Evaluation of Volatility Estimators

Robert Almgren

quantitativebrokers

Jim Gatheral 60th birthday
Oct 15, 2017

Volatility

Theoretical

volatility surface

time structure (rough volatility)

Empirical

estimation from real data (Gatheral/Oomen 2010)

This work joint with Linwei Shang, Baruch student

Quantitative Brokers

Agency algorithmic execution
in futures and interest rate products

- 118 products (futures and cash notes+bonds)
- 7 exchanges (CME, Eurex, LIFFE, etc)
- 8 product types (IR, EQ, etc)

Need quantitative infrastructure:

- volume and volatility forecast curves
- intraday volatility estimator and forecasts

Arrival price benchmark

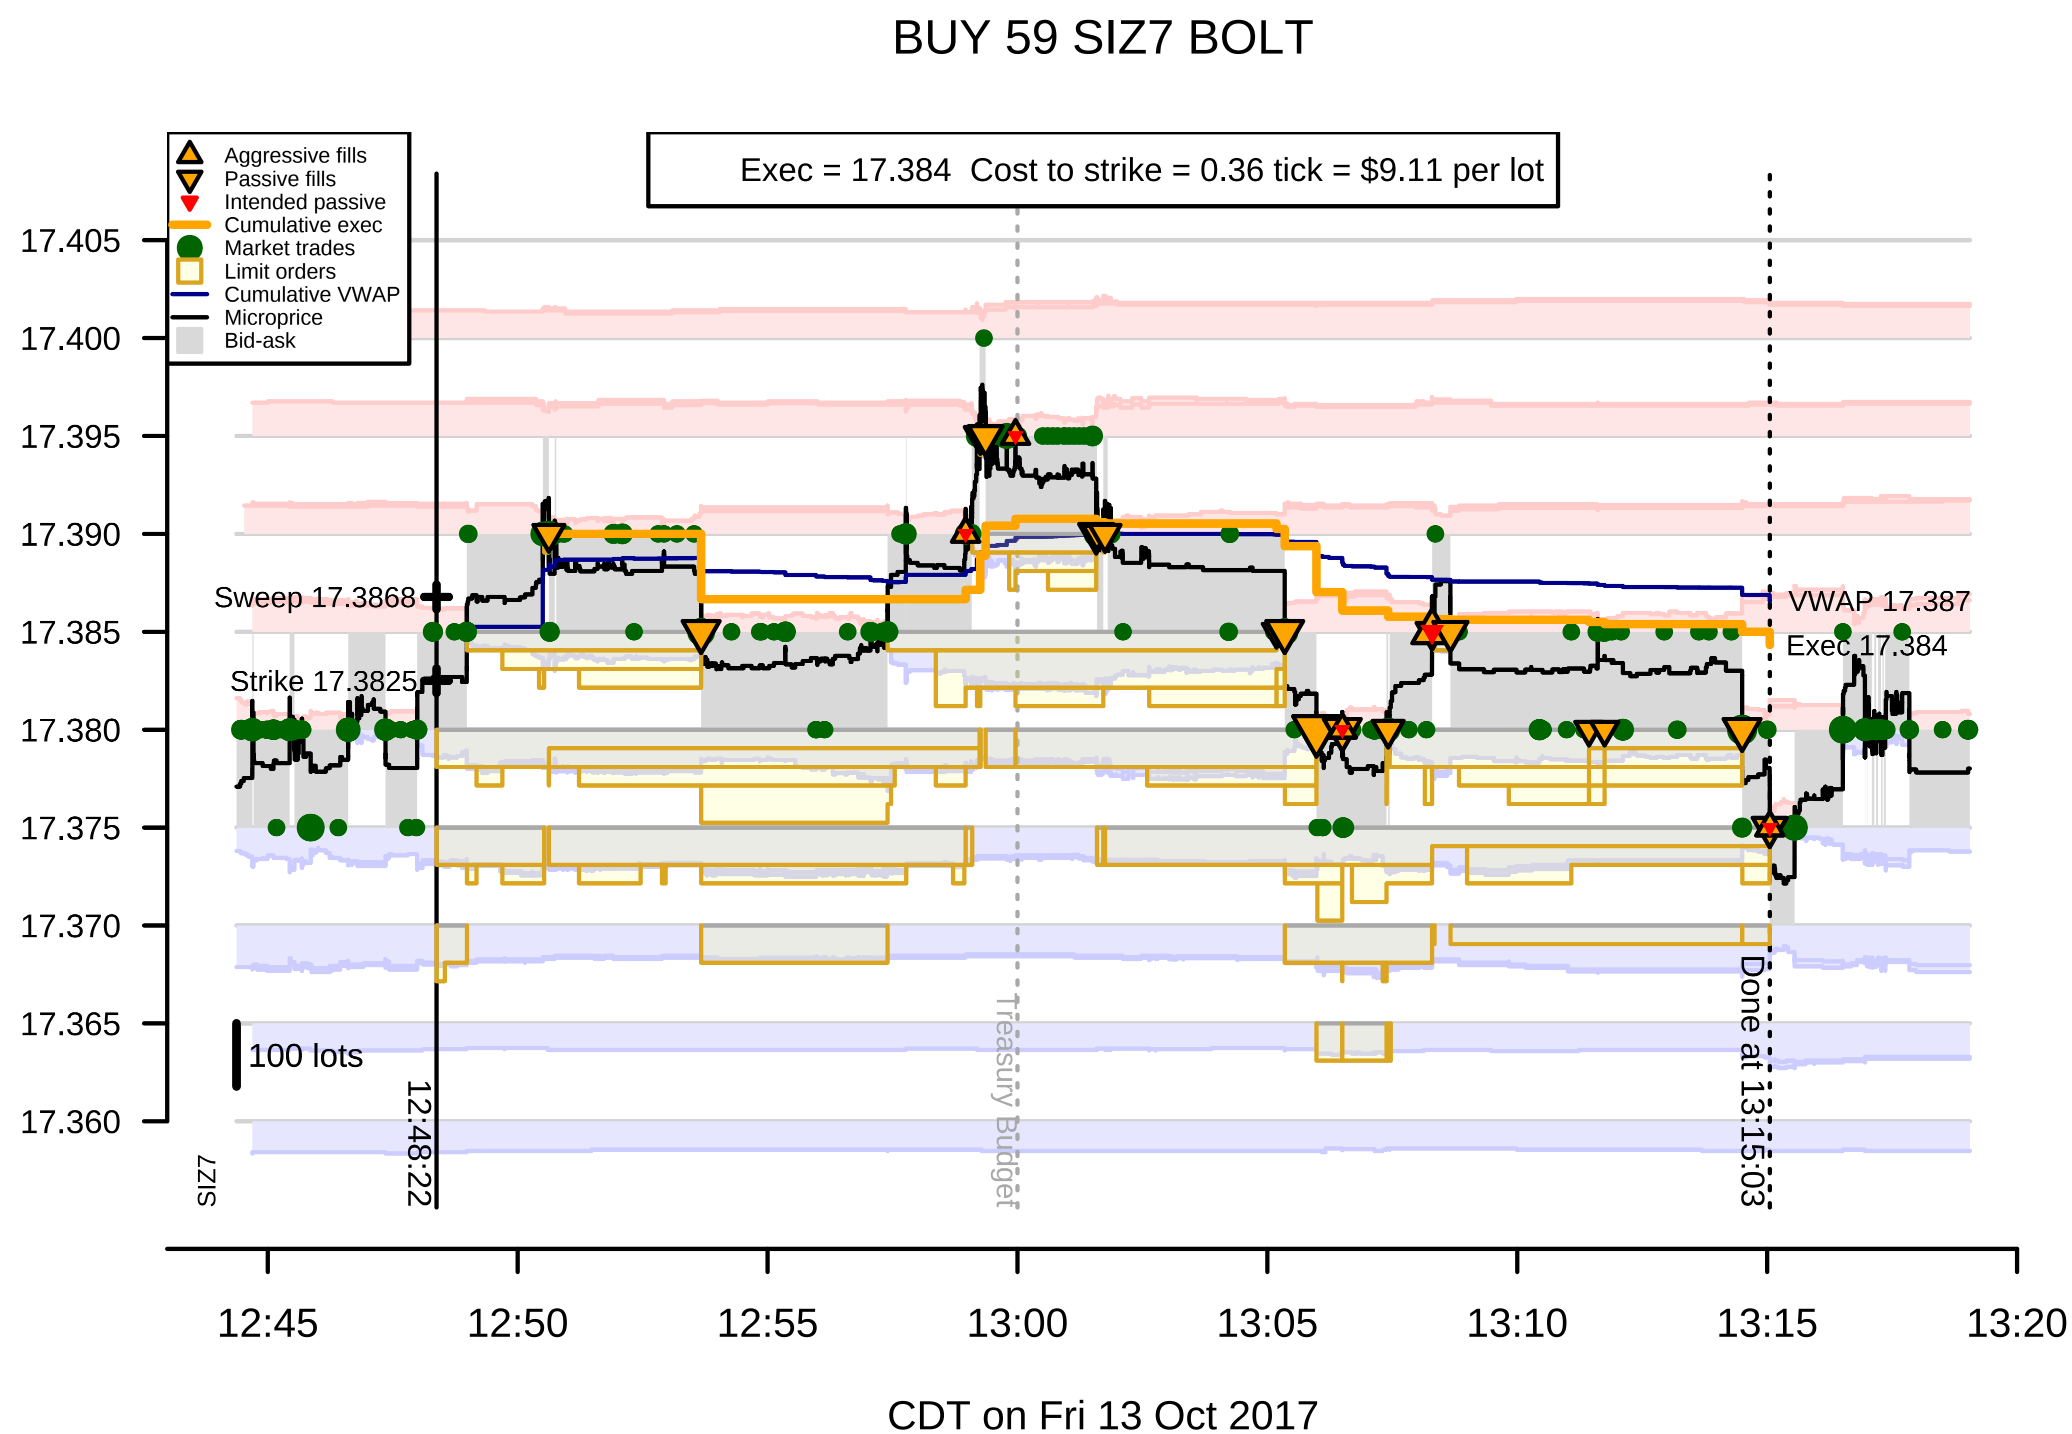

quantitativebrokers

VWAP benchmark

