

THE ESSENTIALS OF RISK MANAGEMENT

Third Edition (EoRM3)

ONLINE CONTENT FOR CHAPTER 5

Chapter 5 [Updated] of EoRM2

From pp. 183–201 of *The Essentials of Risk Management, Second Edition* (EoRM2)

CHAPTER 5 [UPDATED]

A User-Friendly Guide to the Theory of Risk and Return

While risk management is a practical activity, it cannot be understood independently of a body of academic research about risk and reward in financial markets. It's difficult to work out the trade-off between retaining and avoiding risk without reference to the theory of risk valuation; after all, risk management does not mean the complete elimination of risk.

In this chapter, we review five key theoretical models and demonstrate how they relate both to one another, and to the practice of risk management. We'll look at modern portfolio theory (MPT), the capital asset pricing model (CAPM), arbitrage pricing theory (APT), the classic Black-Scholes (BS) approach to pricing an option, and the Modigliani-Miller (M&M) theory of corporate finance. We'll also take a brief look at the fast-developing field of behavioral finance.

Like all theories of risk management, the key theoretical models are based on simplifying assumptions. Real life is complicated and includes many details that models cannot, and maybe should not, accommodate. The role of models is to highlight the most important factors and the relationships among these factors. A "good" financial model is one that helps the analyst separate the wheat from the chaff—that is, to identify the major explanatory variables from a noisy background. We statistically reject models that fail to predict future economic behavior or trends.

As Milton Friedman made clear in his seminal 1953 article "The Methodology of Positive Economics," a model should be evaluated only in terms of its predictive power.¹ It can be simple and yet be judged successful if it helps predict the future and improves the efficiency of the decision-making process. Despite criticism of the use of models in risk management following the 2007–2009 global financial crisis (GFC)—mainly issues of model selection, implementation, and overinterpretation—we strongly believe that models and theories are essential to modern risk management.

1. M. Friedman, "The Methodology of Positive Economics," in *Essays in Positive Economics* (Chicago: University of Chicago Press, 1953).

Harry Markowitz and Portfolio Selection

The foundations of modern risk analysis are to be found in a seminal paper by Harry Markowitz written in 1952, based on his PhD dissertation at the University of Chicago, concerning the principles of portfolio selection.² (Markowitz was awarded the Nobel Prize in Economics for this research in 1990.) Markowitz showed that rational investors select their investment portfolio using two basic parameters: expected profit and risk. While “profit” is measured in terms of the average (mean) rate of return, “risk” is measured in terms of how much returns vary around this average rate of return. The greater the variance of the returns, the riskier the portfolio.


When building a portfolio, investors like to reduce variance as much as possible by diversifying their investments. To put it simply, they avoid putting all their eggs in one basket. Even better, by investing in assets that fluctuate in different directions, investors can actively offset the specific risks inherent in individual stocks. We can see the same behavior in individual firms. Following such a business strategy, for example, Head Corp., which focused initially on ski equipment, diversified into supplying tennis equipment. This strategic move helped to reduce the impact of weather and season on its periodic profits.

According to Markowitz, the result is that investors select financial assets, such as stocks and bonds, for their portfolio based on each asset’s contribution to the portfolio’s *overall* mean and variance. It follows that we must think of the risk of a single investment not in terms of its own variance, but in terms of its interaction with other assets in the portfolio.

Through the power of portfolio diversification, investors can dilute (i.e., reduce) the risk that is specific to an individual stock at virtually no cost. While it is true that the mitigation of risk may also lead to the lowering of expected profits, if assets are selected carefully, then diversification can allow investors to achieve a higher rate of return for a given level of risk.

To the extent that investors succeed in achieving this state, they arrive at the *efficient frontier*, represented by the curved solid line in Figure 5-1. Put formally, this efficient frontier contains all portfolios of assets such that there are no other portfolios (or assets) that for a given amount of risk (in terms of standard deviation of rates of return) offer a higher expected rate of return.

FIGURE 5-1 The Efficient Frontier of Markowitz



2. H. M. Markowitz, “Portfolio Selection,” *Journal of Finance* 7, 1952, pp. 77–91.

For example, portfolio P in Figure 5-1 has the same amount of risk as portfolio A, but P has a higher expected rate of return. There is no portfolio in Figure 5-1 with the same amount of risk as P that also exhibits a higher expected rate of return than P.

