# IPAR 1 to 16 October 2024

| | |
|---|---|
| trades on 16 October 2024 | 4,498 |
| average trading rate | 0.2 per second |
| presample for the diurnal profile | 17 to 30 September 2024 |
| outliers on 16 October 2024, $o$ | 244 (5.42%) |
| price half-life: run / minimizer[1] | 3.0 / 2.2 trades |
| cost of running 3 trades | +1.78% |
| scale half-life: run / minimizer[2] | 240 / 600 s |
| cost of running 240 s | +0.09% |
| direction half-life[3] | 545 s [450, 690] |
| Mandelbrot $a$ minimizing QLIKE | 26 trades |
| $a$ within 0.0005 of the minimum | 26 to 26 |
| linear $\hat{\sigma}$, before $\to$ at the largest residual[4] | 8% $\to$ 12% |
| Huber $\hat{\sigma}$, before $\to$ peak | 8% $\to$ 10% |
| median spot volatility, Huber / linear input[5] | 12.8 / 12.5 % |
| hyperbolic against HExp, correlation of logs | 0.9880 |
| HExp QLIKE advantage, leave-one-day-out[6] | +0.96749 |
| 15-minute signature level | 20 to 41 % |
| Huber signature, 1 min over 1 s, 12 day average | 1.40 |

IPAR over the 12 trading days 1 to 16 October 2024, and on 16 October 2024. [1]the median over $k$ from 2 to 100 of the half-life minimizing the $k$-trade-ahead Huber forecast loss, and the largest cost over those $k$ of running the common 3 trades. [2]the 1-minute-ahead check loss of the noise scale. [3]the half-life maximizing the mean predictive log-likelihood per scored trade on 16 October 2024, with the range within $10^{-4}$ of the maximum. [4]the largest Huber residual after 09:45, at 14:58; before is the median over $t^* - 15$ to $t^* - 3$ minutes and at is the maximum over $t^* - 6$ to $t^* + 12$ minutes. [5]over the 12 days, sampled every 300 trades. [6]hyperbolic minus HExp, so positive favours HExp.

| | Mean of $\lvert\hat{X}_t - \hat{X}_{t-k}\rvert/\sqrt{k}$ in cents | | | | Sqrt of Mean of $(\hat{X}_t - \hat{X}_{t-k})^2/k$ in cents | | | |
|---|---|---|---|---|---|---|---|---|
| $k$ | median | Huber | linear | trades | median | Huber | linear | trades |
| 1 | 1.316 | 1.027 | 0.995 | 4.363 | 2.855 | 1.579 | 1.519 | 8.291 |
| 2 | 1.284 | 1.144 | 1.131 | 3.633 | 2.580 | 1.725 | 1.663 | 6.365 |
| 5 | 1.528 | 1.316 | 1.302 | 2.848 | 2.528 | 1.928 | 1.870 | 4.480 |
| 10 | 1.635 | 1.418 | 1.405 | 2.360 | 2.428 | 2.025 | 1.988 | 3.496 |
| 25 | 1.611 | 1.459 | 1.439 | 1.907 | 2.255 | 2.048 | 2.015 | 2.715 |
| | ACF of $\hat{X}_t - \hat{X}_{t-1}$ at lag $k$ | | | | Hill tail index of $\lvert\hat{X}_t - \hat{X}_{t-k}\rvert$ | | | |
| 1 | −0.18 | 0.19 | 0.20 | −0.41 | 4.38 | 5.19 | 5.05 | 3.74 |
| 2 | 0.10 | 0.10 | 0.11 | −0.04 | 5.04 | 3.59 | 5.23 | 4.26 |
| 5 | −0.03 | 0.00 | 0.02 | −0.02 | 5.07 | 5.07 | 5.88 | 5.37 |
| 10 | 0.02 | 0.01 | 0.00 | 0.00 | 5.04 | 6.29 | 5.86 | 5.45 |
| 25 | 0.04 | −0.01 | 0.00 | −0.01 | 4.95 | 4.74 | 4.53 | 5.37 |

Table 1: Trade-lag differences, `IPAR`, 16 October 2024.

