

Detectable but Not Exploitable: Predictability and Trading Costs on the Casablanca Stock Exchange*

Ahmed Afachtal[†]

September 2026

Abstract

Statistical detectability does not by itself establish economic exploitability. We measure the gap directly, on 73 securities of the Casablanca Stock Exchange, 2010–2026, using a panel laid on the exchange’s own session calendar so that non-trading intervals are observed rather than inferred. Price recovery following such an interval is significant on 21 of 60 securities under standard inference and on 26 of 60 under pairs-bootstrap inference, in both cases more than seven times the false-positive rate (exact binomial $p = 5.7 \times 10^{-13}$ and $p = 1.9 \times 10^{-18}$), yet returns 0.228% over five sessions against a published round-trip fee floor of at least 1.10%. Across 56 testable securities none shows a significant surplus, 32 show a significant shortfall, and the mean shortfall of 0.87% at the published fee floor ($p = 1.0 \times 10^{-16}$) remains significant, at 0.43 points ($p = 2.8 \times 10^{-7}$), even with brokerage at zero and only the non-negotiable components remaining. Published fees, not estimated spreads, anchor the test: the Roll and Corwin–Schultz estimators differ by a factor of four and rank securities inconsistently. The conclusion is unchanged under every market-factor construction and sample window we test, and under inference that assumes nothing about the return distribution or the independence of the securities; the deficit narrows as the holding horizon lengthens and, at 20 sessions, is no longer significant.

Use of generative AI. AI tools assisted with language editing and with drafting and reviewing analysis code. All results come from the code in the replication package, and every number in the paper is checked against it by the verification program of Appendix C. The author is responsible for all content.

JEL classification: G12, G14, C12, C15, C58

Keywords: transaction costs, frontier markets, non-synchronous trading, multiple testing, Casablanca Stock Exchange

*Morocco is classified as a Frontier Market by MSCI, which reclassified it from Emerging in November 2013, and by FTSE Russell. References to frontier-market microstructure in this paper refer to that classification.

[†]National School of Business and Management (ENCG), Sidi Mohamed Ben Abdellah University, Fez, Morocco; and EcoFinance Analytics, Casablanca, Morocco. Email: ahmed.afachtal@usmba.ac.ma. This research was conducted at EcoFinance Analytics. All errors are my own.

1 Introduction

A security that has not traded for several sessions has not incorporated what the market did in the meantime, and when it trades again it recovers part of the move. The effect is mechanical and has been understood since Lo and MacKinlay (1990). Whether the recovery is large enough to pay for the round trip that would capture it is a different question, and it is the one this paper answers.

The two questions come apart on an illiquid market, and the Casablanca Stock Exchange is illiquid in a way that makes the distance between them measurable. The median security in our sample does not trade on 17.5% of sessions and the least liquid on 82.9%, so non-trading intervals are frequent enough to test one security at a time rather than in the aggregate. The exchange also publishes an explicit fee schedule, so the cost of capturing the recovery can be read off rather than estimated, which matters more than it might because the standard estimators of that cost turn out to disagree here by a factor of four.

We build a panel of 73 securities over 2010 to 2026 on the exchange’s own session calendar, so that a run of sessions with no trade is observed rather than proxied by a zero return. The distinction is not a formality: a zero return also occurs on sessions a security did trade, on a median 15.0% of them, and 14.7% of consecutive trades in the panel span more than one session, up to 189. Section 3 sets out the construction and the checks against the exchange’s own record that license it.

The recovery is easy to detect. Of the 60 securities with at least 30 non-trading episodes, 21 show a significant post-thaw recovery against 3 expected by chance, and an exact binomial test rejects at $p = 5.7 \times 10^{-13}$; a pairs bootstrap that assumes nothing about the return distribution raises the count to 26. It is not worth capturing. The mean recovery over the five sessions that follow is 0.228%, against a published round-trip fee floor of at least 1.10%. Across the 56 securities that can be tested against cost, none shows a significant surplus, 32 show a significant shortfall, and no bootstrap resample places the mean above zero. The predictability is real and widespread, and it returns about a fifth of what it costs to collect.

Anchoring that comparison on published fees rather than on an estimated spread is a choice the data force. The Roll and Corwin–Schultz estimators differ by a factor of four on this market, rank the cross-section inconsistently, and one of them reports narrower spreads for securities that trade less often, which is the opposite of what illiquidity should produce. Stated against the Roll spread the mean deficit is 2.49 percentage points; stated against Corwin–Schultz it is 1.28. A magnitude that nearly doubles on the choice of estimator is a researcher degree of freedom of the kind Menkveld et al. (2024) document. The published schedule requires no estimation and is a strict lower bound, since an investor pays the spread in addition to it, and the conclusion survives with brokerage set to zero: on the non-negotiable components alone, at 0.66% round trip, the mean recovery still falls short by 0.43 points.

The paper makes two contributions, each narrower than the literature it builds on. Novy-

Marx and Velikov (2016) sort anomalies by turnover and Lesmond et al. (2004) show that momentum profits accrue to the securities most expensive to trade; neither measures the gap between detectability and exploitability as a quantity in its own right, which is what a market this illiquid permits. Lesmond et al. (2004) triangulate across four cost estimators without reporting how far they disagree; we quantify the disagreement, trace it to the illiquidity the estimators are meant to measure, and propose the published fee floor as the alternative anchor. And Salhi and Ouakil (2026), studying this exchange over an overlapping period, raise the possibility that transaction costs undo their hedging results without measuring them.

The claim the evidence supports is narrower than it may first appear, and stating its scope precisely is part of the contribution. We do not show that statistical detectability and economic exploitability are separate properties of financial signals in general, which a single mechanism on a single exchange could not establish. What the evidence supports is an implication that fails: statistical detectability does not by itself establish economic exploitability. A significant relation can coexist with a return that stays below the cost of trading it, and the paper measures how far below. Nor are the two in opposition: a security's t -statistic and its recovery are positively rank correlated, at $+0.43$ across the 56 securities in the cost test. But the association is far too weak to be useful, and all 10 securities with the strongest statistical evidence remain in deficit, at a mean of -0.60 percentage points.

The remainder of the paper is organised as follows. Section 2 places the paper in the literature on trading costs, cost estimation and researcher degrees of freedom. Section 3 describes the two sources, the panel and the sample. Section 4 measures trading costs and establishes the estimator disagreement. Section 5 reports the central result. Section 6 reports robustness, and Section 7 concludes.

2 Literature Review

2.1 Trading costs and the survival of anomalies

That transaction costs erode documented anomalies is not in dispute; what is disputed is how much, and which ones survive. Novy-Marx and Velikov (2016) provide the reference treatment, sorting 23 anomalies into low-, mid- and high-turnover groups and reporting gross return, turnover, cost and net return for each. Their organising variable is turnover: most anomalies whose one-sided monthly turnover is below 50% still generate a significant net spread when the strategy is designed to mitigate costs, few above it do, and costs reduce the realised spread by more than 1% of monthly one-sided turnover. They also show that mitigation matters as much as selection: a buy/hold spread rule, entering at an extreme decile but exiting only at a second threshold, cuts turnover by roughly 40% and is the single most effective simple technique they test.

Lesmond et al. (2004) make the sharper claim for a single anomaly. Momentum returns, they show, accrue disproportionately to the securities that are most expensive to trade: across size quintiles the extreme-portfolio return falls from 3.91% to 1.00% while the estimated cost falls from 12.92% to 2.32%. Gross returns of 7.8% meet costs of 9.4%, and the strategy that appears profitable before costs is not profitable after them. Their conclusion is that the returns “create an illusion of trading profit opportunities when, in fact, none exist.”

Two features of that paper matter for ours. The first is the mechanism: the signal and its cost are carried by the same securities, so the anomaly cannot be separated from the friction that consumes it. The second is methodological. Lesmond et al. (2004) estimate costs four ways: quoted spreads, effective spreads from execution prices, the Roll estimator, and the limited-dependent-variable model of Lesmond et al. (1999), which infers the round-trip cost from the frequency of zero-return days. They describe the last as the more comprehensive and report results under all four. That the authors of the most influential cost-based refutation in the literature felt obliged to triangulate across four estimators is the background against which Section 4 of this paper should be read.

2.2 Estimating costs without quote data

Where quote-level data are unavailable, as on most frontier exchanges, the spread must be inferred from daily prices. Roll (1984) recovers it from the negative serial covariance that bid–ask bounce induces in transaction prices. Corwin and Schultz (2012) recover it from daily high–low ranges, exploiting the fact that the range reflects volatility, which scales with the interval, and the spread, which does not. Lesmond et al. (1999) take a different route, inferring the round-trip cost from the frequency with which a security posts a zero return, on the reasoning that a security whose expected price move is smaller than its cost of trading will not trade. Amihud (2002) measures price impact rather than spread, as the ratio of absolute return to the value traded.

