

Competition in the Cryptocurrency Exchange Market*

Junyi Hu[†], Anthony Lee Zhang[‡]

March 2024

Abstract

How do cryptocurrency exchanges compete with each other? By analyzing small exchanges' responses to large exchanges' coin listing decisions, we show that small and large crypto exchanges appear to be complements, rather than substitutes, as traditional oligopoly theory would predict. These facts are consistent with a model in which small exchanges behave like *brokerages*, offering captive customer bases a “costly window” to deep liquidity large exchanges, which behave like *inter-broker clearinghouses*. Large exchanges thus play a “leader” role in international cryptoasset markets, which is understated by their market shares of trading volume.

*We thank Andreas Park, Simon Trimborn, Yizhou Xiao, the wagmi chat, and seminar participants at CUHK, European Finance Association 2023 Annual Meeting, the Hong Kong Conference for Fintech, AI, and Big Data, UTS, and Loyola University for helpful comments. We are especially grateful to Jinfei Sheng for comments that greatly improved the paper. We are grateful to Haichuan Wang and Chupei Zhang for excellent research assistance.

[†]Columbia University, Columbia Business School; jh4736@gsb.columbia.edu

[‡]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 14% 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 coin. If exchanges were competing over a fixed customer base, trade volumes of the coin on smaller exchanges should *decrease*, and small exchanges who have not already listed the coin 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 coin, trade volumes of the coin on smaller exchanges *increase*, and small exchanges become *more* likely to list the coin. In other words, large and small exchanges appear to behave like economic complements, rather than economic substitutes.

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

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

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

price correlations across exchanges, and how price correlations relate to the volume-increase and listing-following effects, which we verify empirically. Our results illustrate that small peripheral exchanges can be thought of like *brokerages* for traditional financial assets, and large exchanges like *inter-broker markets*. Moreover, our results imply that large exchanges play a systemically important “leader” role in international cryptoasset markets, which is understated by their market share of trading volume.

We use data on 500 large crypto coins’ prices and trades volumes across 256 exchanges from January 2017 to July 2022. Using this data, we demonstrate three stylized facts. First, when a large exchange lists a new coin, trade volumes of the coin on small exchanges tend to *increase*, by around 38%-76% 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 coin. 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 coin prices across small exchanges by around 7%-22%, suggesting that arbitrage flows across large and small exchanges may play a role in explaining these findings.

We construct a simple conceptual framework to rationalize these results. We model the strategic interactions between a single “central” exchange and a number of “peripheral” exchanges, which have captive customer bases and partially rely on imperfect arbitrage with the central exchange for liquidity provision to their customers. There is a single risky asset, or “coin”, which can be traded. The central exchange has infinite market depth. Each peripheral exchange has a set of captive customers. Consumers may be captive, for example, because only a given peripheral exchange connects to the fiat payment systems they use, or is legal to use in their jurisdiction, or because an exchange offers a differentiated trading experience which is valued by these consumers. Consumers receive inventory shocks for the risky asset; inventory shocks have an aggregate and idiosyncratic component, so customers of a given peripheral exchange may on net want to buy or sell a coin. 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 central 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 central 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 coin 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 central exchange lists the coin, arbitrageurs trade to partially close the price gaps between the peripheral exchange and the central exchange. This effectively gives peripheral exchange customers partial access to central 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 central exchange's entry does not directly cannibalize the peripheral exchange's customers; however, the entry of the central 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 central exchanges' coin listing decisions. The model also predicts that central exchange listings should decrease the cross-sectional dispersion of coin prices across peripheral exchanges, since core-periphery arbitrage trade cause peripheral exchange prices to cluster close to central 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 central 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 central 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 central exchange should have stronger price correlations with the central exchange, larger volume increases when the central exchange lists, and a larger tendency to follow the central exchange's listing decisions. We find empirical evidence supporting both predictions. We also show that our stylized facts are robust to concerns about wash trading and falsified volume, as well as a number of other robustness checks.

Our results suggest a narrative on the nature of crypto exchange market structure. Small crypto exchanges appear to play a role similar to *brokerages* in traditional asset markets: they focus on solving jurisdictional issues and payment connectivity problems, allowing a specific set of consumers to trade fiat for crypto. Small exchanges are then connected by arbitrageurs to large exchanges, which serve a role similar to *inter-broker markets* for traditional assets,

clearing aggregated order flow from a large number of smaller exchanges. Large exchanges thus focus on providing deep global markets for cryptoassets, leaving the task of integrating with specific payments rails, and presenting specific trading interfaces to consumers, largely to arbitrageurs and small exchanges. As some evidence for this narrative, we show that coin price discovery appears to happen largely on large exchanges: when small and large exchanges’ prices deviate, the gap tends to close by small exchanges’ prices moving towards larger exchanges’ prices.

This narrative implies that a small set of central 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 14% 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 coin trading volumes, liquidity, and the decisions of other exchanges whether to list coins 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 1264pp, and Coinbase’s listings increase volumes by around 236pp, within the 10-day window after listing. A large part of this effect is indirect: the volume increase consists of a 275pp (51pp) increase in volume directly on Binance (Coinbase), and a 258pp (122pp) increase on incumbent exchanges, and new exchanges which list the coin following Binance or Coinbase’s listing decisions. Despite the large power central exchanges have to shape market outcomes, central 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 coins they plan to list, and the reasoning for listing these coins.

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 \(2019\)](#) characterizes the exchanges which drive BTC price discovery. This literature has shown that there is meaningful dispersion in prices across exchanges, which is correlated with factors such as capital flows. Relative to this literature, our contribution is to analyze how exchanges strategically interact: how large exchanges’ listing decisions affect small exchange

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

volumes, and how small exchanges’ listing decisions respond to large exchanges. [Shams \(2020\)](#) show that coins which are listed on similar exchanges tend to have greater return comovement. This result is related to, and complementary to, our finding that listing on large exchanges decreases within-coin price dispersion; broadly, both papers shed light on how exchange market structure and exchanges’ listing decisions influence coin prices. [Benedetti and Nikbakht \(2021\)](#) analyze “cross-listings,” defined as the second time a coin is ever listed on an exchange, and find that coin prices, trade volumes, and coin network growth increase after cross-listings. Their analysis is related to, but distinct from, our analysis of the effect of large exchange listings; moreover, [Benedetti and Nikbakht](#) do not analyze listing propensity or coin price dispersion. We also relate to a number of papers that have analyzed the effect of large exchange listings on coin returns ([Ante, 2019](#); [Lemmen, 2022](#)), though we do not focus on returns in this paper.

This paper fits 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 \(2021\)](#) construct 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. [Cong et al. \(2023\)](#) analyze the concentration of activity on the Ethereum blockchain. [Li, Shin and Wang \(2021\)](#) analyzes cryptocurrency pump-and-dump schemes. [von Luckner, Reinhart and Rogoff \(2023\)](#) analyze the use of Bitcoin transactions for capital flight. [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. [Félez-Viñas, Johnson and Putniņš \(2022\)](#) show evidence of systematic insider trading in cryptocurrency markets. 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 discuss our results in Section 6, and conclude in Section 7.

2 Data and Institutional Background

2.1 Institutional Background

Cryptocurrency Exchanges. 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, on any payment rail supported by the exchange. Unlike traditional financial assets, however, cryptoassets like Bitcoin exist on blockchains, allowing users to self-custody assets. Users can “withdraw” custodied assets from an exchange, 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. Once withdrawn into personal crypto wallets, cryptoassets are not jurisdiction-specific: from a logistical perspective, Bitcoins held by consumers in Nigeria are entirely fungible with Bitcoins held by consumers in Vietnam, or any other country. Crypto exchanges thus serve as “on/off-ramps” for crypto, serving as a bridge between fiat and payment rails specific to individual jurisdictions, and the unified international ecosystem of blockchain-based cryptoassets.

As an example of how consumers might use cryptocurrency exchanges, 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 recipient, 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, many market participants purchase cryptocurrencies on centralized exchanges to speculate on crypto price appreciation.

There are a very large number of crypto exchanges: according to [Blockspot.io](https://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 14% 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](#)).

Coin Listing on Crypto Exchanges. Crypto exchanges have essentially full discretion on the coins they “list”, that is, enable their customers to trade on their exchange. From the perspective of a coin issuer, exchange listings increase coin liquidity and the potential buyer base for coins, as well as increasing the public recognition of the coin. From the perspective of an exchange, analogously, trading volume and thus fee revenue increases if exchanges list more coins that their customers want to trade.

An important assumption in our model is that it is costly for an exchange to list a new coin. Listing a crypto coin can involve much higher costs than, for example, an equity exchange listing a new stock for trade. A crypto exchange listing a new coin must build systems to custody the coin, and to meet users’ deposit and withdrawal requests; these processes can be costly, especially if the coin is on a blockchain that the exchange does not already work with.

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

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

Many exchanges also conduct a legal, compliance, and security reviews before listing coins.⁶

When listing coins, crypto exchanges often find market makers to commit to providing liquidity, in order to ensure that listed coins trade at reasonable bid-ask spreads.⁷ Market makers will often commit to maintaining certain levels of market depth or bid-ask spreads.⁸ In our model, we will assume there are arbitrageurs who trade against price deviations across crypto exchanges; this role is likely played largely by professional market makers in practice.

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 weighted by its trading volume, and adding the trade volumes across all hours within a day.

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;⁹ 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

⁶See for example [Coinbase](https://coinbase.com) and [Crypto.com](https://crypto.com). Coinbase states that around 90% of coins reviewed do not meet their legal listing standards.

⁷See, for example, [Binance](https://binance.com) and [Coinbase](https://coinbase.com).

⁸See, for example, [Wintermute](https://wintermute.com).

⁹These 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.

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. After dropping stablecoins and fiat currencies in the final step, the final sample consists of 467 coins across 256 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 coin, incumbent small exchanges which have previously listed the coin experience increases in coin trading volume.*

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

$$\log(\text{Volume}_{c,e,t}) = \sum_{k=-31}^{31} \beta_k \times \text{treat}_{c,k,t} + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t} \quad (1)$$

where c indexes coins, e indexes exchanges, and t indexes days. $\log(\text{Volume}_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t . $\text{treat}_{c,k,t}$, with k from -30 to 30, 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. $\text{treat}_{c,-31,t}$ and $\text{treat}_{c,31,t}$ are dummy

variables that equal 1 when time t is more than 30 days before and after the large exchange listing date, respectively. Coins that have not yet been listed on large exchanges are considered as the control group. If coin c has not experienced a large exchange listing, $treat_{c,-31,t}$ will always equal 1 and other dummy variables always equal 0. To avoid collinearity problem, we set observations that are exactly 30 days before large exchange listings as the reference group, which means that β_{-30} is always 0. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges.

We plot estimated coefficients of $treat_{c,k,t}$ with k from -30 to 31. The results are shown in Figure 2. When a large exchange lists a coin, there is a large increase in trade volumes on small exchanges that have previously listed the coin. 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.¹⁰ Directly after the large exchange lists, we observe a sharp increase in trading volume: volumes increase by around 162% and 111%, 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 30 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 coin, which may tend to cannibalize market share from incumbent small exchanges which have already listed the coin.

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

$$\log(\text{Volume}_{c,e,t}) = \beta_1 \text{Listing}(0-30 \text{ days})_{c,t} + \beta_2 \text{Listing}(> 30 \text{ days})_{c,t} + \beta_3 \text{PreThreedayListing}_{c,t} + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t} \quad (2)$$

$\text{Listing}(0-30 \text{ days})_{c,t}$ and $\text{Listing}(> 30 \text{ days})_{c,t}$ are dummy variables; $\text{Listing}(0-30 \text{ days})_{c,t}$ 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 $\text{Listing}(> 30 \text{ days})_{c,t}$ is one if a large exchange has listed coin c prior to date $t - 30$. $\text{PreThreedayListing}_{c,t}$ 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

¹⁰Binance usually announces coin listings one day before they are available for trading, while Coinbase typically announces them one or two days in advance.

$t + 1$ 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 coin. 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 and time level. We only keep incumbent coin-exchange pairs that list at least 30 days before large exchange listings, or pairs with coins that have not been listed by large exchanges.

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, coin, and exchange fixed effects. Columns (2) and (6) substitute exchange fixed effects with fixed effects for the country in which the exchange operates. Columns (3) and (7) include day and coin-exchange pair fixed effects, controlling for any sources of unobserved heterogeneity that affect all coins traded on a given exchange on a specific day. Columns (4) and (8) further include interactions between exchange and time fixed effects to control for unobserved heterogeneity at the exchange level. The results are broadly consistent with the graphical evidence presented in Figure 2. The $Listing(0-30 \text{ days})$ coefficients obtained across all specifications are positive and significant. Quantitatively, we find that Binance listings increase volumes within 30 days by 38% to 49%, and Coinbase listings increase 30-day volumes by around 69% to 76%. Consistent with Figure 2, the estimated effects are smaller after 30 days, though in most cases they are still positive and significant.¹¹

¹¹A 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. In Appendix Table D.1 and Figure D.1, we provide further evidence that volume falsification is not the main driver of our results.

