Stylized Facts on Market Microstructure II

Sophie LARUELLE

LAMA-UPEC

March 20, 2026

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

Daily traded volume *vs* number of trades

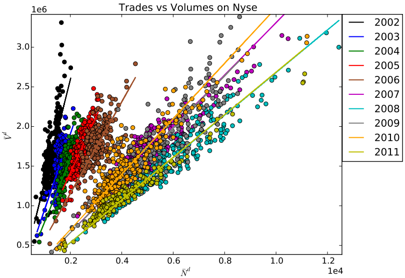

NYSE

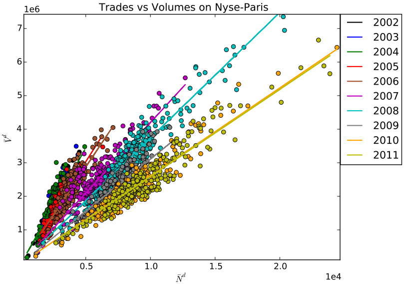

Euronext-Paris

Daily traded volume *vs* number of trades

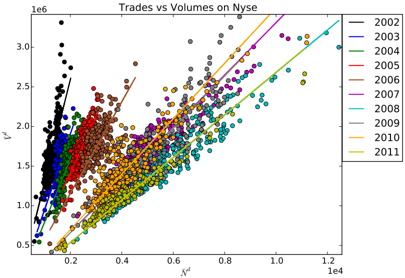

NYSE

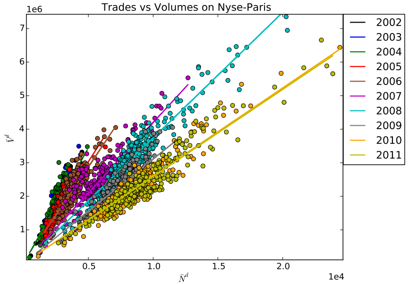

Euronext-Paris

- Globally the average trade size (ATS, namely the slope) has decreased through years because the number of trades regularly raises due to the development of electronic trading and the capability to trade faster and to slice volume across the day.

Daily traded volume *vs* number of trades

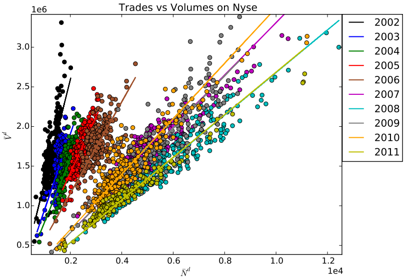

NYSE

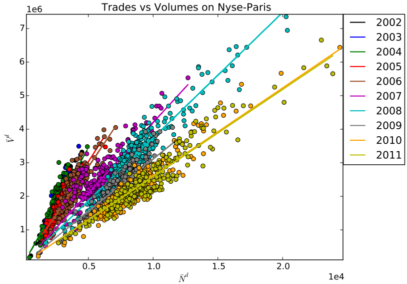

Euronext-Paris

- Globally the average trade size (ATS, namely the slope) has decreased through years because the number of trades regularly raises due to the development of electronic trading and the capability to trade faster and to slice volume across the day.
- Moreover, the total traded volume has decreased on primary markets according to fragmentation.

Daily turnover *vs* number of trades

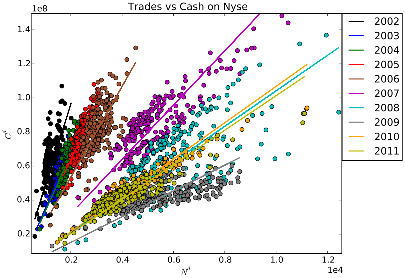

NYSE

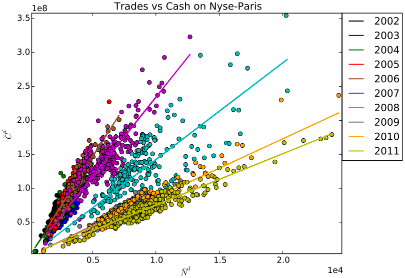

Euronext-Paris

Daily turnover *vs* number of trades

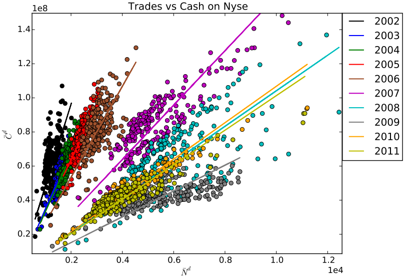

NYSE

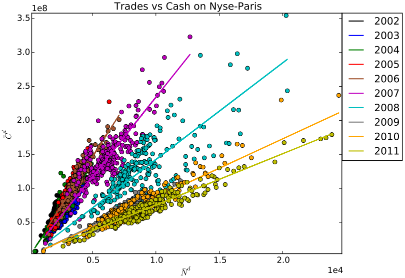

Euronext-Paris

The average trade turnover (ATT) allows us to separate the effect of the fragmentation (shown with the traded volume) and the financial crisis (which implies a downturn of prices).

Daily quantity best quotes *vs* traded volume

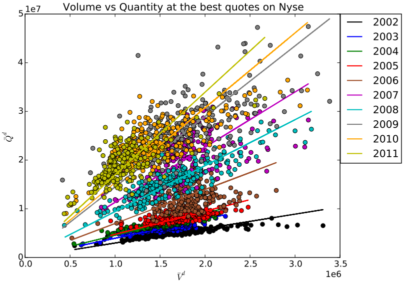

NYSE

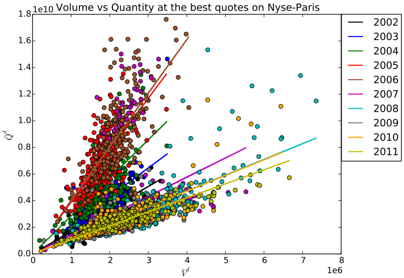

Euronext-Paris

Daily quantity best quotes *vs* traded volume

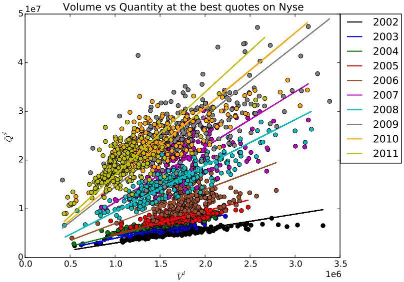

NYSE

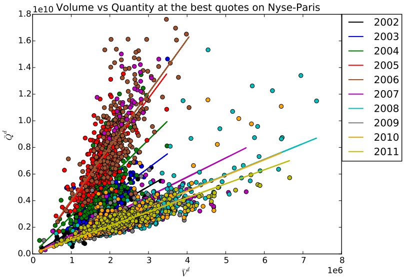

Euronext-Paris

- A measure of liquidity is the available quantity at the best quotes (best ask and best bid) when a market order arrives, it hits the first limit (on the bid side or on the ask side) to consume liquidity in the LOB.

Daily quantity best quotes *vs* traded volume

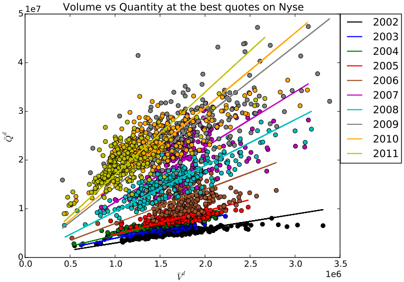

NYSE

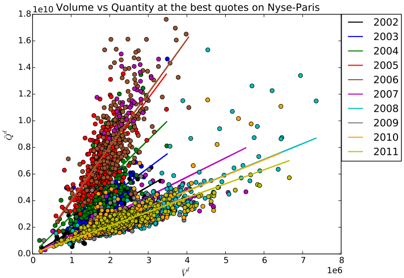

Euronext-Paris

- A measure of liquidity is the available quantity at the best quotes (best ask and best bid) when a market order arrives, it hits the first limit (on the bid side or on the ask side) to consume liquidity in the LOB.
- As tick size defines the grid of prices, it will set the degree of discreteness of traded prices and prices of limit orders.

Daily quantity best quotes *vs* traded volume

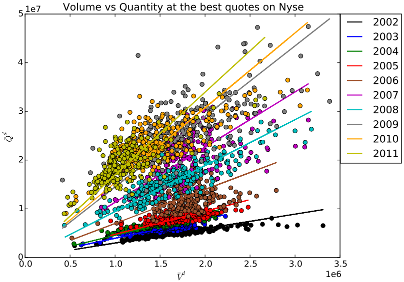

NYSE

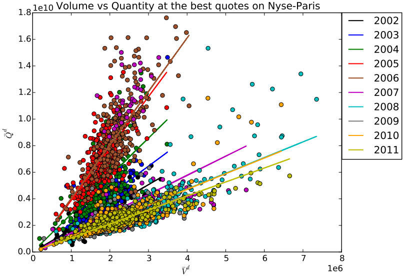

Euronext-Paris

- A measure of liquidity is the available quantity at the best quotes (best ask and best bid) when a market order arrives, it hits the first limit (on the bid side or on the ask side) to consume liquidity in the LOB.
- As tick size defines the grid of prices, it will set the degree of discreteness of traded prices and prices of limit orders.
- A smaller tick size will split limit orders over more price points. The quoted sizes on the best limits are therefore a lot smaller with a small tick size.