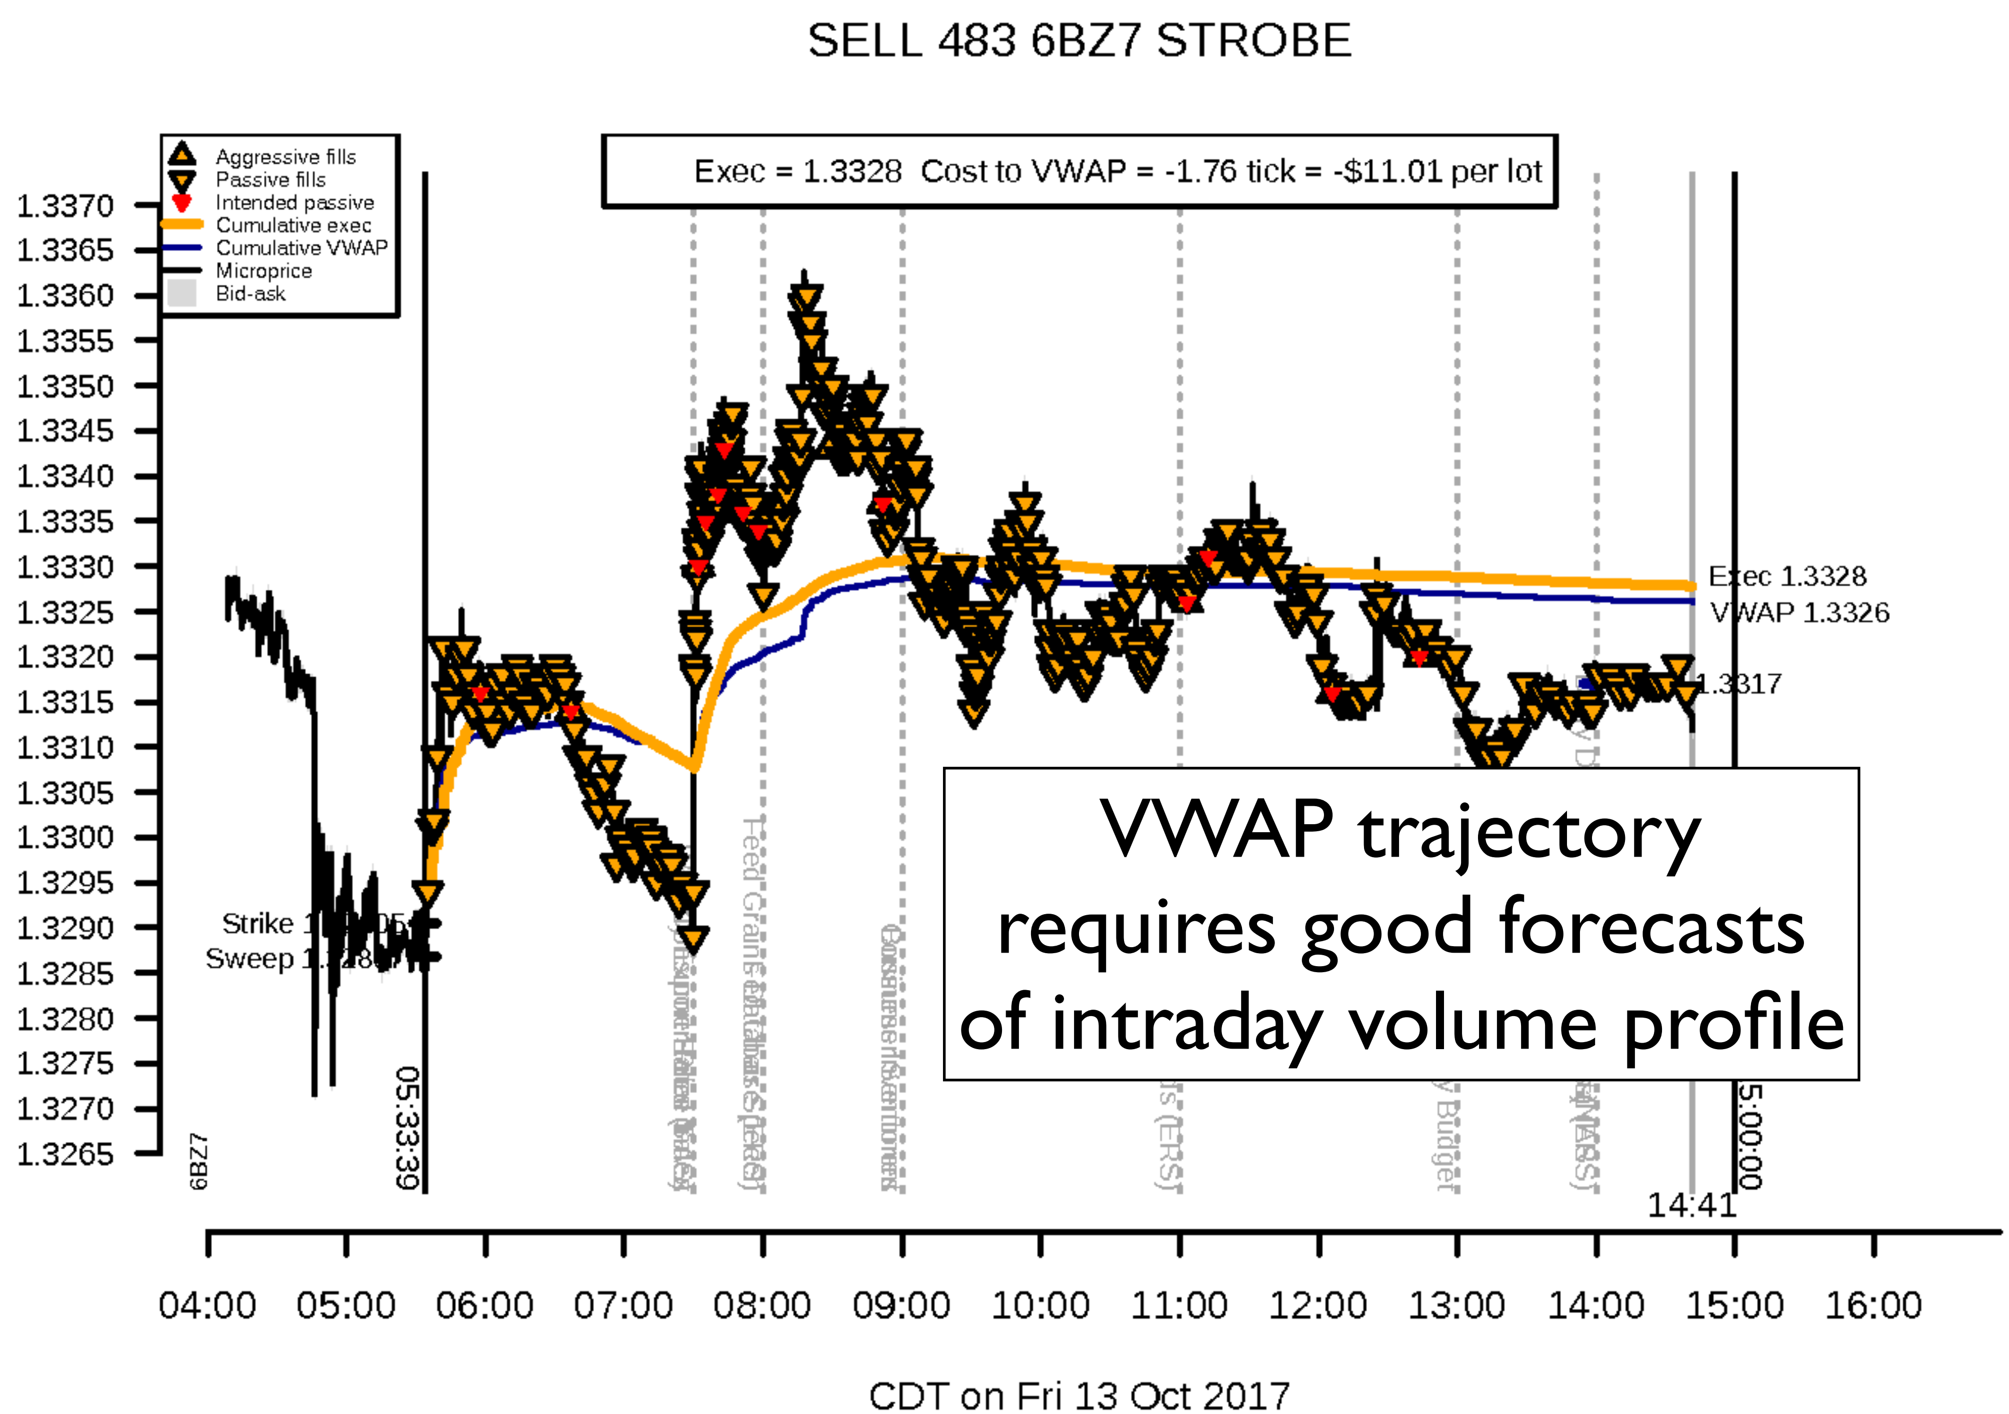

Volume and volatility forecast curves

Volume
Fri 2017-10-13

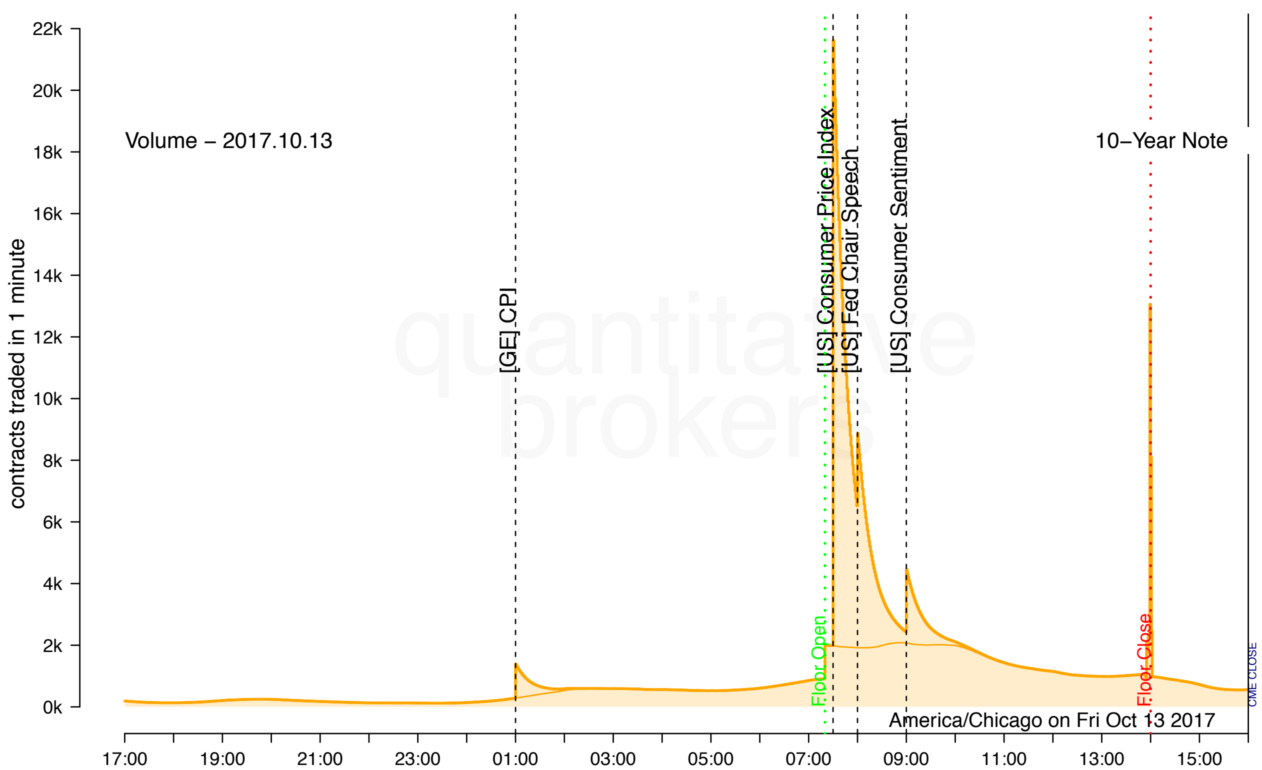

Volatility
Fri 2017-10-13

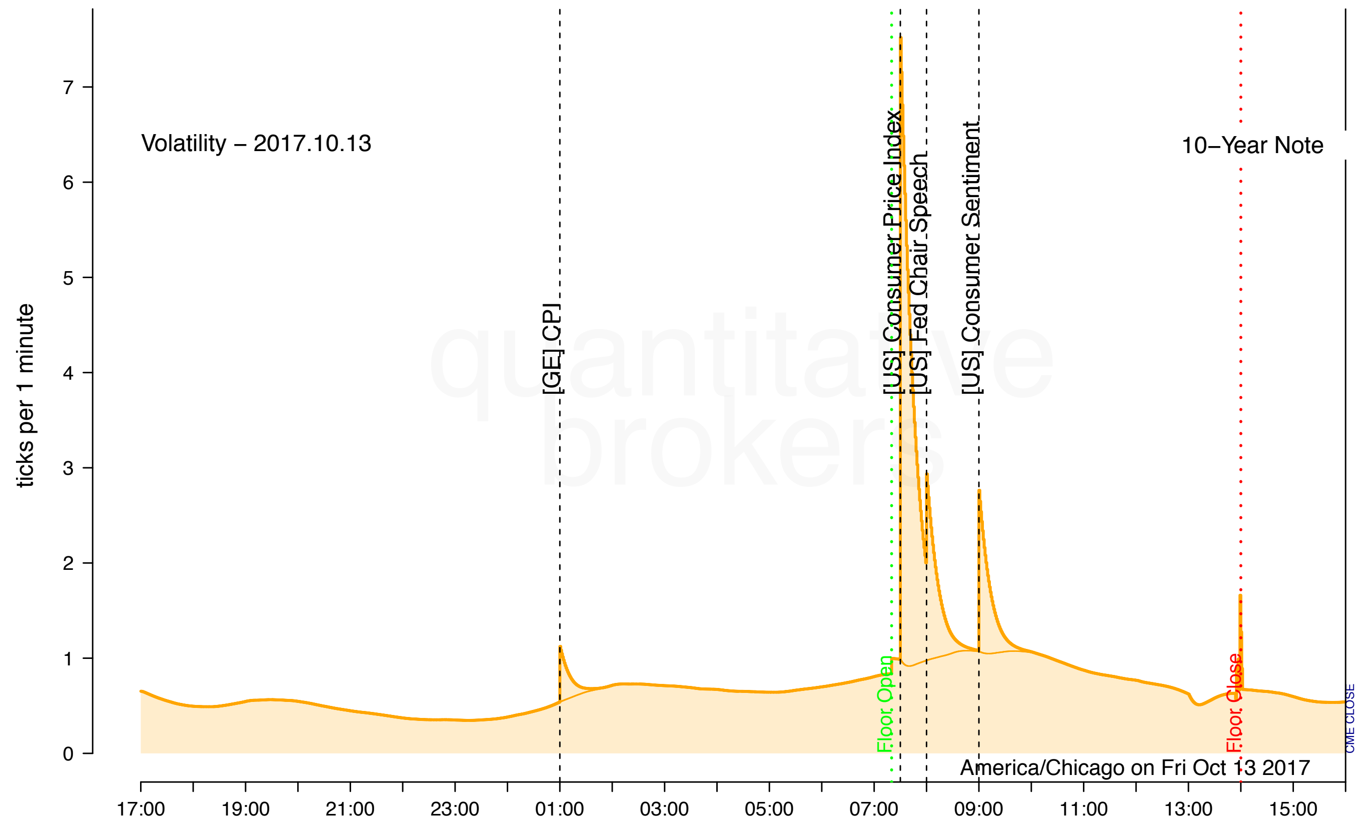

quantitativebrokers

Importance of volatility

Option pricing

realized vs implied

Relation to other market variables

microstructure invariance (Kyle/Obizhaeva):

volatility / volume / trade size

Estimated price change over time intervals

Forward price changes

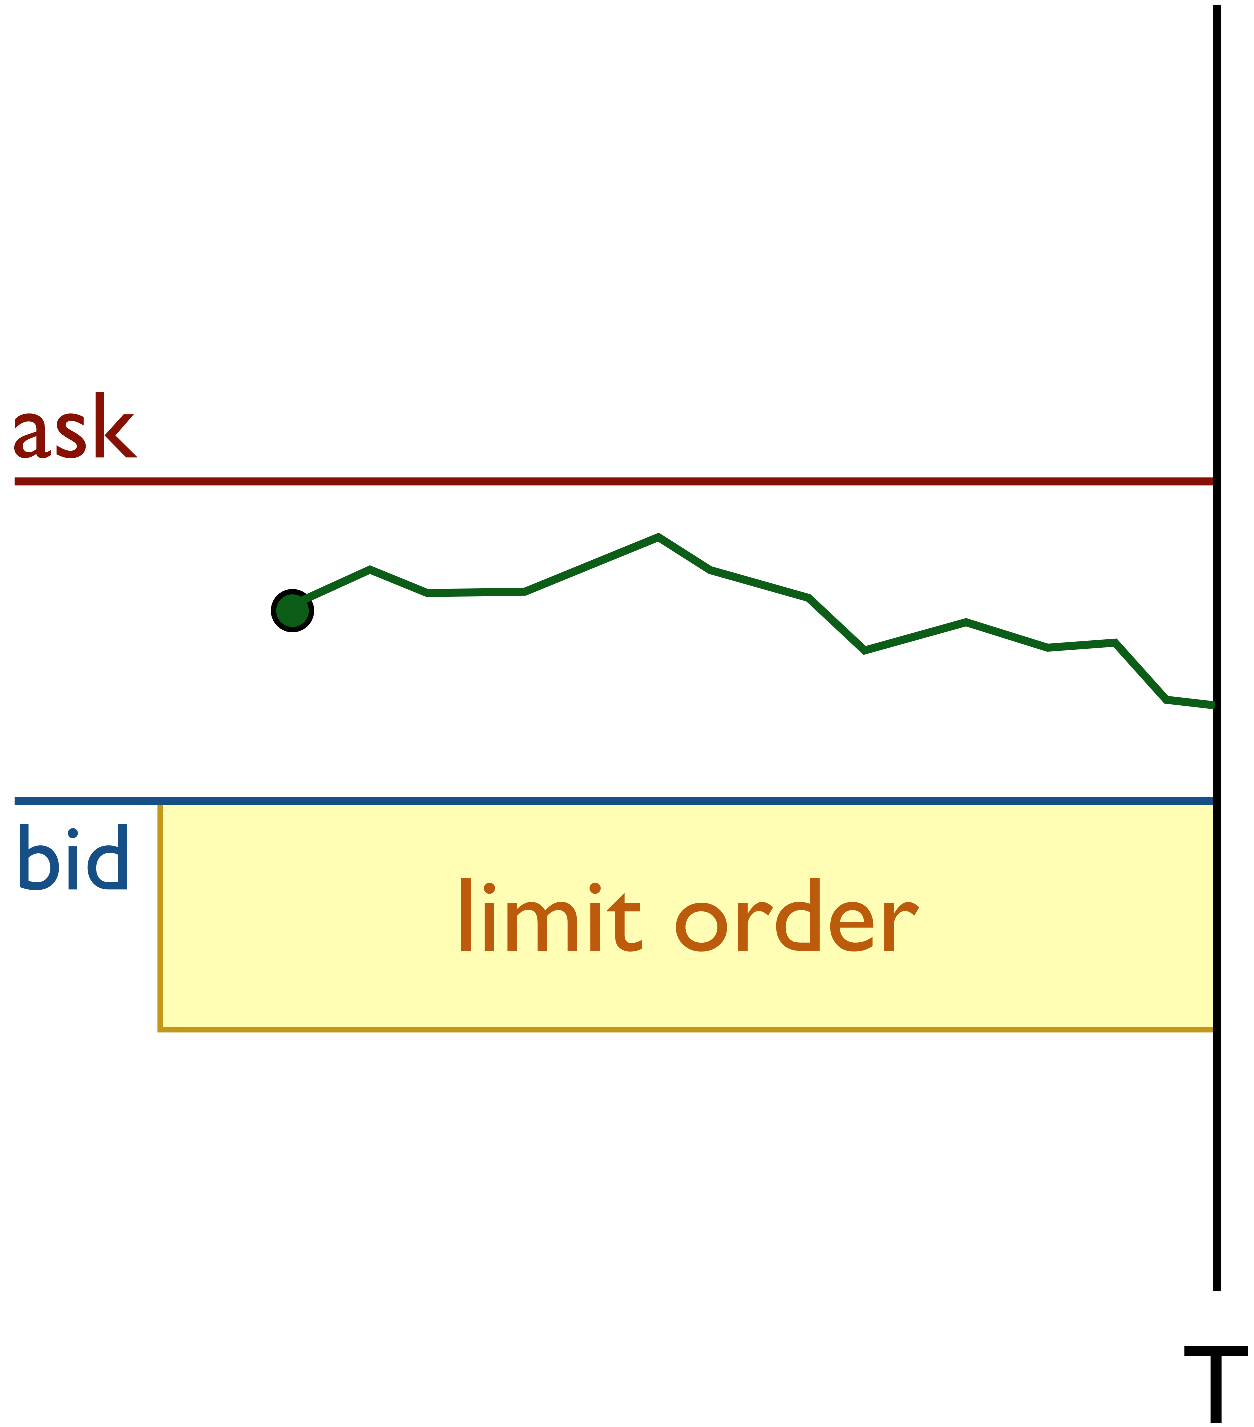

Probability of passive fill before end time T is based on $\sigma\sqrt{T-t}$

Volatility forecast
CME 10-year note futures, Fri Oct 13

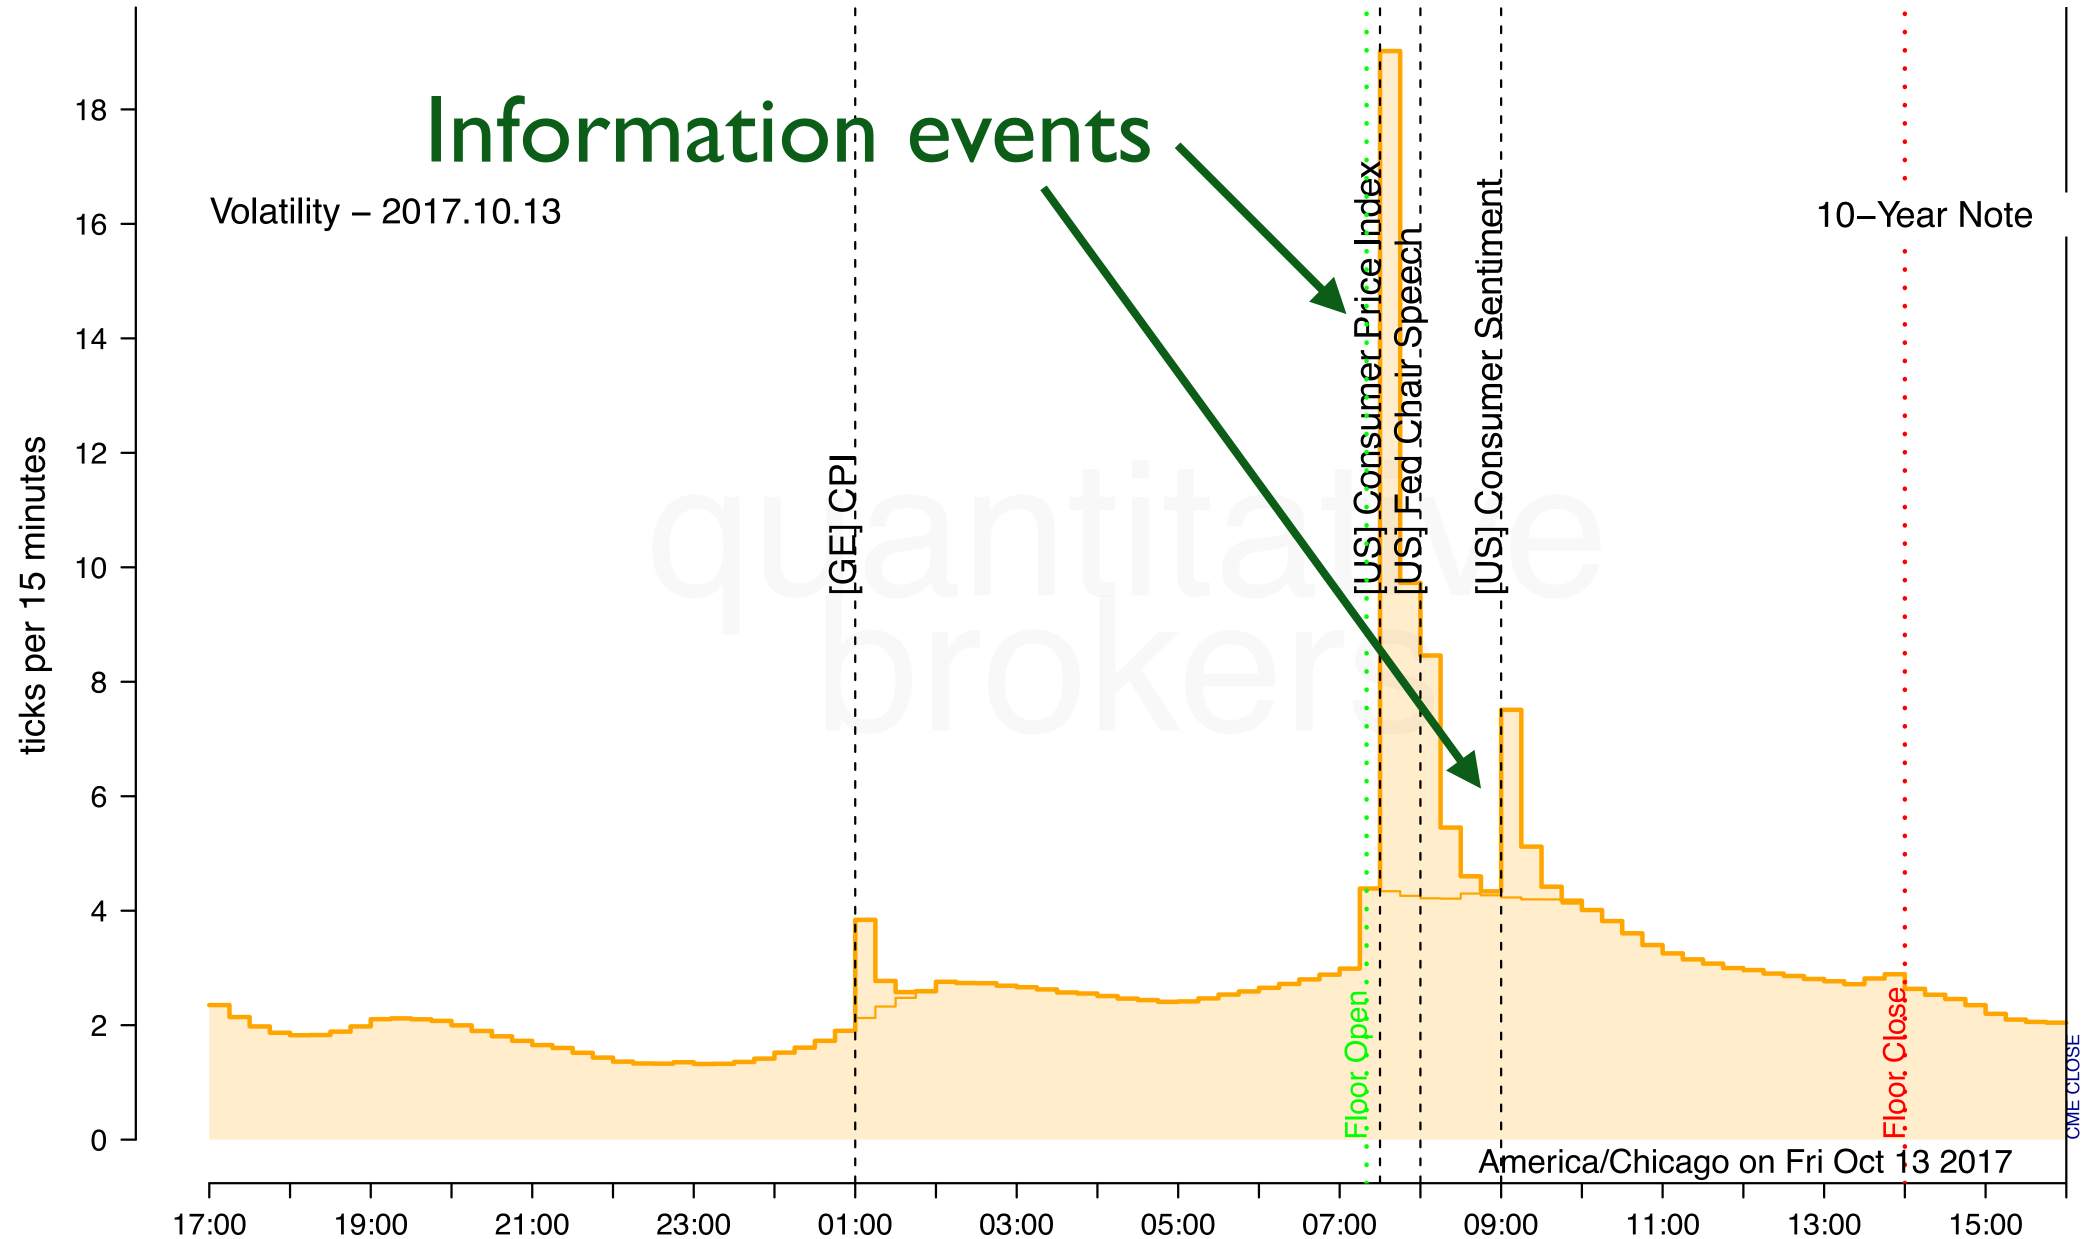

quantitativebrokers

Was this forecast correct?

