

Matching in size: How market impact depends on the concentration of trading

Ilija I. Zovko

Aspect Capital, London
ilija.zovko@aspectcapital.com
zovko@algo-quants.com

In this paper we provide empirical evidence that large orders suffer greater market impact when matched with multiple trade counterparts than when matched with one or a small number of counterparts. This provides some support for using large in scale (LIS) venues to improve execution performance. Filling an order on a venue in which a trader is able to minimise the number of counterparts to match with, all other things being equal, is expected to suffer less price impact than splitting the order and matching with many counterparts on anonymous venues.

To reach this conclusion we introduce the concept of *concentration of trading* and study its effect on market impact. Trading is *concentrated* if a disproportionate volume share of buying or selling in a period of time is done by a few trading firms. It is *dilute* (i.e. not concentrated), if trading is spread across many firms. Buyers as well as sellers can be concentrated and will typically be a consequence of a small number of firms executing large orders.

Using a unique (but unfortunately somewhat dated) dataset from the London Stock Exchange (LSE) we are able to reconstruct the degree of concentration in buying and selling, and investigate its influence on resulting price moves. This is made possible by the record of trade matching details identifying LSE member firms participating in each trade.

Price impact, of course, is to a large degree influenced by the details of order routing, as implemented by, say an SOR, a Smart Order Router. The SOR decisions themselves can be influenced by the size of the order or the trading participation levels. Therefore, in order to isolate the effect of concentration, in our analysis we control for orderflow effects. We take into account the most common orderflow characteristics, such as the number of buyer or seller initiated trades and volumes (i.e., passive and aggressive orders), and the total number of firms buying or selling.

After adjusting for the effects of orderflow on price impact, we still find that concentrated trading results in larger impact than dilute trading.

As a practical and direct test of our findings, we show that large concentrated orders suffer noticeably less market impact in circumstances when they are matched with similarly concentrated trade counterparts on the other side of the market. This suggests that opportunistically trading large orders in size may improve price slippage.

The results presented in this paper imply that while the order routing decisions remain a key determinant of execution outcomes, it is similarly the *number of counterparts* one trades with that influences performance. The fewer the number of counterparts, all other things being equal, the better the performance.

Algo execution | Market impact | LIS

1. Introduction

With the introduction of MIFID II regulation, there has been an emergence of trading venues dedicated to matching in size. I.e., venues with mechanisms facilitating finding counterparts to fill an order with, in as few number of clips, or ideally in its entirety.

While intuitively sound, the performance benefits of matching in size, in comparison to splitting an order into clips and working it on anonymous venues, are less obvious.

We address this question by introducing the concept of *concentration of trading* and investigating its influence on price returns, and consequently execution performance. We start with the trivial observation that the total number of shares bought must be equal to the total number of shares sold. It is how those shares are *partitioned* between the individual traders that determines trading concentration.

Concentrated buying describes a situation where a large proportion of shares are bought by a single or a few exchange member firms. On the other hand, *dilute buying* – a term we use to describe a lack of concentration – describes a situation where the shares bought are spread over a larger group of firms, all contributing comparable amounts to the total. Concentrated and dilute selling are defined analogously for sellers in the market. It is straightforward to see that, say, a large buy order will contribute to the concentration of the buying side of the market.

In the paper we present evidence that concentrated trading produces larger price impact than dilute trading. This means that a large order, when being matched with multiple small orders, adversely moves the market and suffers price slippage. Correspondingly, we find that when a large order is matched with a similarly large order, price impact is significantly reduced.*

As mentioned, concentration is closely related to order size – large orders tend to result in concentrated trading. Order size, on the other hand, tends to influence order routing and orderflow. It is well known that aspects of orderflow strongly affect price impact, with aggressive and passive order placement being the most exemplary one (1–12). Therefore, in order to investigate the effect of concentration on price impact, we must take into account the known price effects

*In the paper, we use the term 'order' to denote the aggregate buying or selling by a single exchange member firm. We do not have data granular enough to recover orders originating from different trading desks, or strategies. However, one can argue that for the purposes of execution, it is the aggregate orderflow that is important anyway.

Executive summary

We show that filling an order with a large number of distinct counterparts incurs additional market impact, as opposed to filling the order with a small number of counterparts. For best execution, therefore, it may be beneficial to opportunistically fill orders with as few counterparts as possible in Large-in-scale (LIS) venues.

This article introduces the concept of concentrated trading, a situation that occurs when a large fraction of buying or selling in a given time period is done by one or a few traders, for example when executing a large order. Using London Stock Exchange data, we show that concentrated trading suffers price impact in addition to impact caused by (smart) order routing. However, when matched with similarly concentrated counterparts on the other side of the market, the impact is greatly reduced. This suggests that exposing an order on LIS venues is expected to result in execution performance improvement.

of order routing and orderflow characteristics:

- the number and GBP volume of aggressive and passive trades, and
- the number of firms buying vs. firms selling.

Concentration is not an instantaneous market property. Like volatility, it is an aggregate property of trading in a time window, for example a day or an hour. We have chosen to perform the analysis on daily intervals. However, the choice of aggregation remains an open question. Instead of calendar time, one can use an aggregation scheme resulting from a fixed number of trades (so called trade time), or resulting from a fixed number of shares traded (volume time). In addition to the daily aggregation, we have also ran the analysis on hourly and intervals resulting from a fixed number of trades. The main conclusions from the daily analysis are roughly borne out on the other tested aggregation schemes. While interesting in its own right, a complete study of the differences between timescales is out of scope of our current work.

