# Pricing Transparency in European Crypto Trading: An Empirical Assessment of Execution Costs Across Major Retail Exchanges

Felix Rihacek / Felix Voigt / Dominik Heilbock

felix.rihacek@tum.de
03747776

TUM School of Management – Technical University of Munich

**Corporate Governance and Capital Markets Law**
**Prof. Dr. jur. Philipp Maume, S.J.D. (La Trobe)**

**23.3.2026**

# 1) Bibliography

Aerts, Senne/Alexandra Born/Zakaria Gati/Urszula Kochanska/Claudia Lambert/Elisa Reinhold/Van Der Kraaij Anton: Just another crypto boom? Mind the blind spots, in: European Central Bank, 21.05.2025, [online] https://www.ecb.europa.eu/press/financial-stability-publications/fsr/special/html/ecb.fsrart202505_01~62255f2625.en.html.

Chainalysis: European Crypto Adoption Highlights a Patchwork of Convergence and Transformation, in: Chainalysis, 23.10.2025, [online] https://www.chainalysis.com/blog/europe-crypto-adoption-2025/.

Deutsche Börse: DAX listed blue chips, in: Deutsche Börse Group, n.d., [online] https://www.cashmarket.deutsche-boerse.com/cash-en/trading/Equities/dax-listed-blue-chips-xetra.

Digital Euro Association: DEA MiCAR Tracker - Track Electronic Money Token (EMT) issuers and Crypto Asset Service Providers (CASPs) across Europe, o. D., [online] https://micatracker.digital-euro-association.de/ (accessed 31.07.2025).

Dyhrberg, Anne H./Sean Foley/Jiri Svec: How investible is Bitcoin? Analyzing the liquidity and transaction costs of Bitcoin markets, in: Economics Letters, vol. 171, 24.07.2018, [online] doi:10.1016/j.econlet.2018.07.032, pp. 140–143.

ESMA: Markets in Crypto-Assets Regulation (MiCAR), o.D., [online] https://www.esma.europa.eu/esmas-activities/digital-finance-and-innovation/markets-crypto-assets-regulation-mica (accessed 22.02.2026).

Kaiko Research: How is crypto liquidity fragmentation impacting markets? - Kaiko - Research, in: Kaiko Research, 22.08.2025, [online] https://research.kaiko.com/insights/how-is-crypto-liquidity-fragmentation-impacting-markets.

Makarov, Igor/Antoinette Schoar: Trading and arbitrage in cryptocurrency markets, in: Journal of Financial Economics, vol. 135, no. 2, 11.07.2019, [online] doi:10.1016/j.jfineco.2019.07.001, pp. 293–319.

# 2) List of Figures



# 3) List of abbreviations

API: Application Programming Interface

MiCAR: Markets in crypto assets regulation

CASPs: Crypto asset providers

CTP: Consolidated Tape Provider

BSDEX: Börse Stuttgart Digital Exchange

MiFID II: Markets in Financial Instruments Directive

RT: Round-Trip

# 4) Introduction

## a) Background: Retail crypto markets & lack of price transparency

While crypto assets have increasingly become a mainstream investment class for European retail investors in recent years[1] [2], the underlying trading infrastructure remains highly fragmented. In stark contrast to traditional equity markets, which benefit from concentrated price discovery and standardized trading conditions the crypto ecosystem lacks such uniformity [3].

Consequently, retail investors typically do not trade directly on centralized order books but instead interact with a diverse array of brokers or hybrid exchange-broker platforms.

This structural difference is significant. In traditional stock trading, both incumbent banks and neobrokers ultimately execute client orders on more transparent and highly liquid venues such as Xetra or other regulated European exchanges, where round-trip implicit transaction costs for blue-chip equities typically range from 3 to 10 basis points[4]. Bid–ask spreads are extremely tight and execution quality is widely understood. Even though fintech neobrokers disrupted the fee structure of stock trading, the underlying trading environment remained the same: near-frictionless price discovery and regulated governance frameworks.

Crypto trading stands in contrast. Each broker or exchange may maintain its own pricing model, liquidity sources, spreads and fee structures. Some platforms show a single “buy price” without revealing the underlying order book, others add hidden markups which bundle spreads and fees in a way that makes it difficult for users to understand the actual cost of a trade. As a result, two platforms may display nearly identical market prices but still deliver meaningfully different amounts of cryptocurrency for the same euro investment, creating a fundamental transparency gap for retail users seeking to compare effective execution costs.

[1] Chainalysis, 2025
[2] Aerts et al., 2025.
[3] Kaiko Research, 2025.
[4] Deutsche Börse, n.d.

## b) Motivation: Sources of opacity in crypto execution costs

The core motivation of this study is the observation that crypto brokers often communicate fees that do not necessarily correspond to what users ultimately pay. While explicit trading fees are usually listed on websites, the effective trading cost is determined by several additional components:

- Bid–ask spreads, which vary significantly across platforms and are often substantially wider than those in traditional equity markets.
- Hidden markups embedded into the displayed execution price, which are not separately disclosed to the user. The price displayed to the user at the point of trade already deviates significantly from the prevailing market price, independent of any separately charged fees.
- Differences in liquidity sourcing, especially for brokers that purchase from external exchanges rather than operating their own order book.

In combination, these factors make crypto trading costs substantially opaquer than those in traditional equity markets. This study aims to quantify this opacity by examining the effective execution quality across major European providers, measured as the amount of cryptocurrency received for a fixed Euro amount, as well as the round-trip cost of buying and selling.

## c) Research questions

Based on this motivation, the project investigates the following research questions:

1. Which European retail crypto brokers provide the best all-in execution prices once spreads, explicit fees and hidden markups are considered?
2. To what extent do published fees account for total execution costs?
3. Does API (Application Programming Interface )-based data collection predict the execution quality that retail users experience in practice?
4. How large are price differences between brokers and how consistent are they across assets?

