Zeros

DOI:

[10.1287/mnsc.2019.3527](https://doi.org/10.1287/mnsc.2019.3527)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Kolokolov, A., Bandi, F., Renò, R., & Pirino, D. (2020). Zeros. *MANAGEMENT SCIENCE*, 66(8), 3466-3479.
<https://doi.org/10.1287/mnsc.2019.3527>

Published in:

MANAGEMENT SCIENCE

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Zeros*

Federico M. Bandi[†] Aleksey Kolokolov[‡] Davide Pirino[§] Roberto Renò[¶]

December 10, 2019

Abstract

Asset prices can be stale. We define price “staleness” as lack of price adjustments yielding zero returns (i.e., zeros). The term “idleness” (resp. “near idleness”) is, instead, used to define staleness when trading activity is absent (resp. close to absent). Using statistical and pricing metrics, we show that zeros are a genuine economic phenomenon linked to the dynamics of trading volume and, therefore, liquidity. Zeros are, in general, not the result of institutional features, like price discreteness. In essence, spells of idleness or near idleness are stylized facts suggestive of a key, omitted market friction in the modeling of asset prices. We illustrate how accounting for this friction may generate sizable risk compensations in short-dated option returns.

JEL classification: G1, G12, G13, C58

Keywords: *Volume, liquidity, short-term options*

*We thank the Editor, Kay Giesecke, an anonymous Associate Editor, two anonymous reviewers and conference participants in the Third International Workshop in Financial Econometrics, Arraial d’Ajuda Ecoresort in Porto Seguro, Brazil, October 8-10 2017, the Quantitative Finance and Financial Econometrics Conference, Marseille, May 30-June 1 2018, the 11th Society for Financial Econometrics Conference in Lugano, June 11-14 2018, the 10th Bachelier World Congress, Dublin, July 16-20 2018, the Conference “High-Voltage Econometrics”, Palermo, October 4-5 2018, and the Conference “Financial Econometrics and New Finance”, Hangzhou, June 10 2019, for discussions. We are particularly indebted to the discussant, Jantje Soenksen, for her comments and suggestions. Her communications with us have lead to a sharper treatment. Davide Pirino acknowledges partial support by the grant RBSI14DDNN, “A new measure of liquidity”, financed within the program “Scientific Independence of Young Researchers” of the Italian Ministry of Education and Research.

[†]Johns Hopkins University, Carey Business School, 100 International Drive, Baltimore MD 20202, USA and Edhec-Risk Institute, France. E-mail: fbandi1@jhu.edu.

[‡]Alliance Manchester Business School, Booth St W, Manchester M15 6PB, UK. E-mail: alexeiuo@gmail.com.

[§]Dipartimento di Economia e Finanza, Università di Roma Tor Vergata, Via Columbia 2, 00133 Roma, Italy and Scuola Normale Superiore, Pisa, Italy. E-mail: davide.pirino@gmail.com.

[¶]Università di Verona, Via Cantarane 24, 37129 Verona, Italy. E-mail: reno@univr.it.

1 Introduction

The ubiquitous semimartingale model for asset prices in continuous time does not appear to be rich enough. Figure 1 displays examples of intra-day price paths, sampled at 30-second intervals, for stocks traded on the New York Stock Exchange (NYSE).¹

Panel A shows the price dynamics of a liquid stock (Exxon Mobil, XOM) with large volume. The graph is visually compatible with the erratic behavior of a semimartingale in continuous time, one which is locally driven by Brownian motion (and, infrequently, by jump-like discontinuities). *Staleness*, defined as the frequency of zero returns, is only 2.1%. *Idleness*, defined as the frequency of time intervals without trading volume, is 0%.² Panel A illustrates an *ideal* situation, supporting semimartingale modeling, one which is – however – far from typical. Specifically, it displays the day for which XOM has the *lowest* level of staleness in the 9-year sample we analyze. For reference, the mean and standard deviation of XOM staleness in our sample are 19.75% and 7.9%, respectively.

Panel B represents a more typical situation for another very liquid stock (Bank of America, BAC). The graph now looks rather different. The price process is considerably stickier. The number associated with idleness raises the question of why market participants were, for nearly one third of the day, electing not to trade this liquid stock. The difference between staleness and idleness (42%) is suggestive of the fact that one can often observe zeros when volume is low, but non-zero. In Panel B, in fact, we observe relatively larger volume being associated with price changes and relatively smaller volume being associated with zeros. What could lead to zeros in the presence of small, but non-zero, volume? We identify and discuss two channels: easy absorption of limited volume without price impacts (near idleness) and/or price discreteness. While the first channel is economically relevant, the second one *may*, in principle, not be. At the NYSE, since 2001, traded prices change by multiples of 1 cent. Zeros observed in conjunction with non-zero trading volume may, therefore, be due to this institutional feature. *Excess staleness*, formally defined below, measures the percentage of zero returns not due to the second channel and can, therefore, be viewed as overall staleness purged of price discreteness. In this case, excess staleness is estimated to be 35.3%, roughly half of total staleness. For XOM, in Panel A, it is 0.6%.

The impact of price discreteness is clearly visible in Panel C. The dynamics of Citigroup (C) are affected by sizable staleness but almost no idleness. Excess staleness, in this case, is estimated to be “only” 15.6%. We note that the price range of 8 cents for that day is small as compared to the minimum allowed price change of 1 cent. In effect, a large number of zeros (leading to limited price variability over the day) occur in the presence of non-zero volume. While price discreteness is the reason for this finding, should one view the effect as being economically uninformative? We argue that price discreteness in the presence of limited volume (like in Panel C) cannot be regarded as leading

¹Our data consists of all trades of 249 NYSE-listed stocks, recorded from 9:30 a.m. to 4 p.m., for the years 2006-2014. We employ the 249 stocks with the largest average traded volume during the period. Given our focus on the NYSE-listed stocks with the largest traded volume, i.e., those whose prices are expected to be the least inactive, our findings will be viewed as being conservative.

²For each stock k (with $k = 1, \dots, N$) and for each day t_k (with $t_k = 1, \dots, T_k$), we sample prices $\bar{P}_{0,t_k}, \dots, \bar{P}_{n,t_k}$ on an evenly spaced grid and compute intra-day returns as $r_{i,t_k} = \ln \bar{P}_{i,t_k} - \ln \bar{P}_{i-1,t_k}$, where $i = 1, \dots, n$. A return is *stale* if $r_{i,t_k} = 0$. It is *idle* if $V_{i,t_k} = 0$, where V_{i,t_k} is the traded dollar volume over the time interval $[(i-1)/n, i/n]$. Of course, idleness is lower than staleness by construction, since absence of trading implies staleness.

Figure 1: Six examples of intra-day stock price dynamics on a 30-second sampling grid. In each panel, the circles represent dollar volume traded in each 30-second interval. The sizes of the circles are proportional to the square root of dollar volume, as indicated in Panel A. Absence of volume in a 30-second interval is signaled by a red cross. Values that prices can take (as multiples of 1 cent) are represented on the horizontal grid. Price data are recorded from 9:30 to 16:00 for the years 2006-2014.

to “spurious” (from an economic standpoint) zeros. It only translates into rounded up prices which may have moved somewhat (but by less than a cent) if they had been allowed to take on a continuum of values. In Panel C, when traded volume increases just before 4 p.m., the number of zeros declines,

due to bid/ask bounce effects and prices reverting on the discretization grid. In essence, then, for price discreteness to represent an institutional feature leading to spurious zeros uninformative about the nature of the trading mechanism, it would need to also be associated with relatively large volume. If this were the case, then zeros would not just be the result of low volume and limited liquidity. They would also be the result of high liquidity, and easily-absorbed high volume, yielding low price impacts and (rounded up, due to price discreteness) stale traded prices. The latter scenario, however, is a rare occurrence in the data.