The analysis is based on a somewhat dated trades and quotes dataset from the on-book trading of the London Stock Exchange (LSE) between 2000 to 2002. What makes this dataset unique is that each trade record contains codes identifying the stock exchange member firms that participated in the trade.[†] While a member firm can make a trade both in a principal (for their own account) or agent capacity, the data shows that in more than 95% of trades the member firms act as principal to the transaction.

We base the analysis on 32 months (674 full trading days) on 46 stocks. Excluding from the dataset days in which a particular share traded less than 500 times, yields a sample of roughly 20 million trades. For this set we have all the trade and quote details, including trade matching details identifying the participants in each trade. The data period is from the beginning of May 2000 until the end of December 2002. During this period, there were on average between 33 to 72 member firms trading on the on-book market of the LSE each day. The average number of firms trading increases during the period and in general there are more firms trading the more active stocks. The typical number of trades is several thousand per day for active stocks. LSE is open for trading from 8:00 to 16:30 but we have discarded data from the first and last half hour of trading to avoid possible opening or closing anomalies.

2. Concentration of trading

Concentrated trading is a situation where a single or a few firms contribute in a disproportionately large proportion to buying or selling in a time period. For example, a firm executing an order with a high percentage of volume (POV) algorithm will contribute to trading concentration. Alternatively, concentration can also be a consequence of a disproportionately large order being worked in the market. Figure 1 illustrates the principle question of the paper: do the three idealised examples have different impact on price movement, while taking into account the different orderflow characteristics expected from trading large orders.

We quantify concentration using metrics of inequality between the volume fractions by which each member firm contributes to the total trade volume. Denote by w_i^b the fraction of shares bought by firm i in a trading session, i.e., the sum of GBP trade sizes in which firm i participated as a

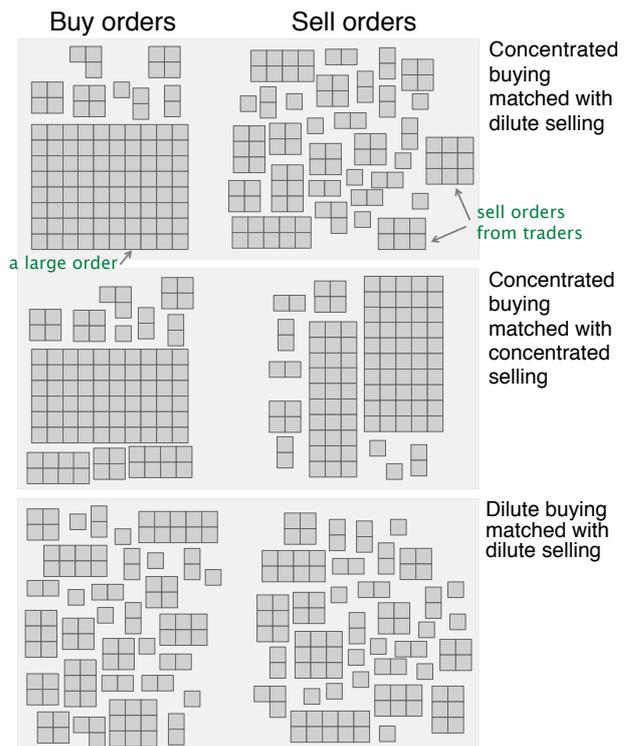


Fig. 1. Three stylised examples of trading where different levels of concentration and dilution may result in different price impact of large orders. Each small square represents a share or a “lot”, which are then grouped by the firm that bought or sold them. In the top example, buying was concentrated as a consequence of a large order on the market. Selling, on the other hand was dilute, with a large number of orders contributing to the aggregate selling. In the middle example, both the buying and selling are concentrated. The bottom example shows trading between dilute buyers and sellers.

In the paper we show that large orders causing concentration suffer adverse price impact, even after taking into account the impact of the orderflow generated by such orders. This suggests there are performance advantages to match a large order with similarly large orders, as opposed to splitting it up and filling with multiple counterparties.

buyer, divided by the total traded volume. Analogously, w_i^s is the fraction of shares sold by firm i (computed analogously). A member firm i , which in a given session both bought and sold, will have w_i^s and w_i^b non-zero.

Volume fractions w_i are positive and sum up to 1, defining a probability distribution.[‡] Hence, we can quantify concentration using any statistical measure for the inequality of a probability distribution. In order to make our results robust to the choice of the measure, we initially use two measures of inequality, the Gini index G and the entropy E .

The sample Gini index can be computed as the half of the relative mean absolute difference between all N samples x_i (13, 14)

$$G \equiv \frac{1}{2} \frac{\sum_{i,j} |x_i - x_j| / N^2}{\sum_i x_i / N} \quad [1]$$

which for volume fractions w_i can be simplified to

$$G \equiv \frac{1}{2N} \sum_{i,j} |w_i - w_j|. \quad [2]$$

The Gini index G_s for the concentration of selling is computed from the sell volume fractions, while G_b for the concentration of buying from the buy fractions.

[‡]We omit the superscript denoting the market side when the expression applies to both sides.

[†]Member firm codes do not identify the firms by name and are scrambled monthly and across stocks for anonymity. This makes it impossible to track firms in time or investigate the properties of their orders with good statistics.

