

The short-term price impact of trades is universal

B. Tóth, Z. Eisler, J.-P. Bouchaud

February 28, 2017

Capital Fund Management, 23 Rue de l'Université, 75007 Paris, France

Abstract

We analyze a proprietary dataset of trades by a single asset manager, comparing their price impact with that of the trades of the rest of the market. In the context of a linear propagator model we find no significant difference between the two, suggesting that both the magnitude and time dependence of impact are universal. This result is important as optimal execution policies often rely on propagators calibrated on anonymous data. We also find evidence that in the wake of a trade the order flow of other market participants first adds further copy-cat trades enhancing price impact on very short time scales. The induced order flow then quickly inverts, thereby contributing to impact decay.

1 Introduction

Trading impacts prices – this is an undisputable empirical statement. However, the interpretation of this observed impact is still debated. Is it observed because trades *forecast* future price changes that would have happened anyway, as Efficient Market theorists would argue? Or should one better think of markets as a kind of physical medium that reacts statistically to all trades, whether informed or not?

This question is key to assess the relevance of most empirical studies of impact, that are based on anonymized order flow, where all orders are *de facto* treated on an equal footing. The average impact extracted from these studies could be misleading if “informed” and “uninformed” trades have completely different impacts. This is especially relevant in the context of popular models of impact, such as Hasbrouck’s VAR model [1, 2] and the “propagator” model [3], where the average reaction to market orders is deduced from a joint analysis of the order flow and the price time series. The standard practice is indeed to calibrate such models using public time series with anonymous orders, and deduce the average lag-dependent impact (or propagator) of market orders, with the hope of using this information to anticipate one’s own market impact in future trading, and of timing executions accordingly [4].

In this study – to our knowledge the first of its kind – we aim to address this question by comparing the impact of the trades of Capital Fund Management (CFM) based on a proprietary dataset with the impact of the rest of the market. Perhaps surprisingly, we find no significant difference between the two. Our results strongly support the “physical medium” picture of markets, that statistically reacts to all perturbations in a similar way. Interestingly, we find that the reaction is first of a “copy-cat” type, where the initial trade is imitated by other market participants on very short time scales, while on longer time scales liquidity refill and/or trades in the opposite direction lead to impact reversion.

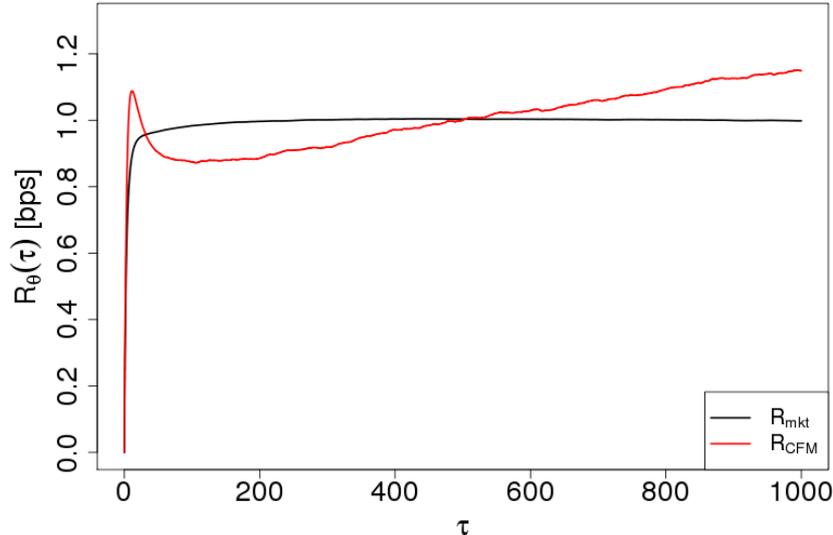


Figure 1: The price response \mathcal{R}_θ for $\theta = \text{mkt}$ (black) and $\theta = \text{CFM}$ (red). \mathcal{R}_{CFM} is stronger than \mathcal{R}_{mkt} for very short times, but \mathcal{R}_{CFM} reverts for $10 < \tau < 100$ and trends afterwards. In the following we will see that the reversion can be related to the reaction of the market, while the longer term trend is due to order splitting.