Panels D and E feature two less liquid stocks (VeriFone Systems, PAY, and Kansas City Southern, KSU) for which the intra-day price range is large with respect to 1 cent. Idleness is now strongly in the data. In the case of PAY, excess staleness is more than twice as large as idleness. As for KSU, virtually all zeros are associated with zero volume. In both cases, very few zeros are estimated to be due to price discreteness. Fluor Corporation (FLR, in Panel F) is extremely stale during the first part of the day and rather erratic (in association with larger volume) during the second part of the day. We note that a jump around 2 p.m. is associated with an increase in trading activity.

To summarize, Figure 1 shows that idleness and near idleness should be an integral part of a realistic data generating process for asset prices in continuous time *before* institutional features (like rounding) are accounted for. The figure raises questions. How do idleness, staleness and excess staleness change dynamically, intra-daily for example? What is the impact of sampling frequency on them? What are their economic determinants and what can be said about economic implications? This work aims at providing some answers.

A successful body of work in empirical finance has offered evidence (using alternative samples and in different markets) for the pervasiveness of zero *daily* returns and their usefulness in obtaining liquidity proxies (see Lesmond et al., 1999, Lesmond, 2005 and the rich set of comparisons in Goyenko et al., 2009). We depart from this stimulating line of work along two main dimensions. First, our emphasis is on the evaluation of staleness at alternative, *intra-daily*, frequencies for a rich cross-section of US stocks. We expect staleness to increase with the sampling frequency (Figure 3, below, supports this intuition). We show that (1) the economics of staleness applies at all scales and is pervasive across stocks (even the most liquid) and (2) the impact of price discreteness (which we confront directly in this paper at all frequencies, the highest in particular) is marginal enough as not to affect economic logic. In this sense, this work provides cross-sectional support for stylized empirical features of the observed price process at different scales. Second, by focusing on intra-daily frequencies, we are able to go “to the limit”, zoom into the granular properties of the observed price dynamics, and illustrate the role which staleness should play in the modeling of asset prices in continuous time. As is well-known, continuous-time modeling through semimartingale price processes driven by Brownian shocks (and, often, jumps) is a fundamental building block of quantitative finance (c.f., Duffie, 2008, Back, 2010, Jarrow, 2018, and the references therein). We provide evidence for an alternative price process (formalized in Eq. (4) of Section 2), one which incorporates aspects of the micro-structure of high-frequency trading and is such that the latent semimartingale (equilibrium) price process is revealed only by trades (possibly up to other market microstructure contaminations, like bid-ask noise). The proposed, alternative, data-generating process has important applied implications. We mention one here, but we return to directions for future work below. A successful recent literature

has emphasized the role of discontinuities in asset prices at high frequencies and their impact on risk assessments, risk management and the pricing of structured products, among other issues (see, e.g., Bajgrowicz et al., 2015). Our proposed data generating process allows for jump-like behavior even in the absence of jumps in the unobservable price process. In our framework, the jumps would be induced by periods of staleness followed by trades at the latent equilibrium price. Not only is the economics of this spurious jump-like behavior different from that of conventional jumps, we expect spurious jump-like behavior to affect existing risk management practices in important ways.

We also depart from the work of Bandi et al. (2017). Their focus is on the theory of a kernel-based, high-frequency, measure of price staleness and its application to an index. The analysis in this paper is cross-sectional, rather than index-based, and centered around the link between (notions of) staleness, volume and liquidity, as proxied by the magnitude of execution costs. In this sense, it can be viewed as providing empirical and economic foundations for the work of Bandi et al. (2017). In addition, while Bandi et al. (2017) handle price rounding by modifying the sample frequency, rounding is confronted directly in this work (in, e.g., Section 2). The result is a novel notion of *excess staleness* and a nomenclature which we trust to be revealing of the economic determinants of price staleness in its various forms.

The paper proceeds as follows. Section 2 provides an econometric methodology to identify excess staleness. The section introduces a data-generating process, for which we offer rich cross-sectional evidence, allowing for repeated prices due to zero, or limited, trading. In Section 3, we put forward an economic rationale for price staleness which we support by studying the empirical relation between excess staleness, volume, execution costs and volatility. Section 4 discusses some of the economic implications of the proposed data-generating process in the context of a timely economic metric, i.e., the pricing of short-dated options. Conclusions and directions for future research are contained in Section 5. Details on derivations are relegated to the Appendix.

2 Staleness and *excess* staleness

Idleness and near idleness are stylized empirical features of recorded asset prices. Modeling prices as continuous semimartingales (before rounding) will, therefore, not capture effects which are strongly in the data. We adopt an alternative data generating process.

We still assume that the *latent* fundamental, or efficient, price (P) follows the local continuous martingale

$$P_t = \int_0^t \sigma_s P_s dW_s, \quad (1)$$

where W_t is Brownian motion. Eq. (1) implies that the logarithmic efficient price process is *locally* Gaussian with instantaneous stochastic volatility σ_t .³

Next, we partition the trading period $[0, T]$, a day (say), into n intervals of length Δ_n . Naturally, Δ_n denotes sampling frequency. Idleness can now be modeled, as in Bandi et al. (2017), using a triangular array of (possibly dependent) Bernoulli variates $\{\mathbb{B}_{i,n}\}_{i=1,\dots,n}$ satisfying

³We are neglecting the drift term (which is reasonable, over small intervals, if the drift is uniformly bounded) and price discontinuities (which is also reasonable, if these discontinuities are rare).

$$\mathbb{P}[\mathbb{B}_{i,n} = 1] = \mathbb{E}[\mathbb{B}_{i,n}] = p_n, \quad (2)$$

for all $i = 1, \dots, n$, and, as $n \rightarrow \infty$,

$$\frac{1}{T} \sum_{i=1}^n \Delta_n \mathbb{B}_{i,n} \xrightarrow{p} p_\infty. \quad (3)$$

In Eq. (2) and in Eq. (3), p_n and p_∞ are the *frequency-specific* probability of repeated prices and the probability of repeated prices at the highest frequency, respectively.

The *observed* price process (before rounding) on the sampling partition follows⁴

$$\tilde{P}_{i\Delta_n} = P_{i\Delta_n}(1 - \mathbb{B}_{i,n}) + \tilde{P}_{(i-1)\Delta_n} \mathbb{B}_{i,n}, \quad (4)$$

with $\tilde{P}_0 = P_0$. Thus, the observed price is the fundamental price if $\mathbb{B}_{i,n} = 0$ (i.e., in the presence of trading). Absent trading, i.e., if $\mathbb{B}_{i,n} = 1$, the previous price is repeated and prices are idle.⁵ Of course, a more nuanced interpretation consistent with empirical evidence would set $\mathbb{B}_{i,n} = 1$ not only when trading is absent (idleness) but, also, when trading is slow (near idleness). We abide by this interpretation in what follows.

