Chapter 3
High-Frequency Data and Financial Econometrics

Jan Novotny and Giovanni Urga

Abstract. This chapter surveys the evolution of financial econometrics from the traditional return-based paradigm to the modern era of high-frequency data. We review the foundations of financial econometrics, including market efficiency, stylised facts, and the development of ARCH/GARCH and stochastic-volatility models, before examining how the electronification of markets transformed empirical analysis through transaction, quote, and limit-order-book data. We discuss realised measures of volatility and their continuous-time foundations, market microstructure, noise-robust estimation, jump detection, multivariate dependence, forecasting, price discovery, and applications to monetary policy identification, systemic risk, machine learning, and cryptocurrency markets. Throughout, we emphasise how high-frequency data shifted the field from modelling latent quantities towards direct measurement, while integrating financial economics, probability, and econometrics into a unified empirical framework. The chapter concludes with a retrospective assessment of which methodological contributions have proved most durable (have aged well), arguing that the evolution of financial econometrics is characterised by a progressive shift from specification-based modelling of latent variables towards measurement-based inference, robust identification, and data-driven empirical analysis.

Jan Novotny
Centre for Econometric Analysis, Bayes Business School (formerly Cass), London, United Kingdom, e-mail: jan@novotny.one

Giovanni Urga ✉
Centre for Econometric Analysis and Faculty of Finance, Bayes Business School (formerly Cass), London, United Kingdom, e-mail: g.urga@citystgeorges.ac.uk

3.1 Introduction - From Low-Frequency to High-Frequency Econometrics

3.1.1 Overview

This chapter provides an integrated overview of financial econometrics and high-frequency econometrics, documenting how the availability of intraday trading data enriched the empirical analysis of asset prices, volatility, dependence, and market efficiency. The chapter begins by positioning high-frequency econometrics within the broader evolution of financial econometrics, outlining how early return-based modelling and stylised facts (fat tails, volatility clustering, leverage effects) shaped the development of benchmark econometric specifications for conditional means and variances. These include ARMA-type models for returns, ARCH/GARCH and asymmetric volatility models, and stochastic volatility frameworks, together with classical (financial) econometric questions surrounding predictability, risk premia, and efficiency. The chapter then highlights the rise of high-frequency data, driven by electronic trading and the availability of tick-by-tick transactions, quotes, and limit-order-book records. This transition motivates a shift from purely parametric modelling to measurement-based econometrics grounded in continuous-time theory. A core emphasis is market microstructure: the trading mechanism simultaneously generates economically informative variables – liquidity, bid-ask spreads, order flow, price impact, and price discovery – and statistical complications such as discreteness, irregular spacing, intraday seasonality, non-synchronous trading, and microstructure noise. Within this setting, the chapter surveys the central methods of high-frequency econometrics. It introduces realised measures of volatility and their interpretation via quadratic variation, discusses noise-robust estimation of integrated variance, and presents jump detection procedures and the decomposition of price dynamics into continuous and discontinuous components. The chapter then expands to multivariate high-frequency econometrics, reviewing realised covariances and correlations under asynchronous sampling, lead-lag effects, and extensions to connectedness and network-based dependence measures. It further highlights how high-frequency information improves forecasting and risk measurement, linking realised volatility and covariation to return-based forecasting models such as HAR-type specifications and to applications in portfolio risk management and systemic risk monitoring. The chapter concludes by revisiting market efficiency and predictability in light of high-frequency evidence, emphasising that predictability is horizon-dependent and tightly linked to liquidity and information flow. Overall, the chapter positions high-frequency econometrics as a methodological pillar of modern financial econometrics, unifying traditional return-based modelling with continuous-time measurement and microstructure-aware inference.

Finally, the chapter takes a deliberately retrospective stance. Beyond surveying the development of models and empirical tools, it asks which contributions of early financial econometrics have *aged well* (remaining central in modern data-rich settings) and which have *aged less well* once confronted with high-frequency evidence and microstructure frictions. This comparison is used to extract a set of *general lessons*: the importance of measurement and identification over specification, the fragility of certain stylised assumptions under temporal aggregation, the centrality of market microstructure to inference, and the need for robustness when moving from low-frequency to high-frequency empirical environments.

3.1.2 From Low-Frequency to High-Frequency Econometrics

The arrival of high-frequency financial data in the late twentieth century reshaped the empirical content of financial econometrics. Where earlier generations of researchers worked with daily, monthly, or annual observations, the gradual electronification of markets opened up tick-by-tick price records, time-stamped order-book messages, and the possibility of observing the price process itself rather than a sparse sample of it. This shift transformed both the questions one could ask and the methods one had to use. Volatility, liquidity, jumps, and information flow, all of which had been treated as latent quantities to be inferred indirectly from low-frequency residuals, became objects of direct measurement. The aim of this chapter is to trace the intellectual and methodological lineage that connects the early stochastic models of asset prices to the data-rich, continuous-time econometrics of the present day, and to ask which of the early ideas have aged well under the weight of subsequent empirical evidence.

The starting point lies more than a century ago. Bachelier (1900), in a doctoral thesis written five years before Einstein's treatment of Brownian motion, introduced the first probabilistic model of speculative prices. Bachelier treated the market as a fair game in which buyers and sellers hold opposing expectations, so the speculator's expected gain at any instant must be zero. From this competitive condition he derived a continuous-time random walk for prices, in which the conditional expectation of the future price coincides with the current price. Modelling small price changes as independent and normally distributed, Bachelier showed that the probability density of the price displacement over an interval is a Gaussian whose spread grows in proportion to the square root of elapsed time – precisely the heat kernel, the fundamental solution of the diffusion equation, derived here independently of physics and governed by a single diffusion coefficient, so that the variance of the displacement accumulates at a constant rate per unit of time. The implied square-root-of-time scaling of price changes underpins essentially all of the subsequent work on integrated and realised volatility that we review in later sections. Bachelier himself went on to apply the framework to options and other speculative instruments and obtained an early closed-form expression for the price of a simple call, but the contribution that lasted is the random-walk model itself rather than the specific pricing formulae.

Bachelier's work lay dormant for half a century. Its rediscovery in the 1950s and 1960s, in part through the edited collection of Cootner (1964), exposed it to empirical scrutiny. Osborne (1959) refined the model by working in logarithms rather than levels, replacing the Brownian motion of prices with the geometric Brownian motion of returns, and produced the first systematic statistical evidence in favour of log-normal returns. Mandelbrot (1963) mounted the most influential dissent. Examining cotton prices over more than a century, including daily data over a shorter span, he documented two features that the Gaussian random walk could not reproduce: tails far heavier than the normal, and a tendency for large moves to cluster in time. As an alternative he proposed the class of *stable Paretian* distributions, a subset of the Lévy-stable family characterised by power-law tails and indexed by a characteristic exponent that lies between zero and two, with the Gaussian recovered only in the special case at the upper limit of two. When the exponent is below two, sums of independent stable Paretian draws scale with the sample size at a faster rate than the square-root scaling of the Gaussian, which implies a self-similarity of returns under temporal aggregation but also infinite population variance and the failure of standard regression inference. Mandelbrot's estimates on cotton data put the characteristic exponent at roughly 1.7, well below two. The stable Paretian model did not become the mainstream parametric workhorse – infinite variance is awkward for portfolio theory – but the empirical critique was robust, and later models from ARCH-type volatility processes to modern stochastic volatility specifications addressed the same heavy-tail and clustering evidence within a finite-variance framework that retained the spirit of Mandelbrot's challenge.

The third strand of the foundation is theoretical. Building on his earlier demonstration that properly anticipated prices fluctuate randomly Samuelson (1965), Samuelson (1973) derived the efficient-markets baseline directly from the internal consistency of competitive markets, without invoking any particular equilibrium asset-pricing model. Writing the discounted price as the raw price scaled by a state-dependent discount factor, and allowing for any discounted intermediate payoff, his central theorem is the martingale relation that the discounted price today equals the expected discounted price next period, inclusive of that payoff. The argument is one of competitive elimination: were the discounted price predictably rising or falling beyond the required return, rational agents would arbitrage the foreseeable change until the opportunity disappeared. The result is a conceptual sharpening of Bachelier. Where Bachelier had derived a random walk from a static, zero-sum balance of buyers and sellers, Samuelson derives the martingale from dynamic optimisation by forward-looking agents. The two routes converge in outcome but differ in motivation, and Samuelson is explicit that the theorem is a statement about price dynamics conditional on the market's information set, not about the quality of that information or the proximity of prices to fundamental values. Subsequent work refined the empirical content of the framework without overturning its theoretical anchor: Lo and MacKinlay (1988), for instance, reject the random-walk hypothesis on weekly US equity data using a variance-ratio test that exploits the fact that, under a random walk, the variance of returns measured over several periods must equal the number of periods multiplied by the variance of one-period returns, and find that the weekly variance ratios lie significantly above unity, indicating positive autocorrelation in short-horizon returns rather than the negative serial dependence that long-horizon mean reversion would imply.

These three threads – Bachelier's diffusion, Mandelbrot's heavy tails, Samuelson's martingale – set the questions that the rest of the chapter addresses. For high-frequency econometrics, each survives in modified form. The diffusion remains the dominant model of the latent efficient-price process; the heavy-tail evidence motivates the non-Gaussian volatility specifications and the jump-diffusion extensions reviewed in Sections 3.5 and 3.7; and the martingale is the long-run benchmark against which the short-horizon predictability induced by microstructure frictions, which we discuss from Section 3.3 onwards, is measured. The remainder of the chapter is organised along the dimensions on which the arrival of high-frequency data extended this foundation: from the empirical regularities of intraday data and the econometrics of trading (Sections 3.3 – 3.4) through the measurement of volatility, microstructure noise, and jumps (Sections 3.5 – 3.7), to multivariate dependence (Section 3.8), risk and forecasting (Sections 3.9 – 3.10), and the reinterpretation of market efficiency in the light of high-frequency evidence (Section 3.11). Sections 3.13–3.17 then survey recent developments and applications, after which a retrospective on which of the foundational contributions have proved durable, including a quantitative complement to the qualitative judgement we develop throughout, is taken up in Sections 3.18 and 3.19. Section 3.20 concludes.

3.2 Financial Econometrics before High-Frequency Data

Before intraday records became available, financial econometrics was built almost entirely on daily, monthly, or annual observations, and this constraint shaped both what could be measured and how it was modelled. Returns were typically treated as independent and identically distributed draws from a single distribution, often Gaussian; volatility, which is not directly observable at these frequencies, was treated as a latent quantity to be inferred indirectly from the dispersion of low-frequency residuals; and market microstructure – the mechanics by which trades generate prices – remained largely a theoretical concern, since the transaction-level data needed to test it empirically did not yet exist. The papers we review in this section established the empirical questions and the statistical tools of this era, and in the process exposed the limitations that the arrival of high-frequency data would later address.

The earliest of these contributions is a methodological warning. Working (1934) showed that a purely random sequence of first differences, once cumulated into levels, produces trajectories with the trends, cycles, and apparent reversals that economists routinely read as economically meaningful. His construction is a random-difference series – a cumulative sum of independent, zero-mean, finite-variance increments, which is simply a random walk in levels. Simulating such series, he found that they reproduced the qualitative features of observed price and production data. Two of his results are striking. First, the population correlation between a deviation of the level from its sample mean and the subsequent first difference converges, regardless of the increment distribution, to approximately minus 0.707 – minus one over the square root of two – as the sample grows, so the apparent "mean reversion" economists had been detecting was an arithmetic artefact of differencing a cumulated series against its own mean. Second, the expected sample autocorrelation of the levels is positive and close to one at short lags and decays only gradually, the same near-unit-root behaviour that the later unit-root literature would formalise. Working's recommendation – to analyse first differences rather than levels, and to compare observed correlations against the random-difference benchmark before attributing them to economic causes – anticipates both the spurious-regression problem and the empirical random-walk programme that Fama (1965) would apply to stock returns three decades later.

That programme reached its fullest early statement in Fama (1965), the most thorough empirical test of the random-walk model of its time. Working with daily closing prices for the thirty Dow Jones Industrial stocks over roughly five years, Fama examined the two components of the hypothesis separately: the independence of successive price changes and the shape of their distribution. On independence, he found that the sample serial correlations of daily log-returns were uniformly small and economically negligible – though a few first-order coefficients were statistically significant in finite samples – and that neither runs tests nor filter rules (buy after a rise of a fixed percentage, sell after a fall of the same percentage) produced returns exceeding a buy-and-hold strategy once transaction costs were accounted for. On the distribution, he documented pronounced leptokurtosis and fitted the stable Paretian family of Mandelbrot (1963) introduced in Section 3.1, estimating a characteristic exponent of about 1.7 across most stocks, well below the Gaussian value of two and consistent with infinite variance. Fama nonetheless concluded that the random walk was an adequate description for practical purposes: price changes were effectively serially uncorrelated, so the past held no exploitable information, which is weak-form efficiency. The tools he used – autocorrelation tests, runs, filter rules, and distributional fitting – were all calibrated to daily data, and the daily frequency forced volatility to be treated as constant within each day. Both the fat-tail evidence and this within-day constancy would later motivate the ARCH models of Engle (1982) discussed in Section 3.5 and the realised-volatility literature, each a response to the inadequacy of the Gaussian, constant-volatility paradigm.

