

Introduction

100+ Years of Financial Risk Measurement and Management

Francis X. Diebold

Abstract

I selectively survey several key strands of literature on financial risk measurement and management. I begin by showing why there's a need for financial risk measurement and management, and then I turn to relevant aspects of return distributions and volatility fluctuations, with implicit emphasis on market risk for equities. I then treat market risk for bonds, focusing on the yield curve, with its nuances and special structure. In addition to market risk measurement and management, I also discuss aspects of measuring credit risk, operational risk, systemic risk, and underlying business-cycle risk. I nevertheless also stress the limits of statistical analysis, and the associated importance of respecting the unknown and the unknowable.

1 Introduction

In this book I collect many of the key papers that contribute to my vision of risk measurement in its role as the key input to successful risk management. I focus mostly on market risk, but I also treat aspects of credit risk, operational risk, systemic risk, business cycle risk, and the special structure of risks associated with the yield curve. Of course any attempt to distill the financial risk measurement literature into a handful of papers borders on the absurd, yet I think that the collection coheres well, and tells the right story. I include mostly well-known papers, but I omit many equally well-known papers, whether to respect space constraints or to maximize coherence with my chosen themes. Conversely, I also include several less-well-known papers deserving of wider attention.

This introductory chapter provides an interpretive overview. I begin in section 2 by emphasizing the *need* for financial risk management, given that, perhaps surprisingly, traditional economic theory suggests no need. The key lies in various real-world “details” ignored by the traditional theory. Those details turn out not to be details at all; rather, they are now recognized as central to the story. I then tell that story in the rest of the chapter. I emphasize the non-normality of financial asset returns in section 3, and I emphasize their fluctuating volatility in section 4. I focus on the more specialized issues associated with market risk in bond yields (that is, the yield curve and the factors that drive it) in section 5. I highlight a variety of issues centered around non-market risk, including credit risk, operational risk, and systemic risk, in section 6. I draw attention to a key macroeconomic fundamental risk driver, the business cycle, in section 7. I emphasize the limits of statistical risk measurement in section 8, and I offer suggestions for further reading in section 9.

2 There is a Role for Financial Risk Measurement and Management

Much of modern economics, including financial economics, begins with the famous Arrow-Debreu model of general equilibrium. The classic statement (under conditions of certainty) is Arrow and Debreu (1954). Arrow (1964), written around the same time as Arrow and Debreu (1954), makes the crucial progression to consideration of general equilibrium in risky environments.¹ Arrow effectively shows how to manage risks using the now-famous

¹As noted in Arrow (1964), “This paper was originally read at the Colloque sur les Fondements et Applications de la Theorie du Risque en Econometrie of the Centre Nationale de la Recherche Scientifique,

theoretical construct of "Arrow-Debreu securities." The i th such Arrow-Debreu security, s_{it} , pays \$1 (say) if contingency i occurs at time t , and \$0 otherwise.

Many standard securities have payoffs that vary with certain contingencies and that therefore might crudely approximate those of certain Arrow-Debreu securities. Much of financial engineering can be interpreted as designing portfolios of Arrow-Debreu securities, or closely-related securities, with desired "payoff profiles." Put and call options are classic examples of derivative assets with such state-contingent payoffs. Moreover, precise Arrow-Debreu securities now trade, under the name of "binary options" (also sometimes called digital, or all-or-nothing options). They are growing in popularity and traded on-exchange.²

In a certain sense, Arrow-Debreu securities are the beginning and the end of the hedging story. Exposure to risk i at time t can be hedged immediately by going long or short the appropriate amount of s_{it} . In the limit, if s_{it} exists for every possible i and t (that is, for all states and times), one can hedge any risk by holding an appropriate portfolio of Arrow-Debreu securities. But there is a crucial caveat: the possibility and efficacy of Arrow-Debreu hedging relies on strong assumptions – essentially complete and perfect capital markets – that provide a natural theoretical benchmark but that may of course fail in reality.

In theory Arrow-Debreu hedging can be done by any economic entity (private agents, firms, governments, ...), and it seems clear that risk-averse private agents may want to avail themselves of the opportunity. But what about firms? Need they be risk averse? Put differently, do firms need to engage in risk management, via Arrow-Debreu securities or any other means? Again working in an idealized setting amounting to perfect capital markets, Modigliani and Miller (1958) answer this question in the negative, showing that a firm's cost of capital is invariant to its risk profile as measured by leverage. The seemingly-counterintuitive – indeed initially-shocking – Modigliani-Miller result is actually quite natural upon reflection. With perfect capital markets, investors can allocate firm risk in any way desired by holding appropriate portfolios of financial assets; hence neither investors nor firms need concern themselves with firm capital structure. Put differently, in the Modigliani-Miller world there is simply no role for risk management by firms.

A key tension therefore arose in theoretical finance circa 1958. It was clearly understood how firms could hedge in principle using Arrow-Debreu securities, but the Modigliani-Miller theorem said there was no need! The tension is resolved by recognizing that, as with Arrow-Debreu hedging, the Modigliani-Miller theorem relies on assumptions – again, essentially

Paris, France, on May 13, 1952 and appeared in French in the proceedings of the colloquium, published by the Centre Nationale under the title, *Econometric*, 1953."

²See, for example, http://www.cboe.com/products/indexopts/bsz_spec.aspx.

perfect capital markets – that fail in reality, due to distortionary taxes, costs of financial distress, and so on.

Hence, despite its invaluable role as a theoretical benchmark, the Modigliani-Miller theorem fails. Indeed a large part of the last half-century of corporate finance theory has been devoted to confronting the failure of Modigliani-Miller, and determining what, if anything, can be said in more realistic environments with imperfect capital markets. Crucially, the failure of Modigliani-Miller creates a role for firm-level risk management. Froot and Stein (1998) is a key contribution in that regard, rigorously illuminating in realistic environments the roles of hedging, capital structure policy, capital budgeting policy, and crucially, their joint endogenous determination in equilibrium.

If finance theorists took some forty years from Modigliani-Miller to Froot-Stein to appreciate the value of firm-level risk management, finance practitioners understood much more quickly. In a sense good hedges were always sought, and by the 1960s at least some options were trading. But even by 1970 option valuation remained rather poorly understood, which limited the use of options and therefore hindered the growth of options markets. Ironically and unfortunately, simultaneously with this bottleneck in the supply of effective options-based hedging vehicles, demand shot upward as the Bretton-Woods exchange rate system began its collapse in the early 1970s, creating the prospect of radically increased exchange rate risk.

Enter Black and Scholes (1973), who solved the option pricing problem in a single beautiful paper (at least under assumptions now recognized as heroic, involving not only perfect capital markets, but also specific dynamics and distributions of the underlying spot assets, the failure of which will concern us for much of the rest of this chapter). The Black-Scholes timing couldn't have been better, as the world moved to floating exchange rates in 1973, the same year that the Black-Scholes paper was published. Firms finally had a key tool to operationalize Arrow-Debreu hedging in the real world – options, priced rigorously – and the derivatives industry was born.

The famous “Black-Scholes formula” relates the value of a European option to the parameters of the option contract (strike price, time to maturity), prevailing general financial market conditions (the interest rate), and prevailing spot market conditions (current spot price, current spot volatility). The Black-Scholes formula for the price of a European call on a non-dividend-paying stock is:

$$C = N(d_1)S - N(d_2)Ke^{-r\tau} \quad (1)$$

where

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}$$
$$d_2 = \frac{\ln(S/K) + (r - \sigma^2/2)\tau}{\sigma\sqrt{\tau}},$$

and where $N(\cdot)$ is standard normal cdf, τ is time to maturity, S is underlying spot price, K is strike price, r is the risk-free rate, and σ is the underlying spot return volatility.

Perhaps surprisingly at first, the last-mentioned option-price determinant (current spot volatility) plays an absolutely central role. The intuitive reasoning is simple: the higher is volatility, the more likely is the spot price to move such that the option is “in the money,” and hence the more valuable is the option. That fact, together with the fact that the other option-price determinants are more-or-less readily and reliably measured, reveals that option prices – and changes in those prices over time – primarily reflect market views on volatility. Effectively, volatility is the “commodity” traded in options markets.³

Indeed, by looking at the market price of an option, one can infer the market’s views about the corresponding path of future spot volatility, at least under the risk-neutral measure relevant for options pricing, and under the previously-mentioned heroic assumptions needed for validity of Black-Scholes. Mathematically, because Black-Scholes gives an option’s price as a function of underlying spot volatility, $C = f(\sigma, \dots)$, we can invert Black-Scholes to find the volatility that rationalizes the market price of the option, $\sigma = f^{-1}(C, \dots)$. This quantity is often called option-implied volatility, or simply *implied volatility*. Implied volatilities on broad equity market indexes have emerged as widely-followed “market fear gauges.” A leading example is the VIX, as described in Whaley (1993). Early implementations (“old VIX”) averaged implied volatilities across strike prices, whereas more recent implementations (“new VIX”) infer implied volatility directly from traded options prices. VIX is now heavily traded on CBOE, as are derivatives on VIX.

