Short-run returns around the Trades of Corporate Insiders on the London Stock Exchange

Sylvain Friederich,* Alan Gregory, John Matatko, and Ian Tonks.†

January 5, 2001

Abstract

Previous work examined the long-run profitability of strategies mimicking the trades of company directors in the shares of their own company, as a way of testing for market efficiency. However, the evidence regarding returns during the month containing the insider trade was ambiguous. The current paper examines patterns in abnormal returns in the days around these trades on the London Stock Exchange.
We find movements in returns that are consistent with directors engaging in short-term market timing. We also report that some types of trades have superior predictive content over future returns. In particular, medium-sized trades are more informative for short-term returns than large ones, consistently with Barclay and Warner's (1993) "stealth trading" hypothesis whereby informed traders avoid trading in blocks.
Another contribution of this study is to properly adjust the abnormal return estimates for microstructure (spread) transactions costs using daily bid-ask spread data. On a net basis, we find that abnormal returns all but disappear.

Keywords: market efficiency, corporate insiders, insider trading, informed trading, London Stock Exchange.

JEL Classification: G 14

*Corresponding author. Tel: (+44-20) 7955-7893. Email: s.j.friederich@lse.ac.uk

†The authors are respectively affiliated with the Financial Markets Group, London School of Economics and Université de Paris - I (Friederich), the School of Business and Management, University of Exeter (Gregory and Matatko), and the Financial Markets Group, LSE and Department of Economics, University of Bristol (Tonks). We thank an anonymous referee as well as Alvaro Almeida, Margaret Bray, Ward Brown, Simon Emrich, Urs Haegler, Marc Henry, Richard Payne, Frank Reilly, Alex Stremme, and seminar participants at the LSE, the University of Bristol, the FMA and the Royal Economic Society 2000 conferences.

1 Introduction

Do the actions of corporate insiders convey private information about their company's prospects to the market? In this paper we investigate this question by examining the behaviour of daily returns, immediately around the trades of company directors. Previous work by Gregory, Matatko, Tonks and Purkis (1994) and Gregory, Matatko and Tonks (1997) identified significant long-run abnormal returns following the trades of directors in UK companies using monthly stock price data. Surprisingly, these papers found no evidence of stock price movements in the month of the directors' trade. But if abnormal returns appear immediately after the trade, using monthly data will miss out on them. The use of higher frequency data will therefore indicate whether relatively long holding periods are required for mimicking strategies to be profitable, and whether earlier results were not taking short-run excess returns into account (and were therefore underestimating total returns) because they were not precisely pinpointing the event date.

The behaviour of stock price returns around the trades of company directors at any frequency is interesting for two reasons. First, the extent to which insiders trade profitably on private information to generate abnormal returns would be a violation of strong-form market efficiency. The ability of insiders to earn money from trading carries welfare and regulatory concerns, and raises issues about corporate governance. Of course, insider trading (defined in the UK as insiders acting on price-sensitive information) is illegal in many countries. Though we may draw a distinction between illegal insider trading and legal trading by insiders: there is nothing illegal about an insider taking a view that their company is misvalued by the stock market, and trading on that basis, provided that the insider does not trade on price-sensitive information.

The second reason for studying price reaction to directors' trading concerns whether outside investors can mimick the actions of insiders to also earn abnormal returns. This second issue would represent a violation of semi-strong form efficiency. Recently Loughran and Ritter (2000) have coined the term "behavioural timing" to refer to the class of anomalies that represent managerial responses to equity undervaluation. Although they emphasise behavioural timing in events like initial public offerings, seasoned equity issues, and takeovers, these events are likely to follow periods of longer term mispricing, by virtue of the fact that such events have a long lead time and are associated with large transactions costs. On the other hand, directors can trade on short term mispricings with minimal transactions costs. If directors are indeed seeking to exploit misvaluation of their firms' equity in the markets, they will tend to buy (sell) following periods of abnormal under- (over) performance of the shares.

We concentrate our attention upon smaller firms, since previous work by Gregory et al. (1997) reported a highly disproportionate amount of directors' trading activity in less liquid stocks. In addition, Loughran and Ritter (1999) argue that undervaluation on which behavioural timing relies is likely to be more common and larger in small firms than large ones: for any given mispricing "there will be a stronger force pushing the price towards fundamental value (and thus limiting the magnitude of any misvaluation) for big stocks."

Of course, markets are only efficient up to the amount of transactions costs incurred through the implementation of a trading strategy. This is important given that the other side of the smaller firms greater misvaluation argument, is larger transactions costs (Shleifer and Vishny 1997). We therefore adjust the profitability of these trading strategies using data on daily bid and ask prices: for an outside investor to earn abnormal returns following a stock purchase by a director, the outsider would buy the stock at the ask price and sell it at the bid at the end of the holding period (and conversely for a director sale).

We examine a large sample of trades by UK directors in "mid-cap" companies over the period October 1986 to November 1994. We report evidence of trading around short-term price changes by corporate insiders over the sample period, and find that directors trade on the basis of market timing: they buy shares after a decline in their company's share price, and sell shares after a run-up in the share price. In addition we find that there are abnormal stock price movements in the days after the directors' trading signals. This suggests that directors have the ability the forecast short-term returns (although they do not imply that they are trading illegally). We also report that some insider trading signals dominate others in terms of predictive contents over future returns. Buy trades are followed on average by larger abnormal returns than sell trades. "Clustered" trades strongly dominate large ones in terms of signal strength. We find that most of these trades are of medium size, and generally report evidence that medium-sized trades as a whole seem more informative than large ones. This is consistent with the "stealth trading" hypothesis of Barclay and Warner (1993). After computing returns from bid and ask prices, to take account of the microstructure costs inherent in a mimicking trading strategy, we find that potential short-term abnormal returns to outsiders are more or less whittled down to zero. We conclude that although stock price patterns after directors' trades are statistically significant, their economic significance is not necessarily a cause of great concern for market authorities.

1.1 Related research

We first review the broader debate on the usefulness and effectiveness of insider trading regulation, and then the literature on the profitability of insider trades to which our work is directly related.

1.1.1 The insider trading debate

Before even considering whether it is *effective*, financial economists are divided over whether insider trading regulation is *useful*. Part of the academic literature since Manne (1966) emphasises the welfare benefits of unregulated trading by company executives, which include alignment of managers' and shareholders' incentives as well as the potential for increased price informativeness: informed traders generally make markets more efficient, and insiders are just seen as a special kind of informed traders whose information has high precision and is acquired at no cost (see Dennert (1991), Hu and Noe (1997) for surveys, and the model by Leland (1992)). On the other hand, the information-based microstructure literature such as Kyle (1985) argues that bid-ask spreads will increase or depth decrease with the number of informed traders in a market, thus emphasising the detrimental effects of their trading activity on market liquidity.

Evidence on the effects of insider trading is more ambiguous. Early work by Meulbroek (1992), who examined cases of illegal insider trading actually prosecuted by the SEC, seemed to indicate that regulations were pointless because of the market's ability to detect genuine insider trading, leading to a quick price adjustment. Garfinkel (1997) reports that earnings announcements appear to be more informative after the passage of the insider trading act of 1986, implying that insider trading did cause price adjustments before the announcements earlier on. Kabir and Vermaelen (1996) similarly report evidence of reduced trading volumes and somewhat slower speed of price adjustment following the introduction of a regulation forbidding corporate insiders to trade two months before an annual earnings announcement on the Amsterdam Stock Exchange. On the other hand, and more in line with the predictions of the microstructure literature, very recent work by Bettis, Coles, and Lemmon (2000) finds for the US market that spreads are narrower during periods when insider trading is forbidden by company charters. Therefore, the evidence is consistent with the predictions of both strands of the theoretical literature: increased price efficiency, but reduced liquidity.

In practice, an important argument for regulatory bodies is that of unfairness: there is no "victimless crime". Less informed or liquidity traders bear the cost of insider trading, and liquidity could suffer because uninformed participants tend to withdraw from the market. Market authorities have as a consequence significantly tightened their insider trading regulations over the past decade, and company directors are prohibited from trading before the release of company information in several countries. Empirical evidence of the effectiveness of these regulations exists. For the US, Bettis, Coles, and Lemmon (2000) find that "blackout" periods successfully suppress insider trading. For the UK, Hiller and Marshall (1998) similarly report that trading by directors drops off to almost zero before earnings announcements.