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 coin, many small exchanges that previously did not list the coin 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. The sample is also filtered to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges. 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:

$$\Delta\#Exchanges_{c,t} = \sum_{k=-31}^{31} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \quad (3)$$

The dependent variable, $\Delta\#Exchanges_{c,t}$, is the net change in the number of exchanges (excluding Binance and Coinbase themselves) which list coin c in time t . All other variables are defined above. ¹² δ_c and η_t are respectively coin and time fixed effects. We filter the sample to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges, eliminating the mechanical effects associated with coins initially listed on large exchanges. 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 at the coin and time level. The results are shown in Figure 4.

¹²Thus, $\Delta\#Exchanges_{c,t}$ 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.

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 coins within 2 days after they are listed by a large exchange. Binance and Coinbase listings lead to at most 0.37 and 0.38 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. The small number translates into a sizable effect, considering that there are only 7 and 17 listings on average when Binance and Coinbase list coins, respectively.

We then estimate the following regression specification:

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

All other variables are defined above, and the sample is also filtered to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges. The results are shown in Table 3. Within 30 days of listing by Binance and Coinbase, the average net listings per day increases by 0.076 and 0.079, respectively. It is equivalent to a cumulative 2.3 and 2.4 listings in the first 30 days after Binance and Coinbase listings, respectively. The effects are not very persistent, though the effect is small but significant for Binance listings. One possible explanation for this is the limited market capacity, suggesting that the influx of listings after being listed by large exchanges could saturate the market.

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 coin, the dispersion of coin 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 incumbent exchanges, and normalize all percentiles by the median. Incumbent exchanges are defined as exchanges that list coin c at least 30 days before its listing on a large exchange, or exchanges whose coins have never been listed by large exchanges. We

then plot the average of the normalized quantiles across coins affected by listings, for a 30-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 600bps for coins Binance lists, and around 310bps for coins Coinbase lists. Visually, dispersion declines substantially after listings, by around 180bps and 100bps respectively for Binance and Coinbase.

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

$$Dispersion_{c,t} = \sum_{k=-31}^{31} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \quad (5)$$

$Dispersion_{c,t}$ is the standard deviation of log prices across all incumbent exchanges for coin c at date t . Incumbent exchanges are defined as exchanges that list coin c at least 30 days before its listing on a large exchange, or exchanges whose coins have never been listed by large exchanges. 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 and time level. Figure 6 presents the results. Immediately after large exchange listings, dispersion shows a gradual decreasing trend. The effect is statistically significant for Binance, but not for Coinbase. Quantitatively, dispersion decreases by approximately 0.03 and 0.009, or 28% and 10% 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 using the same filtered sample:

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

The results are presented in Table 4. The estimates in Column (1) show that Binance listings lead to a 0.02 decrease in price dispersion, relative to coins not affected by listings. This is equivalent to a 22% reduction in dispersion. Unlike our volume results, the effects on price dispersion persist after 30 days. The average daily normalized price dispersion decreases by 0.03, or 34%. For Coinbase listings in Column (2), we also observe a 0.007 and 0.02 decrease in dispersion, or 7% and 19% in percentage terms, within 30 days and 30 days after Coinbase listings, respectively. 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 coin market to become more efficient through increased arbitrage trade, which may increase small exchanges’ trade volumes and thus their profits from listing the coin. 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 Conceptual Framework

We construct a simple model where a single coin is traded on an infinitely deep “central” 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 central 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 “coin”. There is a central exchange, on which the price of the asset is ψ ; the central 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 central 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. An important assumption we make is that users are fully captive: exchange j ’s users have no ability to trade on any other peripheral exchange. One possible microfoundation for users’ captivity is jurisdictional differences: as we discussed in Subsection 2.1, a user in Australia may not be able to use a Vietnam-based crypto exchange, both because the Australian exchange may refuse to work with the consumer for legal reasons, and because the Australian exchange may accept Australian fiat currency and payment rails, which the Vietnamese consumer may not be able to access. Another possible microfoundation is differences in demand for exchanges features across consumers: for example, certain unsophisticated consumers may prefer exchanges which are simpler, or have mobile

applications, whereas sophisticated consumers may prefer exchanges with rich data displays and more efficient and differentiated order submission technology. Sophisticated consumers would then be unwilling to switch to exchanges with features targeted at unsophisticated consumers. The assumption that users are fully captive is stylized. Our main results would still hold if we assumed consumers were only partially differentiated, with the ability to move across peripheral exchanges at some cost, as long as switching costs are sufficiently high; we focus on the fully captive case for expositional simplicity.

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:

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

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:

$$x_{i,0} = \eta_j + \xi_{ij} \quad (8)$$

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:

$$V_i = \underbrace{\psi(z_i + x_{i,0})}_{\text{Fundamental Value}} - \underbrace{p_j z_i}_{\text{Net Payment}} - \underbrace{\frac{\gamma_j}{2}(z_i + x_{i,0})^2}_{\text{Inventory Costs}} - \underbrace{\frac{\tau_j}{2}z_i^2}_{\text{Exchange Fees}} \quad (9)$$

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:

$$\frac{\partial V_i}{\partial z_i} = \psi - p_j - \gamma_j (z_i + x_{i,0}) - \tau_j z_i \quad (10)$$

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

$$z_i = \frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - p_j}{\gamma_j + \tau_j} \quad (11)$$

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

$$Z_{user,j}(p_j) = \int_{-\infty}^{\infty} z_i(x) dF_{x_{i,0}}(x) = \frac{-\gamma_j}{\gamma_j + \tau_j} \eta_j + \frac{\psi - p_j}{\gamma_j + \tau_j} \quad (12)$$

where $F_{x_{i,0}}(x)$ is the cumulative probability function of user's initial inventory.

Each peripheral exchange j has a unit measure of atomistic arbitrageurs, who can trade the risky asset on j as well as the central exchange. We think of these arbitrageurs as modelling professional liquidity providers who make markets on exchanges, as we described in Subsection 2.1. For simplicity, 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 central exchange as much as they sell on the peripheral exchange. Let z_k be the net amount arbitrageur k buys on j and sells on the central exchange. We assume arbitrageurs face quadratic inventory costs for arbitrage: they incur a cost $\frac{\zeta_j}{2} z_k^2$ for arbitraging z_k units of the asset.¹³ We assume arbitrageurs pay the same trading fees as users: if they trade a quantity z_k , they pay fee $\frac{\tau_j}{2} z_k^2$.¹⁴ Arbitrageurs' value for buying z_k units at price p_j on exchange j , and selling at price ψ on the central exchange, is:

$$V_k(z_k) = z_k(\psi - p_j) - \frac{\zeta_j}{2} z_k^2 - \frac{\tau_j}{2} z_k^2 \quad (13)$$

¹³Practically, in order to arbitrage a crypto coin across two exchanges, an arbitrageur must execute a multi-stage trade similar to the process we describe in Appendix A.1: the arbitrageur must send fiat to exchange A; purchase crypto on exchange A; withdraw the crypto and send it to exchange B; sell the crypto for fiat on exchange B; and then withdraw fiat from exchange B. This process ties up capital on both exchanges, and incurs logistical delays which may create price risk from temporary inventory holdings. Technically, inventory costs imply that arbitrageurs do not perfectly eliminate price gaps between central and peripheral exchanges in our model.

¹⁴This assumption may not hold exactly in practice – market participants informed us that exchanges often give fee discounts to their market makers – but relaxing this assumption does not substantively affect our results.

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

$$\frac{\partial V_k}{\partial z_k} = \psi - p_j - \zeta_j z_k - \tau_j z_k \quad (14)$$

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

$$z_k = \frac{\psi - p_j}{\zeta_j + \tau_j} \quad (15)$$

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

$$Z_{arb,j}(p_j) = \frac{\psi - p_j}{\zeta_j + \tau_j} \quad (16)$$

Peripheral exchange j 's profits, if trade volume is z_i for each user, are: $\int_{-\infty}^{\infty} \frac{\tau_j}{2} z_i^2(x) dF_{x_{i,0}}(x)$. We assume exchange j has some cost C_j of listing coins; we discussed some possible microfoundations of these costs in Subsection 2.1.¹⁵ Exchange j will list the risky asset if it anticipates profits greater than C_j from listing. We will model the central exchange's listing decisions as exogenous.

4.1 Equilibrium

In equilibrium, aggregate demand from users and arbitrageurs sums to 0 on each exchange. Thus, adding (12) and (16), market clearing on exchange j requires:

$$Z_{user,j}(p_j) + Z_{arb,j}(p_j) = 0 \quad (17)$$

If the central exchange does not list the coin, arbitrageurs cannot trade, so we have $Z_{arb,j}(p_j) = 0$ and we require aggregate demand from users $Z_{user,j}(p_j)$ to equal 0. The following proposition solves for prices, volumes, and exchange profits when the central exchange does not list the coin.

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

$$p_{j,0}^* = \psi - \gamma_j \eta_j \quad (18)$$

¹⁵Practically, listing a crypto coin for trade can involve much higher costs than, for example, an equity exchange listing a new stock for trade: a crypto exchange listing a new coin must build systems to custody the coin, and to meet users' deposit and withdrawal requests. These processes can be costly, especially if the coin is on a blockchain that the exchange does not already work with. Listing a new coin may also exposes exchanges to increased regulatory risk, depending on local securities laws.

Expected squared trade quantity is:

$$\mathbb{E} [z_{i,0}^{*2}] = \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \sigma_{I,j}^2 \quad (19)$$

Exchange j 's profit from listing the coin is:

$$\pi_{j,0}^* = \frac{\tau_j}{2} \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \sigma_{I,j}^2 \quad (20)$$

Exchange j lists the coin if its cost of listing is lower than (20).

The following proposition solves prices, volumes, and exchange profits when the central exchange does list the coin.

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

$$p_{j,1}^* = \psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j \quad (21)$$

Expected squared trade quantity is:

$$\mathbb{E} [z_{i,1}^{*2}] = \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 (\mu_j^2 + \sigma_{A,j}^2) \right] \quad (22)$$

Exchange j 's profit from listing the coin is:

$$\pi_{j,1}^* = \frac{\tau_j}{2} \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 (\mu_j^2 + \sigma_{A,j}^2) \right] \quad (23)$$

Exchange j lists the coin if its cost of listing is lower than (23).

The intuitions behind Propositions 1 and 2 are as follows.

Prices. Expression (18) states that the price on exchange j , in the absence of the central 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 lower 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 central exchange lists the coin, arbitrageurs trade against this price gap, buy from the peripheral exchange and selling to the CEX at price ψ . Arbitrage cannot perfectly

close the gap, because arbitrageurs also face transaction fees and inventory costs. Comparing (18) and (21), arbitrageurs decrease the effect of inventory shocks on prices to a factor:

$$\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \quad (24)$$

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 central exchange lists.

Trade volumes and exchange profits. In the absence of a central exchange, trade on peripheral exchanges is generated only by the idiosyncratic component of inventory shocks: (19) 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 central exchange enters, trade is generated by both the idiosyncratic and aggregate components of inventory shocks, since arbitrageurs can buy on the central exchange and sell to the peripheral exchange. (22) shows that expected squared trade volume can be cleanly decomposed into the sum of (19), and an extra term reflecting the aggregate shock $\mu_j^2 + \sigma_{A,j}^2$, and the multiplier (24) capturing how active arbitrageurs are. Thus, expected squared trade volume of peripheral exchanges is strictly higher when the central exchange lists the coin. Since profits are proportional to squared trade volume, peripheral exchanges' profits are also higher when the central 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 coin before the central exchange does. These exchanges will experience increases in trading volume for the coin, when the central exchange lists the coin.*

Prediction 1 corresponds to Fact 1. This prediction follows directly from comparing (19) and (22), and the intuition that the aggregate component of inventory shocks also contributes to trade volume after the CEX lists the coin.

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

Prediction 2 corresponds to Fact 3. This prediction follows because, from (18), the variance of j 's prices when the central exchange does not list is

$$\sigma_\psi^2 + \gamma_j^2 \sigma_{A,j}^2 \quad (25)$$

Whereas the variance when the central exchange lists is, from (21), the smaller quantity:

$$\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \gamma_j^2 \sigma_{A,j}^2 \quad (26)$$

Intuitively, arbitrage with the central 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 central exchange, and the lower quantity

$$\left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \gamma_j^2 \sigma_A^2 \quad (27)$$

with the central exchange. Again, arbitrage with the central 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 coin will have a stronger incentive to list, after the central exchange has listed the coin.