These estimators are known to behave differently in thin markets, but the literature rarely quantifies the disagreement for a given sample, and applied work typically selects one without justifying the choice. Section 4 documents what that choice costs on this market, on a cross-section large enough that the comparison is not itself underpowered.

2.3 Non-synchronous trading

Lo and MacKinlay (1990) show that non-synchronous trading induces spurious cross-autocorrelation in portfolio returns. Their concern is inferential, that apparent lead–lag relations may be artefacts of when securities happen to print, but the same mechanism has a tradable reading: a security that has missed a market move will, on average, close part of the gap when it next trades. Whether that reading survives its own transaction costs is an empirical question that the frontier-market literature has largely not asked, in part because the panels available do not

distinguish a session with no trade from a session with an unchanged price. Section 3 sets out how we make that distinction and Section 5 asks the question.

2.4 Researcher degrees of freedom

Menkveld et al. (2024) identify a source of uncertainty that a standard error does not carry. Having 164 research teams test six hypotheses on a single dataset, they find that variation attributable to the teams' evidence-generating processes, their "non-standard errors", is on par with standard errors, correlates only weakly with team expertise or reproducibility, and is systematically underestimated by the researchers themselves. The implication for applied work is that a point estimate understates uncertainty by roughly a factor of two unless the design space is enumerated. Section 4 measures one such degree of freedom directly: on this market the choice between two standard spread estimators moves the reported shortfall by a factor of two.

2.5 Contribution

The contribution is not that costs matter, which Novy-Marx and Velikov (2016) and Lesmond et al. (2004) establish, nor that a frontier market is illiquid, which is definitional. It is twofold and narrow.

First, we measure the gap between statistical detectability and economic exploitability directly, rather than inferring it from a strategy's net return. A predictability significant on 21 of 60 securities under standard inference, and on 26 of 60 under pairs-bootstrap inference, returns 0.228% over the five sessions following an episode, against a published fee floor of at least 1.10%. The quantity of interest is the shortfall, and it is estimable with a precision that a portfolio Sharpe ratio does not offer. The claim it supports is that detectability does not *imply* exploitability, not that the two never coincide: a signal large enough to clear its costs would satisfy both.

Second, we anchor the cost side on published fees rather than on an estimated spread, having established that the available estimators disagree by a factor of four on this market and rank its securities inconsistently. Where a taxonomy in the tradition of Novy-Marx and Velikov (2016) requires one cost number per strategy, and the estimators disagree, a published fee schedule is the more defensible anchor and yields a strict lower bound.

3 Data

3.1 Sources

The Casablanca Stock Exchange publishes no quote-level data and no consolidated historical file. A long panel therefore has to be built from two sources that were produced for different purposes and do not measure the same thing.

The first is a commercial daily series, the Casablanca end-of-day history distributed by Investing.com (Fusion Media Ltd), covering the listed equities from January 2010 to the present: 379,546 security-day observations on 74 securities, with a closing price throughout and open, high, low and traded volume on 52% of them. It is long, and it is derived.

The second is the exchange's own daily record, which reaches back three years: 37,455 rows inside the sample, covering 55 securities over 704 sessions from 22 August 2023. Each row carries the opening price, the last price, the session high and low, the traded value in dirhams, the number of shares traded, the number of trades, and the security's market capitalisation. Two properties make it the reference. It prints a row for *every* listed security on *every* session, so a session on which nothing traded is stated rather than inferred. And it reports the traded value directly, rather than leaving it to be reconstructed from a volume whose units have to be assumed.

The design follows from that division. The commercial series supplies the history; the exchange record establishes, on the three years where the two overlap, that the history can be read the way the tests below read it. Each of the next three subsections states one such check, its result, and what it licenses.

3.2 Prices and corporate actions

A price series is quoted on a basis, and the basis changes when a corporate action does. The panel therefore carries two price columns rather than one. `close_raw` is the price as it printed on the session. `close_adj` is that same price expressed on the basis in force at the end of the sample, obtained by dividing the quote by the cumulative factor of every action that follows it. Both are derived from one series and one dated action table, each period's factor standing on its own rather than being rescaled into place from a neighbouring one. Market capitalisation, tick size and a minimum-commission calculation need the quoted level; a return needs the restated one.

Each commercial export is classified onto a basis against the exchange record rather than assumed to share one, by the median ratio of the two over their common sessions: a ratio of one identifies a raw series, any other stable ratio identifies a restated one and measures the factor. Appendix B reports the classification.

Corporate actions are declared rather than inferred from price ratios. The exchange reports each security's capitalisation alongside its close, so the point-in-time share count is capitalisation divided by price, and a step in that count dates an action to the session. 10 of the 14 actions the panel carries fall in the range where no price-ratio test can separate an action from a price move, which is precisely the range a ratio-based detector misses.

Declaring an action is not the same as adjusting for it. A share count rises for two reasons that look identical in a capitalisation series and are opposite in their effect on price, a split against a capital increase placed for cash, so each action is tested against the quote either

side of it and only those the quote responded to restate the series. Appendix B sets out both cases.

3.3 Trade status

Everything in Section 5 rests on knowing which sessions a security did not trade, so the distinction is made from the record where the record states it and identified where it does not.

Inside the exchange's window the status is given: a listed security with no shares traded on a session did not trade on it. Outside that window a security is traded on the sessions it is quoted on, and not traded on the sessions of the exchange calendar on which no quotation exists.

The second rule is testable against the first on the overlap, and the test is not the trivial one. Of the 31,691 security-sessions on which both sources speak, 85.5% carry a trade, so a rule that simply answered "traded" would score 85.5%. The identification agrees with the exchange's stated status on 99.97%: of 4,609 sessions the exchange records as untraded, the commercial series is silent on 4,604, and of 27,082 it records as traded, the commercial series quotes 27,079. The remaining 8 disagreements are spread across 6 securities.

On a non-trading session the exchange still reports a price, the level carried from the last trade. It is a level, not a transaction. Every such row is flagged, and no return is measured across one: a return runs from one trade to the next and records how many sessions it spans. That matters more than it sounds: consecutive trades span more than one session on 14.7% of occasions, up to 189 of them, and reading two adjacent observations as a one-session return is wrong that often.

3.4 Turnover

Turnover enters the Amihud measure as a denominator, so an error in it is an error in every illiquidity figure the paper reports. Inside the exchange's window the traded value is taken from the exchange, which states it. Outside it, traded shares from the commercial series are valued at the session's quote.

The second is only usable if the two agree where both exist, and they do, on every security that carries both except the single thinnest line in the sample; Appendix B reports the reconciliation. Every row records which of the two sources it came from, and across the study panel 16.4% of traded sessions carry the exchange's own figure.

3.5 Reference-price resets

The exchange caps the variation a quotation may show against the session's reference price at 10%, and no transaction may be executed outside the resulting range (Autorité Marocaine du Marché des Capitaux, 2026b). A consecutive-session move larger than 10% is therefore not an

improbable price move but an impossible one: either the reference price was itself reset, by a distribution or a corporate action, or the basis of the quotation changed. Such moves are flagged rather than deleted, and Section 6 reports the central result with and without them; Appendix B reports how many there are.

3.6 The session calendar and coverage

The panel sits on the exchange session calendar rather than on the observation index, which is what makes a non-trading interval an observation rather than an absence. The calendar is taken from the market index, which prints every session the exchange opens: 4,104 sessions between 4 January 2010 and 30 June 2026, between 245 and 254 in each full year.

The exports were pulled on different days, so each security also carries its own coverage end, the last session any source can speak for it. Sessions past that end are recorded as outside coverage, not as non-trading. The two look identical in the data, since both are a series that stops, and the distinction is the one the whole panel turns on, so it is carried explicitly: a line that ceased to quote is separated from a line whose export simply ended earlier, and the listing event is taken from the public record where one exists. Appendix B reports the breakdown.

Securities that stop quoting during the period are retained and simply end. Defining the universe as the securities quoted at the end of the sample and walking backwards would exclude precisely the securities whose outcomes were worst, so membership is defined per session: a security is in the universe on session t if t lies between its first and last covered session. Survivorship freedom is then a property of the construction rather than a correction applied to it.

3.7 Sample construction

Table 1: Sample construction

Step	Securities	Episodes	Criterion
Listed equity lines in the panel	83		ordinary shares; subscription and attribution rights excluded
Study universe	73		at least 500 traded sessions over 2010–2026
Post-thaw recovery test	60	25,966	at least 30 non-trading episodes
Cost test	56	8,009	at least 20 episodes preceded by an upward market

Note. Each step’s criterion is a power requirement fixed before the corresponding test is run, not a selection on any outcome. The episode column counts the non-trading episodes each test draws on: the recovery test uses every episode, the cost test only those preceded by an upward market. The 500-session threshold removes 10 lines with less than about two years of trading, on which neither a spread estimator nor an episode count has enough to work with; it does not select on the quantity the tests measure: the security that trades least in the universe is absent from 82.9% of sessions and is retained. The 10 excluded lines are listed with their reason in the replication package.