1.1.2 Are the trades of corporate insiders profitable?

At times when directors' trades are authorised, regulators often require directors to disclose them to the market. Using these data, a sizeable literature has developed examining whether corporate insiders seem to benefit from their legal trading and whether outsiders imitating these trades may also reap abnormal returns in the medium to long run. Early work in the US by Jaffe (1974) and Finnerty (1976) identified excess returns in the months after a director's trade, which suggests that insiders are able to predict and exploit future returns. However, this apparent semi-strong form inefficiency was explained away in a later study by Seyhun (1986) in terms of average costs of trading. More recent work by Bettis, Vickrey, and Vickrey (1997) reports that abnormal profits can be made when focussing only on the insiders' block trades (over 10,000 shares, following the definition of blocks used in US markets) as a signal, using again a measure of estimated transactions costs of mean spreads plus mean commissions. Lakonishok and Lee (1998) and Jeng, Metrick, and Zeckhauser (1999) also report long-term excess returns but conclude that they are modest in size.

Empirical work on directors' trading using UK data reports fairly comparable findings. Early work by King and Röell (1988) and Pope, Morris, and Peel (1990) seemed to produce conflicting results: the first study reports positive abnormal returns after director purchases, while the second concludes that significant abnormal returns mostly follow director sales. Further work by Gregory, Matatko, Tonks, and Purkis (1994) and Gregory, Matatko, and Tonks (1997) reconciled those conflicting results by making signal definitions comparable and controlling for size effects. Evidence was found of small pre-transaction costs abnormal returns for some signal definitions.

All of these studies focused on the profitability of strategies imitating the trades of directors in the medium to long-run (several months or years). The evidence on whether corporate insiders may try to exploit short-term price movements and therefore whether these trades may be profitably mimicked in the short-run is unclear. For the US market, earlier studies such as Jaffee (1974) or Seyhun (1986) report some evidence of abnormal returns immediately around the insiders' trades. However, conflicting evidence has been produced in recent work by Lakonishok and Lee (1998) who use a much larger dataset, and conclude that "surprisingly, in spite of the extensive coverage that insider activity receives, the market basically ignores this information when it is reported. Moreover, there is very little action around the time when insiders trade. The magnitude of the returns observed is typically below 0.5 percent."

In UK research by Gregory, Matatko, Tonks, and Purkis (1994) and Gregory, Matatko, and Tonks (1997), although there was evidence of long-run abnormal returns following the trades of corporate insiders, returns during the month containing the trade were found to be not significantly different from zero. However, the exact day of the event was not precisely identified in that research since it was using monthly data, therefore whether there have been short term price movements remains an open question. Also, in the latter study, the authors found that the price reaction in the months after the directors' trades was, surprisingly, inversely related to the strength of the signal. They conjectured that this was because in the case of a strong signal, most of the price reaction occurred within the month of the trade. The current paper will attempt to reconcile these somewhat contradictory findings.

In the next section we give more details on the regulatory environment, the data used and the methodology. Section 3 presents the results. Finally, section 4 provides a summary and conclusion.

2 Regulatory environment, data and methodology

2.1 Regulatory background

In the UK, the 1985 Companies' Act specifies that directors are prohibited from dealing in the securities of their own companies for a period of two months prior to the preliminary announcement of year-end or half-year results, and at other times prior to the announcement of price-sensitive information.[1] The difficulty is to define what "price-sensitive information" consists of: clearly included are dividend, earnings, acquisition or spin-off announcements, board appointments or departures, or security issues. This leaves a large grey area open to interpretation: as the London exchange literature indicates, "there are many events which can trigger significant movements in share prices, such as information on a new product, the fact that sales of a new product are not meeting expectations, or that the company has obtained a large order or embarked into a major redundancy programme," but in general "It is not feasible to define any theoretical percentage movement in a share price which will make a piece of information price-sensitive. Attempts at a precise definition of "price-sensitive" are not possible" (London Stock Exchange (1996), pp. 4 and 2, respectively). The disclosure of business and financial information is necessarily imperfect, and this leaves open the possibility of trading around undisclosed events causing short-term price changes.

The disclosure requirements for directors' trades are as follows: directors must inform their company "as soon as possible after the transaction and no later than the fifth business day" of any transaction carried out for their personal account. In turn, a listed company must inform the Stock Exchange of the transaction "without delay and no later than the end of the business day following receipt of the information by the

[1]Note that, in the UK as in the US, further obligations with respect to director's trading are quite often set out in the charters of individual companies, especially larger ones.

company" (London Stock Exchange (1998), p. 8). The Stock Exchange disseminates this information immediately to data vendors as well as via its own "Regulatory News Service" (the company should also enter this transaction in the Company Register which is available for public inspection within three days of reporting by the insider, but this way of disseminating the information is nowadays much less important).

As a comparison of regulatory requirements, US regulators have taken a different approach: the Securities Exchange Act of 1934 requires insiders to refrain from trading on "material" undisclosed information, and to fill in statements of their holdings in the first ten days of the month following the month in which the trade occurred. Profits made on short-term "swings" in prices (formally, within 6 months) must be surrendered to the company.[2] An important difference with the UK regulatory regime is that in the US, "insiders" are more broadly defined and in particular include large shareholders, who are subject to the same reporting requirements as company officers and directors.

2.2 Data sources and sample selection

Data on the trades of directors for the period October 1986 to end-1990 were obtained on microfiches from the London Stock Exchange. For 1991 to end-1994, the data were provided to us by Directus Ltd, a subsidiary of Barra which re-sells these data to the financial services industry along with investment advice. For all listed companies, this dataset gives details of the identity of the director, the date of the trade, the quantity and direction of the shares traded, and in most cases the transaction price at which the director traded.[3]

The stock price series used to compute returns are obtained from the closing bid and ask quotes of the competing market makers (available from Datastream) and are adjusted for stock splits, stock dividends and issues.

[2] A recent theoretical literature models the welfare effects of these disclosure obligations. Examples are Fishman and Hagerty (1995), in which the trade reporting is used to manipulate the market, while the mandatory disclosure has in Huddart, Hughes, and Levine (1999) the effect of slowing down price discovery. We do not directly address these issues here.

[3] Gregory et al. (1994) found that option-related trades by Directors were insignificantly related to abnormal returns. Therefore, option-related trades were removed from the data. Directus code all non-standard trades. These include all non-beneficial trades, and the take-up or sale of rights. We also exclude such trades from our sample, leaving only normal beneficial trading activity.

As mentioned above, a contribution of this study is to adjust estimates of the profitability of mimicking strategies for microstructure costs. The selection of stocks was therefore governed by the availability of daily bid and ask prices over the sample period, provided in Datastream for most firms in size deciles 1 to 4 of the constituents of the FT-All Share index. We chose not to focus on the most liquid stocks (FTSE-100 companies), because previous work by Gregory, Matatko, and Tonks (1997) showed higher gross abnormal returns in less-liquid securities. Their findings are nonetheless compatible with either a "less-researched firms" effect, or the argument which would suggest that any inefficiencies in pricing are more likely to be found outside the most liquid stocks. One potential problem of studying this group of firms is that any lack of liquidity could impart a bias to our results. We therefore examine the sensitivity of our results to the potential problems this may give rise to. In particular, we adjust betas for thin or non-synchronous trading, perform outlier checks, and conduct further tests on significance levels, all of which we detail below.

Our sample is comparatively homogeneous in terms of firm size. The FTSE-250 index supplies a convenient benchmark for the computation of abnormal returns since it includes the 250 firms immediately below the most liquid ones (FTSE-100) in terms of market values.

A survivorship bias is possible in the sense that prices were not available for dead companies over the period, which includes companies taken over. Our aim is to see whether signals, on average, can be profitably exploited, and not to estimate the profitability either of risk arbitrage strategies, or around any highly unusual event of the kind. Therefore, whether a small number of (possibly very high) returns made by directors whose companies were acquired would significantly bias estimates upward is an open question.

