

Ole Eiler Barndorff-Nielsen and financial econometrics*

Neil Shephard

*Department of Economics and Department of Statistics,
Harvard University*

May 28, 2024

Abstract

This note reviews some of the contributions Ole Eiler Barndorff-Nielsen made to financial econometrics. He was particularly active in that area from the mid-nineties for around a dozen years. His innovations include the NIG Lévy process, the Barndorff-Nielsen-Shephard model, the supOU process, the formalization of realized volatility and realised beta, the introduction of bipower variation, realized semivariance, realized kernels and gradual jumps.

Keywords: High frequency financial econometrics, jumps, in-fill asymptotics, Lévy process, stochastic volatility.

1 Introduction

Ole Eiler Barndorff-Nielsen (1935-2022) was a Danish mathematician and statistician whose research career was dedicated to stochastics and the foundations of statistical inference. From the mid-nineties for a dozen years he mainly focused on research questions which arose from modeling and measurement in financial econometrics, although in parallel he spent considerable time studying models of turbulence in wind fields and quantum probability. This note focuses entirely on his financial work over that period, organized around two themes: Lévy based model building and high frequency financial econometrics.

Ole brought to his work on finance an extraordinarily inventive mind, focus, a powerful mathematical technique and a dedication to supporting young scholars. His past research on, for example, asymptotic approximations (e.g. Barndorff-Nielsen (1983), Barndorff-Nielsen and Cox (1979, 1989, 1994)), infinite divisibility (Barndorff-Nielsen and Halgreen (1977)) and mixture distributions such as the generalized hyperbolic distributions (Barndorff-Nielsen (1979)), had an impact on how he attacked some of the financial problems he addressed.

There is another aspect which is important to keep in mind.

Ole cared not just about the elegance and mathematical insight of his work, but equally that his methods performed empirically — even though he was a pure theorist. This comes across in earlier research on the distribution of blown sand and his work on turbulence of wind fields. He regarded himself as a scientist, not “just” a mathematician.

*This note was written for a keynote talk at a conference held in honor of Ole’s academic career, held at Aarhus University, 29rd-31st May 2024. I am grateful to the conference organizers for the invite and to Torben Andersen, Tim Bollerslev, Frank Diebold, Peter Reinhard Hansen, Nour Meddahi and Per Mykland for their comments on an earlier draft. It is particularly important to note in this case, that I, not them, are fully responsible for the views recorded here.

After this most intense period of work on financial econometrics largely finished around 2008, Ole further developed some of the methods which saw their birth during that decade into new areas, inventing Ambit processes and extending supOU processes as he continued to research and write in the last decade of his life. Some of the work on Ambit processes connected to mathematical finance, but Almut Veraart will be discussing that topic so I will not review it here. A pointer to that work is Barndorff-Nielsen et al. (2013); Barndorff-Nielsen and Veraart (2013); Barndorff-Nielsen et al. (2015a,b, 2010), while a book length review is given in Barndorff-Nielsen et al. (2018).

I met Ole's ideas first on paper, perhaps in 1992 or 1993. By the end of our work together we had jointly coauthored 29 papers. These papers will form the bulk of the work discussed here. Around 10 of those papers are very highly cited. It is obviously impossible for me to be entirely objective about this work, but I will try to place it in context the best I can.

When I first read Ole's work I was at Nuffield College, Oxford. The College was headed by Sir David R. Cox and David and Ole were working on Barndorff-Nielsen and Cox (1994), their second book on asymptotics in statistics. David kindly let me read the manuscript and I gave him some small comments. Later David and Ole organized a Séminaires Européens de Statistiques meeting at Nuffield College on statistical problems in the social sciences and they asked me to give some lectures on financial econometrics (they were published in Shephard (1996)). Ole and his students read my paper closely and gave me considerable comments. The topic was largely on stochastic volatility, which has a rich and complicated history (e.g. Chapter 1 of Shephard (2005)). I think he saw that he could connect how I thought about stochastic volatility with his earlier work on generalized hyperbolic distributions and infinite divisibility, which gave him an in on the field.

2 Lévy based model building

Ole and I started talking about modeling. He would suggest various models and I would trim them and vice versa. We were trying to find models and methods which allow calculations and computations, but would also be empirically appealing.

During this time he produced his first finance related papers, Barndorff-Nielsen (1997, 1998b,a); Barndorff-Nielsen and Prause (2001) (see also Barndorff-Nielsen and Shephard (2001a)), which suggested the normal inverse Gaussian Lévy process as a model for log-prices of speculative assets (which connected back to his earlier work introducing the generalized hyperbolic distribution and on infinite divisibility). This model had the advantage that it could allow quite heavy tails in financial returns and was easy to do calculations with items like aggregational Gaussianity, as well as fit to data (see Cont and Tankov (2004) for a broader development). It had the important drawback that it yielded i.i.d. returns, which is very unrealistic empirically. Related to that work is Barndorff-Nielsen et al. (2012) and Barndorff-Nielsen et al. (2014), which used Lévy processes to build Skellam type models of prices which lived on lattice structures, which is important for very high frequency data.