To answer these questions, the analysis is conducted in two phases:

- **Phase 1 (API Testing):** Automated collection of retail buying prices from Bitvavo, Coinbase and Kraken.
- **Phase 2 (Real-Money Testing):** Execution of 100€ buy–sell trades on Bitvavo, Bitpanda, BSDEX, Bison, Coinbase and Kraken to validate real execution quality.

## d) Project partners and independence

This study was carried out in cooperation with Bitvavo and the Chair of Corporate Governance and Capital Markets Law (Prof. Dr. Maume) at the Technical University of Munich.
Bitvavo provided the initial funds of 5000€ for the real-money execution phase, however, the authors reimbursed the full amount upon concluding all trades, such that no net financial contribution from Bitvavo remained at the completion of the study. All trades on Bitvavo were executed under standard retail conditions, with no preferential fee arrangements or account privileges.
Bitvavo had no influence on the study design, no access to interim results or analyses and no right of review or veto over the final publication. Academic supervision by Prof. Maume ensured full methodological independence, compliance with academic research standards and neutrality in interpreting the results.

## e) Relevance of execution quality for retail consumers

Execution quality is economically meaningful because even small percentage differences in effective price compound with every transaction. A retail investor making monthly purchases of 500€ over five years on a platform that charges 1% more in effective execution costs will pay approximately 300€ in additional costs on the buy side alone. For investors using dollar-cost averaging, a widely adopted retail strategy, each periodic purchase represents a separate exposure to the platform's execution costs, making consistent execution quality critical for long-term outcomes.
The problem is compounded by an information asymmetry. Unlike in equity markets, where standardized market access and minimal spreads allow users to reasonably assume near-identical execution quality across brokers, crypto users cannot rely on such assumptions. Most retail investors take the platform-quoted price at face value, unaware that the displayed execution price may deviate significantly from prevailing market conditions. There is currently no consolidated reference price and no publicly available execution quality comparison across European crypto brokers. This study aims to fill that gap.

# 5) MiCAR and the evolving European regulatory landscape

The introduction of the Markets in Crypto-Assets Regulation (MiCAR)[5] represents a major step toward harmonized crypto regulation in the European Union. MiCAR establishes standardized requirements for crypto-asset service providers (CASPs), including governance rules or prudential safeguards.

Article 78(1) MiCAR imposes a best-execution obligation on CASPs executing orders on behalf of clients: they must take all necessary steps to obtain the best possible result, considering price, costs, speed, likelihood of execution and settlement, size, nature, conditions of custody, and any other consideration relevant to the execution of the order[6].

Article 78(2) further requires CASPs to establish and implement effective execution arrangements, including an order execution policy.

In this respect, the language closely parallels Article 27 of MiFID II, which governs best execution for financial instruments and requires investment firms to obtain the best possible result for retail clients in terms of total consideration, the price of the instrument plus all execution-related costs[7].

However, the two regimes differ substantially in the infrastructure and enforcement mechanisms that give their respective best-execution obligations practical effect.

Under MiFID II, best execution operates within a developed ecosystem: regulated trading venues with consolidated tapes that provide market-wide reference prices[8] and detailed regulatory technical standards specifying how firms must establish and assess the effectiveness of their execution policies[9].

MiCAR lacks this supporting infrastructure in several important respects.

First, while Article 76 MiCAR imposes pre- and post-trade transparency requirements on operators of trading platforms for crypto-assets, these obligations do not extend to CASPs executing orders on behalf of clients under Article 78. The retail investor interacting with a broker-model CASP therefore does not benefit from the same transparency framework that applies to trading platform operators.

---

[5] Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets (OJ L 150, 9.6.2023, p. 40).

[6] Article 78(1) MiCAR.

[7] Article 27(1) MiFID II (Directive 2014/65/EU), as amended by Directive (EU) 2024/790 (MiFID III).

[8] The MiFIR review (Regulation (EU) 2024/791) empowers ESMA to organize selection procedures for consolidated tape providers, beginning with bonds. A CTP for shares is to follow.

[9] See ESMA Final Report on draft RTS specifying criteria for establishing and assessing the effectiveness of investment firms' order execution policies, 10 April 2025.

Second, there is no consolidated tape for crypto-assets under MiCAR. In equity and bond markets, MiFID II/MiFIR provides for consolidated tape providers that aggregate trade data across venues into a single reference source[10], enabling market participants to assess execution quality against a market-wide benchmark. No equivalent mechanism exists for crypto-assets, leaving retail users without a standardized reference price against which to evaluate the prices offered by their CASP.

Third, MiCAR does not require CASPs to disclose the effective spread embedded in their execution prices. Article 78(3) requires CASPs to provide clients with information about their order execution policy, and Article 78(6) mandates periodic assessment of execution quality. However, neither provision requires public or client-facing disclosure of the spread or markup applied to individual transactions. A CASP may therefore comply with MiCAR's disclosure requirements while offering an execution price that deviates significantly from the prevailing market price, without the client having any means of detecting this deviation.

ESMA has recognized the enforcement challenges arising from this structural gap. In its July 2024 Opinion on broker models[11], ESMA noted that when a CASP offers execution across a wide variety of crypto-assets, that relying on a single execution venue is "unlikely" to enable the best possible result for clients (para. 35). ESMA further suggested that CASPs should periodically compare available execution venues, first on price and direct execution costs and then on remaining factors such as speed and settlement conditions (para. 36). However, this Opinion is addressed to national competent authorities as a supervisory convergence tool under Article 29(1)(a) of the ESMA Regulation and does not constitute a binding obligation on CASPs themselves.

The practical consequence is significant for retail investors. Two fully MiCAR-compliant CASPs may deliver substantially different execution prices for the same trade, and the retail user currently has no standardized means of detecting or comparing this difference. Because MiCAR's best-execution obligation operates without the transparency infrastructure needed to make it independently verifiable by consumers, empirical studies like the present one serve an important function: they provide the execution quality comparison that the regulatory framework itself does not yet enable.