Figure 1: The cross-section by year: securities in the point-in-time universe (bars, left axis) and the share of security-sessions on which a trade occurred (line, right axis).

The panel holds 247,319 security-sessions, of which 177,292 carry a trade. The cross-section grows from 54 securities quoted in 2010 to 68 in 2026 (Figure 1). Non-trading is the dominant feature: the share of sessions on which a security does not trade ranges from zero to 82.9% across the universe, with a median of 17.5%. Section 5 is about what that produces.

3.8 Limitations

The panel carries *price* returns, not total returns. No dividend-per-share series exists in either source, and Moroccan payout ratios are high enough that a detachment can reset the reference price outright; Appendix B shows that the falls large enough to do so cluster in the general-meeting season. A detachment lowers the quote by roughly the cash it distributes, so it is a fall no signal predicts and no holder bears. They are flagged, Section 6 reports the result without them, and the direction is known in advance, since a detachment lowers the measured recovery, so removing them can only widen the deficits reported here, which is what it does. A study of long-horizon total returns on this market would nonetheless need a dividend series that neither source contains.

Second, the exchange’s own record reaches back three years. For the years before it, the trade-status distinction rests on the rule stated above, which the overlap tests at 99.97% but cannot prove for years the record does not reach. Section 6 reports the central test on the covered window separately, where no identification is needed at all.

4 Trading Costs

4.1 The measures

With no quote-level data to read a bid–ask spread from, the spread has to be inferred from daily prices. Two estimators are standard. Roll (1984) recovers the spread from the negative serial covariance that bid–ask bounce induces in transaction prices; Corwin and Schultz (2012)

recover it from the daily high–low range, on the argument that the observed range reflects both true volatility, which scales with the interval, and the spread, which does not. Both were designed for continuously traded securities. This section documents that on this market they produce different answers, and that the difference is large enough to matter for the sign of a result.

Four measures are used and each is stated in full, because the comparison between them is the substance of this section. For security i on session t , let H_{it} , L_{it} and C_{it} be the session's high, low and closing price, r_{it} the log return to the close of t from the security's previous trade, and V_{it} its turnover in dirhams. Let $\mathcal{T}_i$ be the sessions of the exchange calendar that fall inside security i 's coverage window, and $\mathcal{D}_i \subseteq \mathcal{T}_i$ the sessions on which it traded and had also traded on the session immediately before, so that every return used spans exactly one session and none spans a carried level.

Non-trading frequency is the share of a security's coverage on which no trade occurred:

$$Z_i = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \mathbf{1}\{\text{security } i \text{ did not trade on session } t\}. \quad (1)$$

Amihud (2002) illiquidity is the average absolute return per million dirhams traded:

$$\text{ILLIQ}_i = \frac{10^6}{|\mathcal{D}_i|} \sum_{t \in \mathcal{D}_i} \frac{|r_{it}|}{V_{it}}. \quad (2)$$

The Roll (1984) spread follows from the first-order autocovariance of those returns:

$$S_i^R = 2\sqrt{-\widehat{\text{Cov}}(r_{it}, r_{i,t-1})}, \quad (3)$$

which is defined only where the autocovariance is negative, as it is for 55 of the 73 securities.

The Corwin and Schultz (2012) spread is estimated on each pair of consecutive traded sessions. Writing $k = 3 - 2\sqrt{2}$,

$$\beta_{it} = \left[\ln \frac{H_{it}}{L_{it}} \right]^2 + \left[\ln \frac{H_{i,t+1}}{L_{i,t+1}} \right]^2, \quad \gamma_{it} = \left[\ln \frac{\max(H_{it}, H_{i,t+1})}{\min(L_{it}, L_{i,t+1})} \right]^2, \quad (4)$$

$$\alpha_{it} = \frac{\sqrt{2\beta_{it}} - \sqrt{\beta_{it}}}{k} - \sqrt{\frac{\gamma_{it}}{k}}, \quad S_{it}^{\text{CS}} = \frac{2(e^{\alpha_{it}} - 1)}{1 + e^{\alpha_{it}}}, \quad (5)$$

and S_i^{CS} is the mean of $\max\{S_{it}^{\text{CS}}, 0\}$ over the pairs available, the censoring at zero being the convention of the original paper. Its overnight adjustment is applied before β and γ are formed: where session $t + 1$ opens clear of session t 's close, both $H_{i,t+1}$ and $L_{i,t+1}$ are shifted by the size of the gap, so that an overnight move is not counted as intraday range.

The restriction of equations (2) and (3) to $\mathcal{D}_i$ is not a refinement. A return measured across a non-trading interval is the accumulation of several sessions, and admitting it would

inflate the numerator of the Amihud ratio and distort the autocovariance the Roll estimator reads. On a market where consecutive trades span more than one session on 14.7% of occasions, that is not a rare correction.

4.2 Rank correlations among the measures

Table 2: Rank correlations among illiquidity and cost measures

	Non-trading	Amihud	Corwin–Schultz	Roll
Non-trading frequency	1.00			
Amihud illiquidity	0.56***	1.00		
Corwin–Schultz spread	−0.70***	−0.08	1.00	
Roll spread	0.44***	0.34**	−0.16	1.00
Median value	17.5%	0.596	0.414%	1.618%
Defined for	73/73	73/73	73/73	55/73

Note. Spearman rank correlations, pairwise complete, over the 73 securities with open–high–low data, 2010–2026. The four measures are defined in equations (1)–(5). Turnover in equation (2) is taken from the exchange’s own traded value where it is available and from traded shares valued at the session’s quote elsewhere. *** denotes $p < 0.01$ and ** denotes $p < 0.05$.




Figure 2: Cost estimator behaviour. Left: the Corwin–Schultz spread declines as the share of sessions with no trade rises. Right: paired estimates by security, sorted by the Roll spread.

Two facts stand out, and they differ in how firmly they are established.

The first is a matter of levels and is not in doubt. The median Roll spread is 1.618% and the median Corwin–Schultz spread 0.414%, a factor of four. This is a direct comparison of estimates, not an inference, and Figure 2 shows the gap is systematic rather than driven by outliers: Roll exceeds Corwin–Schultz for 52 of the 55 securities where both are defined.

The second is a matter of ranking, and Table 2 carries it. Corwin–Schultz is negatively rank-correlated with the non-trading frequency: the less frequently a security trades, the nar-

rower this estimator says its spread is. Roll moves in the opposite direction and agrees with Amihud illiquidity, while Corwin–Schultz is uncorrelated with either. On 73 securities these are not underpowered comparisons. The two estimators rank the cross-section in ways that have nothing to do with each other, and one of them ranks it backwards.

4.3 The source of the disagreement

The usual explanation is a degeneracy on stale sessions: where a security does not trade, the reported high, low and close coincide, the observed range is zero, and the variance the estimator attributes to the spread collapses. That explanation does not survive here. The panel carries no high or low on a session with no trade, so the estimator is already computed on consecutively *traded* sessions and the degeneracy cannot arise mechanically. Restricting the pairing further to sessions one calendar day apart moves the median from 0.414% to 0.452% across the 73 securities where both can be computed, and the negative association with non-trading frequency survives at $\rho = -0.66$.

What does remain is censoring. Across the cross-section the estimator returns a negative spread on a median 16.0% of day-pairs, censored to zero by the usual convention. Roll has its own vulnerability in this setting: stale prices induce positive autocorrelation that offsets bid–ask bounce, which is why it is undefined for 18 of 73 securities here. Each estimator is therefore compromised by the same illiquidity in a different direction, and nothing in the data decides between them.

4.4 A cost floor from the published schedule

Because neither estimator can be defended as the right one, the paper’s central test does not use either. It uses the explicit fee schedule, which is published and requires no estimation.

A Moroccan equity transaction can involve four cost components, each charged on each side of a round trip, with value-added tax applied to the commissions at 10%. The Casablanca Stock Exchange charges 0.10% of the transaction value, the rate it publishes for its own commission (Bourse de Casablanca, 2026) and the rate itemised for an ordinary transaction in the worked retail example of La Vie Éco (2006). The account-keeping institution charges a settlement-and-delivery commission that the exchange puts at about 0.20% and that reaches 0.30% on one 2026 bank schedule (Banque Marocaine pour le Commerce et l’Industrie, 2026). A bank collecting an order on behalf of a client and routing it to a brokerage firm charges a further 0.10%, avoidable by placing the order with the brokerage directly. Only the brokerage commission is negotiable, and it is the one component the exchange does not publish, describing it instead as agreed between the parties; the rates that are published run from 0.20% to a stated maximum of 0.60% (La Vie Éco, 2006; Banque Marocaine pour le Commerce et l’Industrie, 2026).