Fama (1970) turned this empirical tradition into a theoretical framework, defining an efficient market as one in which prices "fully reflect" available information. Making the statement precise requires a model of equilibrium expected returns: given an information set and the expected return that the market's own model assigns conditional on it, the abnormal return – the realised return minus that model-implied expectation – must be a fair game with respect to the information set, having conditional expectation zero, while the stronger submartingale condition that the expected gross return never falls below one rules out any trading rule that beats buy-and-hold. Fama organised the evidence by the breadth of the information set: the weak form, in which it contains only past prices and which the serial-correlation and filter-rule tests of Fama (1965) support; the semi-strong form, in which it contains all public information and which event studies of stock splits and earnings announcements broadly confirm; and the strong form, in which it also contains private information and for which the evidence is mixed. The framework consolidated the random-walk tradition of Bachelier, Working, and Samuelson into a single statement linking price informativeness to equilibrium expected returns, and it carried a caveat that has shaped every subsequent test: efficiency is a joint hypothesis, so a rejection may reflect a misspecified asset-pricing model rather than genuine inefficiency. Shiller (1981) pressed exactly this point, showing through variance-bounds tests that stock prices move far more than the present value of subsequent dividends can justify, a challenge that opened the literature on time-varying discount rates and behavioural alternatives. For high-frequency econometrics, Fama's framework supplies both the benchmark and the puzzle: the hypothesis predicts rapid, pattern-free incorporation of information at any frequency, yet intraday data reveal short-horizon predictability, intraday seasonality, and microstructure-induced autocorrelation. The resolution, developed in Hasbrouck (2007) and revisited in Section 3.11, is that efficiency is horizon-dependent – markets can be efficient at daily frequencies while displaying predictable patterns at the second-by-second level.

The mechanism behind that intraday predictability was first given an econometric form by Roll (1984), who derived an estimator of the effective bid-ask spread from the time-series behaviour of transaction prices alone. Roll keeps the efficient price as a random walk driven by public information but lets the observed transaction price bounce between the bid and the ask by half the effective spread, according to a serially independent buy/sell indicator. The bounce induces negative first-order autocorrelation in observed price changes even when the efficient price is unpredictable, with a first-order autocovariance equal to minus one quarter of the squared spread. This negative autocovariance inverts to an implicit-spread estimator – twice the square root of the negative of the sample first-order autocovariance – defined whenever that autocovariance is negative, as it usually is in transaction data. Applied to NYSE data, the estimator recovers the expected cross-sectional pattern of wider implied spreads for smaller firms, using only transaction prices at a time when quote data were rarely available for research. The model is deliberately spare – a constant spread, i.i.d. order flow, and no inventory or informed-trading effects – assumptions that the structural models of Kyle (1985) and Glosten and Milgrom (1985) in Section 3.4 would later relax by decomposing the spread into adverse-selection, inventory, and order-processing components. Its lasting importance for high-frequency econometrics is conceptual: Roll identified the precise channel through which microstructure injects into observed prices an autocorrelation that is absent from the efficient price, and this bid-ask bounce is one of the main sources of the microstructure noise that contaminates realised-variance estimators at high sampling frequencies, the problem taken up in Aït-Sahalia et al. (2005) and Barndorff-Nielsen et al. (2008) in Section 3.6.

The transition toward measuring volatility from higher-frequency data is already visible in French et al. (1987), who studied the intertemporal relation between the expected market risk premium and volatility. Their economic hypothesis is a positive risk – return trade-off, in which the expected excess market return rises linearly with expected volatility, measured as either its standard deviation or its variance, but the more consequential contribution is methodological: they estimated monthly volatility by summing squared daily returns within the month and adding twice the sum of adjacent daily-return cross-products, this second term correcting for the first-order autocorrelation induced by non-synchronous trading. This is the realised-variance idea in embryo – sampling more finely to measure variance more accurately – and is structurally identical to the realised variance of Andersen et al. (2001) in Section 3.5

once daily returns are replaced by intraday returns. Their empirical findings set several later themes. They documented a positive relation between the expected premium and ex ante volatility, alongside a strong negative relation between return innovations and contemporaneous volatility changes, which they read as indirect evidence of the same trade-off operating through the discount rate. They also confirmed, following French and Roll (1986), that volatility is far higher during trading hours than during non-trading hours of equal length, evidence that information enters prices through trading rather than through the mere passage of time – the channel later formalised in the microstructure models of Kyle (1985) and Glosten and Milgrom (1985). The negative return – volatility relation is the leverage effect that Black (1976) had identified and that the asymmetric GARCH specifications of Nelson (1991) and Glosten, Jagannathan and Runkle (1993) would later parameterise, while Schwert (1989) extended the same volatility-measurement approach to a long historical series, decomposing market volatility into macroeconomic, financial, and trading-activity components without finding any single dominant driver.

Taken together, these papers define both the achievement and the ceiling of pre-high-frequency financial econometrics. They established that returns are close to serially uncorrelated, that their distributions are heavy-tailed, that efficiency is a joint hypothesis testable only against an asset-pricing model, and that the trading process itself leaves a measurable imprint on prices. Yet each result was constrained by the data available: volatility could only be inferred, never observed; microstructure could be modelled but barely tested; and the finest measurements of the period – French, Schwert and Stambaugh's monthly variances built from daily returns – already suggested that the binding constraint was the sampling frequency itself. Relaxing that constraint is the subject of the remainder of the chapter, beginning with the emergence of high-frequency data in Section 3.3.

3.3 Emergence of High-Frequency Financial Data

The history of financial crashes provides an intuitive yardstick for the time scale on which markets operate. The Dutch tulip episode of 1636-37, often described as the first speculative bubble, although Garber (1989) and Goldgar (2007) question how speculative or damaging it actually was, inflated over the better part of a year and collapsed over the course of several weeks: news travelled slowly, contracts settled at infrequent intervals, and the institutional infrastructure for continuous trading did not yet exist. Asset prices in such markets were not "data" in the modern econometric sense but the periodic outcomes of negotiated exchanges among small networks of participants.

As markets industrialised, the time scale of disturbances shortened. The 1929 crash unfolded over weeks; Black Monday in October 1987 compressed a comparable wealth loss into a single trading session; and the 1997 Asian and 1998 Russian/LTCM crises were measured in days rather than months. Academic data sets evolved in parallel: by the 1960s and 1970s daily closing prices were the empirical workhorse of

Fama (1965, 1970) and of the early ARCH literature, and by the late 1990s researchers were routinely working with five-minute returns Andersen and Bollerslev (1998).

The present decade has compressed the scale further still. The Flash Crash of 6 May 2010 saw the Dow Jones Industrial Average fall by roughly nine per cent and largely recover within thirty-six minutes, an episode visible only because intraday data were continuously recorded Kirilenko et al. (2017). Below this, a literature on ultrafast extreme events documents thousands of price dislocations lasting fractions of a second, well below the threshold of human reaction time Johnson et al. (2013). Studying such events requires data at the same resolution as the events themselves: tick-by-tick prices, quote revisions, and order-book messages timestamped to microsecond accuracy. The high-frequency econometrics surveyed in this and the following sections is the methodological response to that convergence – as the time scale of relevant market phenomena has shrunk, the time scale of measurement has had to shrink with it.

The methodological consequence of this convergence is that the electronic, time-stamped record of trading is a different kind of object from the daily closing price. Electronic matching and digital record-keeping changed the data-generating process itself: instead of one observation per asset per day, researchers obtained the full sequence of transactions, quote revisions, and, eventually, complete limit-order-book messages, each time-stamped to the second and later to the microsecond. This abundance brought its own difficulties. Observations arrive at irregular, endogenous times rather than on a fixed grid; activity follows pronounced within-day periodicities; the cross-section of simultaneously recorded assets is large; and one must decide whether to analyse data in calendar time, on equally spaced intervals, or in event time, trade by trade. The first systematic studies of transaction data, such as the analysis of NYSE stocks by Wood et al. (1985), already documented the U-shaped intraday pattern in returns, volume, and volatility that became a defining stylised fact. The papers reviewed here establish the empirical regularities of intraday data and the first econometric tools designed for them.

Hasbrouck (1991) provided the template for using transaction-level data to answer a structural question: how much of a price change conveys new information rather than transient friction. He modelled quote-midpoint revisions and signed trade flow as a stationary bivariate vector autoregression (VAR) – each quote revision depending on its own lags together with current and lagged trade flow, with a companion equation for the trade flow – and measured the information content of a trade by the cumulative impulse response of quote revisions to a unit trade innovation. The permanent component of this response is informational; the transitory component is mechanical. On NYSE data he found that the full price impact of a trade arrives only over several subsequent transactions, that it is a concave function of trade size, and that information asymmetry is greater for smaller firms – empirical confirmation of the strategic-trading and adverse-selection mechanisms of Kyle (1985) and Glosten and Milgrom (1985) taken up in Section 3.4. The permanent-transitory decomposition reappears in the realised-volatility literature, where separating the efficient price from microstructure noise poses an analogous identification problem.

Bollerslev and Domowitz (1993) carried out one of the first systematic analyses of genuinely high-frequency data, using roughly 5,100 Deutsche Mark – Dollar quotes per day from the Reuters interbank network. They documented the regularities that became canonical: a pronounced intraday U-shape in trading intensity and spreads, and strong volatility clustering in five-minute returns. Modelling the conditional variance of those returns as a GARCH process – in which the current conditional variance is a positive intercept plus weighted sums of recent squared innovations and recent conditional variances – they found that these coefficients remained strongly significant at the five-minute frequency, so volatility clustering is not merely an artefact of temporal aggregation. They also flagged the *stale-quote* problem – the lag between when a quote is posted and when it is sampled onto a regular grid – which injects autocorrelation into returns and anticipates the microstructure-noise literature of Section 3.6.

Goodhart and O'Hara (1997) surveyed the methodological issues raised by intraday data before the realised-volatility programme had crystallised. They emphasised the calendar-time versus event-time distinction, the systematic intraday seasonality of volume, volatility, and spreads, and the need to adapt ARCH-type models to irregular spacing. A recurring device is the multiplicative decomposition of the intraday conditional variance into a deterministic seasonal factor, which captures the time-of-day pattern through a flexible Fourier form or time-of-day dummies, and a stochastic component of GARCH or stochastic-volatility type; modelling the two separately is necessary for valid inference about either. The survey set out the agenda – estimating volatility from irregularly spaced noisy data, separating seasonality from persistence, and characterising information flow across assets – that the rest of the chapter pursues.

Madhavan (2000) reviewed the market-microstructure literature at the turn of the century, organising it around price formation, market design, transparency, and applications. Two equilibrium results recur throughout the chapter. In sequential-trade models the bid-ask spread decomposes into an adverse-selection component proportional to the dispersion of fundamental value and an inventory and order-processing component; in continuous-auction models the price is linear in net order flow, with a slope – the price impact per unit of order flow – that is Kyle's measure of illiquidity. The survey made the case that microstructure is not merely about trading mechanics but bears on asset pricing, risk, and regulation, and it catalogued the spread-decomposition, VAR-based price-discovery, and trade-indicator methods that became the standard empirical toolkit.

The final contribution treats the timing of trades as data in its own right. Engle (1996, 2000) observed that the irregular intervals between transactions are informative about liquidity and information arrival, and that resampling onto a fixed grid discards this signal. His Autoregressive Conditional Duration (ACD) model writes the duration between consecutive events as the product of an expected duration and an i.i.d. positive, unit-mean innovation, with the expected duration following a GARCH-like recursion in past durations and past expected durations, the direct analogue of the GARCH model of Bollerslev (1986) applied to durations rather than squared returns. Applied to IBM transactions the model reproduces strong duration clustering and the intraday liquidity pattern, and it connects survival analysis to microstructure by making trade intensity, and hence the information content of *time*, an estimable object – a theme formalised in the event-uncertainty model of Easley and O'Hara (1992) in Section 3.4. The full estimation and specification theory appears in the companion paper of Engle and Russell (1998), and the practitioner-oriented treatment of data cleaning, seasonality, and scaling laws built on the Olsen FX data sets is collected in Dacorogna et al. (2001). Taken together, these papers establish high-frequency data as a distinct data-generating process whose statistical complications and economic content are inseparable: the same trading mechanism that produces irregular spacing, seasonality, and noise also produces the information whose flow the next section sets out to measure.

3.4 Market Microstructure and the Econometrics of Trading

High-frequency data made market microstructure empirically testable for the first time, turning models of how prices form into estimable specifications. In this section we examine the link between order flow, spreads, and price impact on one side and information asymmetry on the other, and we draw the distinction between structural models that derive trading outcomes from the optimisation of informed and uninformed agents and the reduced-form econometrics that estimates their observable consequences. The shared payoff is that price discovery – the rate at which private information is impounded into prices – becomes a measurable quantity.

