



Observatory of
Complex Systems



An Introduction to High Frequency Finance and Market Microstructure

Fabrizio Lillo

University of Palermo (Italy) and Santa Fe Institute (USA)

Pisa, February 11, 2010

Outline


- Overview of high frequency financial data
- An introduction to market microstructure and a strategic model: Kyle
- A look to microstructure data
 - Price impact
 - Order flow
- A microstructure model of price formation
- Inferring strategies from agent's trading data

Financial data

- In the last thirty years the degree of resolution of financial data has increased
 - Daily data
 - Tick by tick data
 - Order book data
 - “Agent” resolved data
- High frequency data are like a powerful microscope -> Change in the nature of social sciences ?
- High frequency data require sophisticated stochastic modeling and statistical techniques

Daily data

- Daily financial data are available at least since nineteenth century
- Usually these data contains opening, closing, high, and low price in the day together with the daily volume
- Standard time series methods to investigate these data




(from Gopikrishnan et al 1999)

Tick by tick data


- Financial high frequency data usually refer to data sampled at a time horizon smaller than the trading day
- The usage of such data in finance dates back to the eighties of the last century
 - Berkeley Option Data (CBOE)
 - TORQ database (NYSE)
 - HFDF93 by Olsen and Associates (FX)
 - CFTC (Futures)

- Higher resolution means new problems
 - data size: example of a year of a LSE stock
 - 12kB (daily data)
 - 15MB (tick by tick data)
 - 100MB (order book data)
 - the total volume of data related to US large caps in 2007 was 57 million lines, approximately a gigabyte of stored data
 - A large cap stock in 2007 had on average 6 transactions per second, and on the order of 100 events per second affecting the order book
 - irregular temporal spacing of events
 - the discreteness of the financial variables under investigation
 - problems related to proper definition of financial variables
 - intraday patterns
 - strong temporal correlations
 - specificity of the market structure and trading rules.


Data size




Irregular temporal spacing



Periodicities



Temporal correlations




- More structured data require more sophisticated statistical tools


- data size:
 - more computational power
 - better filtering procedures
- irregular temporal spacing of events
 - point processes, ACD model, CTRW model...
- the discreteness of the financial variables under investigation
 - discrete variable processes
- problems related to proper definition of financial variables
- intraday patterns
- strong temporal correlations
 - market microstructure
- specificity of the market structure and trading rules.
 - better understanding of the trading process

Order book data

- The next resolution of financial data contains data on all the orders placed or canceled in the market
- Many stock exchanges (NYSE, LSE, Paris) works through a double auction mechanism
- Order book data are fundamental to investigate the price formation mechanisms



Representation of limit order book dynamics



(Ponzi, Lillo, Mantegna 2007)

Efficient Market Hypothesis

Samuelson (1965) stated that in an informationally efficient market, price changes must be unforecastable if they fully incorporate the expectations and information of all market participants

“The market is said to be efficient with respect to some information set if prices would be unaffected by revealing that information to all participants. Moreover, efficiency with respect to an information set implies that it is impossible to make economic profits by trading on the basis of that information set”

(Malkiel, 1992)

Efficient Market Hypothesis

- Therefore, under the efficient market hypothesis, **price changes must be unforecastable**
- The more efficient the market, the more random is the sequence of price changes
- The price is described by a martingale

$$E(P_{t+1} | P_t, P_{t-1}, \dots) = P_t$$

i.e. the best forecast of tomorrow's price is simply today's price

- A widespread model of price that incorporates the efficiency of the market is the **Random Walk Hypothesis**

Random Walk Hypothesis

The simplest model for price dynamics in discrete time is

$$p_t = \mu + p_{t-1} + \epsilon_t$$


where $p_t = \log P_t$, μ a constant, and ϵ_t a noise term consistent with the Efficient Market Hypothesis