| | Mean of $\frac{\lvert\hat{X}_t - \hat{X}_{t,s}\rvert}{\sqrt{\tau_t - \tau_{t,s}}}$ in cents | | | | Annualized volatility: $\hat{\sigma}_{\text{Annual};s}$ | | | |
|---|---|---|---|---|---|---|---|---|
| $s$ | median | Huber | linear | trades | median | Huber | linear | trades |
| 0.1 s | 2.47 | 2.16 | 2.11 | 4.45 | 18.6 | 11.4 | 11.4 | 66.2 |
| 1 s | 2.00 | 1.75 | 1.73 | 3.23 | 17.6 | 11.1 | 11.2 | 62.7 |
| 5 s | 1.57 | 1.40 | 1.38 | 2.51 | 17.5 | 12.3 | 12.1 | 56.2 |
| 10 s | 1.41 | 1.28 | 1.27 | 2.16 | 17.8 | 13.8 | 13.6 | 50.6 |
| 1 min | 1.03 | 0.93 | 0.92 | 1.38 | 18.9 | 16.9 | 16.6 | 36.3 |
| 15 min | 0.89 | 0.87 | 0.87 | 0.91 | 20.2 | 19.7 | 19.5 | 22.0 |
| | $\text{Cor}(\hat{X}_t - \hat{X}_{t,s}, \hat{X}_{t,s} - \hat{X}_{t,2s})$ | | | | Hill tail index of $\lvert\hat{X}_t - \hat{X}_{t,s}\rvert$ | | | |
| 0.1 s | −0.04 | −0.02 | −0.01 | −0.10 | 7.58 | 5.33 | 4.64 | 8.09 |
| 1 s | −0.02 | 0.06 | 0.06 | −0.09 | 7.90 | 5.55 | 4.89 | 8.20 |
| 5 s | −0.07 | 0.03 | 0.02 | −0.23 | 7.14 | 10.49 | 9.53 | 10.19 |
| 10 s | −0.03 | 0.04 | 0.04 | −0.18 | 7.81 | 13.97 | 10.80 | 5.39 |
| 1 min | −0.19 | −0.15 | −0.15 | −0.39 | | | | |

Table 2: Clock-lag differences, `IPAR`, 16 October 2024.

| | median | Huber | linear |
|---|---|---|---|
| Mean of $\lvert\hat{\epsilon}_t\rvert$ in cents | 3.04 | 3.68 | 3.80 |
| s.d. of $\hat{\epsilon}_t$ in cents | 6.14 | 5.80 | 5.77 |
| Tail index of $\psi(\hat{\epsilon}_t)$ | $\infty$ | $\infty$ | 5.02 |
| ACF of $\psi(\hat{\epsilon}_t)$, lag 1 | 0.21 | 0.24 | 0.21 |
| ACF of $\psi(\hat{\epsilon}_t)$, lag 2 | 0.09 | 0.13 | 0.10 |
| ACF of $\psi(\hat{\epsilon}_t)$, lag 5 | 0.01 | 0.02 | 0.01 |
| ACF of $\psi(\hat{\epsilon}_t)$, lag 10 | −0.02 | 0.00 | 0.01 |
| ACF of $\psi(\hat{\epsilon}_t)$, lag 25 | 0.00 | 0.00 | 0.00 |
| 80% C.I. for $\theta_t - X_t$ | — | −2.64,2.42 | −2.73,2.57 |
| 95% C.I. for $\theta_t - X_t$ | — | −5.00,4.57 | −5.31,4.80 |
| 99% C.I. for $\theta_t - X_t$ | — | −8.87,7.93 | −8.84,8.50 |
| s.e. of $\theta_t - X_t$ in cents | — | 2.29 | 2.47 |

Table 3: Residual summaries, `IPAR`, 16 October 2024. Median C.I. and s.e. cells are omitted (price discreteness); the linear s.e. is printed because the Hill index of its score, 5.02, is above two.

| | 16 October 2024 | | 12 days, min to max | |
|---|---|---|---|---|
| | upper | lower | upper | lower |
| All trades (as in the paper) | 3.40 | 3.14 | 1.71 to 3.65 | 1.93 to 3.76 |
| Residual divided by $\delta_t$ | 2.34 | 3.11 | 2.07 to 4.12 | 2.78 to 4.32 |
| Repricing windows removed | 3.40 | 3.14 | 1.72 to 3.65 | 1.93 to 3.99 |
| Calm windows only | 3.39 | 2.95 | | |
| Every 10th trade | — | — | — to — | — to — |
| Six-trade half-life | 3.58 | 3.42 | 1.59 to 4.00 | 2.15 to 4.67 |

Table 4: Hill estimator of the tail index of the Huber filter residual, `IPAR`, 12 analysis days. The index is the mean of Hill estimates over the largest 0.1% to 1% of residuals, each tail separately, which for this stock is an average over 1 to 5 depths a day in row 1; the masked and thinned rows use fewer. Second row: divides the residual by its own threshold $\delta_t$. Third row: drops trades where the filtered price moves more than 6.8 times the day's median $\delta_t$ within $\pm 1$ second, which is 60.4 to 100.6 cents over the days and removes 0% to 2% of a day; the fourth keeps only trades where it moves less than 3.4 times that median within $\pm 5$ seconds, 30.3 cents on 16 October 2024, 96% of that day. Fifth row: keeps every 10th trade. Sixth row: doubles the filter's half-life with the thresholds held fixed. The ranges run over the days on which the row can be estimated.