It took a long time but finally Ole and I had a working paper, Barndorff-Nielsen and Shephard (1998),

which included the so-called Barndorff-Nielsen and Shephard (denoted the BNS model) stochastic volatility model, which was eventually published as Barndorff-Nielsen and Shephard (2001b), a read paper for the Royal Statistical Society. The basic version of this paper had

$$Y(t) = Y(0) + \int_0^t \mu(u)du + \int_0^t \sigma^2(u)dB(u), \quad d\sigma^2(t) = -\lambda\sigma^2(t)dt + dZ(\lambda t), \quad \lambda > 0, \quad t \geq 0,$$

where $\{B(t)\}_{t \geq 0}$ is a Brownian motion, independent from $\{Z(t)\}_{t \geq 0}$ a subordinator (note that an extension of this model to allow for statistical leverage was given but is not discussed here). In finance $\{\sigma(t)\}_{t \geq 0}$ is called the spot volatility process, following that area's tradition to call standard deviation objects volatility. The BNS model had a volatility process which revolved around non-Gaussian OU processes. These OU processes are based on self-decomposability of distributions on the positive half line, which allow for easy calculations as well as interesting math. The paper had a substantial amount of additional material on a superposition of volatility processes, long memory and quadratic variation. The simplest of the BNS models are a special case of the affine models associated with Duffie et al. (2003). The superposition version is important empirically as volatility is not well represented by a Markov process. Instead it approximately has, at least, a large component which is extremely fast mean reverting and a component which has a great deal of memory. In the limit as the number of components goes off to infinity, the superposition model can allow for long memory (a continuous time version on the positive half line of the observation of Granger (1980)) which is empirically important in volatility modeling (e.g. Baillie et al. (1996), Andersen et al. (2003)).

Ole and I did some follow up work on these Lévy based models after this, including Barndorff-Nielsen and Shephard (2003a, 2001a); Barndorff-Nielsen et al. (2002); Barndorff-Nielsen and Shephard (2003b, 2001c), but most of the work carried out on this type of model was written by others. For example, an important early paper was by Nicolato and Venardos (2003) who established an option pricing formula for this model.

Ole went on to establish a rigorous version of the infinite superposition model, which was called the supOU process, in Barndorff-Nielsen (2001), as well as provide a book length treatment of time-change in Barndorff-Nielsen and Shiryaev (2010). For certain choices of the mixing distribution his supOU model yields a continuous time long memory process. This model has a substantial life in the probability literature, e.g. Grahovac et al. (2019, 2020, 2021, 2018). He also developed a multivariate non-Gaussian OU processes and a corresponding stochastic volatility process in Barndorff-Nielsen and Stelzer (2013).

3 High frequency financial econometrics

In terms of finance econometrics, Ole's most impactful papers have been about the pioneering harnessing of high frequency financial data to learn economically interesting objects. I view Ole as founding this subfield, together with two other Danes who received their undergraduate education at Ole's home institution, Aarhus University, Torben Andersen and Tim Bollerslev, and two non-Danes, Francis Diebold and myself (no doubt others could potentially have a different view of this). This subfield is perhaps the second important intellectual

wave of financial econometrics¹, after the initial work on ARCH models by Engle (1982) and Bollerslev (1986) (a review on ARCH type models is given in Bollerslev et al. (1994)). I think it is the most distinctive technical contribution financial econometrics has made to economics. Of course many other brilliant scholars contributed to amplifying that wave, some of their work will be briefly discussed below. I hope these authors will forgive me for focusing here mostly on the work that Ole was involved in.

The harnessing of high frequency data started with what is called realized volatility

$$\hat{\sigma}_i = \sqrt{\sum_{j=1}^n \{Y(i-1+j/n) - Y(i-1+(j-1)/n)\}^2}, \quad i = 1, 2, \dots$$

which used n high frequency returns on the i -th day. The corresponding $\hat{\sigma}_i^2$ is called a realized variance. The idea was to publish the time series of daily volatility measures $\hat{\sigma}_1, \hat{\sigma}_2, \dots$, just as inflation is published every month. At its most basic $\hat{\sigma}_i$ is the square root of the sum of squared financial returns, basically the sample deviation of returns but without dividing by the same size (as returns get smaller when the sample size increases) and ignoring the sample mean (as it is tiny). Given its simplicity, it is not surprising that $\hat{\sigma}_i$ has a very long history in applied finance, before it was formalized to allow time-variation. It is routinely used in financial practice. Early academic references include Rosenberg (1972), Officer (1973), Merton (1980) and Schwert (1989). More formal work came with Drost and Nijman (1993), Nelson and Foster (1994) and Andersen and Bollerslev (1998a), which are in keeping with some of the revolution which was about to appear but focused on using the data to learn about ARCH type models or the (assumed) Markovian spot volatility process $\{\sigma(t)\}_{t \geq 0}$. In passing, one should remark that in the 1990s the financial econometrics community was being inspired by the generosity of Olsen and Associates who made available to academics relatively high frequency foreign exchange data (e.g. Dacorogna et al. (1998)). It was their data which attracted me to finance problems in around 1991, following a question about diurnality posed to me by Charles Goodhart (see Goodhart and Payne (2000)).

The initial high frequency econometrics revolution was really about changing the estimand, which I will write as σ_i . Of course in statistics changing the estimand changes everything! That change was made in concurrent work by Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002a).

To make calculations feasible, that selected estimand had to be placed in the context of an interesting model. The model the above two sets of researchers fixed on was a nonparametric stochastic volatility model (or more broadly semimartingales), which they both argued should be analyzed using an in-fill asymptotic approximation. The choice to use nonparametric models was a dramatic second break from the earlier financial econometric literatures, which used ARCH type models or parameterized discrete or continuous time SV models. The use of the semiparametric models also directly linked in with the mathematical finance literature on models which enforce non-arbitrage (this point appears prominently in Andersen et al. (2001)).

Instead of thinking in terms of ARCH type models, Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002a) used high frequency data, e.g. returns measured every 5 minutes, to estimate particular

¹Some empirical evidence for this claim is to found by looking at the “Journal of Financial Econometrics”, which started publishing in 2003. By May 2024, three of the journal’s top four google cited papers are in the subfield of high frequency financial econometrics (Corsi (2009), Barndorff-Nielsen and Shephard (2004b), and Barndorff-Nielsen and Shephard (2006a)).

increments of quadratic variation, where the increments were over an interval like a day or week (taking out the diurnal feature)

$$\sigma_i^2 = [Y, Y]_i - [Y, Y]_{i-1} = \int_{i-1}^i \sigma^2(u) du,$$

daily integrate variance, while σ_i is called the integrated volatility on the i -th day. This estimand σ_i , which changes each day, was interesting from a practical, empirical and theoretical viewpoint.