Depending on the properties of ϵ_t , we distinguish

- Independent Identically Distributed increments: for example Gaussian distributed (stable laws)
- Independent increments
- Uncorrelated increments: the weaker form implies the vanishing of the linear autocorrelation


Aggregate stylized facts in financial time series

Random walk




(Bachelier 1900)


Efficiency



Fat tails



Clustered volatility



Market microstructure

- Market microstructure “is devoted to theoretical, empirical, and experimental research on the economics of securities markets, including the role of information in the price discovery process, the definition, measurement, control, and determinants of liquidity and transactions costs, and their implications for the efficiency, welfare, and regulation of alternative trading mechanisms and market structures” (NBER Working Group)

Strategic model: Kyle (1985)

- The model describes a case of information asymmetry and the way in which information is incorporated into price.
- It is an equilibrium model
- There are several variants: single period, multiple period, continuous time
- The model postulates three (types of) agents: an informed trader, a noise trader, and a market maker (MM)

- The terminal (liquidation) value of the asset is v , normally distributed with mean p_0 and variance Σ_0 .
- The informed trader knows v and enters a demand x
- Noise traders submit a net order flow u , normally distributed with mean 0 and variance σ_u^2 .
- The MM observes the total demand $y=x+u$ and then sets a price p . All the trades are cleared at p , any imbalance is exchanged by the MM.

- The informed trader wants to trade as much as possible to exploit her informational advantage
- However the MM knows that there is an informed trader and if the total demand is large (in absolute value) she is likely to incur in a loss. Thus the MM protects herself by setting a price that is increasing in the net order flow.
- The solution to the model is an expression of this trade-off

Informed trader

- The informed trader conjectures that the MM uses a linear price adjustment rule $p = \lambda y + \mu$, where λ is inversely related to liquidity.

- The informed trader's profit is

$$\pi = (v - p)x = x[v - \lambda(u + x) - \mu]$$

and the expected profit is

$$E[\pi] = x(v - \lambda x - \mu)$$

- The informed traders maximizes the expected profit, i.e.

$$x = (v - \mu) / 2\lambda$$

- In Kyle's model the informed trader can loose money, but on average she makes a profit

Market maker

- The MM conjectures that the informed trader's demand is linear in v , i.e. $x = \alpha + \beta v$
- Knowing the optimization process of the informed trader, the MM solves

$$(v - \mu) / 2\lambda = \alpha + \beta v$$

$$\alpha = -\mu / 2\lambda \qquad \beta = 1 / 2\lambda$$

- As liquidity drops the informed agent trades less
- The MM observes y and sets

$$p = E[v | y]$$

Solution

- If X and Y are bivariate normal variables, it is

$$E[Y | X=x] = \mu_Y + (\sigma_{XY}/\sigma_X^2)(x - \mu_X)$$

- This can be used to find

$$E[v | y] = E[v | u + \alpha + \beta v]$$

- The solution is

$$\alpha = -p_0 \sqrt{\frac{\sigma_u^2}{\Sigma_0}}; \quad \mu = p_0; \quad \lambda = \frac{1}{2} \sqrt{\frac{\Sigma_0}{\sigma_u^2}}; \quad \beta = \sqrt{\frac{\sigma_u^2}{\Sigma_0}};$$

Solution (II)

- The impact is linear and the liquidity increases with the amount of noise traders

$$p = p_0 + \frac{1}{2} \sqrt{\frac{\Sigma_0}{\sigma_u^2}} y$$

- The informed agent trades more when she can hide her demand in the noise traders demand

$$x = (v - p_0) \sqrt{\frac{\sigma_u^2}{\Sigma_0}}$$

- The expected profit of the informed agent depends on the amount of noise traders

$$E[\pi] = \frac{(v - p_0)^2}{2} \sqrt{\frac{\sigma_u^2}{\Sigma_0}}$$

- The noise traders loose money and the MM breaks even (on average)