For fully concentrated trading, in which all buys or sells are done by a single firm (having a 100% participation ratio), the Gini index has a maximum value equal to 1. For fully dilute trading, in which all firms participate an equal amount to the total trade volume, it takes its minimum value of 0.

The entropy based measure for concentration is defined as

$$E \equiv 1 - \frac{\sum_i w_i \cdot \log(1/w_i)}{\log N}, \quad [3]$$

where we have changed the sign and introduced the normalisation factor of $\log N$ to the usual entropy expression. With these changes, we get the intuitive interpretation that large concentration is associated with large values of the metric, consistent with the Gini metric: for fully concentrated trading the entropy measure is equal to 1, while for fully dilute it is 0.

Finally, we define the *imbalance* in trading concentration between the buying and selling as

$$\begin{aligned} \delta G &\equiv G_b - G_s \text{ and} \\ \delta E &\equiv E_b - E_s. \end{aligned} \quad [4]$$

If the buying is more concentrated than selling, the imbalance will be positive.

One might immediately wonder how often do we see trading situations in which either buying or selling is significantly concentrated. Looking at all stocks in our LSE dataset, on almost half of the days there was at least one firm that traded more than a quarter of all the buying or selling on the day. So a fairly frequent occurrence.

3. Concentration and price returns

We have decided to estimate the effects of concentration and orderflow on price returns using a linear model. After trying out a few nonlinear forms, we did not find clear evidence of nonlinearities. The extremes of the dataset do deviate from the linear assumption, but the noise levels are large and do not warrant additional complexity. We have also found that the Entropy and Gini metrics bear same conclusions, so will restrict the discussion to only using the Entropy metric.

Therefore we will be explaining the daily price returns δP_t by expressions of the form

$$\begin{aligned} \delta P_t &= (\text{concentration imbalance})_t \\ &+ (\text{order routing imbalances})_t \\ &+ (\text{error term})_t. \end{aligned} \quad [5]$$

Price returns are calculated as the percent difference of the volume weighted average price (VWAP) of the last 10% and first 10% of trades on the day[§]. We normalised the returns by the overall market return, subtracting from each daily stock return the FTSE100 return for the day. To make sure that overall stock moves do not impact conclusions, mean return per stock over the entire period is subtracted from the daily stock return. The variance of price returns across stocks is left unchanged.

The order routing effects are captured by the imbalances between buying and selling in

- numbers of aggressive orders M_t ,
- GBP volume of aggressive orders V_t , and
- numbers of firms N_t .

[§] Recall that we discarded 30 minutes of data after the opening auction, and 30 minutes before the close.

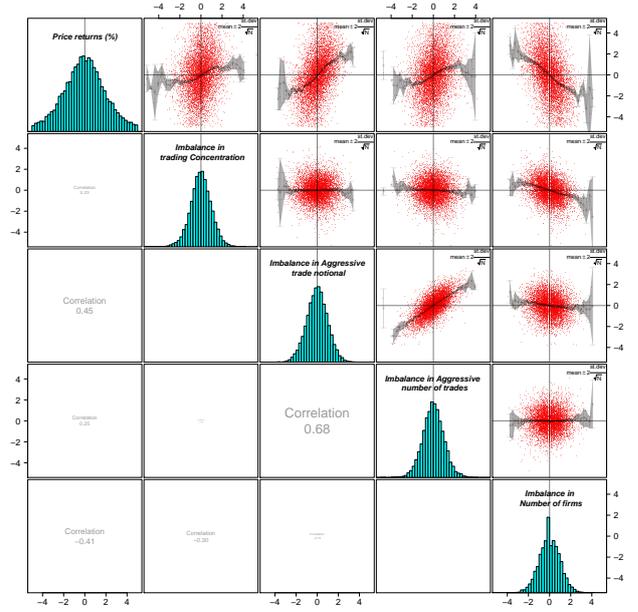


Fig. 2. Scatterplots, histograms and correlations showing the dependency between price returns and imbalances in concentration and orderflow. Matrix diagonal shows the histograms, roughly normally distributed with mean 0 and standard deviation 1. The exception is the distribution of price returns, which in fact are not normalised - price returns are left as percentages. The triangle above the diagonal shows simple scatterplots of daily values (in red), together with a simple binned conditional mean estimate. Values in the matrix below the diagonal are the values of the correlation coefficients between the corresponding variables with the font size proportional to the absolute value of the coefficient.

The charts show a fairly correlated set of variables, which are all in one way or another related to trade volume. It is for this reason that we will have to take special care to properly account for the effects of orderflow in revealing the role of concentration in price returns.

The imbalances are computed using the usual convention that buying gets a positive sign, and using the common normalising form

$$\delta x = \frac{x_b - x_s}{x_b + x_s} \quad [6]$$

where b and s designate buying and selling. The full order-book data we use allows us to calculate these order routing variables precisely.

In summary, for each stock and trading day, we compute the market normalised price return δP_t and the concentration imbalance between buyers and sellers δE_t using Eq.4. Using Eq.6, we compute the imbalance between the number and GBP notional of aggressive buy orders and aggressive sell orders, δM_t and δV_t respectively. And similarly the imbalance in the number of buyers and sellers δN_t . To combine data from different stocks, each imbalance variable is further normalised by the standard deviation for the corresponding stock.