A key implication of the model is that the latent price continues to be diffusive *even* in the presence of idleness. A consequence of this implication, and an interesting prediction of the model, is that the variance of intra-day returns should be larger after idleness than after non-zero returns. In fact, the longer prices have been idle, the larger the variance should be. Figure 2 verifies this prediction with data. For each stock and each day, we compute the relative variance of the intra-day 30-second returns with respect to the overall variance for the day. We then calculate averages conditional on the number of either stale (Panel A) or idle (Panel B) returns observed before the considered return. As can be seen in the figure, the unconditional intra-day (relative) variance is also larger than that associated with returns not preceded by staleness (or idleness).

A portion of staleness is due to price discreteness. We model discreteness by rounding the price $\tilde{P}_{i\Delta_n}$ to the closest price $\bar{P}_{i\Delta_n}$. Observed prices are multiples of a fixed quantity d (e.g., for NYSE-listed stocks d is 1 cent). As is well-known, allowing for rounding when modeling observed prices (given a continuous-time model for latent fundamentals) is a challenging issue (e.g., Delattre and Jacod, 1997). In fact, the probability of a zero return, conditional on the past history of prices, is a non-Markovian problem, even under simple assumptions on the price dynamics.

Having made this point, we show that rounding can explain only a (small) portion of observed zeros. To this extent, we provide an expression for the probability of observing a k -tick movement of the price process for any integer k . For the time being, we de-activate the Bernoulli variates and set them equal to zero. The probability that $\bar{P}_{(i+1)\Delta_n} - \bar{P}_{i\Delta_n} = kd$, conditional on spot volatility

⁴The pricing model in Phillips and Yu (2007) can be written similarly but with independent Bernoullis and a constant, rather than frequency-specific, probability of repeated prices.

⁵One could argue that, in the presence of trading, fundamental values would be revealed to uninformed, or partially-informed, traders since, given Eq. (4), traded prices coincide with fundamental values (ignoring rounding). An easy way to prevent full observability of fundamental values when trading takes place is to assume $\tilde{P}_{i\Delta_n} = P_{i\Delta_n}^*(1 - \mathbb{B}_{i,n}) + \tilde{P}_{(i-1)\Delta_n} \mathbb{B}_{i,n}$, where $P^* = P + \epsilon$, with ϵ representing a market microstructure contamination induced by, e.g., bid/ask bounce effects and order-flow clustering.

Figure 2: We report the relative (with respect to daily variance) variance of intra-day 30-second returns. We do so unconditionally and after conditioning on different numbers of consecutive stale returns (Panel A) and idle returns (Panel B). The relative variance is computed as squared 30-second returns divided by the average of squared 30-second returns for that day and stock. The plots are cross-sectional and daily averages based on 249 NYSE-listed stocks, whose trades are recorded from 9:30 to 16:00, for the years 2006-2014.

$\sigma_{i\Delta_n}$, the rounded price $\bar{P}_{i\Delta_n}$ and the displacement value $x_{i\Delta_n}$, where $\bar{P}_{i\Delta_n} = P_{i\Delta_n} - x_{i\Delta_n}$ with $-d/2 < x_{i\Delta_n} < d/2$, is⁶

$$\begin{aligned}
& \mathbb{P} \left[\bar{P}_{i\Delta_n} + kd - \frac{d}{2} < P_{(i+1)\Delta_n} < \bar{P}_{i\Delta_n} + kd + \frac{d}{2} \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n}, x_{i\Delta_n} \right] \\
&= \mathbb{P} \left[P_{i\Delta_n} + kd - \frac{d}{2} - x_{i\Delta_n} < P_{(i+1)\Delta_n} < P_{i\Delta_n} + kd + \frac{d}{2} - x_{i\Delta_n} \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n}, x_{i\Delta_n} \right] \\
&= \int_{kd - \frac{d}{2} - x_{i\Delta_n}}^{kd + \frac{d}{2} - x_{i\Delta_n}} \frac{e^{-\frac{z^2}{2\sigma_{i\Delta_n}^2(\bar{P}_{i\Delta_n} + x_{i\Delta_n})^2\Delta_n}}}{\sqrt{2\pi\sigma_{i\Delta_n}^2(\bar{P}_{i\Delta_n} + x_{i\Delta_n})^2\Delta_n}} dz \\
&= \int_{kd - \frac{d}{2} - x_{i\Delta_n}}^{kd + \frac{d}{2} - x_{i\Delta_n}} \frac{e^{-\frac{z^2\zeta_{i\Delta_n}^2}{2d^2}}}{\sqrt{2\pi\frac{d^2}{\zeta_{i\Delta_n}^2}}} dz, \tag{5}
\end{aligned}$$

where $\zeta_{i\Delta_n} = d/(\sigma_{i\Delta_n}P_{i\Delta_n}\sqrt{\Delta_n})$ is a “rounding impact ratio” defined as the magnitude of price discreteness (i.e., d) relative to the (approximate) volatility of the return process over Δ_n (i.e., $\sigma_{i\Delta_n}P_{i\Delta_n}\sqrt{\Delta_n}$).

Assume, now, that $x = P - \bar{P}$ is uniformly distributed over $[-d/2, d/2]$. In other words, when we observe the rounded price $\bar{P}$ we do not have any information as to the location of P . Setting, in Eq. (5), $\zeta_{i\Delta_n}$ to the estimable mean value ζ_{Δ_n} ⁷ and integrating x out, we evaluate the probability

⁶The expression in Eq. (5) is based on a Δ_n -discretization of the process in Eq. (1).

⁷The unconditional probability of a k -tick movement is the asymptotic average (over i) of the conditional probabilities in Eq. (5) integrated over the distribution of the displacements x . Technically, then, for each k , ζ_{Δ_n} is not exactly a “mean” value but rather, by Jensen’s inequality, a “mean” value adjusted for nonlinearities in the integrands in Eq. (5).

$p_{\Delta_n}^{k,R}(\zeta_{\Delta_n})$ of observing a k -tick movement over an interval Δ_n by virtue of

$$p_{\Delta_n}^{k,R}(\zeta_{\Delta_n}) = \frac{(k-1)}{2} \operatorname{erf}\left(\frac{(k-1)\zeta_{\Delta_n}}{\sqrt{2}}\right) - k \cdot \operatorname{erf}\left(\frac{k\zeta_{\Delta_n}}{\sqrt{2}}\right) + \frac{(1+k)}{2} \operatorname{erf}\left(\frac{(k+1)\zeta_{\Delta_n}}{\sqrt{2}}\right) \\ + \frac{1}{\sqrt{2}\pi} \left(e^{-\frac{1}{2}(1+k)^2\zeta_{\Delta_n}^2} \left(1 + e^{2k\zeta_{\Delta_n}^2} - 2e^{-\frac{1}{2}(1+2k)\zeta_{\Delta_n}^2} \right) \right) \frac{1}{\zeta_{\Delta_n}},$$

where $k = 0, \pm 1, \dots$ and the symbol $\operatorname{erf}(x)$ defines the Gaussian error function.

