

Bruno Rémillard

HEC Montréal, Montréal, QC

15.1 Introduction

As scientific disciplines, statistics and finance had their first date in 1900. They met through the work of French mathematician Louis Bachelier, whose thesis (Bachelier, 1900) used the familiar bell-shaped curve or Normal distribution to model differences between prices over time. He was also the first to use another well-known probabilistic object, Brownian motion, to evaluate stock options and model asset returns. Five years later, Albert Einstein used the same process to model the displacement of molecules in a liquid.

Statistics and finance have been together ever since. One of their offspring is financial engineering, a multidisciplinary field that combines financial theory with statistical techniques and computational tools. Financial engineers evaluate options, assess risks, manage portfolios and contribute to the development of financial products. Fundamentally, these various tasks all rely on the crucial ability to construct accurate statistical models for predicting the behavior of assets and other financial products. Bad models may lead to wrong decisions and considerable financial losses.

Selecting, fitting and validating models for financial data is key to all research and development activities in financial engineering. This is the hallmark of statistics. But as we shall see, statisticians also make significant contributions to more specialized areas of financial engineering in which estimation and prediction play a critical role, together with the numerical implementation of options and hedging evaluation techniques.

15.2 Modeling

To illustrate the challenges involved in financial data modeling, consider the following simple problem, which is nevertheless typical of the issues faced daily by traders on the stock market. Imagine that the current market value of one share from Apple is \$450 and that for \$18, you can purchase a call option on this asset. If you do, you earn the right to buy one share from Apple for \$460 (the strike price) at any moment of your choice within the next two months (the maturity period). If the price of this share stays below \$460 in that period, you will never exercise the option and you will have wasted \$18. But if the price goes up and you exercise your option when the share is at \$520, say, you will have saved $\$520 - \$460 - \$18 = \42 . Is this an interesting business opportunity?

To answer this question, we must first choose a plausible model for the way in which the underlying asset varies over time. Then we must use previous data to check whether the assumptions of this model hold, at least approximately. If the model is satisfactory, we can then assess whether the price of the asset and the option are “fair” in some way or other. A rational decision to buy the option or not can then be made. But whether the decision is right depends critically on the choice of model and, more specifically, on its appropriateness for the data at hand.

The financial literature abounds with statistical models for option pricing. The most commonly used model was articulated by the American economist Fischer Black and his co-author, Canadian-born financial economist Myron Scholes, in their 1973 paper, “The Pricing of Options and Corporate Liabilities,” published in the *Journal of Political Economy* (Black and Scholes, 1973). American economist Robert Merton later contributed to the refinement of the model.

15.2.1 Black–Scholes–Merton Model

The BSM model due to Black, Scholes, and Merton makes assumptions on the probability distribution of the return of an asset as it varies in time. Suppose for instance that the price of an asset is observed a times $1, 2, \dots$. Denote these prices by S_1, S_2, \dots . The return of the asset at time $k > 1$ is then defined by

$$R_k = \log(S_k/S_{k-1}).$$

In its simplest form, the BSM model then states that R_k has the same bell-shaped or Normal distribution at every time point k in time, and that this distribution depends in no way on past observed values R_1, \dots, R_{k-1} . The latter hypothesis is called serial independence.

The hypotheses of the BSM model are quite strong. In particular, the choice of the bell-shaped curve for the returns means that wild variations in




FIGURE 15.1: Prediction intervals for future prices (left) and returns (right) for the next 21 trading days predicted by the Black-Scholes model when the current price is \$450.

prices are unlikely to occur. The standard deviation is a measure of the spread of a distribution. For a Normal distribution, there is only a .3% probability that a future return will lie more than three standard deviations away from the mean. For other distributions, this probability can be much greater, up to about 10 or 11%.