Kyle model - summary

- The model can be extended to multiple periods in discrete or in continuous time
- The main predictions of the model are
 - The informed agent “slices and dices” her order flow in order to hide in the noise trader order flow
 - Linear price impact
 - Uncorrelated total order flow
 - Permanent and fixed impact

Price impact

- Price impact is the correlation between an incoming order and the subsequent price change
- For traders impact is a cost -> Controlling impact
- Volume vs temporal dependence of the impact
- However price impact may indicate many different quantities
 - Price reaction to individual trades
 - Price reaction to an aggregate number of trades
 - Price reaction to a set of orders of the same sign placed consecutively by the same trader (hidden order)
 - Price reaction in a market to a trade in another market (e.g. electronic market and block market)


What is the origin of price impact?

- Given that in any transaction there is a buyer and a seller, why is there price impact?
 - Agents successfully forecast short term price movements and trade accordingly (i.e. trade has no effect on price and noise trades have no impact)
 - The impact of trades reveals some private information (e.g. Kyle model)
 - Impact is a statistical effect due to order flow fluctuations (zero-intelligence models, self-fulfilling prophecy)


“Orders do not impact prices. It is more accurate to say that orders forecast prices” (Hasbrouck 2007)

Impact of individual trades


NYSE



London Stock Exchange



Paris Bourse
(Potters et al. 2003)



Impact of
individual
transaction is NOT
universal



Individual market impact is a concave function of the volume

Master curve for individual impact

Lillo, Famer, Mantegna, Nature 2003


GROUP A -> least capitalized group

GROUP T -> most capitalized group



$$r = C^{-\gamma} f(V C^{\delta})$$

AGGREGATE MARKET IMPACT



$$R = E[R_t | V, N]$$

$$r_i = \log(p_i / p_{t-1})$$


$$R_t = \sum_{i=t+1}^{t+N} r_i$$

$$V_t = \sum_{i=t+1}^{t+N} v_i$$


Return vs. signed volume for N successive transactions

Vertical offset = -3 for $N = 8$ and -6 for $N = 64$

Structure of a financial market




A typical inventory profile




Credit Agricole trading Santander

Detecting hidden orders



We developed a statistical method to identify periods of time when an investor was consistently (buying or selling) at a constant rate -> **Hidden orders**

Market impact of (statistically identified) hidden orders



$$|\langle R|N \rangle| = A N^\gamma$$

$$\gamma=0.44-0.72$$

Figure 4: Average rescaled market impact R for hidden orders shorter than 1 day as a function of N for the BME (left) and LSE (right). Circles are the results for all hidden orders, while squares are the results when there is a low fraction of market orders ($f_{mo} < 0.2$) and diamonds are for when there is a large fraction of market orders ($f_{mo} > 0.8$). Dashed lines are power law fits $R \sim N^\gamma$. Values of γ are reported in Table II.

All types of price impact are concave
(but each type for a different reason)

Order flow


- One of the key problem in the analysis of markets is the understanding of the relation between the order flow and the process of price formation.
- The order flow is strongly depending on
 - the decisions and strategies of traders
 - the exploiting of arbitrage opportunities
- This problem is thus related to the balance between supply and demand and to the origin of the market efficiency.

Market order flow

- We investigate the London Stock Exchange in the period 1999-2002
- We consider market orders, i.e. orders to buy at the best available price triggering a trade
- We consider the symbolic time series obtained by replacing buy orders with **+1** and sell orders with **-1**
- The order flow is studied mainly in event time

....**+1**,**+1**,**-1**,**-1**,**-1**,**+1**,**-1**,**+1**,**+1**,**+1**,**-1**,**-1**,**+1**,...

Market Order Sign



$$\rho(\tau) \sim \tau^{-0.5}$$

Lillo and Farmer (2004)
independently discovered by Bouchaud et al
(2004)

Long-memory processes

- There is strong evidence that order sign is a long memory process, i.e. a process with an autocorrelation function decaying asymptotically as

$$\rho(\tau) = \langle x(t + \tau)x(t) \rangle \sim \frac{1}{\tau^\gamma} \quad \gamma < 1$$

- Long memory processes lack a typical time scale and are related to anomalous diffusion
- A long memory process depends on its entire past history and has often strong predictability on long time horizons.