3 Asset Returns are Unconditionally Fat-Tailed

Stochastic financial modeling, and in many respects much of the theory of stochastic processes more generally, traces to Bachelier (1900), whose amazing doctoral dissertation was

³There are many flavors of “volatility” and we shall use the term rather loosely, with the meaning clear from context. Depending on the situation, volatility might refer, for example, to conditional variance or standard deviation, unconditional variance or standard deviation, realized variance or standard deviation, or options-implied variance or standard deviation.

largely unnoticed for many decades. In modern parlance and notation, Bachelier proposed using geometric Brownian motion as a dynamic model of spot price P ; that is, he proposed modeling log price as a diffusion:

$$dp = \mu dt + \sigma dW, \quad (2)$$

where $p = \ln P$, μ is instantaneous drift, σ is instantaneous volatility, and dW is an increment of standard Brownian motion. The discrete-time analog is a geometric random walk with Gaussian shocks,

$$\Delta p_t = \mu + \sigma \varepsilon_t \quad (3a)$$

$$\varepsilon_t \sim iid N(0, 1), \quad (3b)$$

where $p_t = \ln P_t$, μ is per-period drift and σ is per-period volatility. In what follows it will prove convenient to define the continuously-compounded return $r_t = \Delta \ln P_t$ and to drop the drift (as it is not central to our concerns), writing

$$r_t = \sigma \varepsilon_t \quad (4a)$$

$$\varepsilon_t \sim iid N(0, 1). \quad (4b)$$

Importantly, Bachelier also explored, both theoretically and empirically, some of the asset-pricing implications of his geometric diffusion assumption for spot price.

Let us think about Bachelier's normality assumption (4b), and its implications for measuring volatility, which is a central concept in risk analysis. (We have already seen the centrality of volatility in the Black-Scholes options pricing formula (1).) The variance of a random variable measures its spread, or dispersion, around its mean. It is well-known that Gaussian random variables are completely characterized by their mean and variance. Hence in Gaussian environments the variance σ^2 (or its square root, the standard deviation σ) contains all information about volatility. It is the uniquely-relevant and all-encompassing volatility measure.

The normal-centric worldview enshrined in Bachelier's model implicitly motivated numerous academics and practitioners to equate risk with standard deviation.⁴ A classic risk

⁴Of course the Bachelier motivation was implicit rather than explicit, as his work went largely unnoticed for more than half a century following its appearance. But Gaussian thinking was very much in the air through much of the first half of the twentieth century, due to the vibrant work in stochastic process theory and general probability and statistical theory undertaken then (e.g., sophisticated central-limit theorems). Indeed central-limit theorems are responsible for converting the increments of standard Brownian motion in

measurement/management example is Markowitz (1952) who works in exclusively Gaussian ("mean-variance") mode in the context of portfolio risk minimization. The idea is that investors like expected returns (μ) but dislike risk (σ), so that for any given expected return investors should seek the smallest risk. Related, investors should not focus on finding portfolios with high expected returns μ *per se*, but rather on finding portfolios with high *risk adjusted* expected returns μ/σ , a quantity closely-related to the so-called Sharpe ratio.

The subsequent and more sophisticated capital asset pricing model (CAPM) of Sharpe (1964) is similar but adopts a risk adjustment different from that of the Sharpe ratio, recognizing that the risk of a portfolio concerns not just its variation, but also its *covariation* with a risk factor. In the CAPM one "controls" for risk via linear regression on a market risk factor, and one measures risk by the corresponding regression coefficient, the famous CAPM "beta".⁵ In multi-factor CAPM extensions, one uses multiple linear regression to control not only for a market risk factor, but also for a few additional risk factors. The key insight for our purposes is that the CAPM and its relatives maintain the thoroughly-Gaussian spirit of Bachelier, Markowitz and Tobin. In fact the CAPM requires normality in the absence of restrictive assumptions regarding preferences, as it is based on linear regressions, and linearity of regression functions is a Gaussian phenomenon. Related, linear regression coefficients like the CAPM beta are simply ratios of covariances to variances, which again are intimately wed to the Gaussian worldview.

It should be obvious by now, but it merits repetition, that Bachelier's assumptions (4), and the linear/Gaussian approach to risk measurement that they implicitly inspired, could be violated for many reasons. Consider in particular the normality in (4b), as we have thus far focused on normality and the associated use of standard deviation for risk measurement. Normality is simply an *assumption*. Of course it can be derived from the continuous-time diffusion (2), but (2) is *itself* an assumption. For example, even if we maintain all other aspects of the diffusion (2), simply relaxing the assumption that it is driven exclusively by standard Brownian motion (as for example if it were driven by a standard Brownian component and a jump component) will invalidate the normality in (4b).

Mandelbrot (1963) and Fama (1965) recognized all this, subjected Bachelier's model to empirical scrutiny, rejected it soundly, and emphasized a new theoretical (stable Paretian) distributional paradigm. They emphasized that real asset returns tend to be symmetrically distributed, but that their distributions have more probability mass in the center and in

(2) into Gaussian discrete-time shocks in (4b), as the shock over any discrete interval is effectively the sum of infinitely many smaller shocks in the interval.

⁵Specifically, excess portfolio returns are projected, or regressed, on excess market returns.

the tails than does the Gaussian. The “fat tails,” in particular, imply that extreme events – crashes and the like, “black swans” in the memorable prose of Taleb (2007) – happen in real markets far more often than would happen the case under normality. In many respects those papers mark the end of “traditional” (linear, Gaussian) risk thinking and the beginning of “modern” (non-linear, non-Gaussian) thinking, even if traditional thinking continues to appear too often in both academics and industry.

Mandelbrot continued, correctly, to emphasize fat tails for the rest of his career, even as much of the rest of the finance profession shifted focus instead to possible dependence in returns, in particular linear serial correlation or “mean reversion” in returns, a violation of the “iid” part of (4b). Reasonable people can – and do – still debate the existence and strength of mean reversion in returns, but no reasonable person can deny the existence of fat tails in distributions of high-frequency returns.

In the relevant non-Gaussian environments emphasized by Mandelbrot and Fama, the standard deviation loses its status as the uniquely-relevant risk measure. It remains of *some* value, as the expected squared deviation of a random variable from its mean always provides *some* risk information, for example via bounds like the classic Chebychev inequality, which remains valid in non-Gaussian environments. But the standard deviation provides an incomplete summary of risk in non-Gaussian environments, and inappropriate use of Gaussian thinking can produce seriously misleading risk assessments. For example, as is well-known, the probability that a Gaussian random variable assumes a value more than two standard deviations from its mean is small (approximately five percent), whereas that same probability for a fat-tailed random variable can be much larger. Indeed events that occur with negligible probability in Gaussian environments (e.g., 3σ events) can (and do) occur frighteningly often in the fat-tailed environments relevant for financial markets.

Hence the Mandelbrot-Fama findings make one skeptical about relying on exclusively on standard deviation for risk assessment, and other approaches have been proposed. For example, extreme-value theory has proved helpful by producing measures of tail fatness in the extreme tails of distributions, and related quantities derived from them, as discussed extensively in Embrechts et al. (1997).

Another useful additional risk measure is so-called “value at risk,” which is the *quantile* of a distribution, typically an extreme quantile related to a tail event. 1% value-at-risk (VaR), for example, is just the first percentile of a payoff or return distribution; that is, 1% VaR is that value such that the probability of a greater loss is just 1%. VaR has achieved notable popularity in applied finance, particularly in industry and regulation, and Duffie and Pan

(1997) provide an insightful introduction.

But even such seemingly-innocuous measures as VaR (say) can wreak havoc when used exclusively. First, and obviously, note that a single VaR (1%, say) is an incomplete risk summary, as complete risk assessment would necessitate monitoring VaR at a variety of levels, which in the limit would trace the entire distribution. That is, one ultimately wants to examine and characterize not just *parts* of distributions, but *entire* distributions.⁶

Second, and less obviously, VaR is also an *incoherent* risk summary, in the sense that it violates certain axioms that reasonable risk measures should arguably satisfy, as shown by Artzner et al. (1999). Unfortunately, however, there is no all-encompassing uniquely-relevant risk measure for non-Gaussian situations. One can and should attempt to assess distributions in a variety of ways, examining a variety of risk measures, each insufficient individually but hopefully aggregating to a well-informed risk assessment.

In closing this section, it is instructive to connect to the Black-Scholes formula (1) for a final take on the perils of uncritical Gaussian thinking. The validity of the Black-Scholes formula requires two sets of heroic assumptions. The first concerns perfect capital markets, with continuous trading possibilities. The second concerns the underlying spot price S : Ironically, Black-Scholes is completely predicated on the maintained assumption that S follows precisely the Bachelier process (4a)-(4b)! As we have emphasized, the Bachelier model fails, due to non-normality of shocks and other, related, reasons that we will discuss shortly. Hence the Black-Scholes formula fails, mis-pricing options in systematic ways across strike prices.⁷

The bottom line is sobering indeed: Although Gaussian assumptions often produce simple, intuitive models, and similarly simple, elegant formulas that serve as useful theoretical benchmarks, their uncritical use in real-world financial risk management is often ill-advised. Real financial asset returns are generally fat-tailed, not Gaussian, so that Gaussian models tend to understate true risk, resulting in mis-priced spot and derivative assets, mis-allocated portfolios, and inadequate risk management strategies.

4 Asset Returns are Conditionally Heteroskedastic

Thus far we have worked in *iid* environments. In that case conditional moments are constant and equal to unconditional moments, and hence there is no need to distinguish them. In particular, we have emphasized the unconditional non-normality first noticed by Mandelbrot.

⁶See Diebold et al. (1998) for background and references on full-density risk assessment and forecasting.

⁷See, for example, Christoffersen (2012), chapter 10.

We could write a generalized Bachelier model as

$$r_t = \sigma \varepsilon_t \quad (5a)$$

$$\varepsilon_t \sim iid(0, 1), \quad (5b)$$

where ε is not necessarily Gaussian.

But time-varying volatility (conditional variance) also turns out to be a key feature of financial asset return data, and risk measurement and management are intrinsically concerned with tracking volatility movements. This leads us to distinguish conditional from unconditional distributions, and in particular, conditional from unconditional variance. With this in mind, we can write a differently-generalized Bachelier model as

$$r_t = \sigma_t \varepsilon_t \quad (6a)$$

$$\varepsilon_t \sim iid N(0, 1). \quad (6b)$$

In particular, we now allow for volatility dynamics but we maintain Gaussian shocks. Different volatility models correspond to different assumptions about the dynamics of the conditional variance, σ_t^2 . The Gaussian shock distribution may seem incongruous with the fat tails emphasized in section 3, but interestingly and importantly, it is not, as will be shown shortly.