Table 3: Explicit transaction costs, Moroccan equities

Component	Rate per side	Set by	Negotiable
Exchange commission	0.10%	the exchange	no
Settlement and delivery	0.20–0.30%	account keeper	no
Order collection	0.10%	collecting bank	avoidable
Brokerage commission	0.20–0.60%	brokerage	yes
Value-added tax	10% of commissions	tax code	no
Minimum charge	per order	brokerage	yes
Custody	0.075% per quarter	account keeper	no
Round trip, non-negotiable only	0.66%	zero brokerage; a bound, not a rate	
Round trip, every rate at its floor	1.10%	0.20% brokerage, 0.20% settlement	
Round trip, one 2026 schedule	2.20%	0.60% brokerage, 0.30% settlement	

Note. Rates as published by the Casablanca Stock Exchange (Bourse de Casablanca, 2026), as itemised in the worked retail example of La Vie Éco (2006), and as set out in the retail tariff schedule of a Moroccan account-keeping bank in force in 2026 (Banque Marocaine pour le Commerce et l'Industrie, 2026). Round-trip figures are $2 \times (\text{sum of applicable commissions}) \times 1.10$ for value-added tax, and all three assume an order placed directly with the brokerage, so none of them carries the order-collection commission. The first row omits brokerage entirely and is an arithmetic bound rather than an available rate. The second takes every component at the bottom of its published range and is the floor used throughout the paper; no rate in it is assumed. The third is a schedule an investor can actually be charged today. Two costs are excluded from every round-trip figure: the per-order minimum that brokerage schedules set, which raises the effective rate on small orders and whose level the exchange does not publish, and custody, which is levied quarterly on the portfolio rather than on a transaction and comes to about 0.30% a year at the smallest band. Both exclusions understate what trading costs.

On the capital gains tax. The *taxe sur les profits de cession de valeurs mobilières* is often described loosely as a transaction levy. It is not: Article 73 of the *Code Général des Impôts* sets it at 15% of the net realised profit on listed shares (Autorité Marocaine du Marché des Capitaux, 2026a). It is therefore excluded from every cost figure in this paper. Its omission is conservative in the direction that matters, since a tax on profit cannot improve a strategy that loses money before tax and reduces by 15% any surplus a strategy does earn.

This floor has three properties that the estimated spreads lack. It is observable rather than inferred. It is a strict lower bound, since an investor pays the spread in addition to it and never instead of it. And it is conservative in the direction that matters: any error in the spread estimates can only widen the deficits reported in Section 5, never narrow them.

The 0.66% figure deserves emphasis even though no investor can trade at it. It is what a round trip would cost if brokerage were free and the account keeper charged the lowest settlement rate observed, and it is the tightest arithmetic bound the published schedules permit. Section 5 reports the central test against it as well, so that the result can be read without any assumption about what a brokerage charges.

The spread between the bound and the schedule an investor is actually quoted is itself the point. The same round trip costs 0.66%, 1.10% or 2.20% depending on where the order is placed and who keeps the account, and none of those three figures requires an estimator. The paper uses the lowest defensible one throughout, which is the assumption least favourable to its own conclusion.

4.5 Implications for the tests

Section 5 reports its central result against the fee floor and then, as robustness, against the non-negotiable components with brokerage removed, against a published 2026 retail schedule, and against the floor plus each median spread estimate in turn. The conclusion is unchanged across all five assumptions, and the estimator disagreement is why they have to be shown: the deficit ranges from 0.43 to 2.49 percentage points depending on which cost number is used, so any single-estimator statement would be reporting a design choice as a magnitude.

The broader implication is for empirical work on markets of this type. Taxonomies of anomalies net of costs, in the tradition of Novy-Marx and Velikov (2016), require a cost number per strategy. Where the implicit estimators disagree by a factor of four and rank securities inconsistently, the choice of estimator becomes a researcher degree of freedom of the kind Menkveld et al. (2024) document. A published fee floor, where one exists, is the more defensible anchor.

5 Post-Thaw Recovery and Its Cost

5.1 Test design

For each security we identify every *episode*: a maximal run of consecutive sessions on which it did not trade. Because the panel sits on the exchange calendar, these runs are observed, not inferred from a zero return, a distinction that matters here because a zero return also occurs on sessions a security did trade, and on this market it does so on a median 15.0% of them. For each episode we record two quantities:

$$s_{i,\tau} = \sum_{t \in \text{episode}} m_t, \quad (6)$$

$$y_{i,\tau} = \sum_{h=0}^{H-1} r_{i,t_\tau+h}, \quad (7)$$

where m_t is the market factor on session t , t_τ is the security's first trade back, and $H = 5$ sessions. The signal is what the security missed; the outcome is what it earns over the five sessions that follow. The market factor is the cross-sectional median of that session's traded log returns; Section 6 reports three alternative constructions and shows that the median is the most conservative of the four.

The horizon is fixed *ex ante* by the settlement cycle. Moroccan equities settle on $J + 3$, stated by the regulator and by the exchange for an ordinary order (Autorité Marocaine du Marché des Capitaux, 2024; Bourse de Casablanca, 2026), so a position opened at the first trade back cannot be closed and the proceeds redeployed inside three sessions; five is the shortest horizon over which a round trip is actually completable, plus an allowance for instruction lag.

Section 6 reports horizons of 1, 3, 10 and 20 sessions.

Three questions are then asked in sequence and kept apart, because the answer to the first does not answer the third:

1. Is there an identifiable relation between s and y ? A hypothesis test.
2. How large is the return it implies? A magnitude.
3. Does that return exceed the cost of capturing it? An economic feasibility question, and the only one an investor faces.

All three are asked of one episode construction: one horizon, one market factor, one definition of an episode. They differ only in which episodes they condition on and in how many securities clear the power requirement each needs, and Table 1 carries both chains.

5.2 Detectability

Of the 73 securities in the study, 60 have at least 30 non-trading episodes. Between them they carry 25,966, a median of 448 each. The episodes are short: the median is 1 session and no security's median exceeds 4, so what is measured is the recovery from a market move missed over a few sessions, not from a long suspension. For each security we compute the correlation between $s_{i,\tau}$ and $y_{i,\tau}$ and its t -statistic.

Table 4: Post-thaw recovery: detectability across the cross-section

	Standard inference	Pairs bootstrap
Securities tested	60	60
Significant	21	26
Expected under the null at 5%	3.0	3.0
Ratio to chance	$7.0\times$	$8.7\times$
Exact binomial p	5.7×10^{-13}	1.9×10^{-18}
Of the significant, positive sign	21 of 21	

Note. Securities with at least 30 non-trading episodes, 2010–2026. Standard inference uses the analytic distribution of a correlation coefficient. The pairs bootstrap resamples episodes with replacement, 2,000 draws per security, and assumes nothing about the shape of the return distribution; a security counts as significant when the 95% percentile interval excludes zero. The binomial test is one-sided against a false-positive rate of 5%.

The recovery is not a marginal statistical artefact. It is identifiable on 21 of the 60 securities that can be tested, seven times the rate chance produces, and the count rises rather than falls when the distributional assumption is removed. Every one of the significant cases has the sign the mechanism predicts.

5.3 Magnitude

Restricting to episodes preceded by an upward market, the ones a strategy would act on, 56 securities have at least 20 such episodes, 8,009 in total, a median of 151 per security. The mean five-session recovery is 0.228% and the median 0.256%, and 41 of the 56 have a positive mean recovery. The 4 securities that clear the episode requirement but not the upward one are S.M Monétique, BCP, CGI and SNEP.

Set beside the security’s own variation rather than beside a fee, that figure is small in a way that explains the rest of the paper. The median security’s one-session return has a standard deviation of 2.43%, so 5.44% over five sessions, and the recovery is 4.2% of it. A return that size is identifiable across a cross-section of 60 securities and invisible inside any one episode.

The recovery is also larger where the mechanism says it should be. Table 5 sorts the 56 securities into terciles of non-trading frequency.

Table 5: Recovery by tercile of non-trading frequency

Tercile	Securities	Non-trading %	Mean recovery %	Mean t
Trades most often	19	9.40	0.087	1.29
Middle	18	25.35	0.179	1.90
Trades least often	19	58.00	0.415	2.04

Note. The 56 securities of the cost test, sorted into equal thirds by the share of sessions on which they do not trade. “Non-trading” is the median of that share within the tercile; “Mean recovery” and “Mean t ” are means across the securities in it of the quantities reported for each security in Appendix A.

Both columns rise monotonically with thinness, which is the ordering non-synchronous trading predicts. The rank correlation across the 56 individual securities is positive but does not reach significance ($\rho = +0.12$, $p = 0.39$), so the pattern is consistent with the mechanism rather than a test of it.