For example, based on recent observations of Apple shares (2009–11), typical values for the mean of the daily returns is .003, while the standard deviation is .02. Under the hypotheses of the BSM model, 99.7% of the future returns should thus lie in the interval $.003 \pm .06$, that is, between $-.057$ and $.063$. Prediction intervals of the prices and returns for the next 21 days are given in Figure 15.1. The interpretation of these intervals is that the probability the stock price, or return, on any given day lies between the limits in the figure is 99.7%. Notice that the prediction intervals for returns all have the same width; this is a consequence of the model assumptions.

Merton and Scholes won the Nobel prize in Economics in 1997 for their work on this subject (unfortunately, Black died in 1995). They were aware that both assumptions in their basic model were debatable; students and colleagues helped them to refine it in later work. The most glaringly simplistic assumption is the hypothesis of serial independence: in practice, financial returns are almost always dependent on previous values, and hence any decision based on a model that ignores this fact is bound to spell trouble eventually.

There is a whole slew of statistical procedures for testing whether a financial series is serially dependent or not. For recent contributions by Canadian statisticians, see, e.g., Genest and Rémillard (2004) and Genest et al. (2007). Having rejected the hypothesis of serial independence, one must turn to time series models. Such models are far better at reflecting the variation of returns and other financial variables in time, but they are also more complex. Their use requires much more sophisticated statistical tools, many of which are still under development.

15.2.2 Extensions of the BSM Model

In most financial time series models, the current value of returns is not only a function of previous values but also of exogenous variables such as interest rates or tax structures. Good models involve parameters that can be adjusted to fit the data at hand, and the residual (hopefully small) fluctuations that remain unexplained are accounted for by an unobservable error term called an innovation. For example, it is typical to assume that the return R_k at time k is of the form

$$R_k = \mu_k + \sigma_k \varepsilon_k, \quad (15.1)$$

where μ_k and σ_k represent the mean and standard deviation of the distribution of R_k . The innovations $\varepsilon_1, \varepsilon_2, \dots$ are generally assumed to be independent and to have the same distribution with mean 0 and standard deviation 1. The formulas for μ_k and σ_k typically involve previous returns and other unknowns.

As an illustration, suppose that we take

$$\mu_k = \mu + \phi R_{k-1}, \quad \sigma_k = \sigma > 0.$$

The first equation amounts to saying that the average return value at time k is a linear function of R_{k-1} with slope ϕ and y -intercept μ . The second equation states that all returns R_1, R_2, \dots have the same standard deviation σ . This model is called the “autoregressive model of order 1” or AR(1) for short. The original Black–Scholes model is an AR(1) model with $\phi = 0$.

When the variability or volatility in the returns is thought to change with time, a popular option is to take

$$\sigma_k^2 = \omega + \alpha \sigma_{k-1}^2 \varepsilon_{k-1}^2 + \beta \sigma_{k-1}^2,$$

in which case the model is called GARCH, as in “generalized autoregressive conditionally heteroscedastic.” American economist Robert Engle won the 2003 Nobel Prize in Economics, sharing the award with British economist Clive Granger, for methods of analyzing economic time series with time-varying volatility.

Figure 15.2 shows 99.7% prediction intervals for the prices and returns for 1 Apple share based on a GARCH model with Normal innovations and parameters $\mu = .0032$, $\omega = .1356 \times 10^{-4}$, $\alpha = .0881$, and $\beta = .87$. Here we assumed $\phi = 0$. These parameter values were chosen on the basis of the same data as for Figure 15.1; see p. 233 in Rémillard (2013) for details.

Comparing Figures 15.1 and 15.2, one can see that the range of values is larger for the GARCH model than for the Black–Scholes model. While the Black–Scholes AR(1) model assumes that the returns are independent and identically distributed, the GARCH extension makes the more realistic assumption that the return at time k depends on the performance of the stock at earlier dates. As a result of this serial dependence, the prediction intervals for the future return values are not all identical; they vary as a function of time. However, as can be seen from Figure 15.2, they stabilize after a while. In other




FIGURE 15.2: Prediction intervals for future prices (left) and returns (right) for the next 21 trading days predicted by the GARCH model when the current price is \$450.

words, the distribution of the long-term return values gradually stabilizes or converges to a so-called stationary distribution. In particular, the long-term prediction becomes a constant, which is usually the mean of the stationary distribution.