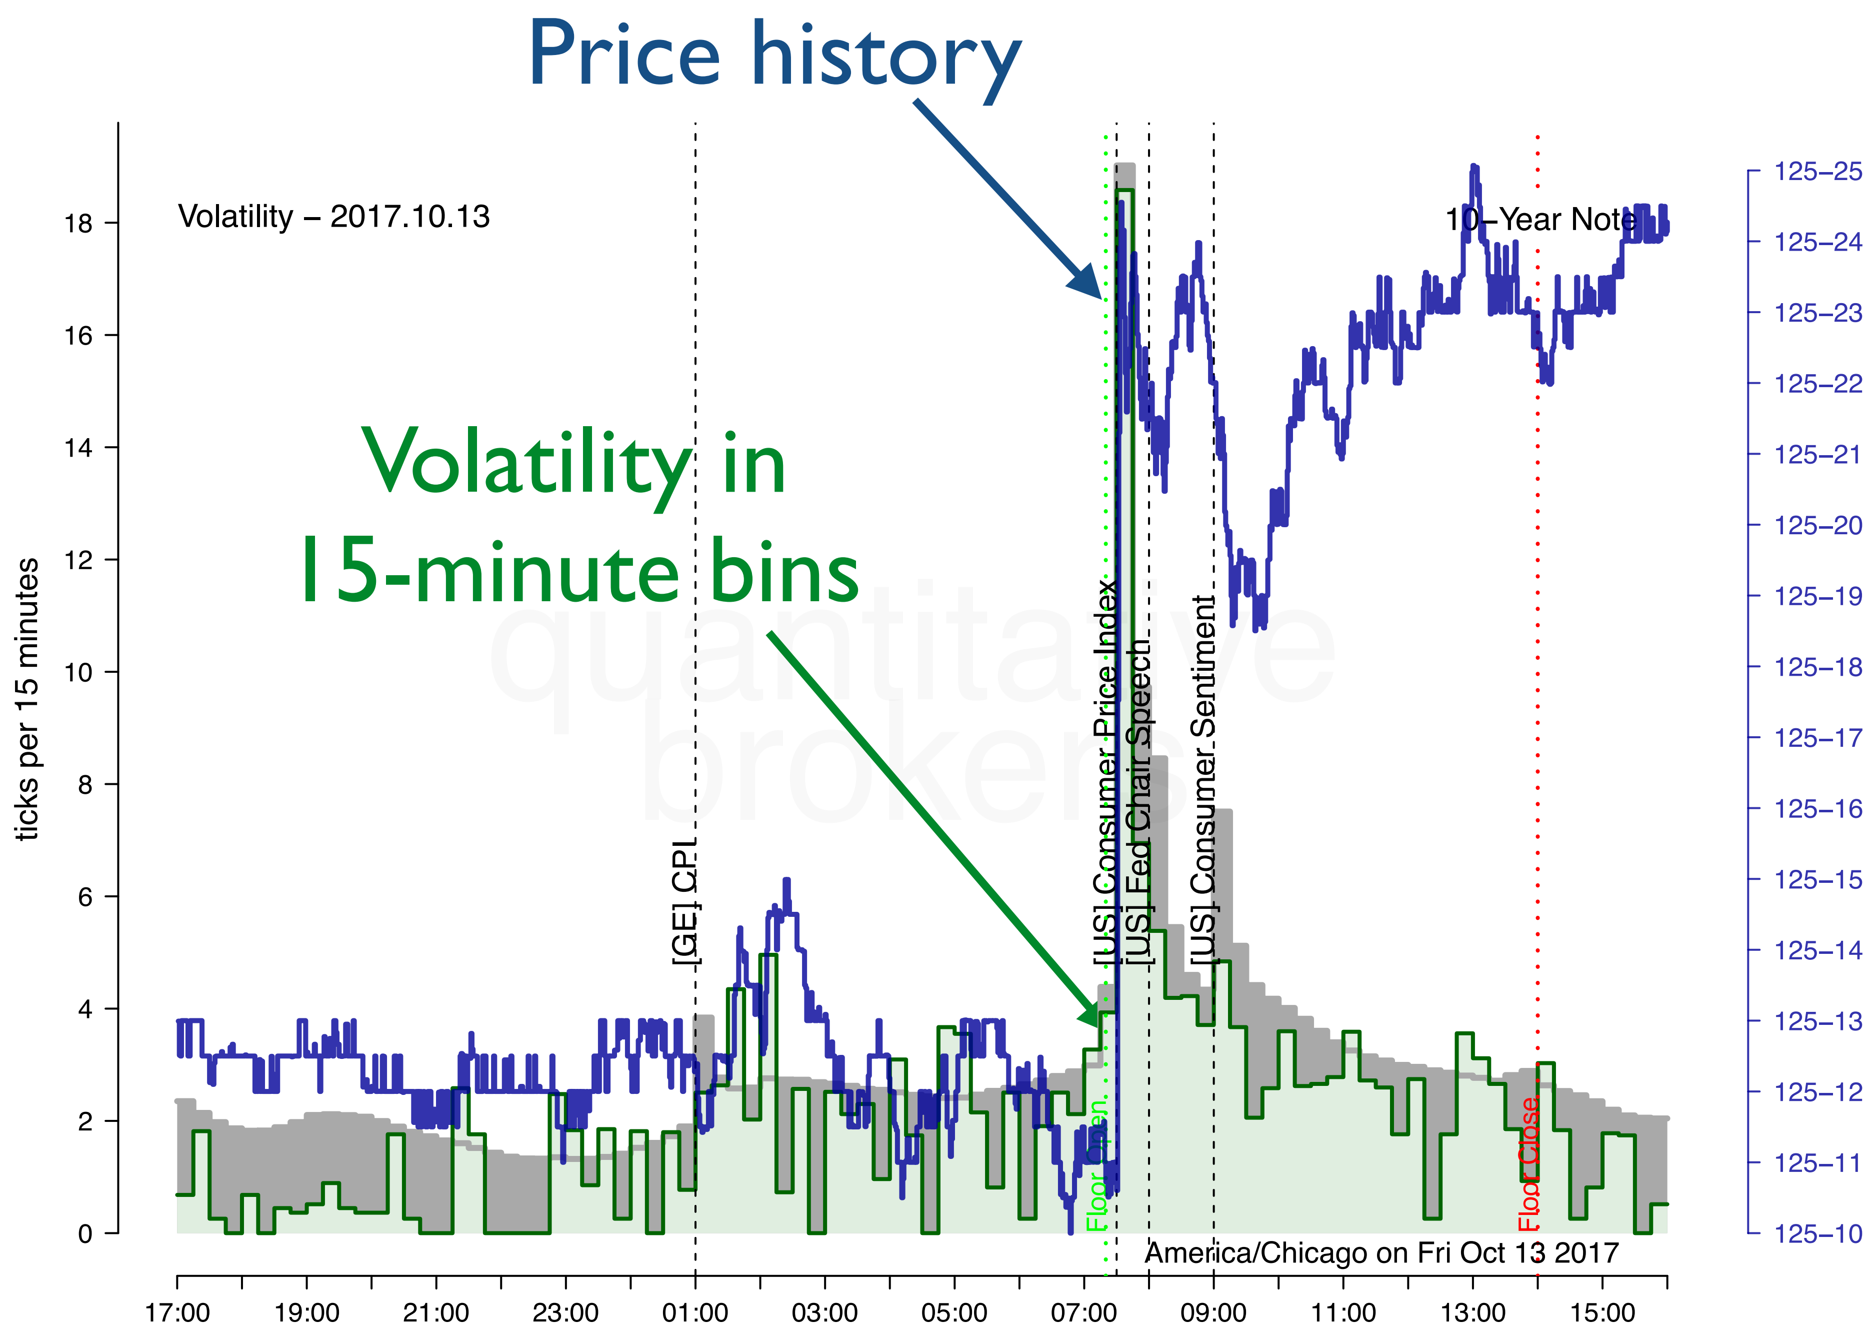

Is this bin volatility "correct" for this price history?

quantitativebrokers

Forecasts in 1-minute bins

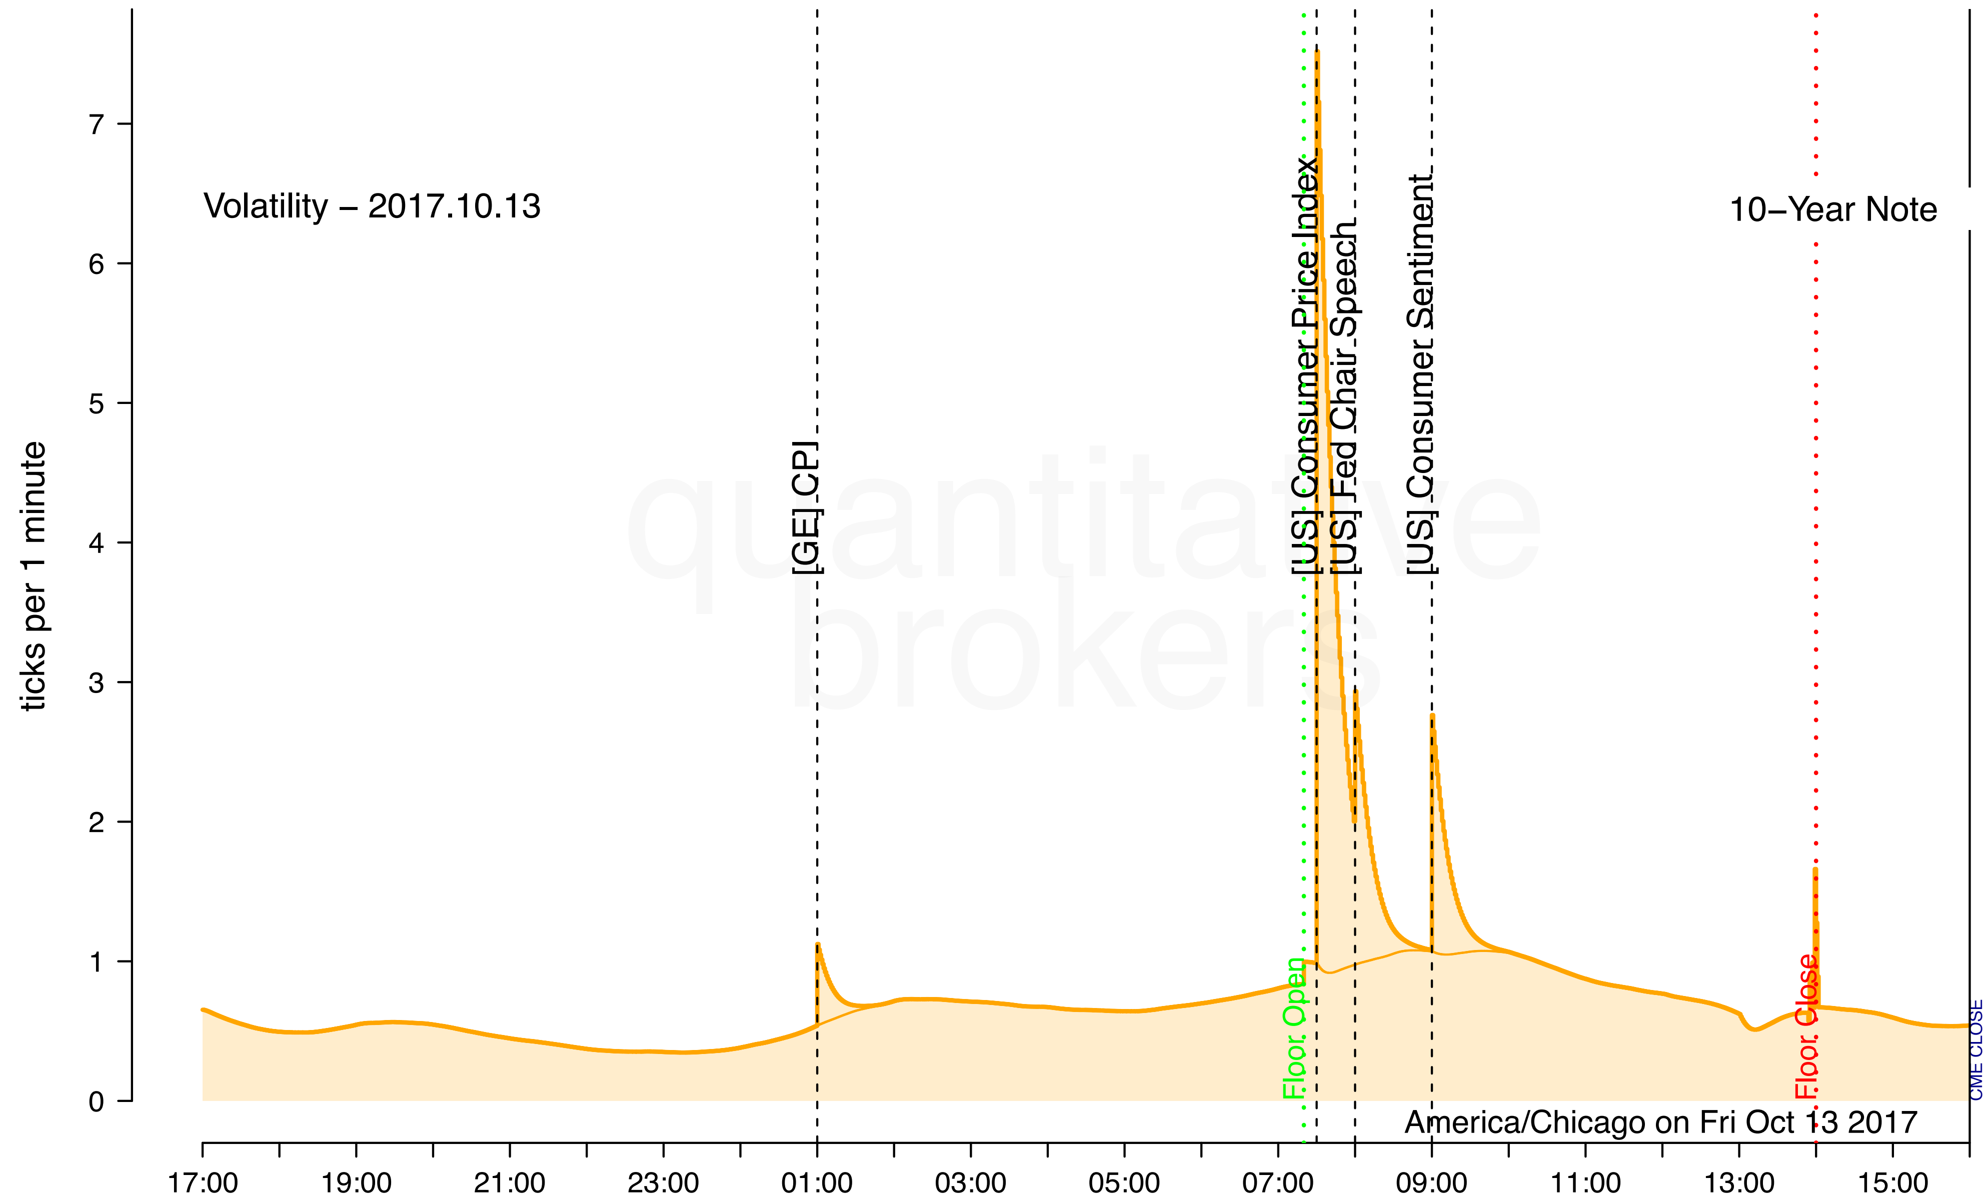

Requires calculation of volatility in 1-minute bins

quantitativebrokers

Challenges in futures products

- Range of products (agricultural, equity, etc)
- Large-tick vs small-tick
- Range of maturities on each product
- Options on futures
- 23-hour trading day
- Event responses