Prediction 3. *Listings will tend to follow the central exchange: when the central exchange lists the coin, some peripheral exchanges which previously did not list the coin will choose to list the coin. Formally, peripheral exchanges' profit with the central exchange, (23), is greater than peripheral exchanges' profit without the central exchange, (20), so the set of peripheral exchanges which lists the coin is strictly larger after the central exchange enters.*

Prediction 3 corresponds to Fact 2. This prediction follows from (20) and (23). When the central exchange enters, expected profits on all peripheral exchanges increase. Thus, once the central exchange lists the coin, 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 coin will find it profitable to list the coin. This prediction essentially implies that the central 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 central exchange's price, and peripheral exchange j 's price, is:*

$$R_{j,CE}^2 = \frac{Cov^2(p_j^*, \psi)}{Var(p_j^*) Var(\psi)} = \frac{\sigma_\psi^2}{\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2} \quad (28)$$

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

$$\begin{aligned} R_{j,j'}^2 &= \frac{Cov^2(p_j^*, p_{j'}^*)}{Var(p_j^*) Var(p_{j'}^*)} \\ &= \frac{\sigma_\psi^2}{\left[\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2\right] \left[\sigma_\psi^2 + \left(\frac{\zeta_{j'} + \tau_{j'}}{\gamma_{j'} + \zeta_{j'} + 2\tau_{j'}}\right)^2 \gamma_{j'}^2 \sigma_{A,j'}^2\right]} \end{aligned} \quad (29)$$

Prediction 4. *We always have:*

$$R_{j,CE}^2 \geq R_{j,j'}^2 \quad (30)$$

That is, the correlation between the central 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 central exchange than they are with each other. This is because, from expression (21), each peripheral exchange's price is equal to the central exchange's price, plus an error term reflecting aggregate inventory shocks on the peripheral exchange which are imperfectly eliminated by arbitrageurs. Thus, $R_{j,CE}^2$ reflects the correlation of the central exchange price ψ , with a price which is ψ plus a noise term, whereas $R_{j,j'}^2$ 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 central exchange should tend to a larger increase in trade volumes when central exchanges list, and should have a stronger tendency to follow the central exchange’s listing decisions. Formally, we have:*

$$\frac{\partial R_{j,CE}^2}{\partial \zeta_j} = \frac{-2\sigma_\psi^2 \sigma_{A,j}^2 \gamma_j^2 (\zeta_j + \tau_j) (\gamma_j + \tau_j)}{\left[\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \gamma_j^2 \sigma_{A,j}^2 \right]^2 (\gamma_j + \zeta_j + 2\tau_j)^3} \leq 0 \quad (31)$$

$$\frac{\partial \Delta \mathbb{E}[z_{i1}^{*2}]}{\partial \zeta_j} = -\frac{2(\gamma_j + \tau_j)^2}{(\gamma_j + \zeta_j + 2\tau_j)^3} \frac{\mu_j^2 + \sigma_{A,j}^2}{\sigma_{I,j}^2} \leq 0 \quad (32)$$

$$\frac{\partial \Delta \pi_j^*}{\partial \zeta_j} = -\frac{2(\gamma_j + \tau_j)^2}{(\gamma_j + \zeta_j + 2\tau_j)^3} \frac{\mu_j^2 + \sigma_{A,j}^2}{\sigma_{I,j}^2} \leq 0 \quad (33)$$

Prediction 5 follows if there are differences in how “connected” peripheral exchanges are to the central 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 central exchange; the central exchange’s listings will tend to increase volumes more; and the central exchange’s listing decisions will increase the peripheral exchange’s profits more, implying that the peripheral exchange has a stronger incentive to “follow” the central 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 central 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 coin. 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 central 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 central exchange listings tend to associate with increased coin 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 central exchange will tend to alleviate this negative price pressure and raise coin prices. The listing pump effect could also be microfounded from a richer multi-period model, in which the entry of the central exchange increases market depth and decreases volatility of the coin, thus pushing prices upwards through a “liquidity premium” effect.

A number of other assumptions are made largely for expositional simplicity. We assume the central exchange has infinite depth; it is sufficient for our effects that the central exchange’s depth is finite, but much greater than peripheral exchanges’ depth. We assume arbitrageurs can only trade the peripheral exchange against the central 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 central 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 central exchange; in practice, there are a number of bigger exchanges which likely behave more like central exchanges, and smaller exchanges which behave more like peripheral exchanges. In our empirical analysis, we treat Binance and Coinbase as central exchanges, and the long tail of smaller exchanges as peripheral.

5 Empirical Tests

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 central exchange, so the correlations between peripheral exchange prices and central exchange prices should be greater than the correlations between different peripheral exchange prices. To test this prediction, we calculate the return correlations of BTC prices 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 CDF of return correlations, separately for exchange pairs between peripheral exchanges and either Binance or Coinbase, and for pairs involving only peripheral exchanges. The CDFs of peripheral-central price correlations lie below and to the right of the CDF of peripheral-peripheral price correlations; hence, as the model predicts, price correlations between peripheral and central exchanges tend to be larger than price correlations between pairs of peripheral exchanges. Quantitatively, the 25th, 50th, and 75th percentiles of price correlations among all exchanges are 0.59, 0.78, and 0.91. The percentiles are 0.76, 0.87, and 0.94 for pairs involving Binance, and 0.64, 0.83, and 0.94 for pairs involving Coinbase.¹⁶

5.2 Interaction between Price Correlations, Volume Increases, and Listing Following

Prediction 5 posits that peripheral exchanges with higher price correlations to the central exchange will see greater volume increases and exhibit a stronger tendency to follow the central 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 central exchange experience larger volume increases when the central exchange lists, we begin with the following flexible DID specification:

¹⁶The average return correlation across all exchanges at day level in our dataset is 0.7. However, the low cross-correlation is largely driven by low ranking exchanges in our representative dataset. Our data produces price correlations very close to those in the literature, once we subset to the same set of exchanges. [Makarov and Schoar \(2020\)](#) use data from Kaiko, and restrict their sample to 17 large exchanges; they measure an average hourly cross-correlation of USD BTC returns of 0.83. After restricting our sample to the same set of exchanges, we measure a very similar average cross-correlation of 0.85. Return correlations are sensitive to the time horizon at which prices are measured: at the daily level, [Makarov and Schoar](#) report a higher average cross-correlation of 0.95.

$$\log(\text{Volume}_{c,e,t}) = \sum_{k=-31}^{31} \left(\beta_k^{\text{low}} \times \text{treat}_{c,k,t} + \beta_k^{\text{high-low}} \times \text{treat}_{c,k,t} \times \text{HighCorrelation}_e \right) + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t} \quad (34)$$

HighCorrelation_e is a dummy variable indicating whether the peripheral exchange e has a BTC price correlation relative to central exchanges higher than the median level. All other variables are defined above. β_k^{low} measures how much coin volume of exchanges with low correlation relative to central exchanges increases after central exchange listing events, compared to coins that did not experience central exchange listings, while $\beta_k^{\text{low}} + \beta_k^{\text{high-low}}$ measures how much coin volume of exchanges with high correlation increases. Still, we set observations that are exactly 30 days before large exchange listings as the reference group, which means that β_{-30}^{low} and $\beta_{-30}^{\text{high-low}}$ are always 0. Similarly, we keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges.

We plot estimated coefficients of β_k^{low} for the low correlation group and $\beta_k^{\text{low}} + \beta_k^{\text{high-low}}$ for the high correlation group with k from -30 to 31. The results are shown in Figure 8. The coefficients for both high and low correlation groups are small in magnitude prior to central exchange listings. Directly after central exchange listings, both groups experience a large increase in trading volume, compared with coin-exchange pairs that did not experience a listing event. Notably, the magnitude differs in the first several days after large exchange listings depending on which group the exchange belongs to. Exchanges with high correlations with central exchanges experience a larger increase in trading volume of coins that are listed by a central exchange, relative to exchanges with low correlations. Quantitatively, volume for high correlation group increases by around 184% and 123%, respectively, for coin-exchange pairs following a coin listing by Binance and Coinbase, and volume for low correlation group increases by around 121% and 95%. It suggests the mechanism of our model, peripheral exchanges with high correlations with central exchanges rely more heavily on arbitraging activity between central and peripheral exchanges, and thus experience a larger volume increase after central exchange listings.

We then estimate the following variant of Specification (2) using the same filtered sample:

$$\begin{aligned}
\text{Log}(\text{Volume})_{c,e,t} = & \beta_1 \text{Listing}(0\text{-}30 \text{ days})_{c,t} + \beta_2 \text{Listing}(0\text{-}30 \text{ days})_{c,t} \times \text{Correlation}_e + \\
& \beta_3 \text{Listing}(> 30 \text{ days})_{c,t} + \beta_4 \text{Listing}(> 30 \text{ days})_{c,t} \times \text{Correlation}_e + \\
& \beta_5 \text{PreThreedayListing}_{c,t} + \beta_6 \text{PreThreedayListing}_{c,t} \times \text{Correlation}_e + \\
& \beta_7 \text{Correlation}_e + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t}
\end{aligned} \tag{35}$$

All variables are defined as in previous specifications. Correlation_e is the correlation between Bitcoin returns between exchange e and a given central 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 central exchange will tend to have larger volume increases within 30 days and 30 days after central exchange listings.

The results are shown in Table 5. We find that β_2 are positive and mostly significant, confirming that exchanges with stronger price correlations with the central exchange experience larger volume increases when the central 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 2.4% and 0.51% increase in daily average volume within 30 days of their listings. However, we do not find that β_4 are significant in most specifications. There are potentially two reasons for this. First, the overall long-term volume increase is small due to entry and cannibalization effects, as shown in Fact 1, and so are the effects for exchanges with differential correlations. Second, exchanges with lower correlations might be able to engage other market makers who have strong connections with large exchanges, thereby reducing the disparity in volume increases between more correlated and less correlated exchanges.

5.2.2 Correlation Structure and Listing Following

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

$$\text{Listing Following Probability} = \frac{\# \text{ Listings within } [0,30] \text{ days window}_e}{\# \text{ Listings}_e} \tag{36}$$

Intuitively, Listing Following Probability_e measures the tendency for exchange *e* to list very quickly (within 30 days) after a central exchange lists. When Listing Following Probability_e is 1, in all cases where *e* lists a coin, the listing occurs from 0 to 30 days after the central exchange’s listing. ¹⁷

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

$$ListingFollowingProbability_e = \alpha + \beta Correlation_e + \epsilon_e \quad (37)$$

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

As shown in Figure 9, *Correlation_e* and *ListingFollowingProbability_e* are positively correlated, whether we treat Binance or Coinbase as the central exchange: that is, exchanges with stronger positive price correlations with a given central exchange tend to follow the central 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 central exchange has a 0.23% (0.08%) higher probability of following the central exchange’s listing decisions, respectively, for Binance (Coinbase).

5.3 Volume Decomposition

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

$$\begin{aligned} \log(Volume_{c,t}[inc]) &= \sum_{k=-31}^{31} \left(\beta_k^{inc} \right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \\ \log(Volume_{c,t}[inc + ent]) &= \sum_{k=-31}^{31} \left(\beta_k^{inc} + \beta_k^{ent} \right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \\ \log(Volume_{c,t}[inc + ent + cen]) &= \sum_{k=-31}^{31} \left(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen} \right) \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t} \end{aligned} \quad (38)$$

¹⁷We only keep coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges. We also drop peripheral exchanges with less than five listings in the sample period.

Where, as in (3), $treat_{c,k,t}$ are dummies for time since a central exchange has listed a coin. We construct $\log(Volume_{c,t})$ in three successively broader ways: using only volume on incumbent exchanges, which list coin c at least 30 days before central exchange listings; using incumbents as well as any new exchanges which enter following the central exchange’s entry, excluding volume on the central exchange itself; and using volume on incumbents, entrants, and the central 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 central exchange entry are generated by volume increases on incumbents, as in Fact 1; the entry of new exchanges which follow the central exchange, as in Fact 2; and the direct effect of trade volumes generated by the central exchange. The results are shown in Figure 10. 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 central exchange listing after k days, into three components: β_k^{cen} , the direct effect of trade volume on the central 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 1.26, 0.33, and 0.31 respectively. Exponentiating the sum and subtract one, $\exp(\beta_k^{inc} + \beta_k^{ent} + \beta_k^{cen}) - 1$, Coinbase’s listing associates with a total volume increase of 568pp, relative to pre-listing incumbent volume. We infer that this effect can be thought of as the net result of a 125pp volume increase among incumbents, a 33pp increase in volume from entrants, and a 31pp increase directly on the central exchange, where the latter two numbers are percentages as a fraction of pre-listing incumbent volume.

We show our estimates of β_k^{inc} , β_k^{ent} , and β_k^{cen} in the bottom row of Figure 10. Taking the exponent of the blue line on the top row, Binance listings are associated with an average 1264pp increase in total volume within first 30 days. From the second row, within first 30 days of listings, the effect for Binance decomposes into a 67pp increase of incumbent volume, a 127pp increase associated with entrants, and a 275pp increase on Binance itself. Coinbase listings are associated with an average 236pp increase in total volumes within first 30 days. Within first 30 days of listing, incumbent volumes increase by 33pp, compared to 70pp for entrants, and 51pp for Coinbase itself.

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

$$LR_k = \frac{\exp(\beta_k^{inc} + \beta_k^{ent}) - 1}{\exp(\beta_k^{cen}) - 1} \quad (39)$$

The ratio (39) can be interpreted as saying, for each dollar of trade volume a central exchange creates by listing a coin, how many dollars of trade volume are generated on other (incumbent and entrant) exchanges. Within first 30 days of listing, we estimate that average LR_k is equal to 0.94 for Binance, and 2.39 for Coinbase: thus, every dollar captured by Binance listing generates \$0.94 of trade on other exchanges, and every dollar captured by Coinbase generates \$2.39 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 central exchanges, a large fraction of the volume increases associated with listings are due to the indirect effects on other exchanges.