The event date is the day of the insider's trade itself. There is evidence that the information on those trades reaches the market in a timely manner, since in many cases, company charters specify that directors must obtain clearance from the board before being able to trade. When the trade is executed, it is often disclosed at the same time. This assumption can also be seen as justified by Meulbroek (1992) who reports in her study of cases of illegal insider trading that the information about the trades of insiders gets quickly detected and incorporated into stock prices even without any disclosure. She concludes that "both the amounts traded by the insider and additional trade-specific characteristics lead to the market's recognition of the informed trading". If it were the case that trades are disclosed with a few days' lag, or that they we not detected by the market, this would only lower estimated abnormal returns.

2.3 Descriptive statistics

Over these eight years and 196 "less-liquid" companies,[4] we observe a total of 4,399 trades (2,558 buy and 1,841 sell transactions).[5]

Descriptive statistics on individual trades by value are given in panel A of table 1: over the sample period, the average buy transaction was worth about £66,000, dwarfed by the average sell of about £343,000. The median buy transaction was £6,650, and the median sell was £32,600. The distributions of both types of trades are clearly skewed to the right, with some very large transactions in both cases: the largest transaction on the buy side was almost £23 million (in 1988), while the largest sell was a staggering £154 million (in 1991). Sell transactions are slightly more infrequent, but much larger.

2.4 Returns and signal definitions

The events in our study are a director's buy or sell trade, computed from the data on individual trades as the *net* quantity of shares traded on an event day. This is standard in this literature (since in some cases more than one director traded on the same day, very occasionally in opposite directions) although it is usually done at the monthly level. A trade is therefore a buy (sell) event if the net traded quantity is positive (negative). Filters were then applied (detailed below) to focus on certain categories of signals. Panel B of table 1 reports descriptive statistics on the distribution of the net buy and sell values traded (removing the sign for sells), for

[4]According to London Exchange terminology, which distinguishes between "liquid" securities (FTSE 100 constituents), "less liquid" securities (FTSE 250 constituents), and "illiquid" securities (the rest).

[5]The actual transaction price was missing for about 300 of these trades, in most cases for the first two years of the sample. For these we extracted the (unadjusted) price data from Datastream. This is not consequential since we are not computing the profitability of the trading strategy to the insider herself, this is only useful when applying signal filters by value.

every year and for the whole sample period. There were 3,409 event-days in total, 1,887 on which directors were net purchasers, and 1,522 when directors were net sellers. Directors as a whole were clearly net sellers of their companies' shares over the sample period.

Using daily closing bid and ask quotes, we compute arithmetic returns on each stock as the percentage change in the midpoint of these quotes (where semi-annual dividend payments were obtained and added back into prices on the ex-dividend dates). We also compute daily returns on the FTSE-250 index, which will be used as a benchmark in abnormal returns computations.

Panel C in table 1 shows the distribution of directors' gross and net trades across firm size-deciles. Firms were ranked by market capitalisation and divided into size-deciles at the beginning of each year. The directors' trades in each firm in that year were then allocated across these size-deciles. This was repeated for each year of the dataset. It can be seen that the distribution of directors' trades across size-deciles is fairly uniform. This is not surprising, given that the firms in our sample are all constituents of the FTSE-250 index and are relatively homogeneous in terms of size.

2.5 Methodology

We examine the short-term movements in returns around the event date to investigate the ability of directors to engage in "market-timing" using an event-study methodology. The use of daily data is central to our aims but also an advantage in estimation terms because the joint hypothesis or "bad-model" problem is much less serious in studies that focus on short return windows. Besides, statistical tests have much greater power with daily data: Brown and Warner (1985) indicate that the "rejection frequencies are roughly three times those reported for monthly data, thus highlighting the substantial gains to more precise pinpointing of an event." The only caveat in the interpretation of the results is that we are not claiming that the event is directly causing any observed pattern in returns, since the directors' trading process is endogenous with respect to the return series (like all market timing). Here, the event is triggered by a realised or expected change in the market value of the security. In turn, mimicking by outsiders after the event may have the potential to move the market in the short-run.

The notation for the modelling of abnormal returns and testing procedures largely follows Campbell, Lo, and MacKinlay (1997) (chapter 4). Event time (a counter) is denoted by τ, with the event date corresponding to $\tau = 0$. The estimation window is defined as the interval from $\tau = T_0 + 1$, to $\tau = T_1$, followed by the event window ($\tau = T_1 + 1$ to $\tau = T_2$). Also let $L_1 = T_1 - T_0$ and $L_2 = T_2 - T_1$ be the length of the estimation and event windows, respectively. In this paper, the event window comprises 20 trading days around the event, while the estimation window is made up of the 200 trading days before this. Therefore, $T_0 = -221$, $T_1 = -21$, and $T_2 = 20$. We compute excess returns in the most standard way, using a market model in the definition of expected returns: Letting $R_{i\tau}$ be the daily observed return on the stock, the returns-generating process for firm i is deemed well-approximated by:

$$R_{i\tau} = \alpha_i + \beta_i Rm_\tau + \varepsilon_{i\tau} \tag{1}$$

where we use the FTSE-250 index (to which a number of our firms actually belong) as a benchmark, since, as mentioned above, a significant size effect was found in Gregory, Matatko, and Tonks (1997). Parameters $\widehat{\alpha}_i$ and $\widehat{\beta}_i$ are estimated by OLS over the estimation window defined above, and excess returns $AR_{i\tau}$ are computed as:

$$AR_{i\tau} = R_{i\tau} - \widehat{\alpha}_i - \widehat{\beta}_i Rm_\tau \tag{2}$$

They are then averaged across events for every day in the event window, and average excess returns are cumulated to yield the familiar cumulative average abnormal return measure centered around the event date, denoted $\overline{CAR}(\tau_1, \tau_2)$:

$$\overline{CAR}(\tau_1, \tau_2) = \sum_{\tau=\tau_1}^{\tau_2} \left(\frac{1}{N} \sum_{i=1}^{N} AR_{i\tau} \right) \tag{3}$$

where N is the number of events and $T_1 < \tau_1 \leq \tau_2 \leq T_2$ (this is used to accommodate different sampling intervals within the event window, e.g. the post-event period only).

3 Results

3.1 Full dataset

Using the full dataset[6] a first run through the data yielded the following results: for director buys, abnormal returns are significantly negative in the twenty days before the net purchase, implying that directors purchase shares on average after a downward run in share prices (in the order of nearly 3%). Over the second half of the event window, the share price clearly recovers and abnormal returns are positive on most days, so that abnormal returns over the 20 days after the director's trade average a significant 1.9% (table 2).[7] The patterns are symmetrical in the case of director sells, though the magnitude of abnormal returns is lower. Directors typically sell shares after a run of positive price movements over twenty days of about 1.25%, and abnormal returns are predominantly negative after the directors' net sale, so that excess returns have averaged about 1.5% twenty days after the event (table 3).

The striking feature of these patterns is that on average, directors appear to be able to time the market in the short-run. These results are in contrast with those reported in Lakonishok and Lee (1998) for the US market, who (in their table 3) report average 5-day $CARs$ of 0.3% and 0.13% (sales and purchases respectively) for medium-sized stocks (therefore about one-third of what we find), and slightly *negative* $CARs$ in the case of large stocks (top three deciles of liquidity).[8]

The second noticeable fact is that larger stock price changes occur around purchases than around sales. These results are made even more striking given that sell trades are on average more than six times larger than buys. If trades of comparable size are considered, the effect is much more pronounced (see below: signal filters). There is a corresponding finding in papers on long-run excess returns following the trades of corporate insiders, such as Lakonishok and Lee (1998), or Jeng et al. (1999), but also in the literature studying the price impact of block trades (e.g. Chan and Lakonishok (1993)). This finding may be related to the regularly-made conjecture in the

[6]Events occurring in the first year of the data are dropped to leave enough days in the estimation window, leaving 1702 buys and 1268 sells.

[7]There are no significant abnormal returns outside this [−20 days,+20 days] window.

[8]The results in Jeng et al. (1999) are less comparable still as they are using a different methodology to assess the profitability of replicating portfolios made up of the actual quantities traded by insiders (therefore value-weighted) at a one-year horizon.

microstructure literature that large buy trades are likely to convey more information on average than large sells (see, Burdett and O'Hara (1987) or Allen and Gorton (1992)) There are more obvious liquidity reasons to sell than to buy.