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.
- In practice, such information is correlated with an order flow that either anticipates or causes successive price changes.

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.
- In practice, such information is correlated with an order flow that either anticipates or causes successive price changes.
- Mechanically, the effect on price of a trade can be of two types according to the tick size:

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.
- In practice, such information is correlated with an order flow that either anticipates or causes successive price changes.
- Mechanically, the effect on price of a trade can be of two types according to the tick size:
  - For a large tick asset, the bid-ask spread is always equal to one tick, hence the average price change induced by a trade is fixed by the tick size and the probability of depleting a queue.

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.
- In practice, such information is correlated with an order flow that either anticipates or causes successive price changes.
- Mechanically, the effect on price of a trade can be of two types according to the tick size:
    - For a large tick asset, the bid-ask spread is always equal to one tick, hence the average price change induced by a trade is fixed by the tick size and the probability of depleting a queue.
    - Conversely, in the regime of small tick size, all the trades trigger a price change, so that the average price change is fixed by the average spread.

Volatility, bid-ask spread and liquidity

- The volatility can be seen as a measure quantifying how much information flows into a price.
- In practice, such information is correlated with an order flow that either anticipates or causes successive price changes.
- Mechanically, the effect on price of a trade can be of two types according to the tick size:
  - For a large tick asset, the bid-ask spread is always equal to one tick, hence the average price change induced by a trade is fixed by the tick size and the probability of depleting a queue.
  - Conversely, in the regime of small tick size, all the trades trigger a price change, so that the average price change is fixed by the average spread.
- In both cases, one has that a trade can induce a price change with positive probability.

Spread *vs* volatility per trade

- It is well-known that spread and volatility are intimately related : while the spread fixes the profit per trade of market makers, the volatility determines the adverse selection that they needs to face.

Spread *vs* volatility per trade

- It is well-known that spread and volatility are intimately related : while the spread fixes the profit per trade of market makers, the volatility determines the adverse selection that they needs to face.

- In any model in which market makers break even, one needs to have an equation fixing the relation between bid-ask spread and volatility. This relation has been the subject of several models trying to justify it microscopically, while empirically this has been analyzed in several works.

Spread *vs* volatility per trade

- It is well-known that spread and volatility are intimately related : while the spread fixes the profit per trade of market makers, the volatility determines the adverse selection that they needs to face.

- In any model in which market makers break even, one needs to have an equation fixing the relation between bid-ask spread and volatility. This relation has been the subject of several models trying to justify it microscopically, while empirically this has been analyzed in several works.

- In the simplest possible model of Madhavan, Richardson and Roomans (MRR), one obtains that

$$\varphi = c\sigma_{mid} N^{-1/2},$$

where φ is the spread, σ_{mid} is the mid-price volatility, N the number of trades and $c > 0$.

Volatility *vs* spread for small tick asssets

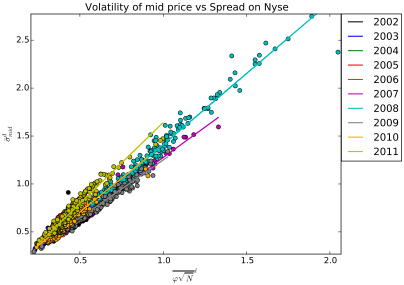

NYSE

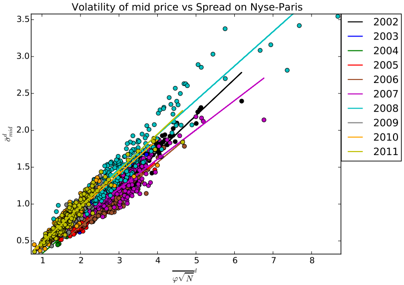

Euronext-Paris

Volatility *vs* spread for small tick asssets

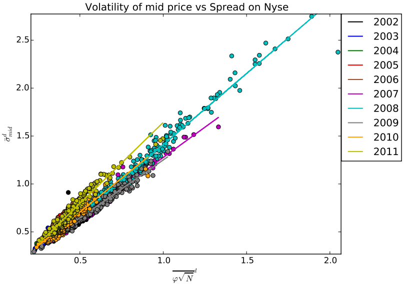

NYSE

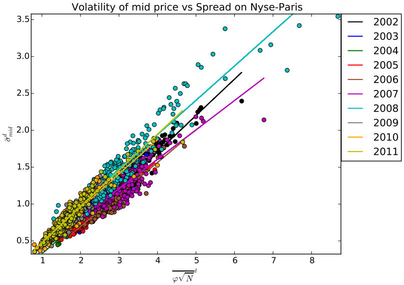

Euronext-Paris

For European and US markets, this relationship is quite stable with one year period 2008 with high volatility and spread (consequences of financial crisis).

Volatility *vs* spread for large tick asssets

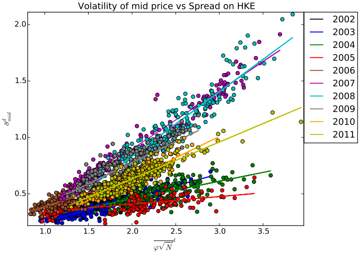

Hong-Kong

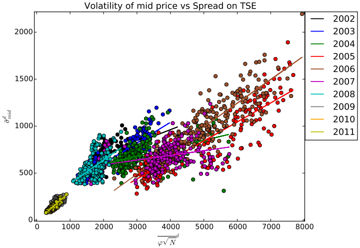

Tokyo

Volatility *vs* spread for large tick asssets

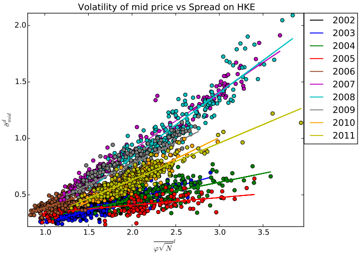

Hong-Kong

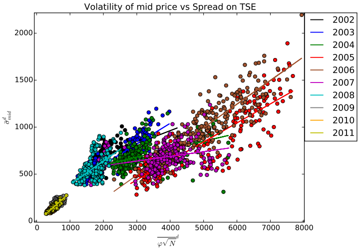

Tokyo

- On Asian markets, we observe an increasing in the slope for HKE until 2008 (the volatility increased more than the bid-ask spread, even if the number of trades increased) with a jump in 2007 due to financial crisis, and a mean reverting effect after 2008.

Volatility *vs* spread for large tick asssets

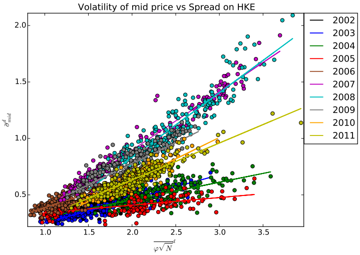

Hong-Kong

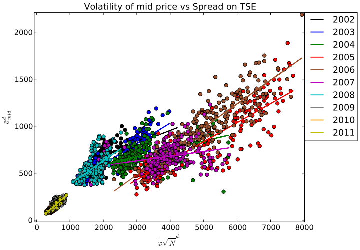

Tokyo

- On Asian markets, we observe an increasing in the slope for HKE until 2008 (the volatility increased more than the bid-ask spread, even if the number of trades increased) with a jump in 2007 due to financial crisis, and a mean reverting effect after 2008.
- For TSE, we identify clusters for each year : In 2005-2006, this period of high volatility and spread could be induce by technical problems of the platform (due to increase of trading) ; in 2009-2011, these variables fell down probably owing to improvements of the electronic platform and fragmentation.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple:

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick. This cutoff effect also implies an extra profit for the market makers due solely to liquidity reasons.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick. This cutoff effect also implies an extra profit for the market makers due solely to liquidity reasons.

Dayri and Rosenbaum developed a new approach that generalizes relationship between bid-ask spread and volatility per trade to the large tick regime.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick. This cutoff effect also implies an extra profit for the market makers due solely to liquidity reasons.

Dayri and Rosenbaum developed a new approach that generalizes relationship between bid-ask spread and volatility per trade to the large tick regime. The model derives from the one with uncertainty zones introduced by Robert and Rosenbaum, in which the tick size α confines the value of a continuous fair price to a discrete price grid.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick. This cutoff effect also implies an extra profit for the market makers due solely to liquidity reasons.

Dayri and Rosenbaum developed a new approach that generalizes relationship between bid-ask spread and volatility per trade to the large tick regime. The model derives from the one with uncertainty zones introduced by Robert and Rosenbaum, in which the tick size α confines the value of a continuous fair price to a discrete price grid. The areas of uncertainty are defined as bands around the middle of the tick size with width $2\eta\alpha$, with $0 < \eta < 1$.

Large tick stocks: The role of the implicit spread

The reason why the proportionality between spread and volatility fails for large tick stocks is simple: Even assuming a proportionality between σ_{mid} and φ, if one imagines to decrease continuously the volatility, the spread eventually will come to a point in which it will be equal to one tick. Then, further diminutions in volatility won't be mechanically able to push the spread below the tick. This cutoff effect also implies an extra profit for the market makers due solely to liquidity reasons.