4.1 ARCH and GARCH

The autoregressive conditional heteroskedasticity (ARCH) model of Engle (1982) was a key breakthrough. ARCH models are tailor-made for time-series environments, in which one often sees volatility clustering, such that large changes tend to be followed by large changes, and small by small, *of either sign*. That is, one often sees persistence, or serial correlation, in volatility, quite apart from persistence (or lack thereof) in conditional mean dynamics. The ARCH process approximates volatility dynamics in an autoregressive fashion; hence the name *autoregressive* conditional heteroskedasticity. The ARCH(1) model, for example, is:

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2. \quad (7)$$

ARCH processes achieve for conditional variance dynamics precisely what standard autoregressive processes achieve for conditional mean dynamics.

In practice the persistence in return volatility is typically so strong that finite-order (let alone first-order) ARCH processes are inadequate. However, the generalized ARCH ("GARCH") model of Bollerslev (1986) solves the problem. The GARCH(1,1) model is

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2, \quad (8)$$

and it has emerged as a canonical benchmark. It simply augments the ARCH(1) process with a direct conditional variance lag. GARCH(1,1) processes achieve for conditional variance dynamics precisely what ARMA(1,1) processes achieve for conditional mean dynamics. Indeed just as an ARMA(1,1) process is AR(∞), back-substitution in (8) reveals that the GARCH(1,1) is ARCH(∞). In particular, current volatility is an exponentially weighted moving average of the entire history of past squared returns.

We are now in a position to make a crucially important observation: *conditional variance dynamics fatten unconditional distributional tails*. Hence, for example, even if ε in (6) is Gaussian, the unconditional distribution of returns will be fat-tailed when the conditional variance evolves in serially-correlated fashion as with ARCH or GARCH. The unconditional fat tails arise because volatility clustering generates disproportionate activity in the center (tranquil times) and tails (volatile times) of the unconditional distribution. If in addition ε is fat-tailed, then the unconditional distribution of returns will be even *more* fat-tailed.

The insight that conditional variance dynamics fatten unconditional distributional tails has an interesting history. Mandelbrot (1963) offered volatility clustering as an "explanation" of fat unconditional tails but didn't pursue it. Engle (1982) pursued a particular model of volatility clustering and showed that it did indeed imply fat tails, but he emphasized the volatility clustering. In any event, the key point is that unconditional fat tails and conditional volatility clustering are not independent phenomena, requiring independent explanations. Instead the GARCH model makes clear that they are intimately-related, two sides of the same coin.

4.2 Stochastic Volatility

In the GARCH model, conditional variance is a deterministic function of conditioning information. That may seem odd at first, but it's not; indeed conditional moments typically *are* deterministic functions of conditioning information.⁸ It is, however, certainly possible

⁸The linear regression model, for example, is a deterministic model of the conditional mean, given by $E(y/X) = x\beta$.

to allow for a separate stochastic shock to volatility, working with a two-shock rather than one-shock model, which takes us to the so-called stochastic volatility (SV) model. The SV(1) model, for example, is given by:⁹

$$\sigma_t^2 = \exp(\beta_0 + h_t) \quad (9a)$$

$$h_t = \rho h_{t-1} + \eta_t \quad (9b)$$

$$\eta_t \sim iid N(0, \sigma_\eta^2). \quad (9c)$$

The stochastic volatility model (6),(9) was proposed and applied empirically by Taylor (1982). The term “stochastic volatility” is a bit odd, insofar as GARCH conditional variance is *also* a stochastic process, but the usage is ubiquitous and so we follow suit.

Taking absolute values and logs, we can re-write the SV model as:

$$2\ln|r_t| = \beta_0 + h_t + 2\ln|\varepsilon_t| \quad (10a)$$

$$h_t = \rho h_{t-1} + \eta_t. \quad (10b)$$

This is a non-Gaussian state-space system, because the measurement error $2\ln|\varepsilon_t|$ can not be Gaussian if ε_t Gaussian. Hence parameter estimation and volatility extraction require filters more sophisticated than the Kalman filter (e.g., the particle filter); see for example Andrieu et al. (2010). Note that any such extractions will be smoothed versions of the volatility proxy (here based on absolute returns) just as in the GARCH case (there based on squared returns).

Although GARCH and SV are similar in many respects, there are also important differences. The key distinction is that, in the parlance of Cox (1981), GARCH is “observation-driven” whereas SV is “parameter-driven.” Observation-driven GARCH refers to the fact that GARCH volatility is driven by *observed* data (the history of r_t^2), whereas parameter-driven SV refers to the fact that SV volatility is driven by *latent* shocks (the history of η_t).

SV’s parameter-driven structure makes it comparatively easy to analyze, using the large amount of powerful theory available for state-space models. SV models are, however, challenging – or at least comparatively tedious – to estimate, as MCMC methods like the particle filter are required to evaluate the likelihood or explore the posterior. The situation is reversed

⁹Note that the exponential in the SV formulation automatically keeps σ_t^2 positive, whereas, for example, GARCH(1,1) requires the parameter constraints $\omega, \alpha, \beta > 0$.

for GARCH models. GARCH models are trivial to estimate, because they are specified directly in terms of one-step conditional densities, and the likelihood can always be written in terms of those densities, so evaluation of the GARCH likelihood is simple.

SV has an interesting and rather deep economic motivation. It embodies time deformation; that is, the idea that “economic time” and “calendar time” need not move together in financial markets. In particular, although calendar time evolves at a constant rate, economic time may speed up or slow down depending on information flow into the market. The varying speed of economic time, corresponding to varying speed of information flow, produces stochastic volatility in calendar-time returns. Calendar-time volatility is high when economic time is running quickly (i.e., when more information is flowing into the market), and low when economic time is running slowly.

The time deformation perspective traces at least to Clark (1973). Let $\exp(\eta_t)$ be a non-negative random variable governing the number of trades per period of calendar time. Then the continuously-compounded return per period of calendar time (r_t) is the sum of the $\exp(\eta_t)$ intra-period returns. That is, $r_t = \sum_{i=1}^{\exp(\eta_t)} r_i$, where (say) $r_i \sim iid(0, \kappa)$. This implies that $\sigma_t^2 = \kappa \exp(\eta_t)$, so that r_t has conditional heteroskedasticity linked to trading volume. Note that this fits precisely the form of the stochastic volatility model (6),(9) with $\beta_0 = \ln \kappa$ and $\rho = 0$ so that volatility is stochastic but not persistent.

Volatility fluctuates in the Clark model in a fashion linked to trading volume, but not in a persistent fashion, because Clark takes volume to be iid. Hence the Clark model is perhaps best thought of as “endogenizing” unconditional fat tails via time deformation, rather than as producing interesting volatility dynamics. It is a simple extension, however, to allow for serial correlation in volume. We write volume as $\exp(h_t)$ where $h_t = \rho h_{t-1} + \eta_t$, in which case volatility becomes persistent, given by $\sigma_t^2 = \kappa \exp(h_t)$, and of course unconditional tails remain fat. This is precisely the stochastic volatility model (6),(9) with $\beta_0 = \ln \kappa$.

4.3 Realized Volatility

Previously we introduced models of discrete-time (e.g., daily) *conditional* variance, effectively $E(r_t^2 | \Omega_{t-1})$, where the precise contents of Ω_{t-1} depended on the precise model studied (e.g., GARCH, SV). That is, we studied conditional *expectations* of discrete-time squared returns. Now we study not expectations, but *realizations*, from a related but somewhat different perspective. We work in continuous time, with continuously-evolving volatility, and we estimate the total quadratic variation over a discrete interval, in a sense dispensing with volatility *forecasts* in exchange for highly-accurate “*nowcasts*”. This realized volatility

facilitates superior risk measurement, which can translate into superior risk management. Moreover, as we shall show, it can also be used to improve discrete-time conditional variance models.

Consider a stochastic volatility diffusion

$$dp = \sigma(t)dW, \quad (11)$$

where as previously p denotes log price and dW is an increment of standard Brownian motion. The change relative to the Bachelier diffusion (2) is that we now allow the instantaneous volatility to be time-varying, writing $\sigma(t)$. A key object of interest for risk measurement, risk management, asset pricing and asset allocation is the so-called *integrated volatility* (quadratic variation) over a discrete interval like a day:

$$IV_t = \int_{t-1}^t \sigma^2(\tau) d\tau. \quad (12)$$

IV is a function of the latent instantaneous volatility, $\sigma(t)$, which would seem to make precise estimation of IV difficult or impossible.

However, following work such as Andersen et al. (2001), Barndorff-Nielsen and Shephard (2002), and Andersen et al. (2003), which builds on the theoretical insights of Merton (1980), we can use high-frequency intra-day data to obtain a highly-accurate estimator of IV , called realized volatility (RV). RV on day t based on returns at intra-day frequency Δ is

$$RV_t(\Delta) \equiv \sum_{j=1}^{N(\Delta)} (p_{t-1+j\Delta} - p_{t-1+(j-1)\Delta})^2, \quad (13)$$

where $p_{t-1+j\Delta} \equiv p(t-1+j\Delta)$ denotes the intra-day log-price at the end of the j th interval on day t , and $N(\Delta) \equiv 1/\Delta$. In principle, RV can be made arbitrarily accurate for IV by sampling intra-day returns finely enough. That is, as $\Delta \rightarrow 0$, corresponding to progressively more finely-sampled returns, $RV_t \rightarrow IV_t$. Hence, the *true* ex-post daily (say) volatility effectively becomes observable, and it does so in an entirely model-free fashion regardless of the underlying process for $\sigma(t)$.