As a first way of testing for the significance of these patterns, we report t-statistics for individual days and cumulative t-statistics over the whole of the event window in tables 2 and 3 (calculated as in Brown and Warner (1985), pp. 7 and 29). As is usual in event studies, the significance of the abnormal returns goes down as we move away from the event. The results for buy trades appear strongly significant for most days taken individually, and their overall significance is also strong. The significance of sell trades is less pronounced, though a window of at least six days around the event is clearly significant. We examine significance issues at length below using alternative, more robust testing methodologies. As a simple robustness check, an alternative specification is a market (as opposed to a market model) adjustment. The results from this simple methodology are reported in panel A of table 4. The $CARs$ for buy and sell signals are similar in magnitude to results in the previous tables.

From these patterns in prices, it is clear why previous work using monthly data found returns in the month containing the trade to be insignificantly different from zero: on average the change in price in the days before the director's trade largely cancels out the price change after the trade.

3.2 Robustness checks

3.2.1 Thin trading

There are a number of zero returns for some securities in the data because of thin trading (stale quotes). Besides the fact that this induces (or increases) autocorrelation, and could pose a problem for significance testing, it might also bias the estimated betas and therefore the abnormal return measures. To adjust for this, the betas were recalculated following the Scholes and Williams (1977) procedure, for the securities for which thin trading is an issue. While Scholes and Williams show that applying this adjustment to actively traded stocks leads to an overestimation of the beta coefficients, there is no clear-cut way of determining a cutoff point beyond which securities are deemed thinly traded. We sorted securities according to the number of zero returns in the data. The betas for the first three quartiles of securities in our sample were estimated in the usual way, while the above adjustment was applied to stocks in the bottom quartile.[9]

Estimated alphas and betas somewhat increased for these stocks and events,[10] abnormal returns estimates were not significantly changed by applying this correction: the results, presented in panel B of table 4 are that for buy trades, 20-day cumulative average abnormal returns stand at 1.92% (with cumulative t-stat from day 0 of 9.67) while for sell trades, 20-day average $CARs$ amount to 1.46% (with cumulative t-stat from day 0 of -6.96). Applying the Scholes-Williams adjustment to half of the securities instead of the ones in the quartile defined above produced very comparable results.

3.2.2 Outlier checks

Very large abnormal returns seemed to appear in a few cases, and we ascertained that our results were not driven by a few influential observations by identifying outliers using the methodology presented in Hadi (1992, 1994). This detected 19 cases of extreme returns after buy trades, and only 3 cases of extreme returns after sell trades. Removing them lowered average $CARs$ after buy transactions to 1.66% and left $CARs$ after sell trades virtually unchanged (at 1.48%). Therefore the impact of this correction, while not negligible in the case of buys, did not significantly alter our findings.

3.3 Significance issues

Besides the "bad model" problem mentioned above, the other major econometric issue in event studies is that the significance of the results itself can be affected by a number of factors. Standard t-tests may reject the null too often in the absence of abnormal performance, mostly because of biased standard errors, or because t-tests have low power. We now consider in turn which of these issues could be the most relevant for our study.

[9] Eight securities in that quartile of liquidity displayed a higher occurrence of both zero and missing returns, and a higher-order adjustment was applied (see Fowler and Rorke (1983)).

[10] From an average of 0.95 to 1.05 for buys events, and from 0.89 to 0.96 for sell events.

3.3.1 Variance changes

A first issue is that the variance of returns in the event window may be different from the variance in the estimation period, which violates the assumption of identically distributed excess returns. This is usually dubbed "event-induced change in variance", although the change of variance in our case could be caused by the trade itself as much as by an underlying company event (creating short term price movements and therefore to an extent triggering the trade).

We also use the test suggested by Boehmer, Musumeci, and Poulsen (1991) (BMP) shown to be robust to event-induced heteroskedasticity. Called the "standardized cross-sectional test", this involves computing the standardised residual on an event day as the estimated abnormal returns divided by their estimated standard deviation (assuming no heteroskedasticity), based on the residual variance from the estimation period ($\widehat{s}_i$), and the fact that they are prediction errors:

$$SR_{i\tau} = \frac{AR_{i\tau}}{\widehat{s}_i\sqrt{1 + \frac{1}{L_1} + \frac{(Rm_\tau - \overline{Rm})^2}{\sum_{\tau=T_0+1}^{T_1}(Rm_\tau - \overline{Rm})^2}}} \tag{4}$$

Then the standard deviation of these standardised excess returns is calculated cross-sectionally in the event period. The significance of the average standardised return is tested using the cross-sectionally estimated standard deviation. The (asymptotically unit normally distributed) test statistic becomes, for a given event day τ:

$$Z = \frac{\frac{1}{N}\sum_{i=1}^{N} SR_{i\tau}}{\sqrt{\frac{1}{N(N-1)}\sum_{i=1}^{N}\left(SR_{i\tau} - \sum_{i=1}^{N}\frac{SR_{i\tau}}{N}\right)^2}} \tag{5}$$

The multi-day version of which is simply constructed by summing the average standardised residual in the denominator above over the event window, divided by

$$\frac{\sum_{\tau=\tau_1}^{\tau_2}\overline{SR}_\tau}{\sqrt{\sum_{\tau=\tau_1}^{\tau_2}\widehat{s}^2(\overline{SR}_\tau)}} \tag{6}$$

Multi-day tests are presented in the second to last column of tables 2 and 3. The significance levels found remain high and consistent with the standard statistic.[11] Variance changes does not seem to be a major problem in these data.

3.3.2 Event clustering

The second issue we were concerned about is a possible clustering of events in the data. This is a problem for inference because the standard errors are not properly estimated if cross-sectional correlation between events is present in the sample. Previous studies such as Seyhun (1992) find quite strong clustering at the monthly level. More generally, there is almost always some event clustering, in the same way that returns on common stocks are never fully independent, though whether this is worth taking into account if the amount of clustering is not extreme (events common to all firms in the sample) has been debated in the econometric literature (see Campbell, Lo, and MacKinlay (1997), chapter 4, and Binder (1998) for recent overviews). From the simulation studies of Brown and Warner (1985), and Bernard (1987), the general conclusions that emerge are that using daily data makes clustering on a single date much less severe than when using monthly data. A simple examination of the data confirms this: even though the number of days for which two signals are recorded is quite large, it only very rarely goes beyond three signals in a day across firms. Given that there are 196 companies in the sample, this does not seem large. Bernard (1987) finds that diversification across industries should further mitigate the correlatedness problem. Our sample is highly diversified in this respect, since most industry sectors are present in our data. The nature of the event is another reason to believe that severe clustering should not be a problem: although there may be correlation in companies' fortunes, it is likely that directors' trades are mostly triggered by company-specific events.

For these reasons, and although partial overlap of event windows is present in the data, the problem is not reckoned to be severe. In the next section, we report results from a testing procedure which should be robust to partial event clustering, as well as non normality and autocorrelation.

[11] As another way of testing for this, the t-statistics in panel A of table 4 are computed using a contemporaneous benchmark instead of pre-event period data to estimate the variance of "normal" returns. This methodology is found in Brown and Warner (1985) to have comparable ability to detect abnormal returns at the daily level. The significance is not noticeably altered.

3.3.3 Non normality and time dependence

Two more issues to consider are that daily returns are not normally distributed for individual securities, and they display a (generally mild) degree of autocorrelation. In the econometric literature, Brown and Warner (1985) present an autocorrelation adjustment and conclude that "The benefits [from autocorrelation adjustment, in hypothesis testing] appear to be limited", while simulations (e.g. in Campbell and Wasley (1993)) show that daily abnormal returns collapse to normality when aggregated over portfolios of 100 stocks or more. However, the characteristics of sample stocks (not the most liquid securities) and the institutional (specifically, dealership) features of the London market may increase non normality and time dependencies: since these are smaller stocks, thin trading and high relative spreads may lead to price adjustment delays and a relatively high incidence of zero returns in the data.

