Journal of Financial Economics 120 (2016) 269–285

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/finec

Discerning information from trade data[☆]

David Easley[a], Marcos Lopez de Prado[b,c], Maureen O'Hara[d,*]

[a] Department of Economics, Cornell University, United States
[b] Guggenheim Partners, United States
[c] Computational Research Division, Lawrence Berkeley National Laboratory, United States
[d] Johnson Graduate School of Management, Cornell University, United States

ARTICLE INFO

Article history:
Received 21 August 2012
Revised 15 June 2015
Accepted 7 July 2015
Available online 22 January 2016

JEL classification:
C02
D52
D53
G14

Keywords:
Trade classification
Bulk volume classification
Flow toxicity
Volume imbalance
Market microstructure

ABSTRACT

How best to discern trading intentions from market data? We examine the accuracy of three methods for classifying trade data: bulk volume classification (BVC), tick rule and aggregated tick rule. We develop a Bayesian model of inferring information from trade executions and show the conditions under which tick rules or bulk volume classification predominates. Empirically, we find that tick rule approaches and BVC are relatively good classifiers of the aggressor side of trading, but bulk volume classifications are better linked to proxies of information-based trading. Thus, BVC would appear to be a useful tool for discerning trading intentions from market data.

1. Introduction

Microstructure models attach a key role to trade data because of their signal value for underlying trading intentions. Traders informed of good news profit by buying and traders informed of bad news profit by selling, thus creating a trade imbalance. Trade imbalance between buys and sells can also signal liquidity pressure in markets, leading to subsequent price movements. Discerning information from trade data, however, has never been straightforward, and a variety of trade classification algorithms in the literature are devoted to this task.[1]

The advent of high frequency markets complicates this endeavor in two fundamental ways. First, the mechanics of trading are radically different than in times past. Trading now takes place largely in electronic markets in which designated liquidity providers need not be present, and practices such as hidden orders make drawing inferences from market data problematic. In US equity markets, trading is fragmented across 13 exchanges and 40 or more alternative trading venues, each reporting trades to the consolidated tape, but at different latencies. Thus, trades on the tape are out of order, compromising trade classification rules

☆ We thank the referee, the editor (Bill Schwert), Robert Almgren, Robert Engle, Craig Holden, Andrew Karolyi, David Leinweber, Pam Moulton, Mark Ready, Riccardo Rebonato, John Wu, and the CIFT (Computational Infrastructure for Financial Technology) group at the Lawrence Berkeley National Laboratory for providing useful comments. We are grateful to Sergey Kosyakov and Steven Jones for their research assistance.

* Corresponding author. Tel.: +1 607 255 3645; fax: +1 607 255 5993.
E-mail address: mo19@cornell.edu, mo19@comell.edu (M. O'Hara).

[1] See, for example, Lee and Ready (1991), Ellis, Michaely, and O'Hara (2000), and Chakrabarty, Moulton and Shkilko (2012).

http://dx.doi.org/10.1016/j.jfineco.2016.01.018
S0304-405X(16)00024-6/

based on upticks or downticks.[2] Order cancellation rates of 98% or more complicate knowing actual quotes, so trade classification algorithms based on proximity to bid and ask quotes are also severely compromised.[3]

A second, and potentially more serious, problem is that the trading process is fundamentally different. Algorithms chop parent orders into numerous child orders, so order flow, not individual orders, relate to trade motivation. Trading is also done dynamically, with direct market access (DMA) allowing participants to strategically place multiple orders at various price levels in the book, monitor the progression of their limit orders in the queue, and cancel and replace orders at different levels. To see why this matters, consider, for example, a trader informed of good news who dynamically trades via limit orders. Instead of hitting the ask price to buy, this trader could post an order at the bid. When that order is hit, the trade appears to be a sell as it is taking place at the lower bid price. To continue to buy via limit orders, the informed trader has to post a higher bid or due to time priority his order simply sits behind the orders at the existing bid. This persistent bidder strategy means that prices are forced up even though the active side of the trade is always the uninformed seller.