Once a portfolio contains only assets that are on the efficient frontier, it can be seen that a higher expected rate of return can be achieved *only* by increasing the riskiness of the portfolio. Conversely, a less risky portfolio can be achieved *only* by reducing the expected return on the portfolio. The lower part of the frontier, which contains all the inefficient assets and portfolios, is represented by a dotted line. It indicates the most inefficient combinations of assets with the lowest possible expected return for a given level of risk.

We can extend this concept to consider the whole investment market. In this framework, if the market is in equilibrium, then portfolio M, the “market portfolio,” will include all risky assets in the economy, each entering the portfolio in a proportion equal to its relative market value. For example, an imperfect but often useful proxy for all the risky equity assets in the economy of the United States is the S&P 500 index, or a wider-based index such as the Russell 2000. In Europe, the Euro Stoxx 50 index of blue-chip stocks can be used.

In this kind of market portfolio, the power of diversification means that the specific, or idiosyncratic, risk of a security is not much taken into account by the market in its pricing of a security. However, in recent years, the role of diversification “à la Markowitz” has been challenged as the average correlation of stock returns has increased dramatically from around 25 percent in the 1970s to nearly 75–80 percent during the financial crisis of 2007–2009. Correlations across asset classes have also increased substantially, even in normal market conditions. One reason invoked for this increase in the covariations of asset prices is the huge development in basket trading and exchange-traded funds (ETFs). Large baskets of assets that compose benchmark indices are traded simultaneously, independently of analyst recommendations concerning the relative performance of these assets.

To adapt to this new environment, quantitative asset management techniques have been proposed that identify risk regimes and optimize portfolio allocations for each risk regime. For example, there may be periods in which market participants are very worried and uncertain about the future. Markets adjust quickly to these situations, giving rise to higher market volatility—for example, higher level of the VIX (the volatility index for the S&P 500)—and also higher credit spreads. These periods tend to be followed by quieter periods, with lower volatility and lower credit spreads. March 2020 was much affected by the outbreak of the COVID-19 pandemic in the United States, with falling markets and major increase in equities’ volatility, while the first half of 2021 is marked by low volatility regime.

Asset managers have been developing techniques to anticipate this kind of change in a market regime. When they anticipate a high-risk regime, they switch their portfolio to a very conservative asset allocation, such as investing only in money market funds; in anticipation of a low-risk regime, they switch to a more aggressive asset allocation (e.g., including equities, emerging markets, commodities, and high-yield bonds). Each asset allocation is optimized to generate the highest return for the regime it is associated with. These approaches combine risk management techniques with optimal portfolio selection in order to control the volatility of investment portfolio returns.

The Capital Asset Pricing Model (CAPM)

In the mid-1960s, William Sharpe and John Lintner took the portfolio approach to risk management one step further by introducing a model based on overall capital market equilibrium.³ For this breakthrough, Sharpe was awarded the Nobel Prize in 1990. (Lintner, a finance professor at Harvard Business School, had passed away many years earlier.) Building on Markowitz, the two professors showed that the risk of an individual asset could be decomposed into two portions:

1. Risk that can indeed be neutralized through diversification (called diversifiable or specific risk)
2. Risk that cannot be eliminated through diversification (called systematic risk)

To build their CAPM, Sharpe and Lintner added the assumption that investors can choose to invest in any combination of a risk-free asset and a “market portfolio” that includes all the risky assets in an economy. Investors therefore weight their personal portfolios as a combination of these two investment vehicles, in various proportions based on their “risk appetite.” This concept allowed Sharpe and Lintner to define the premium that investors demand for taking on the risk of the market portfolio, as opposed to investing in the risk-free asset. This “market risk premium” is simply the difference between the expected rate of return on the risky market portfolio and the risk-free rate.

For example, we might determine the market risk premium by subtracting the interest rate on an asset that is free of default risk, such as a certificate of deposit or U.S. Treasury bill, from the expected return on a market index such as the Dow Jones or S&P 500. (Agreeing on the exact market risk premium evident in real-world data has proved to be a lively area of debate among financial economists, but we won’t allow these technicalities to detain us here.)⁴

Estimates of the market risk premium tell us how much investors have to be paid to take on some notional “average” amount of market risk generated by the complete market portfolio. But how can we estimate the risk, and the risk premium, for an individual asset? According to the CAPM, if the market is in equilibrium, the price (and, hence, the expected return) of a given asset will reflect the relative contribution of that asset to the total risk of the market portfolio. In the CAPM, this contribution is accounted for by means of a factor called beta (β). Beta is often referred to as “systematic risk” in the wider literature.