Kyle (1985) is the canonical model of strategic informed trading. A single insider who knows the asset's liquidation value trades against random noise-trader order flow and competitive, risk-neutral market makers who observe only the aggregate order flow – the sum of the insider's demand and noise-trader demand – and set a break-even price equal to the expected liquidation value conditional on the observed order-flow history, which is linear in the latest aggregate order flow. The unique equilibrium is linear, and in the continuous-auction limit the market-depth parameter is constant: depth rises with noise trading and falls with the precision of private information. Prices follow a Brownian motion, all private information is incorporated by the close, and the insider earns a positive expected profit at the noise traders' expense. The model also yields the empirical workhorse of the field: the price-impact regression of Hasbrouck (1991), in which the market-depth slope is estimated by regressing price changes on signed order flow. Its central message – that noise trading provides the camouflage that lets informed trading be profitable – underlies all subsequent microstructure theory.

Glosten and Milgrom (1985) reach the same conclusion from a different mechanism. In their sequential-trade model the market maker faces a mix of informed traders, who know whether the value is high or low, and uninformed traders, and sets the ask equal to the value expected conditional on an incoming buy and the bid equal to the value expected conditional on an incoming sell, so that a positive spread arises purely from adverse selection whenever any traders are informed, even though the market maker is risk-neutral and earns zero expected profit. Transaction prices form a martingale, beliefs converge to the truth as trading proceeds, and the spread is entirely an information cost. The negative autocovariation that Roll (1984) had attributed to a mechanical bid-ask bounce is shown here to arise from a quite different source – adverse selection alone, even with no order-processing or inventory cost – so that Glosten and Milgrom supply the informational component of the spread that Roll's reduced-form model omits, and the VAR of Hasbrouck (1991) can be read as estimating that permanent, information-related component.

Easley and O'Hara (1992) added the insight that time itself carries information. In standard sequential-trade models the clock advances only when a trade occurs, so a quiet market is uninformative; Easley and O'Hara show instead that the absence of trade is a signal, because trades are more likely when an information event has occurred. With event uncertainty, the spread depends on the time since the last trade and narrows as that interval lengthens, since prolonged quiet makes an information event less likely. Their framework underlies the notion of a Probability of Informed Trading, the fraction of trades that are information-motivated, for which Easley, Kiefer, O'Hara and Paperman (1996) later derived a closed form in terms of the probability of an information event, the informed arrival rate, and the uninformed arrival rate, and provided the maximum-likelihood estimator of these parameters from daily buy/sell counts; PIN became a widely used cross-sectional proxy for information asymmetry. The result that inter-trade durations are informative is the economic counterpart of the ACD model of Engle (1996, 2000) in Section 3.3.

These models established price discovery as the organising concept of empirical microstructure, and the gap between their structural primitives and the reduced-form estimators used to recover them defined the field's methodology. The theoretical synthesis of the inventory, information, and strategic-trading strands is O'Hara (1995), while Huang and Stoll (1997) provide a general empirical framework that decomposes the spread into adverse-selection, inventory, and order-processing components and nests the earlier estimators of Glosten and Milgrom (1985) and Roll (1984). The same permanent – transitory logic carries directly into the measurement of volatility, where the efficient price must be separated from the frictions that surround it.

3.5 Volatility as a Latent Process: The High-Frequency Breakthrough

By the time high-frequency data arrived, the dominant model of return volatility was the ARCH family introduced by Engle (1982) and its generalisations, in which the conditional variance is a deterministic function of past information. These models treat volatility as a latent quantity: it is never observed, only inferred through the lens of a parametric specification. The contribution of the high-frequency literature was to make volatility a measurable object – a model-free estimate of the integrated variance of the price process – and in doing so to bring continuous-time concepts into applied econometrics. In this section we trace that transition, from GARCH as the pre-high-frequency benchmark to realised volatility and its distribution theory.

Engle (1982) introduced the first such specification. The ARCH model makes the conditional variance a linear function of past squared innovations, with a positive intercept and non-negative slope coefficients, so that a large innovation raises the variance of subsequent returns and volatility clusters in time. Making the conditional variance a deterministic function of the past, estimated by maximum likelihood, was the innovation later recognised with the Nobel Memorial Prize; its practical limitation is that matching the slow decay of volatility requires many lags.

Bollerslev (1986) generalised ARCH by letting past conditional variances enter the variance equation, giving the GARCH(p, q) model, in which the conditional variance also depends on its own recent past – the variance analogue of moving from an autoregression to an ARMA model: a parsimonious GARCH(1,1) captures the slow decay of volatility that ARCH could match only with long lag structures. The squared residuals follow an ARMA process whose persistence is the sum of the innovation and variance coefficients, the unconditional variance is the intercept divided by one minus that persistence, and the model generates the leptokurtosis of returns without non-Gaussian innovations. GARCH became the standard tool for daily and lower-frequency volatility, and its asymmetric extensions – the EGARCH of Nelson (1991) and the threshold GARCH of Glosten et al. (1993) – incorporated the leverage effect of Black (1976) within the same parametric framework. Its limitation, from the high-frequency standpoint, is precisely that the conditional variance is latent and must be filtered from a model rather than measured.

Andersen and Bollerslev (1998) demonstrated empirically what later theory would justify. Using a year of five-minute Deutsche Mark-Dollar returns, they decomposed volatility multiplicatively into a deterministic intraday seasonal factor, a scheduled-announcement effect, and a slowly moving daily component; they showed that squared daily returns explain only a few per cent of the variation in the latent daily variance, whereas the sum of squared intraday returns – realised variance – tracks it far more closely. Sampling more finely reduces the measurement error in volatility, a fact established here in practice before it was proved in theory.

Andersen et al. (2001) supplied that theory. Modelling the log-price as a continuous semimartingale driven by a drift and a stochastic-volatility diffusion, the object of interest is the integrated variance – the integral of the instantaneous variance over the day – and the realised variance built from intraday returns, the sum of their squares, converges to it in probability as the sampling frequency rises. Empirically, returns standardised by realised volatility are close to Gaussian and the logarithm of realised variance is approximately normal, so realised volatility is approximately log-normal and exhibits long-memory dependence – results that recast volatility modelling as the modelling of an observed time series.

Barndorff-Nielsen and Shephard (2002), working in parallel and largely independently, established the quadratic-variation and central-limit foundations that make realised variance a basis for inference. For a stochastic-volatility log-price they showed that the realised-variance error, suitably scaled, is asymptotically mixed Gaussian, with a limiting variance set by the integrated quarticity – the integral of the fourth power of instantaneous volatility – itself estimable from high-frequency data. This central limit theorem provides feasible confidence intervals for integrated variance and is the starting point for the noise- and jump-robust theory of later sections; the benchmark continuous-time specification behind this literature is the square-root stochastic-volatility model of Heston (1993). In the discrete-time stochastic-volatility tradition the log-variance is itself a latent autoregressive process, estimated by simulated or quasi-maximum-likelihood and later Markov-chain Monte Carlo methods; realised variance supplies a model-free counterpart to that latent object and a benchmark against which such filters can be assessed.

Andersen, Bollerslev, Diebold and Labys (2003)turned measurement into forecasting. Treating the logarithm of realised variance as an observed series, they modelled it as a fractionally integrated (long-memory) Gaussian vector autoregression that combines a short-memory autoregressive polynomial with a fractional-differencing operator, whose estimated order of about 0.4 reproduces the slow hyperbolic decay of the realised-volatility autocorrelation function, and obtained density and Value-at-Risk forecasts that outperform GARCH and stochastic-volatility benchmarks at all horizons. Direct measurement of volatility, rather than parametric filtering, delivers the forecasting gain.

Finally, Engle (2002) extended the measurement of second moments to the multivariate case. The Dynamic Conditional Correlation (DCC) model factors the conditional covariance matrix into a diagonal matrix of conditional standard deviations from univariate GARCH and a time-varying correlation matrix, the latter driven by a scalar GARCH-type recursion in the outer products of past standardised residuals and then rescaled to unit diagonal. The two-step quasi-maximum-likelihood estimator scales to large cross-sections where earlier multivariate GARCH models did not, and it complements the realised-covariance measures of Section 3.8: where realised covariances give nonparametric snapshots from intraday data, DCC supplies a smooth dynamic filter. The realised-volatility programme reframed volatility as a quantity to be measured rather than assumed; the next section confronts the obstacle that the prices used to measure it are themselves observed with error.

3.6 Microstructure Noise and Econometric Identification

Realised variance converges to integrated variance as the sampling interval shrinks, which suggests sampling as finely as possible. In practice the opposite is observed: at the highest frequencies the estimator deteriorates, because observed prices are contaminated by market microstructure noise – the bid – ask bounce analysed by Roll (1984), price discreteness, and the staleness of quotes. In this section we treat high-frequency volatility estimation as an identification problem under measurement error: the efficient price is latent, only a noisy version is observed, and the econometric task is to recover the variation of the former from the latter. The standard model writes the observed log-price as the sum of a latent efficient price and an additive noise term.

Zhou (1996) was among the first to formalise the problem and propose a correction. Under additive i.i.d. noise the naive realised variance is biased: its expectation is the true integrated variance plus a noise term that grows linearly in the number of sampling intervals, so the estimator diverges as sampling intensifies. Plotting realised variance against frequency therefore produces the upward-sloping *volatility signature plot* later used to diagnose contamination. Zhou's bias-corrected estimator adds twice the first-order autocovariance of returns, whose expectation per term is exactly the negative of the noise variance and so cancels the noise bias – reusing the negative-autocovariance signature that Roll exploited to identify the spread, now turned to volatility measurement.

Aït-Sahalia et al. (2005) characterised the resulting bias-variance trade-off precisely. The naive estimator has a bias that grows with the sampling frequency and a variance whose two components move in opposite directions in the sample size, so the mean-squared-error-optimal number of observations is finite – for a liquid stock, on the order of one observation every few minutes, which justifies the then-common practice of discarding most ticks. Crucially, they show that this is a consequence of ignoring the noise: if the noise is modelled explicitly, likelihood-based estimation restores the optimality of sampling as often as possible, and the Gaussian maximum-likelihood estimator of the variance keeps its asymptotic variance even when the noise is non-Gaussian. The first consistent nonparametric estimator in this vein is the Two-Scale Realised Volatility of L. Zhang et al. (2005), which combines realised variances computed at a fast and a slow sampling frequency and remains consistent, albeit at a slow nonparametric rate; a feasible plug-in rule for the optimal sampling frequency is given by Bandi and Russell (2008).

Barndorff-Nielsen et al. (2008) pushed the rate to its limit with the class of realised kernels. A flat-top kernel forms a weighted sum of the realised autocovariances, and with an appropriately growing bandwidth it converges at the best rate attainable under noise – slower than the rate available without noise, reflecting the information lost to contamination. The smooth, non-flat-top kernels that Barndorff-Nielsen et al. (2008) recommend for empirical work trade a little of this rate for guaranteed positive semi-definiteness and greater robustness to dependent noise. The kernel framework also handles endogenous and dependent noise and yields feasible confidence intervals once the noise variance and integrated quarticity are estimated.

Jacod et al. (2009) reached the same optimal rate by a different and transparent route: pre-averaging. The pre-averaged return is a local weighted average of a growing block of consecutive returns, formed with a weight function that vanishes at both endpoints. This averaging reduces the noise while leaving the signal almost intact; a bias-corrected sum of squared pre-averaged returns is then consistent for integrated variance, attaining the same optimal rate and a mixed-Gaussian limit under weak conditions on the noise, which need not be i.i.d., Gaussian, or independent of the price. Pre-averaging also delivers consistent estimates of the integrated quarticity needed for inference and extends naturally to the multivariate and jump-robust settings. The empirical structure of the noise itself – small in magnitude but serially dependent and correlated with the price – is documented by Hansen and Lunde (2006).

That dependence is the subject of Aït-Sahalia et al. (2011), who extend the two-scale and multi-scale estimators to serially dependent noise by decomposing it into an i.i.d. component and a stationary dependent component, and replacing the simple bias correction with a flat-top weighted sum of autocovariances, so that the bias is removed for any summable noise-autocovariance structure while the optimal multi-scale rate is retained. They make explicit the asymptotic equivalence between the multi-scale and realised-kernel families: both depend on the integrated autocovariance of the noise rather than on its individual lags. The common lesson of the section is methodological – under measurement error, statistical efficiency and robustness pull in opposite directions, and identifying the efficient-price variation requires modelling the noise rather than wishing it away. The same separation of a continuous signal from a contaminating component recurs, in a different guise, in the detection of jumps.

3.7 Jumps, Discontinuities, and Intraday Price Dynamics

The realised-variance theory of Section 3.5 assumes a continuous price path. Asset prices, however, occasionally move discontinuously, and high-frequency data make these jumps directly observable. In this section we distinguish the continuous component of price variation from the discontinuous one, review the estimators and tests that high-frequency data made feasible, and note that jumps are typically tied to the discrete arrival of public information – which alters how one interprets the flow of information into prices.

The theoretical starting point is Merton (1976), who extended the Black-Scholes model by adding a compound-Poisson jump component to the price dynamics, in which jumps arrive according to a Poisson process of constant intensity and carry random sizes. The log-price is then the sum of a Brownian motion and a compound-Poisson process, and with log-normal jump sizes the option price is a Poisson-weighted average of Black-Scholes prices. Merton noted that the continuous hedging argument fails when jumps are present, so jump-diffusion markets are incomplete – a point with lasting implications for risk management. For econometrics, the model supplies the object to be estimated: the separation of integrated variance from the sum of squared jumps.