To examine these issues together the possible event-clustering problem, a non-parametric (rank) testing procedure introduced by Corrado (1989), which does not rely on normality assumptions, was used. It is shown in simulations to be much more robust to thin trading problems and clustering of events. Campbell and Wasley (1993) for instance consider the test to be well-adapted to Nasdaq market data, and the trading system in operation at the London Stock Exchange over our sample period was a dealership system, explicitly modelled on Nasdaq in the mid-1980s, such that we would expect the data examined by Campbell and Wasley to share several features with our own.

The idea behind this statistic is to sort the series of abnormal returns over both the estimation and event windows and transform each observation into its respective rank: $k_{i\tau} = rank(AR_{i\tau})$, for $\tau = T_0 + 1, ..., T_2$. The rank statistic is the ratio of the mean deviation of the securities' day-0 ranks ($k_{i\tau}$) to the estimated standard deviation of the portfolio mean abnormal rank:

$$Z = \frac{\frac{1}{N}\sum_{i=1}^{N}(k_{i\tau} - E(k_i))}{\widehat{s}(k)} \tag{7}$$

Where $E(k_i)$ is the expected rank for security i, equal to $(L_1 + L_2 + 1)/2$. The denominator, $\widehat{s}(k)$, is the estimated standard deviation of the portfolio mean abnormal return rank, again over both estimation and event windows.

$$\widehat{s}(k) = \sqrt{\frac{1}{L_1 + L_2} \sum_{\tau=T_0+1}^{T_2} \left(\frac{1}{N} \sum_{i=1}^{N} (k_{i\tau} - E(k_i)) \right)^2}$$

The Corrado statistic is asymptotically unit normally distributed. In the case of multi-day event windows, the following statistic is formed:

$$\frac{\sum_{\tau=\tau_1}^{\tau_2} \overline{k}_\tau}{\sqrt{\sum_{\tau=\tau_1}^{\tau_2} \widehat{s}^2(\overline{k}_\tau)}} \tag{8}$$

Note that this testing procedure and the previous one complement each other as recent work by Cowan and Sergeant (1996) has questioned the robustness of the Corrado test under conditions of changes in variance around the event.

The estimated test statistics, for each day in the event window as well the cumulative version are presented for the buy and sell returns in the last column of tables 2 and 3. While lower, the significance levels shown by the Corrado test still confirm our finding of trading around short-term price movements. We are therefore confident in the robustness of our results.

3.4 Application of signal filters

3.4.1 Signal definitions

When deciding on which signals to consider, we are faced with a trade-off: on the one hand, it is obvious that in practical trading strategies, traders will apply filters using any relevant information to assess whether the trade is liquidity or information-motivated. The investment advisory services mentioned do not just report the trade as quickly as possible, they also claim to help investors interpret the signal. On the other hand, we want to stick to a limited number of signals which appear widely used to avoid the “data snooping” pitfall when testing for the profitability of a number of trading rules which can be defined by the researcher: by examining a large number of such rules we are bound to find that some of them will yield positive abnormal returns in a given data sample.

One obvious category of signals is based on the value of the director's trade. To illustrate, the Financial Times reports every week the details of trades of directors exceeding £10,000. Similarly, one of the conditions for a director's transaction to be considered "significant" by the Directus service is that its value exceeds £15,000. We will use this second value as a threshold for this first type of signals, keeping buy and sell trades with a value above it.

Alternatively, "contrarian" signals have been suggested as the ones likely to contain the most information. The US manager of a fund copying insider trades defines a strong signal of share undervaluation as a purchase during otherwise declining markets (or a sale during generally rising markets). This action can be interpreted as a bullish (bearish) signal regarding future stock returns. An additional reason for the contrarian trades to be informative is that in bearish (or agitated) markets, there are "flights to quality" towards blue-chip stocks, which depresses the price of smaller companies. Corporate insiders and investors at large may see this as the time to "pick up bargains". Lakonishok and Lee (1998) find that in aggregate, corporate insiders tend to be contrarian investors. We therefore define a second type of signal as a purchase (sell) observed when a moving average of returns in a window of 10 days before the event takes a negative (positive) value.[12]

A third category of signals which is regularly mentioned is based on the observation of repeated (clustered) trades within a short time interval, by (the same or different) insiders. This should provide a clear indication of how bullish a given insider is, or a consensus view among several insiders, in any case an unequivocal signal. We therefore define such a signal as any trade which was preceded by another one in the same stock at most 10 days before.

US studies regularly present abnormal return estimates depending on the type of insider (large shareholders, officers, directors) or their rank within the company, usually reporting that the closer the insider is to the top within the company, the stronger and more reliable the signal is.[13] But compared to US data, which includes various categories of insiders, our dataset is smaller and much more homogeneous,

[12]We tried other window lengths but this did not change results significantly.

[13]This is also pointed out to outside investors explicitly. As an example, Bloomberg News reported on 13 March 1998 that GM managers had been selling quite heavily, although "None of GM's four top executives had sold shares".

containing only directors' transactions.[14] Therefore, this type of signal is not of central relevance in our study.

In all, we therefore evaluate the profitability of three additional types of signals, besides the results obtained using the full dataset.

3.4.2 Results

It has been found in previous work, looking at longer holding periods, that excess returns were more pronounced when applying signal filters.[15] Our results, focusing entirely on short-term returns, are summarised in panel A of table 5. We computed all significance tests statistics introduced before for every one of these signals, and (as could be expected since we are now focusing on "stronger" signals), the significance levels found were higher than what was found for the full dataset with the exception of contrarian signals. We do not report them in detail as this would require 12 more tables, and only cumulative t-stats on the pre- and post-event window average $CARs$ (including the event day in the second case) are presented.

The main points from panel A of table 5 are the following: First, the asymmetry between excess returns around buy and sell trades is apparent for all signals. Second, the pattern in returns across signal definitions is remarkably similar across signals in the pre-event period (with the exception of clustered buy signals, as explained below) but it is different after the event: for the more profitable signals, the price seems to recover almost completely from the pre-event drop or increase, whereas for most other signals, this reversal is only partial. Third, different signals clearly have different strength or predictive contents over future returns: "contrarian signals" do not generate economically significant returns. Indeed, they deliver lower returns than the base case signal (full dataset). Focusing on trades larger than £15,000 (this means keeping only the top 3 deciles of signals, or 534 of the buys) seems to yield larger excess returns twenty days afterwards (2.8% instead of 1.9%), but the type of signal that clearly stands out (in terms of both pre-event price drop and post-event recovery) is the one based on clusters of buy signals. Here, the pre-event drop in prices reaches 6%, while the abnormal returns 20 days after the event are slightly

[14] Large shareholders in the UK are required to disclose their holdings when they reach 3 per cent of corporate equity, but not the individual trades.

[15] As opposed to considering different categories of firms according to their size.

larger than 4.5%. Sell signals, on the other hand, tend to be less clustered than buys, but there were still 174 "clustered" sells in the data (against 264 "clustered" buys.). The same signal definition applied to director sell trades yields abnormal returns that are only marginally larger than those of other sell signals: even though clustered sells are the strongest of sell signals, the asymmetry with the buys is more pronounced than for any other signal (excess returns 20 days after clustered buys are twice those after clustered sells). In the following section, we restrict our analysis to the clustered buy signals only.

3.5 Abnormal returns and trade size after directors' buy signals

Given the magnitude of the price movements around the directors' clustered buy signals, we investigated the pattern around this signal in more detail. In particular, we examined the distribution of individual event CARs according to trade value. In the previous section, we used the Directus definition (over £15,000) of a "significant" trade, though it is difficult to define small, medium-sized or large trades by just looking at the distribution of signals given its strong skewness. We somewhat arbitrarily define a small trade (for an individual investor) as belonging in the [0, £5000) interval, a medium-sized trade in turn being comprised in the [£5000, £70000) interval (£70000 being the 90th percentile), and classify all trades above this value as "large". The average $CARs$ for each category are as follows (panel B of table 5): the 20-day average CAR for the small director trades (607 events) is 1%. For the medium-sized ones (936 events), the same CAR is 2.6%, while for the large trades (159 events) it is only 1.6%. A test of the difference between mean CARs on small and medium trades and on large and medium trades yields highly significant results –with values of 205.07 and 51.41 respectively. In fact, focusing on the larger trades in the medium-sized category (156 events between £20,000 and £70,000) yields an average CAR of 3.7%!.

