Financial Conduct Authority

Research Note

19th November 2024

A revision of our market cleanliness statistic methodology

Chief Economist Department

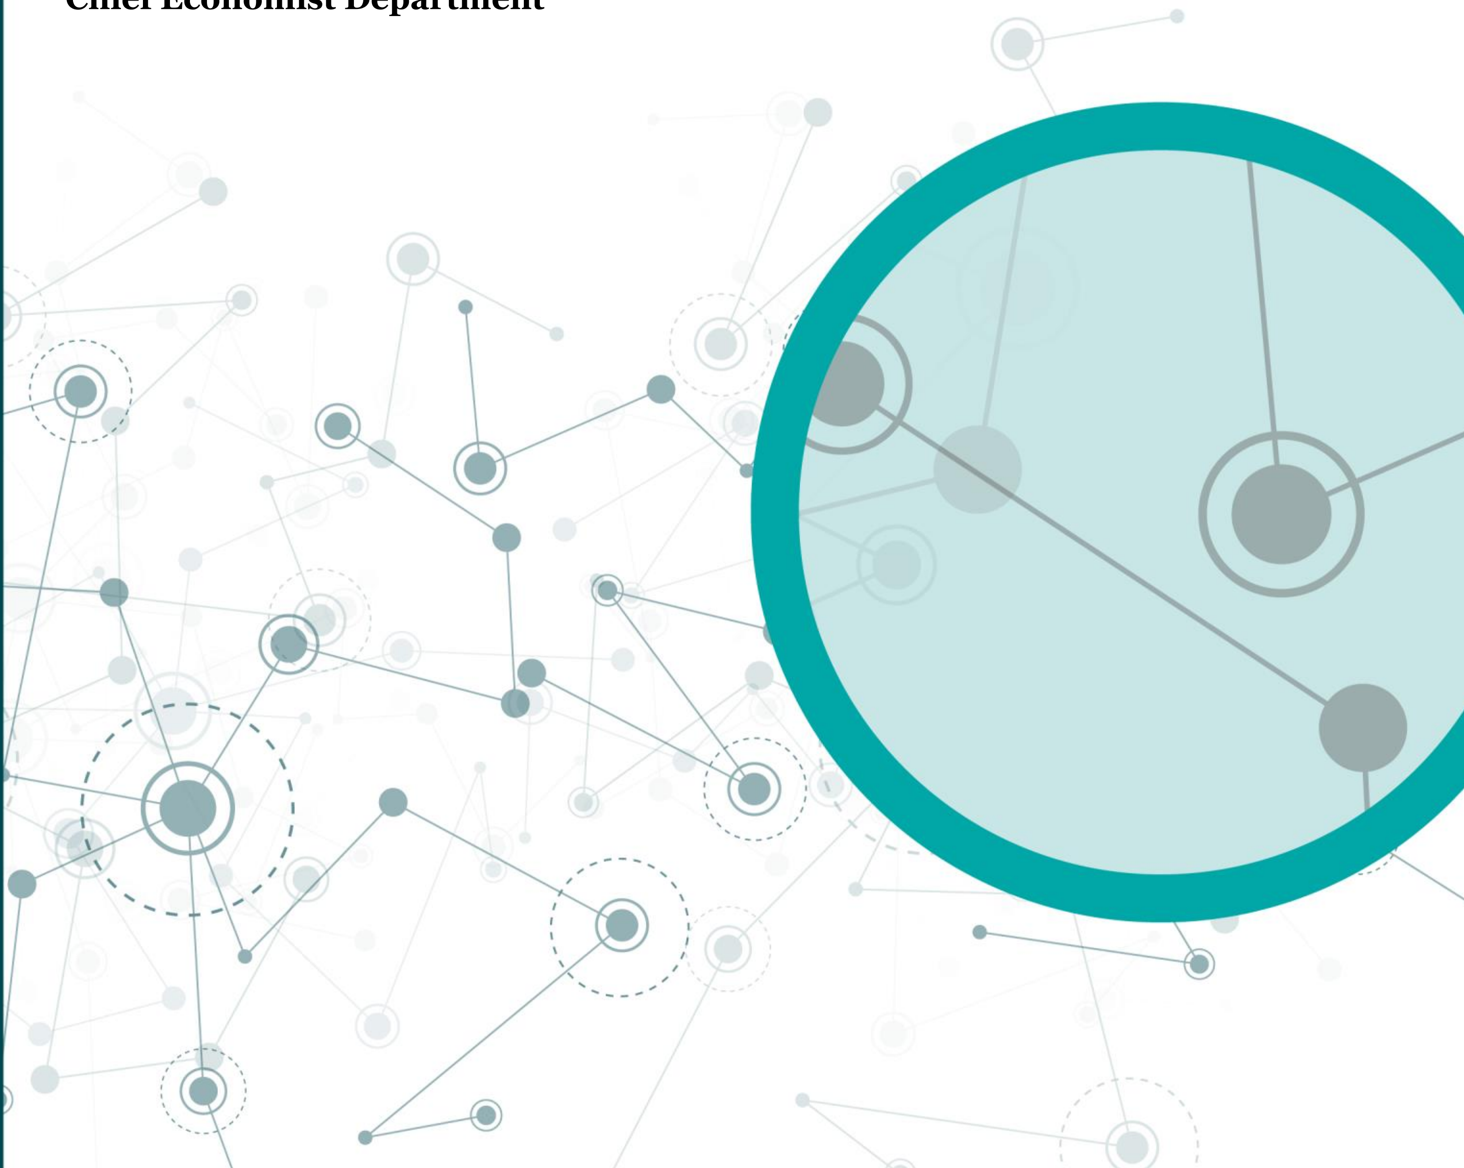

FCA research notes in financial regulation

The FCA research notes

The FCA is committed to encouraging debate on all aspects of financial regulation and to creating rigorous evidence to support our decision-making. To facilitate this, we publish a series of Research Notes, extending across economics and other disciplines.

The main factor in accepting papers is that they should make substantial contributions to knowledge and understanding of financial regulation. If you want to contribute to this series or comment on these papers, please contact Sam Hainsworth (sam.hainsworth@fca.org.uk) or David Stallibrass (david.stallibrass@fca.org.uk).

Disclaimer

Research notes contribute to the work of the FCA by providing rigorous research results and stimulating debate. While they may not necessarily represent the position of the FCA, they are one source of evidence that the FCA may use while discharging its functions and to inform its views. The FCA endeavours to ensure that research outputs are correct, through checks including independent referee reports, but the nature of such research and choice of research methods is a matter for the authors using their expert judgement. To the extent that research notes contain any errors or omissions, they should be attributed to the individual authors, rather than to the FCA.

Authors

Khashayar Rahimi (Technical Specialist), Sam Hainsworth (Lead Associate), Silvia Lozano Guerrero (Senior Associate), Guidogiorgio Bodrato (Senior Associate), Larissa Emde (3rd party, Rozetta Institute and Macquarie University).

Acknowledgements

We would like to thank our academic supervisors for their guidance and support, Gustavo Fruet-Dias (University of East Anglia), Talis Putnins and Vinay Patel (both University of Technology Sydney). We would also like to thank Paul Stocks and Dan O'Donnell (FCA, Enforcement and Market Oversight) and Kieran Keohane (FCA, Economics) for their input and review.

All our publications are available to download from www.fca.org.uk. If you would like to receive this paper in an alternative format, please call 020 7066 9644 or email publications_graphics@fca.org.uk or write to Editorial and Digital Department, Financial Conduct Authority, 12 Endeavour Square, London E20 1JN.

Contents

Executive Summary

The FCA regularly monitors and reports on market cleanliness to contribute to our objective of preventing and reducing harm and strengthening the integrity of UK financial markets. At present, we report three different measures of market cleanliness as indicators of the level of information leakage and potential insider trading in UK equity markets: the Market Cleanliness Statistic (MCS), the Abnormal Trading Volume (ATV) and the Potentially Anomalous Trading Ratio (PATR). Each of these indicators are indirect measures of insider trading. Care is needed when interpreting them, particularly when looking at any statistic in isolation.

The MCS looks for abnormal share price movements before takeover announcements. The FCA has reported the MCS since 2008 in its annual report. The ATV focuses on abnormal trading volumes and the PATR looks at potentially anomalous trading in suspicious accounts. The FCA introduced the ATV and PATR in 2018 to broaden the number of indicators used to assess market cleanliness.

As part of our ongoing effort to improve our performance metrics, in this paper we review the methodology for calculating the Market Cleanliness Statistic (MCS). The FCA last reviewed the methodology in 2014.

The current MCS does not detect intraday abnormal price movements, and it excludes announcements by firms that have been subject to multiple recent takeover offers. We address these limitations using intraday data. In addition, the current MCS may be biased when takeover events coincide with periods of high market volatility. In the new methodology, we have redesigned the statistical test so that it controls for market volatility when determining abnormal price movements.

The new methodology leads to a statistic that is robust to periods of heightened market volatility and incorporates intraday trading activity. The revised measure is higher, reflecting the scope of the statistic now including potential insider trading on the day of an announcement. The change in the statistic is not an indication of a deterioration in market cleanliness.

Overview

The FCA has a statutory operational objective to protect and strengthen the integrity of the financial system in the UK. This objective requires the FCA to ensure that markets function well and are clean (FCA, 2013). Market cleanliness refers to the absence of activities such as insider trading, market manipulation or unlawful disclosure of insider information, implying a transparent trading environment where all participants have equal access to information and opportunities.

As part of our continuous improvement, we have reviewed the headline measure (the Market Cleanliness Statistic), which aims to provide an indication of the extent to which share prices move before takeover announcements for UK traded equities.

We propose changes to improve the coverage and robustness of our statistic. With these adjustments we aim to ensure that the statistic supports the FCA in setting high standards of transparency, reducing the risk of market manipulation, and promoting healthy competition.

Purpose

The Market Cleanliness Statistic (MCS) looks at abnormal returns prior to takeover announcements, which can reflect a potential leakage of information. For each announcement, the current methodology estimates the stocks' cumulative abnormal returns (CAR) in the two days before the announcement compared to a period without an announcement. We calculate the MCS as the share of announcements (in % terms) for which the event window CAR are higher than the 90th percentile of the distribution from the comparison period. We calculate and publish this statistic in the FCA annual report.