Note that the squared daily return is just a particularly bad version of RV_{t-1} with sampling only once per day. To see why sampling only once per day is inadequate, consider a day with large price fluctuations on which, by chance, opening and closing prices are equal. Then the daily return (and hence squared return) is zero, and one would erroneously assess

low volatility for that day based on the squared daily return. In contrast, RV based on high-frequency intra-day sampling would capture the high volatility.

The RV - IV framework can be generalized to allow for jumps. That is, following Barndorff-Nielsen and Shephard (2004), the stochastic volatility diffusion (11) can be generalized to include compound Poisson jumps,

$$dp = \sigma(t)dW + \delta(t)dq(t), \quad (14)$$

where $q(t)$ is a counting process with (possibly) time-varying intensity governing jump arrivals, where $\delta(t) = p(t) - p(t-)$ refers to the size of the discrete jumps. In the compound Poisson case, RV remains consistent for IV , but IV now has a jump and a non-jump component. Barndorff-Nielsen and Shephard (2004) develop a modified version of RV (so-called bi-power variation) that is consistent for the non-jump part of IV , thereby allowing jump and non-jump parts of IV to be disentangled.¹⁰

Because of its superior accuracy relative to IV proxies like daily absolute or squared returns, RV can be used to improve the performance of SV and $GARCH$ models. It can be used in SV to replace the noisier volatility proxy $|r_t|$ in (10), as emphasized by Barndorff-Nielsen and Shephard (2002). It can be used in $GARCH$ to replace the noisier volatility proxy r_t^2 in (8), producing a “GARCH- RV ” model,

$$\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \gamma RV_{t-1}. \quad (15)$$

Hansen et al. (2012) augment the $GARCH$ - RV model with a model for forecasting RV , thereby producing a “closed system” that can be used to forecast volatility at any horizon.

4.4 Extensions and Applications

Multivariate extensions are straightforward, yet tremendously important for financial risk measurement, financial risk management, asset pricing, and portfolio allocation. The key issue is how to enforce parsimony while maintaining fidelity to the data. Fortunately it turns out to be possible. In this section we sketch the rudiments of some popular approaches, and we discuss application to portfolio allocation.

¹⁰See also Andersen et al. (2007).

4.4.1 Multivariate

The multivariate generalization of the Bachelier model with time-varying volatility (6) is

$$R_t = \Omega_t^{1/2} Z_t \quad (16)$$

$$Z_t \sim iidN(0, I),$$

where the $N \times 1$ vector R contains returns, the $N \times N$ matrix $\Omega_t^{1/2}$ is one of the “square-roots” (e.g., the Cholesky factor) of the covariance matrix Ω_t , and the $N \times 1$ vector Z contains shocks.

Without additional restrictions on Ω_t , (16) rapidly becomes unwieldy as N grows. Hence much literature has focused on restrictions on Ω_t that simultaneously achieve parsimony yet maintain fidelity to the data. For example, Diebold and Nerlove (1989) invoke factor structure with orthogonal idiosyncratic factors, so that Ω_t may be written as

$$\Omega_t = B' \Omega_{F,t} B + \Sigma, \quad (17)$$

where B is a matrix of factor loadings, $\Omega_{F,t}$ is a low-dimensional time-varying common-factor covariance matrix, and Σ is a constant diagonal idiosyncratic-factor covariance matrix. Much subsequent literature has followed suit.

An alternative approach proceeds by noting that a conditional covariance matrix may always be decomposed into a conditional correlation matrix pre- and post-multiplied by a diagonal matrix of conditional standard deviations,

$$\Omega_t = D_t \Gamma_t D_t. \quad (18)$$

Motivated by this decomposition, Bollerslev (1990) first proposed treating the conditional correlations as constant, $\Gamma_t = \Gamma$, so that the dynamics in Ω_t are driven solely by the univariate volatility dynamics D_t . The dynamic conditional correlation (DCC) model of Engle (2002) generalizes Bollerslev's approach to allow for time-varying conditional correlations, while still maintaining a simple dynamic structure motivated by GARCH(1,1).

Multivariate modeling of realized volatility (i.e., large realized covariance matrices) presents special challenges, as for example there is no guarantee that separately-estimated realized variances and covariances, when assembled, will produce into a positive-semidefinite realized covariance matrix. Regularization methods have proved useful there, as in Hautsch et al.

(2012) and Lunde et al. (2011).

4.4.2 Volatility Overlays in Dynamic Portfolio Allocation

Thus far our discussion of time-varying volatility has implicitly emphasized risk measurement, asset pricing and risk management. First, the implicit emphasis on risk measurement is obvious, insofar as model-based tracking of time-varying volatility clearly requires models that *admit* time-varying volatility. Second, the implicit emphasis on asset pricing (in particular, options pricing) follows from our recognition that, although Black-Scholes requires a constant-volatility environment, real financial markets feature time-varying volatility, so that Black-Scholes must fail. And indeed it does fail. Appropriate pricing with time-varying volatility depends on its precise form, and tidy closed-form expressions are typically unavailable. Instead, Monte Carlo methods must be used, as discussed in Glasserman (2003) and Christoffersen (2012). Finally, the implicit emphasis on risk management follows from the above-discussed implicit emphasis on options pricing, as options are a key tool of risk management.

Time-varying volatility can also play crucial role in portfolio management. As is well-known, optimal portfolio shares depend on variances and covariances of the assets in the portfolio. In particular, if an investor wants to minimize portfolio return volatility subject to achieving target return μ_p , she must solve $\min_w w' \Sigma w$ subject to $w' \mu = \mu_p$, where w is the vector of portfolio shares, μ is the conditional mean vector of returns, and Σ is the conditional covariance matrix of returns. If Σ is time-varying, then so too are optimal portfolio shares, which then solve $\min_{w_t} w_t' \Sigma_t w_t$ subject to $w_t' \mu = \mu_p$. This leads to the idea of *volatility timing*, or *volatility overlays*, in which portfolio shares are adjusted dynamically to reflect movements (and anticipated movements) in volatility. A key question is how much risk-adjusted excess return is gained from volatility timing. Fleming et al. (2001) find large gains from volatility timing using GARCH-type conditional covariance matrices, and Fleming et al. (2003) find even larger gains when using projections of realized covariance matrices.

5 Bond Market Risks are Special

Here we focus on bond markets rather than equity markets. The focus is still market risk as opposed to credit risk, however, because we consider “riskless” (government) bonds.¹¹ Bond markets are special because bond yield modeling is intrinsically multivariate, and with

¹¹Later in section 6.2 we confront the credit risk associated with defaultable bonds.

a special multivariate structure. That is, in bond markets one is concerned not just with individual yields, but with many yields simultaneously – the entire yield curve. The yields that make up the yield curve, moreover, are not evolving independently; rather, they are linked through their dependence on underlying common factors. For a broad overview, see Diebold and Rudebusch (2013).

5.1 Empirically-Tractable Modeling

At any time, one sees a large set of yields and may want to fit a smooth curve. Nelson and Siegel (1987) fit the curve

$$y(\tau) = \beta_1 + \beta_2 \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} \right) + \beta_3 \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau} \right). \quad (19)$$

Note well that the Nelson-Siegel model as presently introduced is a *static* model, fit to the *cross section* of yields. At first pass, moreover, the Nelson-Siegel functional form seems rather arbitrary – a less-than-obvious choice for approximating an arbitrary yield curve. But Nelson-Siegel turns out to have some very appealing features.

First, Nelson-Siegel desirably enforces some basic constraints from financial economic theory; for example, the implied discount curve is constrained to approach zero as maturity grows. Second, it provides a parsimonious approximation, which is desirable because it promotes smoothness (yields tend to be very smooth functions of maturity), it guards against in-sample overfitting (which is important for producing good forecasts), and it promotes empirically tractable and trustworthy estimation (which is always desirable).

Third, despite its parsimony, the Nelson-Siegel form also provides a flexible approximation. Flexibility is desirable because the yield curve assumes a variety of shapes at different times. Inspection reveals that, depending on the values of the four parameters $(\beta_1, \beta_2, \beta_3, \lambda)$, the Nelson-Siegel curve can be flat, increasing, or decreasing linearly, increasing or decreasing at increasing or decreasing rates, U-shaped, or \cap -shaped.

Fourth, from a mathematical approximation-theoretic viewpoint, the Nelson-Siegel form is far from arbitrary. The forward rate curve corresponding to the Nelson-Siegel yield curve is a constant plus a Laguerre function. Laguerre functions are polynomials multiplied by exponential decay terms and are well-known mathematical approximating functions for non-negative variables.

Now consider dynamic aspects of the yield curve. Bond yields tend to move noticeably together, with factor structure. This classic recognition traces to Litterman and Scheinkman

(1991), and it is echoed repeatedly in the subsequent literature. Typically three factors, which turn out to be interpretable as level, slope and curvature, are all that one needs to explain most variation in sets of government bond yields.

The dynamic factor structure in bond yields is captured by the so-called “dynamic Nelson-Diegel” (DNS) model of Diebold and Li (2006). The DNS model is

$$y_t(\tau) = \beta_{1t} + \beta_{2t} \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} \right) + \beta_{3t} \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau} \right), \quad (20)$$

where τ is maturity, t is time, and β_{1t} , β_{2t} , and β_{3t} are latent level, slope and curvature factors. Operationally, the DNS model (20) is nothing more than the Nelson-Siegel model (19) with time-varying parameters. The interpretation, however, is deep: DNS distills the yield curve into three latent factors (β_{1t} , β_{2t} , and β_{3t}), the dynamics of which determine entirely the dynamics of y_t for any τ , and the coefficients (“factor loadings”) on which determine entirely the cross section of $y(\tau)$ for any t . Hence DNS is a dynamic factor model, in which a high-dimensional set of variables (in this case, the many yields across maturities) is actually driven by much lower-dimensional state dynamics (in this case the three latent yield factors).