These returns are much more sizeable than the ones in Bettis, Vickrey and Vickrey (1997), who report two-week $CARs$ of 0.88% and 0.75% for buys and sells, respectively, but they are difficult to compare as they only examine block trades defined in the US as trades involving 10,000 shares or more, and for firms of all size deciles together.

If directors trade in medium sizes and these are the most significant signals of positive future abnormal returns, we should find, going back to the results of the previous section, that the clustered trades we found to be the most informative are generally medium-sized. Indeed, looking at the distribution by value of these clustered trades, this is exactly the case. Comparing those clustered trades to the ones in the [£5000, £70000] interval, we find that although their means appear quite different, the mean of the clustered trades is pulled up by a handful of large transactions. Once these (the top decile of trades by value, or 27 of them out of 254) are excluded, the means (£13,000 vs £17,000), medians (£7,700 vs £12,000) and standard deviations (15,600 vs 13,300) of the two data subsets become very similar (and much smaller than those of the full dataset).

Therefore medium-sized trades as a whole seem to predict higher returns than large ones. This evidence is consistent with the "stealth trading" hypothesis and findings of Barclay and Warner (1993) who report that the trades which seem to cause most of the total price changes in the price run-ups occurring before a takeover is disclosed to the market are concentrated in the medium-sized category. Our results cannot strictly be interpreted in the same way as it is not clear whether the post-event patterns in prices were or not partly caused by the mimicking by outsiders of the insider's trade, but one interpretation is that directors avoid trading in very large amounts around upcoming events which they expect will be accompanied by sizeable changes in the security's price. This would be revealing to the market that they are informed (especially in a dealership system such as the London Exchange where trading is not anonymous) as well as calling for regulatory scrutiny. Directors can make their trading less conspicuous by using one or several medium-sized trades.

3.6 Inclusion of transaction costs

For a realistic assessment of the actual profitability of these strategies, we correct for spread-induced transactions costs: because of the trade itself or an underlying company event triggering it, spreads may well widen at the time of the trade, removing most or all of the gross profitability. In earlier work, Seyhun (1986) and Bettis, Vickrey, and Vickrey (1997) adjust for these costs by using spread estimates from previous studies (therefore not contemporaneous), and that are averages over time and over portfolios of stocks (for small, medium-sized and large firms). Whereas returns have so far been computed from midquote to midquote, we now use daily bid and ask prices for each security to account for the fact that an outside investor would have to buy (sell) at the market-maker's ask (bid) and do the opposite twenty days later to profit from the price movement.[16] This removes, for each event, the two half-spreads that would have been incurred at the time of purchase or sale from the previously estimated cumulative abnormal returns (from $\tau_1 = 0$ to $\tau_2 = 20$):

$$NetCAR_i(\tau_1, \tau_2) = CAR_i(\tau_1, \tau_2) - (S_{i,\tau_1}/2P_{i,\tau_1} + S_{i,\tau_2}/2P_{i,\tau_2}) \tag{9}$$

where P and S stand for the prevailing closing price and bid-ask spread, respectively, on the day of the director's trade and twenty days later.

The results, presented in panel C of table 5, are that profitability seems on average wiped out by round-trip costs, implying that the higher returns following certain signals also seem compensated by higher spreads. Net returns are slightly negative except for three types of signals, and even then the highest average net return is 1.32% (as expected, for clustered buy trades). The net returns are therefore close to zero or negative, as they should in an efficient market. Even though the net average CAR after clustered buys appears sizeable, it only includes "microstructure" transaction costs and not estimates of "institutional" transaction costs (brokers' commissions). Therefore it seems reasonable to assume that net of commissions, profitability is close to zero, even without taking delays in imitating the trades into account. As a comparison with other results, Seyhun (1986) does not compute figures for net $CARs$ but considers them to be non-positive. To our knowledge, the only recent study which presents short-term $CARs$ with a form of transaction cost correction is Bettis, Vickrey and Vickrey (1997), which reports large negative two-week returns following trades of 10,000 or more shares. But they lump firms of all sizes together, and only

[16] These estimates of transactions costs may be seen as relatively conservative, since they are closing prices and research on patterns in the bid-ask spreads in the London Exchange has documented that they decline at the end of the trading day (presumably for inventory management reasons by market makers). In the case of a small number of very large trades, the mid-point to mid-point returns calculation is arguably preferable, since there is evidence that the execution prices of a sizeable proportion of block trades in London are negotiated and occur somewhere within the quotes or even at the mid-point (Reiss and Werner 1994). But the average director trade in our data is not very large by London Exchange standards, traditionally geared towards institutional investors, such that most of these trades would actually occur at or near the bid and ask quotes.

adjust for estimated transactions costs, averaged both over time and portfolios of stocks by size. Also and perhaps as a result, the negative $CARs$ they report are not statistically significant.

In this study, as well as in the previous ones using lower frequency data, the magnitude of net abnormal returns found after most signals is consistent with market efficiency. It remains to be seen how the excess returns found in the current paper could change the conclusions of previous studies which were using monthly data and were not able to statistically identify short-term excess returns. We leave this for future research but our results generally highlight the need to study events which may constitute market timing in the short as well as in the longer run and at different frequencies.

4 Summary and conclusion

Previous work examining the profitability of the trades of corporate insiders and of strategies mimicking these trades reported some evidence of long-run abnormal returns following these trades. However, the evidence regarding returns during the month containing the insider trade was either contradictory or unreliable. In this paper, we examined the patterns of security returns immediately around the trades of corporate insiders in the shares of their own company and assess the returns to strategies mimicking directors' trades in the days following the trade, after taking transactions costs into account.

We found patterns in abnormal returns in the days around a director's trade that are consistent with short-term market timing by directors and reported positive gross, but not net, abnormal returns to imitating some of the trades of directors. Therefore, although these patterns are statistically significant, their economic significance should not necessarily be a cause of great regulatory concern.

We also examined which types of trades may predict greater future returns. In line with previous work on this topic but also on the price effects of block trades, buy trades are followed by larger abnormal returns than sells. With respect to short-term returns, the strongest signals are the clustered ones, most of which consist of medium-sized trades. Medium-sized trades seem generally more informative than large ones, consistently with Barclay and Warner's (1993) "stealth trading" hypothesis. An interpretation is that informed traders may try to conceal their information by avoiding to trade in blocks, while transactions costs rule out a series of small trades as a strategy for accumulating a significant portfolio position.

Earlier studies on long-run returns in the UK found no evidence of abnormal returns in the month of the insider trade. The results in the current paper can explain this surprising finding. We found that the patterns in daily returns immediately around the insider trade are offsetting, so that in the earlier studies the price effects in the month of the trade were hidden by the lower frequency of the data used. The implication is that the evidence documented earlier on long-run abnormal returns needs to be adjusted upwards to take account of the price movement from the day of the insider trade.

References

Allen, F., and G. Gorton, 1992, "Stock Price Manipulation, Market Microstructure and Asymmetric Information," *European Economic Review*, 36, 624–630.

Barclay, M. J., and J. Warner, 1993, "Stealth Trading and Volatility (Which Trades Move Prices?)," *Journal of Financial Economics*, 34, 281–305.

Bernard, V., 1987, "Cross-Sectional Dependence and Problems in Inference in Market-Based Accounting Research," *Journal of Accounting Research*, 25, 1–48.

Bettis, C., J. Coles, and M. Lemmon, 2000, "Corporate Policies Restricting Trading by Insiders," *Journal of Financial Economics*, 57, 191–220.

Bettis, C., D. Vickrey, and D. W. Vickrey, 1997, "Mimickers of Corporate Insiders Who Make Large Volume Trades," *Financial Analysts Journal*, 53, 57–66.

Binder, J. J., 1998, "The Event Study Methodology Since 1969," *Review of Quantitative Finance and Accounting*, 11, 111–137.

Boehmer, E., J. Musumeci, and A. Poulsen, 1991, "Event-Study Methodology under Conditions of Event-Induced Variance," *Journal of Financial Economics*, 30, 253–272.