A recovery of 0.228% is the whole of the opportunity. A round trip on a Moroccan equity costs at least 1.10% in published fees alone, before any bid–ask spread (Section 4).

5.4 Exploitability

Figure 3 puts the two halves of the result side by side: a distribution of t -statistics displaced well to the right of the conventional threshold, and a set of magnitudes almost entirely to the left of the cheapest fee the exchange’s schedule permits.

The conclusion does not depend on the cost assumption, and it does not depend on what a brokerage charges. Set the brokerage commission to zero, an assumption no investor can obtain, and the mean recovery still falls short of the remaining non-negotiable components by 0.43 percentage points ($t = -5.85$, $p = 2.8 \times 10^{-7}$), with 43 of 56 securities in deficit.




Figure 3: Left: the distribution of recovery t -statistics across the 60 securities tested, with the conventional thresholds marked. Right: each security’s mean five-session recovery against the published fee floor and against the non-negotiable components alone.

Table 6: Net recovery against five cost assumptions

Round-trip cost	Mean net	t	p	Deficits	Wilcoxon p
0.66% non-negotiable components only	−0.43	−5.85	2.8×10^{-7}	43 of 56	1.3×10^{-6}
1.10% published fee floor	−0.87	−11.81	1.0×10^{-16}	53 of 56	1.5×10^{-10}
1.51% floor plus median C–S	−1.28	−17.36	5.6×10^{-24}	56 of 56	7.5×10^{-11}
2.20% a published 2026 schedule	−1.97	−26.70	3.6×10^{-33}	56 of 56	7.5×10^{-11}
2.72% floor plus median Roll	−2.49	−33.75	1.9×10^{-38}	56 of 56	7.5×10^{-11}

Note. Mean net is in percentage points over five sessions, across the 56 securities with at least 20 upward episodes. The t -statistic is from a one-sample test on the 56 security means; the Wilcoxon column is the signed-rank test on the same quantities. “Deficits” counts securities whose mean recovery falls short of the cost. The three schedule levels are the round trips of Table 3; the first of them omits brokerage entirely and is an arithmetic bound rather than an available rate. The two compound levels add to the fee floor the corresponding median spread of Table 2, $1.10 + 0.414$ and $1.10 + 1.618$, over the 73 securities of that table rather than the 56 tested here.

Two further readings sharpen the result. Testing each security separately against the fee floor, *none* of the 56 shows a significant surplus and 32 show a significant shortfall. And bootstrapping the mean over securities, the 95% interval for the net at the floor is $[-1.02, -0.73]$ percentage points; none of 20,000 resamples places it above zero.

The breakeven fee is 0.228% round trip. It is roughly a fifth of the cheapest schedule the exchange’s published commissions permit, and about a third of what a round trip would cost if brokerage were free.

5.5 A balanced panel

The 56 securities do not share a window: some quote from 2010, others begin later, and the cross-sectional spread could in principle reflect securities having lived through different market

regimes rather than differing in liquidity. We therefore repeat the test on the longest window over which a large majority of the universe quotes continuously, 55 securities across 2,609 sessions from January 2016, so that every security is measured on exactly the same period.

Of these, 50 clear the episode requirement and 15 show a significant recovery against 2.5 expected (binomial $p = 1.3 \times 10^{-8}$), and across the 47 securities that clear the cost requirement the mean net at the fee floor is -0.81 percentage points ($t = -9.47$, $p = 2.3 \times 10^{-12}$). The deficit on a common window is within 0.06 points of the deficit on the full sample.

5.6 The three answers together

The relation is identifiable, on 21 securities of 60 under standard inference and on 26 without the distributional assumption. Its magnitude is 0.228% over five sessions, which is 4.2% of what the median security moves in the same time. And the cheapest round trip the published schedule permits is 1.10%, which no security in the cross-section earns back. None of that is in tension. A relation can be real, small, and unprofitable at once, and the three questions are worth keeping apart precisely because it can be. Appendix A reports the cost and recovery figures of Sections 4 and 5 for every security in the universe, so that the cross-sectional aggregates can be read against the securities that carry them.

6 Robustness

Each of the design choices in Section 5 could have produced the result on its own. This section tests them in turn: the holding horizon, the construction of the market factor, dependence across securities and within a security, the distributional assumptions behind the inference, and three variations in what enters the sample.

6.1 Horizon

The five-session horizon was fixed *ex ante* by the $J + 3$ settlement cycle, not chosen because it works. Table 7 reports the whole grid.

The shape is the honest reading of the result and we state it plainly. Detectability holds at every horizon. The shortfall narrows monotonically as the horizon lengthens, from -1.22 points at 1 session to -0.33 at 20, and it is where that narrowing ends that the result has a boundary: the deficit is significant at horizons through 10 sessions and, at 20, no longer is ($p = 0.058$). We report that rather than the horizon grid that would have avoided it. Two things bound how much it concedes. A 20-session horizon is not the mechanism this paper describes, since a post-thaw catch-up that takes a month to accrue is no longer a catch-up. And holding an illiquid Moroccan equity for a month to earn 0.77% gross exposes the position to a month

Table 7: The central result across holding horizons

Sessions	Significant	Binomial p	Gross	Mean net	t	Deficits
1	23 of 60	4.5×10^{-15}	-0.12%	-1.22	-19.98	56 of 56
3	25 of 60	2.8×10^{-17}	0.08%	-1.02	-14.56	56 of 56
5	21 of 60	5.7×10^{-13}	0.23%	-0.87	-11.81	53 of 56
10	19 of 60	5.3×10^{-11}	0.40%	-0.70	-6.39	47 of 56
20	17 of 60	3.7×10^{-9}	0.77%	-0.33	-1.94	29 of 56

Note. Mean net is in percentage points against the 1.10% fee floor. Detectability is significant at every horizon. The deficit narrows as the horizon lengthens, because the recovery accumulates while the round-trip cost is paid once. It is significant at every horizon through 10 sessions and falls short of significance at 20 ($p = 0.058$).

of that security's own volatility, which the test does not charge for, so the 20-session figure overstates what such a strategy nets.

6.2 Construction of the market factor

The market factor enters the signal in equation (6), so its construction could be driving the result. The cross-sectional median is exactly zero on 48.9% of sessions, which is what a market this thin produces and which leaves the factor uninformative on half the calendar. Table 8 reports three alternatives.

Table 8: Four constructions of the market factor

Construction	Tested	Significant	Gross	Mean net
Cross-sectional median (used above)	60	21	0.23%	-0.87
Cross-sectional mean	60	27	0.06%	-1.04
Market index	60	27	0.04%	-1.06
Median of the 20 most-traded securities	60	22	0.14%	-0.96

Note. Mean net is in percentage points against the 1.10% fee floor, five-session horizon. The construction used in Section 5 produces the *smallest* deficit of the four and the lowest count of significant recoveries.

The choice is conservative in both directions that matter. The median gives the largest gross recovery and therefore the smallest shortfall; any of the three alternatives would have let us report a wider deficit. It also gives the lowest detectability count of the four.

6.3 Cross-sectional dependence

The binomial test in Table 4 treats the 60 securities as independent, which they are not: they share a market. We therefore build a null that preserves exactly that dependence. The market factor is circularly shifted by a random number of sessions, which destroys its timing relative to each security's non-trading intervals while leaving its own autocorrelation and every security's

co-movement with every other untouched. Under that null, the count of significant recoveries averages 4.4 with a 95th percentile of 8; the observed count is 21, and no draw in 1,000 reaches it.

For serial dependence in calendar time, we resample the episodes in six-month blocks, 33 of them over the sample, and recompute the mean net at the fee floor within each resample. The 95% interval is $[-1.03, -0.66]$ percentage points, and no resample places it above zero. It is close to the interval from resampling securities independently, $[-1.02, -0.73]$, and it does not change the conclusion.

6.4 Within-security dependence

A security contributes many episodes, and if consecutive episodes were dependent the per-security t -statistics would be overstated. The first-order autocorrelation of the cross-product term entering each correlation averages $+0.009$ across the 56 securities in the cost test, which a one-sample test does not distinguish from zero ($t = 0.54$, $p = 0.59$), and 12 securities have $|\rho_1| > 0.10$.

A moving-block bootstrap within each security’s episode sequence gives the same answer from the other direction. Resampling in blocks of one, five and ten consecutive episodes, the number of securities whose correlation interval excludes zero is 28, 29 and 27 out of 60, stable across the block lengths tested, consistent with independence across episodes rather than the sensitivity that dependence would produce.