[10] Articles 27a–27h MiFIR (as amended by Regulation (EU) 2024/791).

[11] ESMA, Opinion to support the convergent application of MiCA, 31 July 2024, ref. ESMA75-453128700-1048.

# 6) Selection of exchanges

## a) Selection criteria

To ensure that the study reflects the trading environment encountered by typical European retail users, platforms were selected according to three criteria:

**1. MiCAR licensing**
Only platforms that had obtained authorization under MiCAR by July 2025 were considered[12]. While MiCAR does not regulate execution quality or microstructure, as discussed in Section 5, licensing status serves as an indicator of institutional reliability and operational maturity. Restricting the sample to licensed platforms also ensures that all studied brokers operate under the same baseline regulatory obligations, which improves comparability.
**2. Retail-focused trading experience (no pro/advanced interfaces)**
To ensure comparability across platforms, only retail-level trading interfaces were considered. This meant excluding platforms whose primary trading environment is a professional interface with advanced order types, order books, or maker/taker fee tiers inaccessible to typical retail users.
This aligns the study with its practical objective: determining which platform provides the fairest prices to the average retail user, not to algo traders or professional participants.
**3. Market relevance and user reach in Europe**
Exchange popularity was evaluated through publicly observable indicators such as app store rankings[13] and website traffic.
Brokers with limited European presence or niche user bases were not included, as they are unlikely to reflect the trading environment encountered by most EU retail users.

[12] Digital Euro Association, n. D.
[13] On the apple appstore

### b) Final set of platforms included in the study

Based on these criteria, the following platforms were included:

*Table 1: Platforms included in the study by phase and pricing model. Bitpanda, BSDEX, and Bison were excluded from Phase 1 due to the absence of public retail API endpoints.*

| Platform | Phase 1 (API) | Phase 2 (Real-money) | Pricing Model |
|---|---|---|---|
| **Bitvavo** | ✓ | ✓ | Exchange / order-book |
| **Coinbase** | ✓ | ✓ | Broker / spread-based |
| **Kraken** | ✓ | ✓ | Broker / spread-based |
| **Bitpanda** | — | ✓ | Broker / spread-based |
| **BSDEX** | — | ✓ | Exchange / order-book |
| **Bison** | — | ✓ | Broker / spread-based |

## 7) Methodology

### a) Overview of the research design

The methodology is strutured in two complementary phases that capture both theoretical and realized execution outcomes.
Phase 1 queried order book prices via API from Bitvavo, Coinbase and Kraken and computed the quantity of cryptocurrency obtainable for fixed euro amounts after applying each exchange's published fee schedule. This produced time-stamped, comparable datasets enabling controlled cross-venue comparison under near-identical market conditions.
Phase 2 validated the practical relevance of Phase 1 by executing 100€ buy–sell round-trip trades across all six platforms using their standard retail interfaces. This enabled direct measurement of actual spreads, execution price deviations and the divergence between order book conditions and real trading outcomes.
The two phases are deliberately separated: Phase 1 supplies controlled and replicable data at scale, while Phase 2 captures the full practical cost of trading, including any discrepancy between quoted and realized prices. Together, they generate an assessment neither phase could achieve alone.

## b) Conceptual framework and definitions

The methodology rests on three core concepts: coins received, effective execution price and round-trip cost. These allow for a standardized comparison across brokers regardless of differences in fee models or interface design.

**Coins received for a fixed euro amount**

Retail users typically buy crypto using a fixed notional amount (e.g., 100€). Therefore, the most intuitive measure of execution quality is:

$$\text{Coins Received} = \frac{\text{Amount Paid in €} - \text{Fees in €}}{\text{Execution Price €/Coin}}$$

Where the execution price denotes the actual trade price quoted to the user, it includes any spread or price deviation from the market and fees denotes the explicit transaction fees the platform charges. A broker that offers better execution yields more cryptocurrency for the same euro amount.

**Effective execution price**

To standardize comparisons, we compute:

$$\text{Effective Price} = \frac{\text{Amount Paid in €}}{\text{Coins Received}}$$

This converts any fee model (spread-only, fee + spread, markup, hybrid, etc.) into a unified metric. The effective price may exceed the displayed price if the execution price itself deviates from market conditions, as discussed in Section 4b.

**Round-trip cost (buy + sell)**

Execution quality cannot be assessed solely based on buy-side prices. For retail investors, the total cost of trading includes both entering and exiting a position. The round-trip cost is defined as:

$$\text{Round-Trip Cost (\%)} = \left(1 - \frac{\text{Sell Proceeds}}{\text{Buy Expenditure}}\right) \times 100$$

A lower percentage indicates better execution. This metric captures all frictional components such as spreads, fees, execution price deviations and slippage.

## c) Methodology of Phase 1: API-Based Data Collection

**Technical infrastructure**

Data collection was automated using three independent AWS Lambda functions, one per exchange. Each function queried the respective exchange's public order book API to retrieve the current best ask and best bid prices:

- Bitvavo: REST endpoint (api.bitvavo.com/v2/ticker/book)
- Kraken: REST endpoint (api.kraken.com/0/public/Ticker)
- Coinbase: REST endpoint (api.exchange.coinbase.com/products/…/book)

From each ask price, the function computed the quantity of cryptocurrency a retail customer would receive for notional investments of 100€, 500€ and 1,000€ after applying the exchange's published taker fee for the applicable volume tier. The taker fees applied were 0.25% for Bitvavo, 1.00% for Kraken and 1.49% for Coinbase, which corresponds to the default retail tier on each platform. All results were stored with millisecond-precision timestamps in a centralised database (AWS DynamoDB), enabling synchronized cross-venue comparison.