Barndorff-Nielsen and Shephard (2004) made that separation operational with bipower variation. Whereas realised variance converges to the total quadratic variation, including the squared jumps, a suitably scaled sum of products of adjacent absolute returns – bipower variation – converges instead to the integrated variance alone, because two adjacent intervals almost never both contain a jump as the grid shrinks. The difference between realised variance and bipower variation therefore consistently estimates the jump contribution to total variation, and its non-negative truncation gives a feasible jump measure. The construction also generalises to realised power variation of arbitrary order.

Barndorff-Nielsen and Shephard (2006) supplied the distribution theory for testing the null of no jumps. Standardising the scaled relative difference between realised variance and bipower variation by an estimate of the integrated quarticity (built from realised tripower variation) yields a test statistic that is asymptotically standard normal under continuity, with large values signalling a jump. Applied to exchange-rate data the test flags a substantial number of jump days, most of them coinciding with scheduled macroeconomic announcements. Alternative approaches to the same continuous – jump decomposition include the diffusion-versus-jump tests of Aït-Sahalia (2002), which exploit the behaviour of the return density at small and large values, and the threshold estimator of Mancini (2009), which truncates returns above a data-dependent bound and is robust to small jumps.

Lee and Mykland (2008) localised jumps to individual observations rather than detecting them at the daily level. Their statistic standardises each high-frequency return by a local bipower volatility estimate, and accounts for the multiplicity of simultaneous intraday tests through an extreme-value argument: under continuity the maximum absolute standardised return over the day's observations has a Gumbel limit, whose quantiles provide the critical values that control the family-wise error rate. The test delivers the time and size of each jump, and on equity data it shows that index jumps are predominantly tied to macroeconomic news while individual-stock jumps follow firm-specific events – a heterogeneity with direct consequences for option pricing. The decomposition of variation into continuous and jump parts also feeds the forecasting models of Section 3.10: Andersen et al. (2007) show that the continuous component is far more persistent and predictable than the jump component. The same decomposition extends to several assets at once, which is the subject of the next section.

3.8 Multivariate High-Frequency Econometrics

Extending realised measures to several assets confronts an obstacle that has no counterpart in the univariate case: prices are not observed simultaneously, and at very short horizons their measured co-movement collapses. In this section, we treat the estimation of covariances and correlations under asynchronous sampling, beginning with the empirical regularity that motivates it and ending with the connectedness measures that build a network view of dependence on top of realised covariation.

Epps (1979) documented the regularity that bears his name: the sample correlation between returns on related stocks falls sharply as the measurement interval shortens, approaching zero at intervals of a few minutes even for stocks that are strongly correlated at daily horizons. The standard realised-covariance estimator formed from synchronised returns – the sum of cross-products of contemporaneous intraday returns – is biased toward zero as the sampling frequency grows, with the bias governed by the fraction of intervals in which one asset does not trade. The effect arises chiefly from non-synchronous trading and from the negative autocorrelation that microstructure frictions induce in each series. Its resolution became a research programme in its own right.

Hayashi and Yoshida (2005) provided the solution. Rather than forcing the two price series onto a common grid – the step that creates the Epps bias – their estimator sums the products of returns over every pair of *overlapping* observation intervals –

those for which the two assets' sampling intervals intersect – and is consistent and asymptotically normal for the quadratic covariation without any interpolation or data discarding. The estimator was subsequently combined with pre-averaging to handle the simultaneous presence of microstructure noise and asynchrony by Aït-Sahalia et al. (2010), and the realised-kernel approach of Section 3.6 was extended by Barndorff-Nielsen et al. (2011) to a multivariate estimator that is positive semi-definite by construction and robust to noise, asynchrony, and high dimension. Jump-robust covariation follows from the bipower-covariation extension of Barndorff-Nielsen and Shephard (2004) in Section 3.7, completing the set of realised second-moment measures.

Once covariances can be measured, dependence across many markets can be summarised as a network. Diebold and Yilmaz (2012) build such a measure from a vector autoregression by replacing the order-dependent Cholesky variance decomposition with the ordering-invariant generalised forecast-error variance decomposition. Each pairwise entry gives the share of one variable's multi-step forecast-error variance attributable to shocks in another, computed from the moving-average coefficients of the vector autoregression and the innovation covariance matrix. From these elements a total spillover index and directional, net, and pairwise spillovers follow after row-normalisation. Their earlier index, based on Cholesky decompositions, is Diebold and Yilmaz (2009); the 2012 refinement makes the measures comparable across markets and periods without identifying assumptions. Estimated on daily range-based (high-low) volatilities, the framework shows that equity markets are persistent net transmitters of volatility and the total index to rise sharply in crises, so it serves as a gage of systemic stress; the same construction applies directly to the high-frequency measures of Andersen, Bollerslev, Diebold and Labys (2003) where these are available. That use of high-frequency measurement for monitoring risk is the subject of the next section.

3.9 High-Frequency Econometrics and Financial Risk

High-frequency volatility and covariation measures are direct inputs to risk management: a sharper estimate of current volatility yields a sharper forecast of the distribution of future returns. But systemic risk involves dependence in the tails – the tendency of losses to coincide across institutions during extreme events – which the quadratic-variation measures of the preceding sections do not by themselves capture. In this section we review two complementary developments: a sign-resolved refinement of realised variance, and a measure of an institution's contribution to system-wide risk, both of which connect high-frequency measurement to financial stability and regulatory monitoring.

Barndorff-Nielsen, Kinnebrock and Shephard (2010) refine realised variance by separating downward from upward moves. The downside and upside realised semivariances – formed by summing squared negative returns and squared positive returns separately – sum to realised variance, and under a general semimartingale each converges to half the integrated variance plus the jumps of the corresponding sign. The continuous part contributes symmetrically, while jumps load only on the side of their sign, so the difference between the downside and upside semivariances isolates signed jump variation – a diagnostic for downside risk that total realised variance cannot provide. The measure motivates the good-volatility/bad-volatility distinction later exploited in forecasting by Patton and Sheppard (2015), and it matters economically because downside risk is what loss-averse investors and risk managers care about.

Adrian and Brunnermeier (2016) address systemic risk directly with $\Delta CoVaR$. Although it is estimated at lower frequency rather than from intraday data, it enters here as the point at which the chapter's measurement programme meets a problem – tail dependence – that quadratic-variation measures cannot solve. Where the Value-at-Risk of an institution is the loss quantile that measures its risk in isolation, CoVaR is the Value-at-Risk of the financial system conditional on that institution being in distress, and the systemic-risk contribution is the difference between this system Value-at-Risk when the institution is in distress and when it is at its median. Estimated by quantile regression of system returns on institutional returns and lagged state variables, $\Delta CoVaR$ is predicted by leverage, size, and maturity mismatch, and a forward-looking version signalled rising systemic risk a year before the 2007-9 crisis when conventional Value-at-Risk did not. The quantile-regression machinery it relies on is the Conditional Autoregressive Value-at-Risk (CAViaR) approach of Engle and Manganelli (2004), which specifies the dynamics of a conditional quantile directly. Related measures of an institution's systemic contribution include the Systemic Expected Shortfall of Acharya, Pedersen, Philippon and Richardson (2017) and the capital-shortfall measure SRISK of Brownlees and Engle (2017), the latter built on the DCC model of Engle (2002) from Section 3.5 and adopted for near-real-time regulatory monitoring. The lesson for high-frequency econometrics is that realised measures inform the inputs to risk models but that tail dependence requires an additional extreme-value and network perspective. How realised measures improve forecasting more generally is taken up next.

3.10 Forecasting with High-Frequency Data

The realised measures of the previous sections are valuable not only for describing volatility but for forecasting it, and the gains are substantial: a model built on realised variance typically outperforms one built on squared daily returns. In this section we review the two specifications that organise the field – one that captures the persistence of volatility through components defined over different horizons, and one that combines predictors sampled at different frequencies – and situate them against the broader forecasting literature.

Corsi (2009) proposed the Heterogeneous Autoregressive (HAR) model, which forecasts realised variance from its own averages over daily, weekly, and monthly horizons, the weekly and monthly components being simple averages of past daily realised variances. Although the model is just a constrained autoregression, the overlapping averages reproduce the hyperbolic decay of the realised-volatility autocorrelation function, so it approximates long memory without fractional integration and forecasts as well as far more complex ARFIMA models. Its motivation is the heterogeneous-market hypothesis, that traders acting on different horizons each contribute a volatility component. The model is the standard benchmark in volatility forecasting, and its extensions add the jump component of Barndorff-Nielsen and Shephard (2004) from Section 3.7 and the signed semivariances of Patton and Sheppard (2015), the latter showing that negative-jump variation predicts future volatility more strongly than positive-jump variation.

Ghysels, Santa-Clara and Valkanov (2006) approach the frequency-mismatch problem differently, with Mixed Data Sampling regressions that let a low-frequency target be predicted by high-frequency regressors without aggregating either, the lag weights following a parsimonious functional form – such as the exponential Almon polynomial – whose few parameters are estimated jointly with the rest. Comparing predictors, they find that daily realised power – the sum of absolute intraday returns – is the single best forecaster of future volatility, that pushing to intraday predictors adds little once daily realised measures are used, and that one to two months of past data suffice. These results are robust across horizons and hold in and out of sample.

Both models confirm that high-frequency information improves volatility forecasts, but the size of the gain depends on the asset and the benchmark: the comparison of Hansen and Lunde (2005) finds that little beats a GARCH(1,1) for daily exchange rates while more elaborate specifications matter for equities. A reduced-form framework that combines the HAR structure with jump-robust components and a leverage adjustment is developed by Andersen, Bollerslev and Huang (2011), and represents one natural reference point for current forecasting practice. The forecasting gains from high-frequency data return us, finally, to the question of market efficiency with which the chapter began.

3.11 Market Efficiency Revisited

The efficient-market hypothesis of Fama (1970) predicts that prices incorporate information rapidly and without exploitable patterns. High-frequency evidence does not overturn this prediction so much as qualify it: at intraday horizons prices depart from a frictionless random walk in predictable ways, yet the efficient price is restored within seconds. In this section we reinterpret efficiency as a horizon-dependent property and examine the role of high-frequency trading in price discovery.

Hasbrouck (2007) organises empirical microstructure around the decomposition of the observed transaction price into an efficient price that follows a random walk and a transitory deviation due to frictions. This nests Roll's implicit-spread model, the adverse-selection model of Glosten and Milgrom (1985), and VAR-based price discovery in a single framework, and it makes the contribution of competing trading venues to price discovery measurable through the information share introduced by

Hasbrouck (1995). For two markets quoting the same security with uncorrelated reduced-form innovations, the information share of the first simplifies to the ratio of that market's squared scaled innovation to the sum of the two markets' squared scaled innovations, the fraction of efficient-price variance attributable to that market. The picture that emerges is of efficiency as a matter of degree and horizon: the efficient price is restored almost immediately after each trade, but the transaction price fluctuates around it in ways that are predictable at the second-by-second level. High-frequency evidence refines the efficient-market hypothesis rather than rejecting it.

Brogaard, Hendershott and Riordan (2014) bring this question to the high-frequency-trading era using NASDAQ data that identify the order flow of high-frequency traders. Regressing efficient-price innovations on high-frequency and non-high-frequency net order flow separately, and decomposing high-frequency activity into its liquidity-demanding and liquidity-supplying parts, they find that high-frequency traders trade in the direction of permanent price changes and against transitory pricing errors, that this is achieved mainly through their marketable orders, and that their passive limit orders are adversely selected. On balance high-frequency traders facilitate price efficiency by impounding public information quickly, at the cost of imposing adverse selection on slower liquidity providers. The causal effect of automation on liquidity is identified by Hendershott, Jones and Menkveld (2011) through an exogenous shift in NYSE message capacity, the behaviour of a single high-frequency market maker is documented by Menkveld (2013), and the broader changes that speed and complexity have brought to the empirical content of microstructure theory are assessed by O'Hara (2015). The methodological developments that made all of this measurable are the subject of the next section.

3.12 Methodological Synergies in Financial Econometrics

The methods reviewed in this chapter draw on three disciplines at once – econometrics, probability theory, and financial economics – and high-frequency data have intensified their interaction. The realised measures, jump tests, and noise-robust estimators of the preceding sections all rest on the theory of semimartingales and on central limit theorems for functionals of processes observed at discrete times. In this section we review the two bodies of theory that supply that foundation and note the growing links between high-frequency econometrics and the wider statistical and computational literature.

Duffie (2001) develops the asset-pricing theory that motivates the model class. The absence of arbitrage is equivalent, under regularity conditions, to the existence of a state-price density under which discounted prices plus cumulated dividends form a martingale, and the continuous-time machinery – Itô's formula, the Itô isometry, and the change of measure – delivers the stochastic-volatility and jump-diffusion models used throughout. Central to the chapter is the resulting notion of quadratic variation, which for a continuous semimartingale is the integral of the instantaneous variance over time and is precisely the object that realised variance estimates. Duffie supplies the no-arbitrage justification for treating prices as semimartingales; the inferential theory for those semimartingales is the complementary contribution.