We note that the lower volatility (σ_t), the larger the probability $p_{\Delta_n}^{0,R}$ of a zero return due to price discreteness. Low volatility is, however, generally associated with low volume (c.f., Tauchen and Pitts, 1983, and the references therein). Thus, rounded up traded prices (and, as a consequence, zero returns) are hardly in contradiction with slow trading. Once more, price discreteness would represent a spurious effect if it were to occur along with large volume, something which is not typically found in the data. In this sense, $p_{\Delta_n}^{0,R}$, which is silent about volume, may be viewed as a (possibly generous) upper bound on the spurious impact of price discreteness on genuine zeros. We note that the inclusion of bid/ask bounce effects in the data generating process would further reduce $p_{\Delta_n}^{0,R}$.

Re-activating staleness through independent Bernoulli variates, the probability of a k -tick movement becomes

$$\tilde{p}_{\Delta_n}^{k,R}(\zeta_{\Delta_n}, p_n) = p_n \mathbb{I}_{k=0} + \sum_{j=0}^{\infty} (1-p_n)^2 p_n^j \cdot p_{\Delta_n}^{k,R}\left(\zeta_{\Delta_n}/\sqrt{j+1}\right), \quad (6)$$

where $\mathbb{I}$ denotes the indicator function. The formula hinges on the fact that staleness translates into repeated observations of the rounded process backward in time. The Appendix justifies Eq. (6) formally.

We note that the probability in Eq. (6) depends on two parameters, namely the “rounding impact ratio” ζ_{Δ_n} and p_n . Both have economic significance. We have commented on ζ_{Δ_n} . The quantity $\hat{p}_n$ is what we label *excess staleness*, namely the estimate of the implied probability of staleness (p_n) after controlling for rounding. We compute $\hat{\zeta}_{\Delta_n}$ and $\hat{p}_n$ by matching $\tilde{p}_{\Delta_n}^{k,R}(\zeta_{\Delta_n}, p_n)$ with the empirical frequency of k -tick price movements, for all observed k -tick movements.

In Figure 3, we plot $\hat{p}_n$, averaged across stocks, as a function of sampling frequency. Consistent with the model (c.f., Eq. (2)), the probability of excess staleness is frequency-specific and, as expected, increasing with the sampling frequency. According to the model, there could be clustering in excess staleness,⁸ but the assumed dependence structure of the trading indicators (i.e., the Bernoulli variates) is such that their empirical averages converge to the probability of excess staleness at the highest frequency of observation (i.e., p_{∞}), when the observation frequency increases without bound (as $n \rightarrow \infty$), c.f., Eq. (3). If we were to assume that the highest frequency is 30 seconds, then Figure 3 would imply that p_{∞} is about 20%. Equivalently, at 30 seconds, one in five return observations is stale, a sizable magnitude. The corresponding probability of idleness decays at a quicker pace than the probability of excess staleness. The increasing wedge between them suggests that, at low frequencies, all observed zeros occur in the presence of volume which is strictly non-zero, albeit small.

Figure 4 reports average intra-day staleness (black lines), average intra-day idleness (red lines

⁸We model clustering explicitly in Section 4. Eq. (6) hinges on independence of the Bernoulli variates.

Figure 3: We report the average (estimated) probability of excess staleness (blue line with circles) and the average probability of idleness (red line with diamonds) as a function of the sampling frequency. The plots are cross-sectional averages based on 249 NYSE-listed stocks, whose trades are recorded from 9:30 to 16:00, for the years 2006-2014.

with squares) and the *estimated* (using $p_{i\Delta_n}^{0,R}$) upper bound on staleness due to price discreteness (magenta lines with diamonds).⁹ It also reports the sum of average intra-day idleness and staleness due to rounding (blue lines with circles). Both staleness and idleness display a pronounced inverse-U shape mirroring lower volatility and lower trading activity around lunch time. Both quantities are also smaller (on average) every five minutes and, in particular, every half hour, signaling behavioral clustering of trading activity at round time marks. They are, moreover, lower at 10:00 a.m., reflecting the frequent release of macroeconomic announcements at this time. As expected, average staleness and idleness decrease with the sampling frequency. The higher the sampling frequency, the more the adjustment in Eq. (6) fills the gap between staleness and genuine idleness. We interpret this result as suggesting that, at lower frequencies, there is a component of staleness which is not spurious (i.e., due to rounding) and is, also, not due to zero volume, i.e. idleness. Said differently, we may observe genuine staleness in the presence of low (but strictly non-zero) volume. This additional (third) component of staleness, which we called earlier “near idleness”, is relatively more important the lower the sampling frequency.

⁹We note that, in Figure 4, the implementation of Eq. (6) requires estimation of a time-varying spot variance. We identify the latter as follows. For each stock k , first, we compute daily realized variances by summing the squares of the intra-day continuously-compounded returns, i.e., $RV_{t_k} = \sum_{i=1}^n r_{i,t_k}^2$. Second, we define $\hat{\sigma}_{i,t_k}^2 = w_i RV_{t_k}$ for each intra-day return on day t_k , where w_i is a scaling factor designed to account for intra-day effects: $w_i = \frac{1}{T_k} \sum_{t_k=1}^{T_k} \frac{r_{i,t_k}^2}{RV_{t_k}}$.

Figure 4: We report average intra-day staleness (black lines), average intra-day idleness (red lines with squares), estimated intra-day staleness from rounding (magenta lines with diamonds) and the sum of average intra-day idleness and estimated staleness from rounding (blue lines with circles). Prices are sampled every 30 seconds (Panel A), every 1 minute (Panel B) and every 2 minutes (Panel C). The plots are cross-sectional and daily averages based on 249 NYSE-listed stocks, whose trades are recorded from 9:30 to 16:00, for the years 2006-2014.

3 The economics of staleness

What generates staleness and idleness? As pointed out, rounding leads to some staleness. However, at all reported frequencies, staleness is significantly higher than its spurious component due to rounding. Bandi et al. (2017) provide a modeling framework in which the interplay between execution costs and information (about fundamental, or efficient, values) justifies idleness. The framework hinges on micro-structural theories of price formation with transaction costs and asymmetries in information (c.f., Hasbrouck and Ho, 1987, Kyle, 1985, and Glosten and Milgrom, 1985). The logic is as follows: informed traders, by definition, “know” fundamental values. The trading decisions of these market participants should, therefore, be sensitive to the *relative* difference between, on the one hand, the gap between midpoints of bid/ask spreads and fundamental values and, on the other hand, quoted (and perceived) execution costs, inclusive of (half) bid/ask spreads. Generally speaking, these traders will execute trades only when the trades guarantee a profit, net of execution costs. Should execution costs be excessively large, they will be reluctant to trade. Uninformed traders may also respond to the magnitude of execution costs and trade less when these costs are perceived to be too large. Because fundamental values are however unknown to them, their trading decisions will, in general, be affected by the *absolute* magnitude of the cost of transacting (or, more realistically, by *perceived* relative magnitude, given a *prediction* of fundamental value).