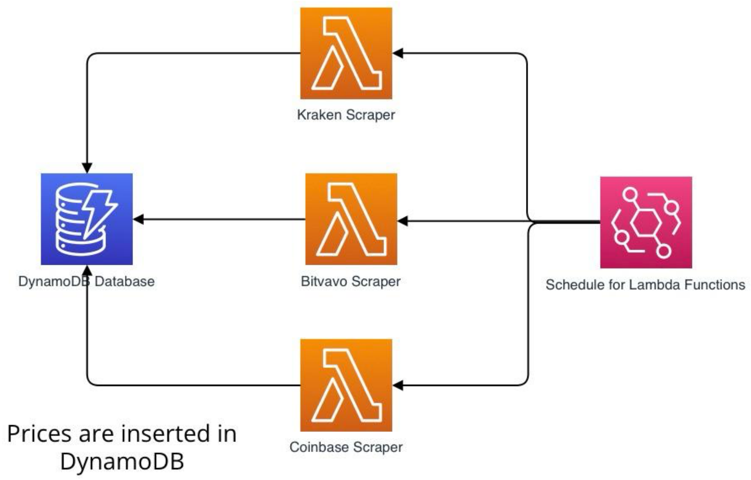


*Figure 1 Technical infrastructure for Phase 1 API-based data collection*

**Scheduling and temporal alignment**

Lambda functions triggered twice daily at 10:00 and 16:00 CET captured market conditions during European and US trading hours respectively. Since all three exchange-specific functions were invoked by the same scheduling event, request timestamps were typically within seconds of each other.

Across the 86 timestamps with observations from all three exchanges, the median inter-exchange timing difference relative to Bitvavo was 3.6 seconds for Kraken and 14.2 seconds for Coinbase.

While sub-second price movements can occur on major pairs during periods of high volatility, these residual differences are unlikely to systematically bias the results in favor of any platform.

However, for the less liquid ALGO/EUR pair, where short-term price volatility is higher relative to tick size, the Coinbase delay warrants consideration. To assess the potential impact, future iterations of this study should compute the average absolute price change at the relevant time horizon for each asset pair and report this as a bound-on timing-induced measurement noise.

**Coverage**

Data was collected over a 62-day period from 26 July to 25 September 2025, yielding 924 individual price observations across three asset pairs (BTC/EUR, ETH/EUR, ALGO/EUR). Bitvavo contributed 128 observations per asset, while Kraken and Coinbase each contributed 90 observations per asset. The difference arises from Bitvavo's Lambda function that got deployed earlier during the technical setup phase.

**Reference price construction**

A neutral reference price was computed for each timestamp and asset as the arithmetic mean of the fee-corrected ask prices across all available API venues at that moment.
Each exchange's raw ask price was recovered by dividing the notional investment by the observed coins received before fee deduction, thereby removing differences in fee schedules from the reference price.
Cross-exchange price convergence serves as a data-quality check. For BTC/EUR, the mean spread between the highest and lowest ask price across the three exchanges was 0.030% (median: 0.024%), which confirms that the three order books reflect near-identical wholesale market conditions. For ETH/EUR, the mean spread was 0.037% (median: 0.029%), and for ALGO/EUR it was 0.061% (median: 0.051%).

**Estimated prices for non-API exchanges**

Bitpanda, BSDEX and Bison were excluded from direct API collection because they do not offer public retail API endpoints suitable for automated price queries. To enable cross-venue comparison, theoretical execution prices for these platforms were estimated by applying their published fee schedules to the reference price at each timestamp.
This approach yields a lower-bound estimate of actual trading costs, since it captures published fees but cannot account for any additional spread that broker-model platforms may embed in the execution price. Phase 2 is specifically designed to detect this additional spread.
Consequently, the estimated figures for Bitpanda, Bison and BSDEX should be interpreted as optimistic projections, not predictions. Where these estimates are presented alongside observed API data (e.g., in figures), they are clearly labelled as estimates and visually distinguished.

## d) Methodology of phase 2: real-money execution testing

API data alone cannot fully capture the practical trading experience, particularly for broker-model platforms that may embed spreads in their displayed prices. Real-money execution tests were conducted to measure the actual cost of trading through each platform's standard retail interface.

**Scope and sample allocation**

Phase 2 covered all six platforms and three asset pairs (BTC/EUR, ETH/EUR, LINK/EUR). A total of 50 buy–sell round-trip trades were executed over 12 trading days between 31 October and 14 November 2025. Note that Phase 2 used LINK/EUR in place of ALGO/EUR (used in Phase 1) because it was not available on all tested platforms.

Cross-phase comparison of the third asset pair is therefore limited to qualitative consistency rather than direct numerical validation. Observations were allocated unevenly across platforms. Bitvavo and BSDEX each received 13 round-trips (approximately 4–5 per asset pair), while the remaining four platforms Bison, Coinbase, Kraken and Bitpanda, each received 6 round-trips (2 per asset pair). The allocation rationale was that exchange-model platforms (Bitvavo, BSDEX) were expected to exhibit lower and more stable costs, which required finer resolution to distinguish genuine execution quality differences from noise. For the broker-model platforms, where cost differences were expected to be large and unambiguous, fewer observations were deemed sufficient to confirm the direction and approximate magnitude of costs.

*Execution procedure.* For each round-trip:

1. Funds were deposited and the account fully verified prior to the testing period.
2. A buy order of 100€ was executed via the standard retail interface (i.e., the default trading view available to a new user, not a professional or "Pro" order book).
3. Immediately after purchase, a sell order for the full quantity of coins received was placed.
4. All timestamps, fee breakdowns, execution prices, and resulting euro proceeds were recorded and can be found in a public Github repository[14].

For BSDEX, which charges a fixed per-trade fee of 0.35€, the initial outlay was 100.35€ to ensure that the net invested amount after fees was comparable to the 100€ invested on other platforms. This adjustment is reflected in the round-trip cost calculations.