Figure 2 shows a matrix of dependencies between the price returns and the imbalance metrics. The upper matrix triangle shows scatter plots (in red) and a simple conditional mean estimate as the solid gray line. The level of correlation (estimated coefficients shown in the lower matrix triangle) is high since none of the orderflow variables can vary fully independently of the others. If a new buyer starts trading, changing the imbalance in the number of firms, this may in turn change the ratio of aggressive sell and buy orders in the market. Depending on the order size, it may also change the concentration of buying.

Trading concentration, on the other hand, is fairly inde-

Number of samples = 15491	Coef.	Error	R_S^2	R_P^2
δE (trading concentration imbalance)	24.9	1.3	0.05	0.02
δN (imb. in num. of firms trading)	-61.2	1.4	0.17	0.16
δV (imb. in aggressive notional)	81.8	1.8	0.20	0.12
δM (imb. in aggressive trade cnt.)	-3.0	1.8	0.06	0.00
Overall	$R^2 = 0.33$			

Table 1. Model fit showing the magnitude and significance of concentration and order routing imbalances on price returns (in basis points). The impact of concentration on the price is adverse, in that positive imbalances corresponding to concentrated buying are associated with positive price returns. The impact of order routing is as expected: an imbalance towards aggressive orders adversely affects the price, as well as the fewer firms trading on one side of the market, the more adverse price impact they can expect, reflecting their order size. Only the imbalance in the number of aggressive trades is not significant, with a similar effect better captured by aggressive notional imbalance.

The effect of concentration is comparable in magnitude (Coef \approx 25 bps) to the impact of routing (Coefs \approx 80 and -60 bps), but the explanatory power of concentration imbalance is less than that of order routing ($R_P^2 \approx$ 0.02 vs. 0.16 and 0.12). In other words, impact predictions based on trading concentration are equally important, but are less reliable and more noisy. R_P^2 is the partial R^2 of the selected variable. It expresses how much the variable explains price returns in addition to the other three variables. R_S^2 is the value of R^2 in a regression where only the selected variable is present in the regression. It expresses how much the variable on its own explains price returns.

pendent and only mildly anti-correlated with the imbalance in the number of firms, e.g., the smaller the number of sellers, the more concentrated selling is. The fact that concentration and orderflow are rather independent is indicative of fairly sophisticated SORs at work. Even when there is a large order being worked (resulting in concentrated trading), the order routing does not to create biases in orderflow.

Ultimately the model to estimate is

$$\delta P_t = \alpha \cdot \delta E_t + \beta \cdot \delta M_t + \gamma \cdot \delta V_t + \tau \cdot \delta N_t + \epsilon_t, \quad [7]$$

where ϵ_t are $\mathcal{N}(\mu, \sigma)$ Normal residuals. We estimate the parameters via a standard OLS on a merged dataset for all stocks, containing 15491 samples corresponding to day sessions. In doing this, we have limited ourselves to sessions containing more than 500 trades and have removed sessions in which the stock price changed by more than 5 percent (to limit the effect of exogenous or news events). The number of samples we removed is around 5% and in reality does not materially change the conclusions. In terms of other robustness checks, we performed the analysis in distinct yearly periods as well as separately for each stock, broadly corroborating the conclusions.

Model estimates and corresponding errors are collected in table 1. The model explains about 30% of overall price return variation (model $R^2 = 0.33$). However, if we focus on large returns and estimate the model on days where there was a price change larger than 1%, the explanatory power increases to 40% ($R^2 = 0.40$). If we further restrict the sample to include only returns between 4% and 5%, the explanatory power increases further to 60%. In general, the model explains larger price returns better than smaller.

In the table we also calculate R_S^2 and R_P^2 . These goodness-of-fit measures describe the contribution of each of the imbalances to the overall model fit. R_S^2 is equal to the r-squared of a model with *only the selected* variable, and no others, included[¶]. R_P^2 is the partial r-squared defined as $R_P^2 = (R^2 - \hat{R}^2)(1 - \hat{R}^2)$ where R^2 is the R-squared of the full model with *all* variables included, and \hat{R}^2 is the R-squared

[¶]As such, it is equal to the square of the correlation between the variable and price returns, shown in figure 2.

of a restricted model with *all except* the selected variable. R_P^2 is the partial contribution to the full r-squared due to the selected variable.

Not surprisingly, order routing explains a large proportion of variation in price returns (large R_P^2). However, that does not mean that order routing is the single “important” contributor to market impact. Loosely speaking, “importance” of a factor is a combination of how *large* an effect the factor predicts, and how *reliable* a prediction it produces. I.e., an unreliable prediction of a large magnitude may be more “important” than a reliable prediction of a small inconsequential effect. R^2 describes how reliable a factor is in predicting the price impact, while the estimated coefficient determines the magnitude of the effect.

In the estimated model, price returns were expressed in basis points. The concentration and orderflow imbalances are without units, however they are normalised by the stock standard deviation. Hence the units of regression coefficients can be read as basis points per unit of typical imbalance. Therefore, while the imbalances in order routing and the number of firms trading explain respectively around R_P^2 (market order volume) = 0.12 and R_P^2 (number of firms) = 0.16 to the overall R^2 . The magnitude of their effects for one standard deviation are estimated to be 80 bps and -60 bps respectively. The contribution of concentration imbalance, is more noisy with R_P^2 (concentration) = 0.02, but has a comparable 25 bps effect on the magnitude of the price impact.