In practice, it sometimes happens that series of financial returns move quite a distance from the mean of their stationary distribution. However, they will eventually come back to it. This so-called “return-to-the-mean” property is a consequence of the famous ergodic theorem, established by the American mathematician George David Birkhoff. It is therefore important for modelers and financial analysts to consider models that have a (unique) stationary distribution. It is not always possible to produce a formula for this distribution but for GARCH models, say, it is relatively easy to simulate it using computer-intensive methods. One could, for example, use the Monte Carlo Markov Chain approach (MCMC) discussed by Jeffrey Rosenthal in Chapter 6.

15.2.3 Choice of Distribution for the Innovations

Another important aspect of the model building process is the choice of distribution for the innovations $\varepsilon_1, \varepsilon_2, \dots$. The predictive properties of the model very much depend on this choice. The Normal curve is only one option. Another popular choice is the Student t distribution, which has fatter tails, meaning that its range of likely values is larger than for the Normal. This choice is often dictated in practice by the presence in the series of a larger proportion of extreme observations than we would expect under the Normal paradigm.

Figure 15.3 shows the impact of replacing the Normal distribution by a Student t_5 in the GARCH model described above. As one can see, the prediction intervals are now wider. Suppose, for example, that the current price of one Apple share is \$450 and that we wish to predict its price in 21 days, based on the GARCH model. If the Normal distribution is used, there is a




FIGURE 15.3: Prediction intervals for future prices (left) and returns (right) for the next 21 trading days predicted by the GARCH model with Student t_5 innovations when the current price is \$450.

99.7% chance that the price will fall between \$366 and \$633. If a Student t_5 distribution is used, however, there is about a 99% chance that it will fall in this interval. The 99.7% prediction interval for the price is now much wider; it extends from \$352 to \$658!

What principles can guide us in choosing an appropriate distribution for the innovations? A major stumbling block is that for models defined by equation (15.1), the variables $\varepsilon_1, \varepsilon_2, \dots$ cannot be observed or measured except in the rare (not to say unrealistic) cases where all model parameters are known. We are thus faced with a situation where we cannot even plot a simple histogram of these innovations to guide us in the choice of their distribution.

The solution to this problem is to replace the unobservable innovations by proxies, that is, quantities that can be computed from the data and which closely resemble or imitate the innovations. The best proxy candidates are the so-called model residuals, denoted e_1, e_2, \dots . The k th residual can be computed as

$$e_k = \frac{R_k - \hat{\mu}_k}{\hat{\sigma}_k}.$$

In this formula, $\hat{\mu}_k$ and $\hat{\sigma}_k$ are estimated values of μ_k and σ_k , respectively. Histograms and other graphical representations of the residuals can then be plotted to help us choose an appropriate distribution.

At this stage, it is tempting to go one step further by trying to apply standard statistical tests to the residuals in order to check various assumptions about the innovations ε_k . As it happens, however, conventional tests were developed for variables that are truly observable and independent of one another. This is not the case for residuals and, as many statisticians have begun to realize, this sometimes makes the tests non-operational, because when residuals are used, the distributions of the statistics are different from those when the mean and standard deviation are known, and not estimated. For

details, see Bai (2003) or Ghoudi and Rémillard (1998, 2004), among others. Therefore, new methodologies adapted to residuals must be developed.

15.2.4 Challenge of Model Validation

Today, the vast majority of articles in financial engineering, including doctoral dissertations, involve financial data modeling. The major journals not only require that the models be thoroughly investigated from a theoretical point of view, but also that they be tried out in applications and that their characteristics be showed to fit the data at hand.