Brown, S., and J. B. Warner, 1985, "Using Daily Stock Returns (The Case of Event Studies)," *Journal of Financial Economics*, 14, 3–31.

Burdett, K., and M. O'Hara, 1987, "Building Blocks," *Journal of Banking and Finance*, 11, 193–212.

Campbell, C., and C. Wasley, 1993, "Measuring Security Price Performance Using Daily NASDAQ Returns," *Journal of Financial Economics*, 33, 73–92.

Campbell, J. Y., A. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton University Press, Princeton, N.J.

Chan, K. C., and J. Lakonishok, 1993, "Institutional Trades and Intraday Stock Price Behavior," *Journal of Financial Economics*, 33, 173–199.

Corrado, C., 1989, "A Nonparametric Test for Abnormal Security Price Performance in Event Studies," *Journal of Financial Economics*, 23, 395–395.

Cowan, A., and A. Sergeant, 1996, "Trading Frequency and Event Study Test Specification," *Journal of Banking and Finance*, 20, 1731–1757.

Dennert, J., 1991, "Insider Trading," *Kyklos*, 44(2), 181–202.

Finnerty, J., 1976, "Insiders and Market Efficiency," *Journal of Finance*, 31, 1141–48.

Fishman, M., and K. Hagerty, 1995, "The Mandatory Disclosure of Trades and Market Liquidity," *Review of Financial Studies*, 8(3), 637–676.

Fowler, D., and H. Rorke, 1983, "Risk Measurement When Shares Are Subject to Infrequent Trading: Comment," *Journal of Financial Economics*, 12(2), 279–283.

Garfinkel, J., 1997, "New Evidence on the Effects of Federal Regulations on Insider Trading: The Insider Trading and Securities Fraud Enforcement Act (ITSFEA)," *Journal of Corporate Finance*, 3(2), 89–111.

Gregory, A., J. Matatko, and I. Tonks, 1997, "Detecting Information from Directors' Trades: Signal Definition and Variable Size Effects," *Journal of Business Finance and Accounting*, 24(3-4), 309–342.

Gregory, A., J. Matatko, I. Tonks, and R. Purkis, 1994, "UK Directors' Trading: The Impact of Dealings in Smaller Firms," *Economic Journal*, 104, 37–53.

Hadi, A., 1992, "Identifying Multiple Outliers in Multivariate Data," *Journal of the Royal Statistical Society*, 54, 761–771.

———, 1994, "A Modification of a Method for the Detection of Outliers in Multivariate Samples," *Journal of the Royal Statistical Society*, 56, 393–396.

Hiller, D., and A. Marshall, 1998, "The Timing of Directors' Trades in the UK and the Model Code," *Journal of Business Law*, 24, 454–467.

Hu, J., and T. Noe, 1997, "The Insider Trading Debate," *FRB Atlanta Economic Review*, 83(4), 34–45.

Huddart, S., J. Hughes, and C. Levine, 1999, "Public Disclosure of Insider Trades, Trading Costs, and Price Discovery," Duke University Fuqua School of Business Working Paper.

Jaffe, J., 1974, "Special Information and Insider Trading," *Journal of Business.*, 47(3), 410–28.

Jeng, L., A. Metrick, and R. Zeckhauser, 1999, "The Profits to Insider Trading: A Performance-Evaluation Perspective," *NBER Working Paper*, No 6913.

Kabir, R., and T. Vermaelen, 1996, "Insider Trading Restrictions and the Stock Market: Evidence from the Amsterdam Stock Exchange," *European Economic Review*, 40, 1591–1603.

King, M., and A. Röell, 1988, "Insider Trading," *Economic Policy*, 7, 163–193.

Kyle, A., 1985, "Continuous auctions and insider trading," *Econometrica*, 53(4), 1315–35.

Lakonishok, J., and I. Lee, 1998, "Are Insiders' Trades Informative?," NBER Working Paper No. 6656.

Leland, H., 1992, "Should Insider Trading Be Prohibited?," *Journal of Political Economy*, 100(4), 859–887.

London Stock Exchange, 1996, "Guidance on the Dissemination of Price Sensitive Information," (Stock Exchange Publication).

——— , 1998, "Continuing Obligations Guide," (Stock Exchange Publication).

Loughran, T., and J. Ritter, 2000, "Uniformly Least Powerful Tests of Market Efficiency," *Journal of Financial Economics*, 55(3), 361–389.

Manne, H., 1966, *Insider Trading and the Stock Market*. Free Press, New York, NY.

Meulbroek, L., 1992, "An Empirical Analysis of Illegal Insider Trading," *Journal of Finance*, 47(5), 1661–1699.

Pope, P., R. Morris, and D. Peel, 1990, "Insider Trading: Some Evidence on Market Efficiency and Directors' Share Dealings in Great Britain," *Journal of Business Finance and Accounting*, 17(3), 359–80.

Reiss, P., and I. Werner, 1994, "Transactions Costs in Dealer Markets: Evidence from the London Stock Exchange," in Andrew Lo, ed., Industrial Organization and Regulation of the Securities Industry (NBER-University of Chicago Press).

Scholes, M., and J. Williams, 1977, "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5, 309–327.

Seyhun, H., 1986, "Insiders' Profits, Costs of Trading and Market Efficiency," *Journal of Financial Economics*, 16, 189–212.

Seyhun, N., 1992, "Why Does Aggregate Insider Trading Predict Future Stock Returns?," *Quarterly Journal of Economics*, 107, 1303–1331.

Shleifer, A., and R. W. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance*, 52(1), 35–55.

Table 1: **Descriptive statistics for all trades and net traded values**

Panel A: All trades						
	N	*p10*	*Median*	*p90*	*Mean*	$\widehat{\sigma}$
Buys	2,558		6,650		66,068.4	652,503.5
Sells	1,841		32,600		343,068.9	3,833,629
Total	4,399					

Panel B: Net Trades						
	N	*p10*	*Median*	*p90*	*Mean*	$\widehat{\sigma}$
Buys						
1986	38	1,125	11,625	695,600	156,999	433,530
1987	211	1,470	8,600	140,040	82,303	289,228
1988	233	1,568	6,500	100,800	195,764	1,569,983
1989	238	1,756	8,600	74,400	32,525	94,201
1990	258	2,060	9,369	69,550	78,430	695,286
1991	218	1,740	6,680	65,000	114,496	1,272,066
1992	296	2,020	7,323	40,750	57,061	402,183
1993	170	2,333	8,806	52,400	42,241	230,441
1994	225	1,330	7,488	32,250	22,627	93,567
Overall	1,887	1,750	7,950	70,000	80,044	776,710
Sells						
1986	33	5,400	48,000	553,500	257,049	750,182
1987	241	5,742	33,500	673,460	343,950	1,230,274
1988	226	4,622	27,150	325,440	166,815	465,732
1989	206	6,160	37,860	373,500	958,938	10,753,229
1990	169	6,440	48,200	647,500	424,028	2,168,552
1991	217	6,440	28,176	742,500	577,511	3,019,809
1992	164	5,396	27,593	391,500	222,227	717,038
1993	180	7,504	30,419	490,000	263,963	845,109
1994	86	6,844	20,865	148,750	70,632	167,798
Overall	1,522	6,150	30,675	475,517	403,173	4,229,704

Panel C: Number of trades in each firm decile				
Firm decile	*Buy trades*	*Sell trades*	*Net buy trades*	*Net sell trades*
1	300	125	222	101
2	246	251	193	225
3	323	198	217	180
4	339	278	242	230
5	160	153	134	125
6	253	225	190	179
7	286	200	206	162
8	249	91	191	80
9	218	132	160	114
10	184	188	132	126
Total	2558	1841	1887	1522

The table reports descriptive statistics on all trades by value in pounds Sterling (panel A) and net traded value (panel B). The net traded value is used because on some days more than one trade occurred on a given day. "N" is the total number of trades. The next columns give some information on the distribution of trades by value, where "p10" and "p90" are the tenth and ninetieth percentiles, respectively. Panel C reports the total number of trades for all deciles of firms by market value, where decile 1 is made up of the largest firms and decile 10 of the smallest ones.