Dayri and Rosenbaum developed a new approach that generalizes relationship between bid-ask spread and volatility per trade to the large tick regime. The model derives from the one with uncertainty zones introduced by Robert and Rosenbaum, in which the tick size α confines the value of a continuous fair price to a discrete price grid. The areas of uncertainty are defined as bands around the middle of the tick size with width $2\eta\alpha$, with $0 < \eta < 1$. The effect of these bands is such that for $\eta < 1/2$ the observed price displays microscopic mean-reversion, for $\eta > 1/2$ it displays microscopic trending patterns while for $\eta = 1/2$ no effect is induced.

Model with uncertainty zones: example of sample path

The efficient price is drawn in blue. The light gray lines drawn at integers form the tick grid of width α. The red dotted lines are the limits of the uncertainty zones of width $2\eta\alpha$. Finally the last traded price is the black stepwise curve. The circles indicate a change in the price when the efficient price crosses an uncertainty zone.

Bid/Ask Zones and Buy/Sell Zone

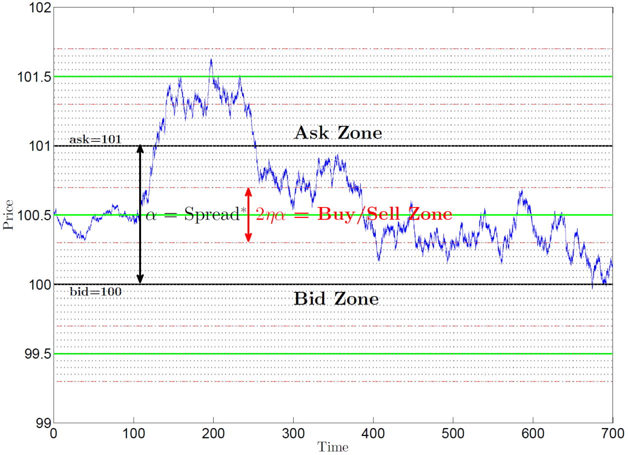

The three different zones when the bid-ask is 100-101 and the tick value is equal to one. The red dotted lines are the limits of the uncertainty zones. The uncertainty zone inside the spread is the buy/sell zone. The upper dotted area is the ask zone and the lower dotted area is the bid zone.

The η parameter can be estimated quite easily by using an estimator of the form

$$\widehat{\eta} = \frac{N^{(c)}}{2N^{(a)}},$$

where $N^{(a/c)}$ are respectively the number of alternations (subsequent price changes in the opposite direction) and continuations (subsequent price changes in the same direction).

The η parameter can be estimated quite easily by using an estimator of the form

$$\widehat{\eta} = \frac{N^{(c)}}{2N^{(a)}},$$

where $N^{(a/c)}$ are respectively the number of alternations (subsequent price changes in the opposite direction) and continuations (subsequent price changes in the same direction).

The volatility of the fair price $(x_t)_t$ is given by

$$\sigma_{fair} = \sqrt{\sum_t (x_t - x_{t-1})^2}, \quad x_t = p_t - \mathrm{sign}(p_t - p_{t-1})(1/2 - \eta)\alpha,$$

where $(p_t)_t$ is the traded price.

The η parameter can be estimated quite easily by using an estimator of the form

$$\widehat{\eta} = \frac{N^{(c)}}{2N^{(a)}},$$

where $N^{(a/c)}$ are respectively the number of alternations (subsequent price changes in the opposite direction) and continuations (subsequent price changes in the same direction).
The volatility of the fair price $(x_t)_t$ is given by

$$\sigma_{fair} = \sqrt{\sum_t (x_t - x_{t-1})^2}, \quad x_t = p_t - \text{sign}(p_t - p_{t-1})(1/2 - \eta)\alpha,$$

where $(p_t)_t$ is the traded price. The model predicts then the following relationship

$$\eta\alpha = \frac{\sigma_{fair}}{\sqrt{N}} + c, \quad N \text{ is the number of trades.}$$

The η parameter can be estimated quite easily by using an estimator of the form

$$\widehat{\eta} = \frac{N^{(c)}}{2N^{(a)}},$$

where $N^{(a/c)}$ are respectively the number of alternations (subsequent price changes in the opposite direction) and continuations (subsequent price changes in the same direction).

The volatility of the fair price $(x_t)_t$ is given by

$$\sigma_{fair} = \sqrt{\sum_t (x_t - x_{t-1})^2}, \quad x_t = p_t - \text{sign}(p_t - p_{t-1})(1/2 - \eta)\alpha,$$

where $(p_t)_t$ is the traded price. The model predicts then the following relationship

$$\eta\alpha = \frac{\sigma_{fair}}{\sqrt{N}} + c, \quad N \text{ is the number of trades.}$$

In fact, the quantity $2\eta\alpha$ can be seen effectively as an implicit spread.

Volatility per trade of fair price *vs* implicit spread

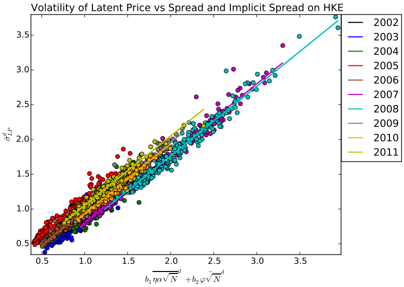

Hong-Kong

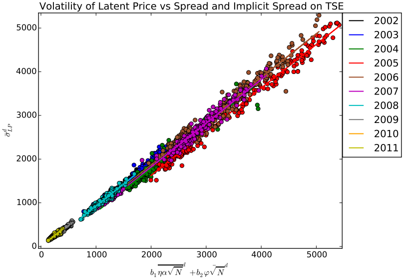

Tokyo

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

Volatility and Signature plot

A "natural" estimator for the integrated variance is

$$\hat{V}_R(\Delta) = \sum_{j=1}^{T/\Delta} (\ln(S_{\Delta j}) - \ln(S_{\Delta(j-1)}))^2 \xrightarrow[\Delta \to 0]{\mathbb{P}} \int_0^T \sigma_s^2 ds.$$

Volatility and Signature plot

A "natural" estimator for the integrated variance is

$$\hat{V}_R(\Delta) = \sum_{j=1}^{T/\Delta} \left(\ln(S_{\Delta j}) - \ln(S_{\Delta(j-1)})\right)^2 \xrightarrow[\Delta \to 0]{\mathbb{P}} \int_0^T \sigma_s^2 ds.$$

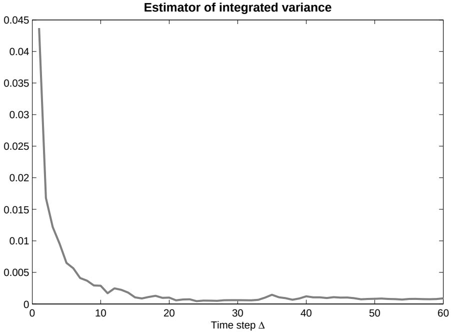

Estimator for the integrated variance computed for Rio Tinto between 03/01/2012 and 19/06/2012 according to the time discretization step.

Volatility estimation

They are many other ways to estimate the volatility: For example by using the minimum and the maximum of the price (see Garman & Klass), by considering multi-scale (see Ait-Sahalia or Ait-Sahalia & Jacod) or by bid-ask modeling (see Robert & Rosenbaum).

Volatility estimation

They are many other ways to estimate the volatility: For example by using the minimum and the maximum of the price (see Garman & Klass), by considering multi-scale (see Ait-Sahalia or Ait-Sahalia & Jacod) or by bid-ask modeling (see Robert & Rosenbaum).

Observations and microstructure. A model for log-price observation X may be $X = M + \varepsilon$, where M is a semi-martingale and ε is the noise. Consequently

$$r_\Delta(j) = M_{\Delta j} - M_{\Delta(j-1)} + \varepsilon_{\Delta j} - \varepsilon_{\Delta(j-1)} = r_\Delta^M(j) + \eta_\Delta(j).$$

Volatility estimation

They are many other ways to estimate the volatility: For example by using the minimum and the maximum of the price (see Garman & Klass), by considering multi-scale (see Ait-Sahalia or Ait-Sahalia & Jacod) or by bid-ask modeling (see Robert & Rosenbaum).

Observations and microstructure. A model for log-price observation X may be $X = M + \varepsilon$, where M is a semi-martingale and ε is the noise. Consequently

$$r_\Delta(j) = M_{\Delta j} - M_{\Delta(j-1)} + \varepsilon_{\Delta j} - \varepsilon_{\Delta(j-1)} = r^M_\Delta(j) + \eta_\Delta(j).$$

Thus the realized variance reads

$$\hat{V}_R(\Delta) = \sum_{j=1}^{T/\Delta} r_\Delta(j)^2 = \sum_{j=1}^{T/\Delta} r^M_\Delta(j)^2 + \sum_{j=1}^{T/\Delta} \eta_\Delta(j)^2 + \sum_{j=1}^{T/\Delta} r^M_\Delta(j)\eta_\Delta(j).$$

Volatility estimation

They are many other ways to estimate the volatility: For example by using the minimum and the maximum of the price (see Garman & Klass), by considering multi-scale (see Ait-Sahalia or Ait-Sahalia & Jacod) or by bid-ask modeling (see Robert & Rosenbaum).