Similarly, financial institutions need to validate the assumptions of the models they rely on to develop new financial products. For obvious reasons, it is increasingly difficult to convince investors of the potential of a new financial product if its relevance and the adequacy of the models on which it is based cannot be justified properly. It is imperative, therefore, to understand the effect of the use of residuals on statistical tests and procedures, and in particular to develop new ways of validating financial models. This is currently the subject of much fundamental research in statistics, both in Canada and abroad.

In order to validate the adequacy of a proposed model for a given dataset, a data analyst must be able to perform the following tasks:

1. *Detect change points in the distribution of innovations over time.* Change points are important characteristics of financial data. In particular, a change in the distribution is likely to occur when the market crashes. As a result, it might be unrealistic to assume that over a long period of time, the distribution of innovations is the same. This is why statistical techniques have to be used to detect change points. However, remember that the data show serial dependence, so change point tests would be applied to residuals. Recent results suggest that this can be done (Rémillard, 2012).
2. *Find a plausible model for the distribution of returns and estimate the parameters as precisely as possible.* Suppose for instance that a GARCH model has been fitted under the assumption that the innovations are Normally distributed. A specification test must then be used to check whether this assumption is supported by the data; see, e.g., Ghoudi and Rémillard (2013) and references therein.
3. *Check whether the innovations are serially independent.* Testing this hypothesis using the residuals is not an easy task. In fact it can be very difficult, especially for GARCH models; see Berkes et al. (2003), among others. Canadians have contributed to the development of new statistical techniques to handle this problem (Genest et al., 2007). Their approach works for simple models but more research is needed to develop appropriate techniques for more complex models like GARCH.

4. *If necessary, model dependence between time series.* When several financial series are involved, as is often the case in practice (options on several securities, etc.), one must also model the dependence between innovations in each series and test the adequacy of the dependence model for these multivariate data. Copulas can be used to this end (Embrechts et al., 2002), but much additional work will be required to develop their use to its full potential. For more information on this topic, see Chapter 4 by Christian Genest and Johanna Nešlehová, and Chapter 8 of Rémillard (2013) for applications in a financial engineering context.

15.3 Applications to Financial Engineering

Once financial data have been modeled correctly, one can go a step further and implement the model. In this section, we will review briefly three important financial applications for these models: portfolio management, option pricing, and risk management.

15.3.1 Portfolio Management

Following Bachelier (1900), the importance of the Normal distribution in finance was reaffirmed by American economist Harry Markowitz. In his doctoral dissertation, Markowitz (1952) used a weighted sum of Normally distributed random variables to model the return of a portfolio. More precisely, if $\omega_1, \dots, \omega_d$ represent the fractions of the wealth invested in $d \geq 2$ possibly risky assets with returns R_1, \dots, R_d , the return of the portfolio with weights $\omega_1, \dots, \omega_d$ is then

$$P_\omega = \sum_{j=1}^d \omega_j R_j.$$

Note here that we are assuming $w_1 + \dots + w_d = 1$; in other words, all the money must be invested in the assets (one of them could be a bank account). The weights are usually positive, but they could be negative too. A negative weight corresponds to a strategy called short-selling; it consists of cashing the present value of the asset without selling it, but the asset must be bought at the end of a given period. It is basically like borrowing the value of a stock.

In financial markets, short-selling is often used to reduce the risk in a portfolio. To illustrate, suppose that at one point in time, an investor has in hand one share of Apple listed at \$450 and that he chooses to be “short” 10 shares of Microsoft listed at \$30 per share. This means that he will cash immediately $10 \times \$30 = \300 , thereby reducing his current investment to \$150.

The weights associated with this investment strategy are $\omega_1 = 450/150 = 3$ for Apple and $\omega_2 = -300/150 = -2$ for Microsoft. This way, $\omega_1 + \omega_2 = 3 - 2 = 1$. In return, however, the investor will need to buy 10 shares of Microsoft at some future point in time, whatever its price may be at the time. One advantage of this strategy is that if the values of the two stocks go down, the investor can still make money if Apple performs better than Microsoft. For example, if both stocks decrease by \$1, the value of the portfolio will be \$459. From the borrowed \$300 obtained by short-selling Microsoft, you just have to reimburse \$290.