6.5 Distributional assumptions

Two of the results above are already non-parametric. The detectability count of 26 in Table 4 comes from a pairs bootstrap that assumes nothing about the return distribution, and it exceeds the parametric count of 21 rather than falling below it. On the cost side, the t -test in Table 6 is corroborated by the Wilcoxon signed-rank test at every cost level ($p \leq 1.5 \times 10^{-10}$ at the fee floor) and by a sign test ($p = 4.1 \times 10^{-13}$). The conclusion does not rest on normality anywhere.

6.6 Sample and definition

Three further variations change what enters the test rather than how it is estimated, and Table 9 reports all three against the baseline.

The first removes every consecutive-session move beyond the exchange’s variation threshold, from the securities and from the market factor alike. These are reference-price resets rather than price moves, and Section 3 flags them; dropping them removes 71 returns from the study universe. The second restricts the test to the window the exchange’s own record covers,

where the trade status is stated and no identification is needed at all. The third moves the end of the sample back to each year-end the panel can support.

Table 9: The central result under alternative samples

Sample	Tested	Significant	Gross	Mean net	t	Deficits
Baseline	60	21	0.23%	−0.87	−11.81	53 of 56
Excluding reference-price resets	60	21	0.22%	−0.88	−11.40	53 of 56
Exchange-covered window only	18	2	0.05%	−1.05	−5.13	9 of 9
Ending December 2021	56	19	0.26%	−0.84	−10.82	49 of 51
Ending December 2022	58	20	0.25%	−0.85	−10.44	51 of 53
Ending December 2023	58	21	0.26%	−0.84	−11.21	52 of 54
Ending December 2024	60	18	0.26%	−0.84	−11.51	52 of 54
Ending December 2025	60	20	0.25%	−0.85	−11.79	52 of 55

Note. Mean net is in percentage points against the 1.10% fee floor, five-session horizon, and the baseline row repeats the five-session row of Table 7. Significance is $|t| > 2$ on the per-security recovery correlation. The exchange-covered window is the 704 sessions from August 2023 on which the exchange states the trade status; 18 of the 48 securities it covers accumulate the 30 episodes the detectability test requires and 9 the 20 upward episodes the cost test requires, so its detectability count is uninformative in either direction while its cost verdict is unanimous.

The deficit is between −0.84 and −1.05 percentage points in every row, and is significant in every row. Removing the observations whose quotation basis is open to question widens it rather than narrowing it. Restricting to the window where the trade status is stated rather than identified widens it further and leaves no security in surplus. And the gross recovery drifts down as the sample lengthens, from 0.26% at the end of 2021 to 0.23% at the end of the sample, so the conclusion does not rest on where the window closes.

7 Conclusion

Statistical detectability does not by itself establish economic exploitability. This paper measures by how much it fails on one market where the measurement can be made precisely: post-thaw recovery on the Casablanca Stock Exchange is significant on 21 of 60 securities and worth 0.228% over five sessions, against a published round-trip fee floor of at least 1.10%, and no security in the cross-section earns a significant surplus.

Two implications follow, and each is wider than the result that produces it.

The first is about cost measurement. Where the standard implicit spread estimators can be shown to disagree, as they do here by a factor of four and with inconsistent rankings, a published fee schedule is the more defensible anchor: it is observable rather than inferred, and it is a strict lower bound, since an investor pays the spread in addition to it. Applied work that reports one estimated spread on a market of this kind is reporting a design choice as a magnitude. The disagreement is not a local curiosity, because it arises from the same illiquidity the anomaly literature is trying to price: a security that trades rarely gives a high–low estimator

little range to work with and a covariance estimator little bounce, and the two errors run in opposite directions.

The second is about how the two questions should be reported. A strategy's net return answers both at once and therefore answers neither: it cannot say whether a signal disappointed because it was not there or because it was too small to pay for itself. Kept apart, the answers here point in opposite directions. The recovery is significant on 21 of 60 securities against 3 expected by chance, and it is worth about a fifth of what collecting it costs. Ranking the securities by how strong the statistical evidence is and ranking them by how large the recovery is gives two orderings that agree only in part, at a rank correlation of +0.43, so neither can be read off the other. Applied work that reports a single net figure on a market of this kind is reporting the sum of a detection result and a cost result without saying which of them moved.

Scope

We do not show that detectability and exploitability are separate properties of financial signals in general; a single mechanism on a single exchange could not. What the evidence supports is narrower: a relation can be significant at $p < 10^{-12}$ across a cross-section and still return a fifth of what it costs to trade. Detection and collection are different questions, and the first does not answer the second.

Three scope conditions are explicit. The deficit narrows as the holding horizon lengthens, from -1.22 percentage points at one session to -0.33 at 20, where it is no longer significant ($p = 0.058$); a 20-session catch-up is in any case not the mechanism described here. The returns are price returns rather than total returns, because no dividend series exists in either source; since dividend detachment lowers measured recovery, correcting for it could only widen the deficits reported here. And the fee floor is the cheapest schedule the published commissions permit, which understates what most investors on this market pay.

Generalisation

The design is portable to any market with a published fee schedule and a thin tape, which describes most frontier exchanges, and three elements travel. Build the panel on the exchange's session calendar, so that non-trading is observed rather than proxied by a zero return: the proxy conflates two different events and here would have been wrong on one consecutive observation in seven. Anchor costs on published fees wherever the estimators can be shown to disagree. And keep the three questions apart rather than collapsing them into a single net-return statistic, which on this evidence is the difference between a shortfall that can be measured and one that cannot be interpreted.

Declarations

Data availability. The construction and analysis code and the derived study database are available at https://github.com/afchtaalahmedquant/cse-microstructure_paper-working. The raw source files are deposited at Zenodo (Afachtal, 2026), under restricted access and available on request. The repository documents both sources, their coverage, and the file registry and validation suite described in Appendix C.

Code availability. All construction and analysis code, in execution order, is provided in the same repository, together with a validation suite that asserts the properties of the data itself and a verification program that checks the manuscript against the results in both directions. Either failing stops the build.

Funding. This research received no grant from any funding agency in the public, commercial or not-for-profit sectors.

Competing interests. The author declares no competing financial or non-financial interests. No position in any security discussed in this paper is or was held by the author.

CRedit author statement. Ahmed Afachtal: Conceptualization; Methodology; Software; Validation; Formal analysis; Investigation; Data curation; Writing – original draft; Writing – review and editing; Visualization.

Use of generative AI. Generative AI tools were used for language editing and for drafting and reviewing analysis code. All empirical results were produced by the code in the replication package, and every number reported in the paper is checked against that code by the verification program described above. The author is responsible for all content.

References

- Afachtal, A. (2026). [dataset] Casablanca Stock Exchange daily record, august 2023 – june 2026. Zenodo. <https://doi.org/10.5281/zenodo.22712772>.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.
- Autorité Marocaine du Marché des Capitaux (2024). Guide pratique des actions. <https://www.ammc.ma/sites/default/files/Guide%20pratique%20des%20actions.pdf>. “Le règlement-livraison s’effectue à J+3.” Accessed 15 September 2026.
- Autorité Marocaine du Marché des Capitaux (2026a). Quel est le taux de la taxe sur la plus-value de cession d’actions à payer? <https://www.ammc.ma/fr/faq/quel-est-le-taux-de-la-taxe-sur-la-plus-value-de-cession-daction-payer>. Article 73, Code Général des Impôts. Accessed 15 September 2026.

- Autorité Marocaine du Marché des Capitaux (2026b). Seuil de variation. <https://www.ammc.ma/fr/seuil-de-variation>. Ten percent of the session's reference price for continuously quoted equities, six percent for securities traded at a fixing; no transaction may be executed outside the resulting range. Accessed 15 September 2026.
- Banque Marocaine pour le Commerce et l'Industrie (2026). Tarification bancaire. <https://www.bmci.ma/wp-content/blogs.dir/sites/2/2024/01/Affiche-tarifaire-2024.pdf>. Retail schedule in force 12 March 2026: brokerage 0.60 percent, settlement and delivery 0.30 percent, custody 0.075 percent per quarter at the smallest band. Accessed 15 September 2026.
- Bourse de Casablanca (2026). Pourquoi et comment investir: commissions et frais de transaction. <https://www.casablanca-bourse.com/fr/pourquoi-et-comment-investir>. Commission de bourse 0.10 percent; collecte des ordres 0.10 percent; règlement-livraison "fixée par les teneurs des comptes à environ 0.20%". On an ordinary order the page states "le compte titres du client est crédité et le compte espèces est débité en j+3". Accessed 15 September 2026.
- Corwin, S. A. and Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67(2):719–760.
- La Vie Éco (2006). Le coût d'une transaction. <https://www.lavieeco.com/argent/le-cout-dune-transaction-5989/>. Published 5 May 2006, updated 22 December 2022. Itemised retail example: exchange commission 0.10 percent, order collection 0.10 percent, settlement and delivery 0.20 percent, brokerage 0.60 percent; value-added tax, then 7 percent, subsequently raised to 10 percent. Accessed 15 September 2026.
- Lesmond, D. A., Ogden, J. P., and Trzcinka, C. A. (1999). A new estimate of transaction costs. *Review of Financial Studies*, 12(5):1113–1141.
- Lesmond, D. A., Schill, M. J., and Zhou, C. (2004). The illusory nature of momentum profits. *Journal of Financial Economics*, 71(2):349–380.
- Lo, A. W. and MacKinlay, A. C. (1990). An econometric analysis of nonsynchronous trading. *Journal of Econometrics*, 45(1-2):181–211.
- Menkveld, A. J., Dreber, A., Holzmeister, F., Huber, J., Johannesson, M., Kirchler, M., Neusüss, S., Razen, M., Weitzel, U., et al. (2024). Nonstandard errors. *Journal of Finance*, 79(3):2339–2390.
- Novy-Marx, R. and Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1):104–147.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4):1127–1139.