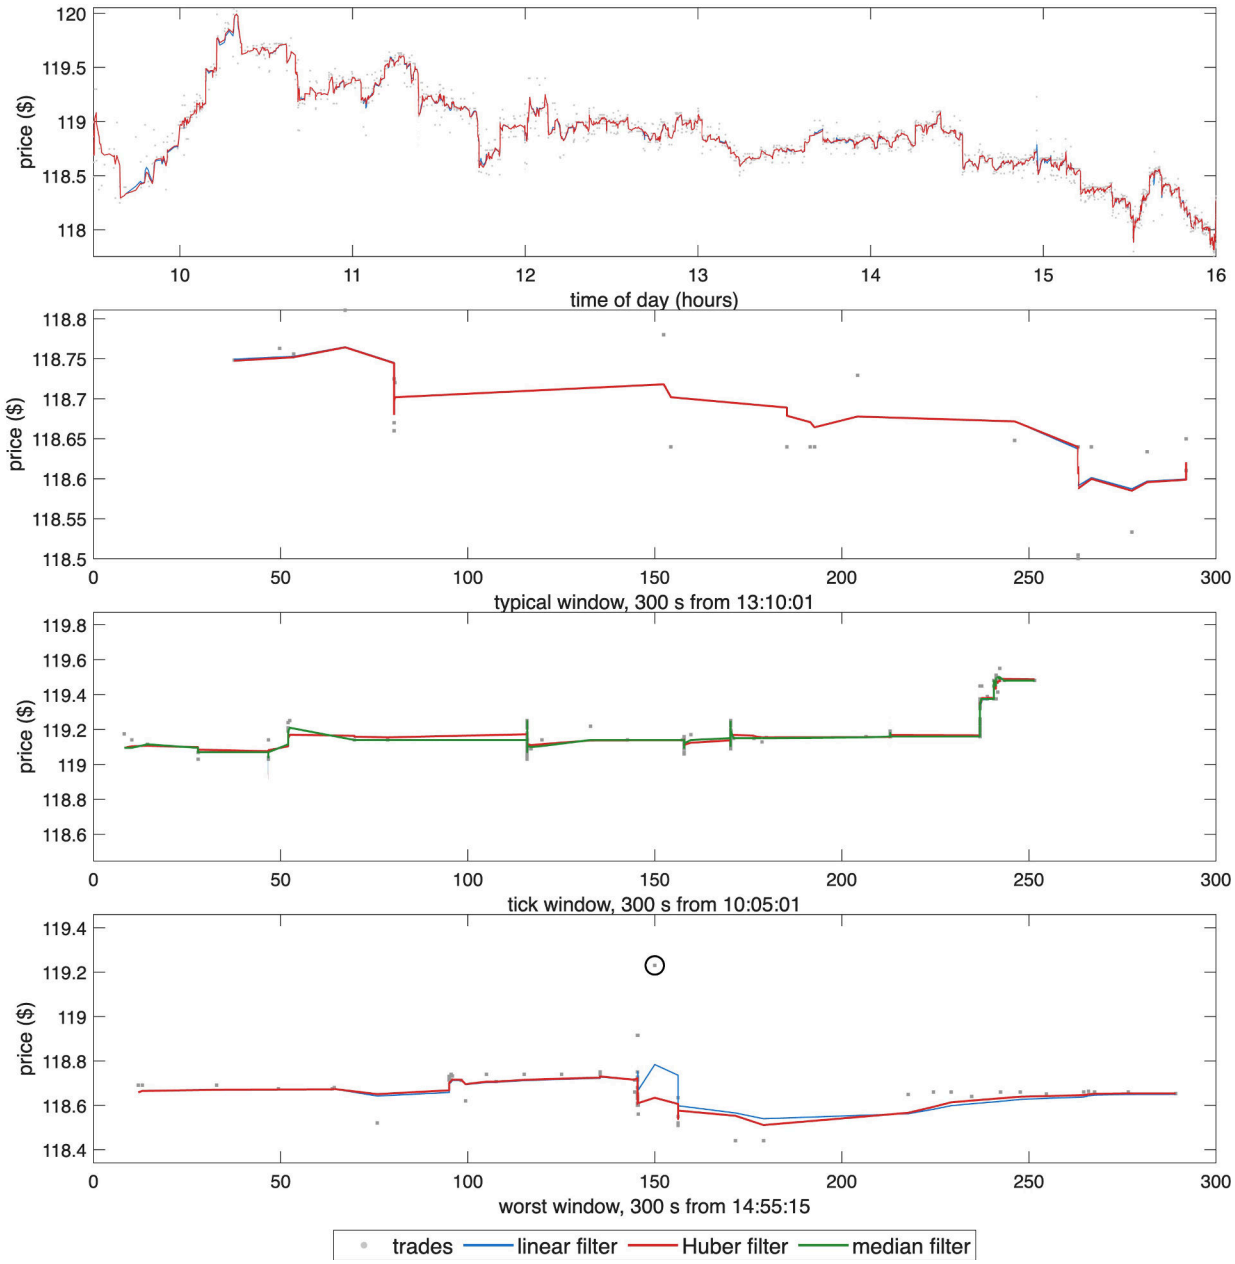


Figure 1: `IPAR` trades (grey points), the linear filter and the Huber filter, 16 October 2024. Top: full trading day. 2nd: 300-second window at the day's median trade count, from 13:10:01. 3rd: 300-second window at 10:05:01 in which the price reverses direction most often, with