In the above example, $\omega_1 = 3$ and $\omega_2 = -2$ are only one possible choice of weights. Can we do better? Using the standard deviation as a measure of the risk of a portfolio P_ω , Markowitz looked at the problem of minimizing this risk for a given average return; he also considered the complementary problem of maximizing the expectation $R_\omega = E(P_\omega)$ for a given level of risk. Assuming only that the weights sum up to 1, Markowitz was able to solve this optimization problem using elementary calculus. It can be shown that the variance (i.e., the square of the standard deviation) of the portfolio P_ω is given by

$$\sigma_\omega^2 = \sum_{i=1}^d \sum_{j=1}^d \omega_i \omega_j \sigma_i \sigma_j \rho_{ij}. \quad (15.2)$$

Therefore, the problem is to minimize σ_ω^2 , subject to the constraints

$$\sum_{j=1}^d \omega_j = 1 \quad \text{and} \quad \sum_{j=1}^d \omega_j \mu_j = \mu.$$

Here μ_j and σ_j are respectively the mean and standard deviation of return R_j , μ is the target mean of the portfolio, and ρ_{ij} is the so-called Pearson correlation coefficient between the returns R_i and R_j , which measures the degree of dependence between them. This leads to the concept of efficient frontier, which is the graph of the maximal expected return as a function of the standard deviation. For this and subsequent work on portfolio management, Markowitz was awarded the Nobel Prize in Economics in 1990.

For example, using the portfolio with weights $\omega_1 = 3$ and $\omega_2 = -2$, we end up with an average return of .2627 and a standard deviation of .0514, as marked by a square on the graph of the efficient frontier in Figure 15.4. This portfolio is optimal for a standard deviation of .0514, but not for a standard deviation of .03, as seen from the graph.

The field of portfolio management, which basically consists in finding optimal weights of a portfolio changing from period to period, is a very active field of research from the statistical point of view. The issue of non-Normal returns is particularly relevant. Furthermore, portfolio managers may be interested either in maximizing a utility function, which represents the value of wealth of the investor, or in minimizing a risk measure, such as the standard deviation.

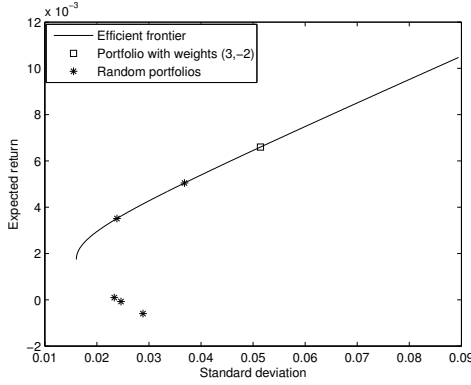


FIGURE 15.4: Graph of the efficient frontier, position of the portfolio with weights $\omega_1 = 3$ and $\omega_2 = -2$, and positions of some portfolios with random weights.

Among the problems of interest for statisticians, one can look at constraints on the weights, discrete time and continuous time models, etc. The computation of the optimal solution for a portfolio composed of a large number of assets is also quite challenging. Several statisticians working in Canadian universities have contributed to the field; see Watier (2003), Vaillancourt and Watier (2005), Labbé and Heunis (2007), and Elliott and Siu (2009), to name a few.

15.3.2 Option Pricing

Going back once again to the case of Apple, assume that we have in hand a theoretical model that has been validated with data. The next challenge then lies is the determination of a “fair” price for this option. A price is considered fair if neither the buyer nor the seller of the option have a chance of making money while being sure of not losing any money in the process.

The most common way of determining a fair price for an option consists of choosing an equivalent probability distribution, called the equivalent martingale measure, under which the actualized value of the asset in the future is a martingale. A martingale has the property that given the current value, the expected value at any future time is the same as the current value. To understand the concept of martingale, suppose for an instant that the value of an asset would correspond to a gambler’s fortune. To say that the fortune is a martingale means that there is no optimal way for the gambler to decide when to stop playing; the average return is the same, whatever the player’s strategy.