To minimize the effect of market movement between buy and sell, trades were executed as quickly as possible, typically within 2–3 minutes. For sessions involving multiple exchanges, all platforms were traded in rapid succession within the same time window (e.g., all BTC/EUR buys

[14] https://github.com/dominikhei/tum-research-markets

placed within a few minutes, followed by all sells). This design ensures that cross-platform comparisons reflect execution quality differences rather than market timing.

**Trade size**

The 100€ notional was chosen to represent a typical retail transaction. At this size, percentage-based fees apply uniformly and fixed fees (particularly Coinbase's 2.99€ per trade) have maximum relative impact, making this a conservative test of retail execution costs.
Phase 1 data, which include computations for 500€ and 1,000€ notionals, show that percentage-based cost differences between exchanges remain stable at higher notional amounts, as expected when fees are proportional. The Coinbase flat fee becomes proportionally less significant at higher notionals, so the cost disadvantage observed for Coinbase at 100€ narrows at larger trade sizes.

## e) Statistical considerations

For Phase 1, the primary metric is the percentage by which each exchange's effective price (including fees) exceeds the cross-venue mean ask price at each timestamp. For Phase 2, the primary metric is the round-trip cost as defined in Section 7b.
Descriptive results are reported as means with standard deviations and 95% confidence intervals, computed using t-distribution critical values. Pairwise differences between exchanges are assessed using the Mann–Whitney U test, a nonparametric test appropriate for small samples that does not assume normally distributed cost differences, at a significance level of $\alpha = 0.05$.
Given 15 pairwise comparisons across six platforms, results are also evaluated against a Bonferroni-corrected threshold of $\alpha_{adj} \approx 0.003$.
Phase 1 and Phase 2 are compared by estimating the hidden retail spread for each platform: the difference between the observed Phase 2 round-trip cost and the predicted cost derived from Phase 1 (2 × one-way API cost). A large hidden spread indicates that the retail interface imposes costs beyond what published fees and order-book prices suggest.
Several design limitations should be noted. Phase 2 sample sizes (N = 6 for four of six platforms, N = 13 for Bitvavo and BSDEX) are sufficient to detect the large cost differences between exchange-model and broker-model categories but cannot resolve finer within-category rankings. The unequal allocation produces narrower confidence intervals for Bitvavo and BSDEX than for the remaining platforms. The buy–sell time gap (mean: 2.6 minutes) introduces symmetric noise of approximately ±0.14% per leg, which is proportionally more significant for low-cost platforms (~24% of Bitvavo's signal) than for high-cost ones (~2% of Coinbase's). The 12-day Phase 2 observation window may not be representative of longer-term conditions.

# 8) Empirical results: integrated findings from API and real-money testing

This section presents the combined empirical findings from both phases of the study. Phase 1 (API testing) provides controlled and synchronized pricing data from exchange order books while Phase 2 (real-money execution) captures the full practical cost of trading through each platform's retail interface. Where relevant, cross-phase comparisons are used to estimate the hidden retail spread, the cost component not visible in order-book data or published fee schedules.

## a) Phase 1: API-based execution quality and theoretical cross-venue predictions

***Cross-exchange price convergence***

As a data-quality check, the raw ask prices (before fees) across Bitvavo, Kraken and Coinbase were compared at each matched timestamp. Table 2 reports the mean and median ask-price spreads by asset pair.

*Table 2: Cross-exchange ask price spread (max – min ask divided by mean ask) across 86 matched timestamps.*

| Asset Pair | N (timestamps) | Mean Spread | Median Spread |
|---|---|---|---|
| BTC/EUR | 86 | 0.030% | 0.024% |
| ETH/EUR | 86 | 0.037% | 0.029% |
| ALGO/EUR | 86 | 0.061% | 0.051% |

The tight convergence confirms that all three order books reflect near-identical wholesale market conditions, which validates the use of cross-venue comparison.

**Execution price comparison across venues**

Phase 1 consisted of two analytical layers: (1) direct API-to-API comparison of coins received after fees across Bitvavo, Kraken and Coinbase, and (2) theoretical estimates for Bitpanda, BSDEX and Bison, computed by applying their published fee schedules to the cross-venue reference price at each timestamp.

**Direct API comparison.** Across all 86 matched timestamps, three asset pairs, and three notional sizes (100€, 500€, 1,000€), Bitvavo provided the highest quantity of coins received at every observation, a 100%-win rate across 774 scoring opportunities (86 timestamps × 3 assets × 3 notionals). This reflects Bitvavo's lower retail fee schedule (0.25% taker) relative to Kraken (1.00%) and Coinbase (1.49%), combined with competitive ask prices. Because the raw ask prices converge tightly (Table 2), the fee differential is the dominant driver of the cost ranking at the order-book level.

When estimated costs for BSDEX (0.35%), Bison (1.25%) and Bitpanda (0.99% for BTC/ETH, 2.49% for ALGO) are added, the ranking remains unchanged: Bitvavo retains the best effective price at virtually every timestamp. BSDEX (est.) ranks second due to its low fixed fee. These estimates represent lower bounds, as they capture published fees but cannot account for any additional spread embedded in broker-model pricing. Figure 1 summarizes the scoring results. Estimated venues are visually distinguished and should be interpreted as optimistic projections.


*Figure 2 Cross-venue execution price comparison across all timestamps, assets and notional sizes.*

Phase 1 thus establishes a clear theoretical baseline: given equivalent wholesale market conditions, Bitvavo offers the lowest effective execution price among the studied platforms, driven primarily by its lower fee structure. Whether this advantage persists when retail-interface effects are included is tested in Phase 2.

## b) Phase 2: Real-money execution results

**Descriptive results**

Table 3 reports the mean round-trip cost per platform, measured as the percentage loss between the initial euro investment and the euro proceeds from the immediate subsequent sell.