As an illustration we can take a sample from the dataset, for example LLOY on $t = 2000-05-09$. The market adjusted return on that day was $\delta P_t = -1.8\%$. Explanatory variables on the day were:

$$\delta E_t = -2.9 \quad \delta V_t = -0.47 \quad \delta N_t = 1.36 .$$

We omit aggressive trade count because its contribution is not significant and likely captured by the aggressive notional. On the day, the concentration of sellers was about 3 times as large as typically, and there was a slightly higher number of buyers than sellers. The number of aggressive sell orders (market orders) was not notably larger than the number of aggressive buy orders. Upon multiplying these values by the coefficients and adding up, the predicted return for the day is

$$\begin{aligned} \delta P_t &= 25 \cdot \delta E_t + 80 \cdot \delta V_t - 60 \cdot \delta N_t \\ -200\text{bps} &= -99 - 23 - 77. \end{aligned} \quad [8]$$

We see that the contribution of concentration to the predicted return is of comparable magnitude as orderflow, albeit with lower reliability. Across the sample of all stocks, the contribution of concentration is between 20% and 30% of the total return.

Checking the statistical specification of the model, the residuals are i.i.d. and very close to normal; all the explanatory variables are exogenous to the model. Therefore, the model coefficients are estimated without bias and are distributed according to a normal distribution. The reported regression coefficient errors and the p-values are calculated in the standard way assuming normality. However, we have also confirmed the estimation errors using a bootstrap test.^{||} Recursive and sliding window estimates are all in line with

^{||}By shuffling the price returns and keeping all other variables intact, we obtain a realisation of the null hypothesis where all the explanatory variables are correlated with themselves but are uncorrelated with the returns. Repeating the shuffling 1000 times and estimating the model on the bootstrapped data, we get a distribution of the coefficients under the null. The standard deviation of the estimates and the p-values obtained in this way coincide with the theoretical values shown in the table.

the overall model fit. In addition, we split the model by year and by stock name. The estimated values are rather constant over the years and do not greatly vary between stocks. We can note however, that higher activity stocks tend to result in better model fits.

While the model and the estimates are reliable, a complication in the interpretation of results is that all of the explanatory variables are, in one way or another, driven and influenced by traded volume. To determine that trading concentration influences price moves in its own right, it is important to determine that the effect of concentration is independent or orthogonal to order routing.

One way of establishing that this is the case, is to split the single model Eq. 7 into, so called, partial regressions, and estimate them via a two stage process. In the first stage we remove the linear effects of order routing from both the price returns *and* trade concentration by fitting models

$$\begin{aligned}\delta P_t &= \alpha_1 \cdot \delta V_t + \beta_1 \cdot \delta N_t + \gamma_1 \cdot \delta M_t + \epsilon_{1,t}, \\ \delta E_t &= \alpha_2 \cdot \delta V_t + \beta_2 \cdot \delta N_t + \gamma_2 \cdot \delta M_t + \epsilon_{2,t}.\end{aligned}\quad [9]$$

The residuals of these fits, $\delta P'_t \equiv \hat{\epsilon}_{1,t}$ and $\delta E'_t \equiv \hat{\epsilon}_{2,t}$, we can term *routing corrected price returns* and *routing corrected concentration*. The second step entails regressing the “corrected” variables on each other

$$\delta P'_t = \eta \cdot \delta E'_t + \epsilon_{3,t}.\quad [10]$$

This procedure removes linear effects of order routing from price returns and trading concentration imbalance. Whatever significance remains in the second step is the effect of trading concentration on price returns – orthogonal to order routing.

A fit of Eq. 10 yields $\hat{\eta} = 0.25 \pm 0.01$ with p-value zero and $R^2 = 0.02$, very much in line with the results in table 1. Fig. 3 shows the fit graphically, together with simple conditional averages and the corresponding standard normal errors of the residuals.

What this shows is that the influence of trading concentration on market returns is indeed orthogonal to order routing, and typically about one quarter in magnitude ($\hat{\eta} = 0.25$).

4. Executing large in scale

The preceding section showed that trading concentration *on average* adversely impacts price moves. Practically speaking however, if one was to trade a large order, what might be the ways of mitigating the additional impact due to concentration? LIS venues aim to reduce impact by matching offsetting interest only when it can be matched in size between a few counterparts. In other words, matching only when both the buy side and the sell side are concentrated. By appropriately partitioning the LSE data we can test this approach and determine if matching concentration with concentration reduces impact.

To do this, we partition the samples into categories, each indexing days from one of the regimes corresponding to illustrations in Fig 1:

- Concentrated buying matched with concentrated selling;
- Concentrated buying matched with diluted selling;
- Dilute buying matched with concentrated selling;
- Dilute buying matched with dilute selling.