The existence of an equivalent martingale measure guarantees that there exists a fair price. In practice, the choice of the equivalent martingale measure is dictated by the previous real market prices of the option. Then the future prices of any option can be theoretically determined.

For example, the value of the call option described at the beginning of Section 15.2 is \$18.04 if we assume that the Black–Scholes model is correct and the annual interest rate is 2%. As it happens, there is only one equivalent martingale measure in this case. Assuming that the model is correct, we can compute the chances that a buyer will not exercise the call option. This probability is approximately 22%. In this case, the buyer will actually lose \$18.10, which is the value of \$18.04 in two months, assuming an annual interest rate of 2%.

When the price of the stock is above the strike price of \$460, the buyer will also lose money if the future price is lower than $\$460 + \$18.04 = \$478.04$. According to the model, this happens 33% of the time. Overall, the negative gain (or loss) turns out to be $-\$15$ while the average positive gain is \$63.

A histogram of the distribution of the net gain is given in Figure 15.5. Note that the average net gain is \$37.03. Therefore, here it seems that the buyer has a net advantage over the seller of the option. However, if the seller could trade continuously on Apple, he could generate exactly the payoff of the option, i.e., what is due to the seller at maturity, by building a portfolio composed only of cash and a fraction of the asset. This was shown in the article of Black and Scholes (1973). It means that the position of the seller is not risky at all, while the position of the buyer is quite risky. In addition, if the price paid for the option is larger than the fair price \$18.04, there is a way for the seller to make money for sure. If the price paid for the option is less than \$18.04, then there is a way for the buyer to make money for sure.

Given an equivalent martingale measure, one cannot find in general an explicit formula for an option price. However, the option price being an average value, or expectation, one can use simulation algorithms to price it. This was first proposed by Canadian-based finance Professor Phelim Boyle (1977). For other applications of Monte Carlo simulation algorithms in finance, see, among others, the books by American mathematician Paul Glasserman (2004) and Canadian statistician Don McLeish (2011) from the University of Waterloo. Monte Carlo methods are particularly effective to demonstrate to potential investors that a proposed product or strategy is interesting for them, as was done in the previous example.

15.3.3 Risk Management

Risk management is another key field of applications of statistical methods in financial engineering. The importance of the field is such that some even believe that risk assessment should be integrated into the business student curriculum (*The Economist*, 2012). An excellent reference on risk management and statistical methods is the book by McNeil et al. (2005).




FIGURE 15.5: Histogram of the net gain for buying the option obtained by simulating 10,000 values of the stock price at maturity.

Since 1988, the Basel Committee on Banking Supervision has issued three supervisory accords — Basel I, II and III — that comprise recommendations on banking regulations. These recommendations are being implemented by many financial institutions worldwide. In particular, the Basel II Accord, initially published in June 2004, called for the computation of risk measures and capital requirements to guard against three important sources of risks: risks of default of debt payments of borrowers (credit risk), risks of losses on the markets (market risk), and risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events (operational risk).

At the heart of risk management is the issue of how to measure risk. Beginning with Markowitz (1952), the standard deviation was the measure of risk in use for a long period. Since the Basel Accords, the so-called Value-at-Risk (VaR for short) has replaced the standard deviation as a measure of market risk. The VaR corresponds to the quantile (usually of order 99.9%) of a loss, meaning that 99.9% of the losses should be smaller than the VaR. Stimulated by the Basel I Accord, Artzner et al. (1999) proposed axioms that should be satisfied by what they called “coherent risk measures.” Today, the issue of risk measurement is by no means settled and many researchers, including statisticians, continue to work on this problem.