DNS blends several appealing ingredients, as discussed above. And its dynamic factor structure is very convenient statistically, as it distills seemingly-intractable high-dimensional yield dynamics into easily-handled low-dimensional state dynamics.¹² Hence DNS has been popular among financial market participants, central banks, and empirically-inclined academic economists.

5.2 Arbitrage-Free Modeling

Vasicek (1977) is the classic early-vintage arbitrage-free model, with the instantaneous short rate determined by a stochastic differential equation,

$$dr_t = k^Q(\theta^Q - r_t)dt + \sigma dW_t^Q,$$

where Q denotes the risk-neutral measure. The longer-maturity yields then feature risk premia constrained in just the right way such that arbitrage possibilities vanish. As we have seen, however, multi-factor models are necessary for capturing yield curve dynamics. Hence

¹²Of course, if one simply *assumed* factor structure but the data did not satisfy it, one would simply have a misspecified model. Fortunately, however, and again as emphasized since Litterman and Scheinkman (1991), bond yields *do* display factor structure.

if Vasicek's ingenious construction launched a now-massive literature on arbitrage-free yield curve modeling, it nevertheless proved too restrictive, as it is effectively a single-factor model.

In a landmark paper, Duffie and Kan (1996) relax the Vasicek single-factor constraint. They work in an affine diffusion environment with a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), Q)$. The state vector, X_t , is defined on a set $M \subset \mathbf{R}^n$ and governed by the stochastic differential equation

$$dX_t = K^Q(t) (\theta^Q(t) - X_t) dt + \Sigma(t) D(X_t, t) dW_t^Q, \quad (21)$$

where W^Q is a standard Brownian motion in \mathbf{R}^n , the information about which is contained in the filtration (\mathcal{F}_t) . The drifts and dynamics $\theta^Q : [0, T] \rightarrow \mathbf{R}^n$ and $K^Q : [0, T] \rightarrow \mathbf{R}^{n \times n}$ are bounded, continuous functions. Similarly, the volatility matrix $\Sigma : [0, T] \rightarrow \mathbf{R}^{n \times n}$ is a bounded, continuous function. The matrix $D : M \times [0, T] \rightarrow \mathbf{R}^{n \times n}$ is diagonal,

$$D(x_t, t) = \begin{pmatrix} \sqrt{\gamma^1(t) + \delta^1(t)X_t} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sqrt{\gamma^n(t) + \delta^n(t)X_t} \end{pmatrix},$$

where $\delta^i(t)$ denotes the i th row of the matrix

$$\delta(t) = \begin{pmatrix} \delta_1^1(t) & \dots & \delta_n^1(t) \\ \vdots & \ddots & \vdots \\ \delta_1^n(t) & \dots & \delta_n^n(t) \end{pmatrix},$$

and $\gamma : [0, T] \rightarrow \mathbf{R}^n$ and $\delta : [0, T] \rightarrow \mathbf{R}^{n \times n}$ are bounded, continuous functions.

The preceding precise mathematical statement of Duffie-Kan state dynamics is a bit tedious, but its intuition and importance are immediate. The admissible Duffie-Kan state dimension and dynamics are very rich, and *much* richer than in the Vasicek model. Duffie and Kan then proceed to specify the instantaneous risk-free rate as an affine function of the state variables X_t ,

$$r_t = \rho_0(t) + \rho_1(t)' X_t, \quad (22)$$

where $\rho_0 : [0, T] \rightarrow \mathbf{R}$ and $\rho_1 : [0, T] \rightarrow \mathbf{R}^n$ are bounded, continuous functions. One can then show that zero-coupon bond prices are exponential-affine functions of the state variables,

$$\begin{aligned} P(t, T) &= E_t^Q \left(\exp \left(- \int_t^T r_u du \right) \right) \\ &= \exp (B(t, T)' X_t + C(t, T)), \end{aligned} \quad (23)$$

where $B(t, T)$ and $C(t, T)$ are the solutions to the system of ordinary differential equations

$$\begin{aligned}\frac{dB(t, T)}{dt} &= \rho_1 + (K^Q)'B(t, T) - \frac{1}{2} \sum_{j=1}^n (\Sigma' B(t, T) B(t, T)' \Sigma)_{jj} (\delta^j)' \\ dC(t, T)dt &= \rho_0 - B(t, T)' \theta^Q - \frac{1}{2} \sum_{j=1}^n (\Sigma' B(t, T) B(t, T)' \Sigma)_{jj} \gamma^j,\end{aligned}$$

with $B(T, T) = 0$ and $C(T, T) = 0$. In addition, the exponential-affine zero-coupon bond pricing functions imply that zero-coupon yields are

$$y(t, T) = -\frac{1}{T-t} \log P(t, T) = -\frac{B(t, T)'}{T-t} X_t - \frac{C(t, T)}{T-t}. \quad (24)$$

In particular, yields at all maturities are affine functions of the state variables, just as with the instantaneous risk-free rate.

Duffie-Kan blends several appealing ingredients. Its rich state dynamics make it very general, its construction precludes arbitrage, and its affine structure makes for elegant simplicity. Hence it has been extremely popular with theoretically-inclined academics.

5.3 Simultaneously Empirically-Tractable and Arbitrage-Free Modeling

Christensen et al. (2011) produce a model that is both empirically-tractable and arbitrage-free, by showing that DNS factor loadings emerge in the Duffie-Kan affine arbitrage-free framework for a particular specification of state dynamics and instantaneous short rate. They call the resulting arbitrage-free version of DNS “arbitrage-free Nelson-Siegel” (AFNS).

AFNS proceeds as follows. Consider a Duffie-Kan model with three state variables, with instantaneous short rate given by

$$r_t = X_t^1 + X_t^2, \quad (25)$$

and with risk-neutral (Q) dynamics given by

$$\begin{pmatrix} dX_t^1 \\ dX_t^2 \\ dX_t^3 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \lambda & -\lambda \\ 0 & 0 & \lambda \end{pmatrix} \left[\begin{pmatrix} \theta_1^Q \\ \theta_2^Q \\ \theta_3^Q \end{pmatrix} - \begin{pmatrix} X_t^1 \\ X_t^2 \\ X_t^3 \end{pmatrix} \right] dt + \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_{22} & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} \end{pmatrix} \begin{pmatrix} dW_t^{1,Q} \\ dW_t^{2,Q} \\ dW_t^{3,Q} \end{pmatrix}, \quad (26)$$

where $\lambda > 0$. Then zero-coupon bond prices are

$$P(t, T) = E_t^Q \left(\exp \left(- \int_t^T r_u du \right) \right) \quad (27)$$

$$= \exp \left(B^1(t, T)X_t^1 + B^2(t, T)X_t^2 + B^3(t, T)X_t^3 + C(t, T) \right),$$

where $B^1(t, T)$, $B^2(t, T)$, $B^3(t, T)$ and $C(t, T)$ are governed by a system of ordinary differential equations derived and solved by Christensen et al. The implied zero-coupon bond yields are

$$y(t, T) = X_t^1 + \frac{1 - e^{-\lambda(T-t)}}{\lambda(T-t)} X_t^2 + \left[\frac{1 - e^{-\lambda(T-t)}}{\lambda(T-t)} - e^{-\lambda(T-t)} \right] X_t^3 - \frac{C(t, T)}{T-t}. \quad (28)$$

This is of precisely DNS form, with a simple “yield-adjustment term” appended.

A key insight is that, although AFNS clearly imposes significant restrictions on dynamics under the Q -measure, it places *no* restrictions on dynamics under the physical measure (P -measure). Indeed the only P -measure difference between DNS and AFNS is the AFNS yield-adjustment term $-\frac{C(t, T)}{T-t}$, which depends only on maturity, not on time. The beauty of AFNS is that it maintains the empirical tractability of DNS while simultaneously imposing the theoretical appeal of absence of arbitrage.

6 Rare Event Risks are Special Too

Thus far we have focused primarily on what might be called “normal-time market risk” in equity and bond markets. By definition, most times are normal times, so a large amount of professional attention has focused on normal-time market risk. There are, however, many other key financial risks to measure and manage. They are typically associated not with normal times, but rather with abnormal times – rare events ranging, for example, from extreme market movements, to information system failures, to bond defaults, to global financial collapse. Because there are (again by definition) comparatively few historical occurrences of such rare events, they are harder to model statistically. Yet they are also disproportionately important, because of the havoc wreaked when they *do* occur, so they must not be neglected. Recent years have witnessed progress in varying degrees. Here we focus on extreme market risk events, credit risk events, operational risk events, and systemic risk events.

6.1 Extreme Market Movements

Standard multivariate volatility models, such as those introduced earlier in section 4.4.1, work well for describing relationships (e.g., correlations) near the center of the joint return distribution; that is, for describing relationships during normal times. But relationships may differ during crises, due to herd behavior as everyone runs for the exit simultaneously. In particular, correlations may increase during crises, so that portfolio diversification is lost when it is most needed.

Multivariate extreme value theory offers hope for describing cross-asset relationships “in the tails,” during crises, whereas such tail relationships are often not well-captured by standard parametric distributions or correlation measures. Longin and Solnik (2001), for example, define and compute extreme correlations between monthly U.S. equity index returns and a number of foreign country indexes. The correlation between extreme negative returns tends to be much larger the correlation between large positive returns. Such strong correlation between negative extremes is a key risk management concern.

6.2 Credit Risk, Operational Risk, and Risk Aggregation

One important non-market risk is credit risk, as treated in the Merton (1974) model. Yields on defaultable bonds have a risk-free yield component and a premium for default risk. That is, credit spreads are driven by bond default probabilities. Merton’s key insight is that default probabilities, and hence risk premia, depend on net asset values (asset value less the value of outstanding debt): if a firm’s net asset value gets too low, it will find it optimal to default on its debt. Empirically, net asset value is unobserved, but observed equity value is a useful proxy. Low equity valuations become more likely when equity volatility is high, so increases in equity volatility should widen credit spreads. Empirical studies do indeed find strong links between credit spreads and equity volatility, as for example in Campbell and Tackler (2003).