Large-tick vs small-tick

Robert-Rosenbaum reversion parameter

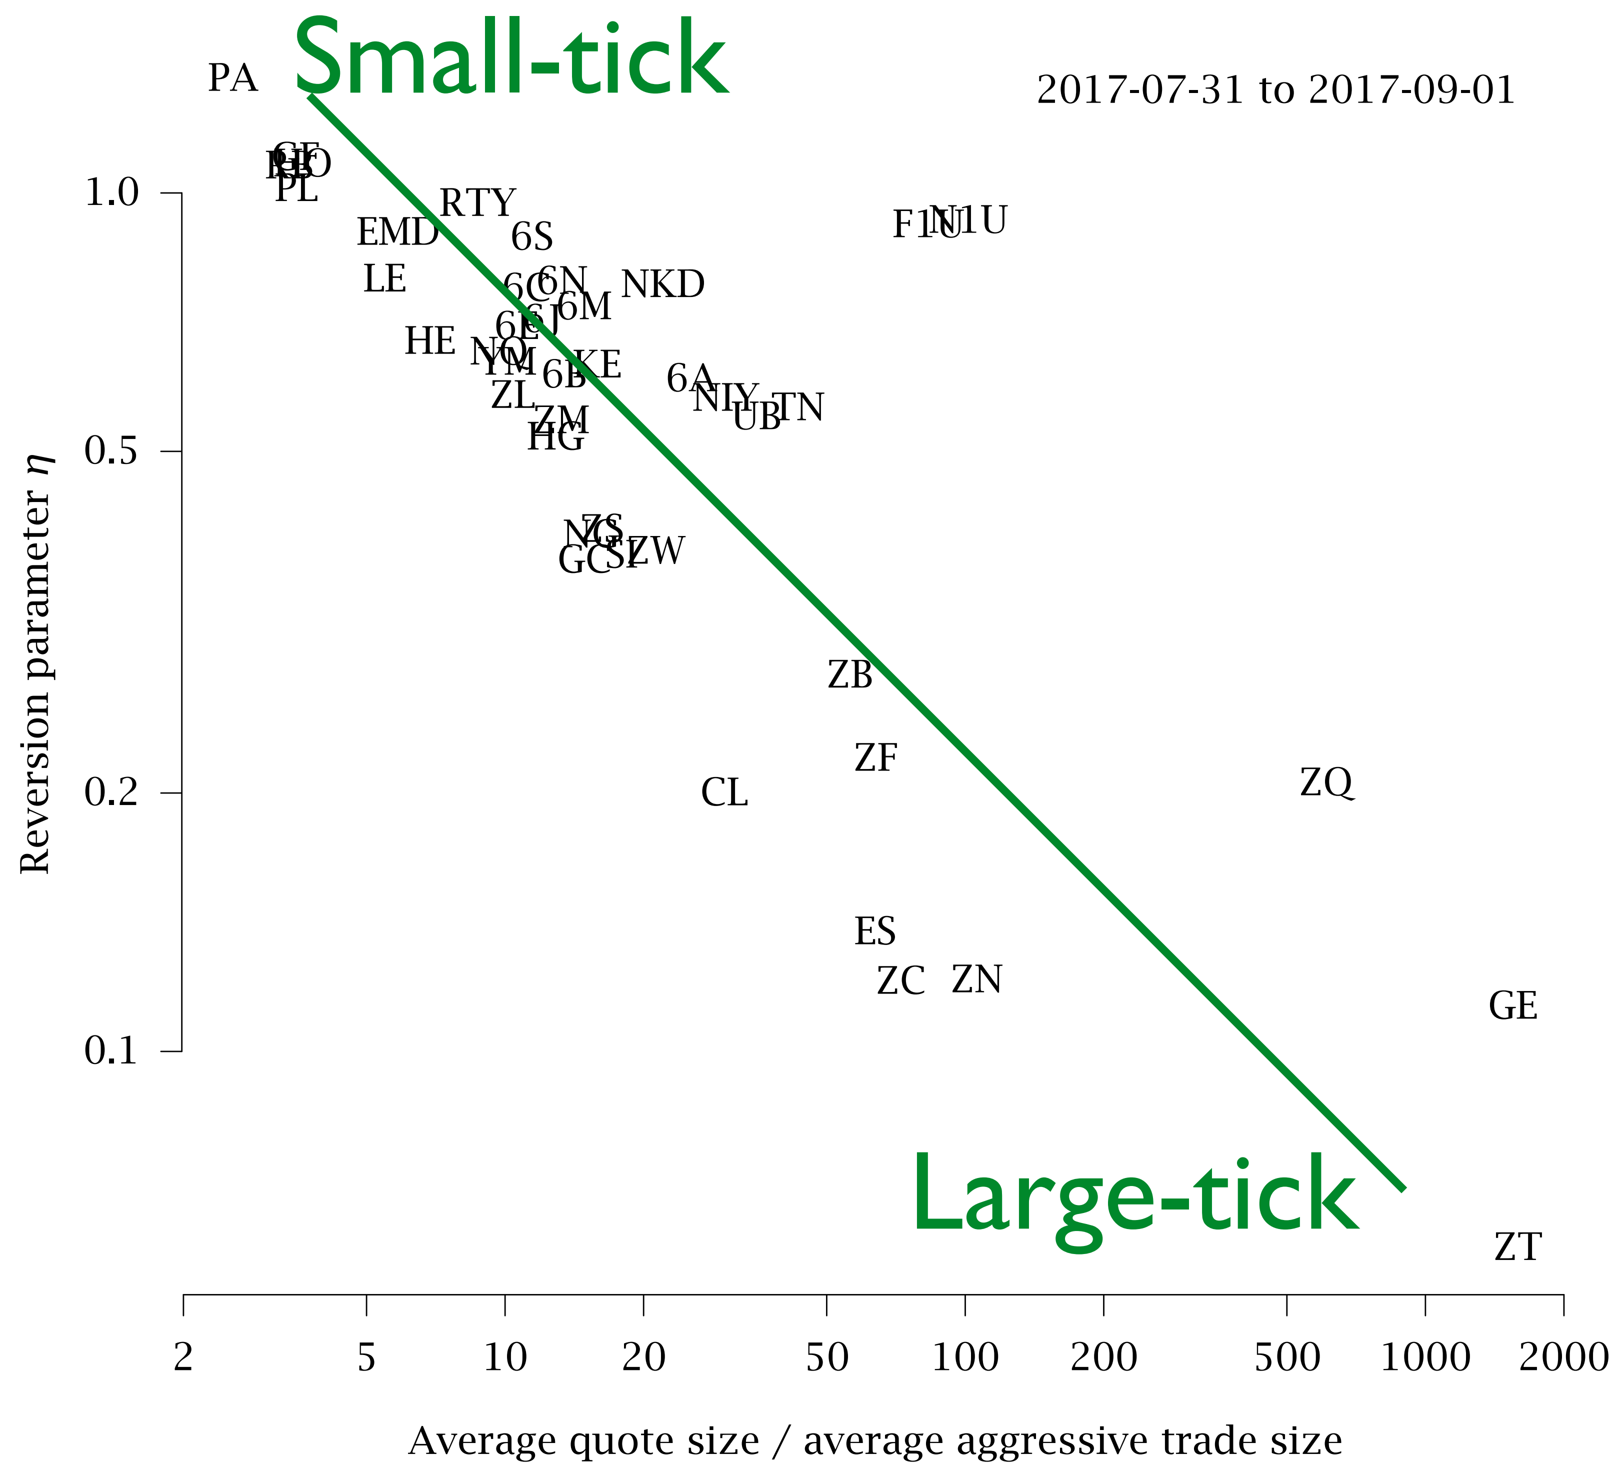

quantitativebrokers

Options on futures: much quote activity, few trades

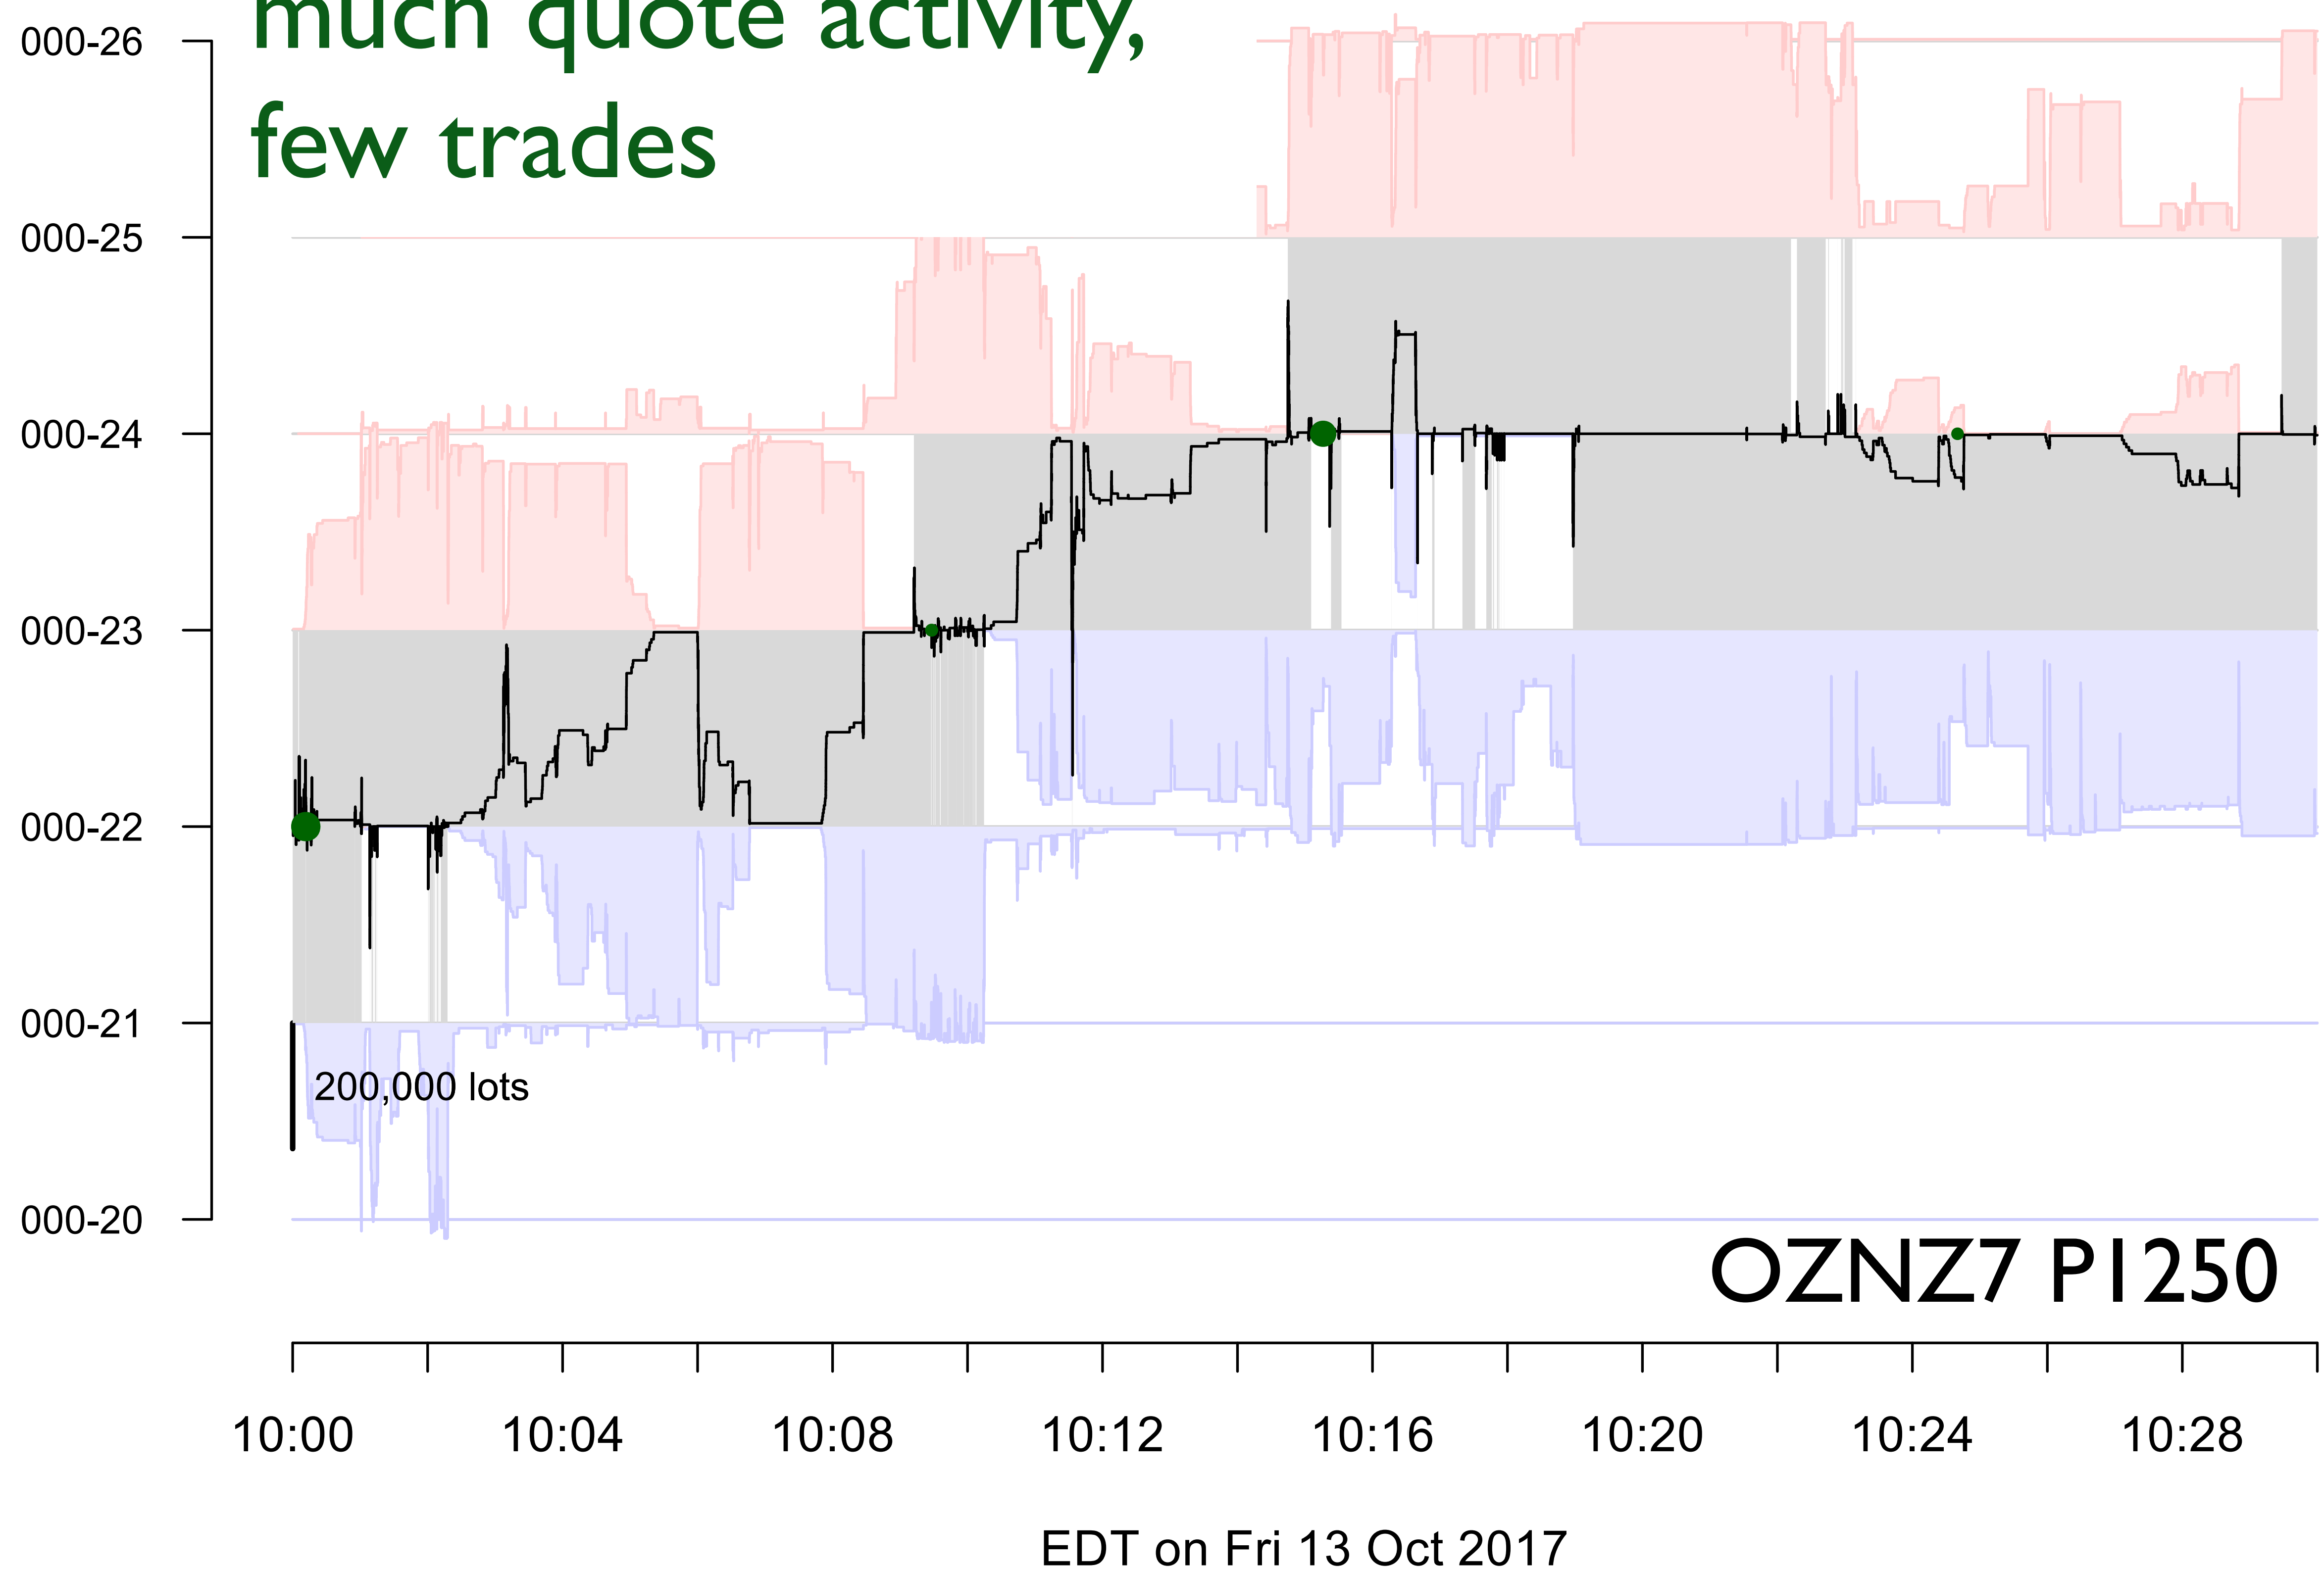

quantitativebrokers

Event reponse

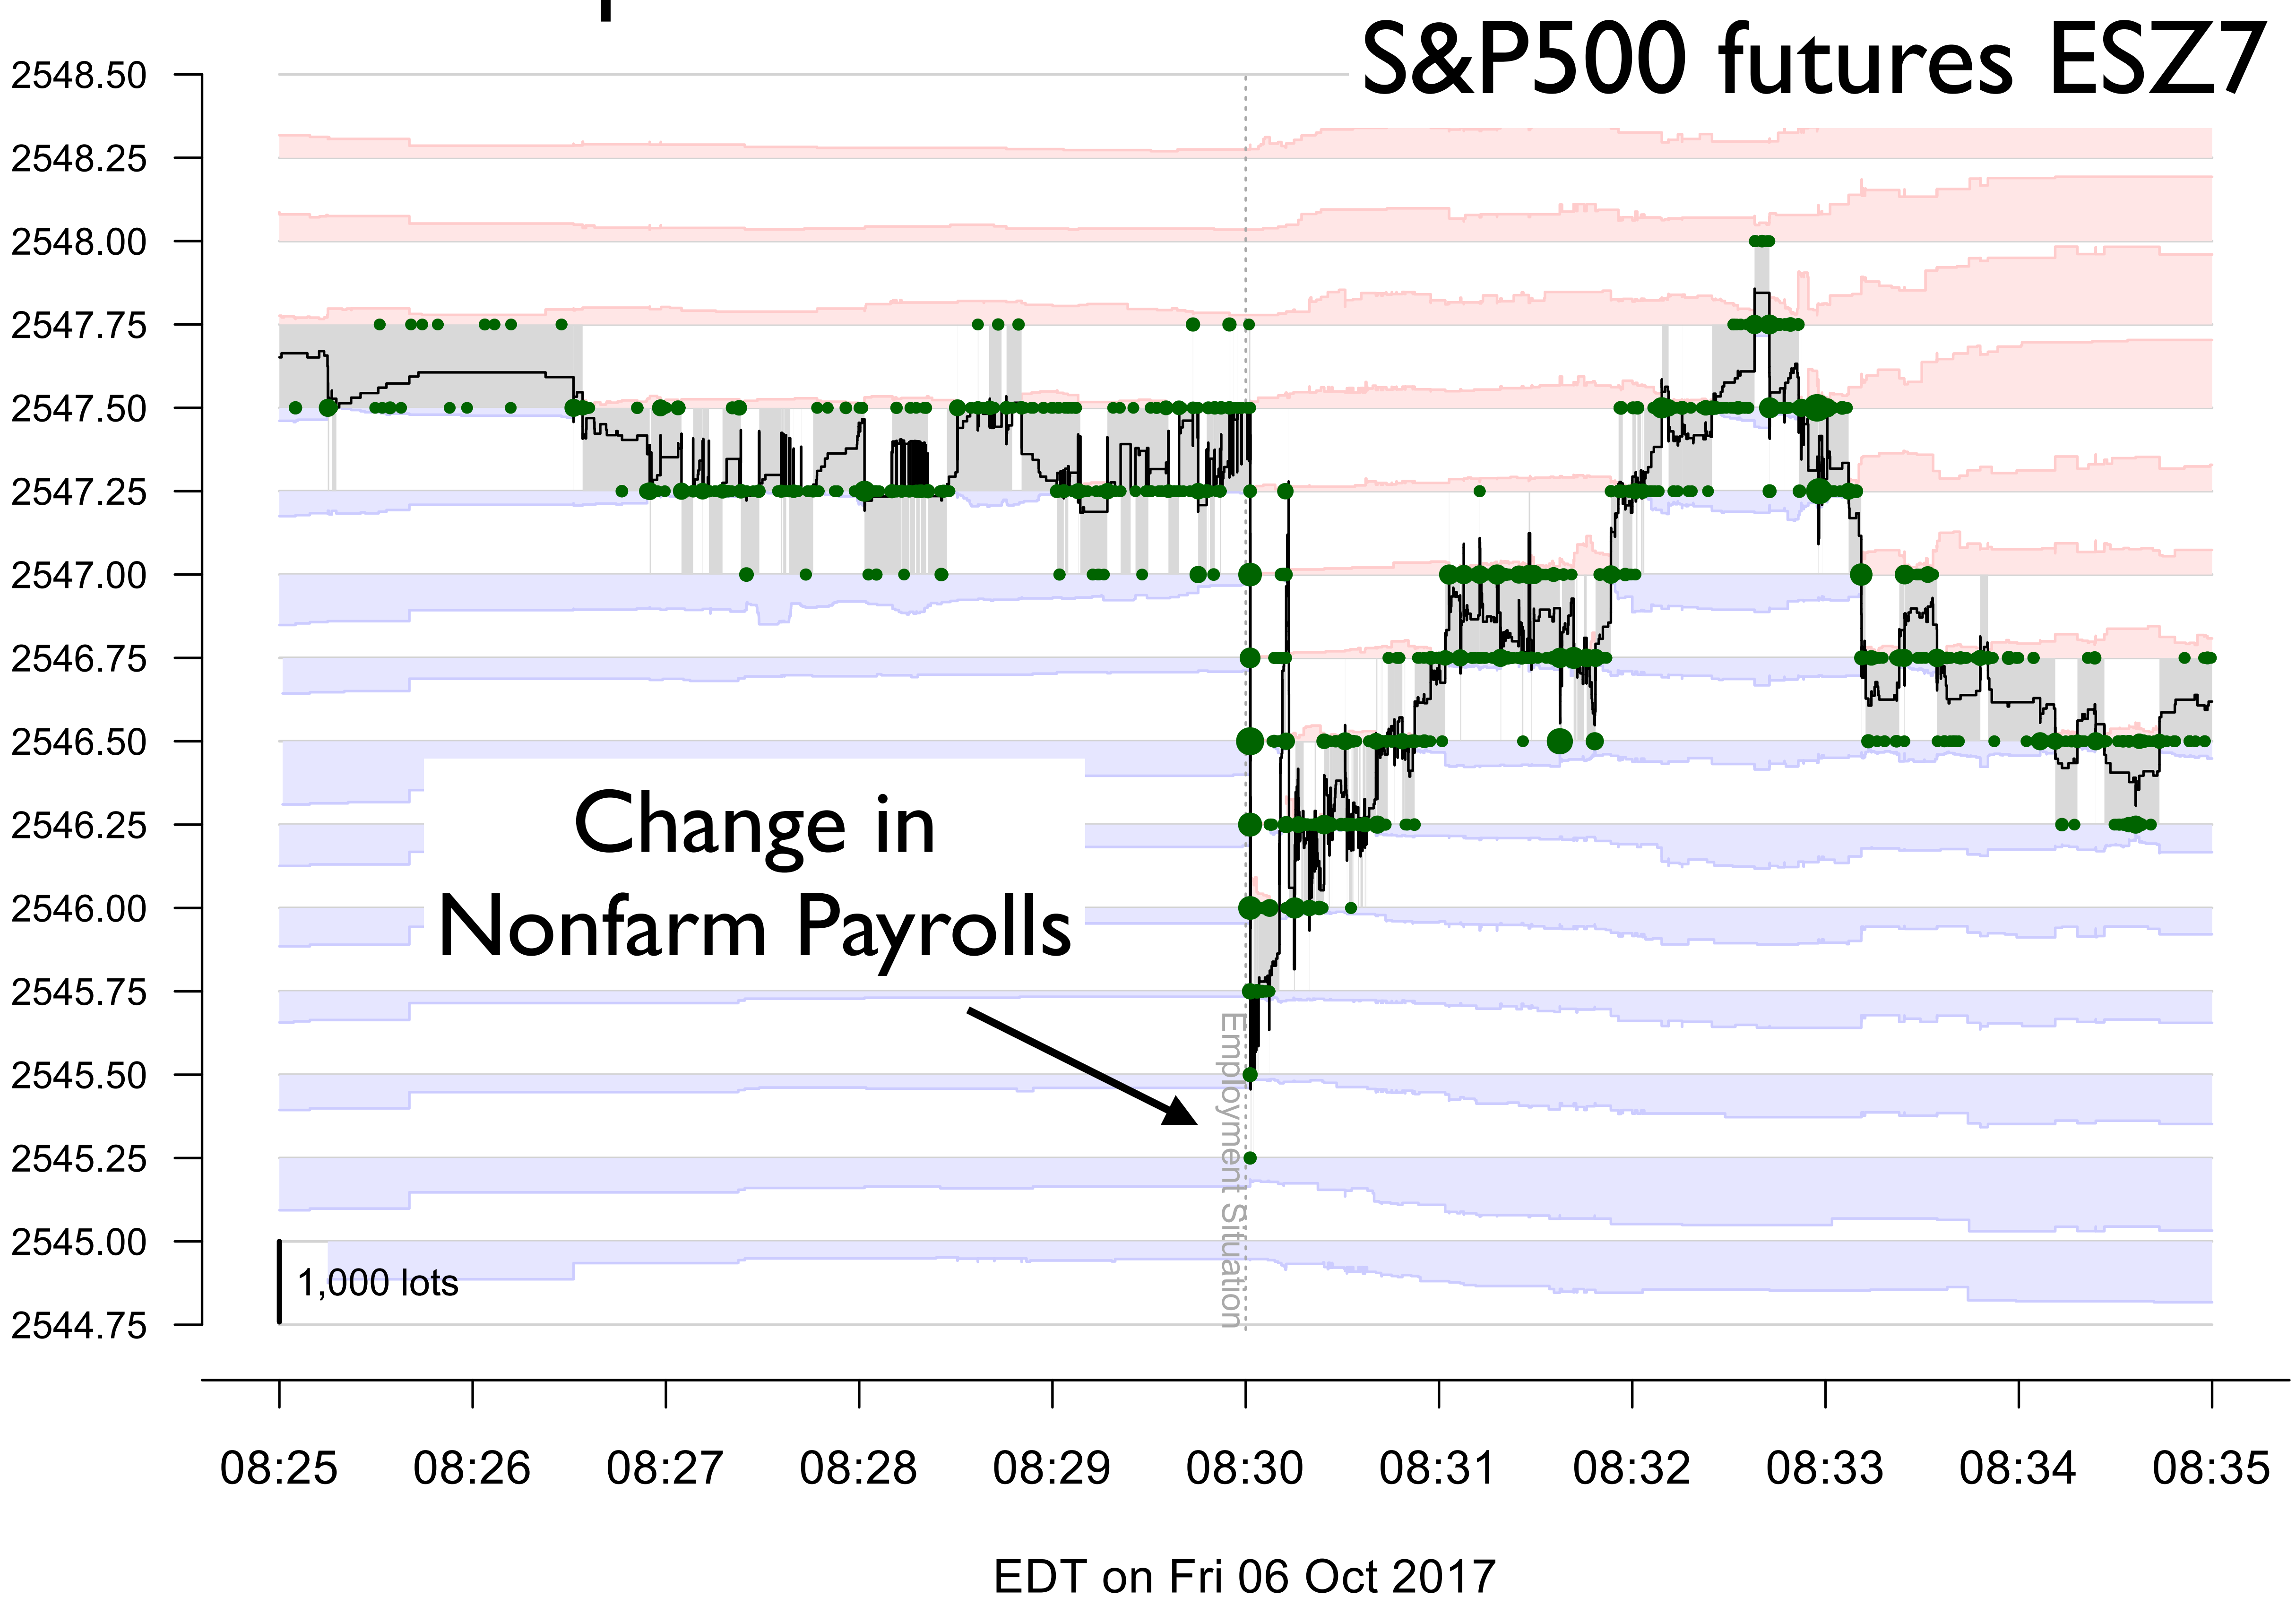

quantitativebrokers

Volatility measurement

Assign realized variance
number to any time interval
(1 minute, say)
using all available market data.
Aggregate for larger intervals.