2 Data and notations

In the following we will analyze data on 98 liquid US stocks in the period 2 January 2014–31 December 2015. This includes all trades executed in the major lit markets. Each trade will be classified via a tag variable θ . Trades by the strategies managed by CFM will be denoted as $\theta = \text{CFM}$, whereas those by all other market participants – which we will consider as a single representative actor – as $\theta = \text{mkt}$. We will consider *trade time*, meaning that the time step counter t is incremented whenever a trade occurs in the stock being studied. All trades on the same stock, reported in the same millisecond, at the same price, in the same direction and bearing the same participant tag are aggregated into one single transaction, because these usually correspond to the same single market order. After aggregation our dataset includes close to 520M trades, of which approx. 0.6% belong to CFM.

We will denote by ϵ_t the side of the market order at time t , $\epsilon_t = +1$ for buyer initiated, and -1 for seller initiated ones. m_t will be the mid price just before the trade t . We will use the indicator function $\mathbb{I}(A)$ defined as $\mathbb{I}(A) = 1$ if A is true and $\mathbb{I}(A) = 0$ otherwise. For example, the expression $\mathbb{I}(\theta_t = \theta)$ is 1 if the trade at time t resulted from a market order by θ . The unconditional probability of the type θ is thus by definition $P(\theta) = \langle \mathbb{I}(\theta_t = \theta) \rangle$. Finally, as discussed in Ref. [5], for any quantity X its average sampled at trades initiated by an actor of type θ is

$$\langle X | \theta_t = \theta \rangle = \frac{\langle X \mathbb{I}(\theta_t = \theta) \rangle}{P(\theta)}. \quad (1)$$

3 Price impact and order flow with labeled data

Let us define the *response function* of trades [3, 6] initiated by actor θ as the average subsequent change of the mid price in the direction of the trade:

$$\mathcal{R}_\theta(\tau) = \frac{\langle (m_{t+\tau} - m_t) \mathbb{I}(\theta_t = \theta) \epsilon_t \rangle}{P(\theta)}. \quad (2)$$

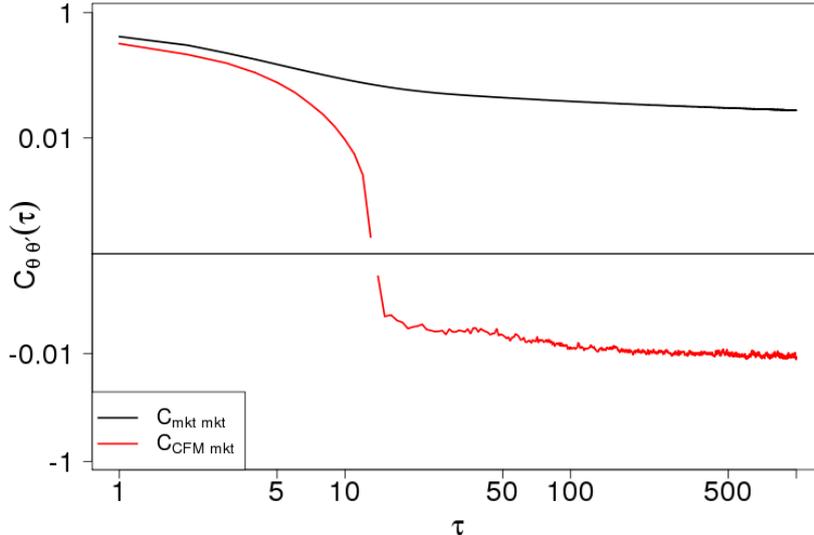


Figure 2: The correlation functions, $C_{\theta, \theta'}$, for $\theta' = \text{mkt}$ and $\theta = \text{mkt}$ (black) and $\theta = \text{CFM}$ (red). Note, that the vertical axis is logarithmic both in the positive and negative planes. We see that while $C_{\text{mkt}, \text{mkt}}$ decays slowly and remains positive up to very long times, $C_{\text{CFM}, \text{mkt}}$ becomes rapidly negative after the initial positive part, and remains weakly negative afterwards.