One of the main challenges in risk management is how to model individual risk factors and account for their interdependence. Two areas of modern statistics are relevant to this problem: copula modeling and extreme-value theory. The latter is particularly important for market and operational risks; see, e.g., Roncalli (2004). As its name indicates, extreme-value theory deals with the

modeling of extremal events, such as sample maxima or losses exceeding high thresholds. Statistical tools developed in that context allow researchers to extrapolate beyond the range of observable data. Such extrapolation is required for the computation of risk measures and financial reserves, and extreme-value techniques are particularly useful because loss distributions typically have considerably fatter tails than the omnipresent Normal distribution.

Extreme-value theory has been an integral part of probability and statistics for over 50 years. Several Canadian statisticians have contributed to it, including UBC Professor Harry Joe and my colleague at HEC Montréal, Debbie Dupuis, but also my frequent coauthors, Christian Genest and Johanna Nešlehová, currently both at McGill. Yet some people are still oblivious to these developments. In his book which popularized the expression “Black Swans,” Lebanese American essayist Nassim Nicholas Taleb (2007) even goes as far as blaming statisticians for not studying extremal events!

Another important risk management challenge is modeling times until default in the context of credit risk. Surprisingly perhaps, models and techniques developed for the analysis of survival data in medical studies come in handy; see, e.g., the book by University of Waterloo professor Jerry Lawless (2003). One specific twist of credit risk modeling is that defaults do not occur often, and in fact you would prefer that they never occur! Simplifying assumptions typically need to be made in order to model default times of firms for which no defaults were ever observed. The articles by Merton (1974) and Jarrow et al. (1997) count among the seminal contributions to this issue. Modeling default times on mortgages and credit cards is also quite a challenge. Yet another problem is how to account for dependence between default times. In fact, such dependencies played a major role in the so-called subprime crisis that led to a financial crisis and a subsequent recession in 2008. In that context, the risk of “contagion,” producing interdependence of defaults, was underestimated by financial institutions and rating agencies.

Risk management is a fledgling field of research. For example, as this chapter is being written, liquidity risk is a hot and emerging topic. This risk has not yet been precisely articulated but in broad terms, it refers to the fact that at some point, there are too few buyers or sellers. How best to measure liquidity risk is a subject of debate and finding appropriate ways of estimating it is one of the future challenges; see, e.g., Jarrow et al. (2012).

15.4 Final Comment

The recent financial crisis has put the use of statistics in finance in a spotlight. The use and abuse of mathematical models and statistical methodology has been the object of much polemic in the popular press. For example, former

French Prime Minister Michel Rocard wrote, in the 2–3 November 2008 edition *Le Monde*,

“On reste trop révérencieux à l’égard de l’industrie de la finance et de l’industrie intellectuelle de la science financière. Des professeurs de maths enseignent à leurs étudiants comment faire des coups boursiers. Ce qu’ils font relève, sans qu’ils le sachent, du crime contre l’humanité.”

The response of senior French probabilists and academicians Jean-Pierre Kahane and Marc Yor is well worth reading (Kahane et al., 2009). In the English language press, British journalist Felix Salmon (2009, 2012) wrote

“A formula in statistics, misunderstood and misused, has devastated the global economy.” [The formula in question is the so-called bivariate Gaussian copula; again, see Chapter 4 by Genest and Nešlehová.]

Are statistical techniques really so dangerous? Of course not! It is the duty of statisticians to educate financial practitioners and make them aware of the limits of their models. The same is true in every other area of application of statistical sciences.

Acknowledgments

I am grateful to Christian Genest for his time and effort in helping me finalize this chapter. I would also like to thank the editor and the associate editor for their helpful comments. This research was funded in part by the Natural Sciences and Engineering Research Council of Canada and by the Fonds de recherche du Québec – Nature et technologies.

About the Author

Bruno Rémillard is a professor of financial engineering at HEC Montréal since 2001. He is a graduate of Université Laval (BSc, MSc) and Carleton University (PhD). His research interests cut across probability, statistics and financial engineering, and he has consulted extensively on finance and investment, helping to develop and implant new quantitative methods for both alternative and traditional portfolios. He is the author of three books, including *Statistical Methods for Financial Engineering* (Chapman & Hall/CRC, 2013).