the median filter $m_t$ added. Bottom: 300-second window around the largest Huber residual after 09:45, circled.

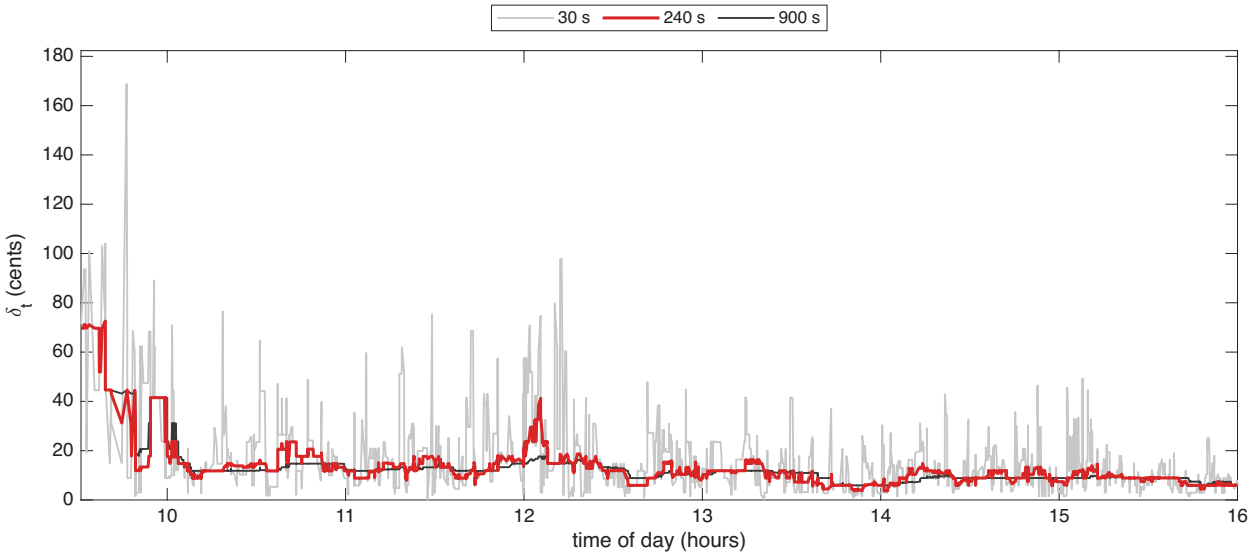


Figure 2: The Huber threshold $\delta_t$ in cents for `IPAR` on 16 October 2024, with half-lives of 30, 240, 900 seconds. The 240-second half-life, in red, is used.