Once we remove the effects of order routing, we are left with the average price move conditional to the four concentration regimes. A way to compute these conditional price moves is to create four, so called, “dummy” variables indexing the four market regimes. The model is very much like Eq. 7, with the difference that the continuous concentration imbalance δE_t

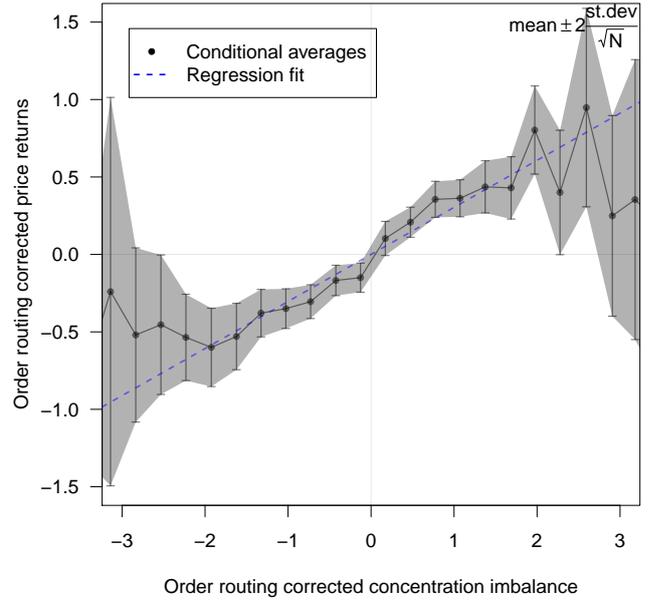


Fig. 3. Figure showing the significant relation of concentration and price impact, after taking order routing effects into account. Days where concentration was high on one side of the market tend to have larger adverse price moves in stock price.

On the x-axis we plot the residual information in the concentration imbalance, once orderflow effects have been regressed out. This is computed as the residuals of the concentration regression in the first of the two stage partial regression process described in the text ($\epsilon_{2,t}$ in Eq.9). Likewise, y-axis shows the residual price return $\epsilon_{1,t}$, after removing the effects of orderflow imbalances. Points in the chart are binned variables averages, the line is the regression line from Eq.10.

is replaced with the four categorical variables. In defining the categories, we consider a market side to be concentrated if the concentration measure is larger than the 70th quantile of the metric, and dilute if it is smaller than the 30th quantile.**

Table 2 displays the model fit containing the four dummy variables. The effect of order routing remains significant and very similar to previously observed. What is interesting here is that of the four conditional means only the ones indexing concentrated with dilute trading are significantly different from zero. Additional price impact due to concentration vanishes in situations where concentrated trading is matched with concentrated counterparts.

Given the above, it seems beneficial (if constraints of the execution allow it) to fill large orders with similarly large opposite interest, minimising the number of trade counterparts. In an anonymous, centrally cleared market it may not be possible to know directly the number of distinct counterparts an order is filled with. However, while far from straightforward, from fill periodicity, size or other patterns, there may be ways to indirectly infer the information. Algos with such logic could, for example, speed up the execution to profit from times when trading with few concentrated counterparties, and slow down when trading with many counterparts.

The other possibility is to advertise interest in large-in-scale (LIS) venues. Arguably, the upstairs voice market of the LSE is a venue designed with such a purpose. The brokers would dial their contacts in search for a few counterparts to fill the order in size. However, trading in the upstairs market brings with it information leakage complications. LIS venues may achieve a similar purpose, but with information leakage tightly controlled and measurable.

**To estimate the quantiles, we merge the concentration metrics for the sell and buy sides.

	Coef.	Error	p-val
Signed volume, δV	82.1	1.8	0.00
No. firms, δN	-65.4	1.4	0.00
No. signed trades, δM	-4.4	1.8	0.01
Concentrated sell, dilute buy	-44.0	4.8	0.00
Concentrated buy, dilute sell	33.0	4.8	0.00
Dilute sell, dilute buy	-3.8	4.3	0.38
Concentrated sell, concentrated buy	4.1	4.3	0.34
Overall		$R^2 = 0.33$	

Table 2. Regression results showing the different effects of concentration on the resulting price move split by different levels of concentration on the two market sides. As before, the concentrated side of the market suffers adverse price moves, but only when matched with a dilute opposite side. When a concentrated order is matched with a concentrated counterpart, the market impact caused by concentration vanishes.

The orderflow effects ($\delta V, \delta N, \delta M$) and coefficients are largely unchanged from before. Likewise, concentrated selling when trading with diluted buying results on average in a price drop (coef = -44 ± 5 bp); concentrated buying when matched with dilute selling results in price appreciation (coef = 33 ± 5 bp). However, when concentrated trading is matched with similarity concentrated counterparts, the impact of concentration vanishes. This observation leads us to expect performance improvements for algos utilising logic to restrict the number of counterparts to trade with.

5. Signalling and time persistence of concentration

This analysis is based on trading in an anonymous and centrally cleared market in which information leakage is minimised. In spite of this, we speculate that it likely is information leakage that is responsible for the observed impact of concentrated trading.

Market impact is known to be a convex function of order size (6, 15, 16), a fact we were able to confirm on the LSE dataset we use as well. The implication of this is that impact per share (or per notional) is larger for small orders than it is for large ones. From a pure mechanical impact view, therefore it follows that when a large order is matched with many small orders, the sum of impacts of small orders should be larger than the impact of the large order. In the language of concentration, the many small orders on the dilute side of the market should – in aggregate – cause more price impact than the large order on the concentrated side of the market. From earlier results, we know the contrary to be the case. Clearly, a simple mechanical explanation does not capture the full story.

We speculate it is likely that signalling and information leakage by the large order allows the “dilute side” firms to profit by adjusting prices as the intent and size of the large order is revealed to the market. For them to be able to do so, a certain level of persistence in concentration is required. Figure 4 shows that, indeed, there is significant autocorrelation in the levels of concentration from one day to the next. Correlations decay roughly as a power function, are strongly significant up to a week, and weakly for multiple weeks. Even the ACF of the concentration imbalance is significant up to a few days.