High frequency data is quite complicated, with extremely strong diurnal features in the volatility (e.g. Andersen and Bollerslev (1997, 1998b,c)) with most of the trading in many markets happening near the opening or closing of the market or around regularly scheduled news arrivals (e.g. monetary announcements by Governments or economic data being published by statistical agencies). As a result (and for many other reasons!) trade by trade modeling is hard and additionally, of course, time-varying. Work in this area includes Engle and Russell (1998), Hasbrouck (1991), Rydberg and Shephard (2003), Bouchaud et al. (2002), Potters and Bouchaud (2003) and the reviews by Bauwens and Hautsch (2009), Russell and Engle (2010) and Hautsch (2012). Of course, for high frequency traders this is essential work, but for academic economists it is not a mainstream goal (an exception is the subfield of “price discovery”, where the causal impact of trading and quoting on future prices is the scientific focus). Hence it is attractive to find an estimand which is not sensitive to diurnal features. Looking at a day of volatility does this, just as traditional inflation measures are defined as yearly differences of consumer price indexes, rather than computed through seasonal, trend and noise models of the underlying consumer price level estimated from monthly consumer price data. The consumer price level data is highly seasonal, just like financial volatility data. It was this inflation analogy which made Ole and I focus on integrated volatility over a day as a useful estimand for financial problems.

Andersen et al. (2001, 2003) were interested in using a history of, say, daily realized volatilities to forecast future, say, daily volatility in the context of a semimartingale (no ARCH model in sight). The empirical evidence suggested that forecasts based on realized volatilities do indeed tend to strongly out perform, especially at short time lags, volatility forecasts compared to ARCH type models, e.g. Corsi (2009).

Barndorff-Nielsen and Shephard (2002a) established a theory of volatility measurement, building a mixed Gaussian central limit theory, finding an estimator of the corresponding standard error of the estimator and implemented confidence intervals using the high frequency financial data. They showed that the realized volatility estimates of σ_i were quite noisy in practice, which was a surprise to many applied researchers. The central limit theorem developed in Barndorff-Nielsen and Shephard (2002a) could have been backed out from a very general and beautiful result on the numerical discretization of stochastic differential equations contained within Jacod and Protter (1998) while the 1994 unpublished working paper of Jacod (2018) had deep, broad and related results which would have helped them, but we did not know the existence of those papers at the time. Some of the probability theory around this area is wonderfully developed and reviewed in Jacod and Protter (2012).

Many fantastic scholars have made contributions to this area after this early breach. They have relaxed assumptions, generalized results, applied them extensively, progressed with new estimands and estimators based on in-fill high frequency arguments and critically evaluated the suggestions of others. Particular work outside

the scope of Ole's work is on high frequency derivative based data (e.g. Viktor Todorov), robust inference in the presence of sophisticated financial noise (e.g. Per Mykland and Lan Zhang), bootstrapping (e.g. Prosper Dovonon, Silvia Goncalves, Ulrich Hounyo and Nour Meddahi), forecasting using high frequency based financial statistics (e.g. Fulvio Corsi and Andrew Patton), managing high dimensional, high frequency financial data (e.g. Yacine Ait-Sahalia, Jianqing Fan, Markus Reiss and Dacheng Xiu), estimating the Blumenthal-Gettoor index of financial jumps (e.g. Yacine Ait-Sahalia and Jean Jacod) and rough volatility (e.g. Jim Gatheral). But there are also powerful stochastic analysts who have importantly generalized results from simple stochastic volatility plus finite activity jump models to semimartingales and then to semimartingales with complicated statistical models of noise (e.g. Jean Jacod, Mark Podolskij and Mathieu Rosenbaum). They have deepened the technical tools available in the financial econometrics field. I hope readers will forgive me in mostly focusing on the high frequency work Ole was involved in. Some additional work on realized volatility he carried out is discussed in Barndorff-Nielsen and Shephard (2002b, 2005a, 2007, 2010, 2006b).

Barndorff-Nielsen and Shephard (2004a) extended the realized volatility work to the multivariate case. Most asset allocation and risk management problems are multivariate so pushing in this direction is practically important. That paper provided the first formal result on using high frequency returns to make nonparametric inference on betas in finance for thickly traded stocks. Later we will discuss the empirically important literature which extends this work to deal with the irregularly spaced data caused by non-synchronous trading in assets. That work was carried out by Ole and me with Aarhus colleague Asger Lunde and yet another Dane, Peter Reinhard Hansen. Ole was understandably very proud of all the brilliant Danes working in this field.

The realized betas of Barndorff-Nielsen and Shephard (2004a) have been very influential work in the field, see also Andersen et al. (2006). Recent inventive work includes Christensen et al. (2023) which extends the realized covariance matrices to being L_1 penalized over large dimensions as well as more traditional high dimensional objects adapted from realized covariances, such as principle components and factor models, e.g. Ait-Sahalia and Xiu (2017), Fan and Kim (2018).

After about 2000, Ole was interested in extending the work on realized volatility to scaled sums of powers of absolute high frequency returns. This is called power variation. He kept on mentioning it in our conversations. There seemed substantial empirical evidence that powers of one or lower were computationally more stable through time and less sensitive to unusual datapoints than realized variance. The difficulty was this leads to estimators of $\int_{i-1}^i \sigma^r(u)du$, where the r -th power of absolute returns is used. That estimand is not very interesting in practice to economists. But it was clear that such estimators would be robust to finite activity jumps, a result which turns out goes back to Lépingle (1976) in stochastic analysis. During a visit by Ole to Oxford, I side stepped the estimand problem by suggesting

$$\sum_{j=2}^n |Y(i-1+j/n) - Y(i-1+(j-1)/n)| |Y(i-1+(j-1)/n) - Y(i-1+(j-2)/n)|,$$

which, when scaled, estimates

$$\sigma_i^2 = [Y^c, Y^c]_i - [Y^c, Y^c]_{i-1} = \int_{i-1}^i \sigma^2(u)du$$

integrated variance σ_t^2 in the presence of jumps — which means quadratic variation can be empirically split into the component due to the continuous evolution of prices and that due of jumps. This bipower statistic is a self-tuning nonparametric estimator: it has no bandwidths for empirical workers to disagree or worry about.