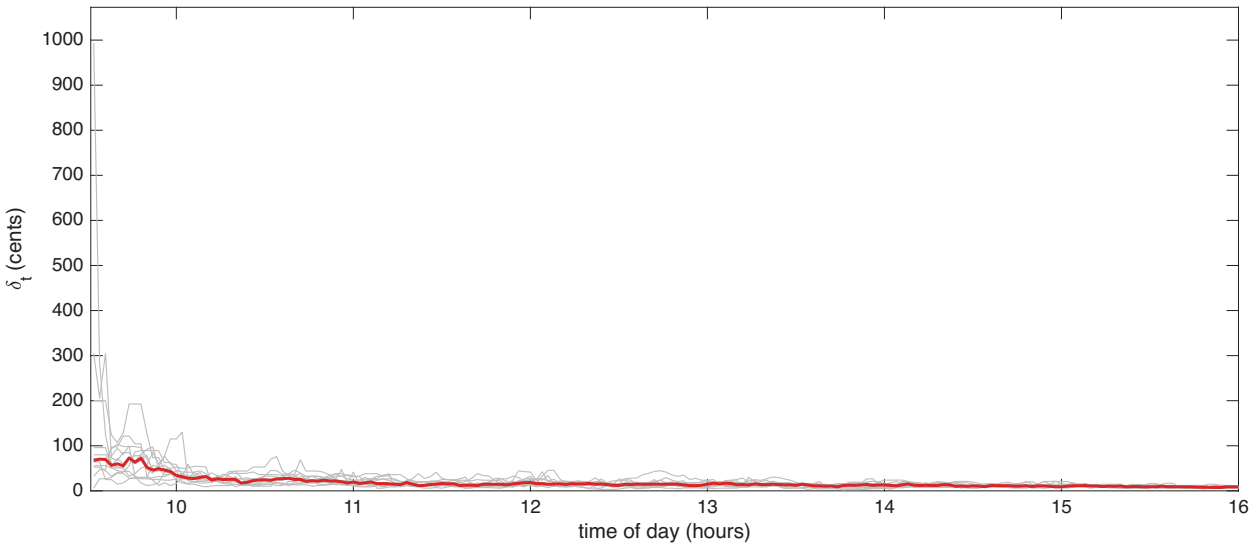


Figure 3: Huber threshold $\delta_t$ in cents through the trading day, for 12 `IPAR` trading days in Oct 2024 (grey lines) and their median (red line).

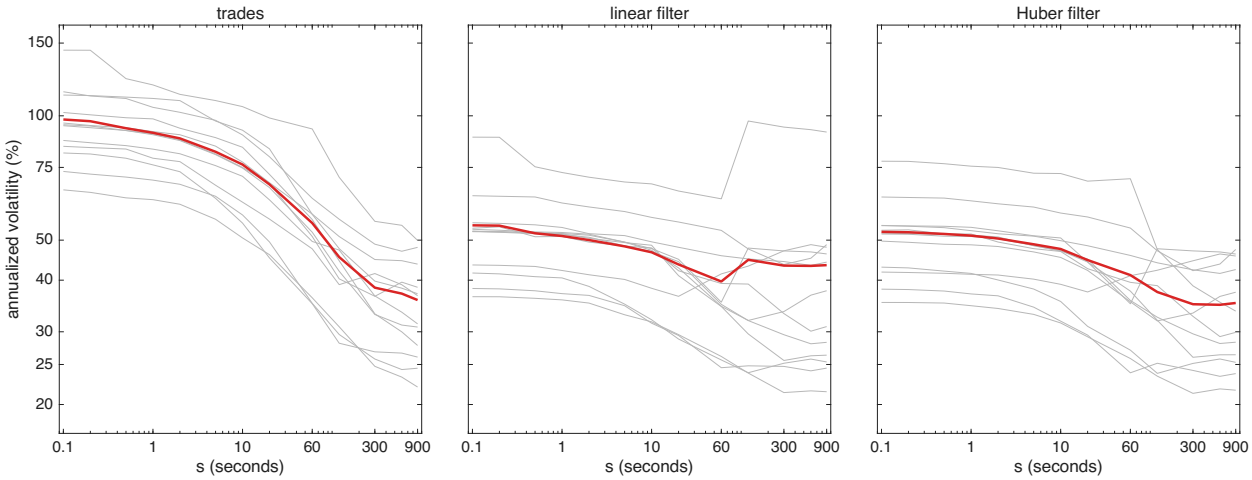


Figure 4: Volatility signature curves for `IPAR` over the 12 trading days in Oct 2024. The annualized vol $\hat{\sigma}100\sqrt{252}$ from the returns $\hat{X}_t - \hat{X}_{t,s}$ of Table 2, against the lag $s$ in seconds, both axes on log scales, for: raw trades, linear filter & Huber filter. Grey lines are the individual trading days & the red line is the square root of the 12 day average of their squares. The linear & Huber returns carry the finite-sample correction $c_{t,s}$ of the paper; the raw trades are unaffected by it.