5.4 Arbitrage and Price Discovery

Finally, we examine the nature of arbitrage through the lens of price gaps between central and peripheral exchanges. Our model is not dynamic, and does not make explicit predictions about how price gaps evolve over time. However, if price gaps tend to reflect idiosyncratic inventory shocks at peripheral exchanges, and if we think of arbitrageurs' inventory capacity as being constrained in the short-run, in the long run arbitrage trade should be able to eliminate price gaps between peripheral and central exchanges. We would thus expect price gaps to mean-revert towards 0 over time. To test this hypothesis, we define the price gap between peripheral exchange e and the central exchange, for coin c at time t , as:

$$PriceGap_{c,e,t} \equiv p_{c,e,t} - p_{c,t}^{Cen}$$

That is, $PriceGap_{c,e,t}$ is the difference in log prices between e and the central exchange for coin c at time t . We expect price gaps to be mean-reverting; hence, we model $PriceGap_{c,e,t}$ as:

$$\begin{aligned} \Delta PriceGap_{c,e,t} &= \delta_{c,e}^{PriceGap} (\mu - PriceGap_{c,e,t-1}) + \epsilon_{c,e,t}^{PriceGap} \\ &= -\delta_{c,e}^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \\ &= \beta_{c,e}^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \end{aligned} \quad (40)$$

Where $\beta_{c,e}^{PriceGap}$ is the coefficient of interest and determines the speed of convergence. We expect $\beta_{c,e}^{PriceGap}$ to range from -1 to 0, with a value of -1 indicating that the price deviation can fully recover in the next period, and 0 indicating that convergence never occurs.

We estimate $\beta_{c,e}^{PriceGap}$ for each coin-exchange pair. The distribution of estimated $\beta_{c,e}^{PriceGap}$ coefficients is shown in the top row of Figure 11. The estimated coefficients concentrates at -1. The median coefficient is -0.87 (-0.9), implying that a 10% price gap on average decays to 1.3% (1%) on average within a day treating Binance (Coinbase) as the central exchange; in other words, it takes roughly 1.12 (1) days for a price gap to tenth. The 75th percentile $\beta_{c,e}^{PriceGap}$ coefficient implies a price gap tenth-life of 2.46 (1.91) days, and the 25th percentile implies a tenth-life of 0.57 (0.41) days, again for Binance (Coinbase) as the central exchange. In other words, these figures show that price gaps revert towards 0 fairly quickly, rarely lasting more than a few days.¹⁸

In an alternative specification, we estimate a single $\beta^{PriceGap}$ pooling across all peripheral exchanges and coins:

$$\Delta PriceGap_{c,e,t} = \beta^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \quad (41)$$

In the context of dynamic and cross-sectional dependent panel data, the OLS estimator is biased and inconsistent. Therefore, we follow the suggestion of (Pesaran, 2006) and estimate a mean group estimator for $\beta^{PriceGap}$. This estimator proves to be consistent when both N and T are large, which is valid for our data. It follows a two-stage estimation. In the first stage, we similarly estimate the individual equation for each coin-exchange pair, but we also include three lagged cross-sectional averages of dependent variables $\sum_{c,e} \Delta PriceGap_{c,e,t-1}$, $\sum_{c,e} \Delta PriceGap_{c,e,t-2}$, and $\sum_{c,e} \Delta PriceGap_{c,e,t-3}$. In the second stage, we use the mean of individual estimators as the estimate for $\beta^{PriceGap}$ and the variance of individual estimators as the estimate for $Var(\beta^{PriceGap})$. These coefficients are reported in Columns (1) and (2) of Table 6. We reach a similar conclusion that price gaps revert towards to 0 very quickly. It takes roughly 1.66 (1.52) days for a price gap to decrease by a factor of 10, treating Binance (Coinbase) as the central exchange, respectively.

Next, we examine the direction of peripheral-central price gap convergence. Price gaps can mean-revert towards 0 in two ways: peripheral exchange prices $p_{c,e,t}$ can revert towards central exchange prices $p_{c,t}^{Cen}$, or vice versa. Once again, our model does not formally make predictions about price dynamics; however, informally, if price gaps are driven by idiosyncratic

¹⁸We are also able to statistically reject that unit roots exist for the vast majority of $PriceGap_{c,e,t}$ series: in the top row of Appendix Figure D.2, we show the distribution of test statistics from a Dickey-Fuller unit root test, and show that we can reject the existence of a unit root for over 91% (94)% of coin-exchange pairs at the 99% confidence interval.

preference shocks on peripheral exchanges, and if central exchanges tend to be more liquid than peripheral exchanges, we would expect reversion in $PriceGap_{c,e,t}$ to be driven mainly by peripheral exchange prices moving towards central exchange prices.

To test this, we simply regress peripheral and central exchange log returns on one-day lagged price gaps for each coin-exchange pair separately:

$$\Delta p_{c,e,t} = \beta_{c,e}^p PriceGap_{c,e,t-1} + \epsilon_{c,e,t} \quad (42)$$

$$\Delta p_{c,t}^{cen} = \beta_{c,e}^{p^{cen}} PriceGap_{c,e,t-1} + \epsilon_{c,e,t} \quad (43)$$

The coefficient estimates are respectively shown in the middle and bottom rows of Figure 11. The middle row shows that the distribution of coefficients from Specification (42) is concentrated at negative values: when $PriceGap_{c,e,t}$ is positive, peripheral exchange prices tend to decrease towards central exchange prices. Quantitatively, taking the median coefficient, a 10% price gap predicts a 6.9% (11%) decrease in return on peripheral exchanges on the following day, treating Binance (Coinbase) as the central exchange, respectively.¹⁹

In an alternative specification, we estimate a single β^p using the mean group estimator; these coefficients are reported in Columns (3) and (4) of Table 6. We arrive at a similar finding that peripheral exchange prices move towards central exchange prices. A 10% price gap predicts a 7.2% (11%) decrease in return on peripheral exchanges on the following day, considering Binance (Coinbase) as the central exchange, respectively.

The bottom row shows the distribution of coefficients from (43). The distribution of estimated $\beta_{c,e}^{p^{cen}}$'s is roughly symmetric around 0 with median value of 0 (-0.21) when treating Binance (Coinbase) as the central exchange; thus, $PriceGap_{c,e,t}$ does not positively or negatively predict returns on central exchanges.²⁰ In an alternative specification, we similarly estimate a single $\beta^{p^{cen}}$ using the mean group estimator; these coefficients are reported in Columns (5) and (6) of Table 6. We similarly conclude that peripheral exchange price gaps have essentially no predictive power for future returns on central exchanges.

Together, these results show that peripheral-central price gaps fairly quickly mean-revert, with peripheral exchange prices moving towards central exchange prices. While not an explicit

¹⁹We are also able to statistically reject that price gap has no predictive power on peripheral exchange prices for most coin-exchange pairs: in the middle row of Appendix Figure D.2, we show the distribution of t-test statistics, and show that we can reject $\beta_{c,e}^p = 0$ for over 60% (70)% of coin-exchange pairs at the 99% confidence interval.

²⁰We are unable to statistically reject that price gap has no predictive power on central exchange prices for most coin-exchange pairs: in the bottom row of Appendix Figure D.2, we show the distribution of t-test statistics, and show that we cannot reject $\beta_{c,e}^{p^{cen}} = 0$ for over 38% (39)% of coin-exchange pairs at the 99% confidence interval.

prediction of our model, these results are intuitively consistent with a world in which price gaps are generally driven by idiosyncratic noise trading shocks to peripheral exchanges, which are arbitrated away over time, and which do not convey substantial information about the future value of coins.

5.5 Robustness Checks

We conduct a number of robustness checks of our results. One concern is that exchanges – especially small exchanges – may be falsifying trade volumes, either through false reporting or wash trading. In order to drive our results, exchanges would not only need to be falsifying volume; they would have to increase the amount they falsify volume sharply after central exchanges list, which we view as implausible. Moreover, volume falsification cannot directly explain why smaller exchanges tend to follow large exchanges’ listing decisions, or why the volume increase and listing following effects are stronger for small exchanges with prices that are more correlated with large exchanges.

To provide further evidence that volume falsification is not the main driver of our results, in Appendix Table D.1 and Figure D.1, we restrict to 29 exchanges analyzed in Cong et al. (2020), and analyze the responses of volumes in their three groups of exchanges.²¹ The classification reflects the intensity, to what extent the exchanges are regulated and thus falsely report their trading volume. The regulated exchanges are supervised by New York State Department of Financial Services (NYSDFS), and they are classified into regulated exchanges. The unregulated exchanges are further classified into two different tiers based on their web traffic. Cong et al. (2020) argue that wash trading should be most prevalent in tier-2 exchanges, less prevalent in tier-1 exchanges, and least prevalent in regulated exchanges. We find similar increases in volume in all 3 groups of exchanges, providing some evidence against the idea that our findings are driven entirely by wash trading or falsified volume.²²

Our main hypothesis is that the volume increase and listing following effects are primarily driven by arbitrage flows between peripheral and central exchanges. An alternative hypothesis, which we call the “attention channel”, is that when central exchanges list coins, the market pays more attention to these coins; this generates spillover effects to trade volumes of the

²¹These exchanges include: (1) regulated exchanges: Bitstamp, Coinbase, Gemini; (2) tier-1 unregulated exchanges: Binance, Bittrex, Bitfinex, Hitbtc, Huobi, Kucoin, Liquid, Okex, Poloniex, Zb; (3) tier-2 unregulated exchanges: Bgogo, Biki, Bitz, Coinbene, Dragonex, Lbank, Mxc, Fcoin, Exmo, Coinmex, Bibox, Bitmart, Bitmax, Coinegg, Digifinex, and Gateio.

²²Similarly, we keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges.

coin on other peripheral exchanges, and thus gives peripheral exchanges increased incentives to list these coins. While we cannot fully rule out the attention channel, it does not explain the evidence in Section 5.2, that the volume increase and listing following effects are larger for peripheral exchanges with stronger price correlations with central exchanges.

To provide further evidence that the attention channel is not the sole driver of our results, we control for a measure of investor attention to coins based on Google Trends. Following studies such as Da, Engelberg and Gao (2011), we collect the Google search volumes for the ticker symbols of coins in our sample. This data provides an index ranging from 0 to 100, based on the search volumes for the keywords during the study period. Then, we calculate the abnormal search frequency as the control variable by taking the difference between the current date log Google search volumes and the median log Google search volumes in the past one week. This gives us a coin-date level measure of attention, which we call $Attention_{c,t}$. In Appendix Table D.2, we repeat the volume specification of Table 2, in the main text, controlling for $Attention_{c,t}$. If the attention channel were fully driving our results, we would expect that coins listed by central exchanges would not experience unusual increases in volume, after controlling for the volume increases associated with increased attention. However, Appendix Table D.2 shows that central exchange listings increase unusual increases in volume even after controlling for attention, providing some evidence that the attention channel is not the sole driver of our results.²³

Our evidence here does not show that the attention channel is fully irrelevant: rather, it only suggests that there are patterns in the data that are consistent with our arbitrage hypothesis, which are difficult to explain using the attention channel. Finally, we note that many of the implications of our results for regulators are similar regardless of whether the attention channel or the arbitrage channel are at work: in either case, large exchanges can influence market outcomes in a manner which is understated by their market shares.

6 Discussion

Our stylized facts allow us to construct a narrative regarding the market structure of the international cryptoasset marketplace. Market participants around the world wish to trade the same cryptoassets; however, market participants reside in jurisdictions with very different fiat payment rails and financial regulators, and who may have different demands regarding the design and complexity of their crypto exchanges. A natural solution is for a plethora

²³Similarly, we keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges.

of “broker”-like entities – peripheral exchanges in our setting – to emerge, each focused on logistically allowing customers with particular preferences in a particular fiat jurisdiction to trade fiat for crypto. But markets are most liquid when combined; it is thus natural for the large number of broker-like entities trading identical cryptoassets to be linked by professional arbitrageurs to large central exchanges, who play a role similar to inter-dealer marketplaces or clearinghouses in other markets. The large central exchanges specialize in pooling liquidity, leaving the task of integrating with specific payment rails, and presenting specific interfaces to customers, to peripheral exchanges and arbitrageurs. The end result of this system is that, after accounting for peripheral exchange fees and arbitrageur spreads, consumers in any given jurisdiction trade with all other jurisdictions in deep international markets for cryptoassets.

Our results imply that large exchanges have significant influence over outcomes in the international market for cryptoassets. Regulators like the SEC determine which assets can legally be traded in their local jurisdictions. In international cryptoasset markets, instead, centralized exchanges play a regulator-like “leader” role, influencing the assets that are traded broadly across exchanges through their listing decisions. From a classical welfarist perspective, central exchanges may list too few coins, being unable to capture the externality profits they generate for peripheral exchanges and arbitrageurs. More broadly, it is unclear central exchange have incentives to list coins satisfying regulators’ other preferences, for example, that coins should be related to legal businesses with relatively transparent disclosure policies. Regulators may wish to monitor, and perhaps restrict, the currently essentially unlimited ability of central exchanges to list coins for trade.