Correlations in trading concentration can be generated in two ways: (i) one or a few firms trading large orders across days, or (ii) different firms entering the market with large orders sequentially.

The former, “same-firm” correlations, are typically a consequence of splitting large orders and trading them sequentially (17, 18). The latter, “cross-firm” correlations can be caused by news events and its different propagation among firms, or by the so called “herding hypothesis” (19) whereupon observing large orders, traders respond by placing own

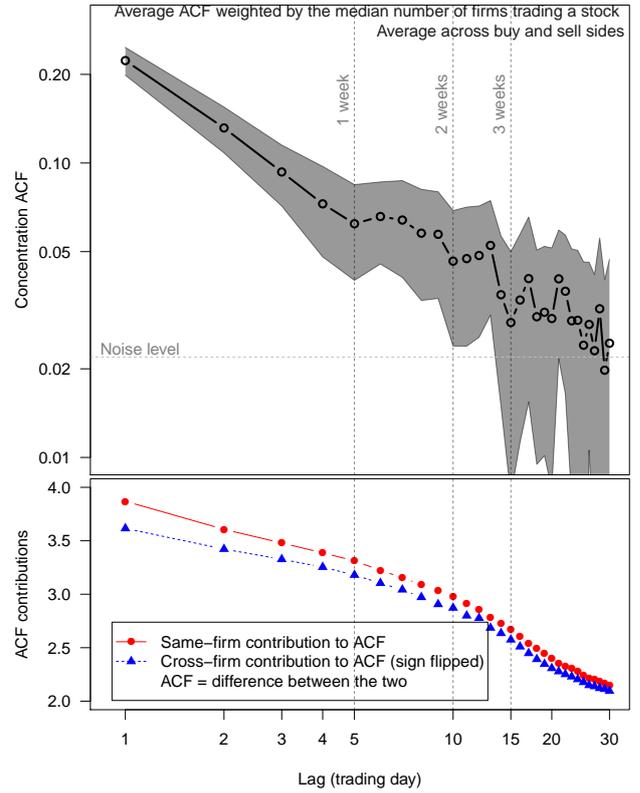


Fig. 4. Autocorrelation function (ACF) of trading concentration showing strong persistence over days (upper chart). We speculate it is this persistence in concentration that allows the dilute side of the market to adjust their behaviour and profit from the knowledge of a large order contributing to concentrated trading. The decay seems to roughly be a power curve (note the log-log scale), and show strong persistence up to a week. The result is obtained by averaging across stocks and both market sides.

The lower chart shows that the persistence in concentration is caused by the same-firms executing large orders across multiple days. We break the aggregate ACF into same- and cross-firm correlations. The effects of the two are opposite, with the same-firm contributing positively to correlations. (We flipped the sign of cross-firm correlations so we can chart it better). The difference between the two curves is the resulting aggregate ACF.

large orders. A “hot-potato” is another variant of the herding hypothesis in which traders sequentially trade a large position between themselves (20).

By splitting the empirical ACF into two components, one contributed by same-firm correlations, and another with cross-firm correlations, we can determine which of the two mechanisms contributes to the persistence in concentration.

The autocorrelation function $\gamma(\tau)$ is defined as

$$\gamma(\tau) = \frac{\langle E(t) \cdot E(t + \tau) \rangle}{\langle E(t) \cdot E(t) \rangle} \quad [11]$$

where the bracket notation $\langle \cdot \rangle$ denotes averaging over days^{††}. We separate the contributions to the ACF by introducing summands $z_i(t)$ so that the concentration on a given day can be written as

$$E(t) = \sum_{i \in \xi_t} \frac{-w_i(t) \log(w_i(t))}{\log N(t)} \equiv \sum_{i \in \xi_t} z_i(t). \quad [12]$$

Substituting the definition of concentration into the expres-

^{††}This expression holds for zero mean variables. Average entropy is not zero, hence we will need to remove the mean prior to computing the decomposition.

sion for the ACF we obtain

$$\gamma(\tau) = \frac{\left\langle \sum_i z_i(t) \cdot \sum_j z_j(t+\tau) \right\rangle}{\left\langle \sum_i z_i(t) \cdot \sum_j z_j(t) \right\rangle} \quad [13]$$

which can be rearranged as

$$\gamma(\tau) = \frac{\left\langle \sum_{i=j} z_i(t) \cdot z_j(t+\tau) \right\rangle + \left\langle \sum_{i \neq j} z_i(t) \cdot z_j(t+\tau) \right\rangle}{\left\langle \sum_{i,j} z_i(t) \cdot z_j(t) \right\rangle}. \quad [14]$$

The first sum takes into account the contributions to the ACF from same-firm, while the second term takes into account cross-firm contributions. It turns out that the second sum is negative, so we will for convenience change the sign of the second sum and write the ACF as

$$\gamma(\tau) = \gamma_{\text{same}}(\tau) - \gamma_{\text{cross}}(\tau). \quad [15]$$

Prior to decomposing the ACF however, we need to ensure $E(t)$ has zero mean which we do by subtracting the mean

$$E'(t) \equiv E(t) - \bar{E} = E(t) - \frac{1}{T} \sum_d E(d). \quad [16]$$

As before, denoting by ξ_d the set of firms or orders present in the market on day d , we can write the above expression as

$$\begin{aligned} E'(t) &= \sum_{i \in \xi_t} z_i(t) - \frac{1}{T} \sum_d \sum_{j \in \xi_d} z_j(d) \\ &= \sum_{i \in \xi_t} \left(z_i(t) - \frac{1}{N_t} \bar{E} \right) \equiv \sum_{i \in \xi_t} z'_i(t) \end{aligned} \quad [17]$$

where N_t again stands for the cardinality of the set ξ_t . We can now compute the components of the ACF using $z'_i(t)$.