Initial results were published in Barndorff-Nielsen and Shephard (2004b), while the first nonparametric test for jumps in finance was reported in Barndorff-Nielsen and Shephard (2006a), which included extending the analysis to a multipower statistic. More substantial theory was reported in Barndorff-Nielsen and Shephard (2005c); Barndorff-Nielsen et al. (2006). A local version of the bipower statistic was suggested by Lee and Mykland (2008). This bipower statistic is not statistically efficient, so some researchers understandably prefer to use the bipower statistic as a preliminary estimator, employed to select the continuous time bandwidth (which depends on the spot volatility) of the important threshold estimator developed by Mancini (2001, 2004), which is based on the modulus of continuity of Brownian motion $|B(t+h) - B(t)| \leq c\sqrt{h \log(1/h)}$. The Mancini threshold estimator is asymptotically efficient and very influential in practice.

Jumps are interesting economically for a number of reasons. In terms of model building, they push researchers to models outside those based on Brownian motion. They add challenges to hedging risky positions. Large jumps are interesting as they suggest shocks to the market at particular times, e.g. such jumps can be potentially used to estimate the causal impact of the Federal Reserve announcements on monetary policy.

The probabilistic aspects of multipower variation was quite stimulating, leading to a number of more theoretical developments, extending the processes under which limit theory could be established and generalizing the types of statistics which could be studied. Example of this work includes Barndorff-Nielsen and Shephard (2005b), Barndorff-Nielsen et al. (2004), Barndorff-Nielsen et al. (2006) and Barndorff-Nielsen et al. (2006). It was a thrill at the time to work with such amazing probabilists pushing new results in econometrics. I do not know another subfield of econometrics which has had a period which has attracted such folk to its challenges.

Later some of this work on multipower variation was extended by Ole and others to additional classes of continuous time processes, such as processes with stationary increments, e.g. Barndorff-Nielsen et al. (2009) and Barndorff-Nielsen et al. (2009, 2011, 2013).

There is a long history in empirical finance of noting the important difference between up and down price moves on codependence and serial dependence, e.g. Roy (1952), Sortino and van der Meer (1991), Sortino and Satchell (2001) and Ang et al. (2006). Some of this appears in models of daily data and how negative returns tend to lead to subsequent spikes in volatility, such as Black (1976), Nelson (1991) and Campbell and Hentschel (1992). This is often called statistical leverage.

Barndorff-Nielsen et al. (2010) split up realized variance into the contributions from positive and negative high frequency returns. They labelled this a realized semivariances. From a stochastic analysis viewpoint, this splits up the increments of quadratic variation into two components. For continuous sample path processes these two quantities are the same. But when there are jumps they differ. From a financial perspective this difference is important.

These realized semivariances have been used extensively in forecasting exercises, see Patton and Sheppard

(2015). Extensions of the realized semivariance to the multivariate case has been developed by Bollerslev et al. (2020) into semicovariances, while Bollerslev et al. (2022) develop semibetas. Both seem to have substantial empirical bite.

The initial work on high frequency financial data used prices recorded every 1 or 5 or sometimes 30 minutes. Prices can be impacted by many short term features, such as price discreteness caused by the rules of the exchange, short physical lags due to the relaying of information, or if the trade is initiated by the buyer (which will tend to push the price up) or the seller (which will tend to push the price down). The latter effects will be particularly large for thinly traded stocks. As prices are studied at finer time intervals, the noise (compared to the movements in the signal) will play a bigger role. In the limit, this becomes its own academic area called market microstructure and what we think of as noise becomes the scientific or engineering point of focus.

In statistics, from the broadest viewpoint, asymptotic analysis plays many different roles, but typically it is used as an intellectual device to form an approximation to the distribution of an estimator based on a finite amount of data (few would study asymptotics if we knew the exact distribution of all estimators). In the context of financial econometrics the literal limit of high frequency financial data as we have finer and finer observations turns us into market microstructure analysts based on data living on a lattice recorded at the times of trades and quotes (e.g. the realized variance would become a sum of a large but finite random number of squared random variables, summing up all the tick by tick squared returns). This connects to the point process type work, pioneered by Engle and Russell (1998) and Bowsher (2007). See the reviews by Bauwens and Hautsch (2009), Bacry et al. (2015) and Hawkes (2018).

A different line of asymptotic argument is to formally introduce a (certainly misspecified) statistical noise model, not to take the literal market microstructure limit but allow more data than the original 1 or 5 minutes, e.g. 1 second returns for a thickly traded asset. There is some early work on this in Zhou (1996, 1998), Bandi and Russell (2008), Hansen and Lunde (2006) and the 2007 working paper version of Phillips and Yu (2023). But in the context of nonparametric analysis using high frequency noisy data, three papers stand out: the two scale estimator of Zhang et al. (2005), the preaveraging estimator of Jacod et al. (2009) and the realized kernel of Barndorff-Nielsen et al. (2008). The realized kernel adds lead and lag effects to the realized variance. All three estimators are quite similar in their core, differing really only in the end effects.