To the extent that informed traders use limit orders, the notion of the active side of the trade signaling underlying information is undermined because informed orders are not crossing the spread. A variety of research (see Bloomfield, O'Hara, and Saar, 2005; Hasbrouck and Saar, 2009; Bouchard, Farmer, and Lillo, 2009; Eisler, Zoltan, Bouchaud, and Kockelkoren, 2012; Kim and Stoll, 2013; Collin-Dufresne and Vos, 2015; O'Hara, 2015) suggests that equating informed traders with aggressive traders is no longer accurate. This has the important implication that buy and sell trades, or imbalances in these trades, might not be good indicators of underlying information.[4]

In this paper, we investigate how best to discern information from trade data. Microstructure research often relies on simple trade classification algorithms to accomplish this task, and we investigate the efficacy of two such approaches for classifying trade data: the bulk volume classification (BVC) methodology and tick rules. Tick rule approaches use simple movements in trade prices (upticks or downticks) to classify a trade as either a buy or a sell. The bulk volume technique, which was first applied in Easley, Lopez de Prado, and O'Hara (2011), aggregates trades over short time or volume intervals and then uses a standardized price change between the beginning and end of the interval to approximate the percentage of buy and sell order flow. Each of these techniques maps observable data into proxies for trading intentions, but how well any of these approaches works in the new high frequency world is unclear.

To understand the differences between these approaches, it is useful to start from a conceptual framework in which we care about some underlying unobservable information (informed buyers or sellers) and we observe data (trade prices) that are correlated with the unobservable information. A Bayesian statistician would start with a prior on the unobservable information, observe the data, and use a likelihood function to update his or her prior to form a posterior on the underlying information. This is not what a tick rule does. It classifies a trade as a buy if the previous price is below the current price, a sell, if it is above. The bulk volume approach, by contrast, can be thought of as assigning a posterior probability to a trade being a buy or sell, an approach closer conceptually to Bayes' rule.

Using a statistical model, we investigate the errors that arise from a tick rule approach and the bulk volume approach, relative to a Bayesian approach. We show that when the noise in the data is low, tick rule errors can be relatively low, and over some regions the tick rule can perform better than the bulk volume approach. When noise is substantial, the bulk volume approach can outperform a tick rule and permit more accurate sorting of the data. Moreover, our model shows that when order flow is imbalanced (as would be the case, for example, when trades are motivated by private information), tick rules based on noisy data underestimate the probability of buys when there is good information and overestimate it when there is bad information.

The underlying information about trading intentions that we seek is not observable, but microstructure theory suggests a variety of proxies for that information. In our empirical work, we test the accuracy of the bulk volume and tick rule approaches using two such proxies: the aggressor side of trading (as given by buy-sell indicator flags in the data) and an estimate of spreads. We show that the BVC and tick rule approaches do reasonably well in matching the aggressor side of trading, with tick rules generally being more accurate. When we consider spread effects using the Corwin and Shultz methodology, however, we find that the BVC approach dominates tick rule measures, which can have perverse correlations with this information proxy. We conclude that BVC appears to be a useful tool for market researchers interested in discerning trading intentions.

An interesting upshot of these results is that the aggressor side of trading appears little related to any underlying information, a decoupling that we argue arises from how trading transpires in modern high frequency markets. Our findings complement recent work by Collin-Dufresne and Vos (2015) who find that standard measures of adverse selection relying on estimates of the persistent price effects

[2] This problem is also particularly acute in the new swap trading markets. The Dodd-Frank Wall Street Reform and Consumer Protection Act currently requires reporting of nonblock trades to the Swap Data Repository, but current reporting rules allow a 30-minute delay. So there is no way to determine the correct order of trades. See also Ding, Hanna, and Hendershot (2014) for evidence on how speed differences between proprietary feeds and the consolidated tape complicate knowing current quotes.

[3] See Hasbrouck (2013) for an excellent analysis of quote volatility and its implications.

[4] These changes also mean that trade information will not be linked with other variables of interest such as trader identity. For example, Lee and Radhakrishnan (2000) and Campbell, Ramdorai, and Schwartz (2008) propose size cutoff rules on trades that they argue identify institutional trading. Even using data from the year 2000, Campbell, Ramdorai, and Schwartz note problems in identification arising from what they suspect was algorithmic trading. With trade sizes now all collapsing to minimum levels, and institutions trading dynamically with limit orders, inferring trader identity from trade size is a daunting task.