How do we know whether
the number is "right"?

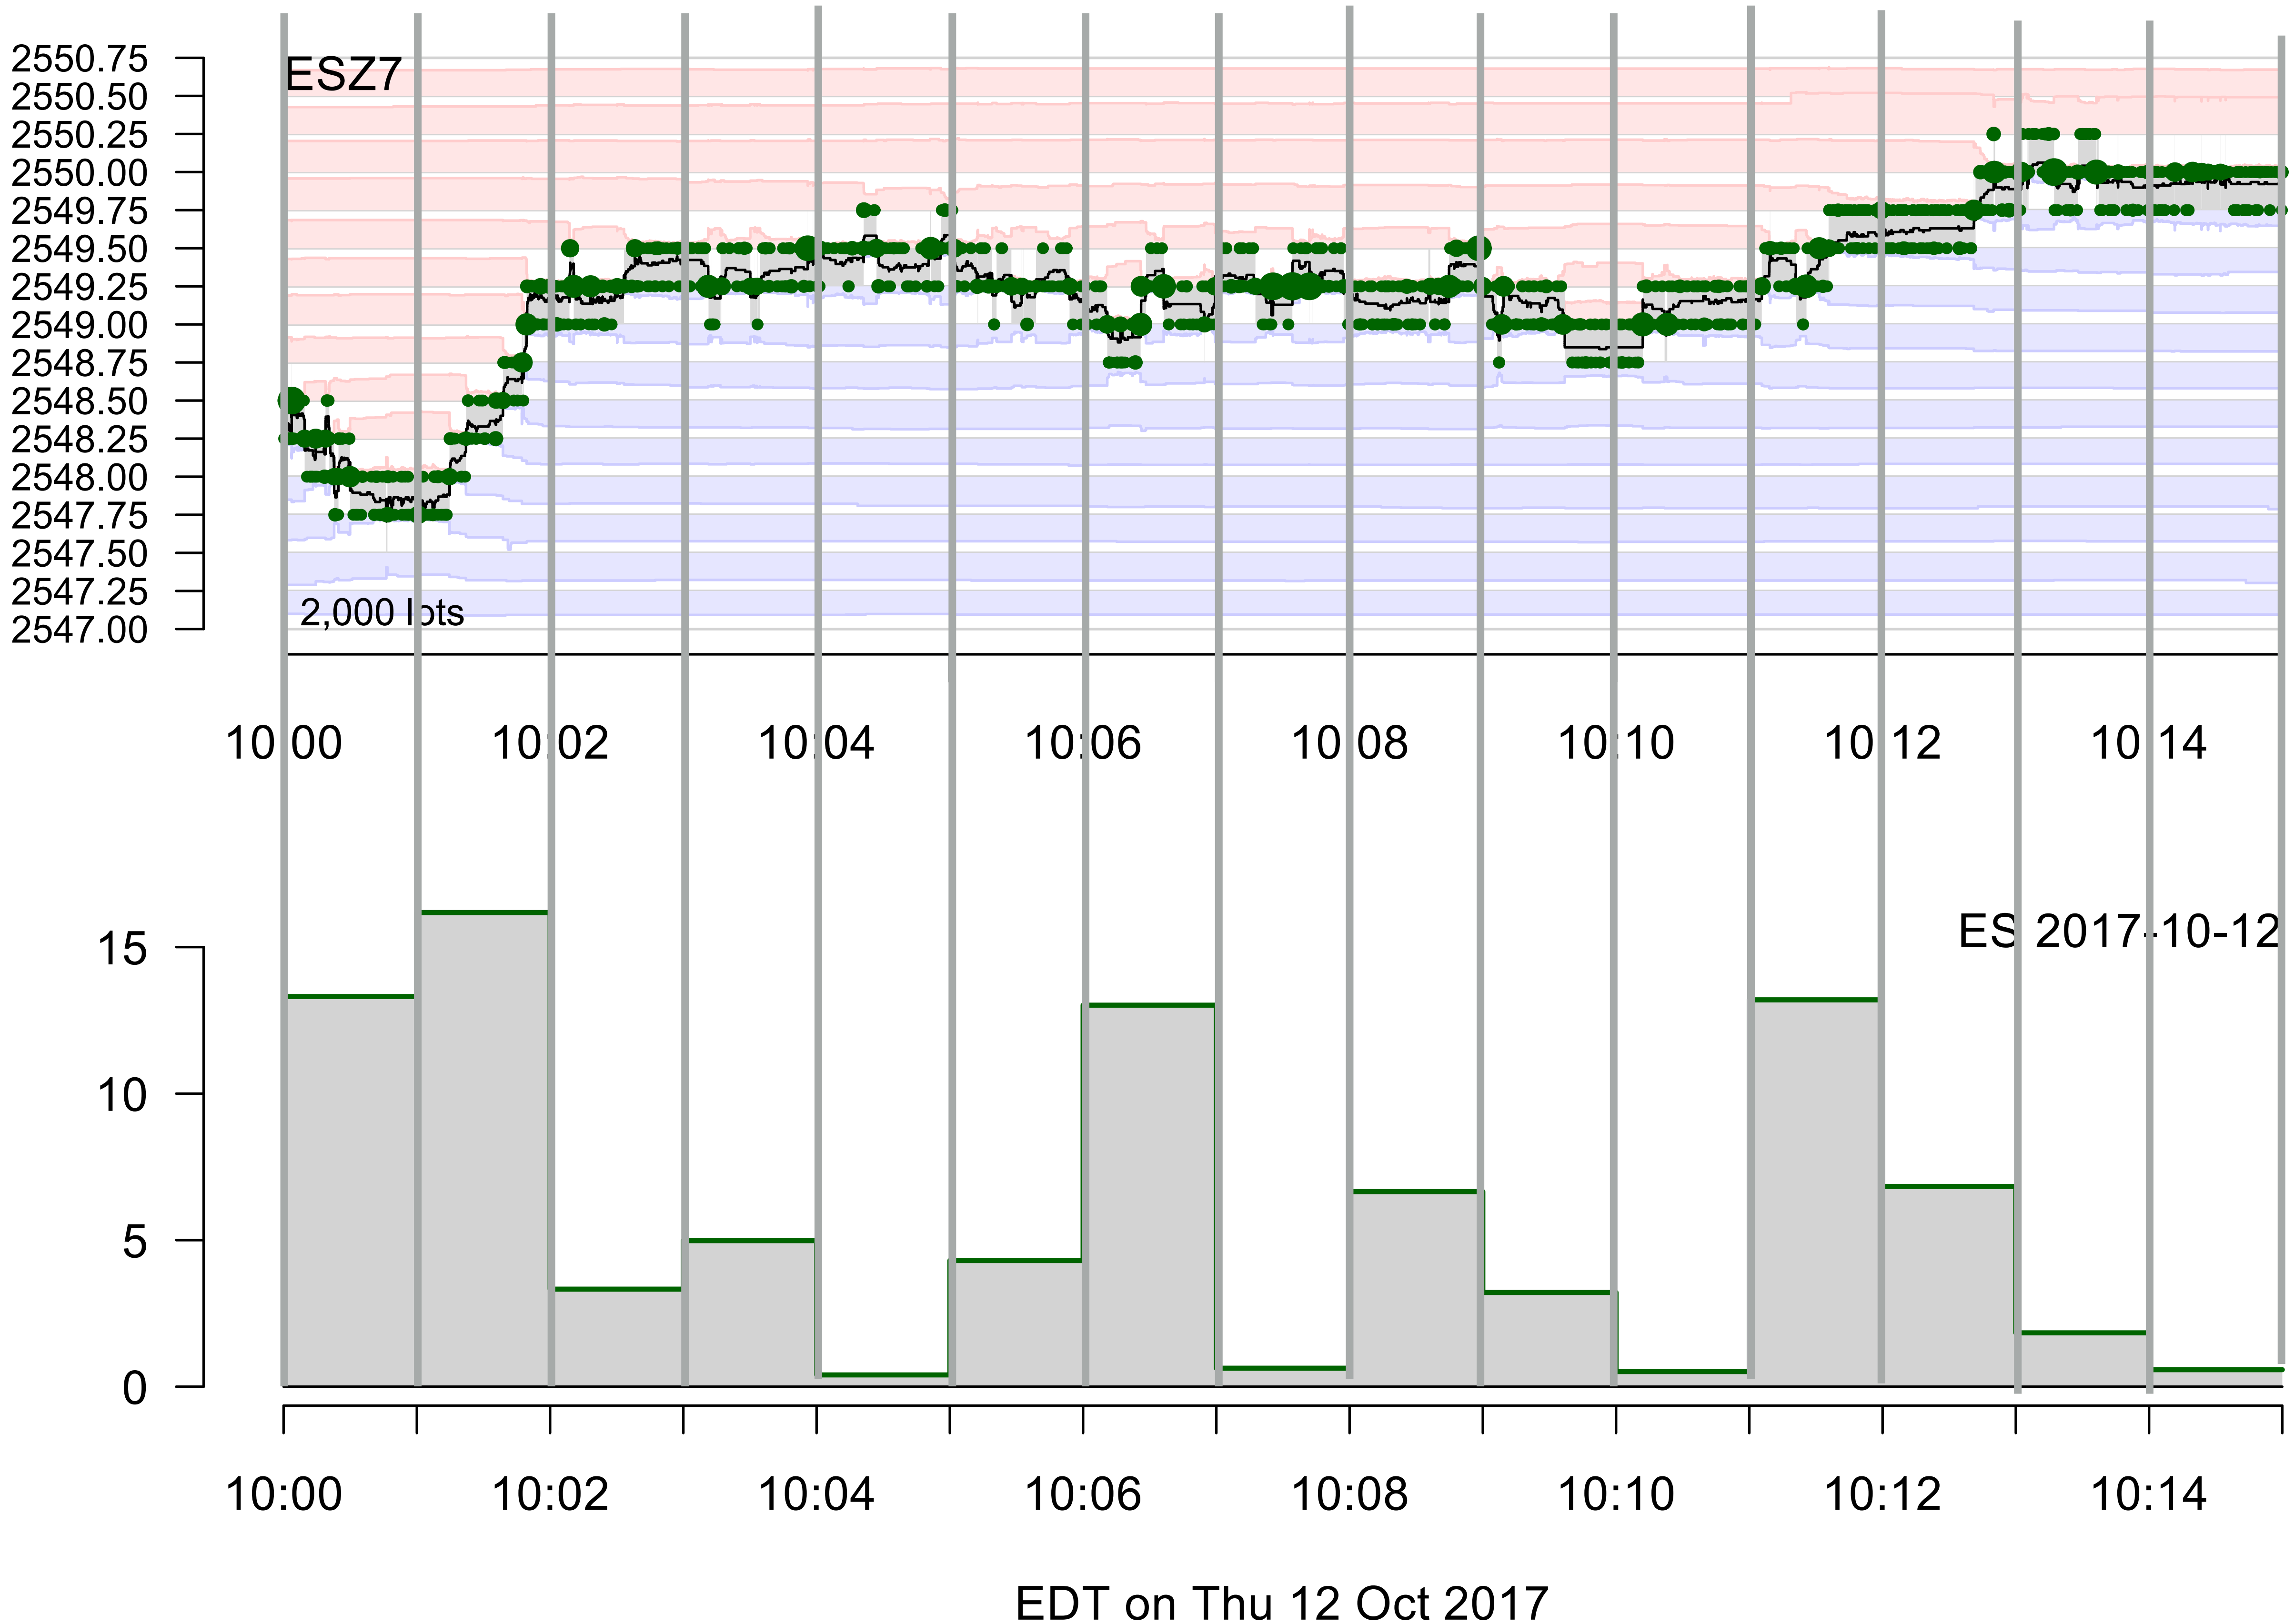

quantitativebrokers

How to evaluate a realized volatility?

Simulated data

- zero-intelligence market (Oomen/Gatheral)
- pure Brownian motion once time scale is longer than microstructure scale

Real data

- unknown price process

Zero-intelligence market

QUANTITATIVE FINANCE VOLUME 3 (2003) 481–514 RESEARCH PAPER
INSTITUTE OF PHYSICS PUBLISHING quant.iop.org

Statistical theory of the continuous double auction

Eric Smith[1], J Doyne Farmer, László Gillemot and Supriya Krishnamurthy

Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

E-mail: desmith@santafe.edu

Received 30 October 2002, in final form 16 July 2003
Published 18 September 2003
Online at stacks.iop.org/Quant/3/481

Abstract

Most modern financial markets use a continuous double auction mechanism to store and match orders and facilitate trading. In this paper we develop a microscopic dynamical statistical model for the continuous double auction under the assumption of IID random order flow, and analyse it using simulation, dimensional analysis, and theoretical tools based on mean field approximations. The model makes testable predictions for basic properties of markets, such as price volatility, the depth of stored supply and demand versus price, the bid–ask spread, the price impact function, and the time and probability of filling orders. These predictions are based on properties of order flow and the limit order book, such as share volume of market and limit orders, cancellations, typical order size, and tick size. Because these quantities can all be measured directly there are no free parameters. We show that the order size, which can be cast as a non-dimensional granularity parameter, is in most cases a more significant determinant of market behaviour than tick size. We also provide an explanation for the observed highly concave nature of the price impact function. On a broader level, this work suggests how stochastic models based on zero intelligence agents may be useful to probe the structure of market institutions. Like the model of perfect rationality, a stochastic zero intelligence model can be used to make strong predictions based on a compact set of assumptions, even if these assumptions are not fully believable.

quantitativebrokers

Zero-intelligence for variance measurement

Finance Stoch (2010) 14: 249–283
DOI 10.1007/s00780-009-0120-1

Zero-intelligence realized variance estimation

Jim Gatheral · Roel C.A. Oomen

Received: 19 March 2007 / Accepted: 4 April 2009 / Published online: 26 January 2010
© Springer-Verlag 2010

Abstract Given a time series of intra-day tick-by-tick price data, how can realized variance be estimated? The obvious estimator—the sum of squared returns between trades—is biased by microstructure effects such as bid–ask bounce and so in the past, practitioners were advised to drop most of the data and sample at most every five minutes or so. Recently, however, numerous alternative estimators have been developed that make more efficient use of the available data and improve substantially over those based on sparsely sampled returns. Yet, from a practical viewpoint, the choice of which particular estimator to use is not a trivial one because the study of their relative merits has primarily focused on the speed of convergence to their asymptotic distributions, which in itself is not necessarily a reliable guide to finite sample performance (especially when the assumptions on the price or noise process are violated). In this paper we compare a comprehensive set of nineteen realized variance estimators using simulated data from an artificial "zero-intelligence" market that has been shown to mimic some key properties of actual markets. In evaluating the competing estimators, we concentrate on efficiency but also pay attention to implementation, practicality, and robustness.