A virtue of these noise robust estimators is that they can consume more data. But, formally, taking into account the noise implies the resulting estimators' rate of convergence is significantly slowed, down from order $n^{1/2}$ to typically order $n^{1/6}$ (careful design of the methods can sometimes yield a $n^{1/4}$ rate and indeed Barndorff-Nielsen et al. (2008) showed the realized kernel can be designed to be as efficient as the Gaussian MLE for this problem: the same result will very likely hold for versions of the multiscale estimator and the preaverage estimator). Hence inference based on these estimators is not particularly precise, even if based on large datasets. Extensions of the realized kernel are given in Barndorff-Nielsen et al. (2009) and Barndorff-Nielsen et al. (2011b).

The preaveraging approach seems the most extendible of the three methods, allowing its use across jump

robust estimators, etc. I have been impressed by the recent paper by Mykland and Zhang (2016) which looks at replacing preaveraging by robust estimators. This seems a practical method which can be widely applied, even possibly automatically in industrial applications.

One of the interesting empirical discoveries reported in Barndorff-Nielsen et al. (2009) was that the realized kernel was sometimes thrown off by what appeared to be a “gradual jump”, a sustained move in the price in one direction over a significant period of time, e.g. 10 minutes. Gradual jumps have been formalized in recent work by others. One approach is a “drift burst,” where the absolute value of the spot drift seems to dominate the spot volatility. Another approach is “persistent noise.” Empirically, gradual jumps seem quite common in practice, and have been used to model, for example, flash crashes. Good places to read about gradual jumps are Christensen et al. (2022), Andersen et al. (2023) and Mancini (2023). This area is technically interesting: the noise is in calendar time not tick time. It is also economically important.

Perhaps the most important form of noise in high frequency analysis of financial data is that due to trades taking place at different times — what is called non-synchronous trading (or indeed quoting). A large empirical literature on this topic goes back to at least Scholes and Williams (1977). One of the most interesting papers on this topic is Hayashi and Yoshida (2005), which yields the maximum likelihood estimator in the multivariate Brownian motion case, where the times of the trading are viewed as being independent from prices. Barndorff-Nielsen et al. (2011a) extend realized kernels to the multivariate case, allowing for both noise and non-synchronous trading. They apply this to a large dataset with some economic applications to asset allocation.

If you want to learn more about high frequency finance, some nice reviews of the field are given in Andersen and Benzoni (2009), Andersen et al. (2010), Hautsch (2012) and Mykland and Zhang (2012). For more mathematical readers the latter will be a good place to start.

4 Conclusion

Ole Barndorff-Nielsen’s research in financial econometrics has been extremely influential, driving one of the largest revolutions in the subject’s history. He started working in that area at the age of around 60. Over a period of around a dozen years his ideas opened up many new research routes that modern researchers have followed. His work was remarkable.

References

- Ait-Sahalia, Y. and D. Xiu (2017). Using principal component analysis to estimate a high dimensional factor model with high-frequency data. *Journal of Econometrics* 201, 384—399.
- Andersen, T., Y. Li, V. Todorov, and B. Zhou (2023). Volatility measurement with pockets of extreme return persistence. *Journal of Econometrics* 237, 105048.
- Andersen, T. G. and L. Benzoni (2009). Realized volatility. In T. G. Andersen, R. A. Davis, J. P. Kreiss, and T. Mikosch (Eds.), *Handbook of Financial Time Series*, pp. 555–575. Springer-Verlag.

- Andersen, T. G. and T. Bollerslev (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T. G. and T. Bollerslev (1998a). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G. and T. Bollerslev (1998b). Deutsche mark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance* 53, 219–265.
- Andersen, T. G. and T. Bollerslev (1998c). Towards a unified framework for high and low frequency return volatility modeling. *Statistica Neerlandica* 52, 273–302.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2010). Parametric and nonparametric measurement of volatility. In Y. Aït-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics*, pp. 67–138. Amsterdam: North Holland.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2001). The distribution of exchange rate volatility. *Journal of the American Statistical Association* 96, 42–55.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and G. Wu (2006). Realized beta: persistence and predictability. *Advances in Econometrics* 20, 1–39.
- Ang, A., J. Chen, and Y. Xing (2006). Downside risk. *Review of Financial Studies* 19, 1191–1239.
- Bacry, E., I. Mastromatteo, and J. Muzy (2015). Hawkes processes in finance. *Market Microstructure and Liquidity* 1.
- Baillie, R. T., T. Bollerslev, and H. O. Mikkelsen (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 74, 3–30.
- Bandi, F. M. and J. R. Russell (2008). Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies* 75, 339–369.
- Barndorff-Nielsen, O. E. (1979). Models for non-Gaussian variation, with applications to turbulence. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 368, 501–520.
- Barndorff-Nielsen, O. E. (1983). On a formula for the distribution of the maximum likelihood estimator. *Biometrika* 70, 343–65.
- Barndorff-Nielsen, O. E. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics* 24, 1–14.
- Barndorff-Nielsen, O. E. (1998a). Probability and statistics: selfdecomposability, finance and turbulence. In L. Accardi and C. C. Heyde (Eds.), *Probability Towards 2000. Proceedings of a Symposium held 2-5 October 1995 at Columbia University*, pp. 47–57. New York: Springer-Verlag.
- Barndorff-Nielsen, O. E. (1998b). Processes of normal inverse Gaussian type. *Finance and Stochastics* 2, 41–68.
- Barndorff-Nielsen, O. E. (2001). Superposition of Ornstein-Uhlenbeck type processes. *Theory of Probability and its Applications* 46, 175–194.