When taken literally, the baseline version of the model only allows for binary choices (trade or not) in information-based trading (as a function of the cost of execution) and, upon trading, for unit volume. A natural extension of the same conceptual framework would allow for alternative levels of volume. The extension would provide economic support for absence of trading and zero volume (particularly at high frequencies) or limited trading on the same side of the market and low volume without price impacts (particularly at lower frequencies) as a source of zero returns, when the cost of execution is large. Since execution costs capture at least one dimension of liquidity (“tightness”, rather than “depth” or “resiliency”), the extended framework would, in the end, provide economic

Figure 5: We report average intra-day volume associated with non-zero returns (red line with squares), average intra-day volume associated with zero returns (blue line with circles), both with 99% confidence bands, and average intra-day volume associated with zero returns with non-zero volume (black line). We consider three sampling frequencies: 30 seconds (Panel A), 1 minute (Panel B) and 2 minutes (Panel C). The plots are cross-sectional and daily averages based on 249 NYSE-listed stocks, whose trades are recorded from 9:30 to 16:00, for the years 2006-2014.

support for a positive relation between illiquidity and zero returns with volume representing the *mediating* variable.

A careful look at volume is, at this juncture, economically revealing. Figure 5 displays the average intra-day volume (with 99% confidence bands) associated with zero returns (blue lines with circles) and non-zero returns (red lines with squares) for the same three frequencies considered in Figure 4. We note that intra-day average volume has a U-shaped pattern which nicely reflects the inverse U-shaped pattern in staleness (c.f., Figure 4). Upward spikes in volume every 5 minutes (and, in particular, every half hour) also mirror downward spikes in staleness at the same times.

Importantly, the average volume associated with stale returns is (statistically) significantly lower than the average volume associated with non-zero returns at all frequencies and, especially, at high frequencies. Because stale returns are often idle, that is with zero volume, we also report the average intra-day volume of stale returns with non-zero volume (black lines). The latter is bound to capture some price discreteness (since non-zero volume can be associated with stale prices when price discreteness plays a role) but also near idleness (since limited trading may lead to staleness). Not surprisingly, as we move to lower frequencies, a larger proportion of stale returns is characterized by strictly positive (i.e., non-zero) volume. Yet, the separation between the lower (possibly non-zero, particularly at lower frequencies) volume associated with stale returns and the higher volume associated with non-stale returns is evident.

We now evaluate the cross-sectional relation between the average probability of excess staleness and average dollar volume. We do so in Figure 6. At all frequencies, the correlations are negative, implying a lower likelihood of staleness in conjunction with higher volume, and remarkably large.

Should price discreteness induce a large percentage of spurious zeros, i.e., zeros associated with relatively large volume, we would expect total staleness to display lower (in absolute value) correlation with respect to volume than excess staleness. The corresponding results are reported in the lower panels of Figure 6. While, in agreement with the confounding effect of price discreteness, we do witness a reduction in correlation (particularly at high frequencies), the relation between staleness

Figure 6: We plot average excess staleness versus average dollar volume for all stocks in our sample. We consider three frequencies, 30 seconds (Panel A), 1 minute (Panel B) and 2 minutes (Panel C). We also plot average total staleness versus average dollar volume for the same three frequencies, Panel D through F. We use a cross-section of 249 NYSE-listed stocks sampled over the years 2006-2014.

Figure 7: We plot the average probability of excess staleness versus the average logarithmic bid/ask spread over three frequencies: 30 seconds (Panel A), 1 minute (Panel B) and 2 minutes (Panel C). We use a cross-section of 249 NYSE-listed stocks sampled over the years 2006-2014.

(total, this time) and volume continues to be strong. Once more, rounded up prices which do not move because of small volume are not in contradiction with our documented negative relation between staleness and volume. In this sense, a (possibly large) component of price discreteness is, just like in the case of excess staleness, a reflection of slow trading.

As stressed above, the reported findings may be suggestive of a positive relation between illiquidity and excess staleness with volume representing the mediating variable. In order to verify this conjecture, we employ the most classical liquidity measure, namely traded logarithmic bid/ask spreads. For the firms in our sample, Figure 7 reports the cross-sectional relation between the average probability of excess staleness and average logarithmic bid/ask spreads associated with three frequencies. In a nutshell, firms with more excess staleness have, on average, larger spreads.

Figure 8: We report average logarithmic bid/ask spreads associated with idle returns (red line with diamonds) and non-idle returns (blue line with circles) across sampling frequencies. The plots are cross-sectional averages based on 249 NYSE-listed stocks sampled over the years 2006-2014.

It is an empirical question as to whether the number of zeros adjusts dynamically to the size of the spreads. Consider a firm with, say, more excess staleness and larger spreads. This firm may simply have more zeros than a firm with lower spreads, but equally large spreads associated with zeros and non zeros. Figure 8 reports the average size of the logarithmic bid/ask spreads associated with idle returns (red line with diamonds) and non-idle returns (blue line with circles) for various sampling frequencies. We report a monotonic increase (over sampling frequencies) in the average size of the spreads associated with idleness. The size of the spreads associated with non-idle returns is, instead, stable. The wedge between spreads associated with idle returns and spreads associated with non-idle returns increases with reductions in sampling frequency.

The evidence in Figure 6 and in Figure 7 is, necessarily, univariate. In order to better evaluate the *joint* cross-sectional dependence of excess staleness on volume, bid-ask spreads and volatility, we regress (across stocks) averages of daily estimates of 30-second excess staleness, $\hat{p}_n$, on averages of daily logarithmic dollar volume, $\log DV$, averages of logarithmic bid-ask spreads, b/a in basis points, and averages of daily logarithmic 5-minute realized variances, $\log RV$. The outcome of the regression (with t -statistics under the parameter estimates) is

$$\hat{p}_n = 0.807 - \frac{0.085}{(7.00)} \log DV + \frac{0.012}{(11.53)} b/a - \frac{0.065}{(-10.90)} \log RV + \hat{\varepsilon}. \quad (7)$$

In the *absence* of volume, the sign and statistical significance of execution costs and variance would not be surprising. We have already commented on the positive relation between the magnitude of execution costs and excess staleness. The impact of variance can also be easily reconciled. Consistent with the realized volatility literature, $\log RV$ is designed to be an estimate of the (average, across days) logarithm of the integrated spot variance (σ^2 in Eq. (1)) of the underlying fundamental price process.

In an economy, such as the one described above, in which informed agents detect larger variability of fundamental prices around quoted midpoints of bid-ask spreads, superior trading opportunities would be available to them, trading would be more active and (excess) staleness would be lower, for every level of execution costs. In the *presence* of volume, because of its mediating role, we would expect execution costs and variance to be considerably less influential than reported in the previous regression. In order to better understand their role, we add a measure of volume variability to the regression. The measure is the average of the logarithm of the daily standard deviation of volume observed on 30-second intervals over DV , i.e., the logarithmic coefficient of variation of volume, $\log CV$. The regression now becomes

$$\hat{p}_n = \underset{(12.98)}{3.389} - \underset{(-23.03)}{0.091} \log DV + \underset{(2.00)}{0.003} \text{b/a} - \underset{(-0.26)}{0.002} \log RV + \underset{(10.63)}{0.276} \log CV + \hat{\varepsilon}. \quad (8)$$

The log coefficient of variation is strongly positively related to excess staleness and reduces drastically the significance of both execution costs and variance. It is intuitive that if, given a level of daily volume, trading is more spread out over the day, daily excess staleness should be lower. This is, however, a case of low volume variability. Conversely, for any level of daily volume, clusters of volume during the day would give rise to high volume variability and, of course, numerous instances of zero returns associated with lack of trades over the day, i.e., higher daily excess staleness. The significance of execution costs (with a positive sign) and variance (with a negative sign) in the regression in Eq. (7) can now be explained by the association between higher execution costs, or lower variance, and more clustering (i.e., more variance) in volume, for every level of daily volume. Said differently, volume variability is a key omitted variable in Eq. (7) correlated with both execution costs and fundamental price variance. In essence, the centrality of volume as a determinant of excess staleness over a period operates *both* through its level and its variability during the same period.