Bibliography

- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9:203–228.
- Bachelier, L. (1900). Théorie de la spéculation. *Annales scientifiques de l'École normale supérieure, 3^e série*, 17:21–86.
- Bai, J. (2003). Testing parametric conditional distributions of dynamic models. *The Review of Economics and Statistics*, 85:531–549.
- Berkes, I., Horváth, L., and Kokoszka, P. (2003). Asymptotics for GARCH squared residual correlations. *Econometric Theory*, 19:515–540.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–654.
- Boyle, P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*, 4:323–338.
- Elliott, R. J. and Siu, T. (2009). On Markov-modulated exponential-affine bond price formulae. *Applied Mathematical Finance*, 16:1–15.
- Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In *Risk Management: Value at Risk and Beyond (Cambridge, 1998)*, pp. 176–223. Cambridge University Press, Cambridge.
- Genest, C., Ghoudi, K., and Rémillard, B. (2007). Rank-based extensions of the Brock–Dechert–Scheinkman test for serial dependence. *Journal of the American Statistical Association*, 102:1363–1376.
- Genest, C. and Rémillard, B. (2004). Tests of independence or randomness based on the empirical copula process. *Test*, 13:335–369.
- Ghoudi, K. and Rémillard, B. (1998). Empirical processes based on pseudo-observations. In *Asymptotic Methods in Probability and Statistics (Ottawa, ON, 1997)*, pp. 171–197. North-Holland, Amsterdam.
- Ghoudi, K. and Rémillard, B. (2004). Empirical processes based on pseudo-observations. II. The multivariate case. In *Asymptotic Methods in Stochastics*, volume 44 of *Fields Institute Communications*, pp. 381–406. American Mathematical Society, Providence, RI.
- Ghoudi, K. and Rémillard, B. (2013). Comparison of specification tests for GARCH models. *Computational Statistics & Data Analysis*, in press.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer, New York.
- Jarrow, R. A., Lando, D., and Turnbull, S. M. (1997). A Markov model for the term structure of credit risk spreads. *The Review of Financial Studies*, 10:481–523.

- Jarrow, R. A., Protter, P., and Roch, A. F. (2012). A liquidity-based model for asset price bubbles. *Quantitative Finance*, 12:1339–1349.
- Kahane, J.-P., Talay, D., and Yor, M. (2009). Finance, politique et mathématiques, quels liens? *Image des Maths, CNRS*, May 2009.
- Labbé, C. and Heunis, A. J. (2007). Convex duality in constrained mean-variance portfolio optimization. *Advances in Applied Probability*, 39:77–104.
- Lawless, J. F. (2003). *Statistical Models and Methods for Lifetime Data*, Second Edition. Wiley, Hoboken, NJ.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7:77–91.
- McLeish, D. L. (2011). *Monte Carlo Simulation and Finance*. Wiley, New York.
- McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, Princeton, NJ.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29:449–470.
- Rémillard, B. (2012). *Nonparametric Change Point Problems Using Multipliers*. Technical Report, SSRN Working Paper Series No. 2043632.
- Rémillard, B. (2013). *Statistical Methods for Financial Engineering*. Chapman & Hall, London.
- Roncalli, T. (2004). *La gestion des risques financiers*. Economica, Paris.
- Salmon, F. (2009). Recipe for disaster: The formula that killed Wall Street. *Wired Magazine*, 17.03.
- Salmon, F. (2012). The formula that killed Wall Street. *Significance*, 9(1):16–20.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Penguin, London.
- The Economist* (2012). Academic view: Underestimating risk. October 31st Edition.
- Vaillancourt, J. and Watier, F. (2005). On an optimal multivariate multiperiod mean-variance portfolio. *Mathematical Reports of the Academy of Science, The Royal Society of Canada*, 27:92–96.
- Watier, F. (2003). *Optimisation stochastique dans le cadre de la gestion de portefeuilles*. Doctoral dissertation, Université de Sherbrooke, Sherbrooke, QC.