Jacod and Protter (2012) systematise the asymptotic theory of inference about continuous-time processes sampled discretely. For an Itô semimartingale – a drift plus a volatility-driven Brownian term plus a jump term – observed on a shrinking grid, power-variation statistics formed by summing a function of the increments obey a law of large numbers and a central limit theorem that is the engine behind every feasible result in the chapter, with convergence that is *stable in law*, the property that allows the random asymptotic variance to be replaced by a consistent estimator and so makes studentised, feasible inference possible. Taking the squared-increment function recovers the realised-variance theory of Barndorff-Nielsen and Shephard (2002); multipower and threshold functionals give the jump-robust estimators of Section 3.7; and the same framework underlies the realised-kernel and pre-averaging results of Barndorff-Nielsen et al. (2008) and Jacod et al. (2009). The applications-oriented companion to this theory is the textbook of Aït-Sahalia and Jacod (2014), the survey of Mykland and Zhang (2012) links the two-scale, multi-scale, and pre-averaging estimators under a common asymptotic framework, and the Realised GARCH model of Hansen, Huang and Shek (2012) bridges the GARCH tradition of Bollerslev (1986) with the realised-volatility literature of Andersen, Bollerslev, Diebold and Labys (2003) by updating the conditional variance with realised measures. High-frequency econometrics has thus become a meeting point of these methodological strands, and increasingly of statistics, computation, and machine learning as well – a confluence that the remaining sections take up, first surveying the recent extensions and applications of the measurement programme and then turning to a retrospective assessment.

3.13 Recent Developments and Further Extensions

In this section we survey four recent extensions of the measurement paradigm: to high-dimensional financial systems, to forecasting models built on realised measures, to network representations of dependence, and to the measurement of systemic risk. Each takes the central lesson of the previous sections – that volatility, covariation, and jump variation can be measured directly from intraday data rather than treated as latent objects requiring parametric modelling – into a setting the original realised-volatility literature did not address. As Aït-Sahalia and Jacod (2014) emphasise, this is a change in econometric methodology as much as in technique: where earlier financial econometrics specified dynamic models for latent quantities, high-frequency econometrics constructs consistent measurements of economically meaningful objects directly from the data.

The first development concerns the transition from small-dimensional systems to environments involving hundreds or thousands of assets. Much of the early realised-covariance literature was developed for small collections of securities; as the cross-section grows, the estimation of integrated covariance matrices becomes harder, because sampling error, microstructure noise, and asynchronous trading accumulate across assets and standard estimators become ill-conditioned. The multivariate realised kernel of Barndorff-Nielsen et al. (2011), which extends the univariate realised-kernel methodology of Section 3.6 to a covariance estimator, addresses this directly: Barndorff-Nielsen et al. (2011) construct estimators that remain robust to microstructure noise and non-synchronous trading while preserving positive semi-definiteness, a property that matters in practice because the resulting matrices feed directly into portfolio allocation and risk management. A related problem is the extraction of common factors from high-dimensional panels of high-frequency returns. The classical approximate-factor literature was developed for low-frequency data, where measurement error is secondary; at high frequency, microstructure noise contaminates the prices and standard principal-component methods break down. Aït-Sahalia and Xiu (2019) develop a principal-component analysis for high-frequency data, showing that latent factor structure can be recovered once the effect of noise is controlled – in practice through sparse sampling – and that the resulting estimators improve the measurement of common variation. Their work connects high-frequency econometrics to the large-dimensional factor-model literature.

The second development integrates realised measures into forecasting models. These measures were initially treated as ex-post estimates of latent variation; later work recognised that they also carry predictive information and can enter forecasting models directly. The Realised GARCH model of Hansen et al. (2012), already introduced in Section 3.12, makes this explicit by adding a measurement equation that links the realised measure to the latent conditional variance, so that returns, volatility, and the realised measure are modelled jointly and the realised measure informs both estimation and prediction. Bollerslev, Patton and Quaedvlieg (2016) address a complementary point: realised variance is itself an estimate whose precision varies over time. Because a noisily measured realised variance enters HAR-type regressions as a regressor, it induces an attenuation bias; their HARQ model lets the autoregressive coefficient vary with realised quarticity, downweighting the realised measure in periods when it is imprecisely estimated. The correction is simple to implement and yields measurable forecasting gains.

The third development views financial markets as networks. Covariance and correlation summarise dependence but say little about how shocks propagate across a system. Building on the variance-decomposition connectedness measures of Diebold and Yilmaz (2012) introduced in Section 3.8, Baruník and Křehlík (2018) develop a frequency-domain decomposition that separates connectedness into short-, medium-, and long-run components, recognising that some shocks produce short-lived spillovers and others persistent ones, and linking connectedness analysis to the frequency-domain methods of time-series econometrics. Demirer et al. (2018) apply the network view to institutions, constructing measures of global bank connectedness and documenting that periods of stress are marked not only by higher volatility but by more intense network linkages. A caveat applies to this strand and the next: these connectedness and systemic-risk measures are typically estimated from daily or range-based data rather than from intraday observations, so they consume the lower-frequency end of the measurement programme rather than tick-level data.

The fourth development is systemic-risk measurement. The global financial crisis exposed the limits of risk measures that treat institutions in isolation and prompted frameworks for system-wide vulnerability. Adrian and Brunnermeier (2016) introduce CoVaR, the value-at-risk of the financial system conditional on an institution being in distress, measured relative to that institution's median state, which shifts attention from stand-alone risk to an institution's marginal contribution to system risk. Acharya et al. (2017) develop the Systemic Expected Shortfall (SES), the expected capital shortfall of an institution conditional on a system-wide crisis, linking individual losses to aggregate stress. Brownlees and Engle (2017) operationalise this idea as SRISK, the expected capital shortfall of an institution in a market downturn, built – as noted in Section 3.9 – on dynamic volatility and correlation estimates and adopted for near-real-time regulatory monitoring. These measures inherit the chapter's theme: they depend on accurate estimates of volatility and dependence, and so on the measurement methods developed for high-frequency data.

The methods now defining the frontier extend the same logic in directions this chapter has only touched on. Looking further ahead, rough and fractional stochastic-volatility models revisit the dynamics of the latent variance, while machine-learning methods are increasingly applied to volatility forecasting and to the order book itself (Section 3.17), where the data are highest-dimensional and least amenable to parametric specification. Together with high-dimensional covariance regularisation, these lines suggest that the next advances will come less from finer sampling than from combining measurement with high-dimensional, network, and learning-based tools.

Taken together, these developments illustrate that the realised-measurement programme has expanded well beyond volatility estimation. The same principle – constructing reliable measurements of economically meaningful quantities directly from increasingly rich data – now underpins high-dimensional covariance estimation, forecasting, systemic-risk measurement, network analysis, and many modern applications reviewed in the following sections. High-frequency econometrics has therefore evolved from a specialised methodology for volatility estimation into a general framework for empirical measurement in finance. The following sections illustrate how this measurement paradigm has expanded into new applications.

3.14 High-Frequency Identification of Macroeconomic and Monetary News

The measurement programme of the preceding sections treats volatility and covariation as the objects of interest. A parallel literature, developed alongside that programme rather than after it, turns the same high-frequency lens on causal identification: by narrowing the observation window around a scheduled macroeconomic release or a monetary-policy announcement to a few minutes, this strategy isolates the surprise component of the news from the slower-moving confounders that contaminate lower-frequency studies. In this section we review how this event-window strategy identifies the response of asset prices to information, first in the monetary-policy setting and then in the broader study of macroeconomic announcements.

The identification problem in the monetary-policy literature is that only the *unexpected* component of a policy action should move asset prices, since the anticipated component is already embedded in them. Kuttner (2001) solves it using the price of Fed funds futures, whose settlement is tied to the monthly average of the realised funds rate. Because a target change partway through the month affects only the remaining days, the surprise component is recovered by scaling up the one-day change in the spot-month futures rate by a factor that undoes the within-month averaging. Regressing changes in market interest rates on this surprise and on the expected component separately, Kuttner finds that rates respond strongly to the unexpected part and barely at all to the anticipated part, direct evidence that policy is priced in advance. Gürkaynak et al. (2005) refine the measurement by moving to intraday windows bracketing the announcement itself and show that the response is not one-dimensional: a first "target" factor captures the surprise in the current rate, while a second "path" factor captures news about the future rate path conveyed by the accompanying statement – an early high-frequency account of forward guidance.

These surprises became inputs to macroeconomic identification. Gertler and Karadi (2015) use high-frequency policy surprises, including surprises in futures dated several months ahead that capture forward guidance, as external instruments in a vector autoregression, and trace the response of output, prices, and especially credit spreads to an exogenous policy shock, finding that movements in credit costs amplify the real effects. Nakamura and Steinsson (2018) construct a "policy news shock" as the first principal component of the changes in several interest-rate futures over a thirty-minute window around scheduled announcements, and document an *information effect*: announcements also reveal the central bank's own assessment of fundamentals, so a tightening surprise can raise private-sector growth expectations and complicate the textbook reading of monetary non-neutrality. Jarociński and Karadi (2020) turn this complication into an identification device, separating a pure monetary-policy shock from a central-bank information shock by the sign of the high-frequency co-movement between interest rates and stock prices: a negative co-movement (rates up, equities down) signals a contractionary policy shock, while a positive co-movement (both rising) signals favourable information about the economy.

The same window-narrowing logic applies to scheduled macroeconomic releases. Fleming and Remolona (1999) document a two-stage adjustment in the U.S. Treasury market: a brief first stage in which a major announcement produces a sharp, near-instantaneous price change accompanied by *reduced* trading volume and a sharply wider bid – ask spread, showing that price reactions to public information do not require trading, followed by a prolonged second stage of surging volume and persistent volatility as positions adjust. To relate price moves to the news itself, Balduzzi et al. (2001) measure the surprise in each release relative to its survey expectation and standardise it – the announced value minus the median survey forecast, divided by the standard deviation of the forecast error – which places surprises measured in different units on a common scale. Regressing intraday Treasury-price changes on these standardised surprises, they find that seventeen releases move prices significantly, the nonfarm payroll report most of all. Andersen, Bollerslev, Diebold and Vega (2003) carry the design to the foreign-exchange market at five-minute frequency, regressing returns on the same standardized news, with each indicator's standardized surprise carrying its own price-impact coefficient (a simplified contemporaneous form of their distributed-lag specification). They show that announcement surprises produce discrete conditional-mean jumps, the return-side counterpart of the price jumps detected in Section 3.7, and that the response is asymmetric, with bad news moving the exchange rate more than good news of equal magnitude, so that high-frequency data resolve a real-time price-discovery process that lower frequencies blur. Finally, Lucca and Moench (2015) document a pattern around the announcement rather than at it: U.S. equities earn large average excess returns in the twenty-four hours *before* scheduled Federal Open Market Committee (FOMC) announcements, a pre-FOMC announcement drift that accounts for a substantial share of realised excess equity returns since the mid-1990s and that standard asset-pricing models struggle to explain, a horizon-dependent departure from efficiency in the spirit of Section 3.11.

These designs share the logic of the realised-measure programme – precise measurement enabled by fine sampling – applied to identification rather than to volatility. We turn next to a related use of high-frequency prices: locating where, across competing venues, information is first impounded.

3.15 Price Discovery Across Markets

When the same security, or a set of closely linked securities, trades in more than one venue, the efficient price is common to all of them but the observed quotes and transaction prices differ. Price discovery asks which venue moves first – where new information enters prices, and how the others follow. In this section we review the econometric measures developed to answer this question, which build on the cointegration and permanent – transitory decompositions of time-series econometrics and on the information-share framework introduced in Section 3.11.

When a security trades in several venues, its prices are tied together by arbitrage and are therefore cointegrated: they share a single non-stationary common factor, the efficient price, around which each venue's prices make stationary excursions. Hasbrouck (1995) builds price discovery on exactly this structure. Writing the vector of market prices as a cointegrated system whose common random-walk component is the efficient price, he attributes the variance of the efficient-price innovation to the individual markets. Given the vector of long-run impacts of each market's quote innovation on the efficient price and the covariance matrix of those innovations, the information share of a market is the fraction of the total efficient-price innovation variance attributable to that venue. The expression holds when the innovations are uncorrelated; with correlated innovations, Hasbrouck reports upper and lower bounds obtained by reordering a Cholesky factorisation of the innovation covariance matrix.

Applied to quotes for the same stock on the NYSE and the regional exchanges, the measure attributes the large majority of price discovery to the NYSE. The efficient price here is the permanent component of Section 3.11, now shared across venues rather than inferred from a single price series.

Two complementary strands fill out the framework. Gonzalo and Granger (1995) provide the econometric foundation, decomposing a cointegrated system into permanent and transitory parts and identifying the common factor as the linear combination of the observed prices that the transitory disturbances do not affect in the long run – the combination whose weights are orthogonal to the error-correction loadings – so that those common-factor weights measure each market's contribution to the permanent component. Harris et al. (1995) apply the error-correction machinery directly to a single security trading on informationally linked exchanges, using the speed-of-adjustment coefficients to locate where price discovery takes place. de Jong (2002) compares the two leading measures and shows that, although the Hasbrouck information share and the Gonzalo–Granger common-factor weight are closely related, only the information share reflects the *variability* of each market's innovations rather than the loadings alone, so the two answer subtly different questions about where information enters prices.