The average of this quantity across all stocks in our dataset is shown in Fig. 1. The overall shape of $\mathcal{R}_{\text{mkt}}(\tau)$ is compatible with that previously reported in the literature. On the other hand, $\mathcal{R}_{\text{CFM}}(\tau)$ is clearly different, with a hump for short times, and then a continuous upward trend. We will discuss the origin of these differences in detail in the following.

It is convenient to define the correlation function of the market order flow corresponding to two market participants θ and θ' as

$$C_{\theta, \theta'}(\tau > 0) = \frac{\langle \mathbb{I}(\theta_t = \theta) \epsilon_t \mathbb{I}(\theta_{t+\tau} = \theta') \epsilon_{t+\tau} \rangle}{P(\theta)P(\theta')}, \quad (3)$$

where conventionally θ is the first to trade. Considering that both θ and θ' can be CFM or mkt, there are a priori four possible combinations. The two of them corresponding to $\theta' = \text{mkt}$ are of particular interest and are shown in Fig. 2. The market-market order flow autocorrelation $C_{\text{mkt}, \text{mkt}}(\tau)$ is well known to be positive and slowly decaying, as a consequence of order splitting [7, 8]. Since liquidity at the best bid and ask in order books is typically small, market participants slice their full quantity they intend to trade into correspondingly small orders. This results in the observed persistency of order flow revealed by $C_{\text{mkt}, \text{mkt}}$.

The correlation $C_{\text{CFM}, \text{mkt}}(\tau)$, on the other hand, measures how the market “reacts” to CFM’s trades. For short time lags we observe a strong positive correlation, i.e. other actors herd on CFM’s trade. This correlation then quickly reverts on the scale of ~ 10 trades and becomes negative for longer time lags, indicating a contrarian behaviour of the rest of the market.

In order to more precisely quantify the causal relationship between individual orders and subsequent price changes, it is customary [3, 5] to assume that the price response to market orders can be described by linear kernels G_θ , i.e.:

$$m_t = m_{-\infty} + \sum_{\theta} \sum_{t' < t} G_\theta(t - t') \mathbb{I}(\theta_{t'} = \theta) \epsilon_{t'} + \sum_{t' < t} \zeta_{t'}, \quad (4)$$

where ζ_t is a zero mean, i.i.d. noise term. In this equation the *propagator* $G_\theta(t - t')$ represents the evolution of the mid price from t' to t given that the actor θ executed an isolated buy

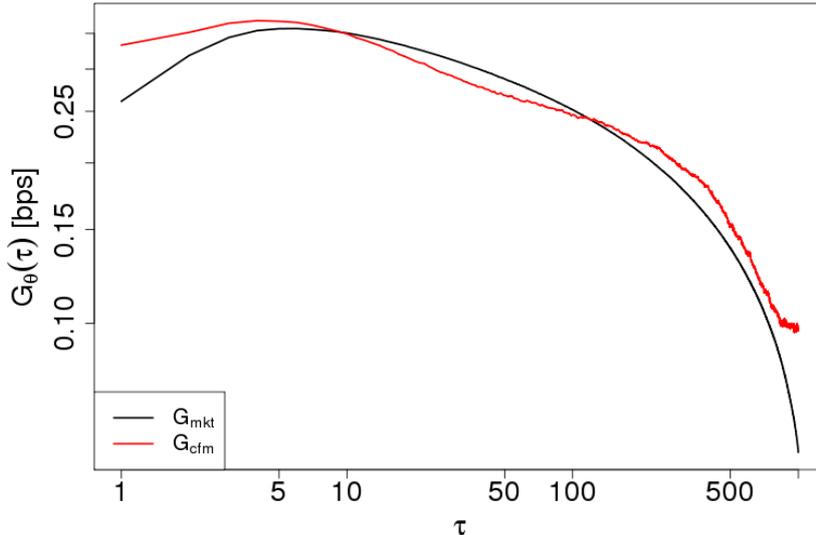


Figure 3: The propagators, G_θ , for $\theta = \text{mkt}$ (black) and $\theta = \text{CFM}$ (red). Apart from the difference for very short times, the two curves are very similar.

order in t' . The model can be calibrated using simple techniques [5]. The resulting propagators $G_{\text{mkt}}(\tau)$ and $G_{\text{CFM}}(\tau)$ are shown in Fig. 3.