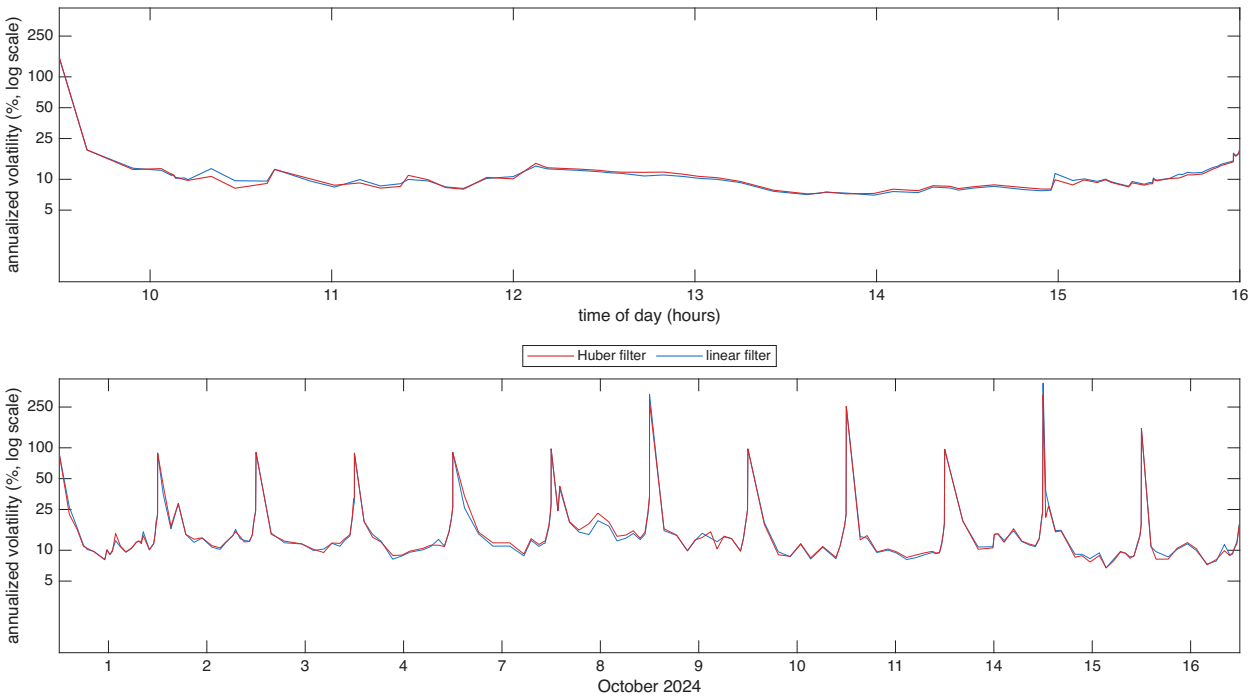


Figure 5: Filtered spot volatility for `IPAR` over 12 trading days in Oct 2024, the diurnally adjusted estimator $\hat{\sigma}(u)100\sqrt{252}$ at $s = 1$ second with hyperbolic weights, Mandelbrot parameter $a = 26$, with the input prices computed using the Huber & linear filters, shown as annualized volatility in percent on a log10 scale. Top: 16 October 2024, bottom: 12 days.