quantitativebrokers

ZI did not work for algo simulation

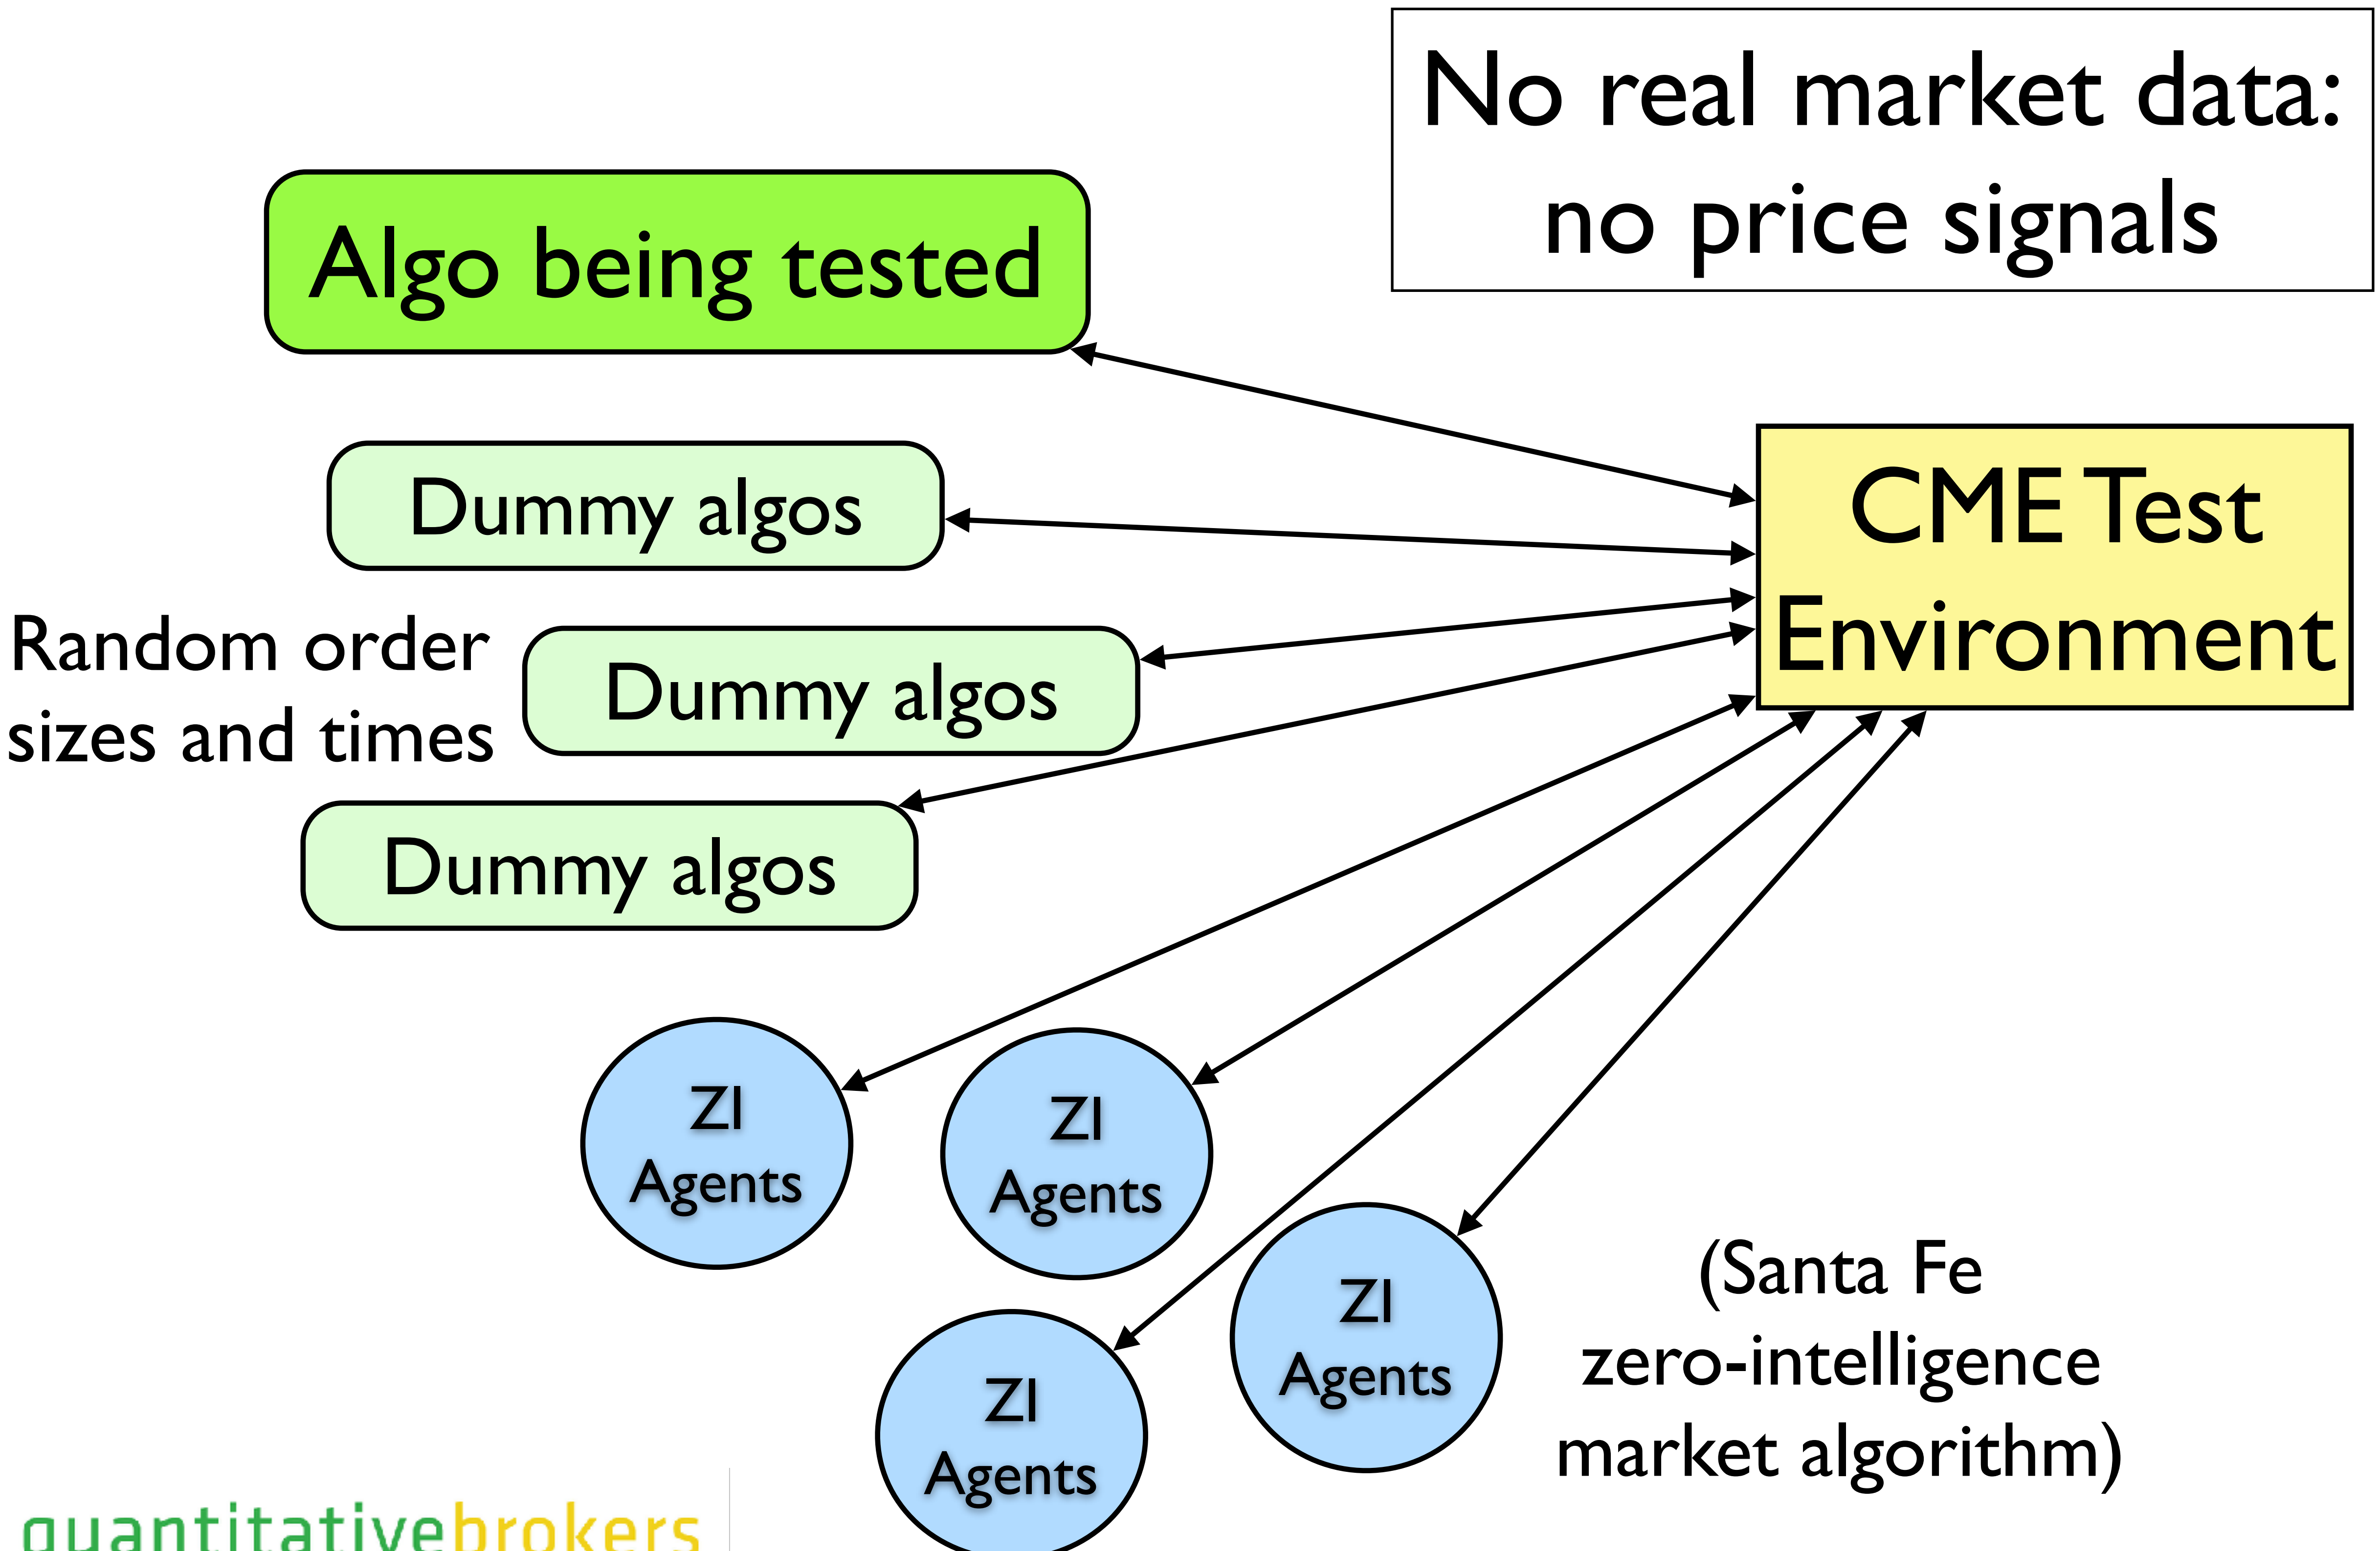

Add some "intelligence" to random market

Simulating and Analyzing Order Book Data: The Queue-Reactive Model

Weibing HUANG, Charles-Albert LEHALLE, and Mathieu ROSENBAUM

Journal of the American Statistical Association
March 2015, Vol. 110, No. 509, Applications and Case Studies

Through the analysis of a dataset of ultra high frequency order book updates, we introduce a model which accommodates the empirical properties of the full order book together with the stylized facts of lower frequency financial data. To do so, we split the time interval of interest into periods in which a well chosen reference price, typically the midprice, remains constant. Within these periods, we view the limit order book as a Markov queuing system. Indeed, we assume that the intensities of the order flows only depend on the current state of the order book. We establish the limiting behavior of this model and estimate its parameters from market data. Then, to design a relevant model for the whole period of interest, we use a stochastic mechanism that allows to switch from one period of constant reference price to another. Beyond enabling to reproduce accurately the behavior of market data, we show that our framework can be very useful for practitioners, notably as a market simulator or as a tool for the transaction cost analysis of complex trading algorithms.

KEY WORDS: Ergodic properties; Execution probability; High frequency data; Jump Markov process; Limit order book; Mechanical volatility; Market impact; Market microstructure; Market simulator; Queuing model; Transaction costs analysis; Volatility.

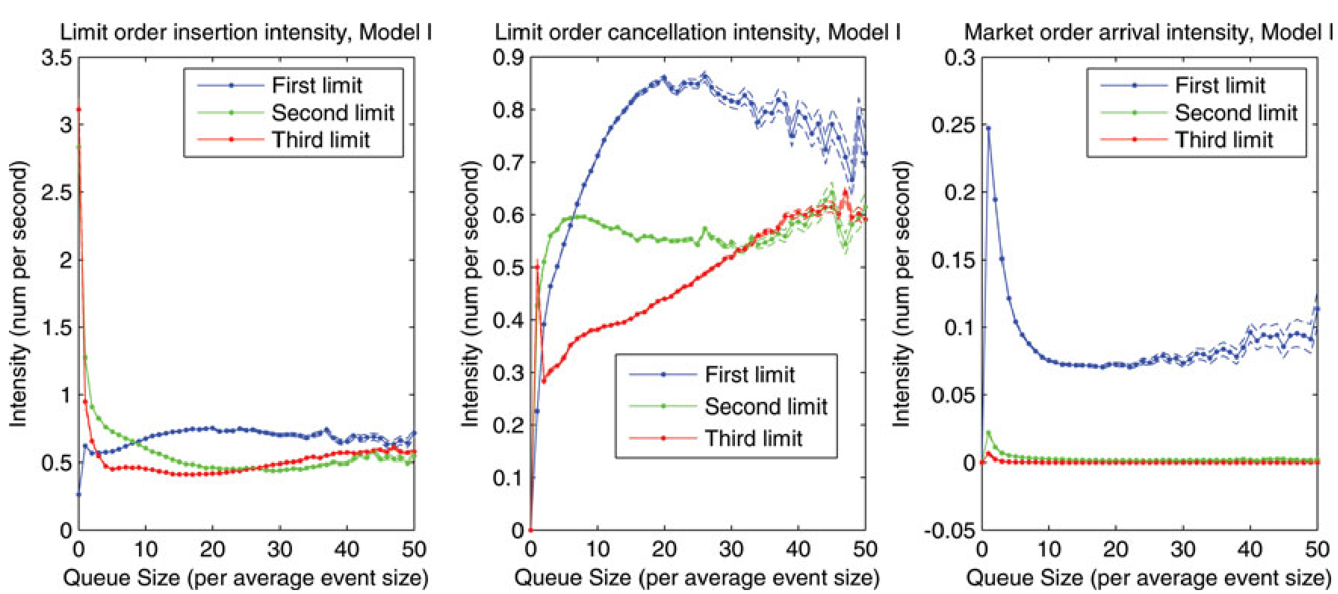

Figure 2. Intensities at $Q_{\pm i}$, $i = 1, 2, 3$, France Telecom.

4. CONCLUSION AND PERSPECTIVES

In this work, we have modeled market participants intelligence through their average behaviors toward various states of the LOB. This enabled us to analyze the different order flows and to design a suitable market simulator for practitioners, allowing notably to investigate the transaction costs of complex trading strategies. To our knowledge, our model is the first one where such pre-trade cost analysis is possible in a simple and efficient way.

Another important public information, the historical order flow, is not considered in this approach. Market order flows have been shown to be autocorrelated in several empirical studies (see, e.g., Toth et al. 2011b). Thus, adding such feature in our framework would probably be relevant. Another possible direction for future research would be to explain the shape of the estimated intensity functions in a more sophisticated way. For example, it would be interesting to design some agent based model where these repetitive patterns of the LOB dynamics would be reproduced, providing an even better understanding of the nature of these intensity curves.

Realistic simulation using real market data

The PLAT Project has developed a trading simulation that merges automated clients with real-time, real-world stock market data. This simulation has been used for three competitions.

The Penn-Lehman Automated Trading Project

Michael Kearns and Luis Ortiz, *University of Pennsylvania*

The Penn-Lehman Automated Trading Project is a broad investigation of algorithms and strategies for automated trading in financial markets. The PLAT Project's centerpiece is the Penn Exchange Simulator (PXS), a software simulator for automated stock trading that merges automated client orders for shares with real-world, real-time order data. PXS automatically computes client profits and losses, volumes traded, simulator and external prices, and other quantities of interest. To test the effectiveness of PXS and of various trading strategies, we've held three formal competitions between automated clients.

We also actively use PXS as a platform for developing novel, principled automated trading strategies (clients). The real-data, real-time nature of PXS lets us examine computationally intensive, high-frequency, high-volume trading strategies (although this last property always presents the challenges of estimating the *market impact*—the effect on prices). We're particularly interested in developing clients that make predictive use of limit order book data, including those using statistical modeling and machine learning. We hope that, over time, the project will generate a library of clients with varying features (trading strategy, volume, frequency, and so on) that can serve to create realistic simulations with known properties.

1094-7167/03/$17.00 © 2003 IEEE
Published by the IEEE Computer Society
IEEE INTELLIGENT SYSTEMS

2003

quantitativebrokers

waters
trading technologies for financial-market professionals

Managed Futures' ALGO CHASE

A potent mixture of in-house, futures commission merchant, and boutique brokerage-provided algorithms now play a part in commodity trading advisors' and managed futures funds' trading activities. Tim Bourgaize Murray examines why a new cadre of simulation tools is helping to organize—and perhaps re-mold—these buy-side specialists' order flow.

QB Simulator

"Skate to where the puck is going to be, not where it has been," Wayne Gretzky once told an interviewer. As the Great One described it, what cuts certain players a level above isn't native instinct alone, so much as endless practice seeing the ice and, frankly, the hard work of getting to where a scoring opportunity will be, before it reveals itself.

Gretzky's advice is one of Robert Almgren's favorite lines—but not because the co-founder of Quantitative Brokers (QB) is a hockey fan. Instead, he says a similar idea applies to the business of algorithmic futures execution: the more you see, the more you

"We do perform reviews on all algos internally using our own simulator, and are always keen to compare these results with those of the provider. If they cannot provide a simulator, it takes a lot longer to see if we believe their story." Murray Steel, AHL

much of the value, they say, derives from what comes before any trades are even made.

April 2014 waterstechnology.com

SALIENT POINTS

- Managed futures specialists are increasingly taking advantage of boutique agency brokers' algorithms, citing their ability to be opportunistic and adjust to markets' behavior, as well as faster speed to implementation and greater alpha realized through price slippage.
- Rates futures, particularly, are ripe for these applications given their correlation and the characteristics of the complexes within which they're traded, and are well-serviced by Quantitative Brokers (QB), among other independent shops. Hedge fund AHL and CTA Revolution Capital Management are among QB's users for rates.
- Another value-added feature at smaller shops like QB is their simulation environments, which mimic the matching engine logic of relevant futures exchange venues and can test new adjustments to algorithms with real-time market data before putting the algos into production.
- Sources expect a greater variety of such brokers to crop up in coming years, while sell-side futures commission merchants (FCMs), sensing greater competition, are also expected to mature their offerings and continue bundling futures algos with other execution and clearing services.