Apart from effects on very short time scales, we find no significant difference between the effect of CFM's trades and those of the rest of the market.¹ In other words, the difference between the response functions reported in Fig. 1 can entirely be attributed to the cross-correlations between order flows, and it does not come from the “bare” reaction of the market encoded in the propagators. This is the central message of our paper.

Let us stress that the curves in Fig. 3 represent an average over all 98 stocks in our sample, so one could wonder whether the agreement reported here still holds for each individual stock. The results are obviously much more noisy for the individual stocks separately, but the conclusion is the same, i.e. there is no systematic difference between the time dependence of the market propagator $G_{\text{mkt}}(\tau)$ and the CFM propagator $G_{\text{CFM}}(\tau)$. The global scale of these propagators can, however, be different for certain stocks. This is related to the fact that the average size of CFM's individual trades can be different from that of the rest of the market.

In order to dwell further on the effects induced by order flow correlations, we deconvolute the correlation functions $C_{\theta,\theta'}$ in terms of flow-reaction kernels $K_{\theta,\theta'}$ defined through the following regression [5]:

$$\mathbb{I}(\theta_t = \theta)\epsilon_t = \sum_{\theta'} \sum_{t' < t} K_{\theta',\theta}(t - t') \mathbb{I}(\theta_{t'} = \theta') \epsilon_{t'} + \mu_t, \quad (5)$$

where μ_t is again a zero mean, i.i.d. noise. Here $K_{\theta',\theta}(t - t')$ is the expected order flow imbalance (probability of buys minus probability of sells) of actor θ at time t conditional to actor θ' sending a buy order at an earlier time t' .

Similarly to correlation functions there are four such kernels, shown in Fig. 4. The self-kernels $K_{\text{mkt},\text{mkt}}$ and $K_{\text{CFM},\text{CFM}}$ are quite similar. From the literature and our own experience we know that both curves are dominated by order splitting, common to most large investors

¹Note that we made no volume corrections in this simplified analysis, i.e. all trades, independently of their volume, are assumed to impact the market in the same way. We expect that the small remaining discrepancy comes from such volume effects, in particular at small lags.

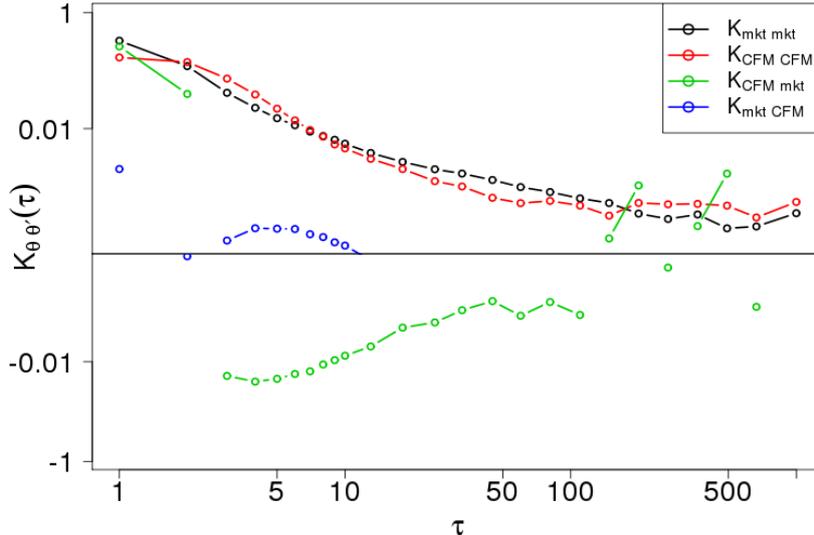


Figure 4: The flow-reaction kernels $K_{\theta, \theta'}$. Note, that the vertical axis is logarithmic both in the positive and negative planes. We find that $K_{CFM, CFM}$ (red) and $K_{mkt, mkt}$ (black) are very similar. $K_{CFM, mkt}$ (green) is positive for very short times and becomes persistently negative later. $K_{mkt, CFM}$ (blue) is much weaker than the other terms.