Another important non-market risk is operational risk. It is distinct from credit risk, but it shares the key feature of rare event arrivals. Operational risk is in certain respects a catch-all for various non-market, non-credit risks, including but not limited to losses from rogue traders, system failures, property damage, etc. It is not clear that such rare risk events are amenable to the same sorts of statistical treatment as, say, those that drive market risk. Work by de Fontnouvelle et al. (2006), however, provides some cause for optimism. A key issue is that short samples of historical data for any single firm are subject to the

“peso problem”; that is, important operational risks may have been present but not realized historically, just by luck.¹³ deFontnouvelle et al. confront such data biases and propose statistical methods for modeling the frequency and severity of operational losses. Among other things, they argue that there may be scope for pooling operational loss data across firms, using a log-exponential distribution.

Confronting different risks – market, credit, operational – forces one also to confront their aggregation, a difficult challenge as they are surely not independent. This leads to the challenging problem of so-called “enterprise risk measurement and management.” Rosenberg and Schuermann (2006) tackle it by using methods like those surveyed already in this chapter to arrive at marginal distributions for various risk exposures, and then invoking various copulas to arrive at a full joint distribution.

6.3 Systemic Risk

There is no single definition of systemic risk, and I will shortly introduce several, but the defining characteristic is that systemic risk – one way or another – involves market-wide movements. Professional attention has turned to systemic risk only recently, following the financial crisis and recession of 2007-2009. Most related research is unpublished, and much is ongoing; hence it would be premature to attempt to select particular papers for this volume. In light of this situation, I now discuss systemic risk and its measurement at some length.

6.3.1 Marginal Expected Shortfall and Expected Capital Shortfall

Marginal expected shortfall (*MES*) for firm j is

$$MES_{T+1|T}^{j|pkt} = E_T[r_{j,T+1}|\mathbb{C}(r_{pkt,T+1})], \quad (29)$$

where $r_{pkt,T+1}$ denotes the overall market return, and $\mathbb{C}(r_{pkt,T+1})$ denotes a systemic event, such as the market return falling below some threshold C . $MES_{T+1|T}^{j|pkt}$ tracks the sensitivity of firm j ’s return to a market-wide extreme event, thereby providing a simple market-based measure of firm j ’s fragility.

Ultimately, however, we are interested in assessing the likelihood of firm distress, and the fact that a firm’s expected return is sensitive to market-wide extreme events – that is, the fact that its *MES* is large – does not necessarily mean that market-wide extreme events are

¹³In addition, small losses may be under-reported for strategic reasons. Firms would, of course, like to avoid reporting large losses as well, but small losses are easier to conceal.

likely to place it in financial distress. Instead, the distress likelihood should depend not only on MES , but also on how much capital the firm has on hand to buffer the effects of adverse market moves.

These distress considerations raise the idea of expected capital shortfall (ECS), which is closely related to, but distinct from, MES . ECS is the expected *additional capital* needed by firm j in case of a systemic market event. Clearly ECS should be related to MES , and Acharya et al. (2010) indeed show that in a simple model the two are linearly related,

$$ECS_{T+1|T}^{j|mkt} = a_{0j} + a_{1j}MES_{T+1|T}^{j|mkt}, \quad (30)$$

where a_{0j} depends on firm j 's "prudential ratio" of asset value to equity as well as its debt composition, and a_{1j} depends on firm j 's prudential ratio and initial capital.

Building on the theory of Acharya et al. (2010), Brownlees and Engle (2011) propose and empirically implement $ECS_{T+1|T}^{j|mkt}$ as a measure of firm j 's systemic risk exposure to the market at time T , with overall systemic risk then given by $\sum_{j=1}^N ECS_{T+1|T}^{j|mkt}$. Implementation of MES (and hence ECS) requires specification of the systemic market event $\mathbb{C}(r_{mkt,T+1})$, or more simply a market return threshold C . Values of $C = 2\%$ and $C = 40\%$ have, for example, been suggested for one-day and six-month returns, respectively.

6.3.2 CoVaR and Δ CoVaR

In the previous section we introduced MES and ECS , which measure firm systemic risk *exposure* by conditioning firm events on market events. Here we introduce $CoVaR$, which works in the opposite direction, measuring firm systemic risk *contribution* by conditioning market events on firm events.

We have already introduced the intuitive concept of value at risk. Now let us flesh it out with some mathematical precision. Firm j 's 1-step-ahead conditional VaR at level p is the value of $VaR_{T+1|T}^{p,j}$ that solves

$$\Pr_T \left(r_{j,T+1} < -VaR_{T+1|T}^{p,j} \right) = p. \quad (31)$$

Similarly, following Adrian and Brunnermeier (2011), one may define firm j 's 1-step-ahead " $CoVaR$ " at level p conditional on a particular outcome for firm i , say $\mathbb{C}(r_{i,T+1})$, as the

value of $CoVaR_{T+1|T}^{ji}$ that solves

$$\Pr_T \left(r_{j,T+1} < -CoVaR_{T+1|T}^{ji} \mid \mathbb{C}(r_{i,T+1}) \right) = p. \quad (32)$$

Because $\mathbb{C}(r_{i,T+1})$ is not in the time- T information set, $CoVaR$ will be different from the regular time- T conditional VaR . The leading choice of conditioning outcome, $\mathbb{C}(r_{i,T+1})$, is that firm i exceeds its VaR , or more precisely that $r_{i,T+1} < -VaR_{T+1|T}^{pi}$. As such, $CoVaR$ is well-suited to measure tail-event linkages between financial institutions.

A closely-related measure, $\Delta CoVaR_{T+1|T}^{ji}$ (read “Delta $CoVaR$ ”), is of particular interest. It measures the difference between firm- j VaR when firm- i is “heavily” stressed and firm- j VaR when firm i experiences “normal” times. More precisely,

$$\Delta CoVaR_{T+1|T}^{ji} = CoVaR_{T+1|T}^{j|VaR(i)} - CoVaR_{T+1|T}^{j|Med(i)}, \quad (33)$$

where $CoVaR_{T+1|T}^{j|VaR(i)}$ denotes firm- j VaR when firm i 's return breaches its VaR , and $CoVaR_{T+1|T}^{j|Med(i)}$ denotes firm- j VaR when firm i 's return equals its median.

A direct extension lets us progress to the more interesting case of firm i 's overall systemic risk contribution, as opposed to just firm i 's contribution to firm j . We simply set $j = mkt$, so that $\Delta CoVaR_{T+1|T}^{mkt|i}$ then measures the difference between market VaR conditional on firm i experiencing an extreme return, and market VaR conditional on firm i experiencing a normal return. Hence $\Delta CoVaR_{T+1|T}^{mkt|i}$ measures the contribution of firm i to overall systemic risk, $\sum_{i=1}^N \Delta CoVaR_{T+1|T}^{mkt|i}$.

6.3.3 Network Connectedness

Connectedness would appear central to modern risk measurement and management, and indeed it is. It features prominently in key aspects of market risk (return connectedness and portfolio concentration), credit risk (default connectedness), counter-party and gridlock risk (bilateral and multilateral contractual connectedness), and not least, systemic risk (system-wide connectedness). Connectedness would also appear central to understanding underlying fundamental macroeconomic risks, in particular business cycle risk (intra- and inter-country real activity connectedness).

The *MES* and *CoVaR* approaches discussed above address certain aspects of connectedness, as they track association between individual-firm and overall-market movements. Although they and various other systemic risk measures are certainly of interest, they measure different things, raising the issue of whether one could construct a unified framework

for conceptualizing and measuring systemic risk.

Interestingly, modern network theory provides just such a framework, as developed in Diebold and Yilmaz (2011).¹⁴ The simplest network is composed of N nodes, where any given pair of nodes may or may not be linked. We represent the network algebraically by an $N \times N$ symmetric adjacency matrix A of zeros and ones, $A = [a_{ij}]$, where $a_{ij} = 1$ if nodes i and j are linked, and $a_{ij} = 0$ otherwise. Because all network properties are embedded in A , any sensible connectedness measure must be based on A . The most important and popular, by far, are based on the idea of a node's *degree*, given by the number of its links to other nodes $\delta_i = \sum_j a_{ij}$, as well as aspects of the degree distribution across nodes. The total degree $\sum_i \delta_i$ (or mean degree $\frac{1}{N} \sum_i \delta_i$) is the key network connectedness measure.

The network structure sketched above is, however, rather too simple to describe the network connections of relevance in financial risk management. Generalization in two key directions is necessary. First, links may be of varying strength, not just 0-1. Second, links may be of different strength in different directions (e.g., firm i may impact firm j more than firm j impacts firm i). Note, for example, that the systemic risk measures introduced above are weighted and directional. For example, $CoVaR_{T+1|T}^{ji}$ tracks effects from i to j , whereas $CoVaR_{T+1|T}^{ij}$ tracks effects from j to i , and in general $CoVaR_{T+1|T}^{ji} \neq CoVaR_{T+1|T}^{ij}$.

It is a simple matter, however, to characterize directed, weighted networks in a parallel fashion. To allow for directionality, we allow the adjacency matrix A to be non-symmetric, and to allow for different relationship strengths we allow A to contain weights $a_{ij} \in [0, 1]$ rather than simply 0-1 entries. Node degrees are now obtained by summing weights in $[0, 1]$ rather than simply zeros and ones. In addition, and importantly, there are now "to-degrees" and "from-degrees," corresponding to row sums and column sums, which generally differ since A is generally non-symmetric. The from-degree of node i is $\delta_i^{from} = \sum_j a_{ij}$, and the to-degree of node j is $\delta_j^{to} = \sum_i a_{ij}$. The total degree is $\delta = \sum_i \delta_i^{from} = \sum_j \delta_j^{to}$.