Observations and microstructure. A model for log-price observation X may be $X = M + \varepsilon$, where M is a semi-martingale and ε is the noise. Consequently

$$r_\Delta(j) = M_{\Delta j} - M_{\Delta(j-1)} + \varepsilon_{\Delta j} - \varepsilon_{\Delta(j-1)} = r_\Delta^M(j) + \eta_\Delta(j).$$

Thus the realized variance reads

$$\hat{V}_R(\Delta) = \sum_{j=1}^{T/\Delta} r_\Delta(j)^2 = \sum_{j=1}^{T/\Delta} r_\Delta^M(j)^2 + \sum_{j=1}^{T/\Delta} \eta_\Delta(j)^2 + \sum_{j=1}^{T/\Delta} r_\Delta^M(j)\eta_\Delta(j).$$

This estimator **diverges** when $\Delta \to 0$.

Volatility Garman-Klass (GK)

They consider consider intraday volatility estimators that are based upon the historical opening, closing, high, and low prices and transaction volume. Since high and low prices require continuous monitoring to obtain, they correspondingly contain superior information content, exploited herein.

Volatility Garman-Klass (GK)

They consider consider intraday volatility estimators that are based upon the historical opening, closing, high, and low prices and transaction volume. Since high and low prices require continuous monitoring to obtain, they correspondingly contain superior information content, exploited herein.

Their model assumes that a diffusion process governs security prices

$$P_t = \Phi(X_t)$$

where P is the security price, t is time, Φ is a monotonic, time-independent transformation, and X_t is a diffusion process with the differential representation

$$dX_t = \sigma dW_t$$

where W_t is the standard Brownian motion and σ is an unknown constant to be estimated.

Volatility Garman-Klass (GK)

"Best" Analytic Scale-Invariant Estimators. They consider estimators depending on : the opening price, the closing price, the high price and the low price.

Volatility Garman-Klass (GK)

"Best" Analytic Scale-Invariant Estimators. They consider estimators depending on : the opening price, the closing price, the high price and the low price.

For the authors, an estimator is "best" when it has minimum variance and is unbiased.

Volatility Garman-Klass (GK)

“Best” Analytic Scale-Invariant Estimators. They consider estimators depending on : the opening price, the closing price, the high price and the low price.

For the authors, an estimator is “best” when it has minimum variance and is unbiased.

They also impose the requirements that the estimators be analytic with price and time symmetries and scale-invariant.

Volatility Garman-Klass (GK)

“Best” Analytic Scale-Invariant Estimators. They consider estimators depending on : the opening price, the closing price, the high price and the low price.

For the authors, an estimator is “best” when it has minimum variance and is unbiased.

They also impose the requirements that the estimators be analytic with price and time symmetries and scale-invariant.

Under regularity condition (analytic in a neighborhood of the origin), they obtain that the estimator must be quadratic in its arguments.

Volatility Garman-Klass (GK)

"Best" Analytic Scale-Invariant Estimators. They consider estimators depending on : the opening price, the closing price, the high price and the low price.

For the authors, an estimator is "best" when it has minimum variance and is unbiased.

They also impose the requirements that the estimators be analytic with price and time symmetries and scale-invariant.

Under regularity condition (analytic in a neighborhood of the origin), they obtain that the estimator must be quadratic in its arguments.

Owing scale invariance and analyticity, they reduce the problem and find the "best" estimator of this form (unbiased with minimum variance).

Volatility Garman-Klass (GK)

$$\sigma_{GK}^2 = \frac{1}{2}(\max S - \min S)^2 - (2\ln 2 - 1)(S_{end} - S_{begin})^2.$$

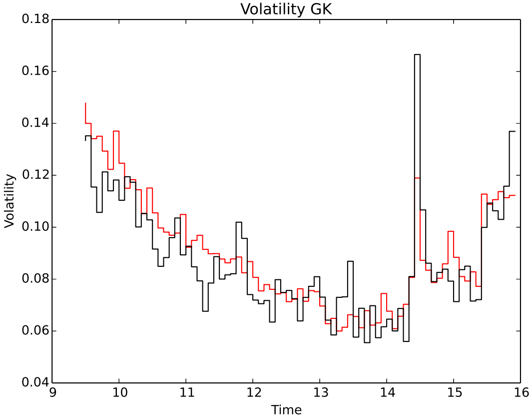

Averaging volatility GK (from Monday to Thursday in red, on Friday in black) for Alstom in 2011.

“Classical” Covariance Estimator

Let $X^1 = \log P^1$ and $X^2 = \log P^2$ be two log-price processes.

$$dX^1 = \mu_t^1 dt + \sigma_{t-}^1 dW_t^1,$$

$$dX^2 = \mu_t^2 dt + \sigma_{t-}^2 dW_t^2,$$

with $d\left\langle W^1, W^2 \right\rangle_t = \rho_t dt$.

"Classical" Covariance Estimator

Let $X^1 = \log P^1$ and $X^2 = \log P^2$ be two log-price processes.

$$dX^1 = \mu_t^1 dt + \sigma_{t-}^1 dW_t^1,$$

$$dX^2 = \mu_t^2 dt + \sigma_{t-}^2 dW_t^2,$$

with $d\left\langle W^1, W^2 \right\rangle_t = \rho_t dt$.

Set for every $j \in 1, \ldots, T/\Delta$, $r_\Delta^1(j) = X_{\Delta j}^1 - X_{\Delta(j-1)}^1$ and $r_\Delta^2(j) = X_{\Delta j}^2 - X_{\Delta(j-1)}^2$ the log-price increments.

"Classical" Covariance Estimator

Let $X^1 = \log P^1$ and $X^2 = \log P^2$ be two log-price processes.

$$dX^1 = \mu_t^1 dt + \sigma_{t-}^1 dW_t^1,$$

$$dX^2 = \mu_t^2 dt + \sigma_{t-}^2 dW_t^2,$$

with $d\left\langle W^1, W^2 \right\rangle_t = \rho_t dt$.

Set for every $j \in 1, \ldots, T/\Delta$, $r_\Delta^1(j) = X_{\Delta j}^1 - X_{\Delta(j-1)}^1$ and $r_\Delta^2(j) = X_{\Delta j}^2 - X_{\Delta(j-1)}^2$ the log-price increments.

Then the integrated covariance between the two assets reads

$$\int_0^T \rho_s \sigma_s^1 \sigma_s^2 ds = \lim_{\Delta \to 0} \sum_{j=1}^{T/\Delta} r_\Delta^1(j) r_\Delta^2(j).$$

"Classical" Covariance Estimator

Thus, an estimator for the realized covariance may be written as

$$\hat{C}_R(\Delta) = \sum_{j=1}^{T/\Delta} r_\Delta^1(j) r_\Delta^2(j).$$

“Classical” Covariance Estimator

Thus, an estimator for the realized covariance may be written as

$$\hat{C}_R(\Delta) = \sum_{j=1}^{T/\Delta} r^1_\Delta(j) r^2_\Delta(j).$$

This estimator is consistent since

$$\hat{C}_R(\Delta) \xrightarrow[\Delta \to 0]{\mathbb{P}} \int_0^T \rho_s \sigma^1_s \sigma^2_s ds.$$

“Classical” Covariance Estimator

Thus, an estimator for the realized covariance may be written as

$$\hat{C}_R(\Delta) = \sum_{j=1}^{T/\Delta} r^1_\Delta(j) r^2_\Delta(j).$$

This estimator is consistent since

$$\hat{C}_R(\Delta) \xrightarrow[\Delta \to 0]{\mathbb{P}} \int_0^T \rho_s \sigma^1_s \sigma^2_s ds.$$

The problem is that we must have synchronous data, but the quotations and transactions are **asynchronous**.

Epps effect

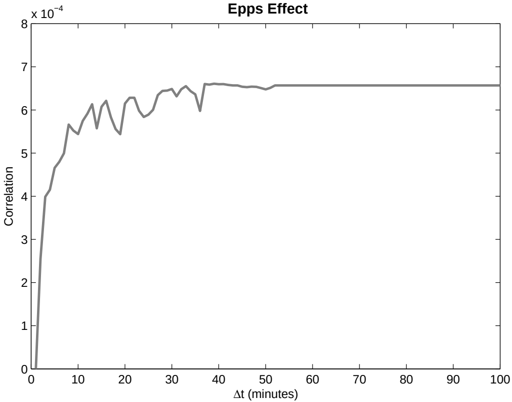

Estimator for correlation between Total and FTE on Euronext Paris from 23/07/2012 to 27/07/2012 as a function of time discretization step.

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

- systematic bias for this estimator,

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

- systematic bias for this estimator,
- "lead-lag" effect for the assets in the same sector,

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

- systematic bias for this estimator,
- "lead-lag" effect for the assets in the same sector,
- asynchronicity of transactions,

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

- systematic bias for this estimator,
- "lead-lag" effect for the assets in the same sector,
- asynchronicity of transactions,
- "tick" effect and other microstructure effects.

Epps Effect

Epps (1979): "Correlations among price changes [...] are found to decrease with the length of the interval for which the price changes are measured."

Many explanations were proposed

- systematic bias for this estimator,
- "lead-lag" effect for the assets in the same sector,
- asynchronicity of transactions,
- "tick" effect and other microstructure effects.