*Table 3: Phase 2 round-trip costs by platform, sorted by mean cost. CIs computed using t-distribution with N−1 degrees of freedom.*

| Platform | Model | N | Mean RT Cost | SD | 95% CI |
|---|---|---|---|---|---|
| Bitvavo | Exchange | 13 | −0.58% | 0.18% | [−0.69%, −0.47%] |
| BSDEX | Exchange | 13 | −0.88% | 0.23% | [−1.01%, −0.74%] |
| Bison | Broker (spread) | 6 | −2.58% | 0.23% | [−2.82%, −2.34%] |
| Kraken | Broker (retail) | 6 | −5.81% | 0.49% | [−6.33%, −5.29%] |
| Bitpanda | Broker (spread) | 6 | −6.23% | 1.37% | [−7.67%, −4.79%] |
| Coinbase | Broker | 6 | −7.49% | 1.00% | [−8.53%, −6.44%] |

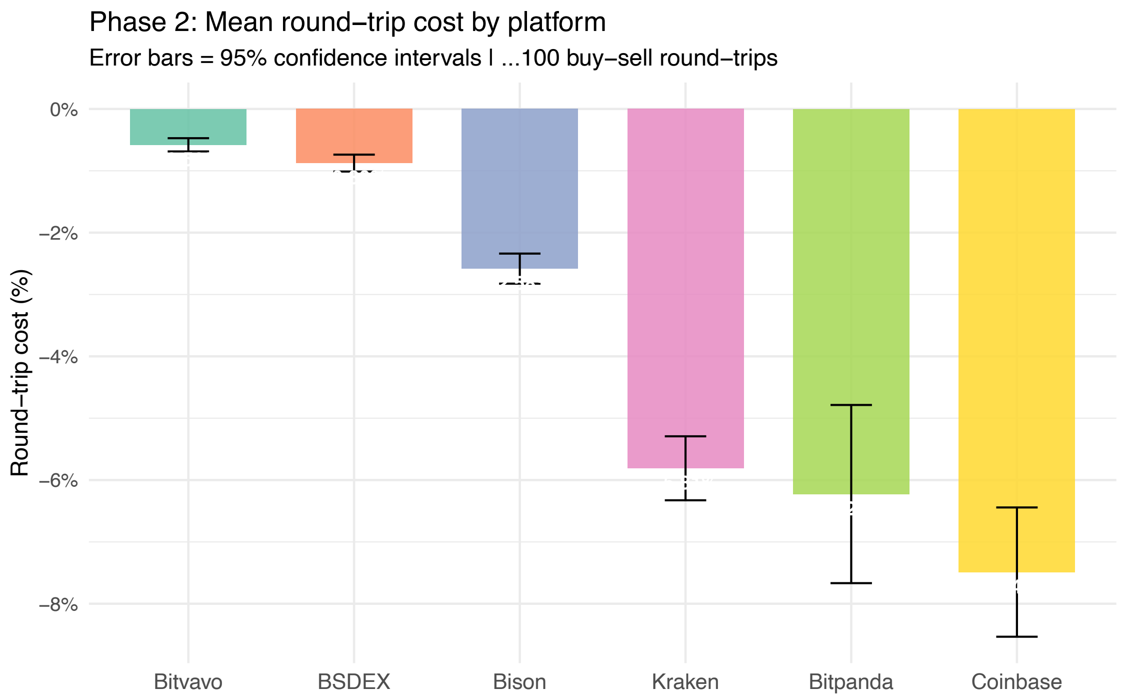


*Figure 3 Mean round-trip cost by platform with 95% confidence intervals. N = 13 for Bitvavo and BSDEX, N = 6 for all others.*

Bitvavo exhibited the lowest mean round-trip cost across all asset pairs (−0.58%), followed by BSDEX (−0.88%). Both platforms employ order-book-based pricing. The remaining platforms showed substantially higher costs, ranging from −2.58% (Bison) to −7.49% (Coinbase). Notably, the confidence interval for Bitpanda is nearly three percentage points wide, reflecting high variance across its asset-specific fee tiers.

**Per-asset breakdown**

Table 4 disaggregates results by asset pair. For Bison, Kraken, Bitpanda and Coinbase, each cell reflects only N = 2 observations and should be treated as illustrative rather than statistically reliable.

*Table 4: Mean round-trip cost by platform and asset pair. Per-cell sample sizes are indicated; N = 2 results are indicative only.*

| Platform | BTC/EUR | ETH/EUR | LINK/EUR |
|---|---|---|---|
| Bitvavo (N=4–5) | −0.62% | −0.52% | −0.59% |
| BSDEX (N=4–5) | −0.86% | −0.75% | −1.01% |
| Bison (N=2) | −2.73% | −2.50% | −2.52% |
| Kraken (N=2) | −5.53% | −5.54% | −6.37% |
| Bitpanda (N=2) | −4.59% | −6.49% | −7.60% |
| Coinbase (N=2) | −6.65% | −7.93% | −7.88% |

The platform ranking is consistent across all three asset pairs. Bitpanda shows the largest cross-asset variation, with BTC/EUR costs substantially lower than ETH/EUR or LINK/EUR, likely reflecting its asset-specific fee schedule.

**Statistical tests**

After Bonferroni correction, the results support a three-tier cost structure:

- **Tier 1 (exchange-model):** Bitvavo (−0.58%) and BSDEX (−0.88%). The difference between these two is significant ($p < 0.001$), but both are substantially cheaper than all other platforms.
- **Tier 2 (intermediate):** Bison (−2.58%). Significantly more expensive than Tier 1 ($p < 0.001$), significantly cheaper than Tier 3 at the nominal level ($p = 0.005$), though this comparison does not survive Bonferroni correction.
- **Tier 3 (broker-model, high cost):** Kraken (−5.81%), Bitpanda (−6.23%), Coinbase (−7.49%). The difference between Kraken and Bitpanda is not statistically significant ($p = 0.589$). The Kraken–Coinbase and Bitpanda–Coinbase differences are borderline ($p = 0.030$) and do not survive Bonferroni correction. These three platforms cannot be reliably ranked relative to each other at current sample sizes.