Crucially, the from- and to-degrees (and of course the total degree) measure aspects of systemic risk. The from- and to-degrees measure systemic risk with respect to particular firms. From-degrees measure exposures of individual firms to systemic shocks *from* the network, in a fashion analogous to $MES_{T+1|T}^{j|mkt}$. To-degrees measure contributions of individual firms *to* systemic network events, in a fashion analogous to $\Delta CoVaR_{T+1|T}^{mkt|i}$. The total degree aggregates firm-specific systemic risk across firms, providing a measure of total system-wide systemic risk.

¹⁴See also Diebold and Yilmaz (2009) and Diebold and Yilmaz (2012).

7 The Business Cycle is a Key Risk Fundamental

The risk models that we have discussed thus far are inherently “reduced form.” They “explain” risk statistically, largely in an autoregressive fashion, as exemplified by the canonical GARCH family. Fortunately, even if the models fail to provide a deep structural understanding of financial risks, they are nevertheless powerful and useful in a variety of contexts. We have obviously emphasized risk measurement and management, with applications to portfolio allocation, spot and derivative asset pricing. Ultimately, however, we aspire to a deeper structural understanding. That is, we aspire to understand the connections between returns and macroeconomic fundamentals.

Asset prices are risk-adjusted discounted claims on fundamental streams, so prices and their properties should ultimately depend on expected fundamentals and associated fundamental risks. The business cycle is probably the most important macroeconomic fundamental of all. Hamilton and Lin (1996) provide a fine statement of the link between the business cycle and equity return volatility. Building on important early work by Schwert (1989), they use regime-switching models of real growth and equity returns, allowing for both high and low real growth states and high and low equity return volatility states, to show that return volatility is significantly higher in recessions. Hence high volatility during bad times is not just a one-off Great Depression phenomenon, but rather a regularly-recurring business cycle phenomenon.

Indeed business cycle effects run throughout all financial market risks, as emphasized in Andersen et al. (2013). For example, the business cycle is related to the equity premium (e.g., Campbell and Diebold (2009)), and to the slope factor of the default-free yield curve (e.g., Diebold et al. (2006)). The business cycle is also related to the risky yield curve, via two channels. First, it is related to default-free yields, as just mentioned. Second, it is related to default risk. That is, for a variety of reasons, debt defaults are more likely in recessions. One route is through the channels emphasized in the earlier-introduced Merton model. The Merton model suggests that credit risk is linked to equity volatility, but as we discussed above, equity volatility is linked to the business cycle. Hence the Merton model suggests that the business cycle is linked not only to equity market risk, but also to the credit risk of defaultable bonds.

The business cycle is also related to systemic risk. We have already discussed *MES* (and Diebold-Yilmaz “connectedness from”) and *CoVaR* (and Diebold-Yilmaz “connectedness to”), and their aggregate across all firms, which is identical whether “from” or “to”. Such aggregate systemic risk is very much related to the business cycle, because in recessions,

particularly in severe recessions, the financial stability of all firms is weakened.

This raises the idea of “stress testing,” and the centrality of business cycle scenarios in stress tests. That is, it is sometimes informative also to consider risk measures that condition not on historical returns, but rather on assumed *scenarios* for particular risk factors. We might, for example, be interested in the firm-specific effects of a market-wide “business cycle” shock. This could be a “country-wide” business cycle shock, or a “global” business cycle shock as in Aruoba et al. (2011).

Note that the conditioning involved in *ECS* is closely related to the idea of stress tests. Indeed the thought experiment embodied in *ECS* is a stress test, with a very important and natural stress (a large market event) being tested. The same is true of *CoVar*, where the stress is failure of a particular firm. Both scenarios are potentially linked to the business cycle. Of course this raises the issue of “which stresses to test” (key among which are various macroeconomic shocks) and how to assign probabilities to them. In a prescient paper, Berkowitz (1999) grapples with that issue, and Rebonato (2010) provides guidance from the vantage point of Bayesian networks.

Ultimately of course the business cycle does not uniquely “cause” financial crises, and financial crises do not uniquely cause macroeconomic crises. Instead, the relationship is complicated, surely with bi-directional causality and feedback. Monetary policy, moreover, may play a key role. Allen and Gale (2000), for example, highlight the role of money and credit in the determination of asset price bubbles, with systemic implications.

8 Beware the Limits of Statistics

In this chapter I have surveyed a variety of issues in risk measurement, all of which feed into risk management, asset pricing and asset allocation. As shown, the literature ranges from the highly developed and sophisticated (e.g., equity market risk), to the much less well-developed (e.g., various event risks). And the least-well-understood risks are also potentially the most important, as they are the ones that can bring down firms, and entire financial systems, as opposed to pushing earnings up or down by a few percentage points. At the same time, great progress has been made, and there is hope.

This chapter and this book focus on that progress. Simultaneously, however, I hasten to emphasize the limits to statistical modeling. In that vein, Diebold et al. (2010) adopt a taxonomy introduced by Ralph Gomory (1995), in which he classifies knowledge into the known, the unknown and the unknowable, for which they adopt the acronym *KuU*.

One way to think of KuU is as risk (K) (possible outcomes known, probabilities known), uncertainty (u) (outcomes known, probabilities unknown), and ignorance (U) (both outcomes and probabilities unknown).

Perhaps the broadest lesson is recognition of the wide applicability of KuU thinking, and the importance of each of K and u and U . KuU thinking spans all types of financial risk, with the proportion of uU increasing steadily as one moves through the spectrum of market, credit, operational and systemic risks. In addition, KuU thinking spans risk measurement and management in all segments of the financial services industry, including investing, asset management, banking, insurance, and real estate. Finally, KuU thinking spans both the regulated and the regulators: regulators' concerns largely match those of the regulated (risk measurement and management), but with an extra layer of concern for systemic risk.

The existing risk management literature focuses largely on K . In part that's reasonable enough: K risks are plentiful and important. But in part that's also silly, like the old joke of "looking for one's lost car keys only under the lamp post, because that's where the light is." That is, a large fraction of real-world risk management challenges clearly fall not in the domain of K , but rather in the domain of u and U . Indeed a cynic might assert that, by focusing on K , the existing literature has made us expert at the least-important aspects of financial risk management. I believe that such assertions are outlandish; K situations are clearly of routine relevance and great importance. But I also believe that u and U are of equal relevance and importance, particularly insofar as many of the killer risks, which can bring firms down, lurk there.

One way to proceed is to attempt to expand the domain of K via new research, a noble undertaking, but one subject to clear limits. Hence one can not escape confronting u and U . The important issues in the world of u and U are more economic (strategic) than statistical, and crucially linked to *incentives*: How to write contracts (design organizations, formulate fiscal and monetary policies, draft regulations, make investments,...) in ways that create incentives for best-practice proactive and reactive risk management of all types of risks, including (and especially) u and U risks. I look forward to the evolution of financial risk management toward confronting K and u and U equally.

9 Additional Reading

Many more extensive or more specialized treatments exist, and I recommended them for further study. I list only books, but the books of course contain references to thousands of

journal articles.

Cootner (1964)'s classic collection of classic papers is as useful now as when it was published long ago. Recent survey papers include Andersen et al. (2006) and Andersen et al. (2013). Christoffersen (2012) provides a broad and insightful textbook treatment. More specialized book-length treatments include Arrow and Hahn (1971) (general equilibrium with risk); Rebonato (2004) (hedging and derivatives pricing); Markowitz (1959) (mean-variance portfolio analysis); Fama (1976) (CAPM and much else); Mandelbrot (1997) (fat-tailed return distributions); Embrechts et al. (1997) (extreme-value methods); Jorion (2007) (value at risk); Engle (1995), Shephard (2005), and Taylor (2007) (volatility); Dacorogna et al. (2001) and Hautsch (2012) (high-frequency data and realized volatility); Glasserman (2003) (Monte Carlo methods in options pricing); Connor et al. (2010) (portfolio management); Engle (2002) (multivariate volatility); Diebold and Rudebusch (2013) (bond markets and the yield curve); Singleton (2006) (empirical asset pricing); Duffie and Singleton (2003) and Lando (2004) (credit risk); Rebonato (2010) (stress testing); Allen and Gale (2007) (systemic risk); Aliber and Kindleberger (2005) (crowd behavior); Diebold et al. (2010) (*KuU* framework – the known, the unknown and the unknowable); and not least Taleb (2007) (healthy skepticism of almost everything).

References

- Acharya, V., L. Pedersen, T. Philippe, and M. Richardson (2010), "Measuring Systemic Risk," Manuscript, Stern School, New York University.
- Adrian, T. and M. Brunnermeier (2011), "CoVaR," Manuscript, Federal Reserve Bank of New York.
- Aliber, R.Z. and C.P. Kindleberger (2005), *Manias, Panics, and Crashes: A History of Financial Crises*, Wiley.
- Allen, F. and D. Gale (2000), "Bubbles and Crises," *Economic Journal*, 110, 236–255.
- Allen, F. and D. Gale (2007), *Understanding Financial Crises*, Oxford University Press.
- Andersen, T. G., T. Bollerslev, and F.X. Diebold (2007), "Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility," *Review of Economics and Statistics*, 89, 701–720.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2006), "Practical Volatility and Correlation Modeling for Financial Market Risk Management," In M. Carey and R. Stulz (eds.), *The Risks of Financial Institutions*, University of Chicago Press for NBER, 513–548.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2013), "Risk Measurement for Market Risk Management," In M. Harris, G. Constantinedes and R. Stulz (eds.), *Handbook of the Economics of Finance*, Elsevier.
- Andersen, T.G., T. Bollerslev, F. X. Diebold, and P. Labys (2003), "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.
- Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42–55.
- Audrieu, C., A. Doucet, and R. Holenstein (2010), "Particle Markov Chain Monte Carlo Methods," *Journal of the Royal Statistical Society B*, 72, 269–342.
- Arrow, K. J. (1964), "The Role of Securities in the Optimal Allocation of Risk-Bearing," *Review of Economic Studies*, 31, 91–96.