7 Conclusion

In this paper, we argued that crypto exchanges do not compete like firms in oligopolistic product markets. Instead, small “peripheral” exchanges behave like brokerages: they are essentially “costly windows” for their customers to access the deeper liquidity available on large “central” exchanges. Our results suggest that the large central 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 central exchange to list a coin induces large indirect trading volume increases, both on incumbent exchanges which have previously listed the coin, and on entrant exchanges which follow the central exchange’s listing decisions. Exchanges currently make their listing decisions with substantial discretion and little oversight; policymakers may wish to monitor or regulate central 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*.
- Benedetti, Hugo, and Ehsan Nikbakht.** 2021. “Returns and network growth of digital tokens after cross-listings.” *Journal of Corporate Finance*, 66: 101853.
- 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, Ke Tang, Yanxin Wang, and Xi Zhao.** 2023. “Inclusion and democratization through web3 and defi? initial evidence from the ethereum ecosystem.” *Working Paper, National Bureau of Economic Research*.
- 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*.
- Da, Zhi, Joseph Engelberg, and Pengjie Gao.** 2011. “In search of attention.” *The journal of finance*, 66(5): 1461–1499.
- Du, Songzi, and Haoxiang Zhu.** 2017. “What is the optimal trading frequency in financial markets?” *The Review of Economic Studies*, 84(4): 1606–1651.
- Félez-Viñas, Ester, Luke Johnson, and Tālis J Putniņš.** 2022. “Insider Trading in Cryptocurrency Markets.” *Available at SSRN 4184367*.
- 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.
- Pesaran, M Hashem.** 2006. “Estimation and inference in large heterogeneous panels with a multifactor error structure.” *Econometrica*, 74(4): 967–1012.
- Shams, Amin.** 2020. “The structure of cryptocurrency returns.” *Fisher College of business working paper*, , (2020-03): 011.
- 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.
- von Luckner, Clemens Graf, Carmen M Reinhart, and Kenneth Rogoff.** 2023. “Decrypting new age international capital flows.” *Journal of Monetary Economics*.
- 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 Trading 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

The figure depicts estimates from Specification (1):

$$\log(\text{Volume}_{c,e,t}) = \sum_{k=-31}^{31} \beta_k \times \text{treat}_{c,k,t} + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t}$$

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, η_t represents day fixed effects, and $\gamma_{e,t}$ represents exchange-time fixed effects. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Observations exactly 30 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin-exchange pair and time 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. We filter the sample to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges, eliminating the mechanical effects associated with coins initially listed on large exchanges. 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):

$$\Delta\#Exchanges_{c,t} = \sum_{k=-31}^{31} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

along with 95% confidence intervals. The outcome variable is the number of exchanges that list coin c between day $t - 1$ and t . δ_c represents coin fixed effects, and η_t represents day fixed effects. We filter the sample to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges, eliminating the mechanical effects associated with coins initially listed on large exchanges. Observations exactly 30 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin and time level. Data source: Cryptotick.




Figure 5: Large Exchange Listings and Price Dispersion

This figure illustrates how the interquartile range of coin prices changes around listing events. For each coin affected by a listing, we measure the median, 25th, and 75th percentile prices across incumbent exchanges, and normalize all percentiles by the median. Incumbent exchanges are defined as exchanges that list coin c at least 30 days before its listing on a large exchange. The figure shows the average of the normalized quantiles across coins affected by listings, for Binance and Coinbase respectively. 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):

$$Dispersion_{c,t} = \sum_{k=-31}^{31} \beta_k \times treat_{c,k,t} + \delta_c + \eta_t + \epsilon_{c,t}$$

along with 95% confidence intervals. The outcome variable is coin-day level price dispersion, calculated as the standard deviation of log prices across incumbent exchanges for coin c at day t . Incumbent exchanges are defined as exchanges that list coin c at least 30 days before its listing on a large exchange, or exchanges whose coins have never been listed by large exchanges. δ_c represents coin fixed effects, and η_t represents day fixed effects. Observations exactly 30 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin and time 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.

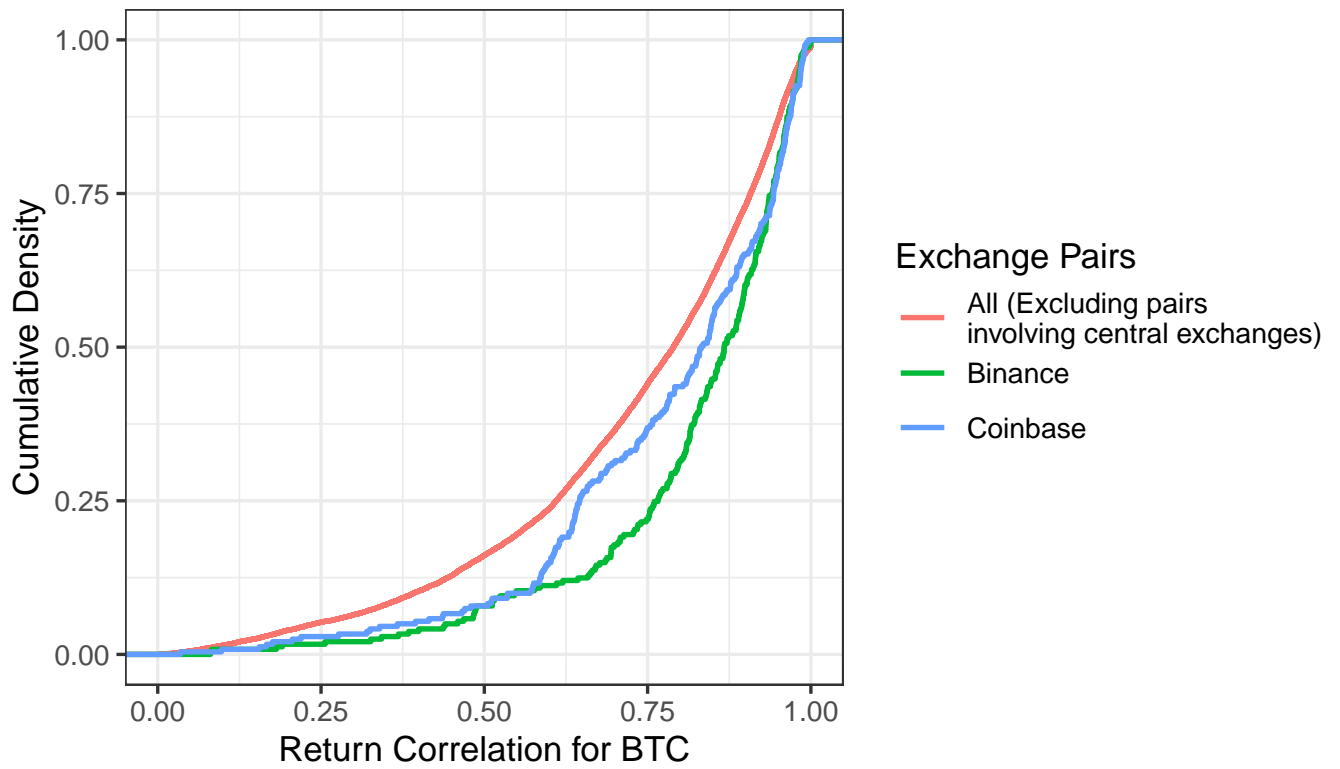


Figure 8: Central Exchange Listings, Trade Volumes, and Price Correlations

The figure depicts estimates β_k^{low} and $\beta_k^{low} + \beta_k^{high-low}$ from Specification (34):

$$\log(\text{Volume}_{c,e,t}) = \sum_{k=-31}^{31} \left(\beta_k^{low} \times \text{treat}_{c,k,t} + \beta_k^{high-low} \times \text{treat}_{c,k,t} \times \text{HighCorrelation}_e \right) + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t}$$

along with 95% confidence intervals. The outcome variable is log coin-exchange-day level logarithmic dollarized volume. *HighCorrelation* is a dummy variable indicating whether peripheral exchange *e* has a higher median Bitcoin return correlation with central exchanges. $\delta_{c,e}$ represents coin-exchange fixed effects, η_t represents day fixed effects, and $\gamma_{e,t}$ represents exchange-time fixed effects. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Observations exactly 30 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin-exchange pair and time level. Data source: Cryptotick.

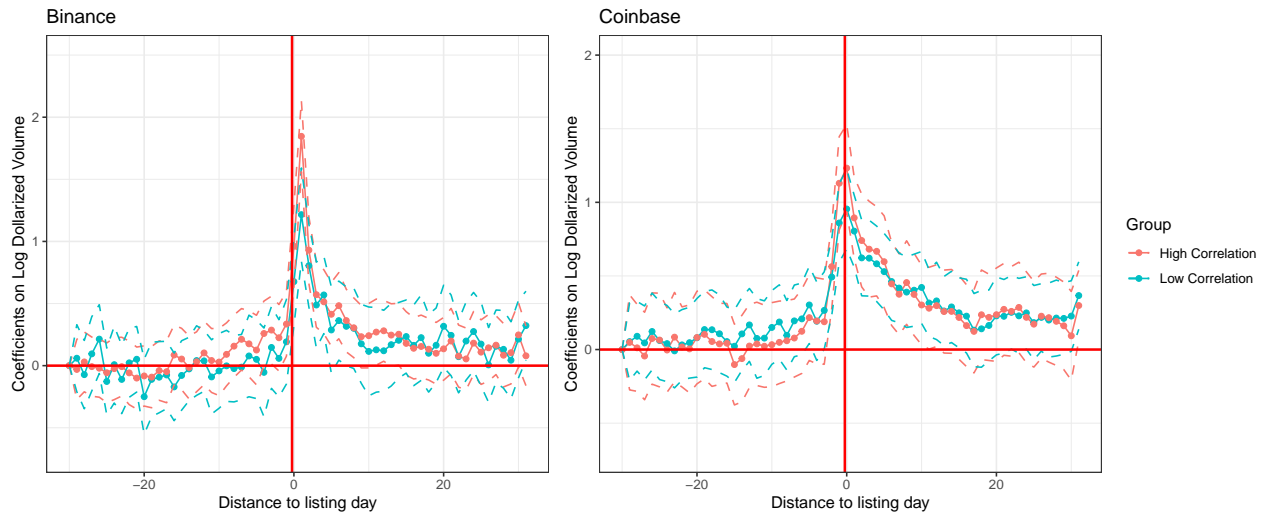


Figure 9: Price Correlations and Listing Following Tendency

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 (36). 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.

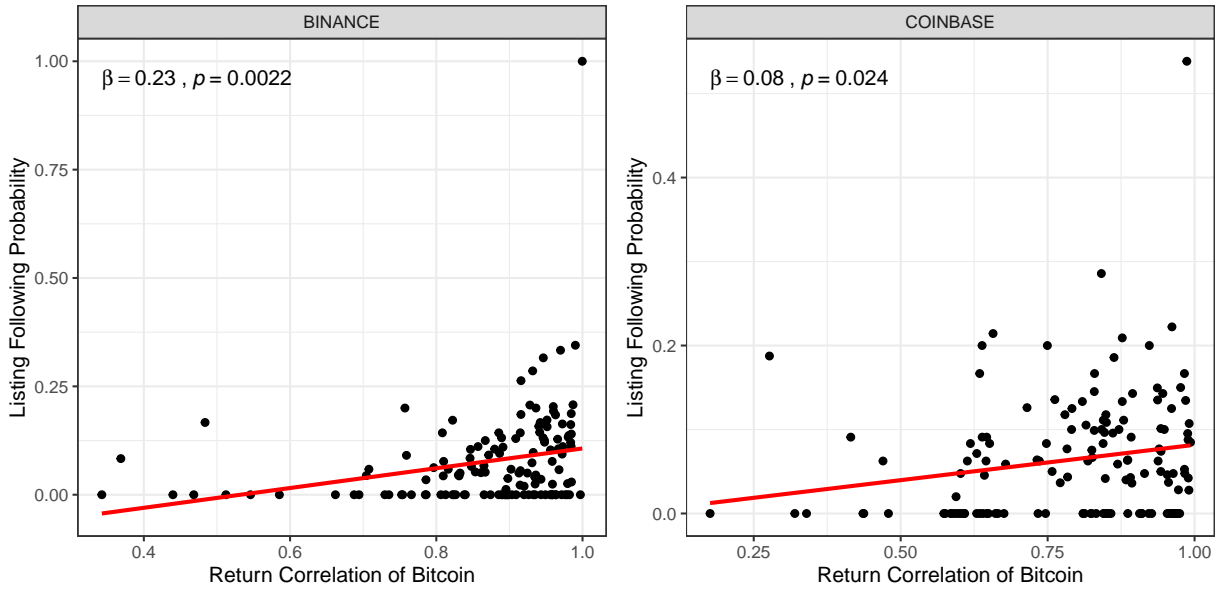


Figure 10: Volume Increase Decomposition

The top row of the figure depicts estimates from Specification (38):

$$\log(\text{Volume}_{c,t}[\text{inc}]) = \sum_{k=-31}^{31} (\beta_k^{\text{inc}}) \times \text{treat}_{c,k,t} + \delta_c + \eta_t + \gamma_{e,t} + \epsilon_{c,t}$$

$$\log(\text{Volume}_{c,t}[\text{inc} + \text{ent}]) = \sum_{k=-31}^{31} (\beta_k^{\text{inc}} + \beta_k^{\text{ent}}) \times \text{treat}_{c,k,t} + \delta_c + \eta_t + \gamma_{e,t} + \epsilon_{c,t}$$

$$\log(\text{Volume}_{c,t}[\text{inc} + \text{ent} + \text{cen}]) = \sum_{k=-31}^{31} (\beta_k^{\text{inc}} + \beta_k^{\text{ent}} + \beta_k^{\text{cen}}) \times \text{treat}_{c,k,t} + \delta_c + \eta_t + \gamma_{e,t} + \epsilon_{c,t}$$

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 30 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 $(\beta_k^{\text{inc}} + \beta_k^{\text{ent}})$, using incumbent plus entrant volume as the dependent variable. The blue line shows estimates of $(\beta_k^{\text{inc}} + \beta_k^{\text{ent}} + \beta_k^{\text{cen}})$, using incumbent, entrant, and central exchange volume as the dependent variable. δ_c represents coin fixed effects, η_t represents day fixed effects, and $\gamma_{e,t}$ represents exchange-time fixed effects. Observations exactly 30 days before central exchange listings are set as the reference group. Standard errors are clustered at the coin and time level. We then recover β_k^{inc} , β_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.