Example. Assume that X^1 and X^2 are two Brownian motions with correlation ρ and the trade times are arrival times of two independent Poisson processes. Then, one can show that

$$\mathbb{E}[\bar{C}_R(\Delta)] \underset{\Delta \to 0}{\longrightarrow} 0.$$

The correlation estimator for high-frequency data is studied in many publications, for example Hayashi & Yoshida or Zhang.

Hayashi-Yoshida Estimator

Let $I_i^1 = [T_i^1, T_{i+1}^1)$, $i \geq 1$, and $I_j^2 = [T_j^2, T_{j+1}^2)$, $j \geq 1$.

Hayashi-Yoshida Estimator

Let $I_i^1 = [T_i^1, T_{i+1}^1)$, $i \geq 1$, and $I_j^2 = [T_j^2, T_{j+1}^2)$, $j \geq 1$.

The cumulative covariance estimator of Hayashi-Yoshida reads

$$U_n = \sum_{i,j} \Delta P^1(I_i^1)\Delta P^2(I_j^2)\mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}.$$

Hayashi-Yoshida Estimator

Let $I_i^1 = [T_i^1, T_{i+1}^1)$, $i \geq 1$, and $I_j^2 = [T_j^2, T_{j+1}^2)$, $j \geq 1$.

The cumulative covariance estimator of Hayashi-Yoshida reads

$$U_n = \sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}.$$

That is, the product of any pair of increments $\Delta P^1(I_i^1)$ and $\Delta P^2(I_j^2)$ will make a contribution to the summation only when the respective observation intervals I_i^1 and I_i^2 are overlapping.

Hayashi-Yoshida Estimator

Let $I_i^1 = [T_i^1, T_{i+1}^1)$, $i \geq 1$, and $I_j^2 = [T_j^2, T_{j+1}^2)$, $j \geq 1$.

The cumulative covariance estimator of Hayashi-Yoshida reads

$$U_n = \sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}.$$

That is, the product of any pair of increments $\Delta P^1(I_i^1)$ and $\Delta P^2(I_j^2)$ will make a contribution to the summation only when the respective observation intervals I_i^1 and I_i^2 are overlapping.

This estimator does not need any selection of Δ and is convergent if the arrival times are independent from the price.

Hayashi-Yoshida Estimator

Then, they define two associated correlation estimators

Hayashi-Yoshida Estimator

Then, they define two associated correlation estimators

- if σ_1 and σ_2 are known,

$$R_n^1 = \frac{1}{T}\frac{\sum_{i,j} \Delta P^1(I_i^1)\Delta P^2(I_j^2)\mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sigma_1\sigma_2},$$

Hayashi-Yoshida Estimator

Then, they define two associated correlation estimators

- if σ_1 and σ_2 are known,

$$R_n^1 = \frac{1}{T} \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sigma_1 \sigma_2},$$

- if σ_1 and σ_2 are known/unknown,

$$R_n^2 = \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sqrt{\sum_i \Delta P^1(I_i^1)^2} \sqrt{\sum_j \Delta P^2(I_j^2)^2}}.$$

Hayashi-Yoshida Estimator

Then, they define two associated correlation estimators

- if σ_1 and σ_2 are known,

$$R_n^1 = \frac{1}{T} \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sigma_1 \sigma_2},$$

- if σ_1 and σ_2 are known/unknown,

$$R_n^2 = \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sqrt{\sum_i \Delta P^1(I_i^1)^2} \sqrt{\sum_j \Delta P^2(I_j^2)^2}}.$$

Under the same conditions as for U_n, R^1 and R^2 are consistent for ρ as $n \to \infty$.

Hayashi-Yoshida Estimator

Then, they define two associated correlation estimators

- if σ_1 and σ_2 are known,

$$R_n^1 = \frac{1}{T} \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sigma_1 \sigma_2},$$

- if σ_1 and σ_2 are known/unknown,

$$R_n^2 = \frac{\sum_{i,j} \Delta P^1(I_i^1) \Delta P^2(I_j^2) \mathbb{1}_{\{I_i^1 \cap I_j^2 \neq \emptyset\}}}{\sqrt{\sum_i \Delta P^1(I_i^1)^2} \sqrt{\sum_j \Delta P^2(I_j^2)^2}}.$$

Under the same conditions as for U_n, R^1 and R^2 are consistent for ρ as $n \to \infty$.

Remarque. This estimator is nevertheless not robust to microstructure effects.

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

- **Market-to-Limit (MtoL).**
  - We send a buying LO at a price that crosses the spread e.g. the best ask.$\Rightarrow$ We hit the ask of the LOB, as in a traditional buy MO.
  - We consume all the liquidity available at the best ask.
  - All non-executed shares will remain as a buy LO. $\Rightarrow$ Our remaining shares at the old best ask become the new best bid.
  - Beware: risk of incomplete execution.
  - We can specify a limit price that consumes several levels of the LOB.

- **Market-to-Limit (MtoL).**
  - We send a buying LO at a price that crosses the spread e.g. the best ask.$\Rightarrow$ We hit the ask of the LOB, as in a traditional buy MO.
  - We consume all the liquidity available at the best ask.
  - All non-executed shares will remain as a buy LO. $\Rightarrow$ Our remaining shares at the old best ask become the new best bid.
  - Beware: risk of incomplete execution.
  - We can specify a limit price that consumes several levels of the LOB.
- **Immediate-or-Cancel (IoC) or Fill-and-Kill (FaK).**
  - We send a MtoL order.
  - We consume all the liquidity available up to the specified price.
  - All non-executed shares will be immediately canceled.
  - Beware: risk of incomplete execution.

- **Market-to-Limit (MtoL).**
  - We send a buying LO at a price that crosses the spread e.g. the best ask.$\Rightarrow$ We hit the ask of the LOB, as in a traditional buy MO.
  - We consume all the liquidity available at the best ask.
  - All non-executed shares will remain as a buy LO. $\Rightarrow$ Our remaining shares at the old best ask become the new best bid.
  - Beware: risk of incomplete execution.
  - We can specify a limit price that consumes several levels of the LOB.
- **Immediate-or-Cancel (IoC) or Fill-and-Kill (FaK).**
  - We send a MtoL order.
  - We consume all the liquidity available up to the specified price.
  - All non-executed shares will be immediately canceled.
  - Beware: risk of incomplete execution.
- **Fill-or-Kill (FoK).**
  - We send a MtoL order.
  - We wait for a confirmation that our order can be fully executed.
    - If the order can be fully executed $\Rightarrow$ OK.
    - If the order cannot be fully executed $\Rightarrow$ we cancel the order.

  $\Rightarrow$ FoK = FaK + 100

- **Peg order.** It is a LO that always has price priority.
  - We send a buy LO at the best bid b_0.
  - Suppose someone sends a LO and improves the best bid, say to $b_0 + \delta$:
    - We have just lost the price priority.
    - To recover it, we cancel our buy LO at b_0 and send a new buy LO at $b_0 + \delta$.
  - We repeat the process until full execution.
  - Beware: risk of incomplete execution.

- **Peg order.** It is a LO that always has price priority.
  - We send a buy LO at the best bid b_0.
  - Suppose someone sends a LO and improves the best bid, say to $b_0 + \delta$:
    - We have just lost the price priority.
    - To recover it, we cancel our buy LO at b_0 and send a new buy LO at $b_0 + \delta$.
  - We repeat the process until full execution.
  - Beware: risk of incomplete execution.
- **Iceberg order.**
  - An iceberg order of (say) 500 shares consists of two parts:
    - A visible part of (say) 100 shares.
    - A hidden part of (say) 400 shares.
  - We place the visible part in the LOB via LO at price p_0.
  - When the visible part is executed, we send a second LO of 100 shares at p_0.
  - When the second LO is executed, we send a third LO of 100 shares at p_0.
  - We repeat the iceberg game until full execution: In our example, until the fifth LO of 100 shares at p_0 is executed.
  - Beware: risk of incomplete execution.

- **Stop-Loss.** It is the most basic algorithm for risk management of a position:
  - We build up our long position.
  - We set a maximal loss for our portfolio. $\Rightarrow$ Set a minimal price or lower bound for our portfolio.
  - Portfolio value hits lower bound $\Rightarrow$ we get out of the position: we sell everything.

- **Stop-Loss.** It is the most basic algorithm for risk management of a position:
  - We build up our long position.
  - We set a maximal loss for our portfolio. $\Rightarrow$ Set a minimal price or lower bound for our portfolio.
  - Portfolio value hits lower bound $\Rightarrow$ we get out of the position: we sell everything.

- **Take-Profit.** It is the second most basic algorithm for risk management of a position:
  - We build up our long position.
  - We set a maximal gain for our portfolio. $\Rightarrow$ Set a maximal price or upper bound for our portfolio.
  - Portfolio value hits upper bound $\Rightarrow$ we get out of the position.