The current methodology for the statistic relies on the approach developed in Dubow and Monteiro (2006) and in Monteiro et al. (2007). The methodology was assessed and consolidated by Goldman et al. (2014). In this review, we propose improvements to the known limitations of the methodology:

1. **Currently, the MCS does not capture potential insider trading activity on the day of the announcement.** We calculate the statistic using end-of-day prices up to the day before the announcement, excluding trading on the same day of an announcement that occurs within market hours. This potentially underestimates the number of firms showing Abnormal Pre-announcement Price Movements (APPMs). For reference, in 2023, over a third of all takeover announcements (35%) happened during market hours and may have information leakages on the same day, that aren't captured in the current methodology.

2. **Firms with multiple recent takeover announcements are dropped from the sample.** Under the current methodology, the statistic uses data from the prior year up to ten days before each announcement for the estimation window (i.e., comparison period covering 240 full trading days). We exclude

announcements from the sample where the firm was subject to a takeover offer within the estimation window. The rationale behind this exclusion is that a previous takeover offer can contaminate the distribution of cumulative abnormal returns in the estimation window, biasing the result of our statistical test. Consequently, a reduced number of events are considered: in 2023, we removed almost 8% of takeover announcements from the sample with the current methodology.

3. **MCS results may be biased if a takeover event coincides with periods of high market volatility.** Events such as the COVID-19 pandemic or the Russian invasion of Ukraine can influence the headline statistic. Since the current methodology does not control for market volatility, these macroeconomic events can cause a stock's price to move abnormally before an announcement. This is not necessarily due to the leakage of firm' specific news. Heightened volatility can generate false positives or false negatives (e.g., if a takeover announcement falls close to a macroeconomic shock that affects the event window or estimation window CAR). This means that the statistic under the current methodology can be biased.

Key improvements

To address the three issues presented above, we update the current MCS methodology (Goldman et al., 2014) as follows:

- We move from end-of-day prices to intraday prices. This means we **include pre-announcement price movements on the same day as a takeover offer during market hours**.
- We shorten our estimation window from 240 days (using daily price observations) to 60 days (using 5-minute interval price observations). This allows us to **consider firms with more than one takeover offers in a year**, provided that any two announcements are at least one quarter apart.
- We replace the current methodology to **include a cross-sectional Market Comparison test to account for market volatility near the time of an announcement**. The intuition behind this test is that when many stock prices in the market move abnormally, we are unable to attribute abnormal returns to firm-specific information leakages. Under this new methodology, when market movements are sufficiently large, Abnormal Pre-announcement Price Movements (APPMs) are not considered to be due to leaked information.

Key findings

From implementing the changes listed above, we find:

- **We detect more APPMs using intraday price data**. This pushes the revised statistic upwards, especially in years with a higher proportion of announcements during market trading hours.
- **We exclude fewer announcements from the calculations of the statistic** by shortening the estimation window**.** From 2020 to 2023, 5 (2%) additional announcements are included compared to the existing methodology, helping to improving the representativeness of the statistic.

- **By implementing a cross-sectional Market Comparison test, we control for the risk of potential false positives and false negatives** in the statistic. Over the sample period (2020-2023), this slightly reduces the volatility of the statistic compared to the current methodology. However, we need more years of analysis to confirm if this trend persists.

Figure 1: Market Cleanliness Statistic (Annual % of APPM detected)

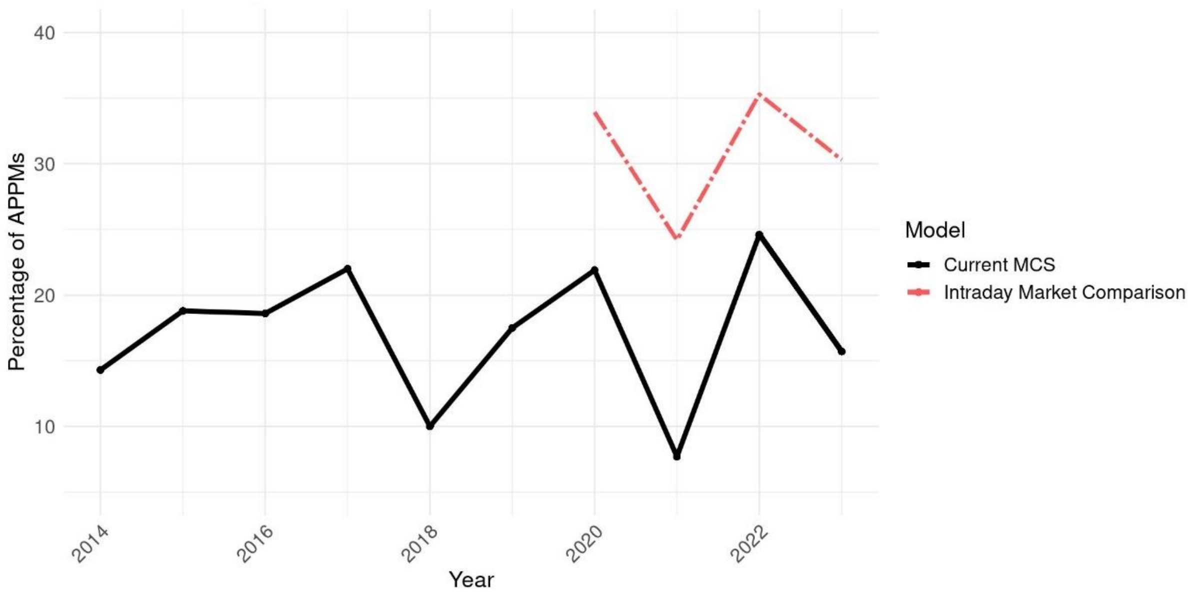

Figure 1 shows the MCS based on the new methodology (hereafter Market Comparison) using intraday price observations (red line) compared to the current methodology from Goldman et al. (2014) which uses end-of-day price observations (black line).

The resulting Market Comparison MCS is higher than the current methodology. This does not mean that UK markets have become less clean. The new Market Cleanliness Statistic has a different scope to previous estimates; therefore, the numbers aren't directly comparable. The increase is mostly driven by the introduction of intraday data, with a small smoothing effect from the Market Comparison test. The impacts of the methodological changes are explained in detail in the Results section (and in Annex 2).

Next steps

From 2024, the FCA will implement the Market Comparison method using intraday data to estimate the Market Cleanliness Statistic. We welcome feedback on our methodology and findings.

Research context

Market Cleanliness

Since 2008, the FCA's (then FSA) annual report included a measure of market cleanliness as an indicator of the level of potential insider trading in UK equity markets. Market cleanliness refers to the absence of fraudulent activities, such as insider trading or market manipulation, implying a transparent trading environment where all participants have equal access to information and trading opportunities. The Market Cleanliness Statistic (MCS hereafter) was first developed in Dubow and Monteiro (2006) and in Monteiro et al. (2007), and then revised by Goldman et al. (2014). The MCS analyses whether price movements before a takeover announcement reflect a potential information leakage.

The FCA is not the only regulator that publishes this statistic yearly. The Australian Securities and Investments Commission (ASIC) calculates a similar market cleanliness statistic for the Australian equity market (ASIC, 2016). Aside from regulators, academics have also considered the problem of insider trading in equity markets such as Patel and Putnins (2020) for the United States.

The MCS measures the percentage of takeover announcements where we observe abnormal returns immediately prior to such announcements. The assumption is that, in the absence of new information, we should not observe significant abnormal price movements prior to an unanticipated announcement. If instead, we do observe abnormal price movements, this can reflect a potential leakage of information and trading based on inside information.

As a result, the MCS focuses on Cumulative Abnormal Returns (CAR) for a period immediately prior to the announcement, which is referred to as the "event window". This is compared to the distribution of CAR calculated using data from the "estimation window". Figure 2 shows an illustration of the concept from Goldman et al. (2014). The more we observe price movements just before takeover announcements, the more a market is considered unclean.

In addition to the MCS, the FCA produces two further measures, namely the "potential anomalous trading ratio" (PATR) and "abnormal trading volume" (ATV). The PATR looks at the characteristics of accounts' trading activity to determine which share of the total trading value comes from accounts demonstrating anomalous behaviour with respect to their past. The ATV detects abnormal increases in trading volumes just before an unanticipated announcement. Both indicators look at any price sensitive announcements, and not just takeover offer announcements.

Figure 2: Example of clean versus unclean market

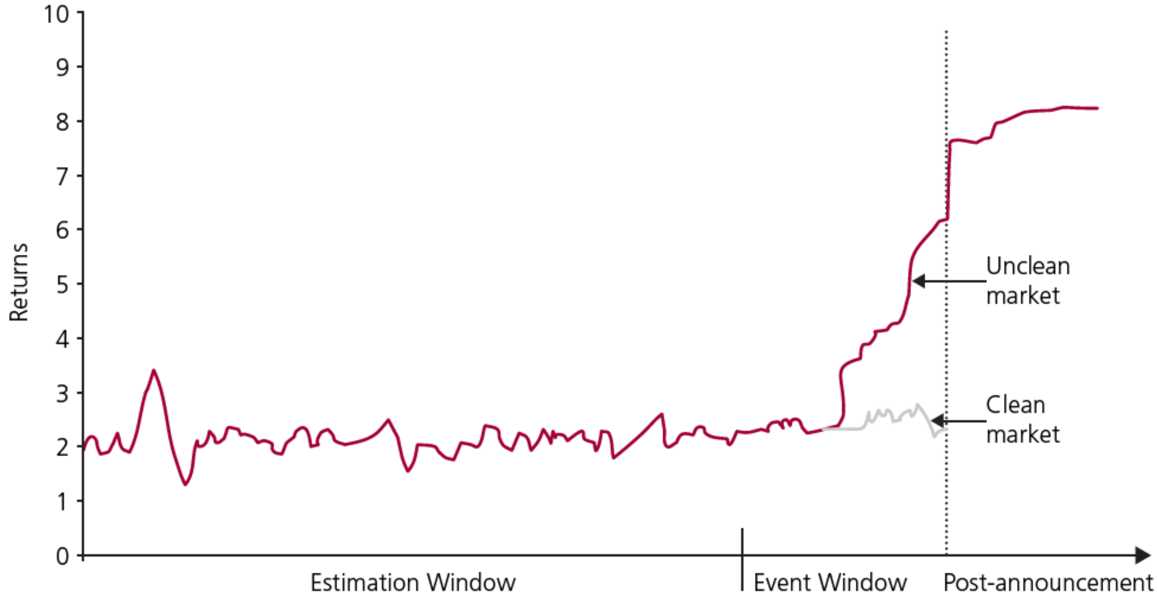

Source: Goldman et al. (2014)

The FCA Market Cleanliness Statistics

The MCS was first introduced by Dubow and Monteiro (2006) in the Occasional Paper "Measuring market cleanliness" to measure the level of potential insider trading in the UK's equity market by analysing (significant) abnormal share price movements prior to regulatory and takeover announcements.