- Barndorff-Nielsen, O. E., F. E. Benth, and A. Veraart (2013). Modelling electricity futures by ambit fields. *Advances in Applied Probability* 46, 719–745.
- Barndorff-Nielsen, O. E., F. E. Benth, and A. Veraart (2015a). Cross-commodity modelling by multivariate ambit fields. In R. Aid, M. Ludkovski, and R. Sircar (Eds.), *Commodities, Energy and Environmental Finance*, Volume 74, pp. 109–148. Springer.
- Barndorff-Nielsen, O. E., F. E. Benth, and A. Veraart (2018). *Ambit Stochastics*. Heidelberg: Springer.
- Barndorff-Nielsen, O. E., F. E. Benth, and A. E. D. Veraart (2010). Modelling electricity forward markets by ambit fields. CREATES Research Paper 2010–41.
- Barndorff-Nielsen, O. E., F. E. Benth, and A. E. D. Veraart (2015b). Recent advances in ambit stochastics with a view towards tempo-spatial stochastic volatility/intermittency. *Banach Center Publications* 104, 25–60.
- Barndorff-Nielsen, O. E., J. M. Corcuera, and M. Podolskij (2009). Power variation for Gaussian processes with stationary increments. *Stochastics Processes and Their Applications* 119, 1845–1865.
- Barndorff-Nielsen, O. E., J. M. Corcuera, and M. Podolskij (2011). Multipower variation for Brownian semistationary processes. *Bernoulli* 17, 1159–1194.
- Barndorff-Nielsen, O. E., J. M. Corcuera, and M. Podolskij (2013). Limit theorems for functionals of higher order differences of Brownian semistationary processes. In A. Shiryaev, S. R. S. Varadhan, and E. Presman (Eds.), *Prokhorov and Contemporary Probability Theory*, Volume 33. Berlin: Springer.
- Barndorff-Nielsen, O. E., J. M. Corcuera, M. Podolskij, and J. Worner (2009). Bipower variation for Gaussian processes with stationary increments. *Journal of Applied Probability* 46, 132–150.
- Barndorff-Nielsen, O. E. and D. R. Cox (1979). Edgeworth and saddlepoint approximations with statistical applications. *Journal of the Royal Statistical Society, Series B* 41, 279–299.
- Barndorff-Nielsen, O. E. and D. R. Cox (1989). *Asymptotic Techniques for use in Statistics*. London: Chapman & Hall.
- Barndorff-Nielsen, O. E. and D. R. Cox (1994). *Inference and Asymptotics*. London: Chapman & Hall.
- Barndorff-Nielsen, O. E., S. E. Graversen, J. Jacod, M. Podolskij, and N. Shephard (2006). A central limit theorem for realised power and bipower variations of continuous semimartingales. In Y. Kabanov, R. Lipster, and J. Stoyanov (Eds.), *From Stochastic Analysis to Mathematical Finance, Festschrift for Albert Shiryaev*, pp. 33–68. Springer.
- Barndorff-Nielsen, O. E., S. E. Graversen, J. Jacod, and N. Shephard (2006). Limit theorems for realised bipower variation in econometrics. *Econometric Theory* 22, 677–719.
- Barndorff-Nielsen, O. E., S. E. Graversen, and N. Shephard (2004). Power variation and stochastic volatility: a review and some new results. *Journal of Applied Probability* 41A, 133–143.
- Barndorff-Nielsen, O. E. and C. Halgreen (1977). Infinite divisibility of the hyperbolic and generalised inverse Gaussian distributions. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 38, 309–311.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica* 76, 1481–1536.

- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009). Realised kernels in practice: trades and quotes. *Econometrics Journal* 12, C1–C32.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2011a). Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics* 162, 149–169.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2011b). Subsampling realised kernels. *Journal of Econometrics* 160, 204–219.
- Barndorff-Nielsen, O. E., S. Kinnebrouck, and N. Shephard (2010). Measuring downside risk: realised semivariance. In T. Bollerslev, J. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, pp. 117–136. Oxford University Press.
- Barndorff-Nielsen, O. E., A. Lunde, N. Shephard, and A. E. D. Veraart (2014). Integer-valued trawl processes: A class of stationary infinitely divisible processes. *Scandinavian Journal of Statistics* 41, 693–724.
- Barndorff-Nielsen, O. E., E. Nicolato, and N. Shephard (2002). Some recent developments in stochastic volatility modelling. *Quantitative Finance* 2, 11–23.
- Barndorff-Nielsen, O. E., D. G. Pollard, and N. Shephard (2012). Integer-valued Lévy processes and low latency financial econometrics. *Quantitative Finance* 12, 587–605.
- Barndorff-Nielsen, O. E. and K. Prause (2001). Apparent scaling. *Finance and Stochastics* 5, 103–113.
- Barndorff-Nielsen, O. E. and N. Shephard (1998). Aggregation and model construction for volatility models. Unpublished discussion paper: Nuffield College, Oxford.
- Barndorff-Nielsen, O. E. and N. Shephard (2001a). Modelling by Lévy processes for financial econometrics. In O. E. Barndorff-Nielsen, T. Mikosch, and S. Resnick (Eds.), *Lévy Processes – Theory and Applications*, pp. 283–318. Boston: Birkhäuser.
- Barndorff-Nielsen, O. E. and N. Shephard (2001b). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics (with discussion). *Journal of the Royal Statistical Society, Series B* 63, 167–241.
- Barndorff-Nielsen, O. E. and N. Shephard (2001c). Normal modified stable processes. *Theory of Probability and Mathematical Statistics* 65, 1–19.
- Barndorff-Nielsen, O. E. and N. Shephard (2002a). Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B* 64, 253–280.
- Barndorff-Nielsen, O. E. and N. Shephard (2002b). Estimating quadratic variation using realised variance. *Journal of Applied Econometrics* 17, 457–477.
- Barndorff-Nielsen, O. E. and N. Shephard (2003a). Integrated OU processes and non-Gaussian OU-based stochastic volatility. *Scandinavian Journal of Statistics* 30, 277–295.
- Barndorff-Nielsen, O. E. and N. Shephard (2003b). Realised power variation and stochastic volatility. *Bernoulli* 9, 243–265. Correction published in pages 1109–1111.
- Barndorff-Nielsen, O. E. and N. Shephard (2004a). Econometric analysis of realised covariation: high frequency covariance, regression and correlation in financial economics. *Econometrica* 72, 885–925.