Stop-Loss & Take Profit

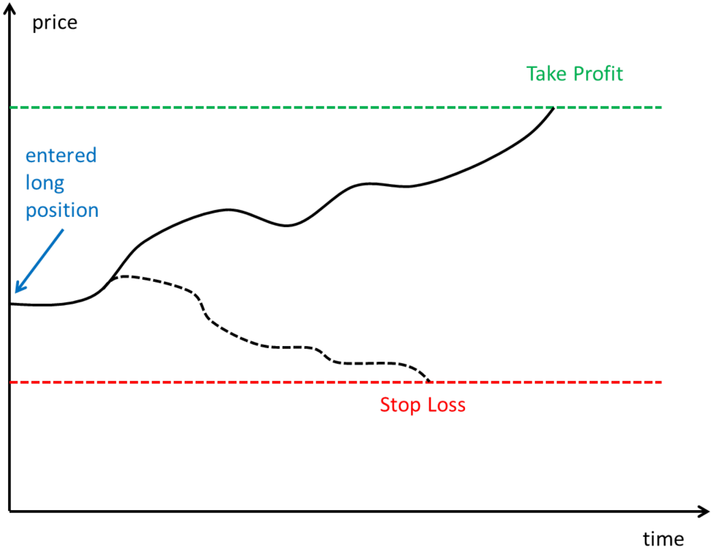

Sketch of Stop-Loss and Take-Profit.

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

Order Slicing

What to do when my order is too big?

Order Slicing

What to do when my order is too big?

- Say we wish to clear a portfolio of 1% of total market cap of a stock S:
  But $0.1 - 1\%$ is the daily turnover for highly liquid stocks.
  $\Rightarrow$ Impossible to dump that quantity directly on the market.

Order Slicing

What to do when my order is too big?

- Say we wish to clear a portfolio of 1% of total market cap of a stock S:
  But $0.1 - 1\%$ is the daily turnover for highly liquid stocks.
  $\Rightarrow$ Impossible to dump that quantity directly on the market.
- There are two ways of dealing with that:
  - Block order: we find a counterparty via a broker to clear the portfolio.
  - Slice and dice: we cut the original order into several child orders, and send them one by one.

Order Slicing

What to do when my order is too big?

- Say we wish to clear a portfolio of 1% of total market cap of a stock S:
  But $0.1 - 1\%$ is the daily turnover for highly liquid stocks.
  $\Rightarrow$ Impossible to dump that quantity directly on the market.
- There are two ways of dealing with that:
  - Block order: we find a counterparty via a broker to clear the portfolio.
  - Slice and dice: we cut the original order into several child orders, and send them one by one.
- There are pros and cons for each strategy:
  - With block orders your broker works for you, but you are revealing information: ex-ante to the broker, ex-post to the market.
  - With slice and dice you are hiding your information from the market, but you have to do the trades yourself.

Order Slicing

What to do when my order is too big?

- Say we wish to clear a portfolio of 1% of total market cap of a stock S:
  But $0.1 - 1\%$ is the daily turnover for highly liquid stocks.
  $\Rightarrow$ Impossible to dump that quantity directly on the market.
- There are two ways of dealing with that:
  - Block order: we find a counterparty via a broker to clear the portfolio.
  - Slice and dice: we cut the original order into several child orders, and send them one by one.
- There are pros and cons for each strategy:
  - With block orders your broker works for you, but you are revealing information: ex-ante to the broker, ex-post to the market.
  - With slice and dice you are hiding your information from the market, but you have to do the trades yourself.

$\Rightarrow$ Algorithmic trading and optimal execution: what is the best way to slice and dice?

Time-Weighted Average Price (TWAP)

- **Definition of the algorithm.** Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.
  - We choose how many orders we want to send, say one every 10 minutes.
    $\Rightarrow$ We divide the time period from 09:00 to 17:30 into 51 small periods of 10 minutes each.
  - Every 10 minutes we send an order to buy $20,000 \approx 51 \times 392$ shares.

Time-Weighted Average Price (TWAP)

- **Definition of the algorithm.** Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.
  - We choose how many orders we want to send, say one every 10 minutes.
    $\Rightarrow$ We divide the time period from 09:00 to 17:30 into 51 small periods of 10 minutes each.
  - Every 10 minutes we send an order to buy $20,000 \approx 51 \times 392$ shares.

- **Properties of the algorithm.**
  - The TWAP is extremely predictable: same order (buy 392 shares) periodically sent to the market (every 10 minutes).
  - To avoid being detected by the market, one can add some randomness:
    - ⋆ Random number of shares in every trade: $392 \pm \varepsilon$ shares.
    - ⋆ Each of the 51 time periods of random time length: $10 \pm \tau$ minutes.

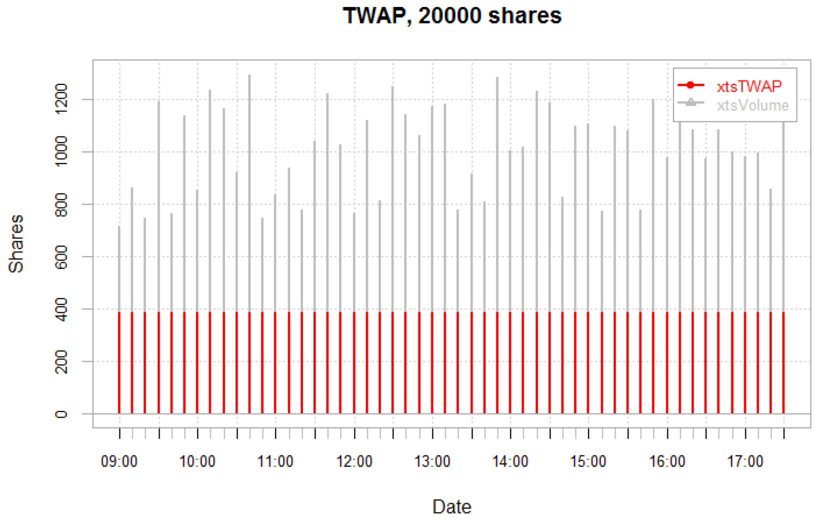

TWAP. Execution is constant throughout the whole day.

Percentage of Volume (PoV)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, starting at 09:00.

Percentage of Volume (PoV)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, starting at 09:00.

- We compute the intraday volume curve: For each day we compute the number of traded shares every 10 minutes from 09:00 to 17:30 (51 periods); we compute the average on each period for the last month.

Percentage of Volume (PoV)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, starting at 09:00.

- We compute the intraday volume curve: For each day we compute the number of traded shares every 10 minutes from 09:00 to 17:30 (51 periods); we compute the average on each period for the last month.
- We set a market share, *i.e.* the percentage of the whole market volume we will trade every 10 minutes. Let us suppose this is 50%.

Percentage of Volume (PoV)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, starting at 09:00.

- We compute the intraday volume curve: For each day we compute the number of traded shares every 10 minutes from 09:00 to 17:30 (51 periods); we compute the average on each period for the last month.
- We set a market share, *i.e.* the percentage of the whole market volume we will trade every 10 minutes. Let us suppose this is 50%.
- We trade 50% of the (historical) average volume every 10 minutes.

Percentage of Volume (PoV)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, starting at 09:00.

- We compute the intraday volume curve: For each day we compute the number of traded shares every 10 minutes from 09:00 to 17:30 (51 periods); we compute the average on each period for the last month.
- We set a market share, *i.e.* the percentage of the whole market volume we will trade every 10 minutes. Let us suppose this is 50%.
- We trade 50% of the (historical) average volume every 10 minutes.
- The final time of the algorithm is not defined a priori: it is part of the solution.

Percentage of Volume (PoV)

Properties of the algorithm.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve: Robust statistical measures are needed.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve: Robust statistical measures are needed.
- It is a static algorithm: the number of shares to trade is set before the execution and does not change.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve: Robust statistical measures are needed.
- It is a static algorithm: the number of shares to trade is set before the execution and does not change.
- To add some dynamic adaptation, add price targets and exible market shares:
  - A range of market share $50 \pm \varepsilon\%$.
  - A price range $p_0 \pm x$ EUR.
  - A dynamic intensity: more (less) participation at low (high) prices.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve: Robust statistical measures are needed.
- It is a static algorithm: the number of shares to trade is set before the execution and does not change.
- To add some dynamic adaptation, add price targets and exible market shares:
  - A range of market share $50 \pm \varepsilon\%$.
  - A price range $p_0 \pm x$ EUR.
  - A dynamic intensity: more (less) participation at low (high) prices.
- *Example*: suppose S opened the day at 100 EUR.
  - If $p \in [98; 102] \Rightarrow 50\%$ participation.
  - If $p < 98 \Rightarrow 75\%$ participation.
  - If $p > 102 \Rightarrow 25\%$ participation.

Percentage of Volume (PoV)

Properties of the algorithm.

- The PoV is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve: Robust statistical measures are needed.
- It is a static algorithm: the number of shares to trade is set before the execution and does not change.
- To add some dynamic adaptation, add price targets and exible market shares:
  - A range of market share $50 \pm \varepsilon\%$.
  - A price range $p_0 \pm x$ EUR.
  - A dynamic intensity: more (less) participation at low (high) prices.
- *Example*: suppose S opened the day at 100 EUR.
  - If $p \in [98; 102] \Rightarrow 50\%$ participation.
  - If $p < 98 \Rightarrow 75\%$ participation.
  - If $p > 102 \Rightarrow 25\%$ participation.
- *Beware*: the dynamic PoV has a risk of incomplete execution.