Dubow and Monteiro (2006) define the MCS as the ratio of detected Informed Price Movements (their term for Abnormal Pre-announcement Price Movements) to the total number of significant announcements. Significant announcements are identified as those where we observe large price movements both before and after the announcement.

Monteiro et al. (2007) further improve the MCS in a following Occasional Paper "Updated measurement of Market Cleanliness". To estimate expected returns for regulatory announcements, they apply an extended market model and normalise abnormal returns to account for issues such as serial correlation (i.e., when returns on nearby days are not independent, resulting in correlated errors terms over time) and heteroskedasticity (i.e., when the variance of abnormal returns is not constant over time). These issues can lead to misinterpretation of the significance of the cumulative abnormal returns before the event, and result in an inaccurate assessment of announcements. They normalise abnormal returns by the conditional variance over the estimation window and they address serial correlation by including lagged variables to the model. When neither serial correlation nor heteroskedasticity are detected, authors estimate a one-factor market model (as in Dubow and Monteiro, 2006). For takeover announcements, the authors use a simple mean model to estimate expected returns. See Annex 1 for specifications on these models, which are used in our analysis for robustness checks.

Goldman et al. (2014) investigate the reasons behind a decline in the market cleanliness statistic from 2009 to 2013, falling from ~30% to ~15%. Their study focuses exclusively on takeover announcements in the calculation of the market cleanliness statistic. They use a simple mean model to calculate the expected return. As a robustness check, they repeat the MCS calculation using the Capital Asset Pricing Model (CAPM) instead of a simple mean model (i.e., where expected returns are proxied by the mean return over the estimation window) to show that the difference in the results is statistically non-significant. They point out, however, that "particularly in times of high stock market volatility, there may be cases where market movements could be driving our statistic."

The authors exclude methodological choices as a reason for this decrease in market cleanliness. They control for different length event windows, confidence levels for the significance of cumulative abnormal returns, and apply alternative models to compute the expected returns. The findings support the perception of cleaner markets, suggesting this might be due the increased regulatory actions immediately after the Global Financial Crisis.

The ASIC Market Cleanliness Statistics

Like the FCA, the Australian Securities and Investments Commission measures market cleanliness in the Australian equity market as a signal of "possible insider trading and information leakage ahead of material, price-sensitive announcements" (or MPSAs). ASIC (2016) calculates the measure as the percentage of statistically significant abnormal pre-announcement price movements (APPMs) out of the total number of MPSAs.

To estimate the expected return, ASIC (2016) use the Capital Asset Pricing Model (CAPM). A significance level of 1% is set for the event window CAR "to be large enough that the probability they were driven by random volatility in the price of the security was extremely low". Like Dubow and Monteiro (2006) and Monteiro et al (2007), the MCS methodology from ASIC (2016) accounts for the direction of APPMs.

ASIC (2016) set out further requirements for their sample of announcements which include "announcement proximity", that excludes MPSAs in the 10 trading days after another MPSA. The authors also control for factors that can move the stock's price e.g., volatility, liquidity, or size of the issuer.

ASIC (2016) analyse price-sensitive announcements over the period from November 2005 to October 2015 ("five years before and five years after the transfer of market supervision to ASIC") and find an improvement in market cleanliness in the Australian equity market over this period.

Limitations to the current methodology

In this paper, we present a methodology which addresses three different potential issues with the FCA's current methodology to calculate Market Cleanliness Statistic:

- **Potential insider trading activity on the same day of the announcement is not captured.** Following Monteiro et al. (2007), the FCA uses a 2-day event window to test whether a significant price movement occurred before an announcement. Using end-of-day prices up to the day before, the current methodology includes announcements happening during market hours but does not consider any trading on the same day of the announcement. Someone reacting to leaked information hours before the announcement can still affect the price of a stock. This could lead us to underestimate the number of abnormal pre-announcement price movements (APPMs). For reference, in 2023 35% of announcements happened during trading hours.

- **Firms with multiple recent takeover announcements are dropped from the sample.** Under the current methodology, when a firm had already made an announcement within the estimation window, the firm's next announcement is excluded from the calculation of the MCS. This is because the second announcement's estimation window, the period ranging from 250 days to 11 days before the event, is likely to be affected by the first announcement. The requirement of one year of unimpacted data in the estimation window reduces the sample of announcements used to calculate the statistic. For reference, in 2023 6 takeover announcements (close to 8% of the total in that year) were excluded from the sample for this reason.
- **MCS results may be biased if a takeover event coincides with periods of high market volatility**. Goldman et al. (2014) use robustness checks for assessing the impact of market volatility (i.e., they estimate the CAPM model for obtaining measures of expected returns) and find no significant impact on the MCS. In recent years, we observe several episodes of high volatility in stock markets due to macroeconomic events such as the COVID-19 pandemic or the Russian invasion of Ukraine, that could have impacted the MCS. The MCS itself has also become more volatile in recent years. A price move before a takeover announcement can be identified as abnormal during these stress periods due to high market volatility, rather than information leakage and insider trading. Similarly, abnormal price movements before an announcement are possibly not detected if the event is after a period of high market volatility occurring during the estimation window. Therefore, not controlling for market volatility can result in overestimation or underestimation of the MCS.

Research design

To improve the current methodology, we reviewed the literature on price movement detection and high-frequency event studies. We conducted an extensive review of the literature and engaged with academics to understand which strands are the most relevant for our revision.

We propose to move from using end-of-day daily price observations to intraday 5-minute interval price observations, and to employ a new methodology to detect potential APPMs. Using higher frequency data, we can consider price movements on the day of the announcement. We can also include firms with multiple takeover offer announcements in the same year, as long as the respective announcements are at least one quarter apart. Our new method for detecting APPMs accounts for the impact of market volatility, which reduces the risk of potential false negatives and false positives.

In the following sections we describe the data that we use, the current methodology and the changes we propose.

Data

Our analysis considers 302 takeover announcements ranging from 2020 to 2023 based on a list of takeover announcements provided by the Market Oversight team in the FCA. We consider announcements that occur before the market opens, after it closes, and during market trading hours. Announcements where the company is delisted, not traded in the UK market or where the price information is not available are excluded from our analysis. After this selection step, our sample includes 285 announcements. Figure 3 below presents the number of announcements for each quarter, grouped by the time of the day.

Figure 3: Number of announcements per quarter by timing of announcement

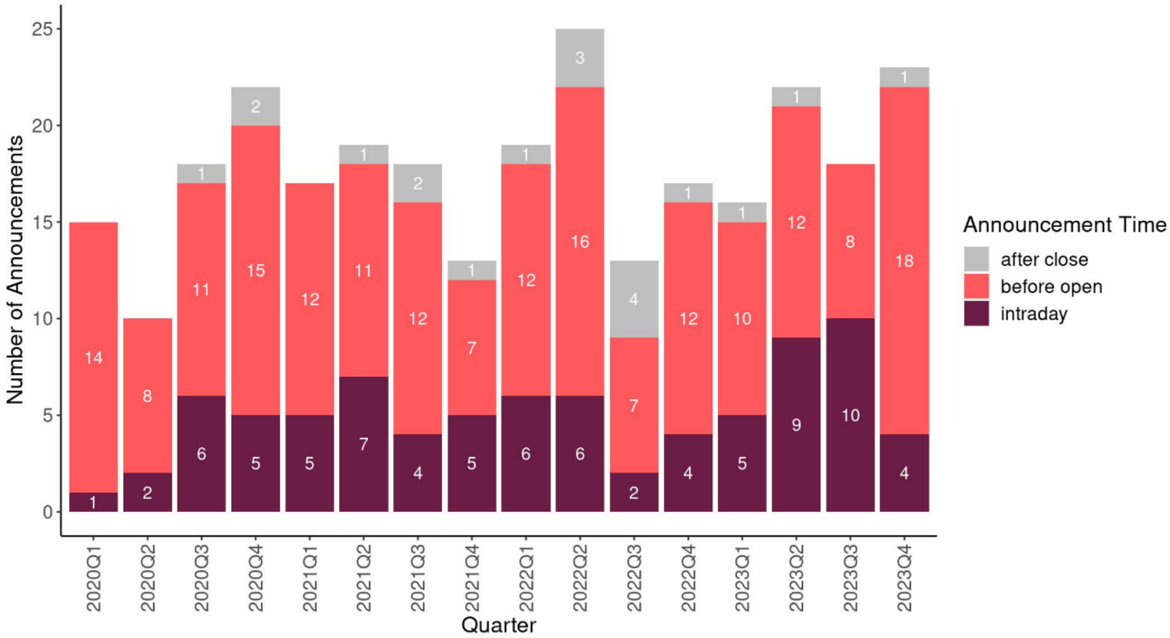

We collect daily end-of-day prices from Refinitiv Eikon. Intraday prices (i.e., the last trade price) are sourced from Datascope Tick History Summaries data and they are aggregated at 5-minute intervals. Given contractual constraints and data coverage, Datascope was selected over other data providers. For each announcement, we extract intraday price data for the 70 days preceding the announcement.

The use of intraday prices has two main benefits:

- We perform the analysis over a larger number of datapoints,
- We observe price movements up to 5 minutes before an announcement, which improves our detection of abnormal price movements on the same day as an announcement.

One caveat of using 5-minute interval observations is that not all the stocks are liquid and traded frequently. For illiquid stocks, there is the possibility that within a 5-minute interval we observe no trades. For an announcement to be included in the statistic, we require at least 100 price observations (the same requirement that we were applying in our current method) spread over at least 20 different trading days in the estimation window. This ensures that expected returns are not estimated based on a concentrated set of information. When no information exists for a given 5-minute interval, we carry the last available price forward (in line with Refinitiv process for daily data). If less than 100 observations are available or the stock was not traded for at least 20 out of the 60 days of the estimation window, we remove the announcement from the calculations of the statistic. As a result, 26 announcements were excluded (9% of the sample) from the statistic. These conditions are introduced to ensure that a minimum number of datapoints are available to inform the analysis.

Figure 4 displays the percentage of zero return observations in the estimation window for each decile of the sample. Each decile represents 10% of the firms in each sample, ordered based on the total number of zero returns within the estimation window. Using intraday data means there are more occurrences of no change in price between data points, generating zero returns. However, the number of data points differs substantially between daily and intraday data and therefore a higher percentage of zero returns with intraday data is not necessarily problematic.

These patterns limit the models we can use to estimate expected returns, as explained in detail in the limitations section.