The common thread is that a shared efficient price, observed through several noisy channels, can be recovered and each venue's contribution to it quantified – a multivariate counterpart to the signal-extraction problem of Section 3.6. The same high-frequency resolution that makes price discovery measurable also sharpens the comparison between realised and option-implied measures of risk, to which we turn next.

3.16 Volatility, Risk, and Connectedness

This section draws together two extensions of high-frequency measurement that reach beyond the realised variance of a single asset. The first compares the realised volatility of Sections 3.5–3.6 with the volatility implied by option prices and studies the wedge between them – the variance risk premium. The second moves from pairwise covariation to the system level, representing a panel of assets or institutions as a network whose linkages can themselves be measured and monitored, extending the connectedness measures introduced in Section 3.8 and the systemic-risk discussion of Section 3.9.

The realised measures of Sections 3.5–3.6 estimate volatility from the price path under the physical measure. Option prices offer a second, forward-looking estimate under the risk-neutral measure, and the gap between the two is itself informative. Britten-Jones and Neuberger (2000) show that, when the price follows a continuous process and interest rates are zero (equivalently, when prices are expressed as forwards), the risk-neutral expected integrated variance is pinned down by the cross-section of option prices alone, with no parametric volatility model, as a strike-weighted integral of European call prices across all strikes at the relevant maturity. This model-free implied variance is the fair value of a variance swap and is the construction underlying the VIX index. Jiang and Tian (2005) extend the result to processes with jumps and show that the resulting model-free implied volatility is an efficient forecast of future volatility that subsumes both the Black – Scholes implied volatility and the historical realised volatility. Carr and Wu (2009) use the same portfolio of options to synthesise the variance swap rate directly and define the variance risk premium as the difference between realised variance and this risk-neutral rate, documenting that it is large and on average negative, so that investors pay to hedge variance risk. Bollerslev et al. (2009) make the premium an asset-pricing variable. Measuring the *ex-ante* variance risk premium as the difference between the risk-neutral implied variance and the conditional expectation of next period's realised variance, the latter computed from high-frequency intraday returns, they find that it predicts aggregate stock returns, with the predictability strongest at the quarterly horizon, where it dominates standard predictors such as the price – earnings ratio. The result depends on measuring the realised component from intraday rather than daily data, tying the premium back to the measurement programme of the chapter.

A second extension moves from the second moments of a single series to the dependence structure of an entire system, represented as a network. Diebold and Yilmaz (2014) build such a network from the generalised forecast-error variance decomposition introduced in Section 3.8: the pairwise variance-decomposition shares are read as the weighted, directed edges of a connectedness network among financial firms, from which node-level directional connectedness and a system-level total connectedness index follow as simple sums, and standard network statistics such as the degree distribution acquire an econometric meaning. Billio et al. (2012) construct connectedness differently, from principal-components analysis and pairwise Granger-causality tests among banks, broker-dealers, insurers, and hedge funds, and show that it rises ahead of and during crises, providing an early-warning signal of systemic risk. Barigozzi and Brownlees (2019) place network estimation on a formal footing with NETS (network estimation for time series), which fits a vector-autoregressive system whose autoregressive matrices and innovation concentration matrix are assumed sparse, so that the long-run and contemporaneous partial-correlation networks are recovered jointly by a LASSO penalty, a high-dimensional counterpart to the connectedness measures that scales to large panels of institutions.

Whether read as a premium for bearing variance risk or as a map of how shocks travel through a financial system, these measures rest on the same foundation – accurate, high-frequency estimates of second moments. We close the survey by turning to two frontiers on which that foundation meets new tools and new markets.

3.17 Frontiers: Machine Learning and New Markets

We close with two frontiers on which high-frequency econometrics is currently expanding. The first is methodological: machine-learning methods are increasingly applied to asset pricing and to the limit order book itself, where the data are highest-dimensional and least amenable to parametric specification. The second is empirical: cryptocurrency markets, which trade continuously and record every transaction on a public ledger, provide a new and unusually complete high-frequency laboratory. In this section we review both.

Machine learning enters financial econometrics as a flexible way to estimate conditional expectations from many, possibly redundant, predictors. Gu et al. (2020) provide the canonical demonstration, casting the measurement of equity risk premiums as a prediction problem in which next period's return is an unknown, non-parametric function of a large vector of firm characteristics and macroeconomic predictors, plus noise. Comparing regularised linear models, regression trees, and neural networks across a wide cross-section of U.S. stocks, they find that the flexible methods deliver large out-of-sample gains, and trace those gains to nonlinearities and interactions that linear models cannot capture. The same tools are applied directly to the high-frequency record. Sirignano and Cont (2019) train deep neural networks on a large limit-order-book data set and find a striking *universality*: a single model estimated by pooling across stocks forecasts short-horizon price moves better than models fitted asset by asset, evidence that the mapping from order flow to price changes is a stable, nonlinear, and largely common feature of price formation. Z. Zhang et al. (2019) specialise the architecture to the order book with DeepLOB, a convolutional and recurrent (LSTM) network applied to the raw book that delivers stable out-of-sample accuracy across instruments. Machine learning also reaches the hedging problem: Buehler et al. (2019) replace the model-based delta with a neural-network trading strategy chosen to minimise a convex risk measure of the hedged payoff – the derivative payoff net of the strategy's cumulative trading gains and transaction costs – so that hedging under frictions becomes a data-driven optimisation rather than a closed-form formula.

The second frontier is a new market. Cryptocurrencies trade continuously and record every transaction on a public ledger, providing an unusually complete high-frequency laboratory. Makarov and Schoar (2020) document large, recurrent arbitrage spreads for the same coin across exchanges, much wider across countries than within them, which they link to capital controls and the slow movement of arbitrage capital, and decompose exchange-level order flow into common and idiosyncratic components. Liu and Tsyvinski (2021) show that cryptocurrency returns form a distinct asset class: they are not explained by stock-market or macroeconomic factors but are predicted by a strong own-momentum effect and by proxies for investor attention and user adoption. Cong et al. (2021) provide a matching valuation theory in which a token's price is set by aggregating heterogeneous users' transactional demand rather than by discounting cash flows, so that adoption dynamics and network externalities drive prices, a structural complement to the reduced-form evidence.

These frontiers inherit the chapter's central lesson while extending it. Machine learning advances the estimation rather than the measurement leg of the programme – letting the data, not a parametric form, determine the conditional expectation – yet it shares the chapter's guiding principle that the object of interest should be recovered from the data rather than imposed by assumption. Progress continues to come from disciplined inference about economically meaningful objects, now with richer tools and in new markets.

3.18 High-Frequency Econometrics in Perspective

Viewed as a whole, the development surveyed in this chapter is a movement from data scarcity to data abundance, and the methodological lessons follow from that shift. When observations were few and widely spaced, progress came from specification: volatility, liquidity, and information were latent quantities, and the econometrician's task was to write down a parametric model from which they could be inferred. When the price process can be observed almost continuously, the emphasis moves to measurement and identification – constructing model-free estimators of well-defined population objects such as integrated variance and quadratic covariation, and understanding the conditions under which the data identify them. The realised-volatility programme is the clearest example: a quantity that GARCH could only filter became something one could measure, with a distribution theory attached. For a textbook treatment of the econometric methods surveyed in this chapter, see Hautsch (2012).

A second lesson is that market microstructure is not a nuisance to be removed but a structural feature of the data. The same trading mechanism that generates bid – ask bounce, irregular spacing, and intraday seasonality also generates the order flow, price impact, and price discovery that carry economic information; noise and signal are produced together and must be separated rather than assumed away. A third lesson is the centrality of robustness – to noise, to jumps, to asynchronous observation – which has shaped the estimators that survive in practice. The enduring impact on econometrics as a discipline is that the high-frequency literature reoriented applied work toward continuous-time models, nonparametric measurement, and inference under measurement error, and made these the default rather than the exception. Beyond the window this chapter surveys, the same measurement-first orientation now shapes the current frontier: rough and fractional stochastic-volatility models that exploit the fine sampling of realised measures, and machine-learning methods for volatility forecasting and order-book dynamics. The next section makes part of this assessment quantitative.

3.19 A Data-Driven View of What Ages Well

A retrospective chapter such as this one inevitably involves a judgement about which methodological contributions remain in active use and which have given way. Such judgements are typically qualitative. To complement them, we adopt an explicitly algorithmic procedure: we collect per-year citation counts for each paper listed in the original outline of this chapter, together with the more recent contributions surveyed in Sections 3.13–3.17, and visualise their reception trajectories side by side. The intuition is straightforward – a paper whose citation rate is large *and* stable across a decade and a half, rather than peaking and decaying, can be said to have aged well, in the sense that current researchers continue to find it useful at the same rate as the previous generation did.

Approach

For each citation key in this list, we resolve the paper against the OpenAlex API, with Semantic Scholar as a fallback when a record is missing. We retrieve the `counts_by_year` field, assemble a long-format data set indexed by paper and citation year, and normalise each row by its paper-specific maximum so that the colour scale represents *relative* reception across the paper's own history. The result is Figure 3.1: bright cells mark years close to a paper's peak; dark cells mark years substantially below it. The colour map (*cividis*) is monotonic in luminance and remains readable in black-and-white print.[1]

Outcomes

Several distinct patterns emerge. Many recent papers display the canonical *inverted-U* trajectory: a build-up of citations in the years immediately following publication, a peak roughly five to ten years later, and a slow decline. This is the expected shape for an actively-cited research-frontier contribution and characterises, for example, the realised-kernel and pre-averaging papers in Section 3.6 and the systemic-risk literature in Section 3.9.

Of more interest, given the question of what ages well, are the papers whose citation rate is large *and remains close to its peak* throughout the observation window. These are contributions that current researchers continue to cite at essentially the same rate as researchers writing five or ten years ago. Five papers stand out on this criterion. The first is the ARCH paper of Engle (1982), which accumulates roughly eight hundred citations per year on a near-flat distribution from 2010 onwards; time-varying volatility is now taught at undergraduate level, and ARCH-type specifications remain the default in regulatory and applied work. Its GARCH generalisation, Bollerslev (1986), is the most heavily cited paper in the outline by a wide margin, with a rate that is both large – around nine hundred per year – and remarkably stable, GARCH(1,1) having become a piece of empirical infrastructure rather than a research topic. Kyle (1985) shows the longest and flattest reception profile in the microstructure block:

[1] The Python implementation, `citations-per-year` (`cpy`), is available at https://github.com/jannovotny/citations-per-year (placeholder; private repository to be opened on publication). The tool was prototyped iteratively using Claude Code (Anthropic, `claude-opus-4-8`); the full prompt-and-iteration history is preserved in the repository's commit log.

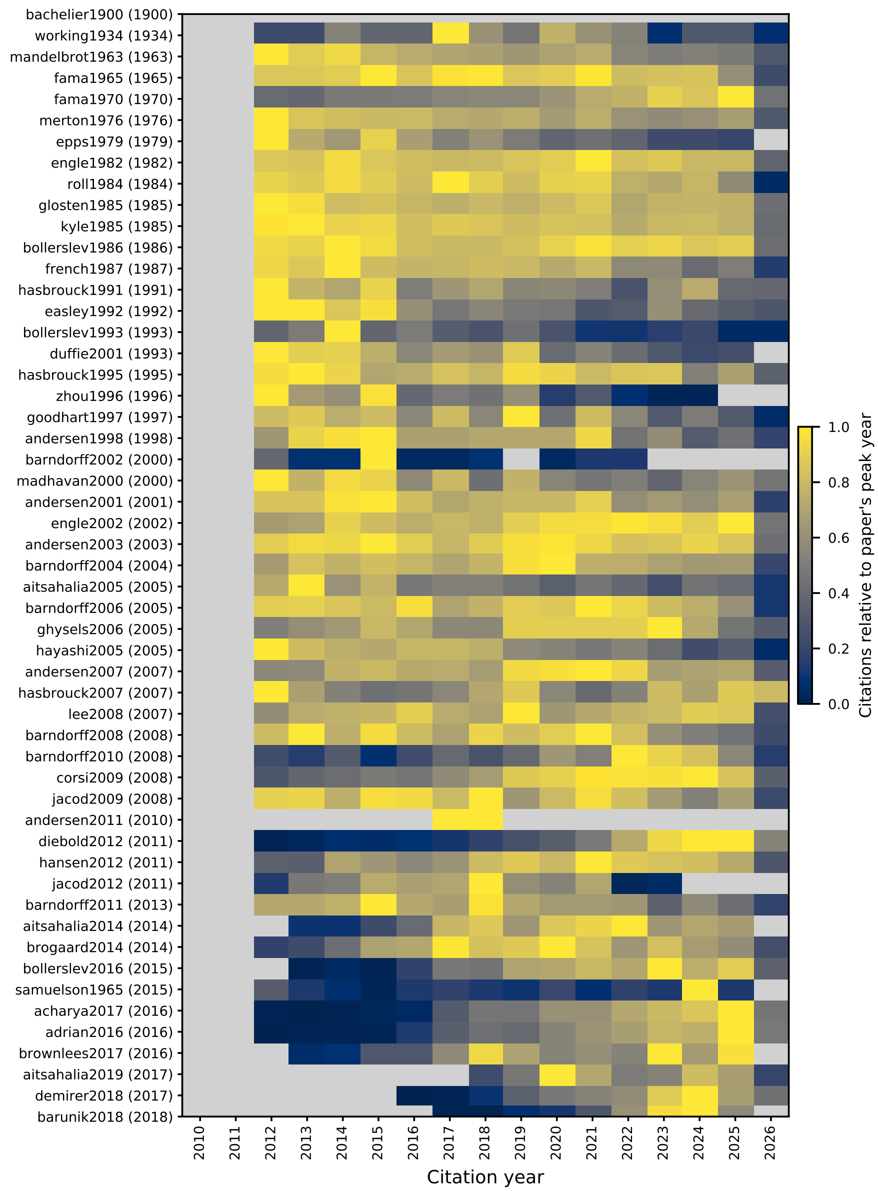

Fig. 3.1: Per-paper-normalised citation heatmap for the references of this chapter, spanning both the original outline and the recent developments of Sections 3.13–3.17. Rows are sorted by publication year (oldest at top); the colour scale is *cividis*, chosen for black-and-white print legibility. Each row is normalised independently: the colour at a given year shows the citation count relative to that paper's peak year. Grey cells indicate years for which OpenAlex returns no data. One early citation (Bachelier 1900) falls outside the indexable window and appears as an empty row.