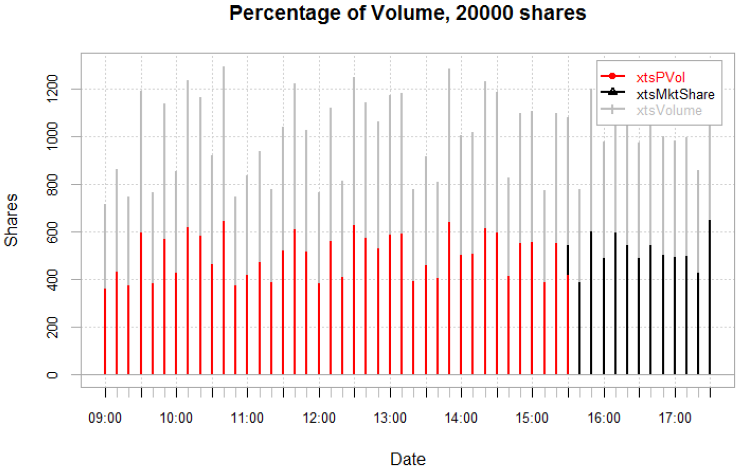

PoV with market share at 50%. The ending time 15:30 is part of the solution.

Volume-Weighted Average Price (VWAP)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.

Volume-Weighted Average Price (VWAP)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.

- We compute the intraday volume curve every 10 minutes :

$$(V_1, V_2, \dots, V_N), \quad N = 51.$$

Volume-Weighted Average Price (VWAP)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.

- We compute the intraday volume curve every 10 minutes :

$$(V_1, V_2, \dots, V_N), \quad N = 51.$$

- We compute the weights at every period n:

$$w_n := \frac{V_n}{\sum_{n=1}^{N} V_n}, \quad n = 1, \dots, N, \quad \sum_{n=1}^{N} w_n = 1.$$

Volume-Weighted Average Price (VWAP)

Definition of the algorithm. Suppose we want to buy 20,000 shares of S in Europe, from 09:00 to 17:30.

- We compute the intraday volume curve every 10 minutes :

$$(V_1, V_2, \dots, V_N), \quad N = 51.$$

- We compute the weights at every period n:

$$w_n := \frac{V_n}{\sum_{n=1}^{N} V_n}, \quad n = 1, \dots, N, \quad \sum_{n=1}^{N} w_n = 1.$$

- We define the number of shares to buy at every period n as:

$$v_n := 20,000 w_n, \quad n = 1, \dots, N.$$

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.
- It can be more dynamic with price targets.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.
- It can be more dynamic with price targets.
- The dynamic VWAP has a risk of incomplete execution.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.
- It can be more dynamic with price targets.
- The dynamic VWAP has a risk of incomplete execution.

However, it is more flexible than PoV:

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.
- It can be more dynamic with price targets.
- The dynamic VWAP has a risk of incomplete execution.

However, it is more flexible than PoV:

- The final time N of the execution is a parameter we can play with.

Volume-Weighted Average Price (VWAP)

Properties of the algorithm. The VWAP shares some of the properties of the PoV:

- It is less predictable than the TWAP.
- It is extremely dependent on the intraday volume curve.
- It is a static algorithm.
- It can be more dynamic with price targets.
- The dynamic VWAP has a risk of incomplete execution.

However, it is more flexible than PoV:

- The final time N of the execution is a parameter we can play with.
- Different $N \Rightarrow$ different trading pattern.

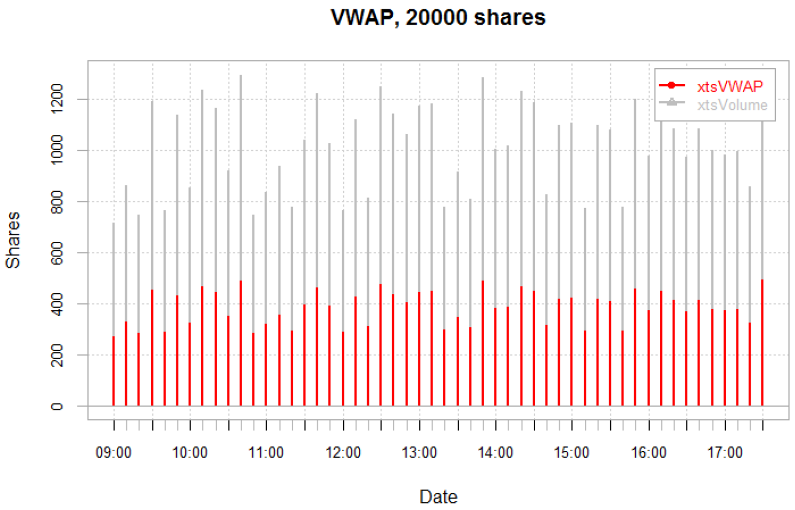

VWAP. Execution follows the historical (average) volume curve.

Examples of trading algorithms

Examples of trading algorithms

- **Time-Weighted Average Price (TWAP)** : same size at a chosen frequency $\Longrightarrow$ extremely predictable. To avoid being detected by the market, one can add some randomness (on both frequency and size).

Examples of trading algorithms

- **Time-Weighted Average Price (TWAP)** : same size at a chosen frequency $\Longrightarrow$ extremely predictable. To avoid being detected by the market, one can add some randomness (on both frequency and size).
- **Percentage of Volume (PoV)** : trade a fixed % of the market volume traded in a slice basing on historical volume curves $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, dynamic adaptation, risk of incomplete execution.

Examples of trading algorithms

- **Time-Weighted Average Price (TWAP)** : same size at a chosen frequency $\Longrightarrow$ extremely predictable. To avoid being detected by the market, one can add some randomness (on both frequency and size).
- **Percentage of Volume (PoV)** : trade a fixed % of the market volume traded in a slice basing on historical volume curves $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, dynamic adaptation, risk of incomplete execution.
- **Volume-Weighted Average Price (VWAP)** : send on every slice of the day (*e.g.* 10 minutes) a % of the whole volume equals to the % of market volume traded on this slice $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, risk of incomplete execution but more flexible than PoV.

Examples of trading algorithms

- **Time-Weighted Average Price (TWAP)** : same size at a chosen frequency $\Longrightarrow$ extremely predictable. To avoid being detected by the market, one can add some randomness (on both frequency and size).
- **Percentage of Volume (PoV)** : trade a fixed % of the market volume traded in a slice basing on historical volume curves $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, dynamic adaptation, risk of incomplete execution.
- **Volume-Weighted Average Price (VWAP)** : send on every slice of the day (*e.g.* 10 minutes) a % of the whole volume equals to the % of market volume traded on this slice $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, risk of incomplete execution but more flexible than PoV.
- **Implemented Shortfall** : the benchmark is the price at the moment when the execution starts $\Longrightarrow$ Almgren-Chriss, optimal control.

Examples of trading algorithms

- **Time-Weighted Average Price (TWAP)** : same size at a chosen frequency $\Longrightarrow$ extremely predictable. To avoid being detected by the market, one can add some randomness (on both frequency and size).
- **Percentage of Volume (PoV)** : trade a fixed % of the market volume traded in a slice basing on historical volume curves $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, dynamic adaptation, risk of incomplete execution.
- **Volume-Weighted Average Price (VWAP)** : send on every slice of the day (*e.g.* 10 minutes) a % of the whole volume equals to the % of market volume traded on this slice $\Longrightarrow$ dependent on the intraday volume curve, static algorithm, risk of incomplete execution but more flexible than PoV.
- **Implemented Shortfall** : the benchmark is the price at the moment when the execution starts $\Longrightarrow$ Almgren-Chriss, optimal control.
- **Target Close** : the benchmark is the price at the moment when the execution ends.

Outline

1. Relationship between daily indicators
2. Estimation of volatility and correlation in microstructure
3. Trading Algorithms
   - Simple trading algorithms
   - Complex trading algorithms
4. Flash Crash: May 6, 2010

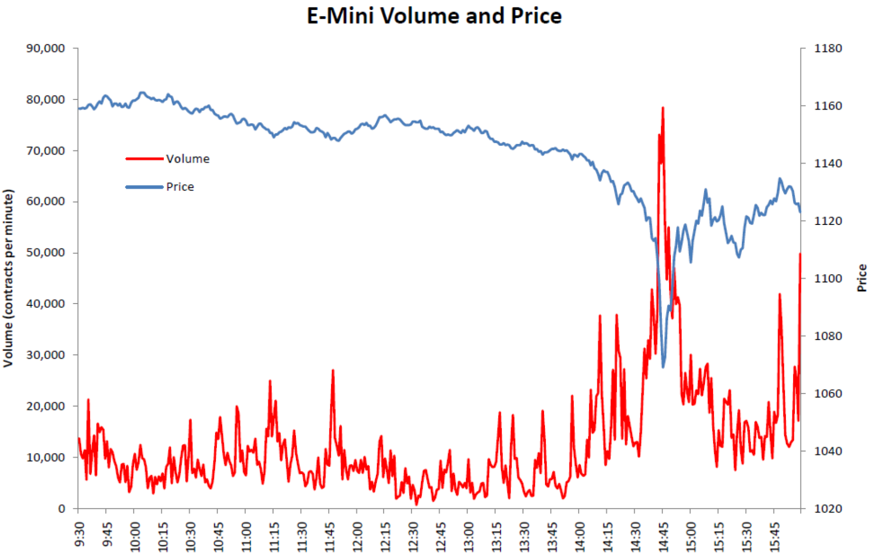
E-Mini Volume and Price
Volume
Price
Volume (contracts per minute)
90,000
80,000
70,000
60,000
50,000
40,000
30,000
20,000
10,000
0
1180
1160
1140
1120
1100
1080
1060
1040
1020
9:30
9:45
10:00
10:15
10:30
10:45
11:00
11:15
11:30
11:45
12:00
12:15
12:30
12:45
13:00
13:15
13:30
13:45
14:00
14:15
14:30
14:45
15:00
15:15
15:30
15:45

Source: SEC.