Figure 4: Incidence of zero returns (daily vs intraday data)

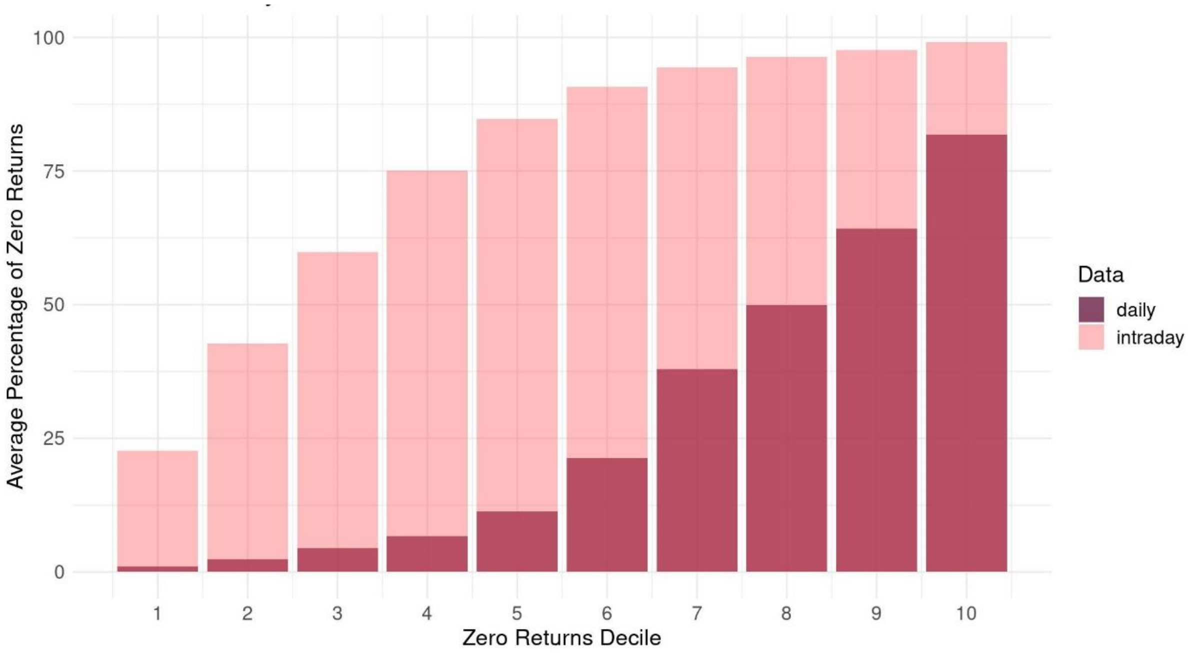

Note: Deciles divide each sample into 10 equal-sized groups, in order of the proportion of zero-returns. Decile 1 contains the lowest 10% and Decile 10 the highest 10%.

If there are no trades in the event window, the APPM is set to 0, which signifies that significant abnormal returns are not observed for this stock (refer to "Current methodology" and "Proposed new methodology" sections for more details). This correction applies to only 3 announcements in the sample.

Over 65% of the stocks in our sample are small capitalisation firms (i.e., with a market value of less than £200 million). This group has the largest proportion of stocks not traded each day in their estimation window. For more information about the composition of the sample (i.e., including breakdown between sectors, stock's market capitalisation and the number of days traded), please refer to Annex 3.

Current methodology

Potential insider trading can be proxied by a statistically significant price increase immediately prior to a takeover announcement. We define abnormal returns as the difference between observed returns and expected returns at each point of time t, as set out in the equation below:

$$ABR_t = r_t - E(r_t)$$

Under the current method, to detect abnormal returns, we construct the test statistic (a statistic calculated from sampled data), as the cumulative abnormal return (CAR) over n-days before the event:

$$CAR_{n,t} = \sum_{h=1}^{n} ABR_{t-h}$$

where n is the size of the event window and ABR_t is the abnormal return on day t (when $t = T$ is the event date) calculated as the difference between the actual return and the average return over a clean estimation window (expected return). We evaluated various methods to measure expected returns which have been previously used in the literature. We explain these more in detail in the Robustness checks section and in Annex 1.

We identify price moves that are unlikely to be normal by comparing the test statistic with the distribution of $CAR_{n,t}$ from the estimation window. We perform a one-sided test based on following hypothesis:

$$H_0: CAR_{n,T} \leq CV$$

$$H_I: CAR_{n,T} > CV$$

Where the critical value, CV, is the 90th percentile of the cumulative abnormal returns' distribution over the estimation window; and $CAR_{n,T}$ represents the test statistic. If the test statistic is larger than the CV, then we reject that the cumulative abnormal returns in the n-days before the event are normal. As a result, our abnormal pre-announcement price movement (APPM) takes value as 1 and we consider that an information leakage prior to the announcement could have caused an abnormal spike in the stock's price. Otherwise APPM takes the value 0. Note that this is a one-sided test, and that it will only detect abnormal price movements which cause positive abnormal returns.

The share of total events for which we reject the null hypothesis each year is the final market cleanliness statistic.

Proposed new methodology

The Market Comparison test

With the Market Comparison test, we analyse whether a stock's price is moving significantly more than other stocks' prices in the same period. The intuition behind this test is that abnormal price movements preceding an event should be firm-specific to be considered as potential insider trading. If we look at the cross-section of stocks in the same period and we observe their returns varying significantly more or less than their recent average, then we cannot tell if movements in stock prices are due to firm-specific or market volatility. Market volatility can be problematic because it can inflate abnormal returns for a given stock. To control for market volatility, we build a Market Comparison test in line with a similar test employed by Comerton-Forde and Putnins (2011 and 2014).

We extract price observations at 5-minute intervals for 500 randomly sampled publicly traded stocks in the FTSE All Share market index. We use the 500 randomly selected constituents of FTSE ALL Share as a representative comparison sample because it includes firms with different market capitalisations, and it is representative of the stocks which are the subject of our analysis. To implement the Market Comparison test, we take the following steps:

- Let K be the number of stocks in our comparison sample, including our stock of interest. For each k in $1,2,\ldots,K$ we calculate the CAR_k over the event window, as well as its series from the estimation window.

- Then, we rank each stock's CAR_k with respect to its own estimation window's distribution to obtain the percentile value P_k, i.e., each stock is given a percentile score value based on its CAR_k compared to its sample of cumulative abnormal returns from the estimation window. This gives us a normalised score that can be compared across all stocks.
- After this normalisation step, we rank the percentile scores from all comparison stocks and the stock of interest. Then we assign a second percentile score, S_k to each P_k based on its cross-sectional distribution against all other comparison stocks and the stock of interest. This can be considered as the percentile of percentiles.

Let j be the stock of interest from the takeover sample, then our Market Comparison test hypothesis is constructed as:

$$\mathrm{H_0: S_j \leq 90}$$

$$\mathrm{H_A: S_j > 90}$$

If we fail to reject the null hypothesis, then we cannot tell apart firm-specific movements before an announcement from market volatility and so we restrict the APPM value to 0 for the event. In other words, we consider APPMs (equal to 1) only for those stocks with an associated score in the top decile of the distribution of P_k across all stocks.

Comparing stock price movements during the event window with other stocks in the market, we account for market price movements. Under the current methodology, higher price variation within the estimation window results in a higher distribution of CAR_k and our may mean the test fails to detect potential APPMs. Similarly, if there is higher price variation within the event window due to market volatility, the existing methodology does not differentiate between this and potential insider trading. Hence, we can observe both false negatives and false positives.

The Market Comparison test overcomes these issues in two steps:

1. for each stock in the market, it ranks price movements during the event window compared to its own history. If a market is experiencing a period of high volatility during the estimation window, a considerable number of stocks should be given a low score for their price movements within the event window[1];
2. in the cross-section of these normalised scores, $\mathrm{S_j}$ still can still be relatively high (or low) in comparison to other stocks. However, we should be able to determine whether for stock j, price moves are significantly higher than other stocks.

Estimation and Event window lengths

With intraday data, we use a 60-day estimation window, from 70 days before the event date up to 10 days before the event date (-70, -11, both included). We now calculate the cumulative abnormal returns in the estimation window for a 48-hours window (i.e., equal to the length of the event window) at every 5-minute interval to estimate our distribution. Similarly, the event window at intraday runs for a fixed period 48 hours preceding the announcement, from 5 minutes before the announcement is made.

[1] Similarly, if during the event window, market is experiencing a higher price variation, then considerable number of stocks should get a high score for their $CAR_{n,T}$

Alternatives considered

A hybrid approach

Our proposed new methodology uses intraday data at 5-minute intervals. Other options were considered, such as using a combination of daily and intraday price data. In this approach, we estimated cumulative abnormal returns using end-of-day prices data for events when the announcement takes place before or after market hours, and intraday price data for those events when the announcement happens during market hours. However, this approach was ruled out as it requires potentially inconsistent datasets and estimation windows, with only marginal computational improvements.

Different expected return estimators

For the estimation of expected returns, we tested different models and specifications (including those used by previous authors) and discussed the results with our academic reviewers. These are discussed in more detail in the Robustness Checks section and Annex 1.

Results and limitations

Results

In the following section, we discuss the key results from reviewing the current methodology and implementing our new approach. We discuss how this methodological change impacts the Market Cleanliness Statistic and explain the motivation for this change.

The current methodology vs Market Comparison using daily data

To check the performance of our new method, we use daily data to compare the results with the current methodology. Figure 5 shows the annual MCS estimated by the current methodology and the Market Comparison test (dashed line), both with daily data. The Market Comparison test results in a higher MCS but follows the same trend as the current method, meaning that we still have higher MCS in 2020 and 2022.

In the current methodology, when there is a period of high volatility during the estimation window, price variations in the event window are less likely to be detected as abnormally high and the model is less likely to point to evidence of insider trading. However, when using the Market Comparison test, we may still find abnormally high price movements in the event window. We label cases as "false negative" if they were not previously flagged as APPM in the current methodology but are now flagged as APPM in the Market Comparison model. Similarly, "false positive" cases can happen when there is a high market volatility period during the event window. In these situations, price variations over the event window are detected to be abnormally high (and so using the current methodology we reject the null of no insider trading), but we may see that these variations are not abnormally high when comparing with other stocks in the market.

Table 1 outlines how the new method (i.e., using the Market Comparison test) compares to the current methodology. The upper panel of the table shows the number of "false negative" cases. For each year the second column shows the range of percentile values, P_k, of cumulative abnormal return of the event window for the stock of interest. In all "false negative" cases P_k is below 90, which explains why the current methodology fails to detect potential APPMs. The last column shows the range of percentile of percentiles scores when comparing the announcing firm's stock with a market sample of other stocks. A high score demonstrates that our stock of interest had a relatively higher price movement compared to the market.