## c) Cross-phase validation and structural interpretation

For the three platforms with both Phase 1 (API) and Phase 2 (real-money) data, Table 5 compares the predicted round-trip cost, estimated as twice the one-way retail fee applied in Phase 1, with the actual Phase 2 round-trip cost. The difference represents the hidden retail spread: additional cost embedded in the retail trading interface beyond what is visible in order-book prices and published fees.

*Table 5: Cross-phase validation for API-observed platforms. Hidden spread = actual RT cost minus predicted RT cost from Phase 1 fee schedules.*

| Platform | Phase 1 Fee (one-way) | Predicted RT (2× fee) | Actual RT (Phase 2) | Hidden Spread |
|---|---|---|---|---|
| Bitvavo | 0.25% | 0.50% | 0.58% | 0.08% |
| Kraken | 1.00% | 2.00% | 5.81% | 3.81% |
| Coinbase | 1.49% | 2.98% | 7.49% | 4.51% |

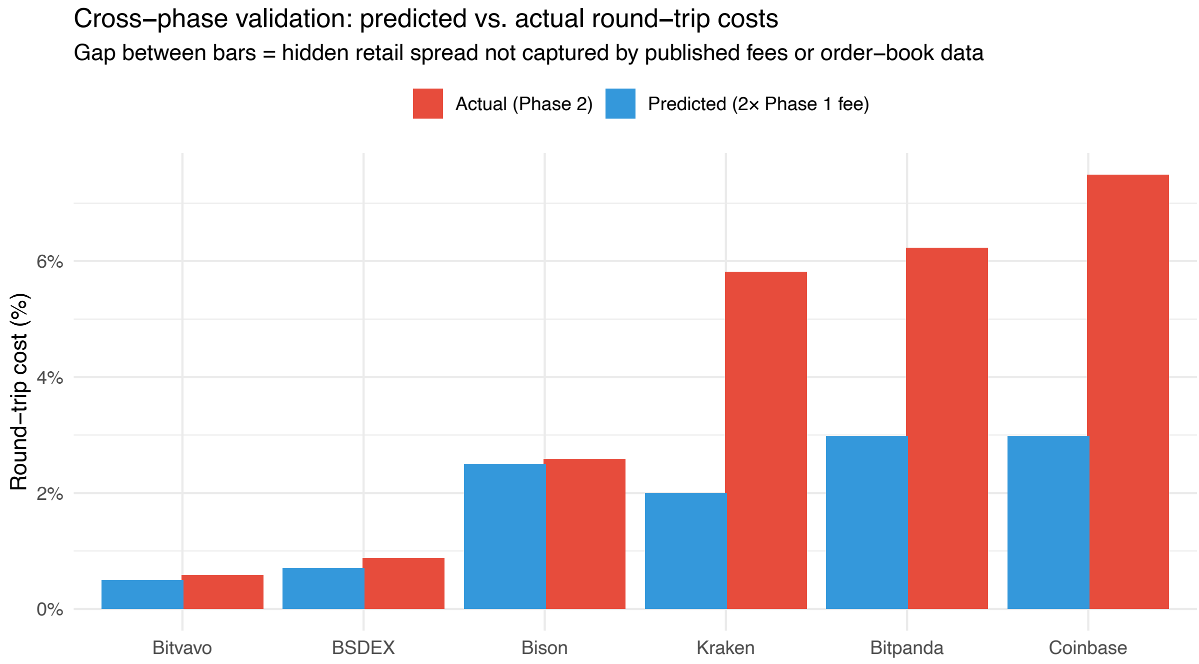


*Figure 4 Predicted (2× Phase 1 fee) versus actual round-trip cost. The gap represents the hidden retail spread.*

For Bitvavo, the hidden spread is 0.08 percentage points, which is negligible and likely attributable to bid–ask spread fluctuations during the buy–sell interval. This indicates that Bitvavo's retail interface closely tracks its wholesale order book and that published fees are essentially the full cost of trading.

For Kraken and Coinbase, the hidden spreads are 3.81 and 4.51 percentage points respectively. These gaps are too large and too consistent across assets to be explained by market movement or measurement noise. They reflect a structural feature of these platforms' retail interfaces: the execution price presented to the user incorporates a substantial markup over the wholesale order-book price, above and beyond the explicitly charged fee.

For the non-API platforms a similar pattern emerges. Phase 1 lower-bound estimates (based on published fees alone) understate actual Phase 2 costs for Bison (estimated 2.50% vs. actual 2.58%) and dramatically for Bitpanda (estimated ~1.98% vs. actual 6.23%). BSDEX's estimate (0.70%) understates its actual cost (0.88%) by 0.18 percentage points, a gap consistent with normal bid–ask spread variation.

# 9) Discussion

The empirical results reveal systematic and economically meaningful differences in execution quality across the six platforms studied. This section interprets the key findings and addresses limitations of the study.

## a) The hidden retail spread as the central finding

The cross-phase validation produces the study's most informative result. For Bitvavo, the actual round-trip cost (0.58%) closely matches the prediction derived from its published fee schedule (0.50%), with a hidden spread of just 0.08 percentage points. For Kraken and Coinbase, however, actual costs (5.81% and 7.49%) exceed fee-based predictions (2.00% and 2.98%) by 3.81 and 4.51 percentage points respectively. These gaps are too large and too consistent across assets to reflect measurement noise. They indicate that the retail interfaces of these platforms embed a substantial, undisclosed markup in the execution price, a cost component that is invisible to users and absent from published fee schedules.