- Arrow, K.J. and G. Debreu (1954), "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, 22, 265–290.
- Arrow, K.J. and F.H. Hahn (1971), *General Competitive Analysis*, North-Holland.
- Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999), "Coherent Measures of Risk," *Mathematical Finance*, 9, 203–228.
- Aruoba, S.B., F.X. Diebold, M.A. Kose, and M.E. Terrones (2011), "Globalization, the Business Cycle, and Macroeconomic Monitoring," In R. Clarida and F. Giavazzi (eds.), *NBER International Seminar on Macroeconomics*, University of Chicago Press.
- Bachelier, L. (1900), *Theory of Speculation*, Doctoral Dissertation, Sorbonne, Paris. Translated and reprinted in P.H. Cootner (ed.) (1964), *The Random Character of Stock Market Prices*, Cambridge, Mass.: MIT Press.
- Barndorff-Nielsen, O. and N. Shephard (2002), "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society*, 64, 253–280.
- Barndorff-Nielsen, O.E. and N. Shephard (2004), "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics*, 2, 1–48.
- Berkowitz, J. (1999), "A Coherent Framework for Stress Testing," *Journal of Risk*, 2, 5–15.
- Black, F. and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–654.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.
- Bollerslev, T. (1990), "Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model," *The Review of Economics and Statistics*, 498–505.
- Brownlees, C. and R.F. Engle (2011), "Volatility, Correlation and Tails for Systemic Risk Measurement," Manuscript, NYU.
- Campbell, J.Y. and G.B. Tackler (2003), "Equity Volatility and Corporate Bond Yields," *Journal of Finance*, 58, 2321–2350.

- Campbell, S.D. and F.X. Diebold (2009), "Stock Returns and Expected Business Conditions: Half a Century of Direct Evidence," *Journal of Business and Economic Statistics*, 27, 266–278.
- Christensen, J.H.E., F.X. Diebold, and G.D. Rudebusch (2011), "The Affine Arbitrage-Free Class of Nelson-Siegel Term Structure Models," *Journal of Econometrics*, 164, 4–20.
- Christoffersen, P. (2012), *Elements of Financial Risk Management*, Second Edition, Academic Press.
- Clark, P.K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41, 135–155.
- Connor, G., L.R. Goldberg, and R.A. Korajczyk (2010), *Portfolio Risk Analysis*, Princeton University Press.
- Cootner, P. (1964), *The Random Character of Stock Market Prices*, MIT Press.
- Cox, D.R. (1981), "Statistical Analysis of Time Series: Some Recent Developments," *Scandinavian Journal of Statistics*, 8, 93–115.
- Dacorogna, M., R. Gencay, U. Muller, R. Olsen, and O. Pictet (2001), *An Introduction to High-Frequency Finance*, Academic Press.
- de Fontnouvelle, P., E.S. Rosengren, and J.S. Jordan (2006), "Implications of Alternative Operational Risk Modeling Techniques," In M. Carey and R. Stulz (eds.), *The Risks of Financial Institutions*, University of Chicago Press for NBER, 475–505.
- Diebold, F.X., N.A. Doherty, and R.J. Herring (2010), *The Known, the Unknown and the Unknowable in Financial Risk Management*, Princeton University Press.
- Diebold, F.X., T.A. Gunther, and A.S. Tay (1998), "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39, 863–883.
- Diebold, F.X. and C. Li (2006), "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, 130, 337–364.
- Diebold, F.X. and M. Nerlove (1989), "The Dynamics of Exchange Rate Volatility: A Multivariate Latent-Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1–22.

- Diebold, F.X. and G.D. Rudebusch (2013), *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*, Princeton University Press, in press.
- Diebold, F.X., G.D. Rudebusch, and S.B. Aruoba (2006), "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach," *Journal of Econometrics*, 131, 309–338.
- Diebold, F.X. and K. Yilmaz (2009), "Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets," *Economic Journal*, 119, 158–171.
- Diebold, F.X. and K. Yilmaz (2011), "On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms," Manuscript, University of Pennsylvania.
- Diebold, F.X. and K. Yilmaz (2012), "Better to Give than to Receive: Predictive Directional Measurement of Volatility Spillovers," *International Journal of Forecasting*, 28, 57–66.
- Duffie, D. and R. Kan (1996), "A Yield-Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.
- Duffie, D. and J. Pan (1997), "An Overview of Value at Risk," *Journal of Derivatives*, 4, 7–49.
- Duffie, D. and K.J. Singleton (2003), *Credit Risk: Pricing, Measurement, and Management*, Princeton University Press.
- Embrechts, P., C. Kluppelberg, and T. Mikosch (1997), *Modelling Extremal Events*, Springer.
- Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987–1007.
- Engle, R.F. (1995), *ARCH: Selected Readings*, Oxford University Press.
- Engle, R.F. (2002), "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, 20, 339–350.
- Fama, E.F. (1965), "The Behavior of Stock-Market Prices," *Journal of Business*, 38, 34–105.
- Fama, E.F. (1976), *Foundations of Finance*, Basic Books.

- Fleming, J., C. Kirby, and B. Ostdiek (2001), "The Economic Value of Volatility Timing," *Journal of Finance*, 56, 329–352.
- Fleming, J., C. Kirby, and B. Ostdiek (2003), "The Economic Value of Volatility Timing using Realized Volatility," *Journal of Financial Economics*, 63, 473–509.
- Froot, K.A. and J.C. Stein (1998), "Risk Management, Capital Budgeting, and Capital Structure Policy for Financial Institutions: An Integrated Approach," *Journal of Financial Economics*, 47, 55–82.
- Glasserman, P. (2003), *Monte Carlo Methods in Financial Engineering*, Springer.
- Hamilton, J.D. and G. Lin (1996), "Stock Market Volatility and the Business Cycle," *Journal of Applied Econometrics*, 11, 573–593.
- Hansen, P.R., Z. Huang, and H. Shek (2012), "Realized GARCH: A Complete Model of Returns and Realized Measures of Volatility," *Journal of Applied Econometrics*, in press.
- Hautsch, N., L.M. Kyj, and R.C.A. Oomen (2012), "A Blocking and Regularization Approach to High Dimensional Realized Covariance Estimation," *Journal of Applied Econometrics*, in press.
- Hautsch, Nikolaus (2012), *Econometrics of Financial High-Frequency Data*, Springer Verlag, New York, *Econometrics of Financial High-Frequency Data*, Springer.
- Jorion, P. (2007), *Value at Risk: The New Benchmark for Managing Financial Risk*, Third Edition, McGraw-Hill.
- Lando, D. (2004), *Credit Risk Modeling: Theory and Applications*, Princeton University Press.
- Litterman, R.B. and J. Scheinkman (1991), "Common factors affecting bond returns," *The Journal of Fixed Income*, 1, 54–61.
- Longin, F. and B. Solnik (2001), "Extreme Correlation of International Equity Markets," *Journal of Finance*, 56, 649–676.
- Lunde, A., N. Shephard, and K. Sheppard (2011), "Econometric Analysis of Vast Covariance Matrices Using Composite Realized Kernels," Manuscript, University of Aarhus.

- Mandelbrot, B.B. (1963), "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394–419.
- Mandelbrot, B.B. (1997), *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, Springer.
- Markowitz, H. (1952), "Portfolio Selection," *Journal of Finance*, 7, 77–91.
- Markowitz, H.M. (1959), *Portfolio Selection: Efficient Diversification of Investments*, Wiley.
- Merton, R.C. (1974), "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29, 449–470.
- Merton, R.C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.
- Modigliani, F. and M.H. Miller (1958), "The Cost of Capital, Corporation Finance and the Theory of Investment," *American Economic Review*, 48, 261–297.
- Nelson, C.R. and A.F. Siegel (1987), "Parsimonious modeling of yield curves," *Journal of business*, 60, 473–489.
- Rebonato, R. (2004), *Volatility and Correlation: The Perfect Hedger and the Fox*, Wiley.
- Rebonato, R. (2010), *Coherent Stress Testing: A Bayesian Approach to the Analysis of Financial Risk*, Wiley.
- Rosenberg, J.V. and T. Schuermann (2006), "A General Approach to Integrated Risk Management with Skewed, Fat-tailed Risks," *Journal of Financial Economics*, 79, 569–614.
- Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance*, 44, 1115–1153.
- Sharpe, W.F. (1964), "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance*, 19, 425–442.
- Shephard, N. (2005), *Stochastic Volatility: Selected Readings*, Oxford University Press.
- Singleton, K.J. (2006), *Empirical Dynamic Asset Pricing*, Princeton University Press.
- Taleb, N.N. (2007), *The Black Swan: The Impact of the Highly Improbable*, Random House.

- Taylor, S.J. (1982), "Financial Returns Modelled by the Product of Two Stochastic Processes: A Study of Daily Sugar Prices, 1961-1979," In O.D. Anderson (ed.), *Time Series Analysis: Theory and Practice, I*, North-Holland.
- Taylor, S.J. (2007), *Asset Price Dynamics, Volatility, and Prediction*, Princeton University Press.
- Vasicek, O. (1977), "An equilibrium characterization of the term structure," *Journal of financial economics*, 5, 177-188.
- Whaley, R.E. (1993), "Derivatives on Market Volatility: Hedging Tools Long Overdue," *Journal of Derivatives*, 1, 71-84.