The lower panel shows similar results for the "false positive" cases, when our current method detects APPM (high P_k), but these price movements were common in the market sample (low score of S_k).

Figure 5 shows the resulting MCS from applying the Market Comparison test with daily data versus the current methodology. The difference in results comes from addressing "false negative" and "false positive" cases, which results in a net increase in the number of APPMs.

Figure 5: The current methodology vs The Market Comparison (daily data)

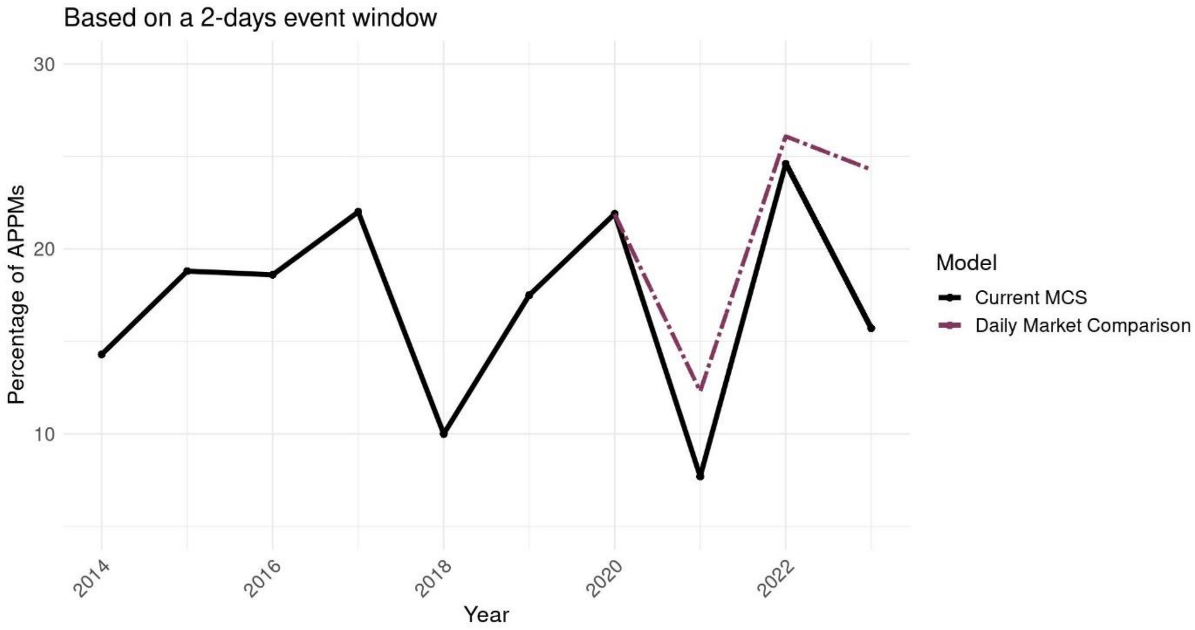

Table 1: False Negatives and False positives (daily, annual)

Year	False Negatives		
	Number of events	P_k range	Score (percentile) range of P_k in market sample, (S_k)
2020	4	66-88	92-95
2021	3	79-89	91-95
2022	6	78-89	94-98
2023	8	67-89	91-96

Year	False Positives		
	Number of events	P_k range	Score (percentile) range of P_k in market sample, (S_k)
2020	-4	94-98	60-89
2021	0	NA	NA
2022	-5	90-99	63-82
2023	-2	98	77-87

Market Comparison using intraday data

Figure 6 shows the results from the Market Comparison test, using intraday data, against the current annual MCS with daily data. The Market Comparison MCS is higher than the current methodology.

Figure 6: The current methodology vs The Market Comparison (intraday data)

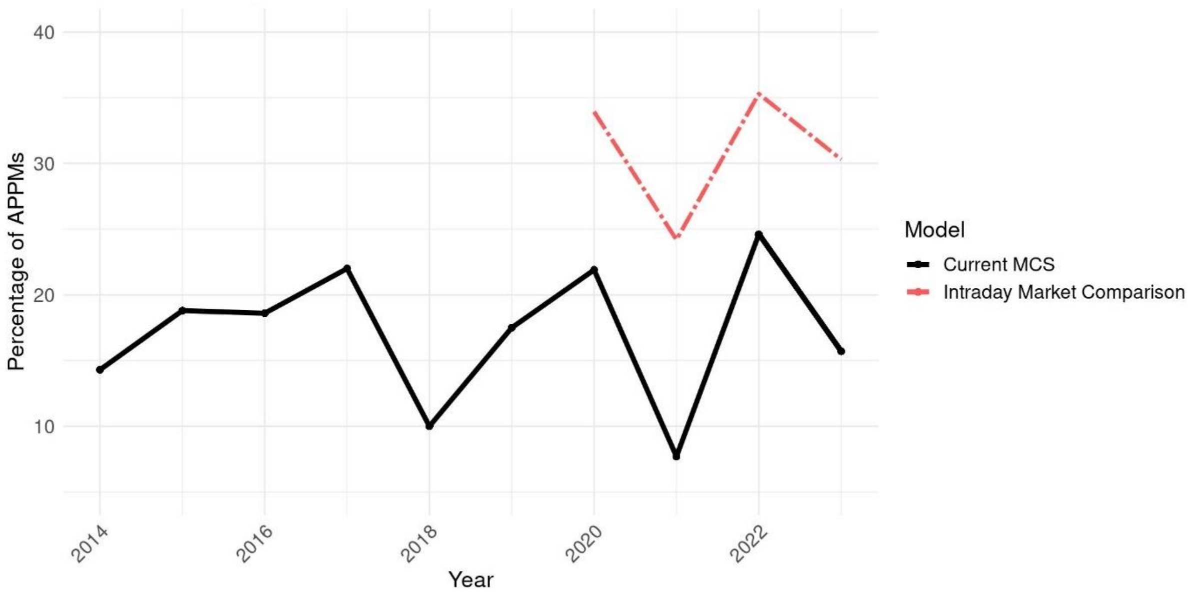

Using intraday stock prices instead of end-of-day prices one intuitively expects an increase in the number of APPMs, because price movements occurring on the day of an announcement during market hours are now included in the event window[2].

The increase in the market cleanliness statistic reflected by the red line in Figure 6 may also be due of the different data sources and estimation window lengths:

- The data used for daily and intraday analysis differs. We move from taking Refinitiv daily end-of-day price observations to taking Datascope 5-minute intervals last traded prices. In our intraday analysis, we only include observations within market hours, excluding the end-of-day auctions. This means that the latest available price each day in Datascope may differ from the end-of-day price used in previous analysis. These discrepancies might contribute to different results.
- The estimation window that we use in the intraday Market Comparison test is now shortened to 60 days. If the stock had idiosyncratically less or more volatile returns during that shortened period, this can also affect the result.

Additionally, implementing the Market Comparison test allows us to control for market volatility as explained in the previous section. This is evidenced in Annex 2, which contains a detailed comparison between the Market Comparison test and the current methodology when using intraday data. Whilst the new statistic is slightly less volatile than the current methodology, it is sensitive given the relatively small sample of takeovers in each year. We need more years of analysis to confirm if this trend persists.

[2] However, not all the events that were identified with APPM by the Market Comparison methodology using intraday data (and which had not been identified with APPM by the current methodology with daily data) are intraday events, as shown in Table 2 in Annex 2.

Robustness checks

Whilst reviewing the methodology of the MCS, we assessed a range of methods and models which account for market volatility. In addition, we tested the robustness of the results to changes in the main assumptions (e.g., length of the event window). A summary of the results from these checks is presented below. More detailed information on how each of these methods is specified and estimated can be found in Annex 1.

We implement the following econometric models in addition to the Goldman et al. (2014) current methodology to calculate expected returns for each event:

1. One-factor Market Model
2. Extended Market Model, including lagged returns
3. Capital Asset Pricing Model (CAPM)
4. CAPM with time-varying beta

All these methods use a measure of market returns to control for wider market volatility. Across these models, we found a very low correlation between our stocks of interest and the market (approximated by the return of the FTSE350).

In addition to these models, we implemented a bootstrap method to approximate a distribution of cumulative abnormal returns to compare with the test statistic and control for potential issues related to serial correlation. More details about how we implemented the bootstrap can be found in the Annex 1. Implementing the bootstrap did not result in any significant changes in the MCS.

Figure 7 shows the current published MCS in black, as before. The grey shaded area around the black line represents the maximum and minimum MCS estimate from fitting the above models plus the current methodology using a bootstrap. This shows that implementing these methods has an inconsequential impact on the MCS.

Figure 7: Alternative methodologies considered (Annual % of APPM)

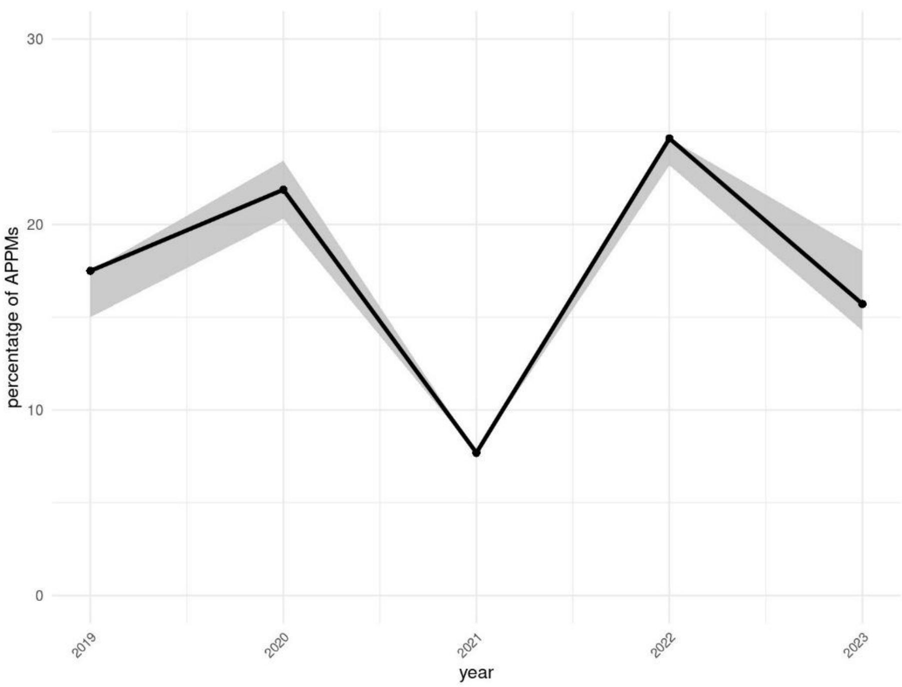

Zero return observations are frequent due to sparsely traded illiquid stocks in our sample. This gives low estimates for the correlation between the stock's returns and the market returns. As a result, expected returns, and consequentially abnormal returns, are not significantly different from those using the current methodology (as already remarked in the first revision of the market cleanliness statistic by Goldman et al.). It can be argued that if a few large capitalisation stocks are unresponsive to market news, a FTSE index will not move significantly even when most public traded stocks do. While a market model works for large and mid-cap stocks, it does not perform equally well for small stocks. Hence, we are replacing the current Goldman et al. (2014) methodology with the Market Comparison test as it works consistently for all stocks.