its primitives, adverse selection, market depth, and Kyle's lambda, have entered the vocabulary of the field, and the paper is still cited as the reference for them. Fama (1970) remains the obligatory citation for any work on return predictability, its rate high and broadly constant despite half a century of accumulated qualifications. Finally, Andersen, Bollerslev, Diebold and Labys (2003) underpins both the modelling and the forecasting literatures of Section 3.10, and its high, stable citation rate reflects that realised-volatility measurement has been absorbed as a tool rather than treated as an open research question. That ARCH and GARCH should head this list is not in tension with the chapter's emphasis on measurement over specification: both endure as low-frequency infrastructure, applied where high-frequency data are unavailable or unnecessary, so that parametric filtering and direct measurement coexist by horizon rather than competing.

By contrast, several papers show citation patterns that taper off in the most recent years. This is not necessarily evidence that the underlying contribution has been superseded. For more recent papers such as those of Adrian and Brunnermeier (2016) and Brogaard et al. (2014), the citation curve is still rising and the 2025–26 counts are typically lower because OpenAlex's index lags publications by twelve to eighteen months. The recent contributions surveyed in Section 3.13 sit squarely in this growth phase: the systemic-risk and connectedness measures of Acharya et al. (2017), Demirer et al. (2018), and Baruník and Křehlík (2018) show the steepest rising trajectories in the figure, with citation rates that climb steadily through the end of the sample, while the realised-measure forecasting models of Hansen et al. (2012) and Bollerslev et al. (2016) have settled into steadier, mature rates. For older empirical contributions such as those of Roll (1984) and French et al. (1987), the gradual decline reflects that the specific findings have been absorbed into a more general framework and survive primarily as historical references rather than as active building blocks; the multivariate realised kernel of Barndorff-Nielsen et al. (2011) shows the same absorption on a shorter horizon, its citation rate easing back from a mid-2010s peak as it became part of the standard covariance-estimation toolkit introduced in Section 3.8.

The application-oriented sections added to this survey – the high-frequency identification of macroeconomic and monetary news in Section 3.14, price discovery in Section 3.15, the connectedness and variance-risk literature of Section 3.16, and the machine-learning and digital-asset frontier of Section 3.17 – divide along the same lines, and they supply the two steepest trajectories in the entire figure. The network-connectedness measure of Diebold and Yilmaz (2014) rises almost without interruption from around sixty citations in 2015 to over eight hundred a decade later, and the machine-learning asset-pricing study of Gu et al. (2020) climbs from negligible counts before 2020 to several hundred per year by the end of the sample; both sit squarely in the growth phase, and their reception has not yet turned over. A second group of recent applications – the monetary-policy surprises of Gertler and Karadi (2015), the information-shock decomposition of Jarociński and Karadi (2020), the cryptocurrency-market studies of Makarov and Schoar (2020) and Cong et al. (2021), and the deep-hedging approach of Buehler et al. (2019) – shows the same upward profile on a smaller scale, several having risen to a high recent plateau.

Against these, a handful of the older application papers display the flat, high profile that we associated above with ageing well: Kuttner (2001) and Gürkaynak et al. (2005) on monetary-policy surprises, Hasbrouck (1995) on price discovery, and Bollerslev et al. (2009) on the variance risk premium each hold a large and broadly stable citation rate across the window, having become standard references in their applied literatures rather than active research topics. The foundational announcement and price-discovery papers they built on – Fleming and Remolona (1999), Balduzzi et al. (2001), and Andersen, Bollerslev, Diebold and Vega (2003), together with the model-free implied-variance work of Britten-Jones and Neuberger (2000) – show instead the gentle decline of contributions absorbed into a now-routine empirical practice.

Caveats

The exercise has two limitations worth flagging. First, OpenAlex's per-year citation history typically covers the last twelve to fifteen years, so the heatmap describes the post-2010 reception window rather than the full lifetime. Several foundational references – Bachelier (1900), Samuelson (1965), the very oldest microstructure work – are absent or sparsely covered for this reason, and our conclusions about "ageing well" are conditional on that window. Second, citation counts measure visibility in the subsequent literature, not scientific correctness, methodological soundness, or current usefulness in practice. The heatmap is therefore best read as a coarse, replicable complement to – rather than a substitute for – the methodological judgements developed in the preceding sections.

3.20 Conclusions

Financial econometrics has undergone one of the most profound methodological transformations in applied economics. The discipline evolved from analysing relatively sparse observations of asset prices at daily or lower frequencies to exploiting millions of intraday observations generated by modern electronic markets. This expansion in data availability has not merely increased the precision of empirical analysis; it has fundamentally altered the questions that can be asked and the methods by which they are answered.

The historical perspective developed throughout this chapter shows that many of the central ideas of financial econometrics have proved remarkably durable. Bachelier's diffusion model remains the benchmark for continuous-time price dynamics, Samuelson's martingale formulation continues to underpin modern asset-pricing theory, Mandelbrot's emphasis on heavy tails and non-Gaussian behaviour anticipated later developments in stochastic volatility and jump models, while Fama's efficient-market framework continues to provide the natural benchmark against which predictability is assessed. Rather than replacing these foundations, high-frequency econometrics has substantially refined their empirical interpretation.

Perhaps the most important methodological development has been the transition from specification-based to measurement-based econometrics. For much of the twentieth century, economically relevant quantities such as volatility, covariation, liquidity and information arrival were treated as latent variables whose dynamics had to be inferred through increasingly sophisticated parametric models. The availability of transaction-level data, together with advances in continuous-time probability theory, allowed many of these quantities to become directly measurable. Realised volatility, realised covariance, jump measures, and liquidity indicators exemplify this transition. The emphasis has therefore shifted from specifying increasingly complex stochastic models towards constructing statistically robust estimators capable of extracting information from noisy, irregularly spaced observations.

A second lesson concerns the central role of market microstructure. Early financial econometrics often regarded the trading mechanism as a nuisance that obscured the efficient-price process. Modern high-frequency econometrics instead recognises that market microstructure is simultaneously a source of statistical complications and a source of economically valuable information. Bid-ask spreads, order flow, limit-order books, and trading intensity contain information about liquidity provision, information asymmetry, price discovery, and systemic risk. Consequently, successful econometric methods must account explicitly for the institutional environment within which prices are generated.

A third lesson is the increasing integration of financial econometrics with neighbouring disciplines. Modern high-frequency methods combine continuous-time stochastic calculus, nonparametric statistics, financial economics, market microstructure theory, computational methods, and increasingly machine learning. Likewise, applications now extend well beyond volatility measurement to monetary-policy identification, macroeconomic announcement effects, systemic-risk monitoring, connectedness analysis, algorithmic trading, portfolio allocation, and digital asset markets. High-frequency econometrics has therefore evolved from a specialised branch of financial statistics into a general empirical framework for analysing information transmission in modern financial markets.

The retrospective perspective adopted in this chapter also suggests that not every methodological innovation has aged equally well. Some early assumptions—including constant volatility, Gaussian returns, and frictionless trading—have proved increasingly difficult to reconcile with high-frequency evidence. By contrast, concepts such as arbitrage-free pricing, the martingale property, volatility persistence, and the importance of information arrival have remained central, although they now operate within richer econometric frameworks that explicitly recognise market frictions, heterogeneous information, and complex dependence structures. The evolution of financial econometrics has therefore been characterised less by the replacement of earlier theories than by their progressive refinement.

Looking ahead, several developments are likely to shape the next generation of research. The increasing availability of complete limit-order-book data, alternative trading venues, cryptocurrency markets, decentralised finance, and ultra-high-frequency transaction records continues to expand both the opportunities and the econometric challenges associated with financial data. At the same time, machine-learning methods and artificial intelligence offer increasingly powerful tools for prediction and pattern recognition, although their successful application will continue to depend on rigorous statistical inference, careful identification strategies, and economically interpretable models. Rather than replacing econometric theory, these developments reinforce the importance of combining flexible computational methods with sound probabilistic foundations.

Ultimately, the history reviewed in this chapter illustrates that advances in financial econometrics have been driven not simply by improvements in statistical methodology but by the interaction between theory, data, and technology. As financial markets have become faster, more interconnected, and increasingly digital, econometric methods have evolved from modelling latent quantities under strong parametric assumptions towards robust measurement, identification, and inference based on rich transaction-level information. This evolution has transformed financial econometrics from a discipline primarily concerned with explaining asset returns into one capable of measuring, in real time, the complex dynamics of modern financial markets.

References

Acharya, V. V., Pedersen, L. H., Philippon, T. & Richardson, M. (2017). Measuring systemic risk. *Review of Financial Studies*, *30*(1), 2–47. doi: 10.1093/rfs/hhw088

Adrian, T. & Brunnermeier, M. K. (2016). Covar. *American Economic Review*, *106*(7), 1705–1741. doi: 10.1257/aer.20120555

Aït-Sahalia, Y. (2002). Telling from discrete data whether the underlying continuous-time model is a diffusion. *Journal of Finance*, *57*(5), 2075–2112. doi: 10.1111/1540-6261.00489

Aït-Sahalia, Y., Fan, J. & Xiu, D. (2010). High-frequency covariance estimates with noisy and asynchronous financial data. *Journal of the American Statistical Association*, *105*(492), 1504–1517. doi: 10.1198/jasa.2010.tm10163

Aït-Sahalia, Y. & Jacod, J. (2014). *High-frequency financial econometrics*. Princeton: Princeton University Press. doi: 10.1515/9781400850327

Aït-Sahalia, Y., Mykland, P. A. & Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies*, *18*(2), 351–416. doi: 10.1093/rfs/hhi016

Aït-Sahalia, Y., Mykland, P. A. & Zhang, L. (2011). Ultra high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics*, *160*(1), 160–175. doi: 10.1016/j.jeconom.2010.03.028

Aït-Sahalia, Y. & Xiu, D. (2019). Principal component analysis of high-frequency data. *Journal of the American Statistical Association*, *114*(525), 287–303. doi: 10.1080/01621459.2017.1401542

Andersen, T. G. & Bollerslev, T. (1998). Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, *53*(1), 219–265. doi: 10.1111/0022-1082.85732

Andersen, T. G., Bollerslev, T. & Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics*, *89*(4), 701–720. doi: 10.1162/rest.89.4.701

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, *96*(453), 42–55. doi: 10.1198/016214501750332965

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, *71*(2), 579–625. doi: 10.1111/1468-0262.00418

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Vega, C. (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, *93*(1), 38–62. doi: 10.1257/000282803321455151

Andersen, T. G., Bollerslev, T. & Huang, X. (2011). A reduced form framework for modeling volatility of speculative prices based on realized variation measures. *Journal of Econometrics*, *160*(1), 176–189. doi: 10.1016/j.jeconom.2010.03.029

Bachelier, L. (1900). Théorie de la spéculation. *Annales scientifiques de l'École Normale Supérieure*, *17*, 21–86.