Waters Technology
Tim Murray
April 2014

quantitativebrokers

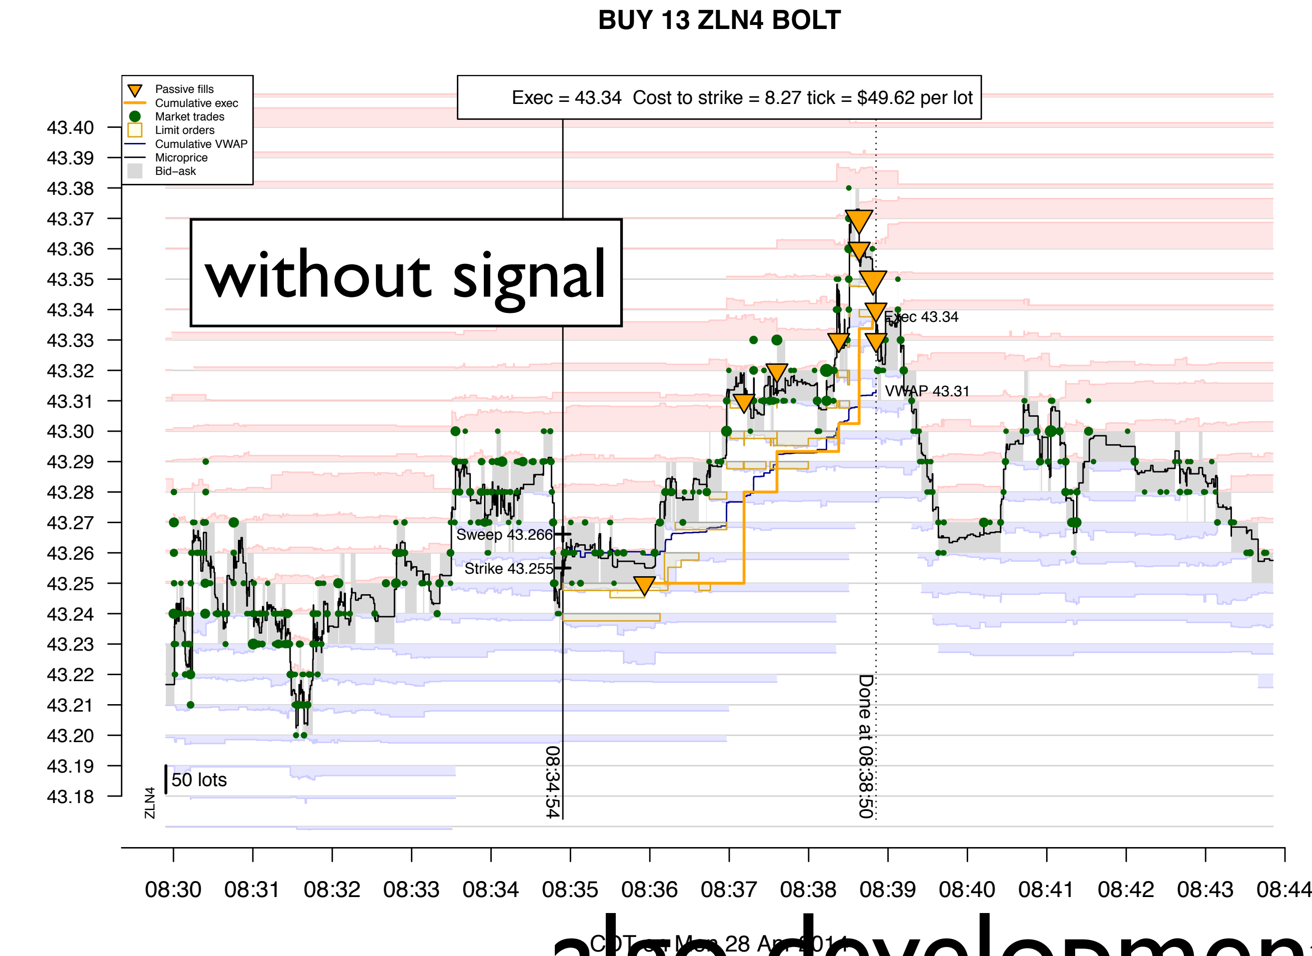

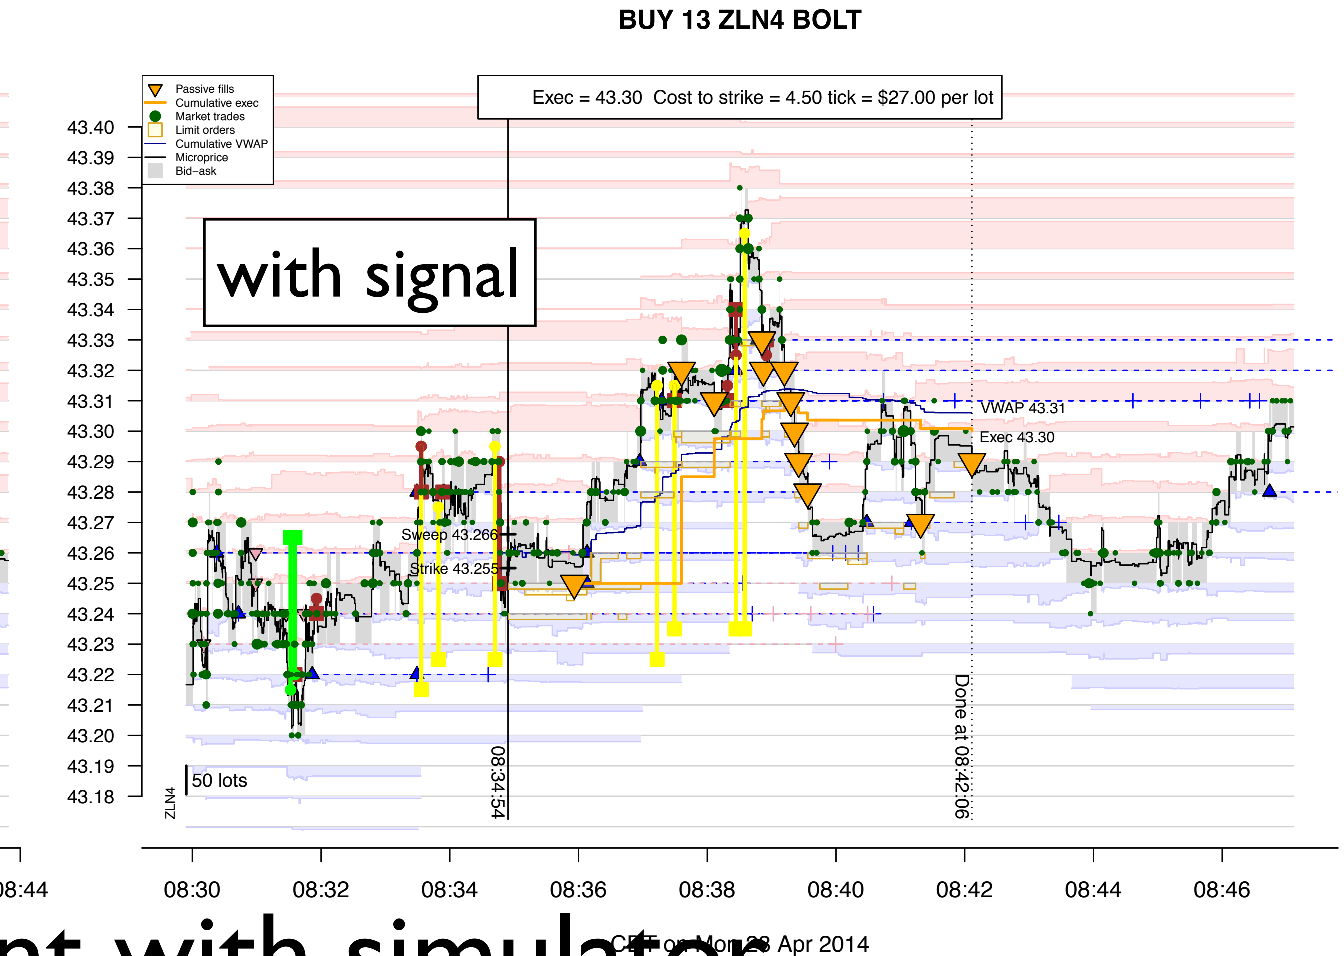

algo development with simulator on real market data

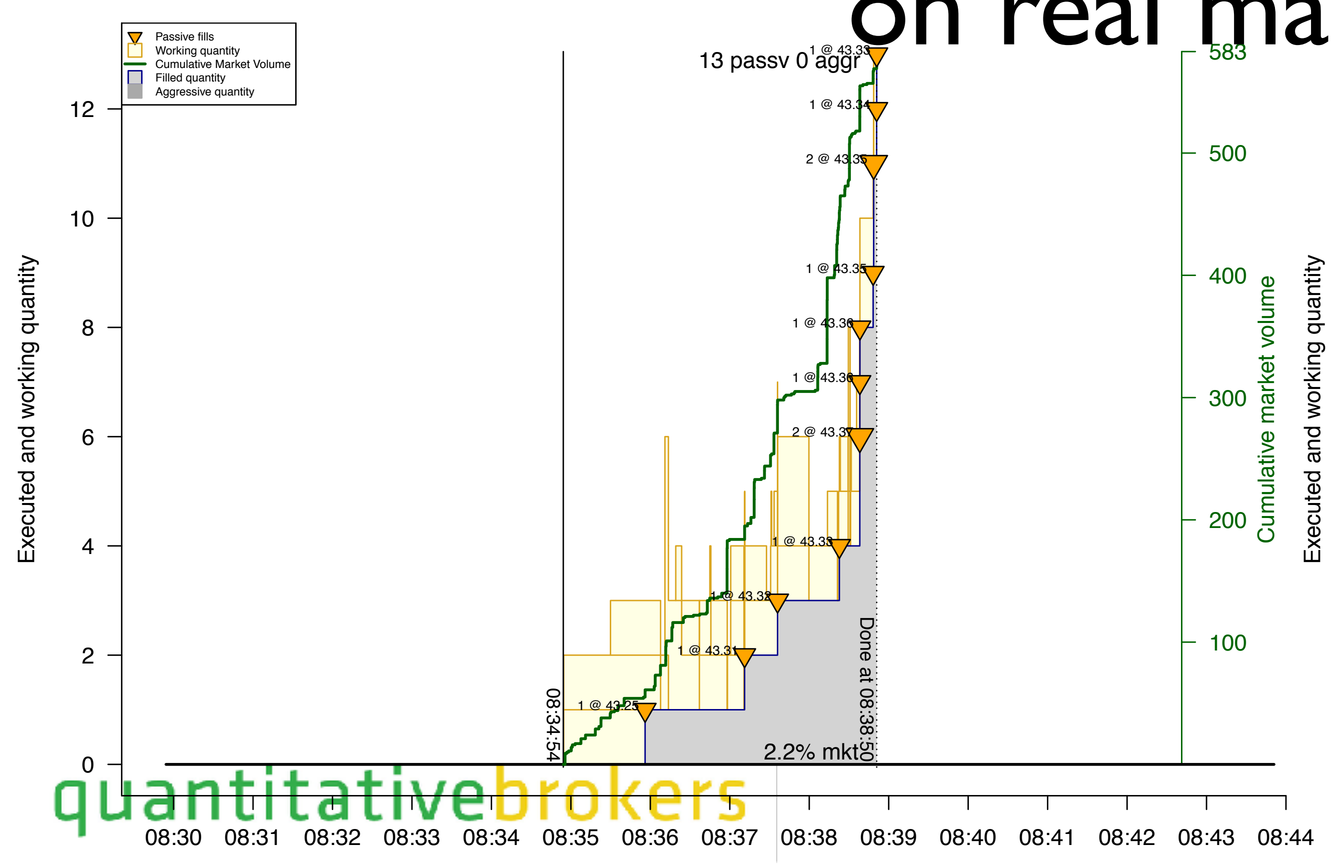

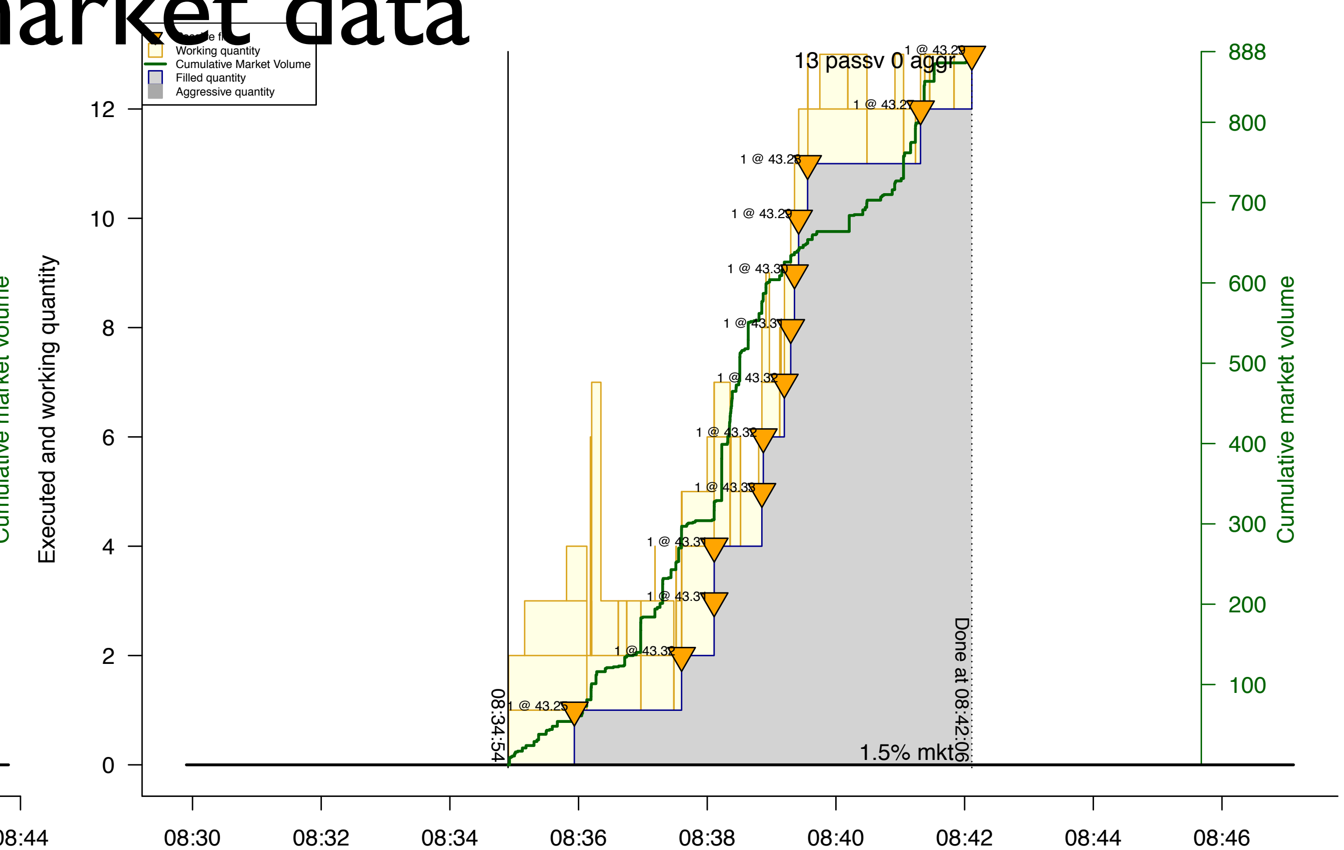

quantitativebrokers

Produced by QB from CME and internal data

We want to test estimators on real data

Problem: do not know real volatility

Method: see if we can normalize price changes over finite time intervals

Realized variance

Price process $$dX(t) = \sigma(t)\, dB(t)$$

$\sigma(t)$ random or deterministic

Integrated variance $$Q(t_{\mathrm{L}}, t_{\mathrm{R}}) = \int_{t_{\mathrm{L}}}^{t_{\mathrm{R}}} \sigma(t)^2\, dt$$

For given realization of $\sigma(t)$,

$$X(t_{\mathrm{R}}) - X(t_{\mathrm{L}}) \sim \mathcal{N}\big(\, 0,\ Q(t_{\mathrm{L}}, t_{\mathrm{R}})\,\big)$$

$$\frac{X(t_{\mathrm{R}}) - X(t_{\mathrm{L}})}{\sqrt{Q(t_{\mathrm{L}}, t_{\mathrm{R}})}} \text{ is unit normal}$$

quantitativebrokers

More generally

Any X(t) is
time-changed
Brownian motion

$$X(t) = B(\tau(t))$$

$X(t_\mathrm{R}) - X(t_\mathrm{L})$ is highly non-normal

$$\frac{X(t_\mathrm{R}) - X(t_\mathrm{L})}{\sqrt{\tau(t_\mathrm{R}) - \tau(t_\mathrm{L})}} \text{ is unit normal}$$

quantitativebrokers

Idea:

Measure $Q(t_\mathrm{L}, t_\mathrm{R})$ on many historical intervals

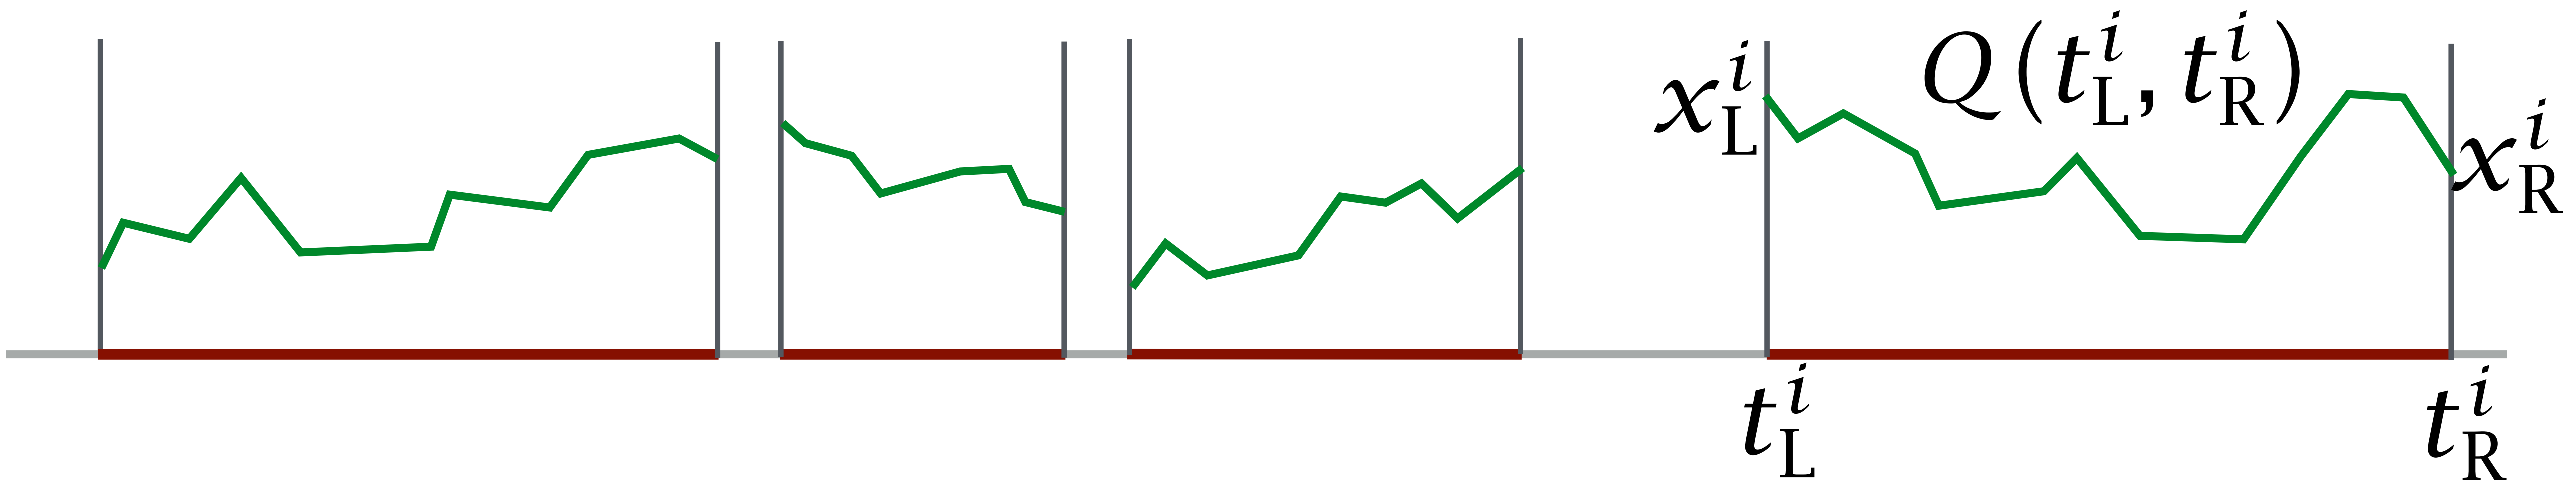

See if scaled price changes reduce to a unit normal population

$$z_i = \frac{x_\mathrm{R}^i - x_\mathrm{L}^i}{\sqrt{Q(t_\mathrm{L}^i, t_\mathrm{R}^i)}}$$

This avoids the problem of forecasting or modeling Q

Realized variance measurement

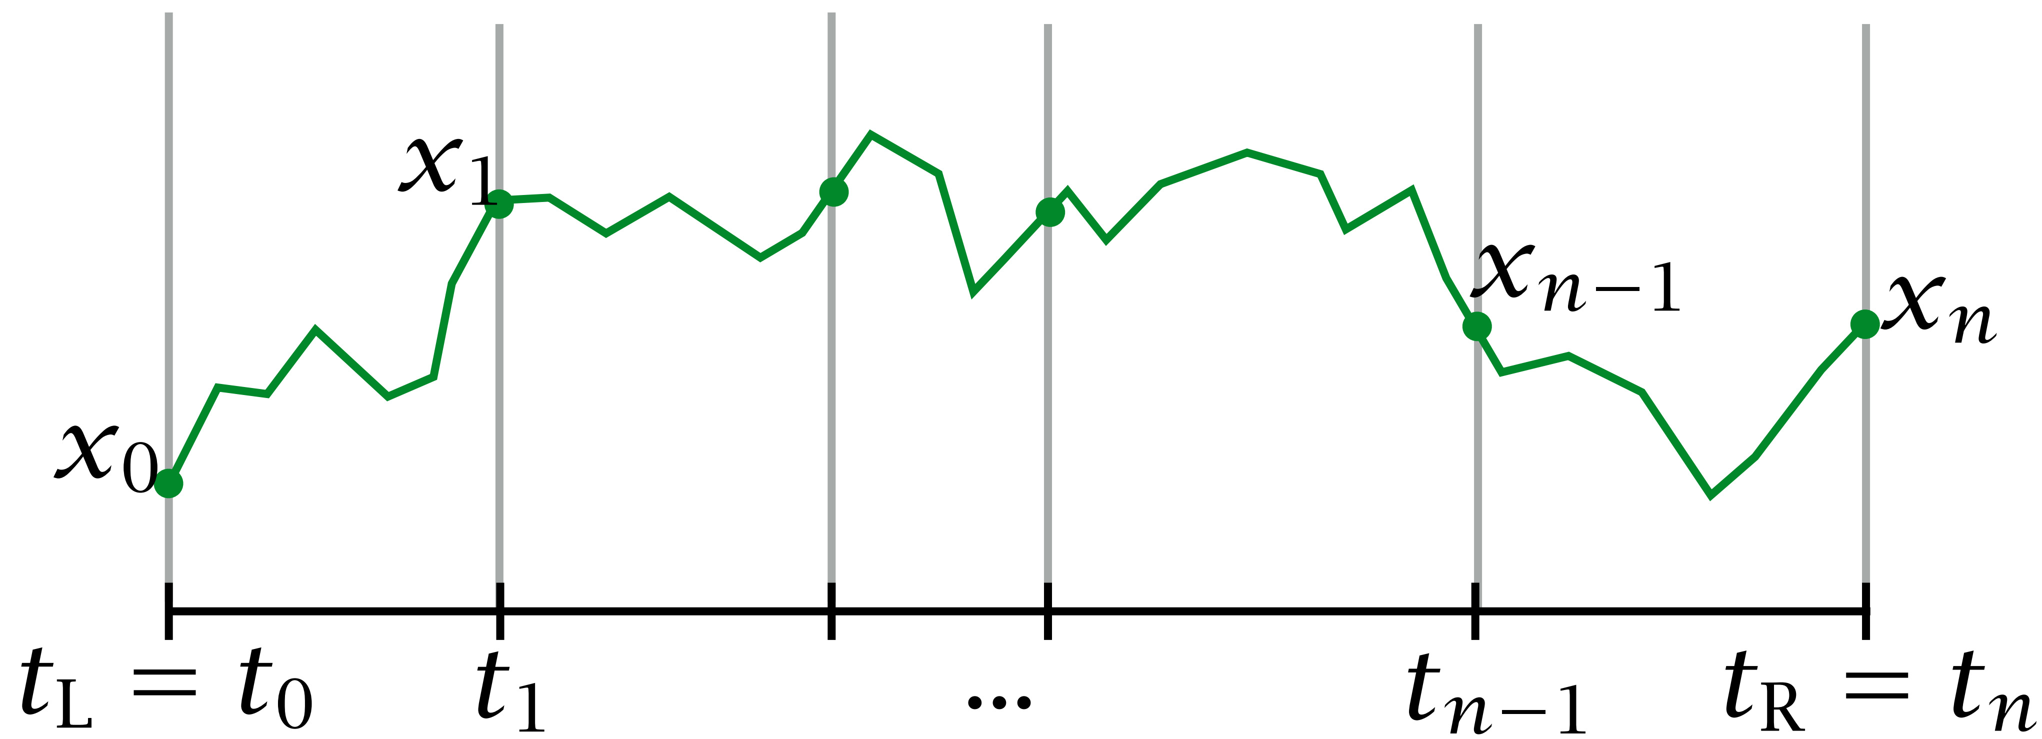

Simplest estimator for Q:
Realized variance

$$RV(t_L, t_R) = \sum_{j=1}^{n} (x_j - x_{j-1})^2$$

Susceptible to microstructure noise (serial correlation)

Design of estimators

What time points to sample at?

quote updates, trades, etc

What prices to use at those times?

trade prices, bid-ask midpoint, microprice, etc

How to correct for serial correlation?