Additionally, we test the robustness of our results to changes in the length of the event window. We calculate the MCS using both a 3 and 5 day event window and show the results in Figure 8 below. The final statistic using 3 and 5 day event windows changes but follows a broadly similar trend. A relatively short event window is less likely to capture events other than the announcement and, as noted by Dubow and Monteiro (2006), with wider event windows it is more difficult to detect statistically significant cumulative abnormal returns. We show that the interpretation of the results is robust to these changes. This gives confidence to the findings from our proposed methodology and does not provide a clear rationale for changing the event window length from the current 2-day.

Figure 8: New methodology with alternative event windows (Annual % of APPM)

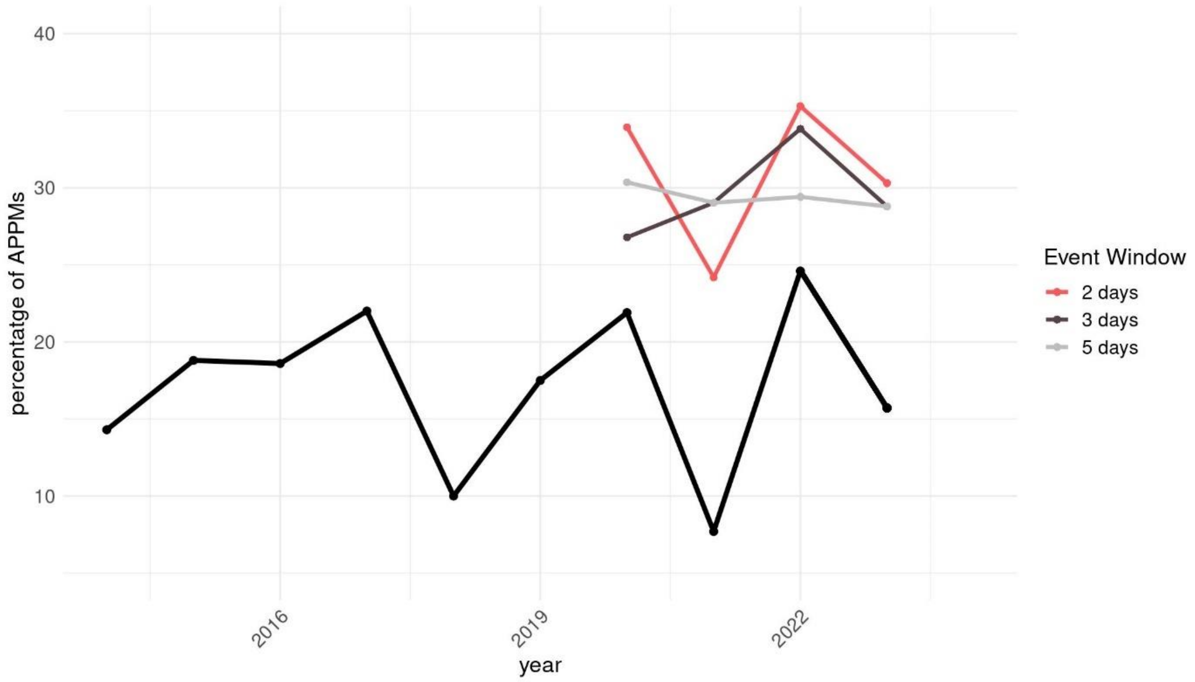

Limitations

Although we test our results for robustness, the analysis has some limitations. First, we only investigated intraday stock price data from mid-2019. This limits the ability to compare the results across an extended period of time. This makes it difficult to determine whether the trends we observe will be persistent or may be unique to these years.

When checking for robustness, we did not implement Market Models and CAPM model with intraday data. In addition to significantly higher computational costs, implementing those models at higher frequency requires additional assumptions (such as constant beta estimates over a day for CAPM models). However, we do not expect a significant difference between the Market Model or CAPM and the current methodology at intraday frequency for the same reasons listed in previous section (i.e., predominately, the frequency at which some of the stocks are traded, and the incidence of zero returns in the estimation window).

When moving to intraday data, we did not repeat the analysis sampling prices at lower frequencies (e.g., 15-minute or 30-minute intervals). The main motivation to move to higher frequencies is to capture as much information as possible in the statistic. However, as some of the stocks we analyse are illiquid (as noted in previous sections and in Annex 3), different sampling intervals could have been considered. We do not expect large variation in the MCS if a different interval had been selected.

Our statistic only captures abnormal price movements and is not a direct measure of the level of insider trading in the UK equities market. Abnormal price movements can be caused by information leakage or other factors such as financial analysts or the media correctly predicting the likelihood of an imminent takeover, leading to legitimate trades ahead of an announcement. In practice, the MCS is an informational tool, among many others, used by the FCA to detect and act on potential insider dealing.

A limitation of the MCS is that it only covers abnormal price movements in equities markets. Empirical studies document the presence of informed trading through derivatives (e.g., Kacpercyzk and Pagnotta, 2019; Patel et al., 2020; Bohmann and Patel, 2022).

Finally, we do not investigate whether prices moved in the expected direction of the announcement. The MCS only identifies announcements where inside information could have been potentially used, independently of whether a trader profited from that information[3]. In addition, our test only captures abnormal positive price moves, assuming all announcements have a positive impact on the stock’s price. In practice, this may not always be the case.

[3] The Potentially Anomalous Trading Ratio (PART) filters for portfolios where a participant traded significantly more in the direction of the announcement.

Conclusion

This review addresses known limitations in the Market Cleanliness Statistic methodology. It improves the identification of abnormal pre-announcement price movements (APPMs) and corrects for potential bias during periods of heightened market volatility.

Using intraday price data, we identified more cases of potential insider trading. This occurs because the statistic now captures APPMs for which the abnormal price movement took place on the same day as an announcement, where the announcement occurs during market hours. We also reduce the length of the estimation window from 240 to 60 days before an announcement. This means less announcements are excluded from the analysis. By introducing a Market Comparison test, we control for market volatility and reduce the number of false positives and false negatives.

Implementing these methodological changes has important benefits. We adapt our methodology so that our statistic is less affected by market volatility: this means the statistic now more accurately reflects the APPMs that are due to firm-specific information leakages.

The revised measure is higher, reflecting the scope of the statistic now including potential insider trading on the day of an announcement, as well as in the two days prior. The change does not indicate a deterioration in market cleanliness.

From 2024, the FCA will implement the Market Comparison method using intraday data to estimate the Market Cleanliness Statistic. We welcome feedback on our methodology and findings, in particular on:

- The Market Comparison test,
- Alternative models to estimate expected returns,
- Using intraday price data for estimating cumulative abnormal returns.

Annex 1

Alternative models and techniques considered

We considered five alternative models to the one proposed by this paper. Models 1-5 introduce different approaches to expected return estimation only. While these models did not return a significantly different MCS from the current figure, we thought it was sensible to check the effect of introducing some well-known models to our MCS analysis for robustness.

In next sections we outline how we estimated expected returns for models 1-5. In each model, expected returns were used to estimate measures of abnormal returns, based on which the test statistic and Critical Value are computed. We also include some commentary around how Abnormal Pre-announcement Price Movements (APPMs) are detected, based on which we calculate the final yearly MCS.

Model 1: The current methodology (Goldman et al. 2014)

Consider P_t to be the stock price of a given firm at day t where $t \in (0, \dots, T)$, and its return as $r_t = \frac{P_t - P_{t-1}}{P_{t-1}}$. In our current method, expected returns are simply the average return calculated over the estimation window.

Model 2: One-factor Market Model

For each t (day) r_t can be modelled using One-factor Market Model as

$$r_t = \alpha + \beta^{MM} R_{m,t} + \varepsilon_t \ (1)$$

where $R_{m,t}$ is return to a FTSE350 and ε_t is stock specific variation. Provided some weak assumption and absence of arbitrage, we estimate the expected return $E(r_t)$ by,

$$E(r_t) = \alpha + \hat{\beta}^{MM} R_{m,t}$$

Parameters of the model in (1) are estimated by Ordinary Least Squares (OLS) estimator over the estimation window and all expected returns (the estimation window, event window and cooling off period) are calculated using same estimated $\hat{\beta}^{MM}$.

Model 3: Extended Market Model

We have also used an extended version of Market Model in which we control for return's serial correlation by including lags of r_t in (1). Hence, the expected returns are modelled as:

$$r_t = \alpha + \beta_1 r_{t-1} + {\beta_2}^{MM} R_{m,t} + {\beta_2}^{MM} R_{m,t-1} + \varepsilon_t$$

As in Model 2 above, the parameters of the model are estimated via OLS over the estimation window and all expected returns (for the estimation window, event window and cooling off period) are calculated using same estimated parameters.

Model 4: Capital Asset Pricing Model

Another approach to calculate expected returns is to employ a version of CAPM model, such as

$$E(r_t) = R_{f,t} + \beta^{CAPM}\left(R_{m,t} - R_{f,t}\right) \quad (3)$$

with $R_{f,t}$ being the risk-free rate and $R_{m,t}$ a measure of the market returns. This one factor version of the CAPM model is the one used to calculate the equities Market Cleanliness Statistic by the Australian Securities and Investments Commission (ASIC, 2016). Some studies add extra terms to the model to capture fat tail variations or jumps observed in the data (see Todorov and Bollerslev, 2010 for example).

There are various methods to estimate the "stock beta" or β^{CAPM} in the related literature. The Australian Market Cleanliness statistics method takes the ratio of covariation between the market and stock returns over variance of stock returns as an estimator of β^{CAPM}. Thus, Model 3 is the CAPM model with constant β^{CAPM} for stock i as:

$$\hat{\beta}_i^{CAPM} = \frac{COV_{i,m}}{VAR_i}$$

Where $\hat{\beta}_i^{CAPM}$ is estimated over the estimation window and used to calculate expected returns. In our analysis, we use Overnight Index Swap daily rates as $R_{f,t}$ and returns on FTSE350 for $R_{m,t}$.