Next, we report the outcome of an asset pricing exercise whose aim is to shed additional light on the economics of staleness. In order to do so, we construct an *excess staleness factor* in a traditional way. At time t , we sort stocks into deciles using the excess staleness observed over the previous month. We then construct equally-weighted decile portfolios. The excess staleness factor is the difference between the return on the top-decile portfolio and the return on the bottom-decile portfolio, which we label as $R_{ES,t}$. We use monthly rebalancing (22 days) and regress the monthly returns of this long-short strategy on a state-of-the-art, 5-factor, Fama-French model (c.f., Fama and French, 2015) in which risk is captured by the market ($R_{M,t} - R_{F,t}$, with $R_{F,t}$ denoting the risk-free rate), size (SMB_t), value (HML_t), profitability (RMW_t) and investment (CMA_t). The output of the regression, again with t -statistics under the parameter estimates, is as follows

$$\begin{aligned} R_{ES,t} - R_{F,t} = & \underset{(3.37)}{0.0100} - \underset{(-0.56)}{0.0499} (R_{M,t} - R_{F,t}) + \underset{(3.64)}{0.8242} SMB_t + \underset{(1.33)}{0.2498} HML_t \\ & - \underset{(-1.60)}{0.6174} RMW_t + \underset{(2.10)}{0.8418} CMA_t + \hat{\varepsilon}_t. \end{aligned} \quad (9)$$

The interpretation is conventional. The positive and significant value of the intercept suggests that excess staleness leads to yet another anomaly hardly explained by well-accepted risk factors. Differently from Liu (2006), who employs an illiquidity measure based on daily zeros, excess staleness

exploits the full potential of high-frequency data. Should one view the reported excess staleness anomaly as an illiquidity anomaly, our evidence points to an illiquidity premium even in the short (for pricing purposes) sample, and for the extremely liquid NYSE stocks, used in this study.

In light of the dependence between excess staleness and volume (levels and variability), we now relate the returns on the long-short excess staleness portfolio to the returns on a long-short volume portfolio (i.e., a portfolio long stocks in the high volume decile and short stocks in the low volume decile) and the returns on a long-short volume-*CV* portfolio (i.e., a portfolio long stocks in the high volume-*CV* decile and short stocks in the low volume-*CV* decile):

$$R_{ES,t} - R_{F,t} = 0.0019 - 0.5419 R_{DV,t} + 0.2341 R_{CV,t} + \hat{\varepsilon}_t. \quad (10)$$

(0.61) (-7.13) (3.88)

Consistent with the excess staleness regression in Eq. (8), the excess staleness factor is negatively correlated with the volume factor and positively correlated with the volume-*CV* factor.¹⁰ Also, when controlling for the volume factor and the volume-*CV* factor, the intercept is small and statistically insignificant. We note that a very similar outcome would be obtained if, rather than regressing the excess staleness factor on separate volume and volume-*CV* factors, we regressed the excess staleness factor on a *joint* volume/volume-*CV* factor. The latter is obtained by double-sorting with respect to the volume dimension and the volume-*CV* dimension into quintiles before going long high volume - low volume-*CV* stocks and going short low volume - high volume-*CV* stocks:¹¹

$$R_{ES,t} - R_{F,t} = 0.0038 - 0.3669 R_{DV-CV,t} + \hat{\varepsilon}_t. \quad (11)$$

(1.18) (-8.16)

In essence, even when seen through an asset pricing lens, excess staleness appears to be negatively correlated with volume levels and positively correlated with volume variability. This conclusion is, once more, in agreement with the economic logic we put forward earlier: excess staleness is hardly affected by rounding and is intimately associated with trading activity.

Our focus in this section has been on the *economic determinants* of staleness. The proposed analysis of a long-short excess staleness portfolio is, however, also suggestive of the pricing potential of excess staleness, a sufficient statistic for moments of the volume distribution and, hence, for trading dynamics. Differently put, our analysis is suggestive of important *economic implications* other than those discussed in the next section. A richer investigation of such a potential is a topic better left for future research.

¹⁰The positive correlation with the volume-*CV* factor also reflects a positive premium associated with high volume-*CV* stocks. A positive compensation for variability in illiquidity, in addition to one for illiquidity levels, is in keeping with traditional economic logic. It is, however, counter to the negative premium found by Chordia et al. (2001) using an earlier, lower-frequency, sample.

¹¹Controlling, as in Eq. (9), for the 5 Fama-French factors in Eq. (10) and in Eq. (11) would not modify the reported results in any meaningful way.

4 Risk-neutral staleness and short-dated options

The market for weekly options (and options with less than a week to expiration) has seen a substantial increase in traded volume over the last few years. As emphasized by Andersen et al. (2017), the proportion of volume of these instruments on the CBOE was about 10% in 2010 and it is now well over 30%. Short-dated at-the-money options are known to respond to (spot) volatility risk, short-dated out-of-the-money options are, instead, particularly responsive to jump risk (c.f., Andersen et al., 2017). In both cases, the microstructure of the observed price process matters. In particular, we show that illiquidity-induced staleness decreases option prices. High aversion to illiquidity, leading to increased levels of the *risk-neutral* probability of staleness, would enhance this effect and yield high expected (option) returns (as compensation for perceived illiquidity risk). The channel is intuitive: the higher the probability of staleness and its risk compensation, the tighter the risk-neutral distribution of the price process, the lower option prices.

To clarify, Figure 9 presents a 3-step tree which incorporates staleness. The *latent* (efficient) price process evolves on a binomial tree, with a risk-neutral probability q_u of moving up and a risk-neutral probability q_d of moving down. The *observed* price process has a risk-neutral probability of staleness equal to p_n^* and can, consistent with Eq. (4), either stay constant (on the dotted branches) or revert to the efficient price. Naturally, pricing is based on observed prices. On every branch, the figure reports their conditional probabilities of transitioning from one node to the other. It also reports the unconditional probabilities (P_{node} on nodes $\in \{E, F, G, H, I\}$) associated with the final nodes. This simple representation illustrates the impact of spurious jumps, discussed earlier. Assume the observed price process remains stale between node A and node C. Should it move next, the traded price has a probability (p_{CE} or p_{CI}) of reaching node E and node I, respectively, thereby skipping two, rather than one, step (c.f. Figure 2 for empirical evidence).

We work with efficient and traded prices which are both risk-neutral. In order to do so, we set the risk-free rate equal to zero, which is reasonable in current economic environments and is certainly reasonable for short-dated options. The zero-rate assumption also guarantees that put options written on a multi-node extension of the above tree (our focus below) have zero time value of money and would, therefore, never be exercised early. Thus, the reported prices refer to both American and European-style put options.