Table 2: **Abnormal returns and significance tests (Buy trades)**

Days	$\overline{AR}$	t	$\overline{CAR}$ $(-20, 20)$	$Cumul.\ t$	$\overline{CAR}$ $(0, 20)$	$Cumul$ BMP	$Cumul.$ $Corrado$
-20	-0.000932	-2.092	-0.00093	-2.092		-1.978	-0.889
-15	-0.001168	-2.619	-0.00484	-4.433		-4.166	-2.35
-10	-0.000919	-2.062	-0.01004	-6.794		-6.437	-3.929
-8	-0.001026	-2.303	-0.01262	-7.851		-6.340	-4.691
-6	-0.002484	-5.573	-0.01638	-9.488		-6.750	-4.900
-4	-0.003491	-7.832	-0.02166	-11.787		-8.428	-5.913
-3	-0.002080	-4.666	-0.02374	-12.555		-9.450	-6.448
-2	-0.002496	-5.599	-0.02624	-13.504		-9.880	-6.822
-1	-0.002317	-5.199	-0.02855	-14.325		-7.504	-6.802
0	0.001514	3.397	-0.02704	-13.238	0.00151	-7.424	-6.178
1	0.002747	6.163	-0.02429	-11.620	0.00426	-9.310	-5.339
2	0.001999	4.485	-0.02229	-10.430	0.00626	-9.272	-4.648
3	0.001729	3.878	-0.02056	-7.206	0.00799	-8.130	-4.113
4	0.000884	1.983	-0.01968	-9.418	0.00887	-8.708	-3.660
6	0.001022	2.292	-0.01724	-7.445	0.01131	-7.430	-2.817
8	0.001279	2.869	-0.01539	-6.412	0.01316	-5.623	-2.397
10	0.000400	0.897	-0.01392	-5.609	0.01463	-5.922	-1.698
15	0.000351	0.787	-0.01049	-3.924	0.01806	-3.510	-0.554
20	0.000868	1.947	-0.00892	-3.126	0.01963	-2.992	-0.299

The table reports abnormal returns on selected days around a director's buy trade. Column 2 lists average daily abnormal returns computed from equation 2. Columns 4 and 6 list average cumulative abnormal returns from equation 3 from the beginning of the event window and from the day of the trade, respectively. T-statistics on individual days' average abnormal returns (column 3) and on average CARs (column 5) are computed as in Brown and Warner (1985), p. 7 and 29, respectively. Column 7 presents the Boehmer, Musumeci and Poulsen test statistic on the cumulative abnormal returns computed as in equation 6. Column 8 reports the multi-day version of the non-parametric test statistic of Corrado from equation 8.

Table 3: **Abnormal returns and significance tests (Sell trades)**

Days	$\overline{AR}$	t	$\overline{CAR}$ $(-20, 20)$	$Cumul.$ t	$\overline{CAR}$ $(0, 20)$	$Cumul.$ BMP	$Cumul.$ $Corrado$
-20	0.000561	1.215	0.00056	1.215		1.917	0.971
-15	-0.000296	-0.640	0.00040	0.352		0.491	1.656
-10	0.000485	1.050	0.00312	2.040		3.264	3.201
-8	0.000467	1.011	0.00394	2.363		3.763	3.364
-6	0.000970	2.100	0.00591	3.305		4.676	4.131
-4	0.001321	2.860	0.00812	4.267		5.909	4.851
-3	0.001327	2.873	0.00945	4.824		7.048	5.342
-2	0.001124	2.433	0.01057	5.253		8.140	5.751
-1	0.001755	3.800	0.01233	5.970		8.048	6.213
0	-0.000099	-0.214	0.01223	5.779	-0.000099	7.345	5.912
1	-0.001653	-3.580	0.01058	4.883	-0.001752	7.828	5.074
2	-0.001585	-3.432	0.00899	4.060	-0.003337	6.433	4.259
3	-0.001140	-2.469	0.00785	3.471	-0.004477	5.900	3.926
4	-0.000183	-0.395	0.00767	3.322	-0.004660	5.915	3.918
6	-0.001447	-3.134	0.00497	2.071	-0.007361	3.480	3.095
8	-0.001110	-2.404	0.00331	1.332	-0.009017	2.224	2.353
10	-0.000887	-1.921	0.00163	0.634	-0.010700	0.876	1.898
15	-0.000101	-0.219	0.00078	0.280	-0.011553	0.511	2.051
20	-0.000541	-1.172	-0.00232	-0.786	-0.014654	-1.388	1.342

The table reports abnormal returns on selected days around a director's sell trade. Column 2 lists average daily abnormal returns computed from equation 2. Columns 4 and 6 list average CARs from equation 3 from T_1, the first day in the event window and the day of the trade, respectively. T-statistics on individual days' average abnormal returns (column 3) and on cumulative abnormal returns (column 5) are computed as in Brown and Warner (1985), p. 7 and 29, respectively. Column 7 presents the Boehmer, Musumeci and Poulsen test statistic on the cumulative abnormal returns computed as in equation 6. Column 8 reports the multi-day version of the non-parametric test statistic of Corrado from equation 8.

Table 4: **20-day CARs using a market-index and a thin trading adjustment**

Signal definition	*N°obs*	$\overline{CAR}(-20,-1)$	*cumul. t*	$\overline{CAR}(0,20)$	*cumul. t*
Panel A: CARs after market adjustment					
Buys	1702	-3.08%	-12.25	1.78%	8.02
Sells	1268	1.26%	6.76	-1.32%	-8.55
Panel B: CARs after thin trading adjustment					
Buys	1702	-2.78%	-14.01	1.92%	9.67
Sells	1268	1.24%	5.92	-1.46%	-6.96

The first panel in the table reports 20-day average CARs when estimated using a simple market index adjustment. Panel B reports the same CARs after market-model coefficients are adjusted for thin trading using the Scholes and Williams (1977) procedure.

Table 5: **CARs after various types of signals and after transactions costs adjustments**

Panel A: average CARs around different signals

Signal definition	$N°obs$	$\overline{CAR}(-20,-1)$	*cumul. t*	$\overline{CAR}(0,20)$	*cumul. t*
Buy signals					
All Buys	1702	-2.85%	-14.32	1.96%	9.84
Large Buy	534	-3.32%	-9.24	2.8%	7.79
"Contrarian" Buy	835	-3.41%	-12.75	1.01%	3.79
Clustered Buy	264	-6.00%	-11.43	4.52%	8.47
Sell signals					
All sells	1268	1.23%	5.97	-1.46%	-7.09
Large sell	1042	1.13%	5.04	-1.66%	-7.37
"Contrarian" sell	716	1.60%	6.00	-1.11%	-4.14
Clustered sell	174	1.21%	2.44	-2.41%	-4.85

Panel B: average CARs and buy trade size

Signal definition	$N°obs$	$\overline{CAR}(-20,-1)$	*cumul. t*	$\overline{CAR}(0,20)$	*cumul. t*
Small buy	607	-2.1%	-6.16	0.98%	12.86
Medium-sized buy	936	-2.62%	-9.86	2.59%	9.72
Large buy	159	-4.54%	-7.85	1.57%	2.71

Panel C: average post-event CARs after transactions costs adjustment

Signal definition	*Net* $\overline{CAR}(0,20)$	*Signal definition*	*Net* $\overline{CAR}(0,20)$
Buy signals		**Sell signals**	
All buys	-0.66%	All sells	-0.55%
Large buys	0.285%	Large sells	-0.23%
Contrarian buys	-1.68%	Contrarian sells	-0.86%
Clustered buys	1.32%	Clustered sells	0.366%

The table reports average cumulative abnormal returns computed from equation 3 around various types of signals (panels A and B) and net of transactions costs (panel C). In panel A, results are reported for both buy and sell signals for the full dataset, for large trades (defined as exceeding £15,000 in value), for "contrarian" trades (buys in bearish markets, sells in bullish markets) and "clustered" trades (preceded by another trade in the ten days or less before). Panel B reports average CARs after directors' buy trades of different sizes. Small buys are defined as the ones having value of less than £5,000. Medium-sized trades belong in the [£5000, £70000) interval while any trade of more than £70,000 is a large one Panel C reports average CARs after removing "round-trip" transaction costs (the half-spreads incurred at the time of trading) as in equation 9.