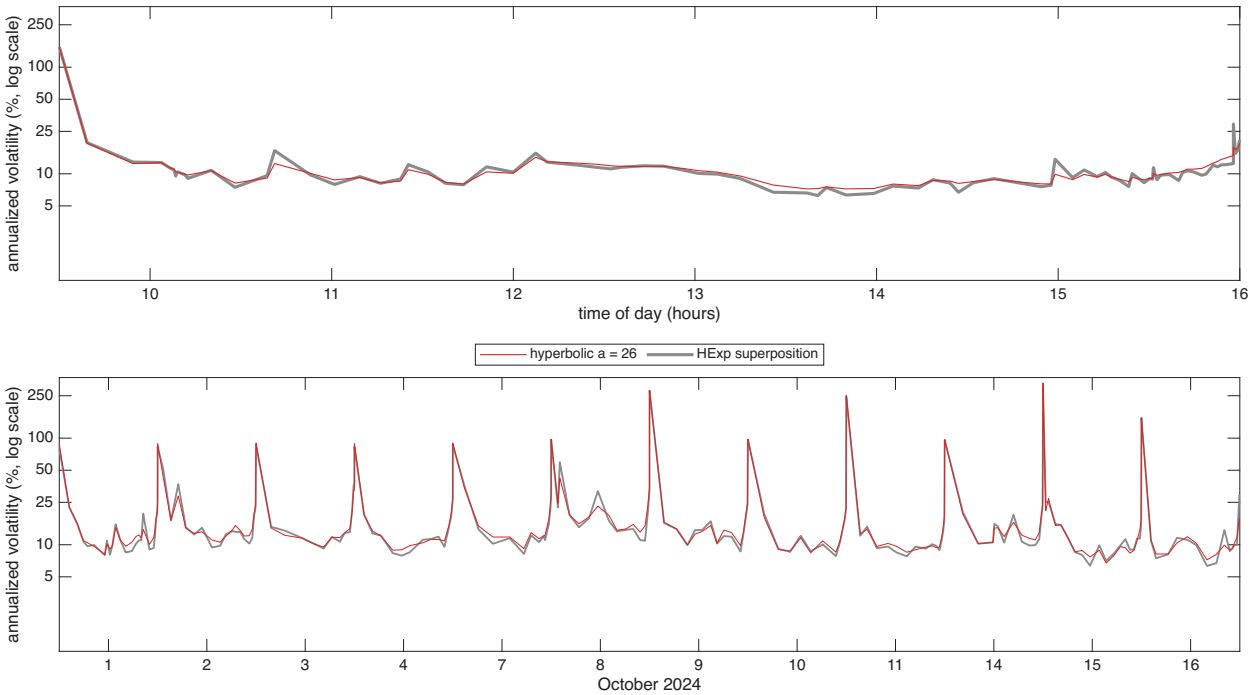


Figure 6: Hyperbolic against superposition weights for `IPAR` over 12 trading days in Oct 2024, the diurnally adjusted $\hat{\sigma}(u)100\sqrt{252}$ at $s = 1$ second from the Huber filter, shown as annualized vol in % on a log10 scale. The only difference between the 2 paths is the weight function: the hyperbolic weights with $a = 26$ in red, drawn on top, and the `HExp` superposition in grey, 4 exponential components with half lives of 10, 100, 1000, 10000 seconds. Hyperparameters minimize the QLIKE objective. Top: 16 October 2024, bottom: 12 days.

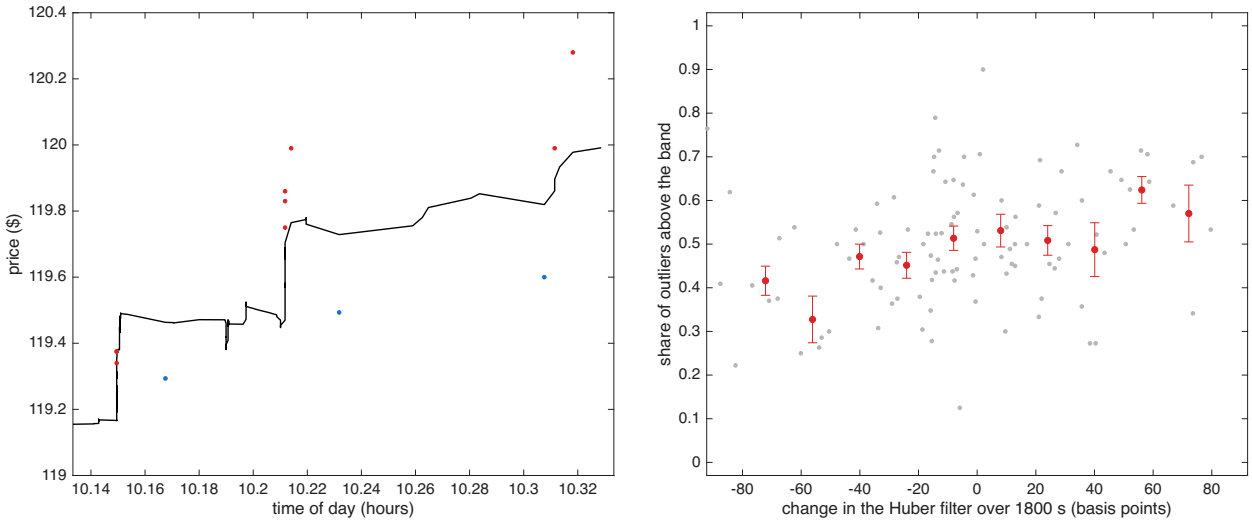


Figure 7: Outliers & robust price for `IPAR` on 16 October 2024. Left: the 12 minutes with the largest move in the Huber filter that day; outliers above $Y_t - \theta_t > \delta_{t-1}$ in red, outliers below $Y_t - \theta_t < -\delta_{t-1}$ in blue, Huber filter in black. Right: over the 12 trading days, the 1800 second empirical conditional probability that outliers are above (rather than below) the Huber band against the 1800 second change in the Huber filter in basis points: light points are single bins, the line & bars are means & s.e.