We focus on daily puts ($T = 1/252$) and set the intra-daily frequency ($\Delta_n = T/n$) equal to 5 minutes. Specifically, we generalize the tree in Figure 9 to $n = 78$ nodes¹² so that the corresponding intra-daily frequency is 5 minutes. On the tree, following Cox et al. (1979), the up moves are expressed as $u = \exp(\sigma\sqrt{\Delta_n})$ and the down moves are expressed as $d = \exp(-\sigma\sqrt{\Delta_n})$, where σ (the annualized volatility of the efficient price process) is set equal to 20%. The risk-neutral probabilities of up and down moves for the efficient price process are, therefore, given by $q_u = \frac{1-d}{u-d}$ and $q_d = 1 - q_u$. The initial price S_0 is set equal to \$1 and the strike K is equal to $\$(1 - 2\sigma\sqrt{T}) = \0.9748 . The resulting

¹²The tree illustrated in Figure 9 (with $n = 3$, solely for conciseness) continues to be treatable for any n since it is possible to recombine states. A state is characterized by the efficient price, the traded price and whether the previous return was zero or not. The number of states grows linearly with n , as in the standard binomial model. The code that generates the tree (and the corresponding probabilities) and computes the derivative price for any payoff (and both European and American-style instruments) is available upon request.

$$\begin{aligned}
P_{CE} &= (1 - p_n^*)q_u^2 \\
P_{CI} &= (1 - p_n^*)q_d^2 \\
P_{CG} &= p_n^*(q_u^2 + q_d^2) + 2q_uq_dp_n^* + 2q_uq_d(1 - p_n^*) \\
P_E &= (1 - p_n^*)^2q_u^2 + (1 - p_n^*)p_n^*q_u^2 \\
P_F &= p_n^*(1 - p_n^*)q_u \\
P_G &= (p_n^*)^2(q_u^2 + q_d^2) + 2(p_n^*)^2q_uq_d + 2(1 - p_n^*)^2q_uq_d + 2p_n^*(1 - p_n^*)q_uq_d \\
P_H &= p_n^*(1 - p_n^*)q_d \\
P_I &= (1 - p_n^*)^2q_d^2 + (1 - p_n^*)p_n^*q_d^2
\end{aligned}$$

Figure 9: Example of a 3-step tree to value options when the efficient price process is partially observable due to staleness.

put is, therefore, two standard deviations out of the money.

Figure 10 (black line with circles) reports percentage price changes with respect to the no-staleness case as a function of the unconditional risk-neutral probability of staleness. In other words, it reports the percentage pricing error that one would make when ignoring staleness in the risk-neutral dynamics.

Figure 10: For a put option with daily maturity and strike two standard deviations below the initial price S_0 , we report percentage price changes (with respect to the no-staleness case) as a function of the unconditional risk-neutral probability of staleness for three different levels of the persistence (in staleness) parameter π_n^* . We use the following parameter values: $S_0 = 1$, $\sigma = 20\%$, $K = 0.9748$, $T = 1/252$, $\Delta_n = T/n$ with $n = 78$.

The percentage price drop is economically large, and particularly so for high values of p_n^* , that is for highly illiquid stocks and/or for high aversion to illiquidity risk reflected in high aversion to staleness. In essence, drastic spikes of illiquidity leading to a high risk-neutral likelihood of intra-daily (over 5 minutes, in this study) staleness may generate a substantial illiquidity premium in expected put option returns. Because after a period of staleness prices revert to the underlying efficient values, we expect the impact of staleness on pricing to be magnified when pricing is more heavily affected by the *dynamics* of the observed price process. This is the case for path-dependent options but it is also the case in instances in which staleness is *persistent*, a case to which we now turn.

The data-generating process in Eq. (4) allows for dependence in the driving Bernoulli variates. We add persistence in the staleness dynamics by using the following (risk-neutral) specification:

$$\mathbb{B}_{i,n} = \mathbb{B}_{i-1,n} \mathbb{B}_{i,n}^{(2)} + (1 - \mathbb{B}_{i,n}^{(2)}) \mathbb{B}_{i,n}^{(1)} \quad i \geq 2, \quad (12)$$

with $\mathbb{B}_{1,n} = \mathbb{B}_{1,n}^{(1)}$, where the $\mathbb{B}_{i,n}^{(1)}$ s are iid Bernoulli variates with risk-neutral probability p_n^* and the $\mathbb{B}_{i,n}^{(2)}$ s are iid Bernoulli variates independent of $\mathbb{B}_{i,n}^{(1)}$ with risk-neutral probability π_n^* . In agreement with the baseline case without persistence, $\mathbb{E}[\mathbb{B}_{i,n}] = \mathbb{E}[\mathbb{B}_{i,n}^{(1)}] = p_n^*$, so that p_n^* continues to be the unconditional risk-neutral probability of staleness. The parameter π_n^* drives persistence. When $\pi_n^* = 0$, the Bernoulli variates $\mathbb{B}_{i,n}$ are, of course, iid.

Figure 10, blue line with crosses and red line with squares, provides results for two levels of π_n^* . As persistence increases, even moderate levels of (risk-neutral) staleness associated with liquid stocks may generate sizable illiquidity risk premia and, consequently, sizable drops in option prices.

Some observations are in order. First, this section is illustrative. Future work should estimate the risk-neutral probability of staleness, as well as the risk-neutral persistence in staleness, using short-

and long-maturity option data by focusing on assets with various liquidity properties. As reported by Lesmond (2005), in certain markets the daily proportion of zeros is extremely high. Translating his results into our nomenclature, the *daily* “objective” probability of staleness is extremely high. In the same market, the *intra-daily* “risk-neutral” probability is expected to be even higher. Second, our approach connects the market microstructure of price formation to both the pricing of the underlying and the pricing of derivatives, the latter, in particular, being a rather under-explored area of research. Finally, we view the discussion in this section as being of independent interest. The same structure as in Figure 9, and its extension with $n > 3$ nodes, may be employed to price real options, e.g., options to delay, in which the underlying (the expected future stream of income of the project) is latent, due to asymmetries in information, and only revealed with some probability at every step.

5 Conclusions

This paper provides theoretical and empirical foundations for the relevance of staleness (and excess staleness) in high-frequency asset prices, a novel stylized fact shown to be linked to the distribution of volume.

Institutional effects, such as price discreteness, do not represent a first-order, spurious determinant of staleness. Rounding can not hide sizable price impacts due to large volatility and large volume. On the contrary, rounding will lead to more zeros in the presence of low volatility, low volatility states being, in general, associated with low volume. In this sense, rounding-induced zeros are not in contradiction with our economic logic. Our proposed correction for rounding (leading to a notion of excess staleness) may, therefore, be viewed as very conservative.¹³

Our evidence supports an alternative data generating process for asset prices, one which deviates in important ways from classical semimartingale modeling and, instead, allows for clusters of inactivity. Because we have focused on the NYSE-listed stocks with the largest traded volume over the period, i.e., those whose prices are expected to be the least inactive, our findings should be viewed, once more, as being conservative.

Econometrically, the presence of zeros is bound to alter the inferential properties of moment-based statistics relying on high-frequency data, from jump tests (see, e.g., Bajgrowicz et al., 2015 and the references therein) to higher-order moment estimates (see, e.g., Amaya et al., 2015). Both have been central to the high-frequency finance literature and its applications in asset pricing and asset management. We will report on the impact of zeros on key inferential problems in future work.