The long memory of order sign is statistically robust and persists across several days

This raises two questions:


- Why is the order flow a long-memory process ?
- How can the market be efficient in the presence of long memory of order sign ?

WHAT CAUSES LONG-MEMORY OF SUPPLY AND DEMAND FLUCTUATIONS?

- Herding (LeBaron and Yamamoto, Physica A 2007)
- Strategic hiding of supply and demand (Mike, Lillo, Farmer, Physical Review E, 2005):
 - To avoid revealing true intentions, large investors break their trades up into small pieces, and trade incrementally. (Kyle, Econometrica, 1985). We simply assume order splitting and compute its consequences
 - Converts heavy tail of volume distribution into long-memory in signs of realized trade series.

AUTOCORRELATION OF VOLUME SIGN (VS. TRANSACTION NUMBER)

Autocorrelation of trade signs
for same vs. different broker codes



Order splitting is the most likely explanation for correlated order flow

Model of order splitting

- We assume that the distribution of initial hidden order size is a Pareto distribution

$$p(L) = \frac{\alpha}{L^{\alpha+1}} \quad L \geq 1 \quad \alpha > 1$$



This assumption relates to the Pareto distribution of mutual fund size

We prove that for our model the time series of the signs of the revealed order has an autocorrelation function decaying asymptotically as

$$\rho(\tau) \sim \frac{N^{\alpha-2}}{\alpha} \frac{1}{\tau^{\alpha-1}} \longrightarrow \gamma = \alpha - 1$$

Volume distribution


The volume of on-book and off-book trades have different statistical properties




- The exponent $\alpha=1.5$ for the hidden order size and the market order sign autocorrelation exponent γ are consistent with the order splitting model which predicts the relation $\gamma=\alpha-1$.

Is it possible to identify directly hidden orders?


Distributional properties of hidden orders



Investment horizon



Number of transactions



$$\gamma \cong \alpha - 1$$

The long memory of order sign is statistically robust and persists across several days

This raises two questions:

- Why is the order flow a long-memory process ?
- How can the market be efficient in the presence of long memory of order sign ?

Long memory and efficiency

- How can the long memory of order flow be compatible with market efficiency?
- In the previous slides we have shown two empirical facts
 - Single transaction impact is unconditionally and on average non zero and given by

$$E[r|v] = \text{sign}(v)f(v) = \varepsilon f(v)$$

- The sign time series is a long memory process

$$E[\varepsilon_t \varepsilon_{t+\tau}] \approx \tau^{-\gamma}$$

Naïve model

- Consider a naïve random walk model of price dynamics

$$p_{t+1} - p_t \equiv r_t = \varepsilon_t f(v_t) + \eta_t$$

- It follows that

$$E[r_t r_{t+\tau}] \propto E[\varepsilon_t \varepsilon_{t+\tau}] \approx \tau^{-\gamma}$$

- If market order signs ε_t are strongly correlated in time, price returns are also strongly correlated, prices are easily predictable, and the market inefficient.

- It is not possible to have an impact model where the impact is both fixed and permanent
- There are two possible modifications
 - A fixed but transient impact model (Bouchaud et al. 2004)
 - A permanent but variable (history dependent) impact model (Lillo and Farmer 2004, Gerig 2007, Farmer, Gerig, Lillo, Waelbroeck)
- For simplicity we shall neglect volume fluctuations, i.e. we assume that all orders have unit volume

Propagator model (Bouchaud et al 2004)

