plugging in for p_{j0}^* using (26) and simplifying, we have:

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \int_{-\infty}^{\infty} (\eta_j - x_{i,0})^2 \, dF_{x_{i,0}}(x)$$

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \sigma_{I,j}^2$$

Exchange profits. The exchange's profit from user i is simply $\frac{\tau_j}{2} z_i^2$; hence, the exchange's profit over all users is:

$$\pi_{j,0}^* = \int_{-\infty}^{\infty} \frac{\tau_j}{2} z_i^{*2}(x) \, dF_{x_{i,0}}(x) = \frac{\tau_j}{2} \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \sigma_{I,j}^2$$

C.2 Proof of Proposition 2

Prices. When the CEX lists the token, arbitrageurs can trade the risky asset on j as well as the central exchange. Market clearing requires aggregate demand from all users and arbitrageurs on exchange j to equal 0. Hence, from (10) and (11), we need:

$$Z_{user,j}(p_j) + Z_{arb,j}(p_j) = \left(\frac{-\gamma_j}{\gamma_j + \tau_j} \eta_j + \frac{\psi - p_j}{\gamma_j + \tau_j} \right) + \frac{\psi - p_j}{\zeta_j + \tau_j} = 0$$

Solving for p_j, we have:

$$p_{j,1}^* = \psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j \tag{27}$$

This is (15).

Trade quantities. To solve for expected squared trade quantity, note that user i's trade quantity is (9). Plugging in for $\psi - p_j$ using (27), we have:

$$z_{i1}^* = \frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \frac{\gamma_j \eta_j}{\gamma_j + \tau_j}$$

Taking the expectation over all users, we have:

$$\mathbb{E}\left[z_{i1}^{*2}\right] = \int_{-\infty}^{\infty} \left(\frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - p_{j,1}^*}{\gamma_j + \tau_j} \right)^2 dF_{x_{i,0}}(x)$$

Plugging in for prices using (27), we have: Algebra:

$$= \int_{-\infty}^{\infty} \left[\frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - \left(\psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j\right)}{\gamma_j + \tau_j} \right]^2 dF_{x_{i,0}}(x)$$

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \int_{-\infty}^{\infty} \left(-x_{i,0} + \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\eta_j \right)^2 dF_{x_{i,0}}(x)$$

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \int_{-\infty}^{\infty} \left[x_{i,0}^2 - 2\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\eta_j x_{i,0} + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \eta_j^2 \right] dF_{x_{i,0}}(x)$$

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \left[\left(\sigma_{I,j}^2 + \eta_j^2\right) - 2\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\eta_j^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \eta_j^2 \right]$$

$$= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \eta_j^2 \right]$$

Exchange profits. The exchange's profit from user i is simply $\frac{\tau_j}{2} z_i^2$; hence, the exchange's profit over all users is:

$$\pi_{j,1}^* = \int_{-\infty}^{\infty} \frac{\tau_j}{2} z_i^{*2}(x)\, dF_{x_{i,0}}(x) = \frac{\tau_j}{2} \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \eta_j^2 \right]$$

C.3 Proof of Proposition 3

Here we assume the aggregate inventory shock at peripheral exchange η_j is independent of the efficient price ψ and aggregate inventory shock at other peripheral exchanges. The coefficient of determination R^2 between the central exchange's price, and peripheral exchange j's price, is:

$$\begin{aligned} R_{j,CE}^2 &= \frac{Cov^2\left(p_j^*, \psi\right)}{Var\left(p_j^*\right) Var(\psi)} \\ &= \frac{Cov^2\left(\psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j, \psi\right)}{Var\left(\psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j\right) Var(\psi)} \\ &= \frac{Cov^2(\psi, \psi)}{\left[Var(\psi) + Var\left(-\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j\right)\right] Var(\psi)} \\ &= \frac{\sigma_\psi^2}{\sigma_\psi^2 + \left(\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2} \end{aligned} \tag{28}$$