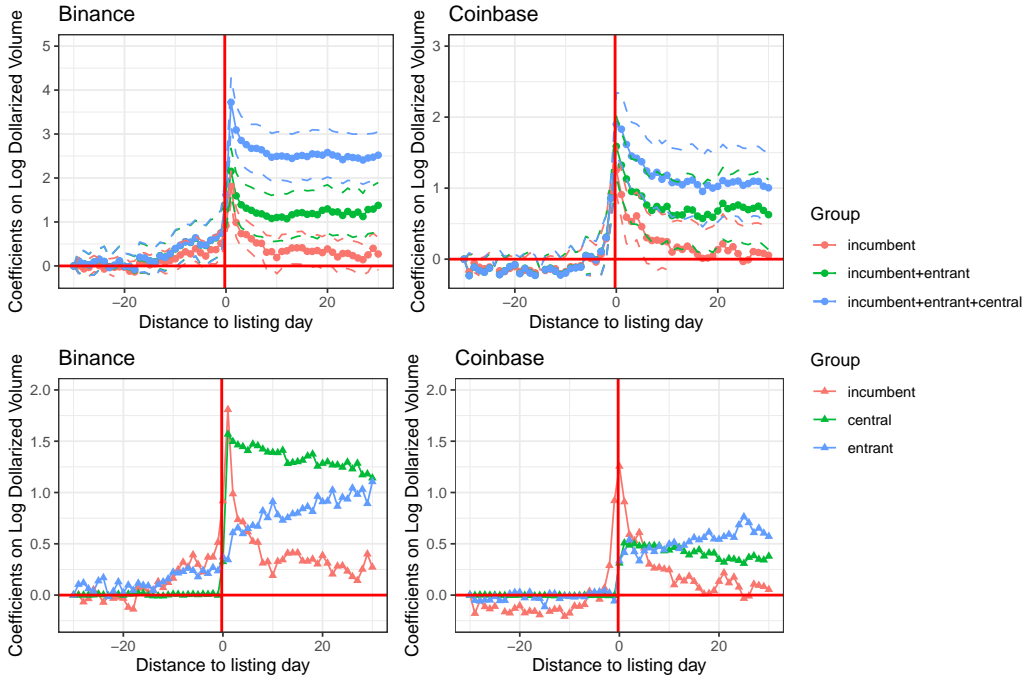


Figure 11: Price Gap Reversion

This figure displays, from the top to bottom row, the distribution of estimates from Specifications (40), (42), and (43) for all coin-exchange pairs:

$$\begin{aligned}\Delta PriceGap_{c,e,t} &= \beta_{c,e}^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \\ \Delta p_{c,e,t} &= \beta_{c,e}^p PriceGap_{c,e,t-1} + \epsilon_{c,e,t} \\ \Delta p_{c,t}^{cen} &= \beta_{c,e}^{p^{cen}} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}\end{aligned}$$

In all cases, one data point is one peripheral exchange-coin pair. The top row shows the distribution of estimated $\beta_{c,e}^{PriceGap}$, where the outcome variable is the change in price gap of coin c between peripheral exchange e and central exchange (either Binance or Coinbase) at day t . The red vertical line shows $\beta^{PriceGap} = -1$, the zero-persistence benchmark, where price gaps on a given day fully revert on the following day. The middle row shows the distribution of estimated $\beta_{c,e}^p$, where the outcome variable is the price change in coin c on peripheral exchange e at day t . The bottom row shows the distribution of estimated $\beta_{c,e}^{p^{cen}}$, and the outcome variable is the price change in coin c on central exchange (either Binance or Coinbase) at day t . The red vertical lines in the middle and bottom row denote $\beta_{c,e}^p = 0$ and $\beta_{c,e}^{p^{cen}} = 0$, respectively. Data source: Cryptotick.

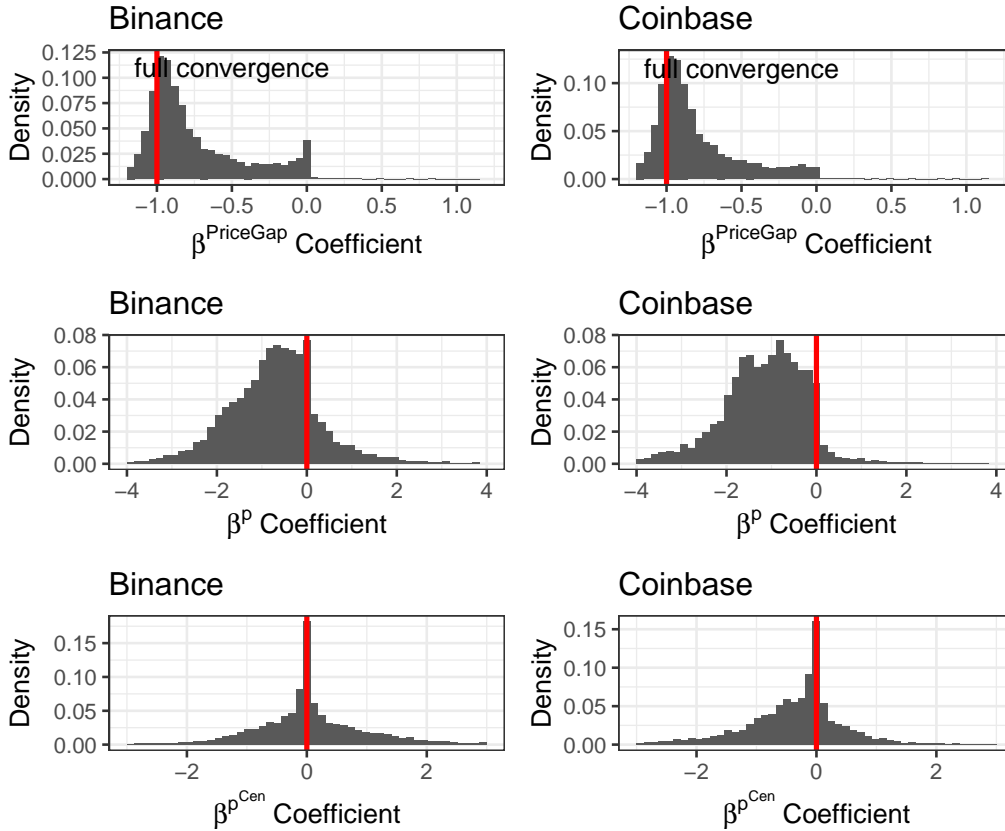


Table 1: Summary Statistics

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
	<i>BINANCE</i>						<i>COINBASE</i>					
BTC Return Correlation	242	0.81	0.2	0.76	0.87	0.94	243	0.77	0.24	0.64	0.83	0.94
Listing Following Prob	156	0.089	0.11	0	0.06	0.14	155	0.07	0.081	0	0.059	0.11

Panel B: Coin Level

	Obs.	Mean	SD	p25	p50	p75
Price Dispersion	389413	0.092	0.15	0.011	0.027	0.11
Net Listings	445453	0.015	0.24	0	0	0
Log (Volume)	453856	14	4.4	13	15	17
Listing (0-30 days) Binance	453856	0.018	0.13	0	0	0
Listing (> 30 days) Binance	453856	0.46	0.5	0	0	1
Pre Three-day Listing Binance	453856	0.0013	0.036	0	0	0
Listing (0-30 days) Coinbase	453856	0.0092	0.096	0	0	0
Listing (> 30 days) Coinbase	453856	0.14	0.34	0	0	0
Pre Three-day Listing Coinbase	453856	8.80E-04	0.03	0	0	0

Panel C: Coin-Exchange Level

	Obs.	Mean	SD	p25	p50	p75
Log (Volume)	5511168	12	3.9	9.9	12	15
Listing (0-30 days) Binance	5511168	0.012	0.11	0	0	0
Listing (> 30 days) Binance	5511168	0.72	0.45	0	1	1
Pre Three-day Listing Binance	5511168	6.20E-04	0.025	0	0	0
Listing (0-30 days) Coinbase	5511168	0.013	0.11	0	0	0
Listing (> 30 days) Coinbase	5511168	0.4	0.49	0	0	1
Pre Three-day Listing Coinbase	5511168	0.001	0.032	0	0	0

Table 2: Large Exchange Listings and Trade Volumes

This table presents estimates from Specification (2):

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

$\log(\text{Volume}_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t . $\text{Listing}(0-30 \text{ days})_{c,t}$ and $\text{Listing}(> 30 \text{ days})_{c,t}$ 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. $\text{PreThreedayListing}_{c,t}$ 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 $t + 3$. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Columns (1) to (4) are results 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 and time level. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	Log Dollarized Volume							
	Binance				Coinbase			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	0.45*** (0.08)	0.49*** (0.09)	0.39*** (0.07)	0.50*** (0.06)	0.77*** (0.05)	0.73*** (0.06)	0.69*** (0.05)	0.72*** (0.04)
Listing (> 30 days)	0.28*** (0.09)	0.24** (0.11)	0.15 (0.10)	0.35*** (0.08)	0.69*** (0.06)	0.61*** (0.08)	0.61*** (0.06)	0.69*** (0.05)
Pre Three-day Listing	0.41*** (0.09)	0.44*** (0.10)	0.33*** (0.08)	0.42*** (0.07)	1.0*** (0.08)	0.97*** (0.08)	0.94*** (0.08)	0.95*** (0.07)
Coin FE	Yes	Yes	No	No	Yes	Yes	No	No
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exchange FE	Yes	No	No	No	Yes	No	No	No
Country FE	No	Yes	No	No	No	Yes	No	No
Coin-Exchange Pair FE	No	No	Yes	Yes	No	No	Yes	Yes
Exchange FE \times Day FE	No	No	No	Yes	No	No	No	Yes
Adjusted R ²	0.64	0.47	0.79	0.85	0.61	0.43	0.78	0.85
Observations	2,008,881	1,955,789	2,008,881	2,008,881	4,013,259	3,882,427	4,013,259	4,013,259

Table 3: Large Exchange Listings and Small Exchange Listings

This table shows estimates from specification (6):

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

$\Delta\#Exchanges_{c,t}$ is the net change in the number of exchanges (excluding Binance and Coinbase themselves) which list coin c in time t . $Listing(0-30\text{ days})_{c,t}$ 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 $t + 3$. We filter the sample to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges, eliminating the mechanical effects associated with coins initially listed on large exchanges. Columns (1) and (2) are results based on listings on Binance and Coinbase, respectively. Standard errors are clustered at the coin and time 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
Model:	(1)	(2)
Listing (0-30 days)	0.08*** (0.01)	0.08*** (0.01)
Listing (> 30 days)	0.01*** (0.003)	0.0005 (0.005)
Pre Three-day Listing	0.02* (0.01)	0.04* (0.02)
Coin FE	Yes	Yes
Day FE	Yes	Yes
Adjusted R ²	0.09	0.10
Observations	418,600	476,858

Table 4: Large Exchange Listings and Price Dispersion

This table shows estimates from specification (6):

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

$Dispersion_{c,t}$ is the standard deviation of log price across exchanges of coin c at time t , across exchanges. $Listing(0-30 \text{ days})_{c,t}$ 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 $t + 3$. We filter the sample to include only coins that were listed on any exchange for at least 30 days before being listed on large exchanges or were not listed by large exchanges, eliminating the mechanical effects associated with coins initially listed on large exchanges. Columns (1) and (2) are results based on listings on Binance and Coinbase, respectively. Standard errors are clustered at the coin and time 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
Model:	(1)	(2)
Listing (0-30 days)	-0.02 (0.01)	-0.007 (0.008)
Listing (> 30 days)	-0.03** (0.01)	-0.02* (0.010)
Pre Three-day Listing	-0.005 (0.01)	0.004 (0.008)
Coin FE	Yes	Yes
Day FE	Yes	Yes
Adjusted R ²	0.51	0.47
Observations	280,202	378,296