Salhi, Z. and Ouakil, H. (2026). Time-varying connectedness and hedging properties of commodities in the moroccan stock market sectors. *Finance Research Open*, 2(1):100105.

A The Cross-Section by Security

Table 10 reports every security in the study universe, sorted by the share of sessions on which it does not trade. It is the full version of the summary statistics used in Sections 4 and 5, and it lets a reader see which securities carry each result rather than taking the cross-sectional aggregates on trust.

Three columns deserve a word. The non-trading and zero-return columns measure different things and are reported separately: the first is the share of exchange sessions on which the security did not trade, the second the share of its *traded* one-session returns that were exactly zero. Roll is blank where the return autocovariance is non-negative and the estimator is undefined. The episode, t and recovery columns are blank where the security has fewer than 30 non-trading episodes and so does not enter the detectability test; the net columns are blank where it has in addition fewer than 20 episodes preceded by an upward market and so does not enter the cost test. A recovery shown against an empty net column is therefore a mean over too few episodes to test, reported for completeness and used nowhere.

Table 10: Illiquidity, cost and post-thaw recovery, by security

Security	Sec.	Trades	No-tr. %	Zero %	Amihud	C-S %	Roll %	Ep.	t	Recov. %	Net pts	t
AKDITAL	SAN	874	0.0	6.5	0.008	0.530	—	—	—	—	—	—
CFG BANK	BAN	622	0.0	10.5	0.008	0.856	0.888	—	—	—	—	—
DOUJA PROM ADDOHA	IMM	1244	0.0	4.7	0.012	0.724	—	—	—	—	—	—
ATTIJARIWABA BANK	BAN	4097	0.2	11.3	0.002	0.380	0.545	—	—	—	—	—
ITISSALAT AL-MAGHRIB	TCO	4094	0.2	10.0	0.001	0.283	—	—	—	—	—	—
TGCC S.A	BMC	1104	0.2	7.4	0.043	0.662	—	—	—	—	—	—
SODEP-MARSA MAROC	SDT	2463	0.3	9.3	0.015	0.487	—	—	—	—	—	—
STOKVIS NORD AFRIQUE	DIS	700	0.6	4.3	10.538	1.044	—	—	—	—	—	—
ALLIANCES	IMM	4077	0.7	15.7	0.081	0.869	—	—	—	—	—	—
BCP	BAN	4062	1.0	16.1	0.029	0.389	0.559	32	0.63	0.269	—	—
RESIDENCES DAR SAADA	IMM	2835	1.0	7.6	0.331	0.825	—	—	—	—	—	—
MUTANDIS SCA	AGR	1840	1.6	14.7	0.117	0.524	1.253	—	—	—	—	—
SNEP	CHI	4032	1.8	6.9	0.562	0.996	0.961	52	-1.39	-4.299	—	—
DISTY TECHNOLOGIES	LSI	958	1.9	8.9	0.274	0.663	0.303	—	—	—	—	—
BANK OF AFRICA	BAN	4010	2.3	22.5	0.120	0.512	1.116	86	1.76	1.063	-0.037	-0.08
ATLANTASANAD	ASS	4008	2.3	11.6	0.648	0.928	1.982	90	1.61	1.417	0.317	0.39
MANAGEM	MIN	3954	3.7	11.9	0.138	0.483	0.929	112	-1.01	-1.277	-2.377	-1.68
AFRIC INDUSTRIES SA	BMC	675	4.1	17.0	0.979	0.484	1.504	—	—	—	—	—
LAFARGEHOLCIM MAROC	BMC	3877	5.5	18.8	0.026	0.618	1.214	190	2.46	0.913	-0.187	-0.35
TAQA MOROCCO	ELC	2934	5.6	13.8	0.101	0.458	1.319	147	1.72	0.719	-0.381	-0.95
ARADEI CAPITAL	SPI	1297	5.9	15.7	0.234	0.459	1.658	71	1.72	0.559	-0.541	-1.37
IMMORENTE INVEST	SPI	1865	7.5	18.7	0.596	0.355	1.336	129	1.20	-0.496	-1.596	-3.78
AFMA	ASS	2418	7.7	44.3	1.383	0.028	1.213	89	-1.52	-0.706	-1.806	-2.15
DELTA HOLDING	SPH	3768	8.2	12.9	2.629	0.812	1.972	258	2.11	-0.104	-1.204	-3.08
SONASID	BMC	3718	9.4	10.3	0.224	0.572	—	286	0.61	-0.786	-1.886	-2.65
CIH	BAN	3709	9.6	11.8	0.279	0.516	1.840	317	3.99	0.828	-0.272	-0.89
TOTALENERGIES MARKETING MAROC	PEG	2481	9.8	12.6	0.106	0.403	1.035	216	1.48	1.231	0.131	0.25
MED PAPER	SP	3694	10.0	6.4	5.588	1.133	1.056	196	1.15	-0.679	-1.779	-1.30
JET CONTRACTORS	BMC	3170	12.3	8.9	0.638	0.651	1.263	275	1.25	0.087	-1.013	-1.68