The R^2 between the prices of exchanges j and j' is:

$$\begin{aligned}
R^2_{j,j'} &= \frac{Cov^2\left(p^*_j, p^*_{j'}\right)}{Var\left(p^*_j\right)Var\left(p^*_{j'}\right)} \\
&= \frac{Cov^2\left(\psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j, \psi - \frac{\zeta_{j'}+\tau_{j'}}{\gamma_{j'}+\zeta_{j'}+2\tau_{j'}}\gamma_{j'}\eta_{j'}\right)}{Var\left(\psi - \frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j\right)Var\left(\psi - \frac{\zeta_{j'}+\tau_{j'}}{\gamma_{j'}+\zeta_{j'}+2\tau_{j'}}\gamma_{j'}\eta_{j'}\right)} \\
&= \frac{Cov^2\left(\psi,\psi\right)}{\left[Var\left(\psi\right) + Var\left(-\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\gamma_j\eta_j\right)\right]\left[Var\left(\psi\right) + Var\left(-\frac{\zeta_{j'}+\tau_{j'}}{\gamma_{j'}+\zeta_{j'}+2\tau_{j'}}\gamma_{j'}\eta_{j'}\right)\right]} \\
&= \frac{\sigma^2_\psi}{\left[\sigma^2_\psi + \left(\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2\gamma_j^2\sigma^2_{A,j}\right]}\frac{\sigma^2_\psi}{\left[\sigma^2_\psi + \left(\frac{\zeta_{j'}+\tau_{j'}}{\gamma_{j'}+\zeta_{j'}+2\tau_{j'}}\right)^2\gamma_j^2\sigma^2_{A,j'}\right]}
\end{aligned} \tag{29}$$

The R^2 between the prices of exchanges j and j' is simply the product of the R^2 between the prices of exchanges j and the central exchange, and the R^2 between the prices of exchanges j' and the central exchange. Therefore, we always have:

$$R^2_{j,CE} \geq R^2_{j,j'}$$

C.4 Proof of Prediction 5

The prediction that the correlation between the central exchange's price and peripheral exchange j's price is decreasing in the arbitrage costs ζ_j follows directly from (28):

$$\frac{\partial R^2_{j,CE}}{\partial \zeta_j} = \frac{-2\sigma^2_\psi\sigma^2_{A,j}\gamma_j^2\left(\zeta_j+\tau_j\right)\left(\gamma_j+\tau_j\right)}{\left[\sigma^2_\psi + \left(\frac{\zeta_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2\gamma_j^2\sigma^2_{A,j}\right]^2\left(\gamma_j+\zeta_j+2\tau_j\right)^3} \leq 0 \tag{30}$$

The volume increase of peripheral exchanges after the central exchange lists is defined as:

$$\begin{aligned}
\Delta\mathbb{E}\left[z_{i,1}^{*2}\right] &= \frac{\mathbb{E}\left[z_{i1}^{*2}\right] - \mathbb{E}\left[z_{i0}^{*2}\right]}{\mathbb{E}\left[z_{i0}^{*2}\right]} \\
&= \frac{\left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2 \eta_j^2\right] - \left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2 \sigma_{I,j}^2}{\left(\frac{\gamma_j}{\gamma_j+\tau_j}\right)^2 \sigma_{I,j}^2} \\
&= \left(\frac{\gamma_j+\tau_j}{\gamma_j+\zeta_j+2\tau_j}\right)^2 \left(\frac{\eta_j}{\sigma_{I,j}}\right)^2
\end{aligned} \tag{31}$$

The volume increase of peripheral exchanges is also decreasing in the arbitrage costs ζ_j from (31):

$$\frac{\partial\Delta\mathbb{E}\left[z_{i,1}^{*2}\right]}{\partial\zeta_j} = -\frac{2(\gamma_j+\tau_j)^2}{(\gamma_j+\zeta_j+2\tau_j)^3}\left(\frac{\eta_j}{\sigma_{I,j}}\right)^2 \leq 0 \tag{32}$$

Similarily, the exchange's profit from user i is simply $\frac{\tau_j}{2} z_i^2$. Therefore, the profits increase of peripheral exchanges is also decreasing in the arbitrage costs ζ_j.