- Barndorff-Nielsen, O. E. and N. Shephard (2004b). Power and bipower variation with stochastic volatility and jumps (with discussion). *Journal of Financial Econometrics* 2, 1–48.
- Barndorff-Nielsen, O. E. and N. Shephard (2005a). How accurate is the asymptotic approximation to the distribution of realised volatility? In D. W. K. Andrews and J. H. Stock (Eds.), *Identification and Inference for Econometric Models. A Festschrift in Honour of T.J. Rothenberg*, pp. 306–331. Cambridge: Cambridge University Press.
- Barndorff-Nielsen, O. E. and N. Shephard (2005b). Multipower variation and stochastic volatility. In A. Shiryaev, M. Grossinho, P. Oliveira, and M. Esquivel (Eds.), *Stochastic Finance*, pp. 73–82. Springer.
- Barndorff-Nielsen, O. E. and N. Shephard (2005c). Power variation and time change. *Theory of Probability and Its Applications* 50, 1–15.
- Barndorff-Nielsen, O. E. and N. Shephard (2006a). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4, 1–30.
- Barndorff-Nielsen, O. E. and N. Shephard (2006b). Impact of jumps on returns and realised variances: econometric analysis of time-deformed Lévy processes. *Journal of Econometrics* 131, 217–252.
- Barndorff-Nielsen, O. E. and N. Shephard (2007). Variation, jumps and high frequency data in financial econometrics. In R. Blundell, T. Persson, and W. K. Newey (Eds.), *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress*, Econometric Society Monographs, pp. 328–372. Cambridge University Press.
- Barndorff-Nielsen, O. E. and N. Shephard (2010). Volatility. In R. Cont (Ed.), *Encyclopedia of Quantitative Finance*, Volume IV, pp. 1898–1901. Chichester: Wiley.
- Barndorff-Nielsen, O. E., N. Shephard, and M. Winkel (2006). Limit theorems for multipower variation in the presence of jumps. *Stochastic Processes and Their Applications* 116, 796–806.
- Barndorff-Nielsen, O. E. and A. N. Shiryaev (2010). *Change of Time and Change of Measure*. Singapore: World Scientific.
- Barndorff-Nielsen, O. E. and R. Stelzer (2013). The multivariate supOU stochastic volatility model. *Mathematical Finance* 23, 275–296.
- Barndorff-Nielsen, O. E. and A. Veraart (2013). Stochastic volatility of volatility and variance risk premia. *Journal of Financial Econometrics* 11, 1–46.
- Bauwens, L. and N. Hautsch (2009). Modelling financial high frequency data using point processes. In T. G. Andersen, R. A. Davis, J. P. Kreiss, and T. Mikosch (Eds.), *Handbook of Financial Time Series*, pp. 953–979. Springer-Verlag.
- Black, F. (1976). Studies of stock price volatility changes. pp. 177–181.
- Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics* 51, 307–327.
- Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). ARCH models. In R. F. Engle and D. McFadden (Eds.), *The Handbook of Econometrics, Volume 4*, pp. 2959–3038. Amsterdam: North-Holland.
- Bollerslev, T., A. Patton, J. Li, and R. Quaedvlieg (2020). Realized semicovariances. *Econometrica* 88, 1515–1551.

- Bollerslev, T., A. J. Patton, and R. Quaadvleig (2022). Realized semibetas: Disentangling “good” and “bad” downside risks. *Journal of Financial Economics* 144, 227–246.
- Bouchaud, J.-P., M. Mezard, and M. Potters (2002). Statistical properties of stock order books: empirical results and models. *Quantitative Finance* 2, 251–256.
- Bowsher, C. G. (2007). Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics* 141, 876–912.
- Campbell, J. Y. and L. Hentschel (1992). No news is good news: an asymmetric model of changing volatility in stock returns. *Journal of Financial Economics* 31, 281–318.
- Christensen, K., M. S. Nielsen, and M. Podolskij (2023). High-dimensional estimation of quadratic variation based on penalized realized variance. *Statistical Inference for Stochastic Processes* 26, 331–359.
- Christensen, K., R. C. Oomen, and R. Reno (2022). The drift burst hypothesis. *Journal of Econometrics* 227, 461–497.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. London: Chapman and Hall.
- Corsi, F. (2009). A simple long memory model of realized volatility. *Journal of Financial Econometrics* 7, 174–196.
- Dacorogna, M. M., U. A. Müller, R. B. Olsen, and O. V. Pictet (1998). Modelling short term volatility with GARCH and HAR. In C. Dunis and B. Zhou (Eds.), *Nonlinear Modelling of High Frequency Financial Time Series*, pp. 161–176. Chichester: Wiley.
- Drost, F. C. and T. E. Nijman (1993). Temporal aggregation of GARCH processes. *Econometrica* 61, 909–927.
- Duffie, D., D. Filipovic, and W. Schachermayer (2003). Affine processes and applications in finance. *Annals of Applied Probability* 13, 984–1053.
- Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica* 50, 987–1007.
- Engle, R. F. and J. R. Russell (1998). Forecasting transaction rates: The autoregressive conditional duration model. *Econometrica* 66, 1127–1162.
- Fan, J. and D. Kim (2018). Robust high-dimensional volatility matrix estimation for high-frequency factor model. *Journal of the American Statistical Association* 113, 1268–1283.
- Goodhart, C. A. E. and R. Payne (2000). *The Foreign Exchange Market: Empirical Studies with High-Frequency Data*. Palgrave Macmillan.
- Grahovac, D., N. N. Leonenko, and M. S. Taqqu (2018). Intermittency of trawl processes. *Statistics and Probability Letters* 137, 235–242.
- Grahovac, D., N. N. Leonenko, and M. S. Taqqu (2019). Limit theorems, scaling of moments and intermittency for integrated finite variance supOU processes. *Stochastic Processes and their Applications* 129, 5113–5150.
- Grahovac, D., N. N. Leonenko, and M. S. Taqqu (2020). The multifaceted behavior of integrated supOU processes: The infinite variance case. *Journal of Theoretical Probability* 33, 1801–1831.
- Grahovac, D., N. N. Leonenko, and M. S. Taqqu (2021). Intermittency and infinite variance: the case of integrated supOU processes. *Electronic Journal of Probability* 26, 1–31.