multiscale estimators, etc

Sample time points

Times of trades
Times of trades at new price
Times of new midquote price

quantitativebrokers

Difficulty with using all trades: many repeated trades at same price

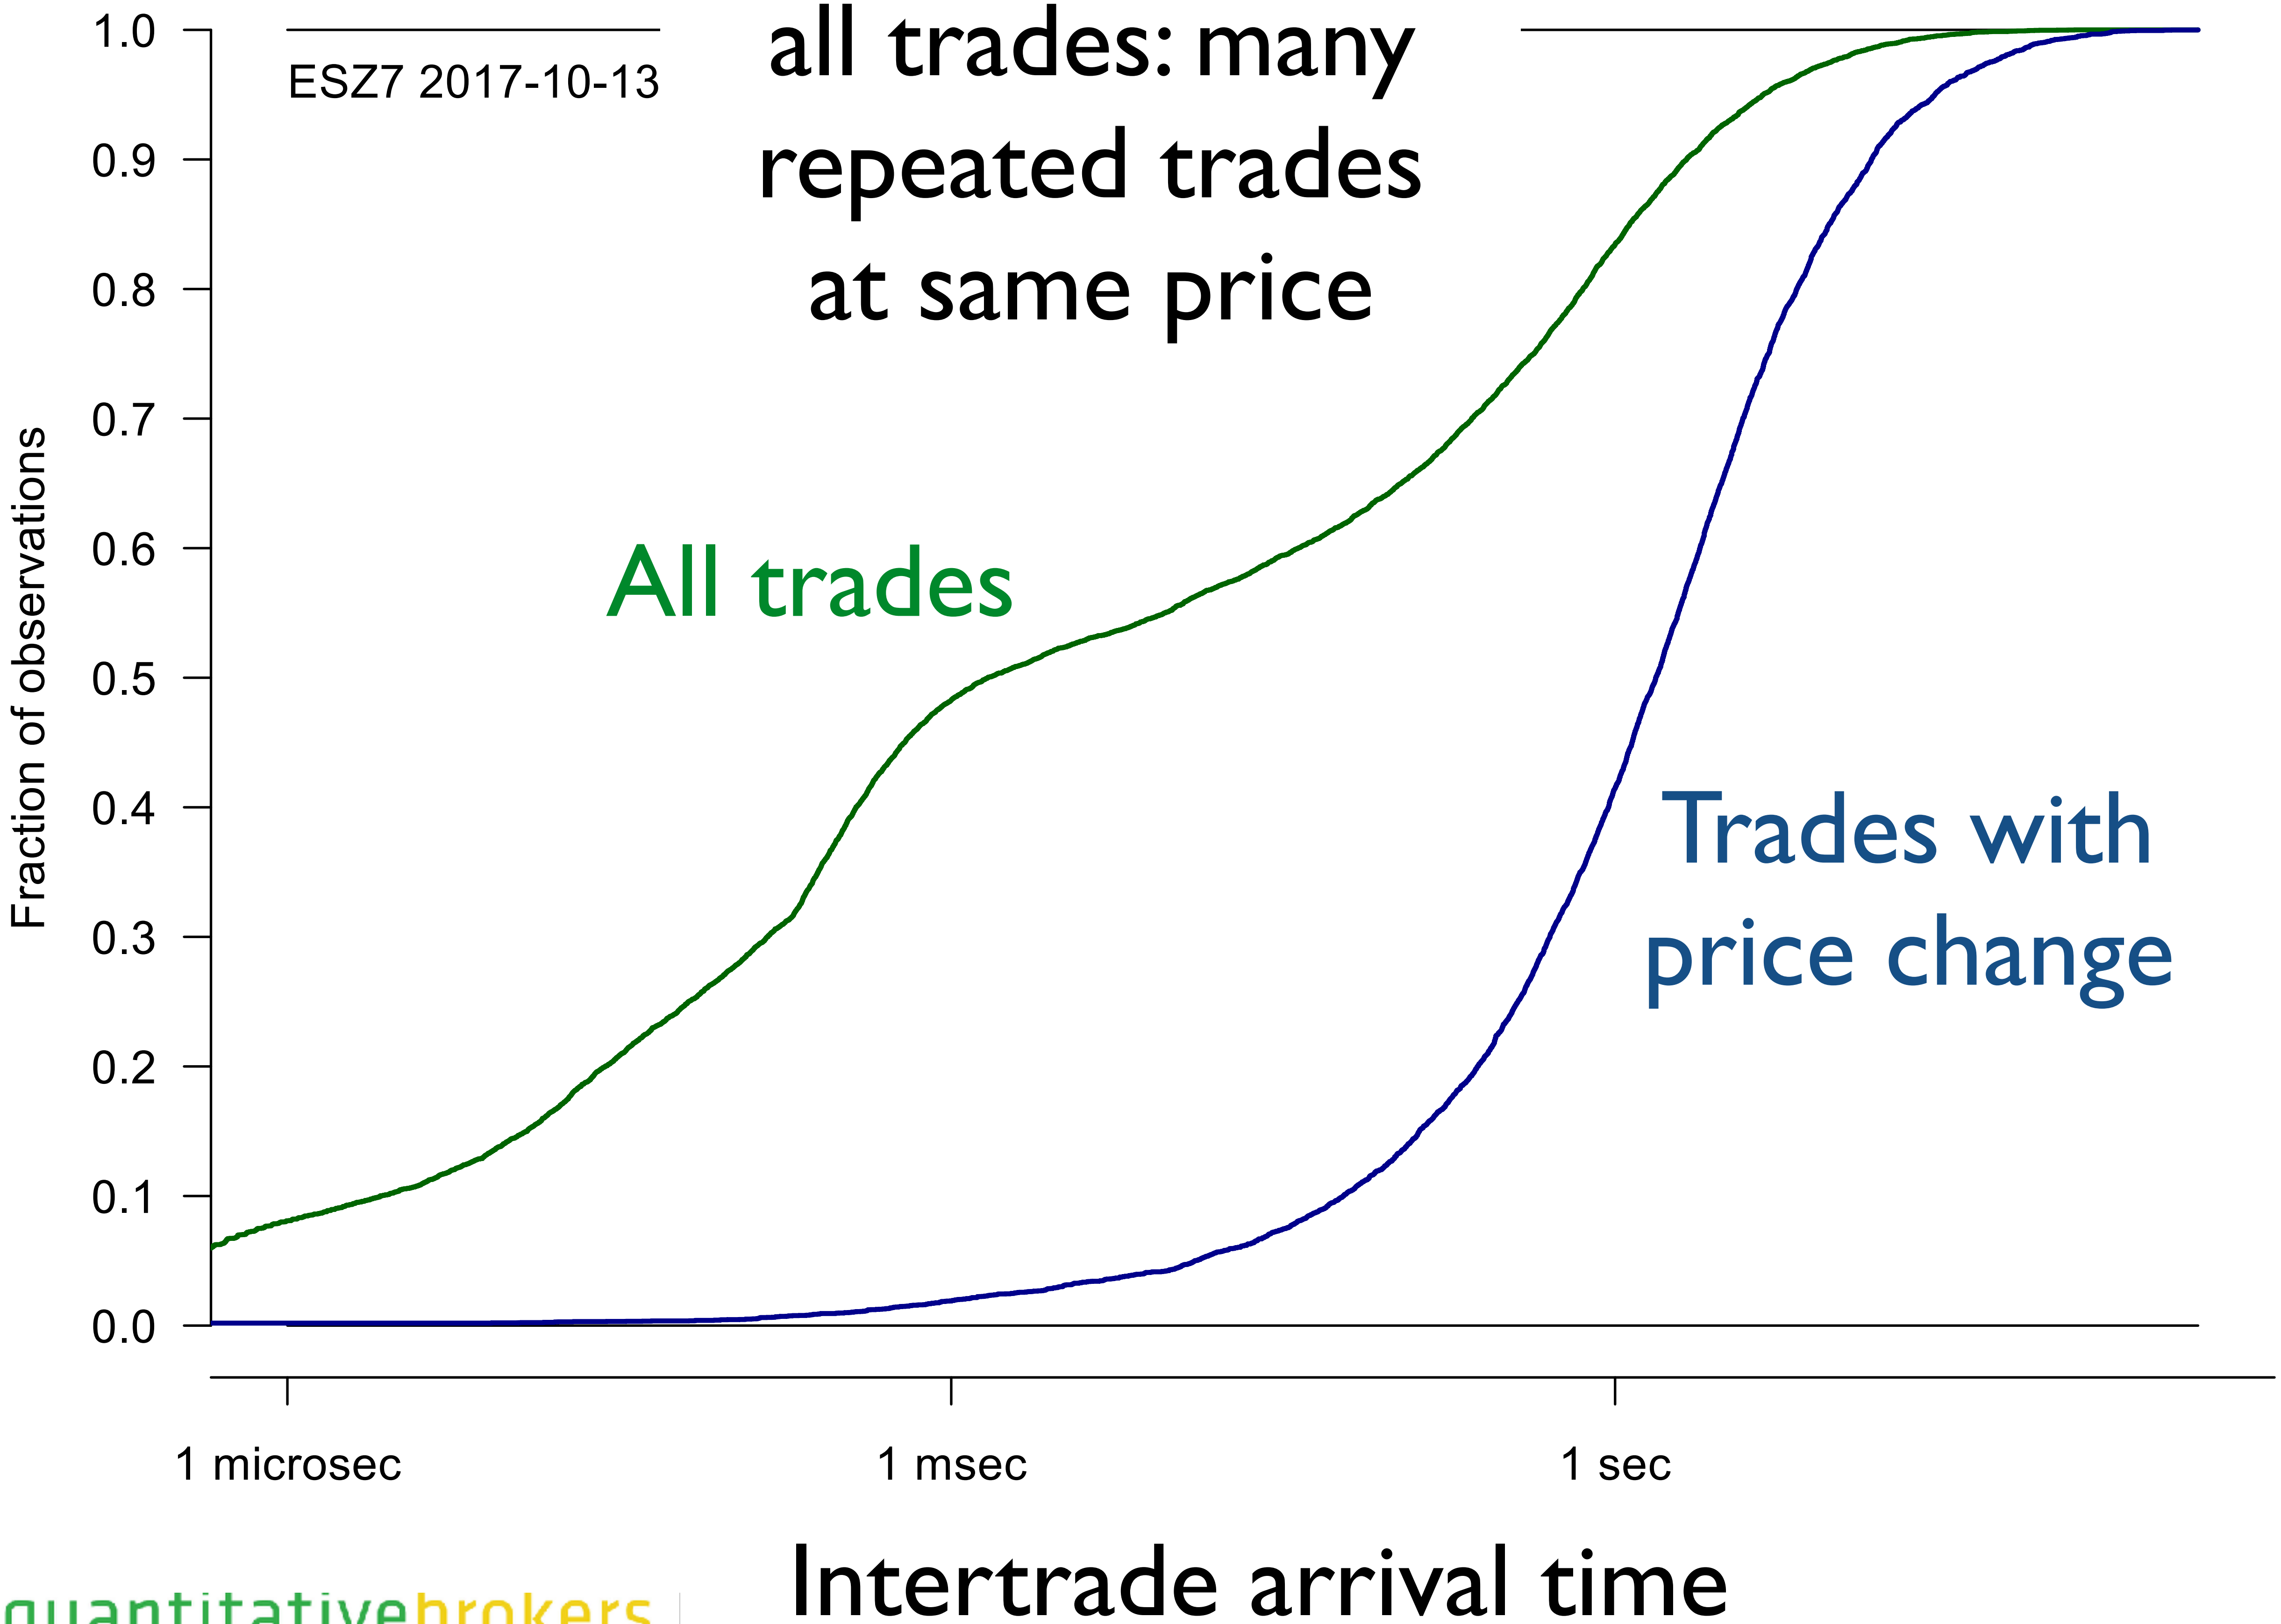

quantitativebrokers

Sample prices

Trade price

- bid-ask bounce

Mid-quote price

- also has negative autocorrelation but less

Microprice

- adds random noise

Time/price pairs

AT	all trade prices
TC	trade prices when change
MQC	mid-quote prices when change
MQBT	mid-quote before trade time
MQBTC	mid-quote before trade at new price

RV estimators

RV	realized variance
AR	AR(1) model on return
Zhou	Bias-corrected RV of Zhou (1996)
ZMA	2-scale RV Zhang, Mykland, Ait-Sahalia 2005
MSRV	multi-scale RV
KRV-cubic	kernel estimator, cubic kernel
KRV-th2	Tukey-Hanning kernel TH_2
KRV-th16	Tukey-Hanning kernel TH_{16}

Set all lags arbitrarily to q=5

3 CME futures products

10-year T-Note (ZN)

large tick, low activity

E-mini SP500 (ES)

large tick, high activity

Henry Hub Natural Gas (NG)

small tick, mid activity

Time intervals

Data on calendar year 2016 (252 days)
On each day:

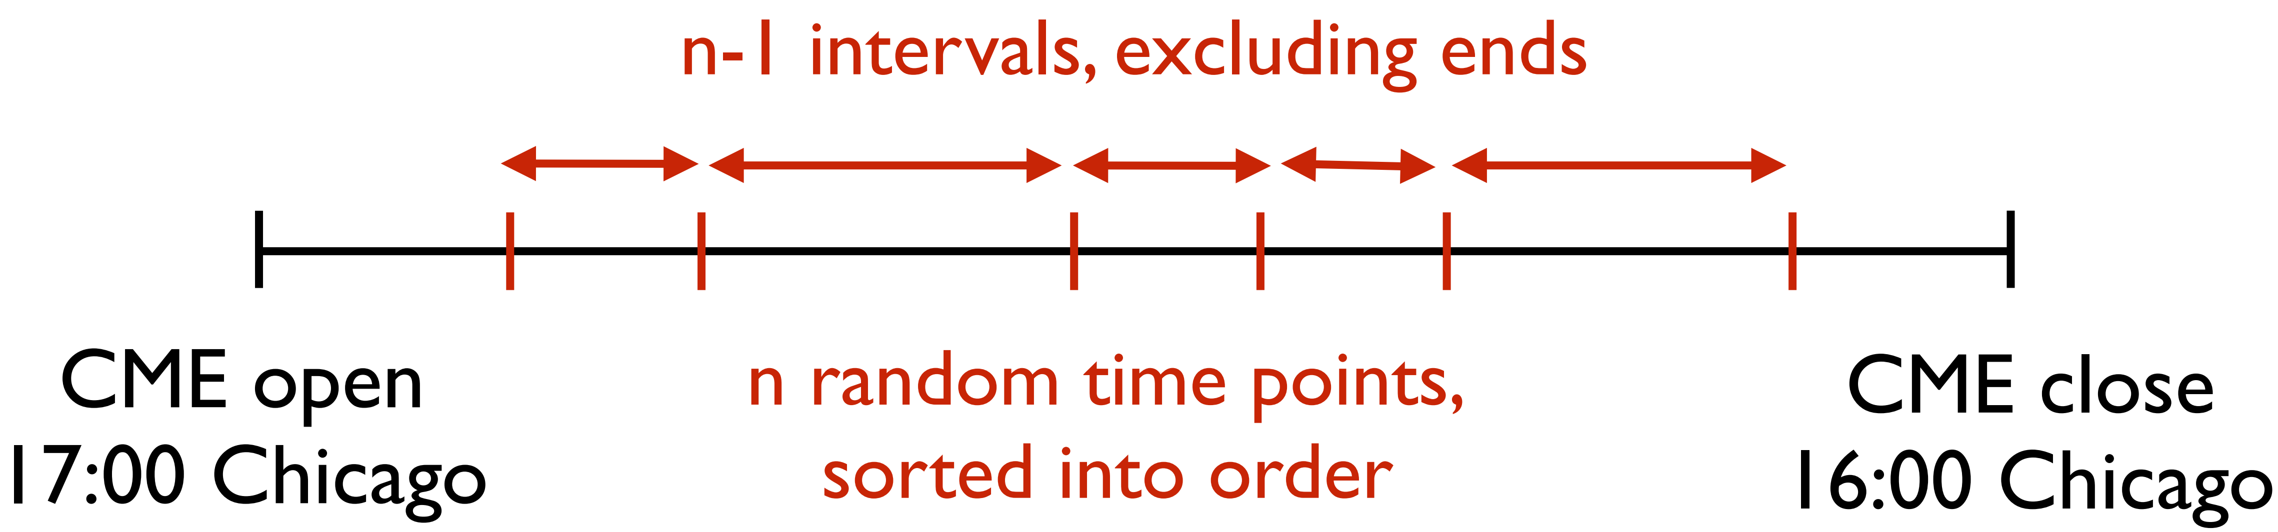

Interval lengths approximately Poisson,
non-overlapping.
Do not exclude special events, etc

Evaluate estimator on 1-minute bins

Aggregate into selected intervals

Ask

- Are the scaled price changes normal?
- Is the variance equal to 1?

Normality test

Shapiro-Wilks test

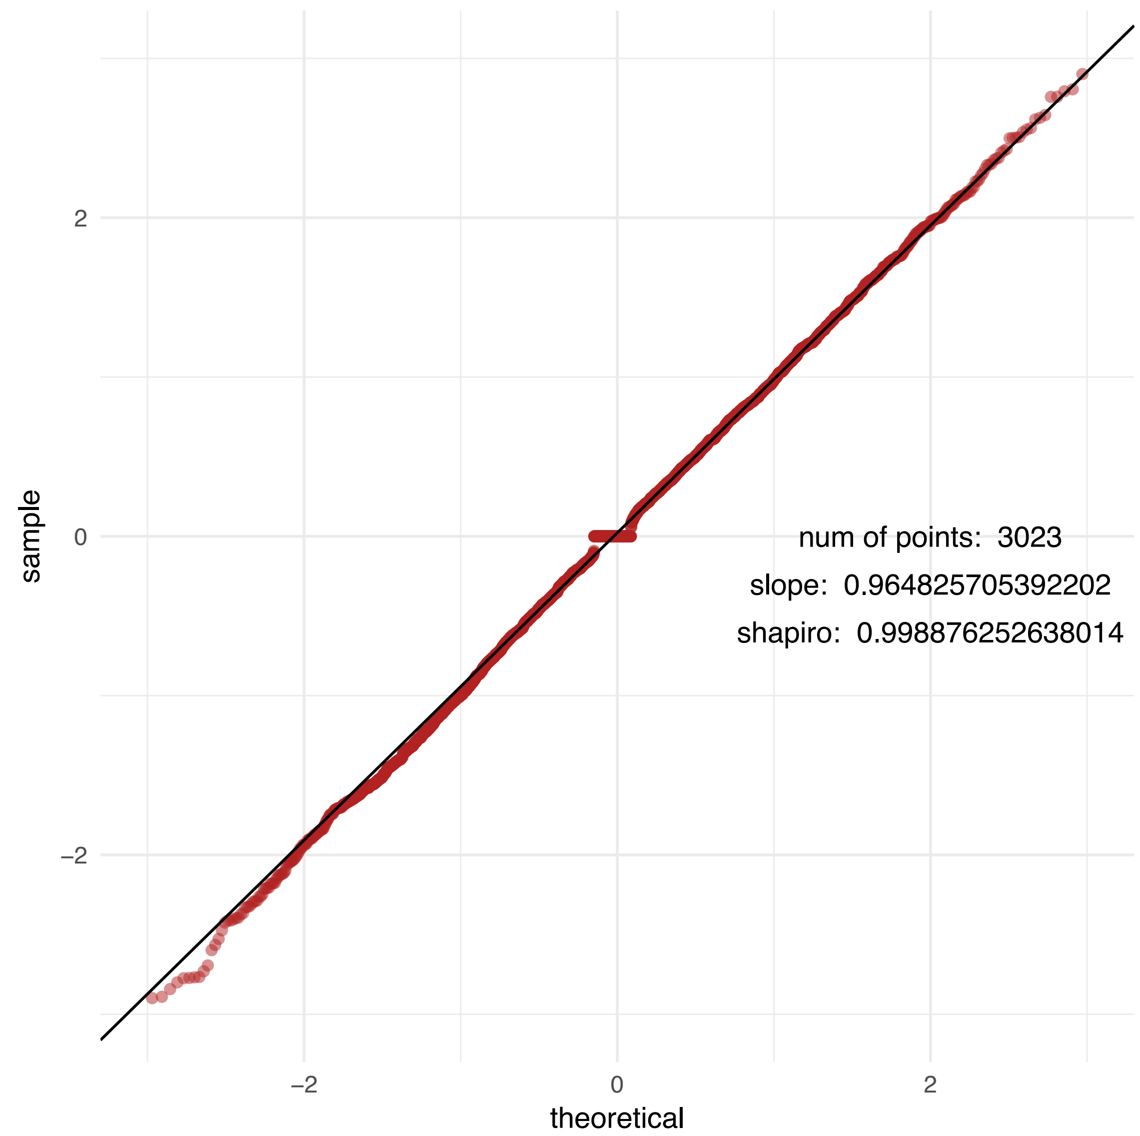

Are they normal?

ZN

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.1201	0.0419**	0.0143**	0.0361**	0.1541
ar	0.0664*	0.1190	0.0015**	0.0122**	0.1973
zhou	0.1461	0.4158	0.1632	0.0441**	0.2450
zma	0.3179	0.2967	0.2741	0.1282	0.2680
msrv	0.1814	0.2403	0.1393	0.0680*	0.2271
krv-cubic	0.3026	0.1781	0.3624	0.2367	0.2624
krv-th2	0.2801	0.2030	0.2287	0.1730	0.2095
krv-th16	0.0718*	0.3117	0.1042	0.0704*	0.1448

ES

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.0467**	0.0862*	0.0106**	0.0076***	0.5416
ar	0.0910*	0.1594	0.0046***	0.0302**	0.6021
zhou	0.3308	0.2040	0.0803*	0.2345	0.5456
zma	0.6161	0.1758	0.5367	0.2989	0.4415
msrv	0.4095	0.4183	0.1959	0.0969*	0.4559
krv-cubic	0.4549	0.5646	0.3428	0.1627	0.4534
krv-th2	0.4710	0.1677	0.1525	0.1531	0.4901
krv-th16	0.3770	0.4143	0.0944*	0.1138	0.3238

NG

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.5087	0.3480	0.1506	0.2835	0.3901
ar	0.4174	0.2925	0.0533*	0.2233	0.4033
zhou	0.3857	0.0057***	0.0654*	0.2069	0.3941
zma	0.3082	0.4556	0.3542	0.1589	0.3001
msrv	0.2839	0.4123	0.3311	0.3140	0.2508
krv-cubic	0.3064	0.1279	0.3030	0.3635	0.2593
krv-th2	0.3825	0.2586	0.2347	0.1888	0.2530
krv-th16	0.3995	0.0006***	0.0297**	0.2371	0.3261

quantitativebrokers

Are the variances equal to 1?

ZN

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.3612	0.3635	0.4752	0.3802	0.8648
ar	0.3870	0.4994	0.4770	0.3842	0.8587
zhou	0.7262	1.4523	0.5155	0.4492	0.8148
zma	0.8904	1.4062	1.0422**	0.5837	1.0189**
msrv	0.8079	1.1137	1.2769	0.5243	0.9273*
krv-cubic	0.8317	0.9715**	1.2433	0.5429	0.9295*
krv-th2	0.7856	1.1748	1.3443	0.4991	0.8988
krv-th16	0.7301	1.4914	0.5182	0.4525	0.8114

ES

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.3668	0.3676	0.5227	0.4330	0.8446
ar	0.3936	0.4871	0.5209	0.4430	0.8400
zhou	0.6883	1.5516	0.5782	0.5356	0.8008
zma	0.8712	1.3647	1.0040***	0.7172	0.9969***
msrv	0.7782	1.0630*	1.0368**	0.6374	0.8826
krv-cubic	0.7984	0.9386*	1.0625*	0.6801	0.8985
krv-th2	0.7506	1.1239	1.0472**	0.6132	0.8877
krv-th16	0.6800	1.5051	0.5762	0.5392	0.8017

NG

estimator	at	tc	mqc	mqbt	mqbtc
rv	0.6643	0.6659	0.7589	0.7400	0.9551**
ar	0.6672	0.6883	0.7718	0.7492	0.9574**
zhou	0.8569	1.1062	0.9669**	0.8374	0.9177*
zma	0.9989***	1.0573*	1.0302**	1.0000***	1.0488**
msrv	0.8891	0.9317*	0.9273*	0.8881	0.9459*
krv-cubic	0.8961	0.9415*	0.9190*	0.9035*	0.9418*
krv-th2	0.8875	0.9419*	0.9313*	0.8812	0.9309*
krv-th16	0.8672	1.1214	0.9583**	0.8511	0.9040*

quantitativebrokers

Conclusion

Use ZMA

with mid-quote changes

or mid-quote before trade price changes

not before all trades

Gives good prediction of real price changes