[7, 8]. Note however, that $K_{mkt, mkt}$ contains both a contribution from order splitting and a contribution from reaction trades among other participants. The kernel $K_{CFM, mkt}$ – describing purely how the market reacts to CFM’s trades – is more interesting. Its shape supports what was expected based on the sign correlation itself. Notably, on a short time scale (in this case up to lag 2) trades by CFM induce copy-cat trades by the market. This effect then reverts and gives way to a weaker, but much more persistent (up to at least 100 trades) flow of contrarian orders by others. This direct observation of the interaction of aggressive order flows from different actors is another noteworthy result of our study. It is also clearly related to the previously documented passive limit order counter-flow (“liquidity refill”), which manifests itself as the decay of the propagator G [5, 6, 9].

4 The total impact of a single trade

By definition, the impact of an isolated trade, stripped off from its “cloud” of correlated trades, is measured by the propagator $G_{mkt}(\tau) \approx G_{CFM}(\tau)$. However, we now know that even if a given market participant chooses to execute a single trade, it will be followed by other “reaction” trades, with a probability encoded by the cross kernel K [5]. The total impact, denoted $G_{CFM}^*(\tau)$, of a single market order sent by CFM is therefore given by

$$G_{CFM}^*(\tau) = G_{CFM}(\tau) + \sum_{0 < \tau_1 < \tau} K_{CFM, mkt}(\tau_1) G_{mkt}^*(\tau - \tau_1), \quad (6)$$

where $G_{mkt}^*(\tau)$ is itself dressed by its own reaction trades. The latter can be represented as

$$G_{mkt}^*(\tau) = G_{mkt}(\tau) + \sum_{0 < \tau_1 < \tau} K_{mkt, CFM}(\tau_1) G_{CFM}^*(\tau - \tau_1) + \sum_{0 < \tau_1 < \tau} K_{mkt, mkt'}(\tau_1) G_{mkt}^*(\tau - \tau_1), \quad (7)$$

where we introduced a new function: $K_{mkt, mkt'}(\tau)$ describes the reaction to an order by a given market participant, by all other market participants *excluding* CFM. Exactly as the propagator

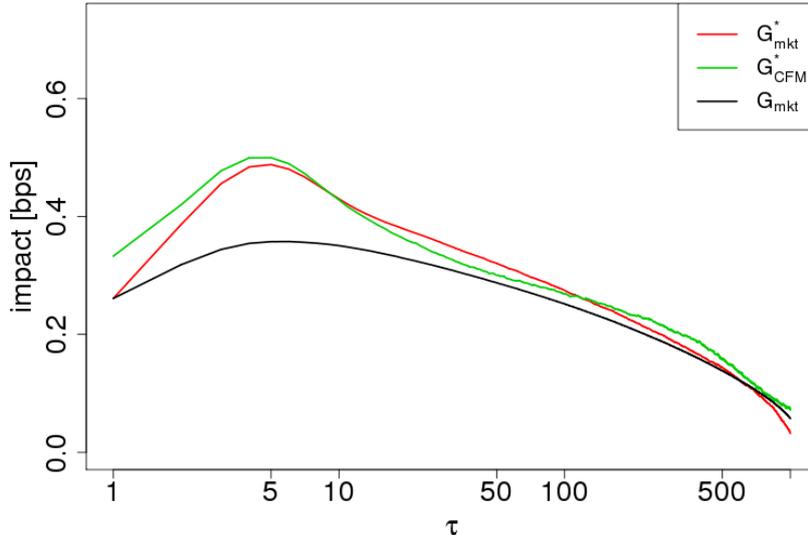


Figure 5: Comparison of G_{mkt} (black), G_{mkt}^* (red) and G_{CFM}^* (green). The total impact of a single trade is enhanced at short times by copy-cat trades, but afterwards decays back close to the bare propagator G . G_{CFM}^* is very similar to G_{mkt}^* in line with our assumption that CFM trades are not statistically different from those of the market.