- Granger, C. W. J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 14, 227–238.
- Hansen, P. R. and A. Lunde (2006). Realized variance and market microstructure noise (with discussion). *Journal of Business and Economic Statistics* 24, 127–218.
- Hasbrouck, J. (1991). The summary informativeness of stock trades: An econometric analysis. *The Review of Financial Studies* 4, 571–595.
- Hautsch, N. (2012). *Econometrics of Financial High-Frequency Data*. Springer.
- Hawkes, A. G. (2018). Hawkes processes and their applications to finance: a review. *Quantitative Finance* 18, 193–198.
- Hayashi, T. and N. Yoshida (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli* 11, 359–379.
- Jacod, J. (2018). Limit of random measures associated with the increments of a Brownian semimartingale. *Econometric Theory* 16, 526–569. Preprint number 120, Laboratoire de Probabilités, Université Pierre et Marie Curie, Paris, 1994.
- Jacod, J., Y. Li, P. A. Mykland, M. Podolskij, and M. Vetter (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and Their Applications* 119, 2249–2276.
- Jacod, J. and P. Protter (1998). Asymptotic error distributions for the Euler method for stochastic differential equations. *Annals of Probability* 26, 267–307.
- Jacod, J. and P. Protter (2012). *Discretization of Processes*. Heidelberg: Springer.
- Lee, S. and P. A. Mykland (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies* 21, 2535–2563.
- Lépingle, D. (1976). La variation d’ordre p des semi-martingales. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 36, 295–316.
- Mancini, C. (2001). Disentangling the jumps of the diffusion in a geometric Brownian motion. *Giornale dell’Istituto Italiano degli Attuari* LXIV, 19–47.
- Mancini, C. (2004). Estimation of the characteristics of jump of a general Poisson-diffusion process. *Scandinavian Actuarial Journal* 1, 42–52.
- Mancini, C. (2023). Drift burst test statistic in the presence of infinite variation jumps. *Stochastic Processes and their Applications* 163, 535–591.
- Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8, 323–361.
- Mykland, P. and L. Zhang (2016). Between data cleaning and inference: Pre-averaging and robust estimators of the efficient price. *Journal of Econometrics* 194, 242–262.
- Mykland, P. A. and L. Zhang (2012). The econometrics of high frequency data. In M. Kessler, A. Lindner, and M. Sørensen (Eds.), *Statistical Methods for Stochastic Differential Equations*, pp. 109–190. Chapman & Hall/CRC Press.

- Nelson, D. B. (1991). Conditional heteroskedasticity in asset pricing: a new approach. *Econometrica* 59, 347–370.
- Nelson, D. B. and D. P. Foster (1994). Asymptotic filtering theory for univariate ARCH models. *Econometrica* 62, 1–41.
- Nicolato, E. and E. Venardos (2003). Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type. *Mathematical Finance* 13, 445–466.
- Officer, R. R. (1973). The variability of the market factor of the New York stock exchange. *Journal of Business* 46, 434–453.
- Patton, A. J. and K. K. Sheppard (2015). Good volatility, bad volatility: signed jumps and persistence of volatility. *Review of Economics and Statistics* 97, 683–697.
- Phillips, P. C. B. and J. Yu (2023). Information loss in volatility measurement with flat price trading. *Empirical Economics* 64, 2957–2999.
- Potters, M. and J.-P. Bouchaud (2003). More statistical properties of order books and price impact. *Physics A: Stat. Mech. Appl.* 324, 133–140.
- Rosenberg, B. (1972). The behaviour of random variables with nonstationary variance and the distribution of security prices. Working paper 11, Graduate School of Business Administration, University of California, Berkeley. Reprinted in Shephard (2005).
- Roy, A. D. (1952). Safety first and the holding of assets. *Econometrica* 20, 341–449.
- Russell, J. R. and R. F. Engle (2010). Analysis of high-frequency data. In Y. Ait-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics: Tools and techniques*, pp. 383–426. Elsevier.
- Rydberg, T. H. and N. Shephard (2003). Dynamics of trade-by-trade price movements: Decomposition and models. *Journal of Financial Econometrics* 1, 2–25.
- Scholes, M. and J. Williams (1977). Estimating betas from nonsynchronous trading. *Journal of Financial Economics* 5, 309–27.
- Schwert, G. W. (1989). Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In D. R. Cox, D. V. Hinkley, and O. E. Barndorff-Nielsen (Eds.), *Time Series Models in Econometrics, Finance and Other Fields*, pp. 1–67. London: Chapman & Hall.
- Shephard, N. (2005). *Stochastic Volatility: Selected Readings*. Oxford: Oxford University Press.
- Sortino, F. and S. E. Satchell (2001). *Managing Downside Risk in Financial Markets*. Butterworth-Heinemann.
- Sortino, F. A. and R. van der Meer (1991). Downside risk. *The Journal of Portfolio Management* 17, 27–31.
- Zhang, L., P. A. Mykland, and Y. Ait-Sahalia (2005). A tale of two time scales: determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* 100, 1394–1411.
- Zhou, B. (1996). High-frequency data and volatility in foreign-exchange rates. *Journal of Business and Economic Statistics* 14, 45–52.
- Zhou, B. (1998). Parametric and nonparametric volatility measurement. In C. L. Dunis and B. Zhou (Eds.), *Nonlinear Modelling of High Frequency Financial Time Series*, Chapter 6, pp. 109–123. New York: John Wiley Sons Ltd.