Economically, we expect the presence, magnitude and dynamic properties of the number of zeros to be informative about the microstructure of price formation mechanisms in modern financial markets, from the extent of asymmetric information, to the dynamic evolution of unobserved fundamental values, to learning about fundamental values. The structural analysis in Bandi et al. (2017), justified by the empirical work in this paper, is only a first step in this area.

¹³The link between tick size and liquidity provision is the subject of a considerable amount of work, see, e.g., Goldstein and Kavajecz (2000). The impact on staleness and excess staleness of changes in tick size around events like the 2016 SEC “Pilot” is of interest to further study the effects of rounding on the dynamics of a quantity, like excess staleness, which is designed to capture genuine liquidity effects, net of price discreteness.

References

- Amaya, D., P. Christoffersen, K. Jacobs, and A. Vasquez (2015). Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics* 118(1), 135–167.
- Andersen, T. G., N. Fusari, and V. Todorov (2017). Short-term market risks implied by weekly options. *Journal of Finance* 72(3), 1335–1386.
- Back, K. (2010). *Asset pricing and portfolio choice theory*. Oxford University Press.
- Bajgrowicz, P., O. Scaillet, and A. Treccani (2015). Jumps in high-frequency data: Spurious detections, dynamics, and news. *Management Science* 62(8), 2198–2217.
- Bandi, F., D. Pirino, and R. Renò (2017). EXcess Idle Time. *Econometrica* 85(6), 1793–1846.
- Chordia, T., A. Subrahmanyam, and V. R. Anshuman (2001). Trading activity and expected stock returns. *Journal of Financial Economics* 59(1), 3–32.
- Cox, J. C., S. A. Ross, and M. Rubinstein (1979). Option pricing: A simplified approach. *Journal of Financial Economics* 7(3), 229–263.
- Delattre, S. and J. Jacod (1997). A central limit theorem for normalized functions of the increments of a diffusion process, in the presence of round-off errors. *Bernoulli*, 1–28.
- Duffie, D. (2008). *Dynamic asset pricing theory*. Princeton University Press.
- Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics* 116(1), 1–22.
- Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14(1), 71–100.
- Goldstein, M. A. and K. A. Kavajecz (2000). Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the nyse. *Journal of Financial Economics* 56(1), 125–149.
- Goyenko, R., C. Holden, and C. Trzcinka (2009). Do liquidity measures measure liquidity? *Journal of Financial Economics* 92(2), 153–181.
- Hasbrouck, J. and T. Ho (1987). Order arrival, quote behavior, and the return-generating process. *Journal of Finance* 42(4), 1035–1048.
- Jarrow, R. A. (2018). *Continuous-Time Asset Pricing Theory: A Martingale-Based Approach*. Springer.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica* 53(6), 1315–1335.
- Lesmond, D., J. Ogden, and C. Trzcinka (1999). A new estimate of transaction costs. *Review of Financial Studies* 12(5), 1113.
- Lesmond, D. A. (2005). Liquidity of emerging markets. *Journal of Financial Economics* 77(2), 411–452.
- Liu, W. (2006). A liquidity-augmented capital asset pricing model. *Journal of Financial Economics* 82(3), 631–671.
- Phillips, P. C. and J. Yu (2007). Information loss in volatility measurement with flat price trading. *Cowles Foundation Discussion Paper No. 1598*.
- Tauchen, G. and M. Pitts (1983). The price variability-volume relationship on speculative markets.

Econometrica 51, 485–505.

A Appendix: rounding and implied excess staleness

Let us consider the case $k = 0$. Define the event “observed staleness in the absence of true staleness” as:

$$R_{i,i+1} = \{\bar{P}_{(i+1)\Delta_n} - \bar{P}_{i\Delta_n} = 0\} \mid \{\mathbb{B}_{i+1,n} = 0\}.$$

Note that

$$\mathbb{P}[R_{i,i+1} \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n}] \quad (13)$$

$$= \int_{-d/2}^{d/2} \mathbb{P}[\bar{P}_{i\Delta_n} - d/2 < P_{(i+1)\Delta_n} < \bar{P}_{i\Delta_n} + d/2 \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n} = \tilde{P}_{i\Delta_n} - x] dx \quad (14)$$

$$= \int_{-d/2}^{d/2} \mathbb{P}[\bar{P}_{i\Delta_n} - d/2 < P_{(i+1)\Delta_n} < \bar{P}_{i\Delta_n} + d/2 \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n} = P_{(i-K_{i,n})\Delta_n} - x] dx, \quad (15)$$

where, for all $i = 1, \dots, n$, the number of consecutive flat trades before instant $i\Delta_n$ is

$$K_{i,n} = \min \{j \in \{0, \dots, i\} \mid \mathbb{B}_{i,n} = 1, \mathbb{B}_{i-1,n} = 1, \dots, \mathbb{B}_{i-j+1,n} = 1, \mathbb{B}_{i-j,n} = 0\}.$$

The key is to recognize that, in the conditioning statement in Eq. (14), the observed price $\bar{P}$ is not a function of the fundamental price P but, rather, it is a function of the stale price $\tilde{P}$. Hence, Eq. (5) cannot be used directly before writing the problem as in Eq. (15). Because $\mathbb{P}[K_{i,n} = j] = (1 - p_n) p_n^j$, we have

$$\begin{aligned} & \mathbb{P}[R_{i,i+1} \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n}] \\ &= \sum_{j=0}^{\infty} \int_{-d/2}^{d/2} \mathbb{P}[\bar{P}_{i\Delta_n} - d/2 < P_{(j+1)\Delta_n} < \bar{P}_{i\Delta_n} + d/2 \mid \sigma_{(i-j)\Delta_n}, \bar{P}_{i\Delta_n} = P_{(i-j)\Delta_n} - x] dx \mathbb{P}[K_{i,n} = j] \\ &= \sum_{j=0}^{\infty} \int_{-d/2}^{d/2} \mathbb{P}[\bar{P}_{i\Delta_n} - d/2 < P_{(j+1)\Delta_n} < \bar{P}_{i\Delta_n} + d/2 \mid \sigma_{(i-j)\Delta_n}, \bar{P}_{i\Delta_n} = P_{(i-j)\Delta_n} - x] dx (1 - p_n) p_n^j. \end{aligned} \quad (16)$$

Define, now, the event “observed staleness in the presence of true staleness” as:

$$\bar{R}_{i,i+1} = \{\bar{P}_{(i+1)\Delta_n} - \bar{P}_{i\Delta_n} = 0\} \mid \{\mathbb{B}_{i+1,n} = 1\}.$$

It is clear that

$$\mathbb{P}[\bar{R}_{i,i+1} \mid \sigma_{i\Delta_n}, \bar{P}_{i\Delta_n}] = 1. \quad (17)$$

Finally, Eq. (16) and Eq. (17) lead to Eq. (6) after recalling that $\mathbb{P}[\mathbb{B}_{i+1,n} = 0] = 1 - \mathbb{P}[\mathbb{B}_{i+1,n} = 1] = 1 - p_n$. The case $k > 0$ now follows immediately.