Balduzzi, P., Elton, E. J. & Green, T. C. (2001). Economic news and bond prices: Evidence from the U.S. Treasury market. *Journal of Financial and Quantitative Analysis*, *36*(4), 523–543. doi: 10.2307/2676225

Bandi, F. M. & Russell, J. R. (2008). Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies*, *75*(2), 339–369. doi: 10.1111/j.1467-937x.2008.00474.x

Barigozzi, M. & Brownlees, C. (2019). NETS: Network estimation for time series. *Journal of Applied Econometrics*, *34*(3), 347–364. doi: 10.1002/jae.2676

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. & Shephard, N. (2008). Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, *76*(6), 1481–1536. doi: 10.3982/ecta6495

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. & Shephard, N. (2011). Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics*, *162*(2), 149–169. doi: 10.1016/j.jeconom.2010.07.009

Barndorff-Nielsen, O. E., Kinnebrock, S. & Shephard, N. (2010). Measuring downside risk: Realised semivariance. In *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle* (pp. 117–136). (DOI points to the SSRN working paper; the published version is a chapter in Bollerslev, Russell & Watson (Eds.), Oxford University Press.) doi: 10.2139/ssrn.1262194

Barndorff-Nielsen, O. E. & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B*, *64*(2), 253–280. doi: 10.1111/

1467-9868.00336

Barndorff-Nielsen, O. E. & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, *2*(1), 1–37. doi: 10.1093/jjfinec/nbh001

Barndorff-Nielsen, O. E. & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, *4*(1), 1–30. doi: 10.1093/jjfinec/nbi022

Baruník, J. & Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, *16*(2), 271–296. doi: 10.1093/jjfinec/nby001

Billio, M., Getmansky, M., Lo, A. W. & Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, *104*(3), 535–559. doi: 10.1016/j.jfineco.2011.12.010

Black, F. (1976). Studies of stock price volatility changes. In *Proceedings of the 1976 meetings of the american statistical association, business and economic statistics section* (pp. 177–181). (No DOI available (ASA proceedings paper).)

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. doi: 10.1016/0304-4076(86)90063-1

Bollerslev, T. & Domowitz, I. (1993). Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance*, *48*(4), 1421–1443. doi: 10.1111/j.1540-6261.1993.tb04760.x

Bollerslev, T., Patton, A. J. & Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, *192*(1), 1–18. doi: 10.1016/j.jeconom.2015.10.007

Bollerslev, T., Tauchen, G. & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, *22*(11), 4463–4492. doi: 10.1093/rfs/hhp008

Britten-Jones, M. & Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance*, *55*(2), 839–866. doi: 10.1111/0022-1082.00228

Brogaard, J., Hendershott, T. & Riordan, R. (2014). High-frequency trading and price discovery. *Review of Financial Studies*, *27*(8), 2267–2306. doi: 10.1093/rfs/hhu032

Brownlees, C. & Engle, R. F. (2017). SRISK: A conditional capital shortfall measure of systemic risk. *Review of Financial Studies*, *30*(1), 48–79. doi: 10.1093/rfs/hhw060

Buehler, H., Gonon, L., Teichmann, J. & Wood, B. (2019). Deep hedging. *Quantitative Finance*, *19*(8), 1271–1291. doi: 10.1080/14697688.2019.1571683

Carr, P. & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, *22*(3), 1311–1341. doi: 10.1093/rfs/hhn038

Cong, L. W., Li, Y. & Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies*, *34*(3), 1105–1155. doi: 10.1093/rfs/hhaa089

Cootner, P. H. (Ed.). (1964). *The random character of stock market prices*. Cambridge, MA: MIT Press.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, *7*(2), 174–196. doi: 10.1093/jjfinec/nbp001

Dacorogna, M. M., Gençay, R., Müller, U., Olsen, R. B. & Pictet, O. V. (2001). *An introduction to high-frequency finance*. San Diego: Academic Press. doi: 10.1016/b978-0-12-279671-5.x5000-x

de Jong, F. (2002). Measures of contributions to price discovery: A comparison. *Journal of Financial Markets*, *5*(3), 323–327. doi: 10.1016/S1386-4181(02)00015-9

Demirer, M., Diebold, F. X., Liu, L. & Yilmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, *33*(1), 1–15. doi: 10.1002/jae.2585

Diebold, F. X. & Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal*, *119*(534), 158–171. doi: 10.1111/j.1468-0297.2008.02208.x

Diebold, F. X. & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, *28*(1), 57–66. doi: 10.1016/j.ijforecast.2011.02.006

Diebold, F. X. & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, *182*(1), 119–134. doi: 10.1016/j.jeconom.2014.04.012

Duffie, D. (2001). *Dynamic asset pricing theory* (3rd ed.). Princeton, NJ: Princeton University Press.

Easley, D., Kiefer, N. M., O'Hara, M. & Paperman, J. B. (1996). Liquidity, information, and infrequently traded stocks. *Journal of Finance*, *51*(4), 1405–1436. doi: 10.1111/j.1540-6261.1996.tb04074.x

Easley, D. & O'Hara, M. (1992). Time and the process of security price adjustment. *Journal of Finance*, *47*(2), 577–605. doi: 10.1111/j.1540-6261.1992.tb04402.x

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, *50*(4), 987–1007. doi: 10.2307/1912773

Engle, R. F. (1996). *The Econometrics of Ultra-High-Frequency Data* (NBER Working Paper No. 5816). National Bureau of Economic Research. doi: 10.3386/w5816

Engle, R. F. (2000). The econometrics of ultra-high-frequency data. *Econometrica*, *68*(1), 1–22.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics*, *20*(3), 339–350. doi: 10.1198/073500102288618487

Engle, R. F. & Manganelli, S. (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics*, *22*(4), 367–381. doi: 10.1198/073500104000000370

Engle, R. F. & Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, *66*(5), 1127–1162.

doi: 10.2307/2999632

Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, *74*(366), 291–298. doi: 10.1080/01621459.1979.10482508

Fama, E. F. (1965). The behavior of stock-market prices. *Journal of Business*, *38*(1), 34–105. doi: 10.1086/294743

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, *25*(2), 383–417. doi: 10.2307/2325486

Fleming, M. J. & Remolona, E. M. (1999). Price formation and liquidity in the U.S. Treasury market: The response to public information. *Journal of Finance*, *54*(5), 1901–1915. doi: 10.1111/0022-1082.00172

French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, *19*(1), 3–29. doi: 10.1016/0304-405x(87)90026-2

Garber, P. M. (1989). Tulipmania. *Journal of Political Economy*, *97*(3), 535–560. doi: 10.1086/261615

Gertler, M. & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76. doi: 10.1257/mac.20130329

Ghysels, E., Santa-Clara, P. & Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, *131*(1-2), 59–95. doi: 10.1016/j.jeconom.2005.01.004

Glosten, L. R., Jagannathan, R. & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, *48*(5), 1779–1801. doi: 10.1111/j.1540-6261.1993.tb05128.x

Glosten, L. R. & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, *14*(1), 71–100. doi: 10.1016/0304-405x(85)90044-3

Goldgar, A. (2007). *Tulipmania: Money, honor, and knowledge in the dutch golden age*. Chicago: University of Chicago Press.

Gonzalo, J. & Granger, C. W. J. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, *13*(1), 27–35. doi: 10.1080/07350015.1995.10524576

Goodhart, C. A. E. & O'Hara, M. (1997). High frequency data in financial markets: Issues and applications. *Journal of Empirical Finance*, *4*(2-3), 73–114. doi: 10.1016/s0927-5398(97)00003-0

Gu, S., Kelly, B. & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, *33*(5), 2223–2273. doi: 10.1093/rfs/hhaa009

Gürkaynak, R. S., Sack, B. & Swanson, E. T. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, *1*(1), 55–93.

Hansen, P. R., Huang, Z. & Shek, H. H. (2012). Realized GARCH: A joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, *27*(6), 877–906. doi: 10.1002/jae.1234

Hansen, P. R. & Lunde, A. (2005). A forecast comparison of volatility models:

Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, *20*(7), 873–889. doi: 10.1002/jae.800

Hansen, P. R. & Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business and Economic Statistics*, *24*(2), 127–161. doi: 10.1198/073500106000000071

Harris, F. H. d., McInish, T. H., Shoesmith, G. L. & Wood, R. A. (1995). Cointegration, error correction, and price discovery on informationally linked security markets. *Journal of Financial and Quantitative Analysis*, *30*(4), 563–579. doi: 10.2307/2331277

Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, *46*(1), 179–207. doi: 10.1111/j.1540-6261.1991.tb03749.x

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, *50*(4), 1175–1199. doi: 10.1111/j.1540-6261.1995.tb04054.x

Hasbrouck, J. (2007). *Empirical market microstructure*. doi: 10.1093/oso/9780195301649.001.0001

Hautsch, N. (2012). *Econometrics of financial high-frequency data*. Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-21925-2

Hayashi, T. & Yoshida, N. (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli*, *11*(2), 359–379. doi: 10.3150/bj/1116340299

Hendershott, T., Jones, C. M. & Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *Journal of Finance*, *66*(1), 1–33. doi: 10.1111/j.1540-6261.2010.01624.x

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, *6*(2), 327–343. doi: 10.1093/rfs/6.2.327

Huang, R. D. & Stoll, H. R. (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies*, *10*(4), 995–1034. doi: 10.1093/rfs/10.4.995

Jacod, J., Li, Y., Mykland, P. A., Podolskij, M. & Vetter, M. (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications*, *119*(7), 2249–2276. doi: 10.1016/j.spa.2008.11.004

Jacod, J. & Protter, P. (2012). *Discretization of processes*. doi: 10.1007/978-3-642-24127-7

Jarociński, M. & Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, *12*(2), 1–43. doi: 10.1257/mac.20180090

Jiang, G. J. & Tian, Y. S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, *18*(4), 1305–1342. doi: 10.1093/rfs/hhi027

Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J. & Tivnan, B. (2013). Abrupt rise of new machine ecology beyond human response time. *Scientific Reports*, *3*, 2627. doi: 10.1038/srep02627

Kirilenko, A., Kyle, A. S., Samadi, M. & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *Journal of Finance*, *72*(3), 967–998. doi: 10.1111/jofi.12498

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, *47*(3), 523–544. doi: 10.1016/S0304-3932(01)00055-1

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, *53*(6), 1315–1335. doi: 10.2307/1913210

Lee, S. S. & Mykland, P. A. (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies*, *21*(6), 2535–2563. doi: 10.1093/rfs/hhm056

Liu, Y. & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, *34*(6), 2689–2727. doi: 10.1093/rfs/hhaa113

Lo, A. W. & MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies*, *1*(1), 41–66. doi: 10.1093/rfs/1.1.41

Lucca, D. O. & Moench, E. (2015). The pre-FOMC announcement drift. *Journal of Finance*, *70*(1), 329–371. doi: 10.1111/jofi.12196

Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, *3*(3), 205–258. doi: 10.1016/s1386-4181(00)00007-0

Makarov, I. & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, *135*(2), 293–319. doi: 10.1016/j.jfineco.2019.07.001

Mancini, C. (2009). Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, *36*(2), 270–296. doi: 10.1111/j.1467-9469.2008.00622.x

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, *36*(4), 394–419. doi: 10.1086/294632

Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, *16*(4), 712–740. doi: 10.1016/j.finmar.2013.06.006

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, *3*(1-2), 125–144. doi: 10.1016/0304-405x(76)90022-2

Mykland, P. A. & Zhang, L. (2012). The econometrics of high-frequency data. In M. Kessler, A. Lindner & M. Sørensen (Eds.), *Statistical methods for stochastic differential equations* (pp. 109–190). Boca Raton: Chapman & Hall/CRC. doi: 10.1201/b12126-3

Nakamura, E. & Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, *133*(3), 1283–1330. doi: 10.1093/qje/qjy004

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, *59*(2), 347–370. doi: 10.2307/2938260

O'Hara, M. (1995). *Market microstructure theory*. Cambridge, MA: Blackwell.

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*, *116*(2), 257–270. doi: 10.1016/j.jfineco.2015.01.003

Osborne, M. F. M. (1959). Brownian motion in the stock market. *Operations Research*, *7*(2), 145–173. doi: 10.1287/opre.7.2.145

Patton, A. J. & Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, *97*(3), 683–697. doi: 10.1162/rest_a_00503

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, *39*(4), 1127–1139. doi: 10.1111/j.1540-6261.1984.tb03897.x

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, *6*(2), 41–49.

Samuelson, P. A. (1973). Proof that properly discounted present values of assets vibrate randomly. *Bell Journal of Economics and Management Science*, *4*(2), 369–374. doi: 10.2307/3003046

Schwert, G. W. (1989). Why does stock market volatility change over time? *Journal of Finance*, *44*(5), 1115–1153. doi: 10.1111/j.1540-6261.1989.tb02647.x

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, *71*(3), 421–436.

Sirignano, J. & Cont, R. (2019). Universal features of price formation in financial markets: Perspectives from deep learning. *Quantitative Finance*, *19*(9), 1449–1459. doi: 10.1080/14697688.2019.1622295

Wood, R. A., McInish, T. H. & Ord, J. K. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, *40*(3), 723–739. doi: 10.2307/2327796

Working, H. (1934). A random-difference series for use in the analysis of time series. *Journal of the American Statistical Association*, *29*(185), 11–24. doi: 10.1080/01621459.1934.10502683

Zhang, L., Mykland, P. A. & Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, *100*(472), 1394–1411. doi: 10.1198/016214505000000169

Zhang, Z., Zohren, S. & Roberts, S. (2019). Deeplob: Deep convolutional neural networks for limit order books. *IEEE Transactions on Signal Processing*, *67*(11), 3001–3012.

Zhou, B. (1996). High-frequency data and volatility in foreign-exchange rates. *Journal of Business and Economic Statistics*, *14*(1), 45–52. doi: 10.2307/1392098

Index