The analysis reveals that same-firm and cross-firm correlations are of a different sign and work in opposite directions (lower panel of Fig 4). Same-firm correlations increase overall ACF and are very closely matched in magnitude by the reduction due to the anti-correlation of cross-firms. The small difference in magnitude between the opposing forces results in the overall concentration ACF.

While this analysis does not rule out that the ACF decomposition is purely a mechanical effect due to how we compute entropy, it does show that the persistence in concentration is due to large order splitting. As such it fits well in the common narrative that large orders leak information during execution.

6. Conclusions

This paper introduces the concept of trading concentration and investigates its effect on price impact. When a firm executes a large order, it contributes to the concentration of that side of the market. In such a situation, in addition to price impact due to the way how the order is traded (SOR order routing), concentration will have an adverse effect on price impact, in magnitude roughly 30% as the effects of orderflow. A way to reduce this additional impact is to selectively fill the order with comparably large opposing interest.

This can either be facilitated by expressing interest in LIS venues, or by inferring patterns of fills when trading in an anonymous market, which may signal the presence

of large opposite interest. In such situations, it is advantageous to speed up trading to minimise the number of trade counterparts.

We speculate that large orders, spread over hours or days, over time leak their intent even when trading in anonymous markets. This allows smaller traders to adjust prices and profit from the large order fills.

This paper does not unequivocally find support for trading in large in scale venues. This would be impossible only with data from anonymous trading. What we do show is that even in a standard continuous double auction market, taking into account orderflow effects, price impact of an order is increased when the order is filled with a large number of counterparts.

1. Plerou V, Gopikrishnan P, Gabaix X, Stanley HE (2002) Quantifying stock price response to demand fluctuations. *Physical Review E* 66(2):Article no. 027104.
2. Berger DW, Chaboud AP, Howorka E, Wright J, Chernenko SV (2006) Order flow and exchange rate dynamics in electronic brokerage system data. *International Finance Discussion Papers* 830.
3. Carlson J, Lo M (2006) One minute in the life of the dm/usd: Public news in an electronic market. *Journal of International Money and Finance*.
4. Evans M, Lyons R (2002) Order flow and exchange rate dynamics. *Journal of Political Economy* 110(1):170–180.
5. Payne R (2003) Informed trade in spot foreign exchange markets: An empirical investigation. *Journal of International Economics* 61(2):307–29.
6. Farmer JD, Patelli P, Zovko I (2005) The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America* 102(6):2254–9.
7. Kaniel R, Saar G, Titman S (2008) Individual investor trading and stock returns. *Journal of Finance* 63(1):159–196.
8. Kumar A, Lee CM (2005) Retail investor sentiment and return comovements. *Journal of Finance* 61:2451–2486.
9. Gopikrishnan P, Plerou V, Gabaix X, Stanley HE (2000) Statistical properties of share volume traded in financial markets. *Physical Review E* 62(4):R4493–R4496. Part A.
10. Gabaix X, Gopikrishnan P, Plerou V, Stanley HE (2003) Understanding the cubic and half-cubic laws of financial fluctuations. *Physica A-Statistical Mechanics and Its Applications* 324:1–5.
11. Maslov S, Mills M (2001) Price fluctuation from the order book perspective - empirical facts and a simple model. *Physica A-Statistical Mechanics and Its Applications* (299):234–246.
12. Solomon S, Richmond P (2001) Power laws of wealth, market order volumes and market returns. *Physica A* (299):188–197.
13. Dixon PM, Weiner J, Mitchell-Olds T, Woodley R (1987) Bootstrapping the gini coefficient of inequality. *Ecology* 68(5):1548–1551.
14. Sen A (1973) *On Economic Inequality*. (Oxford Clarendon Press).
15. Potters M, Bouchaud JP (2003) More statistical properties of order books and price impact. *Physica A* 324:133–140.
16. Lillo F, Farmer JD, Mantegna RN (2003) Master curve for price impact function. *Nature* 421(1619):129–130.
17. Lillo F, Farmer JD (2004) The long memory of the efficient market. *Studies in Nonlinear Dynamics and Econometrics* 8(3).
18. Bouchaud JP, Gefen Y, Potters M, Wyart M (2004) Fluctuations and response in financial markets: The subtle nature of "random" price changes. *Quantitative Finance* 4(2):176–190.
19. Lakonishok J, Shleifer A, Vishny R (1992) The impact of institutional trading on stock prices. *Journal of Financial Economics* 31:13–43.
20. Lyons R (1997) A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics* 42:275–298.

ACKNOWLEDGMENTS. The author thanks J.Doyne Farmer for providing the LSE data used in the analysis and early discussions. Thanks to Roel Oomen for reading a draft and providing useful comments.