Fast and Furious: Preface

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.
- The error was triggered by an execution algorithm X of a (traditional) fund. X activated a huge sell order: 75,000 futures, aprox. 4.1bn USD.

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.
- The error was triggered by an execution algorithm X of a (traditional) fund. X activated a huge sell order: 75,000 futures, aprox. 4.1bn USD.
- The asset was the E-Mini, a S&P500 future issued by the Chicago Mercantile Exchange (CME).

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.
- The error was triggered by an execution algorithm X of a (traditional) fund. X activated a huge sell order: 75,000 futures, aprox. 4.1bn USD.
- The asset was the E-Mini, a S&P500 future issued by the Chicago Mercantile Exchange (CME).
- X was of PoV type, $i.e.$ percentage-of-volume: it sends an order of 9% of the total traded volume of the last minute. $\Longrightarrow$ No control on price, time or market impact.

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.
- The error was triggered by an execution algorithm X of a (traditional) fund. X activated a huge sell order: 75,000 futures, aprox. 4.1bn USD.
- The asset was the E-Mini, a S&P500 future issued by the Chicago Mercantile Exchange (CME).
- X was of PoV type, $i.e.$ percentage-of-volume: it sends an order of 9% of the total traded volume of the last minute. $\Longrightarrow$ No control on price, time or market impact.
- Normally, an intraday execution of this size takes at least 5 hours to be filled.

Fast and Furious: Preface

- The Flash Crash lasted only 36 minutes, but it erased 862bn USD of market capitalisation.
- The Dow Jones lost almost 1,000 points (9%) but minutes later it recovered.
- The error was triggered by an execution algorithm X of a (traditional) fund. X activated a huge sell order: 75,000 futures, aprox. 4.1bn USD.
- The asset was the E-Mini, a S&P500 future issued by the Chicago Mercantile Exchange (CME).
- X was of PoV type, *i.e.* percentage-of-volume: it sends an order of 9% of the total traded volume of the last minute. $\Longrightarrow$ No control on price, time or market impact.
- Normally, an intraday execution of this size takes at least 5 hours to be filled.
- But X executed it in only 19 minutes $\Longrightarrow$ 15 times faster.

Fast and Furious: Exposition

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.
- But the more liquidity the market offered to X, the more X accelerated its execution: PoV algo without upper cap.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.
- But the more liquidity the market offered to X, the more X accelerated its execution: PoV algo without upper cap.
- The market liquidity at the bid was not enough to counter the selling pressure of X.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.
- But the more liquidity the market offered to X, the more X accelerated its execution: PoV algo without upper cap.
- The market liquidity at the bid was not enough to counter the selling pressure of X.
- Between 2:41pm and 2:44pm, all HFTs capitulated: they sold aggressively 2,000 E-Mini contracts to reduce the inventories they built up by buying to X.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.
- But the more liquidity the market offered to X, the more X accelerated its execution: PoV algo without upper cap.
- The market liquidity at the bid was not enough to counter the selling pressure of X.
- Between 2:41pm and 2:44pm, all HFTs capitulated: they sold aggressively 2,000 E-Mini contracts to reduce the inventories they built up by buying to X.
  - HFTs not only stopped providing liquidity, they started to consume it.

Fast and Furious: Exposition

- At 2:32pm the first sell child order of X hit the market.
- X consumed the market liquidity at a fast rate $\Longrightarrow$ HFTs and (cross-market) arbitrageurs spotted it and started to offer more liquidity.
- But the more liquidity the market offered to X, the more X accelerated its execution: PoV algo without upper cap.
- The market liquidity at the bid was not enough to counter the selling pressure of X.
- Between 2:41pm and 2:44pm, all HFTs capitulated: they sold aggressively 2,000 E-Mini contracts to reduce the inventories they built up by buying to X.
  - HFTs not only stopped providing liquidity, they started to consume it.
  - Chain reaction: the market liquidity at the bid disappeared, hence prices went into a free fall.

Fast and Furious: Complication

Fast and Furious: Complication

- Between 2:32pm and 2:45pm (13 minutes), with prices falling down, X sold 35,000 E-Mini contracts (1.9bn USD) out of its total of 75,000.

Fast and Furious: Complication

- Between 2:32pm and 2:45pm (13 minutes), with prices falling down, X sold 35,000 E-Mini contracts (1.9bn USD) out of its total of 75,000.
- During this period there were 80,000 contracts sold vs 50,000 contracts bought, *i.e.* a net balance of -30; 000 traded contracts.

Fast and Furious: Complication

- Between 2:32pm and 2:45pm (13 minutes), with prices falling down, X sold 35,000 E-Mini contracts (1.9bn USD) out of its total of 75,000.
- During this period there were 80,000 contracts sold vs 50,000 contracts bought, *i.e.* a net balance of -30; 000 traded contracts.
- Comparing with the former three days of trading: 15 times more sell volume and 10 times more buy volume.

Fast and Furious: Complication

- Between 2:32pm and 2:45pm (13 minutes), with prices falling down, X sold 35,000 E-Mini contracts (1.9bn USD) out of its total of 75,000.
- During this period there were 80,000 contracts sold vs 50,000 contracts bought, *i.e.* a net balance of -30; 000 traded contracts.
- Comparing with the former three days of trading: 15 times more sell volume and 10 times more buy volume.
- At 2:45:28pm trading in the E-Mini was suspended for 5 seconds: the CME Stop Logic Functionality triggered automatically to stop prices from further loss.

Fast and Furious: Resolution

Fast and Furious: Resolution

- After the pause, an important flow of buy orders appeared: opportunistic and fundamental buyers.

Fast and Furious: Resolution

- After the pause, an important flow of buy orders appeared: opportunistic and fundamental buyers.
- These buy orders not only neutralised the fall in prices, they even countered them: prices started to climb.

Fast and Furious: Resolution

- After the pause, an important flow of buy orders appeared: opportunistic and fundamental buyers.
- These buy orders not only neutralised the fall in prices, they even countered them: prices started to climb.
- But at least half of the market participants had left.

Fast and Furious: Resolution

- After the pause, an important flow of buy orders appeared: opportunistic and fundamental buyers.
- These buy orders not only neutralised the fall in prices, they even countered them: prices started to climb.
- But at least half of the market participants had left.
- Between 2:45pm and 3:08pm (23 minutes) more than 110,000 contracts were sold and more than 110,000 contracts were bought, but the net balance was zero.

Fast and Furious: Resolution

- After the pause, an important flow of buy orders appeared: opportunistic and fundamental buyers.
- These buy orders not only neutralised the fall in prices, they even countered them: prices started to climb.
- But at least half of the market participants had left.
- Between 2:45pm and 3:08pm (23 minutes) more than 110,000 contracts were sold and more than 110,000 contracts were bought, but the net balance was zero.
- At 3:08pm, 36 minutes after the algo X entered the market, the intervention of opportunistic and fundamental buyers pushed prices back to their pre-crash levels.

References

K. Dayri and M. Rosenbaum, *Large tick assets: implicit spread and optimal tick size*. preprint, 2012.

T. W. Epps, Comovements in Stock Prices in the Very Short Run, *Journal of the American Statistical Association*, Vol. 74, No 366, pp291–298, 1974.

M.B. Garman, M.J. Klass, On the Estimation of Security Price Volatilities from Historical Data, *Journal of Business*, Vol. 53, No 1, pp67-78, 1980.

L. Glosten, P. Milgrom, Bid, ask and transaction prices in an specialist market with heterogeneously informed traders, *J..Fin.Econ.*, Vol. 14, pp71–100, 1985.

A. Kyle, Continuous auctions and insider trading, *Econometrica*, Vol.53 No.6 ,pp1315–1336, 1985.

S. L. et C.-A. Lehalle, *Market Microstructure in Practice*, World Scientific, 2018.

A. Madhavan, M. Richardson, and M. Roomans, *Why do security prices change? a transactionlevel analysis of nyse stocks*. Review of Financial Studies, 10(4):1035–64, 1997.

C. Y. Robert and M. Rosenbaum, *A new approach for the dynamics of ultra-high-frequency data: The model with uncertainty zones*. Journal of Financial Econometrics, 9(2):344–366, 2010.

M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between Bid-Ask spread, impact and volatility in double auction markets. Quantitative Finance, 8(1):41–57, 2008.