This finding directly addresses the study's core research question: published fees are a reliable indicator of total execution cost on exchange-model platforms but capture only a fraction of the effective cost on broker-model platforms. On Coinbase, for instance, the posted fee accounts for roughly 40% of the actual trading cost, the remaining 60% is hidden in the execution price.

## b) Pricing model and execution quality

The data supports a three-tier cost structure. The two exchange-model platforms (Bitvavo at −0.58%, BSDEX at −0.88%) cluster well below all broker-model platforms, with pairwise differences significant after Bonferroni correction. Bison (−2.58%) occupies an intermediate position. Kraken (−5.81%), Bitpanda (−6.23%) and Coinbase (−7.49%) form a high-cost tier that cannot be reliably distinguished internally at current sample sizes (Kraken vs. Bitpanda: $p = 0.589$).
However, the sample contains only two exchange-model platforms, precluding generalization. The observed advantage may partly reflect these specific platforms' competitive strategies or liquidity conditions rather than a universal property of the exchange model.

## c) Alternative explanations

**Retail interface vs. exchange product.**
Kraken and Coinbase both operate professional order-book interfaces alongside their retail simple-buy products. Phase 1 queried the professional API, Phase 2 tested the retail interface. The hidden spread may therefore not reflect platform-wide pricing but rather the specific cost of using the simplified retail product. Users who access the exchange's order book directly would likely achieve substantially better execution. This has an important implication: the cost disadvantage observed for Kraken and Coinbase is partly a product-design choice (subsidizing simplicity with higher spreads) rather than a structural deficiency. Whether typical retail users are aware that a cheaper execution path exists on the same platform is itself a transparency question.

**Trade size effects.**
All Phase 2 trades used a 100€ notional. Coinbase's flat fee of 2.99€ per trade imposes a 2.99% one-way cost at this size, which would fall to 0.60% at 500€ and 0.30% at 1,000€. The ranking among broker-model platforms may therefore shift at higher trade sizes, even if the exchange-model advantage persists.

**Temporal specificity.**
Phase 2 covers 12 trading days in a single market regime. Execution costs on broker-model platforms may vary with liquidity conditions or internal pricing-engine calibration. The observed cost levels should be understood as estimates for the specific period studied, not as stable platform characteristics.

## d) Regulatory implications

The findings are directly relevant to the enforcement of MiCAR's best-execution obligation (Article 78). As discussed in Section 5, MiCAR requires CASPs to obtain the best possible result for clients but lacks the transparency infrastructure, consolidated tapes, standardised spread disclosure and execution quality reporting, needed to make this obligation independently verifiable. The present study illustrates the practical consequence: two fully MiCAR-compliant platforms can deliver execution prices that differ by more than six percentage points on a round-trip basis, with no standardized mechanism for the retail user to detect this difference.
Possible regulatory responses include mandatory disclosure of the effective spread embedded in retail execution prices, periodic publication of execution quality metrics analogous to SEC Rule 605/606, or the development of a consolidated tape for crypto-assets. Each carries implementation costs and potential unintended consequences, such as reduced platform diversity or disadvantages for smaller venues, that require further analysis.

## e) Limitations

Beyond the statistical limitations discussed in Section 7e, several broader constraints apply. First, the study examines six platforms, results should not be generalized to the wider European CASP landscape. Second, the study measures execution costs but does not assess other dimensions of service quality (custody security, user experience, asset coverage, withdrawal speed) that may legitimately factor into platform selection. Third, the cooperation with Bitvavo, while methodologically safeguarded, creates an appearance of conflict that readers should weigh when interpreting results where Bitvavo ranks first. Fourth, the classification of platforms into 'exchange-model' and 'broker-model' categories is based on the interface tested and may not capture the full complexity of each platform's market structure.

# 10) Conclusion

This study evaluated the effective execution costs faced by retail investors on six major European crypto trading platforms, combining API-based price monitoring with real-money round-trip testing. Three principal findings emerge.

First, execution costs vary dramatically across platforms. Round-trip costs ranged from −0.58% (Bitvavo) to −7.49% (Coinbase), a difference of nearly seven percentage points. Platforms employing transparent, order-book-based pricing (Bitvavo, BSDEX) consistently outperformed those relying on spread-based or markup-based retail interfaces.

Second, published fee schedules are a poor predictor of total execution cost on broker-model platforms. The hidden retail spread, the gap between fee-predicted and actual round-trip costs, was negligible for exchange-model platforms (0.08% for Bitvavo) but exceeded four percentage points for Coinbase and nearly four for Kraken. This means that most trading costs on these platforms is not disclosed to users.

Third, the current regulatory framework does not enable retail users to detect these differences. MiCAR's best-execution obligation (Article 78) operates without the transparency infrastructure, consolidated tapes, spread disclosure requirements, execution quality reporting, that gives the analogous MiFID II obligation practical effect. Empirical studies like this one currently serve a function that the regulatory framework itself does not yet provide.

From a consumer perspective, platform selection matters substantially. For an investor making monthly 500€ purchases, the difference between the best and worst platforms in this study amounts to approximately 200€ per year in excess execution costs on the buy side alone. For investors using periodic purchasing strategies, each transaction represents a separate exposure to the platform's cost structure, making consistent execution quality critical for long-term outcomes.

Key limitations include the small Phase 2 sample sizes (N = 6 for four platforms), a short observation window, a single trade size (100€) and the Bitvavo cooperation, which prevented reliable ranking among the three highest-cost platforms. Future research should expand sample sizes, test multiple trade sizes, extend the observation period and include additional platforms.

Despite these constraints, the core finding is robust: transparent, exchange-based pricing delivers materially lower execution costs than opaque, spread-based retail interfaces and the gap between advertised and actual costs on broker-model platforms is large enough to warrant both regulatory attention and greater consumer awareness.