for CFM trades is nearly identical to that of any market trade, it is reasonable to assume that in an anonymous market the same is true for the reaction of the order flow itself. In other words, one expects that

$$K_{\text{mkt},\text{mkt}'}(\tau) \approx K_{\text{CFM},\text{mkt}}(\tau). \quad (8)$$

Furthermore, from the blue curve in Fig. 4 we see that the reaction of CFM's order flow to market trades is negligible, i.e. $K_{\text{mkt},\text{CFM}}(\tau) \approx 0$. This final observation along with Eq. (8) allows us to simplify Eq. (7) to a closed, self-consistent equation for $G_{\text{mkt}}^*(\tau)$ in terms of observable quantities:

$$G_{\text{mkt}}^*(\tau) \approx G_{\text{mkt}}(\tau) + \sum_{0 < \tau_1 < \tau} K_{\text{CFM},\text{mkt}}(\tau_1) G_{\text{mkt}}^*(\tau - \tau_1). \quad (9)$$

This equation can be solved numerically, leading to the result plotted in Fig. 5. As expected from the short-term herding effect unveiled in the last section, we see that the total impact of a single trade is enhanced at short times by copy-cat trades, but afterwards decays back close to the bare propagator G . Comparing Eq. (9) with Eq. (6) furthermore yields $G_{\text{CFM}}^*(\tau) \approx G_{\text{mkt}}^*(\tau)$, in line with our assumption that CFM trades are not statistically different from those of the market.

5 Conclusion

Based on our proprietary dataset, we have shown that there is no significant difference between the short term price impact of CFM trades and that of the rest of the market. This is a strong indication of the universality of impact across all market participants, bolstering the idea that short term impact has no relation whatsoever to the information content of trades, but is rather a property of the market itself. In an anonymous market we indeed expect the same collective behavior to follow the actions of any market participant. Our findings empirically support

this theory, which is important as optimal execution policies often rely on it as an implicit assumption. The interaction of aggressive order flow among actors contributes to price moves in a small, but significant way, and its calibration on anonymous data is warranted by the present results.

Acknowledgements

We thank J. Bonart, X. Brokmann, J. Donier, M. Gould, J. Kockelkoren and I. Mastromatteo for many useful discussions on these topics.

References

- [1] Joel Hasbrouck. Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179–207, 1991.
- [2] Joel Hasbrouck. *Empirical Market Microstructure*. Oxford University Press New York, 2007.
- [3] Jean-Philippe Bouchaud, Yuval Gefen, Marc Potters, and Matthieu Wyart. Fluctuations and response in financial markets: the subtle nature of random price changes. *Quantitative Finance*, 4(2):176–190, 2004.
- [4] Jim Gatheral, Alexander Schied, and Alla Slynko. Transient linear price impact and fredholm integral equations. *Mathematical Finance*, 22(3):445–474, 2012.
- [5] Zoltán Eisler, Jean-Philippe Bouchaud, and Julien Kockelkoren. Models for the impact of all order book events. *Market Microstructure: Confronting Many Viewpoints*, pages 113–135, 2013.
- [6] Bence Tóth, Zoltán Eisler, Fabrizio Lillo, Julien Kockelkoren, Jean-Philippe Bouchaud, and J Doyne Farmer. How does the market react to your order flow? *Quantitative Finance*, 12(7):1015–1024, 2012.
- [7] Fabrizio Lillo, J Doyne Farmer, et al. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics*, 8(3):1, 2004.
- [8] Bence Tóth, Imon Palit, Fabrizio Lillo, and J Doyne Farmer. Why is equity order flow so persistent? *Journal of Economic Dynamics and Control*, 51:218–239, 2015.
- [9] Jean-Philippe Bouchaud, Julien Kockelkoren, and Marc Potters. Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance*, 6(2):115–123, 2006.