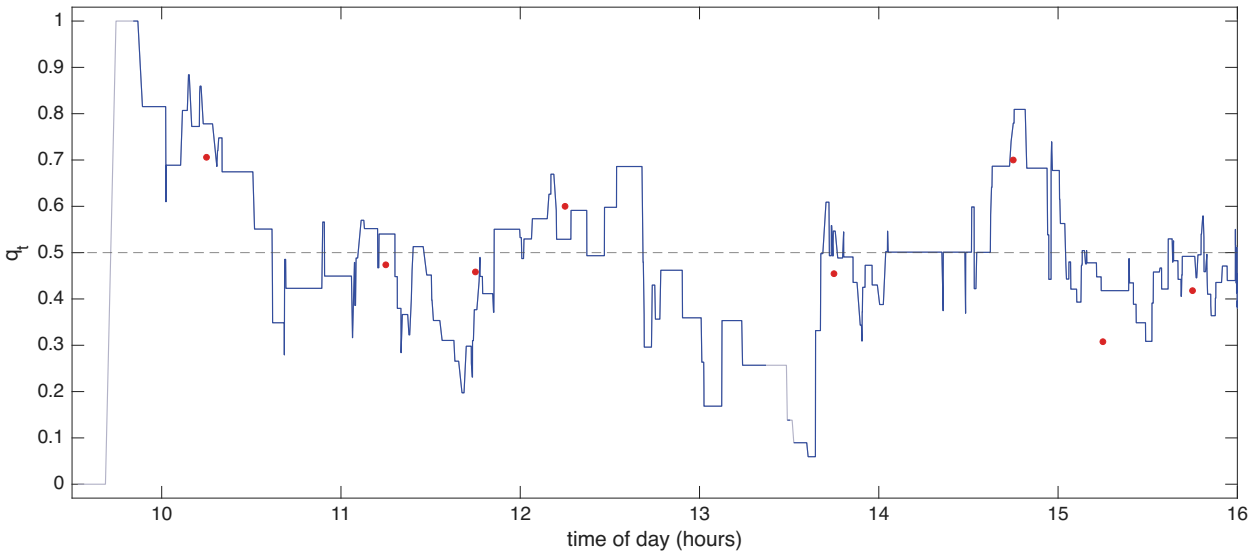


Figure 8: The directional conditional probability for `IPAR`, 16 October 2024. The filtered conditional probability $q_t$ that an outlier lies above the robust Huber band, under exponential weights decaying in clock time with half life of 545 seconds; where fewer than two outliers remain in the weights it is drawn in grey. The dashed line marks 0.5. The red points are the empirical conditional probability using the 1800 second bins of the previous figure.

14 of the 22 days built (10 presample and 12 analysis) have fewer than 5,000 trades.
The direction half-life grid was extended beyond its usual end to 1500 seconds, because the log-likelihood was still rising there.
Selecting a on single days gives values from 4 to 26112 trades, against 26 trades pooled over the 12 days; the 10-second estimate of the spot volatility behind the QLIKE objective averages roughly 2 trades for this stock.
The Flexible Fourier fit of the diurnal profile has R-squared below 0.2 on 1 day(s) (lowest 0.18).
On 12 analysis day(s) the opening noise scale rests on fewer than 100 nonzero residuals from the previous openings (fewest 25).
The spot volatility filter has no return to work with for the first 25 seconds of the slowest opening, because the first continuous trade comes late and its reference must be 1 second old.
On 2 October 2024, 37 per cent of the day's largest one per cent of residuals fall inside its first minute, against at most 29 per cent on any other day. Every number that day contributes to carries that burst.
The tail index of the residual cannot be estimated for every 10th trade, on none of the 12 days: what the row's mask or thinning leaves is too few residuals for the largest 0.1 to 1 per cent to hold enough order statistics.
Where the tail index can be estimated it averages over as few as 1 depth on one day, because the largest 0.1 to 1 per cent of 3,152 residuals is a short stretch of the tail. It is a much less precise estimate than for a heavily traded name.

Diagnostics.