Table 5: Coin Volume, Exchange Correlations, and Listing Decisions

This table presents estimates of Specification (35). The dependent variables are log trading volume. $Listing(0-30 \text{ days})_{c,t}$, $Listing(> 30 \text{ days})_{c,t}$, 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. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Columns (1) to (4) are results 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 and time 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			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	-0.16 (1.3)	-0.36 (1.6)	-3.4*** (0.96)	-1.8** (0.83)	-1.0** (0.51)	-1.2* (0.63)	-0.77* (0.42)	0.28 (0.38)
Listing (0-30 days) \times Correlation	0.66 (1.4)	0.96 (1.7)	4.1*** (1.0)	2.4*** (0.89)	2.1*** (0.58)	2.2*** (0.72)	1.7*** (0.48)	0.51 (0.43)
Listing (> 30 days)	0.005 (1.3)	0.34 (1.5)	-3.1*** (1.2)	-0.24 (0.94)	0.30 (0.79)	0.06 (0.91)	-1.2** (0.55)	0.93* (0.52)
Listing (> 30 days) \times Correlation	0.32 (1.4)	-0.006 (1.6)	3.5*** (1.3)	0.63 (1.0)	0.46 (0.90)	0.65 (1.0)	2.1*** (0.63)	-0.28 (0.58)
Pre Three-day Listing	-0.83 (1.6)	-1.4 (1.9)	-3.9*** (1.2)	-2.5** (1.0)	-1.1** (0.57)	-1.6** (0.66)	-0.76 (0.46)	0.25 (0.42)
Pre Three-day Listing \times Correlation	1.3 (1.7)	2.0 (2.0)	4.5*** (1.3)	3.1*** (1.1)	2.5*** (0.64)	2.9*** (0.76)	2.0*** (0.53)	0.82* (0.47)
Correlation		8.3*** (0.94)				9.1*** (0.49)		
Coin FE	Yes	Yes	No	No	Yes	Yes	No	No
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exchange FE	Yes	No	No	No	Yes	No	No	No
Country FE	No	Yes	No	No	No	Yes	No	No
Coin-Exchange Pair FE	No	No	Yes	Yes	No	No	Yes	Yes
Exchange FE \times Day FE	No	No	No	Yes	No	No	No	Yes
Adjusted R ²	0.64	0.49	0.79	0.85	0.61	0.47	0.78	0.85
Observations	1,994,658	1,941,566	1,994,658	1,994,658	4,006,407	3,875,575	4,006,407	4,006,407

Table 6: Price Gap Reversion

This table presents estimates of Specification (40), (42), and (43) for all coin-exchange pairs:

$$\begin{aligned}\Delta PriceGap_{c,e,t} &= \beta^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \\ \Delta p_{c,e,t} &= \beta^p PriceGap_{c,e,t-1} + \epsilon_{c,e,t} \\ \Delta p_{c,t}^{cen} &= \beta^{p^{cen}} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}\end{aligned}$$

We follow the suggestion of (Pesaran, 2006) and estimate a consistent mean group estimator for $\beta^{PriceGap}$, β^p , and $\beta^{p^{cen}}$. It follows a two-stage estimation. In the first stage, we similarly estimate the individual equation for each coin-exchange pair, but we also include three lagged cross-sectional averages of dependent variables $\sum_{c,e} \Delta PriceGap_{c,e,t-1}$, $\sum_{c,e} \Delta PriceGap_{c,e,t-2}$, and $\sum_{c,e} \Delta PriceGap_{c,e,t-3}$. In the second stage, we use the mean of individual estimators as the estimate for $\beta^{PriceGap}$ and the variance of individual estimators as the estimate for $Var(\beta^{PriceGap})$. Columns (1) to (2) are results with change in price gap as the dependent variable. Columns (3) to (4) are results with peripheral exchange price change as the dependent variable. Columns (5) to (6) are results with central exchange price change as the dependent variable. Odd (even) number columns are results treating Binance (Coinbase) as the central exchange. Standard errors in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick.

Dependent Variables:	$\Delta PriceGap$		Δp		Δp^{Cen}	
	Binance	Coinbase	Binance	Coinbase	Binance	Coinbase
Model:	(1)	(2)	(3)	(4)	(5)	(6)
$PriceGap_{c,e,t-1}$	-0.75*** (0.01)	-0.78*** (0.01)	-0.72*** (0.02)	-1.1*** (0.03)	0.04 (0.03)	-0.33*** (0.03)
R ²	0.3	0.67	0.15	0.59	0.08	0.52
Observations	4028134	2156482	4028134	2156482	4028134	2156482

Internet Appendix

A Supplementary Material for Section 2

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 coins, and to direct the network to take actions such as transfer coins to other wallets. Crypto “miners”, which build the blockchain by inserting proposed transactions in new “blocks”, are incentivized to mine by newly minted crypto coins 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 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.²⁴

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

²⁴One class of exceptions to this rule is that the administrators of certain coins, such as the Circle (USDC) and Tether (USDT) USD stablecoins, include code in the “smart contracts” governing their coins which allows them to freeze the funds of certain “blacklisted” wallets. These coin 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 coin, administrators include the capability to blacklist coins, and the majority of crypto coins do not have built-in blacklist functionality.

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.²⁵

²⁵See [Devon Zuegel](#).

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 ²⁶; 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 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. We drop stablecoins and fiat currencies.
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. ²⁷We

²⁶These 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.

²⁷Specifically, 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 18% from 5,511,168 to 6,511,290.

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 coin, arbitrageurs have no activity. Market clearing requires aggregate demand from all users on exchange j to equal 0. Hence, from (12), we need:

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

Solving for p_j , we have:

$$p_{j,0}^* = \psi - \gamma_j \eta_j \quad (44)$$

This is (18).

Trade quantities. To solve for expected squared trade quantity, note that user i 's trade quantity is (11). Plugging in for $\psi - p_j$ using (44), we have:

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

Thus, we have:

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

Plugging in for $p_{j,0}^*$ using (44) and simplifying, we have:

$$\begin{aligned} &= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \int_{-\infty}^{\infty} (\eta_j - x_{i,0})^2 dF_{x_{i,0}}(x) \\ &= \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \sigma_{I,j}^2 \end{aligned}$$

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 coin, 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 (12) and (16), 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:

$$p_{j,1}^* = \psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j \quad (45)$$

This is (21).

Trade quantities. To solve for expected squared trade quantity, note that user i 's trade quantity is (11). Plugging in for $\psi - p_j$ using (45), 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:

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

Plugging in for prices using (45), we have:

$$\begin{aligned} &= \mathbb{E} \left[\int_{-\infty}^{\infty} \left[\frac{-\gamma_j}{\gamma_j + \tau_j} x_{i,0} + \frac{\psi - \left(\psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j \right)}{\gamma_j + \tau_j} \right]^2 dF_{x_{i,0}}(x) \right] \\ &= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \mathbb{E} \left[\int_{-\infty}^{\infty} \left(-x_{i,0} + \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \eta_j \right)^2 dF_{x_{i,0}}(x) \right] \\ &= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \mathbb{E} \left[\int_{-\infty}^{\infty} \left[x_{i,0}^2 - 2 \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \eta_j x_{i,0} + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \eta_j^2 \right] dF_{x_{i,0}}(x) \right] \\ &= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \mathbb{E} \left[\left(\sigma_{I,j}^2 + \eta_j^2 \right) - 2 \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \eta_j^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \eta_j^2 \right] \\ &= \left(\frac{\gamma_j}{\gamma_j + \tau_j} \right)^2 \mathbb{E} \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \right)^2 \eta_j^2 \right] \end{aligned}$$

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

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:

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

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:

$$\begin{aligned} R_{j,CE}^2 &= \frac{Cov^2(p_j^*, \psi)}{Var(p_j^*) Var(\psi)} \\ &= \frac{Cov^2\left(\psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j, \psi\right)}{Var\left(\psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j\right) Var(\psi)} \\ &= \frac{Cov^2(\psi, \psi)}{\left[Var(\psi) + Var\left(-\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j\right)\right] Var(\psi)} \\ &= \frac{\sigma_\psi^2}{\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2} \end{aligned} \tag{46}$$

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

$$\begin{aligned}
R_{j,j'}^2 &= \frac{\text{Cov}^2(p_j^*, p_{j'}^*)}{\text{Var}(p_j^*) \text{Var}(p_{j'}^*)} \\
&= \frac{\text{Cov}^2\left(\psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j, \psi - \frac{\zeta_{j'} + \tau_{j'}}{\gamma_{j'} + \zeta_{j'} + 2\tau_{j'}} \gamma_{j'} \eta_{j'}\right)}{\text{Var}\left(\psi - \frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j\right) \text{Var}\left(\psi - \frac{\zeta_{j'} + \tau_{j'}}{\gamma_{j'} + \zeta_{j'} + 2\tau_{j'}} \gamma_{j'} \eta_{j'}\right)} \\
&= \frac{\text{Cov}^2(\psi, \psi)}{\left[\text{Var}(\psi) + \text{Var}\left(-\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j} \gamma_j \eta_j\right) \right] \left[\text{Var}(\psi) + \text{Var}\left(-\frac{\zeta_{j'} + \tau_{j'}}{\gamma_{j'} + \zeta_{j'} + 2\tau_{j'}} \gamma_{j'} \eta_{j'}\right) \right]} \\
&= \frac{\sigma_\psi^2}{\left[\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2 \right]} \frac{\sigma_\psi^2}{\left[\sigma_\psi^2 + \left(\frac{\zeta_{j'} + \tau_{j'}}{\gamma_{j'} + \zeta_{j'} + 2\tau_{j'}}\right)^2 \gamma_{j'}^2 \sigma_{A,j'}^2 \right]}
\end{aligned} \tag{47}$$

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:

$$R_{j,CE}^2 \geq R_{j,j'}^2$$

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 (46):

$$\frac{\partial R_{j,CE}^2}{\partial \zeta_j} = \frac{-2\sigma_\psi^2 \sigma_{A,j}^2 \gamma_j^2 (\zeta_j + \tau_j) (\gamma_j + \tau_j)}{\left[\sigma_\psi^2 + \left(\frac{\zeta_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \gamma_j^2 \sigma_{A,j}^2 \right]^2 (\gamma_j + \zeta_j + 2\tau_j)^3} \leq 0 \tag{48}$$

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

$$\begin{aligned}
\Delta \mathbb{E} [z_{i,1}^{*2}] &= \frac{\mathbb{E} [z_{i1}^{*2}] - \mathbb{E} [z_{i0}^{*2}]}{\mathbb{E} [z_{i0}^{*2}]} \\
&= \frac{\left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \left[\sigma_{I,j}^2 + \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 (\mu_j^2 + \sigma_{A,j}^2) \right] - \left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \sigma_{I,j}^2}{\left(\frac{\gamma_j}{\gamma_j + \tau_j}\right)^2 \sigma_{I,j}^2} \\
&= \left(\frac{\gamma_j + \tau_j}{\gamma_j + \zeta_j + 2\tau_j}\right)^2 \frac{\mu_j^2 + \sigma_{A,j}^2}{\sigma_{I,j}^2}
\end{aligned} \tag{49}$$

The volume increase of peripheral exchanges is also decreasing in the arbitrage costs ζ_j from (49):

$$\frac{\partial \Delta \mathbb{E} [z_{i,1}^{*2}]}{\partial \zeta_j} = -\frac{2(\gamma_j + \tau_j)^2}{(\gamma_j + \zeta_j + 2\tau_j)^3} \frac{\mu_j^2 + \sigma_{A,j}^2}{\sigma_{I,j}^2} \leq 0 \quad (50)$$

Similarly, 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 .

D Robustness Checks

Figure D.1: Large Exchange Listings and Trade Volumes: Wash Trading

The figure depicts estimates $\beta_k^{regulated}$, $\beta_k^{regulated} + \beta_k^{tier1-regular}$, and $\beta_k^{regulated} + \beta_k^{tier2-regular}$ from Specification (1):

$$\begin{aligned} \log(\text{Volume}_{c,e,t}) = & \sum_{k=-31}^{31} (\beta_k^{regulated} \times \text{treat}_{c,k,t} + \beta_k^{tier1-regular} \times \text{treat}_{c,k,t} \times \text{Tier1}_e \\ & + \beta_k^{tier2-regular} \times \text{treat}_{c,k,t} \times \text{Tier2}_e) \\ & + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t} \end{aligned}$$

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, η_t represents day fixed effects, and $\gamma_{e,t}$ represents exchange-time fixed effects. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Observations exactly 30 days before large exchange listings are set as the reference group. Standard errors are clustered at the coin-exchange pair and time level. Data source: Cryptotick.