Model 5: CAPM with Time-variant Betas

Some advanced methods use higher frequency data to better estimate the beta (see Todorov and Bollerslev, 2010). In this occasion, we have employed a simpler approach, namely the "rolling window", in which $\hat{\beta}_i^{CAPM}$ is estimated on a rolling window over the estimation period.

Let the estimation window have $T = s + R + P$ daily observations. For each $t \in \{1, ..., P\}$ the $\hat{\beta}_{t,i}^{CAPM}$ is estimated using observations from $s + t$ to $s + R + t$. We considered a window of $R = 30$ in our analysis. $\hat{\beta}_{t,i}^{CAPM}$ for the event window was then predicted by an AR (1) model for the cooling off and event window.

After getting expected returns, the abnormal returns for each t were calculated as a difference between observed returns and expected returns:

$$ABR_t = r_t + E(r_t).$$

Bootstrap method

To test whether there is abnormal price movement prior to the takeover announcement, we check our test statistic, $CAR_{n,T}$, against the 90th percentile of its distribution (CV) in the estimation window. To approximate the cumulative abnormal returns distribution, we use a block bootstrap method. Consider the test hypothesis as the following:

$\boldsymbol{H_0}$*: normal cumulative returns during the event window*

$\boldsymbol{H_A}$*: abnormal cumulative returns during the event window*

For each $t \in \{T - 250, \ldots, T - 10\}$ (i.e. each day in estimation window) we compute the $CAR_{n,t}$ and implement the block bootstrap (with replacement) on vector of CAR as the following:

- Step 1: we randomly draw a day index t.
- Step 2: our first block of observations then is $\{t+1, t+2, t+3, \dots, t+b\}$, where $b = 5$ is the length of each block.,
- Step 3: we repeat step 1 and 2 for l=20,000 number of times such that $b \times l = 100{,}000$.
- Step 4: we rank all these 100,000 CAR and compute the CV at 90th percentile.

We reject the null hypothesis at 10% when $CAR_{n,T} > CV$. Note that the rejection rule is based on a one-sided test in which we are only concerned with insider trading associated with "good" news.

Annex 2

Using Intraday Data

As described in the current methodology section, to detect abnormal price increases, we construct the test statistic as the cumulative abnormal return (CAR) over n-days before the event. Some changes were implemented to estimate the current methodology with intraday data.

Firstly, for intraday analysis, this window is of a duration of 60 days. Secondly, when replicating the current methodology with intraday data, we measure expected returns as the mean return over the estimation window by taking only returns from 5-minute intervals when a trade took place (i.e., effectively excluding created returns observations by carrying last price forward when calculating the mean return). However, CAR are calculated considering returns from all 5-minute intervals.

Lastly, we require not only a minimum of 100 observations but also that these are spread over 20 trading days for an event to be included in the statistic. Figure 9 shows the Market Cleanliness Statistic resulting from implementing the current methodology with intraday data. The resulting statistic (brown line) is significantly higher. Most of this increase is explained by the incidence of "intraday" events amongst those newly identified with an APPM after using intraday data (See Table 2).

Figure 9: The current methodology – Daily vs Intraday Data

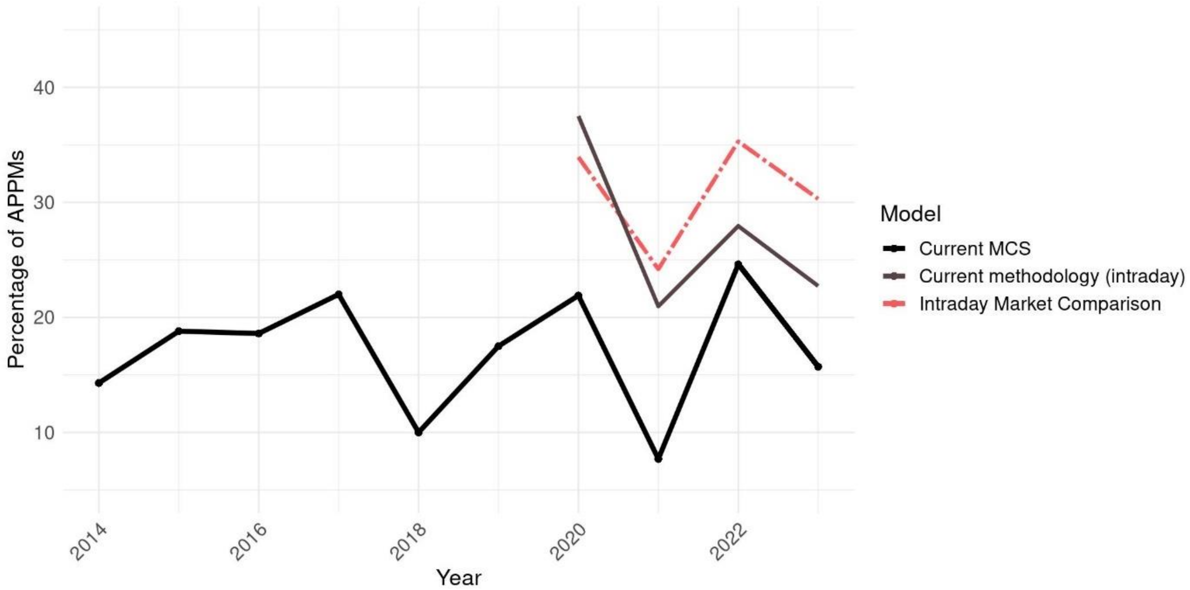

Using the Market Comparison methodology, the market cleanliness statistic is higher for all years except 2020. Table 3 shows the number of 'false positives' and 'false negatives' which are corrected by the cross-section Market Comparison test. For example, 2020 is the year with most false positives. A significantly higher number of APPMs are observed for the last quarter of 2020 when implementing the current methodology in intraday data, shown in Figure 9. When comparing the stocks of interest with others in the market, a significant proportion of events that had been detected as APPM are marked down (i.e., 5 events). A similar correction was not observed in the daily data analysis. One potential explanation is simply that market volatility is better reflected through intraday data. However, further investigation and years of samples may be needed before reaching a conclusion.

Table 2: Events newly found with an abnormal price movement by current methodology using intraday data.

Year	New APPMs	Of which, are intraday	% intraday announcements
2020	12	7	58%
2021	8	7	88%
2022	7	6	86%
2023	8	6	75%

Table 3: False Negatives and False Positives (intraday, annual)

Year	False Negatives	False Positives
2020	3	-5
2021	3	-1
2022	7	-2
2023	7	-2

Annex 3

The sample of firms with takeover offer announcements

As described in the Research Design - Data section, our analysis considers 302 events (takeover announcements) from 2020-2023 based on a list of takeover announcements provided by the Secondary Market Oversight team in the FCA. Announcements where the company is delisted, or not traded in a UK based market or where price information is not available are excluded from the analysis. This results in 285 firms with announcements in our sample.

As shown in Figure 10, most of the firms from the sample are in the Consumer Discretionary, Financials and Industrials sectors. Across all sectors, most of the firms subject to announcements are small cap (i.e., with a market value of less than 200 million pounds).

Figure 10: Number of firms per Sector and Market Value

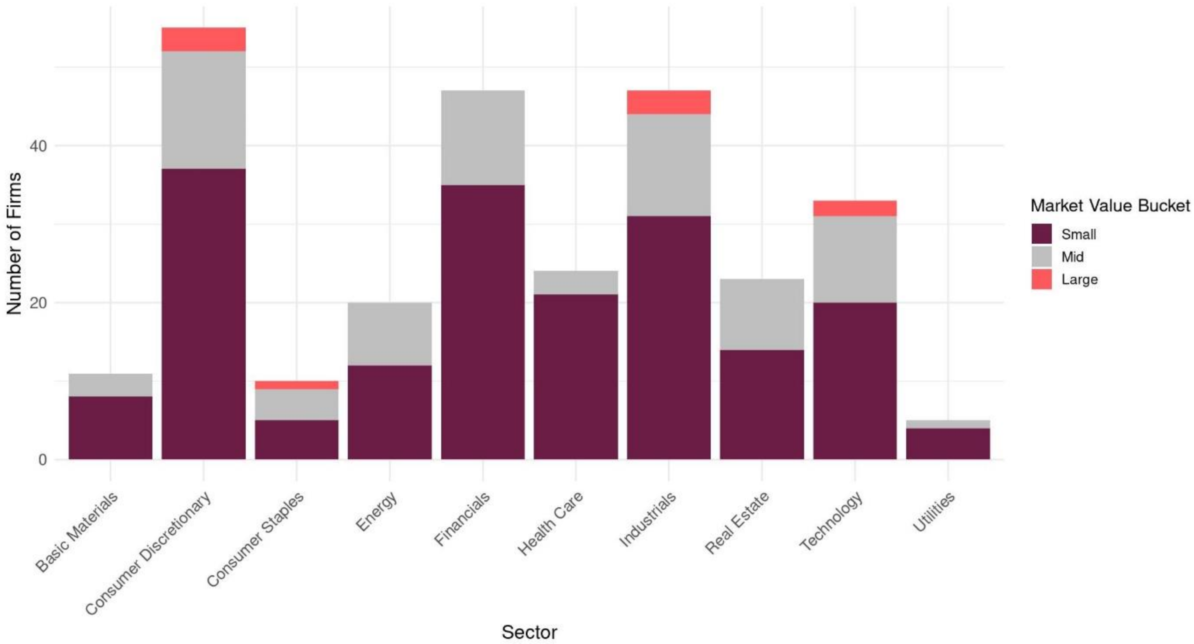

In Figure 11, each dot represents a firm with a takeover offer announcement. On the y-axis, the number of firms (per sector) and number of days in which a stock was traded within the estimation window are presented. The maximum number of days a stock can be traded within the estimation window is 60 days (which is the length of the window). Across all sectors there are stocks that are not always traded. In some sectors, exclusions are triggered (i.e., when the stock has not reached a minimum number of 20 days traded). However, when looking at the distribution of the number of traded days per market value bucket (i.e., Figure 12), there is a clear pattern. Most of the firms that we analyse are small cap. It is within this group where we see stocks being traded less frequently.