continued

Table 10, continued

Security	Sec.	Trades	No-tr.	Zero	Amihud	C-S	Roll	Ep.	t	Recov.	Net	t
MINIERE TOUISSIT	MIN	3549	13.5	12.9	0.157	0.357	—	400	1.28	0.150	-0.950	-2.41
AUTO HALL	DIS	3552	13.5	14.5	1.463	0.544	1.826	387	1.94	-0.040	-1.140	-3.63
RISMA	LH	3532	13.9	12.9	1.106	0.624	1.478	383	0.55	0.085	-1.015	-2.78
DISWAY	LSI	3525	14.1	11.8	0.671	0.518	1.923	421	1.00	-0.438	-1.538	-3.27
STROC INDUSTRIE	IBE	3179	14.7	9.2	10.943	0.827	—	351	1.24	-0.866	-1.966	-3.11
COSUMAR	AGR	3457	15.7	22.7	0.300	0.419	1.194	351	2.75	0.380	-0.720	-2.59
CGI	IMM	1160	16.4	13.7	0.204	0.615	1.239	45	0.12	-4.299	—	—
COLORADO	BMC	3383	17.5	14.1	4.465	0.525	1.940	475	3.11	0.557	-0.543	-1.63
CIMENTS DU MAROC	BMC	3345	18.5	16.6	0.114	0.302	1.841	505	2.47	0.720	-0.380	-1.19
ENNAKL	DIS	3190	19.6	9.7	3.980	0.691	1.652	397	0.96	-0.432	-1.532	-3.70
MICRODATA	LSI	3283	20.0	14.4	1.418	0.372	1.802	557	2.17	0.350	-0.750	-2.20
HOLCIM MAROC (pre-2016)	BMC	1245	20.1	15.5	0.133	0.282	1.644	220	1.60	0.003	-1.097	-1.81
LESIEUR CRISTAL	AGR	3177	22.5	21.6	1.035	0.328	1.711	543	1.75	0.229	-0.871	-4.32
S.M MONETIQUE	LSI	545	22.6	7.4	3.249	0.677	1.781	74	1.18	1.922	—	—
REALISATIONS MECANQUES	DIS	531	24.6	7.1	7.882	0.815	—	80	1.26	-0.360	-1.460	-0.90
HPS	LSI	3064	25.3	17.2	1.600	0.414	1.774	567	1.88	0.352	-0.748	-2.03
FENIE BROSSETTE	DIS	3060	25.4	9.1	2.492	0.703	—	578	2.63	-0.356	-1.456	-2.88
BMCI	BAN	2891	29.6	19.1	0.455	0.329	1.906	706	0.87	0.008	-1.092	-3.74
LABEL VIE	DIS	2877	29.9	22.1	0.093	0.328	1.885	650	1.29	0.224	-0.876	-4.10
CARTIER SAADA	AGR	2818	31.3	16.7	9.804	0.498	2.135	633	-0.10	0.193	-0.907	-3.21
SANLAM MAROC	ASS	2592	33.1	19.0	0.381	0.411	2.612	644	2.47	0.478	-0.622	-2.16
CTM	TRP	2741	33.2	15.5	1.004	0.298	1.540	747	4.09	0.688	-0.412	-1.59
SALAFIN	SFA	2740	33.2	29.5	0.533	0.321	2.127	679	2.08	-0.402	-1.502	-5.91
SMI	MIN	2642	35.6	10.3	0.280	0.404	1.001	581	2.18	0.292	-0.808	-2.26
M2M GROUP	LSI	2628	35.9	12.1	2.009	0.294	2.152	699	0.72	0.300	-0.800	-2.03
INVOLYS	LSI	2625	36.0	15.0	3.549	0.444	1.505	712	2.75	0.102	-0.998	-2.78
WAFI ASSURANCE	ASS	2538	38.2	21.7	0.088	0.132	2.177	838	1.78	0.851	-0.249	-1.04
IB MAROC.COM	LSI	2484	39.5	12.6	15.284	0.401	—	667	2.49	-0.059	-1.159	-2.44
SOTHEMA	PHA	2372	42.1	52.6	1.266	0.022	0.984	361	1.16	1.207	0.107	0.26
AFRIQUIA GAZ	PEG	2236	45.5	21.2	0.136	0.216	1.761	791	3.23	0.693	-0.407	-1.83
ALUMINIUM DU MAROC	BMC	2088	49.1	22.7	0.317	0.176	1.507	753	1.11	0.405	-0.695	-2.88
SOCIETE DES BOISSONS DU MAROC	BOI	1872	54.3	24.2	0.351	0.210	2.315	763	3.08	0.420	-0.680	-2.57
CDM	BAN	1872	54.4	18.9	0.656	0.289	1.618	722	3.66	0.360	-0.740	-2.67
MAGHREB OXYGENE	CHI	1725	57.9	19.3	4.340	0.127	1.954	743	0.61	0.282	-0.818	-2.57
EQDOM	SFA	1721	58.0	24.0	0.340	0.068	1.516	778	4.00	0.092	-1.008	-3.70
NEXANS MAROC	IBE	1111	62.6	23.2	3.055	0.058	1.711	542	3.71	0.890	-0.210	-0.50
TIMAR	TRP	1158	67.7	27.5	2.773	0.098	—	531	1.44	0.516	-0.584	-1.21
AGMA	ASS	1067	74.0	25.4	0.627	0.003	0.902	672	1.65	0.536	-0.564	-2.42
MAGHREBAIL	SFA	1001	75.6	24.8	1.932	0.036	2.197	556	2.05	0.688	-0.412	-1.37
ZELLIDIA S.A	SPH	997	75.7	16.9	8.417	0.001	—	506	2.54	0.159	-0.941	-1.70
MAROC LEASING	SFA	979	76.1	23.5	4.597	0.136	1.157	505	0.69	-0.026	-1.126	-3.50
CENTRALE DANONE	AGR	706	77.3	29.4	0.958	0.038	—	399	1.37	0.190	-0.910	-1.64
OULMES	BOI	752	81.6	21.4	3.921	0.004	—	489	0.52	0.025	-1.075	-2.41
UNIMER	AGR	702	82.9	35.5	7.101	0.086	1.777	419	0.93	0.551	-0.549	-1.67

Note. Sector codes follow the exchange’s own classification. “Trades” is the number of sessions on which the security traded over 2010–2026; “No-tr.” the share of sessions in its listed span on which it did not; “Zero” the share of its one-session returns equal to zero. Amihud is $|r_t|/\text{dirham turnover} \times 10^6$. “Ep.” is the number of non-trading episodes; “Recov.” the mean five-session return following episodes preceded by an upward market; “Net” that quantity less the 1.10% published fee floor, and the last column its t -statistic.

B Panel Construction

Section 3 states the rules the panel is built on and the checks that license them. This appendix records the counts behind those rules.

Price bases. Each commercial export is classified against the exchange record by the median ratio of the two over their common sessions. Of the 155 series so classified, 96 are raw, 4 are restated, and 55 have too little overlap with the exchange record to be classified; the last group is used only where a classified series does not reach. One security is carried on both bases at once, which is what makes the classification worth doing rather than assuming a common basis: pooled, its ratio is bimodal and identifies no basis at all.

Corporate actions. The panel carries 14 dated actions over the covered period, each declared from a step in the point-in-time share count. Of these, 10 fall between 2% and 50%, the range in which no price-ratio test can separate an action from a price move.

A share count rises for two reasons that look identical in a capitalisation series and are opposite in their effect on price: a split or a bonus issue divides the price by the factor, while a capital increase placed for cash leaves the price where it was. Adjusting a return series for the second kind manufactures a move that never happened. Each action is therefore tested against the quote either side of it, and only those the quote responded to restate the series.

Turnover reconciliation. Across the 45 securities that carry both a commercial share count and the exchange's own, the two agree at a median ratio of 1.0006, and the value implied by the commercial count lands within 5% of the exchange's stated traded value for 44 of them. The commercial count is reported to a resolution of 10 shares, immaterial on a line trading in thousands and coarse on one trading in single figures; the single security outside that band is the thinnest in the sample and trades 5 to 10 shares on the sessions concerned.

Reference-price resets. Of the 152,613 consecutive-session returns the panel carries before the universe rule of Table 1 is applied, 72 exceed the exchange's 10% variation threshold, 0.05% of the total; 71 of the 72 fall on securities the rule retains, which is the count Section 6 drops. They are flagged and retained in the panel itself.

Dividend detachments. The panel carries price returns, and a detachment can reset the reference price outright. Of the moves just counted, 25 are falls, and 15 of those occur between May and July, the general-meeting season, which is the seasonal signature a detachment leaves. Since a detachment lowers the measured recovery, removing these moves can only widen the deficits the paper reports.

Coverage ends. Coverage runs to the sample end for 64 of the 73 securities in the study. Of the 9 that stop earlier, 5 ceased to quote inside the sample (a buyout, an absorption, a mandatory withdrawal offer and two delistings) and 4 are exports pulled between 2 and 21 sessions before the sample closes. The listing event is recorded from the public record where one exists and left blank where none does.

C Replication

The replication package contains the database, the code in execution order, and a verification program.

Database. The panel is distributed as an SQLite file. Its tables are the study universe with sector, identifier and coverage; the daily panel of 247,319 security-sessions with prices on both the restated and the quoted basis, trade status, session-gap-aware returns and microstructure; the exchange session calendar; the market factor and the index; and the published fee schedule. Every number in the paper is computed from this file and from nothing else.

Code. Two layers, run in order. The data layer is eight programs: the two sources are loaded and each commercial export classified onto a price basis; the session calendar and the corporate-action table are derived; the price panel is assembled, carrying the quoted and the restated price as separate columns; trade status, carried levels, returns and turnover are attached; the universe rule is applied once and the securities it removes are listed with their reason; the market index is aligned to the calendar and cross-checked; the study extract is written; and a validation suite of 41 assertions is run against the result, which exits non-zero and stops the build on any failure. The analysis layer is seven programs: the illiquidity and cost measures of Section 4, the post-thaw test of Section 5, the robustness of Section 6, the figures, this appendix table, the quantity ledger described next, and the verification program.

Verification. Two failures are possible and they are different: the pipeline can produce a number the paper does not state, and the paper can state a number the pipeline never produced. The verification program checks both. Every quantity the paper asserts is computed in one place and written to a ledger; the program recomputes each of them from the result files and compares. It then reads every numeric literal out of the manuscript source and matches it against that ledger, so that a figure typed into the text which no computation produced is reported rather than overlooked. The published fee schedule is not exempted from this: every rate in Table 3, every round-trip figure and every cost level of Table 6 is read from the fee table in the database and from the cost panel, so the paper's cost anchors are checked like any other result. What remains exempt, being horizons, sample thresholds and calendar years, is enumerated explicitly with the reason each is there, so the exemption list is an inventory of what is not a result rather than a way of silencing the check. The per-security table of Appendix A is generated rather than typed, so it is checked differently: its rows are rebuilt from the result files and compared line by line, which catches a table left stale by a pipeline re-run. The program exits non-zero unless every check passes.

Data provenance. The panel is built from two sources: a commercial daily series covering 2010–2026 and the exchange’s own daily record for the most recent three years. An explicit file registry names every input file and the security it carries, and the loader raises on any file the registry does not name, so a new or renamed file stops the build rather than disappearing from the panel. Appendix B reports the construction counts; the raw files are deposited as stated under Data availability.

Repository. Code and the derived study database: https://github.com/afchtal-ahmedquant/cse-microstructure_paper-working. The exchange’s own daily record: <https://doi.org/10.5281/zenodo.22712772>.