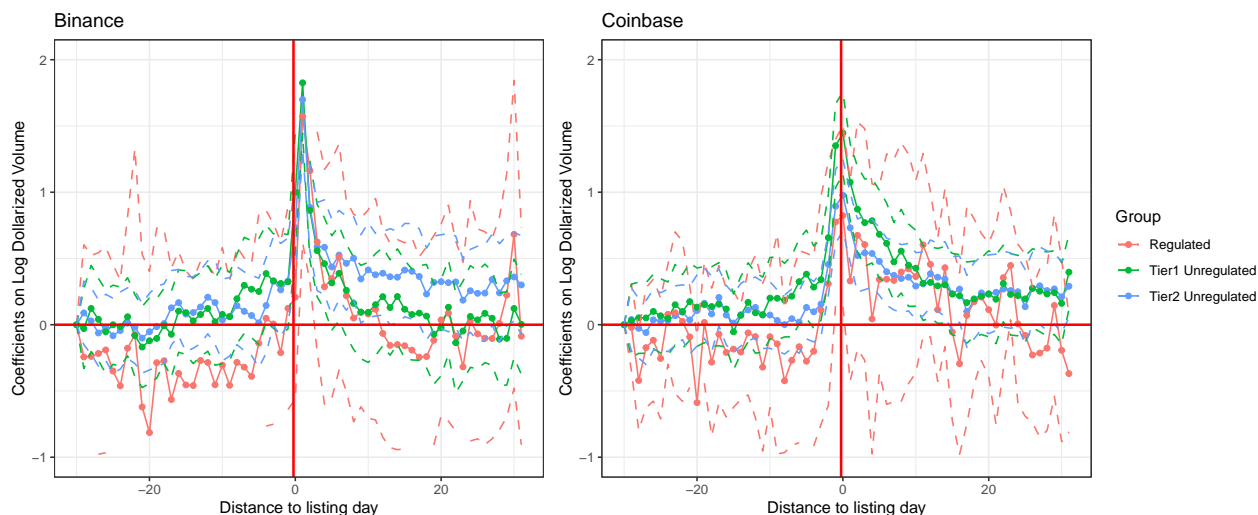


Figure D.2: Price Gap Reversion: Statistical Power

The figure displays the distribution of statistical power from Specification (40), (42), and (43) for all coin-exchange pairs:

$$\begin{aligned}\Delta PriceGap_{c,e,t} &= \beta_{c,e}^{PriceGap} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}^{PriceGap} \\ \Delta p_{c,e,t} &= \beta_{c,e}^p PriceGap_{c,e,t-1} + \epsilon_{c,e,t} \\ \Delta p_{c,t}^{cen} &= \beta_{c,e}^{p^{cen}} PriceGap_{c,e,t-1} + \epsilon_{c,e,t}\end{aligned}$$

In all cases, one data point represents one peripheral exchange-coin pair. The top row displays the distribution of the statistical power of the estimated $\beta_{c,e}^{PriceGap}$ from a Dickey-Fuller unit root test. The outcome variable is the change in the price gap of coin c between peripheral exchange e and central exchange (either Binance or Coinbase) at day t . The gray area to the right of the red vertical line indicates the critical value range where we can reject the existence of a unit root at a 99% confidence interval when the sample size is 30. The middle row shows the distribution of the t-statistics of the estimated $\beta_{c,e}^p$, and the outcome variable is the price change in coin c on peripheral exchange e at day t . The bottom row shows the distribution of t-statistics of the estimated $\beta_{c,e}^{p^{cen}}$, and the outcome variable is the price change in coin c on central exchange (either Binance or Coinbase) at day t . The gray area around the red vertical line indicates the critical value range where we can reject that $\beta_{c,e}^p$ and $\beta_{c,e}^{p^{cen}}$ are equal to 0 at a 99% confidence interval when the sample size is 30. Data source: Cryptotick.

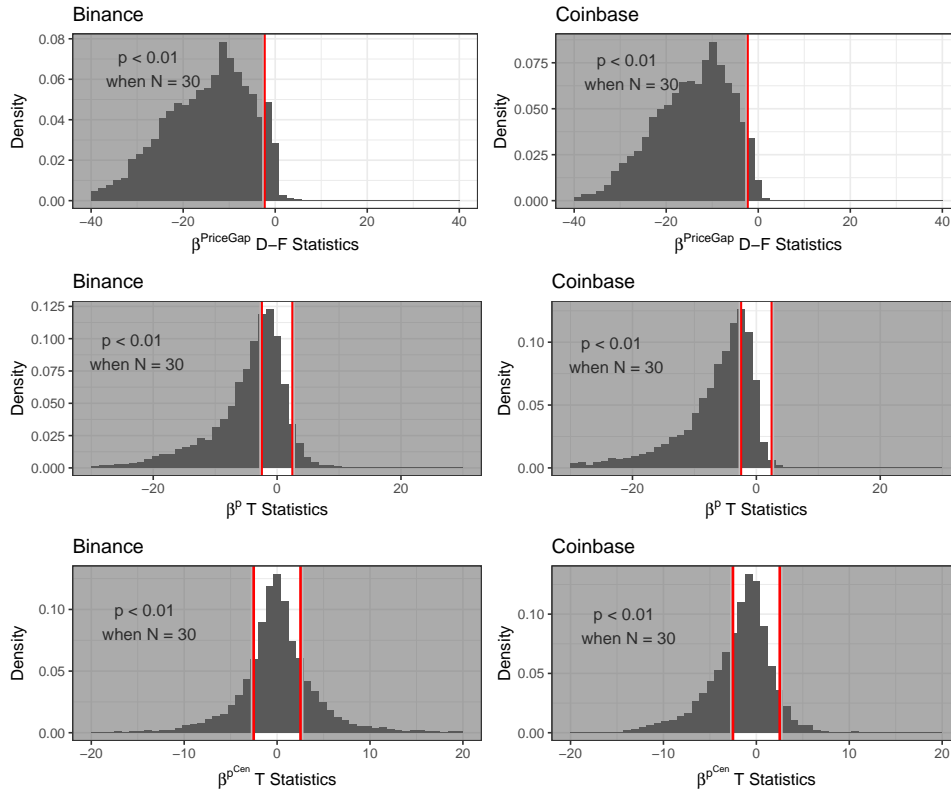


Table D.1: Large Exchange Listings and Trade Volumes: Wash Trading

This table presents estimates from Specification (2):

$$\begin{aligned} \log(\text{Volume}_{c,e,t}) = & \beta_1 \text{Listing}(0\text{-}30 \text{ days})_{c,t} + \beta_2 \text{Listing}(0\text{-}30 \text{ days})_{c,t} \times \text{Regulated}_e + \beta_3 \text{Listing}(0\text{-}30 \text{ days})_{c,t} \times \text{Tier1}_e \\ & + \beta_4 \text{Listing}(> 30 \text{ days})_{c,t} + \beta_5 \text{Listing}(> 30 \text{ days})_{c,t} \times \text{Regulated}_e + \beta_6 \text{Listing}(> 30 \text{ days})_{c,t} \times \text{Tier1}_e \\ & + \beta_7 \text{PreThreedayListing}_{c,t} + \beta_8 \text{PreThreedayListing}_{c,t} \times \text{Regulated}_e + \beta_9 \text{PreThreedayListing}_{c,t} \times \text{Tier1}_e \\ & + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t} \end{aligned}$$

$\log(\text{Volume}_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t . $\text{Listing}(0\text{-}30 \text{ days})_{c,t}$ and $\text{Listing}(> 30 \text{ days})_{c,t}$ 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. $\text{PreThreedayListing}_{c,t}$ 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 $t + 3$. Regulated_e and Tier1_e are dummy variables that indicate whether exchange e belongs to the regulated or Tier 1 group of exchanges, as defined in Cong et al. (2020). We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Columns (1) to (4) are results 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 and time 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			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	0.60*** (0.19)	0.62*** (0.19)	0.56*** (0.19)	0.68*** (0.16)	0.76*** (0.13)	0.79*** (0.13)	0.85*** (0.11)	0.78*** (0.10)
Listing (> 30 days)	0.42** (0.20)	0.41** (0.20)	0.54** (0.23)	0.54*** (0.20)	0.59*** (0.13)	0.66*** (0.14)	0.89*** (0.13)	0.72*** (0.10)
Pre Three-day Listing	0.42** (0.19)	0.45** (0.19)	0.43** (0.18)	0.55*** (0.15)	0.87*** (0.14)	0.90*** (0.14)	0.99*** (0.12)	0.94*** (0.10)
Listing (0-30 days) × Regulated	1.4*** (0.51)	1.4** (0.58)	-0.20 (0.32)	-0.34 (0.24)	0.70 (0.45)	0.49 (0.47)	0.15 (0.34)	0.13 (0.32)
Listing (0-30 days) × Tier1	-0.09 (0.24)	-0.23 (0.25)	0.13 (0.23)	-0.04 (0.20)	0.08 (0.16)	0.01 (0.17)	-0.08 (0.15)	0.03 (0.12)
Listing (> 30 days) × Regulated	1.2*** (0.41)	1.3*** (0.49)	-0.23 (0.52)	-0.42 (0.32)	0.27 (0.49)	0.02 (0.52)	-0.43 (0.44)	-0.41** (0.18)
Listing (> 30 days) × Tier1	0.04 (0.24)	-0.07 (0.25)	-0.23 (0.29)	-0.13 (0.25)	0.26* (0.16)	0.18 (0.17)	-0.28 (0.18)	0.06 (0.14)
Pre Three-day Listing × Regulated	1.3** (0.55)	1.4** (0.58)	-0.29 (0.31)	-0.37 (0.27)	0.75 (0.50)	0.52 (0.52)	0.29 (0.35)	0.13 (0.28)
Pre Three-day Listing × Tier1	0.19 (0.25)	0.10 (0.26)	0.35 (0.23)	0.18 (0.20)	0.29 (0.18)	0.23 (0.20)	0.15 (0.15)	0.24* (0.13)
Regulated		2.1*** (0.58)				1.3*** (0.47)		
Tier1		-0.62 (0.39)				-0.08 (0.32)		
Coin FE	Yes	Yes	No	No	Yes	Yes	No	No
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exchange FE	Yes	No	No	No	Yes	No	No	No
Country FE	No	Yes	No	No	No	Yes	No	No
Coin-Exchange Pair FE	No	No	Yes	Yes	No	No	Yes	Yes
Exchange FE × Day FE	No	No	No	Yes	No	No	No	Yes
Adjusted R ²	0.59	0.57	0.74	0.80	0.55	0.52	0.74	0.79
Observations	795,315	746,335	795,315	795,315	1,625,092	1,498,509	1,625,092	1,625,092

Table D.2: Coin Volume and Listing: Controlling for Attention

This table presents estimates from Specification (2):

$$\log(\text{Volume}_{c,e,t}) = \beta_1 \text{Listing}(0\text{-}30 \text{ days})_{c,t} + \beta_2 \text{Listing}(> 30 \text{ days})_{c,t} + \beta_3 \text{PreThreedayListing}_{c,t} + \beta_4 \text{Attention}_{c,t} + \delta_{c,e} + \eta_t + \gamma_{e,t} + \epsilon_{c,e,t}$$

$\log(\text{Volume}_{c,e,t})$ denotes the log of dollarized coin trading volume for coin c and exchange e at day t . $\text{Listing}(0\text{-}30 \text{ days})_{c,t}$ and $\text{Listing}(> 30 \text{ days})_{c,t}$ 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. $\text{PreThreedayListing}_{c,t}$ 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 $t + 3$. $\text{Attention}_{c,t}$ measures the difference between the log Google search volumes at date t and the median log Google search volumes from $t - 7$ to $t - 1$. We keep only those incumbent coin-exchange pairs that were listed at least 30 days before their listing on large exchanges, or pairs with coins that have not been listed by large exchanges, in order to identify the listing effect on incumbent exchanges. Columns (1) to (4) are results 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 and time level. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively. Data source: Cryptotick, Google Search.

Dependent Variables:	Log Dollarized Volume							
	Binance				Coinbase			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Listing (0-30 days)	0.48*** (0.09)	0.52*** (0.10)	0.35*** (0.08)	0.47*** (0.07)	0.80*** (0.06)	0.76*** (0.07)	0.73*** (0.05)	0.76*** (0.04)
Listing (> 30 days)	0.24** (0.10)	0.19 (0.12)	0.02 (0.11)	0.22** (0.09)	0.69*** (0.07)	0.62*** (0.08)	0.61*** (0.07)	0.69*** (0.05)
Pre Three-day Listing	0.41*** (0.10)	0.45*** (0.11)	0.28*** (0.09)	0.36*** (0.09)	1.0*** (0.08)	0.99*** (0.09)	0.96*** (0.06)	0.96*** (0.08)
Attention	0.19*** (0.02)	0.18*** (0.02)	0.17*** (0.01)	0.16*** (0.01)	0.23*** (0.01)	0.23*** (0.01)	0.23*** (0.007)	0.23*** (0.01)
Coin FE	Yes	Yes	No	No	Yes	Yes	No	No
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exchange FE	Yes	No	No	No	Yes	No	No	No
Country FE	No	Yes	No	No	No	Yes	No	No
Coin-Exchange Pair FE	No	No	Yes	Yes	No	No	Yes	Yes
Exchange FE \times Day FE	No	No	No	Yes	No	No	No	Yes
Adjusted R ²	0.67	0.47	0.80	0.86	0.62	0.44	0.78	0.85
Observations	1,493,239	1,451,244	1,493,239	1,493,239	3,059,751	2,955,470	3,059,751	3,059,751