Figure 11: Trading frequency of announcing firms' stock, per sector

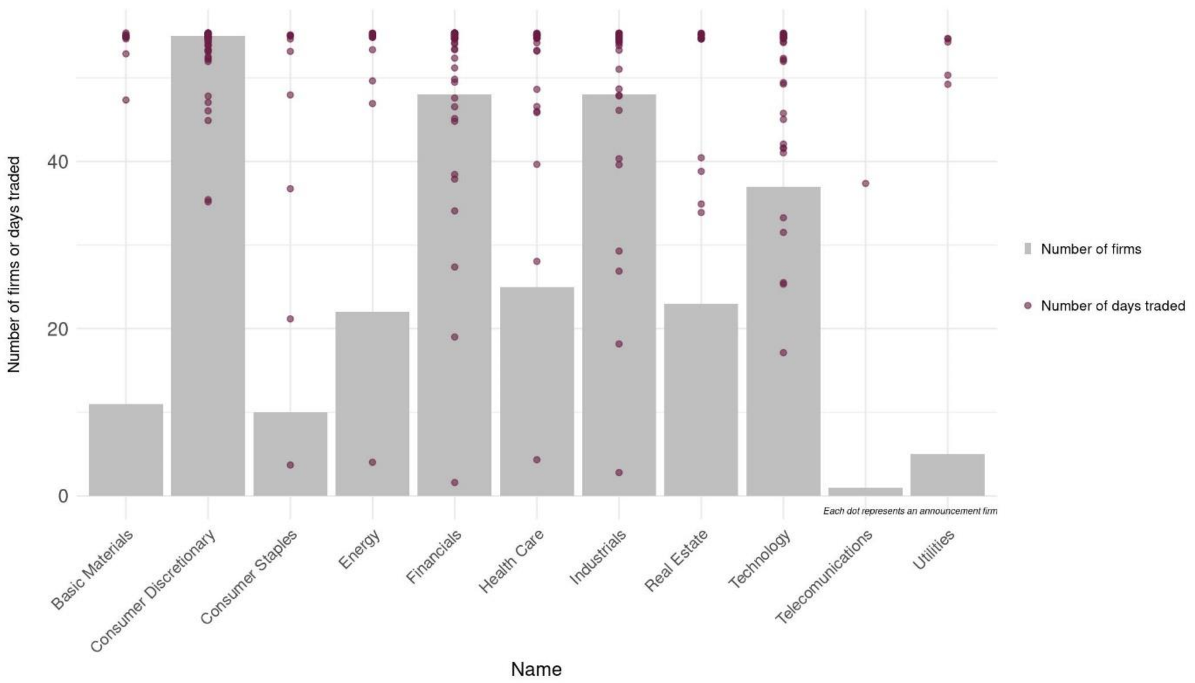

Figure 12: Trading frequency of stocks, per market value

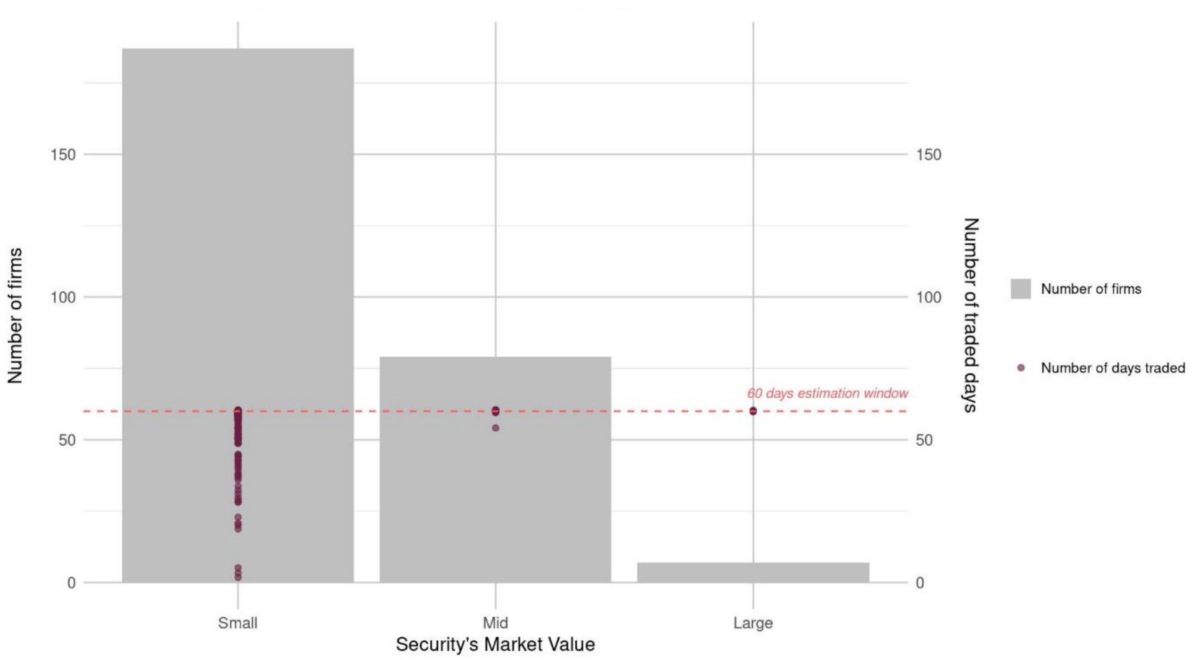

We checked if illiquid stocks (i.e., those that are not traded as frequently) contribute most to the total number of announcements classified as APPMs. In Figure 13, we split the stocks according to the number of days they were traded within their estimation window (i.e., allocating them across 4 quantiles). The most illiquid stocks are those allocated to the first and second quantiles. What we observe is that the most illiquid stocks contribute least to the pool of firms found with APPM (i.e., only 20%). This percentage is even lower once that stocks not meeting the minimum observations and traded days requirements are excluded from the sample.

Figure 13: Number of stocks per traded days' quantile, and contribution to the total number of APPMs across the sample.

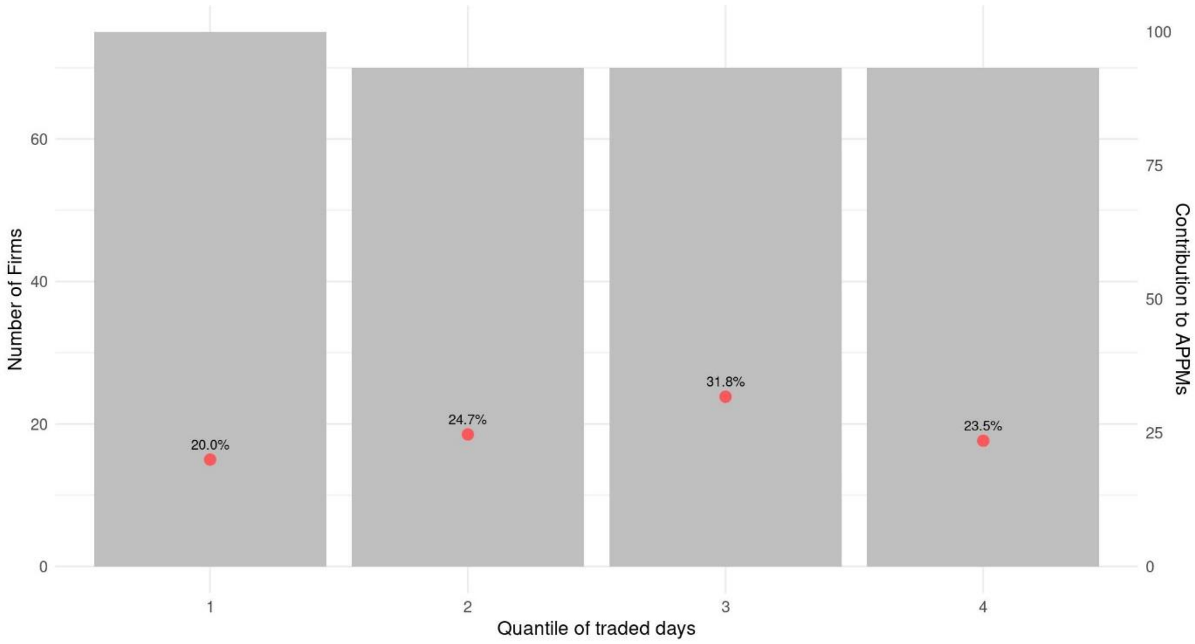

References

Australian Securities and Investment Commission (2016). Review of Australian equity market cleanliness*, ASIC Report 487*.

Austin, J. (2015). Unusual trade or market manipulation? How market abuse is detected by securities regulators, trading venues and self-regulatory organizations. *Journal of Financial Regulation*, *1*(2), 263-283.

Austin, J. (2016). What exactly is market integrity: an analysis of one of the core objectives of securities regulation. *Wm. & Mary Business Law Review*, *8*, 215.

Bohmann, M. J. M., Patel, V., (2022). Informed trading prior to FDA announcements. Journal of Business Finance and Accounting 49, 1211-1236.

Comerton-Forde, C., & Putniņš, T. J. (2011). Measuring closing price manipulation. *Journal of Financial Intermediation*, *20*(2), 135-158.

Comerton-Forde, C., & Putniņš, T. J. (2014). Stock price manipulation: Prevalence and determinants. *Review of Finance*, *18*(1), 23-66.

Dubow, B., & Monteiro, N. B. (2006). Measuring market cleanliness. *Available at SSRN 1019999*.

Financial Conduct Authority (2023). Market cleanliness statistics 2022/23. *https://www.fca.org.uk/data/market-cleanliness-statistics-2022-23*

Financial Conduct Authority (2013). Financial Conduct Authority Handbook, https://www.handbook.fca.org.uk/

Kacperczyk, M., Pagnotta, E., (2019). Chasing private information. Review of Financial Studies 32, 4997–5047.

Goldman, J., Hunt, S., Minter, P., Suntheim, F., & Zaman, Q. (2014). Why has the FCA's market cleanliness statistic for takeover announcements decreased since 2009? *FCA Occasional Paper*, (4).

Monteiro, N. B., Zaman, Q., & Leitterstorf, S. (2007). Updated measurement of market cleanliness. *FSA Occasional Paper, (25).*

Patel, V., & Putniņš, T. J. (2020). How Much Insider Trading Happens in Stock Markets? In *American Finance Association (AFA) Annual Meeting*.

Patel, V., Putniņš, T. J., Michayluk, D., Foley, S., (2020). Price discovery in stock and options markets. Journal of Financial Markets 47, 100524.

Putniņš, T. J. (2012). Market manipulation: A survey. *Journal of Economic Surveys*, *26*(5), 952-967.

12 Endeavour Square, London E20 1JN
Telephone: +44 (0)20 7066 1000 Website: www.fca.org.uk