- The model postulates a fixed but transient impact described by the equation

$$p_t = p_{-\infty} + \lambda \sum_{k=1}^{\infty} G(k) \varepsilon_{t-k} + noise$$

i.e. for returns

$$r_t \equiv p_{t+1} - p_t = \lambda G(1) \varepsilon_t + \lambda \sum_{k=1}^{\infty} [G(k+1) - G(k)] \varepsilon_{t-k} + noise$$

- Since $E[\varepsilon_t \varepsilon_{t+\tau}] \approx \tau^{-\gamma}$, returns are uncorrelated if

$$G(k) \approx k^{-\beta} = k^{-(1-\gamma)/2}$$

- The impact of a transaction for large times is finite

History dependent, permanent impact model

- We assume that agents can be divided in three classes
 - Directional traders (liquidity takers) which have large hidden orders to unload and create a correlated order flow
 - Liquidity providers, who post bid and offer and attempt to earn the spread
 - Noise traders
- The strategies of the first two types of agents will adjust to remove the predictability of price changes

Model for price diffusion

We neglect volume fluctuations and we assume that the naïve model is modified as

$$p_{t+1} - p_t \equiv r_t = \lambda(\varepsilon_t - \hat{\varepsilon}_t) + \eta_t \qquad \hat{\varepsilon}_t = E_{t-1}[\varepsilon_t | \Omega]$$

where Ω is the information set of the liquidity provider.

It is the “surprise” in order flow that impacts price

Ex post there are two possibilities, either the predictor was right or wrong.

Let p_t^+ (p_t^-) be the probability that the next order has the same (opposite) sign of the predictor and r_t^+ (r_t^-) are the corresponding price change

- The efficiency condition $E_{t-1}[r_t | \Omega] = 0$ can be rewritten as



$$p_t^+ r_t^+ - p_t^- r_t^- = 0$$

- The market maker has expectations on p_t^+ and p_t^- given her information set Ω and adjusts r_t^+ and r_t^- in order to make the market efficient

-----> MARKET EFFICIENCY

ASYMMETRIC LIQUIDITY MODEL

Empirical evidence of asymmetric liquidity



A linear model

The history dependent, permanent model is completely defined when one fixes

- the information set Ω of the liquidity provider
- the model used by the liquidity provider to build her forecast $\hat{\varepsilon}_t$

As an important example we consider the case in which

- the information set is made only by the past order flow
- the liquidity provider uses a finite or infinite order autoregressive model to forecast order flow

$$r_t = \lambda \left(\varepsilon_t - \sum_{i=1}^K a_i \varepsilon_{t-i} \right) + \eta_t$$

If the order flow is long memory, i.e. $E[\varepsilon_t \varepsilon_{t+\tau}] \approx \tau^{-\gamma}$ and described by an autoregressive model, the optimal parameters of the autoregressive model scales as

$$a_k \approx k^{\frac{\gamma-3}{2}} \equiv k^{-\beta-1}$$

and the optimal number of lags K in the model is infinite.

More realistically, K is finite and the optimal parameters of the autoregressive model follows the same scaling behavior with k

Under these assumptions and if K is infinite the linear model becomes mathematically equivalent to the fixed-temporary model (or propagator) model by Bouchaud et al.

$$G(k) = 1 - \sum_{j=1}^{k-1} a_j$$

Impact of hidden orders

The above model allows to make quantitative prediction on the impact of an hidden order

Assume an hidden order of length N is placed by a liquidity taker by using a slice and dice strategy which mixes the trades with the flow of noise traders with a constant participation rate π

The impact of the hidden order is

$$E[p_N] - p_0 \approx \varepsilon \theta \left(1 + \frac{2^{\beta-1} \pi^\beta}{1-\beta} \left[(2N-1)^{1-\beta} - 1 \right] \right) \approx \pi^\beta N^{1-\beta}$$

An empirical value $\gamma=0.5$, gives $\beta=0.25$, which in turn implies that the impact of an hidden order should grow as the $\frac{3}{4}$ power of its size. Moreover, as expected, the impact is smaller for slower execution (i.e. smaller π)

- Asymmetric liquidity depends on the information set Ω .
- This model predicts the existence of two classes of traders that are natural counterparties in many transactions
 - Large institutions creates predictable component of order flow by splitting their large hidden orders
 - Hedge funds and high frequency traders removes this predictability by adjusting liquidity (and making profit)
- This ecology of market participants is empirically detectable?
- What is the interaction pattern between market participants?

Daily inventory variation time series

We quantify the trading activity of a firm in a given time period τ by introducing the inventory variation


$$v_i(t) \equiv \sum_{s=t}^{t+\tau} \epsilon_i(s) p_i(s) V_i(s)$$

signpricevolume



Diagram illustrating the components of the inventory variation formula:

- $\epsilon_i(s)$ is labeled "sign" (indicated by an arrow).
- $p_i(s)$ is labeled "price" (indicated by an arrow).
- $V_i(s)$ is labeled "volume" (indicated by an arrow).

- Inventory variation is a measure of the net buy/sell position of agent i



Cross correlation matrix of inventory variation





min=-0.53

max=0.75

Trading activity is significantly cross correlated among firms

Origin of collective behavior



- The first eigenvalue is not compatible with random trading and is therefore carrying information about the collective dynamics of firms.
- The corresponding factor is significantly correlated with price return.
- There are groups of firms having systematically the same position (buy/sell) as the other members of the group they belong to.

The role of size

Few large trending firms


JP Morgan,
Merryl Lynch,
Credit Suisse,,
Credit Agricole,
BNP Paribas,
UBS Warburg,
Société General

Undetermined
firms

BBVA
Santander


Renta 4,
Mercavalor

Many heterogeneous reversing firms



Size = average daily fraction of volume


Inventory variation
correlation matrix
obtained by sorting
the firms in the rows
and columns
according to their
correlation of
inventory variation
with price return



Correlated order flow

Inventory variation is long range correlated


$$\langle v_i(s+\tau) v_i(s) \rangle \approx \tau^{-\alpha}$$



Granger causality

$$I(X \rightarrow Y) = \begin{cases} 1 & \text{if } X \text{ Granger-causes } Y \\ 0 & \text{if } X \text{ does not Granger-cause } Y \end{cases}$$

- Returns cause inventory variations
- Inventory variations does not cause returns



Two models

- Roll model (1984): Distinction between price components due to fundamental security value and those attributable to the market organization and trading process
- Kyle model (1985): (strategic) Asymmetric information model describing how information is incorporated into prices

Roll model (1984)

- Random walk model of the efficient price

$$m_t = m_{t-1} + u_t$$

where u_t is an i.i.d. noise with variance σ_u^2

- All trades are conducted through a dealer who posts bid and ask prices
- Dealer incurs a noninformational cost c per trade (due to fixed costs)
- The bid and the ask are

$$b_t = m_t - c; \quad a_t = m_t + c$$

and thus the spread is $2c$

- The transaction price p_t is

$$p_t = m_t + q_t c$$

where $q_t = +1$ (-1) if the customer is buyer (seller).

Moreover q_t are assumed serially independent.


- The variance of transaction price increments is

$$Var(\Delta p_t) = 2c^2 + \sigma_u^2$$

- The autocovariance of transaction price increments is

$$\gamma_1 = Cov(\Delta p_{t-1} \Delta p_t) = -c^2$$

and it is zero for lag larger than one



Time series perspective

- The autocovariance structure of transaction price increments in the Roll model is the same as in a MA(1) process
- The Roll model is a structural model, while its time series representation is a statistical model
- By assuming covariance stationarity we can use the Wold theorem: any zero-mean covariance stationary process can be represented as

$$x_t = \sum_{j=0}^{\infty} \theta_j \varepsilon_{t-j} + \kappa_t$$

where ε_t is a zero mean white noise process and κ_t is a linearly deterministic process (can be predicted arbitrarily well by a linear projection on past observation of x_t)