More formally, an asset’s beta is a measure of the covariance between that asset’s return and the return on the market index, divided by the market’s variance. So, it is a relative measure of the risk of an asset normalized by the total market risk.

From an investor’s perspective, beta represents the portion of an asset’s total risk that cannot be neutralized by diversification in a portfolio of risky assets and for which some compensation must be demanded. Put another way, the more beta risk a portfolio manager assumes by investing in higher-beta securities, the higher the risk and also the higher the expected future rate of return of the portfolio.

3. W. F. Sharpe, “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk,” *Journal of Finance* 19, 1964, pp. 425–442; J. Lintner, “Security Prices, Risk and Maximal Gains from Diversification,” *Journal of Finance* 20, 1965, pp. 587–615.


4. For an updated empirical estimation of the market risk premiums of different countries, see the website of Professor Aswath Damodaran: http://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/ctryprem.html.

Beta is the key to working out the expected return on an individual asset. We can think of this expected return as consisting of the interest on the riskless asset (invested over the same time period as the holding period of the asset), plus the market risk premium adjusted by beta. This can be represented more formally as

$$\text{Expected rate of return on security} = \text{risk-free interest rate} + \text{beta risk} \times (\text{expected rate of return on the market portfolio} - \text{risk-free interest rate})$$

Figure 5-2 is based on Sharpe's work. It shows the *market line*, which depicts the linear relationship between the expected rate of return on any asset and its systematic risk as measured in relation to assets that are beta-efficient.

FIGURE 5-2 The Market Line Connecting the Beta Risk of Individual Assets and Their Expected Rates of Return



In Figure 5-2, the intersection with the vertical axis yields the risk-free interest rate, R_F . This rate of return reflects the yield on an asset with zero beta. Assets B and C are beta-efficient, since they lie on the market line; C is riskier than B, and therefore is expected to yield a higher return. M, the market portfolio, is also beta-efficient, and its beta is 1, by definition. Asset A is inferior, since it lies under the line, meaning that another asset (or a portfolio of assets) can be found with the same amount of beta risk but a higher expected rate of return. Asset D is a “winner,” since it is expected to yield more in relation to its risk than assets on the market line. But, if participants in the financial market realize that D is superior, they will increase the demand for D, putting pressure on its price. As the price of this asset rises, its rate of return can be expected to fall until D lies on the market line.

So how does beta vary across different kinds of securities? As a baseline, we can think of the beta factor for a risk-free asset as zero, since the returns of that asset are indifferent to fluctuations in the capital market. Likewise, the beta of the complete market portfolio is 1, since by definition the market portfolio expresses the average beta risk for the whole market and thus requires no adjustment to take into account the specific risk of the portfolio. The beta of an individual stock (or other financial asset) can have any positive or negative value, depending on its characteristics.


Let us illustrate this last point with a numerical example based on U.S. historical data. The average rate of return on the New York Stock Exchange CRSP (Center for Research in Security Prices) index over 70 years is approximately 10 percent. The average risk-free rate on risk-free U.S. government bonds is approximately 4 percent. Hence the market risk premium, on average, is 6 percent.

Now, if a given stock has a beta estimated at 0.8, its expected rate of return is 4 percent + 0.8(6 percent) = 8.8 percent. If the past predicts the future, then we can expect the average rate of return on the market index to be 10 percent and the average rate of return on the specific stock to be 8.8 percent.

In this example, beta is positive but rather lower than the market average of 1. If the beta of a stock is higher than 1, the stock is considered “aggressive,” or riskier than the market portfolio. Conversely, if the beta is lower than 1, the stock is considered “defensive,” as it will have a mitigating effect on the total risk of the portfolio.

Figure 5-3, taken from Bloomberg L.P., shows the estimate of beta for IBM stock, based on weekly rates of return, for the period July 20, 2019, to July 19, 2020. The beta is estimated as the slope of a regression line of the rates of return for IBM on the rates of return for the market index. The regression line points to a raw, unadjusted beta of 0.574.

FIGURE 5-3 Raw Beta Computation for IBM



The most interesting class of assets is that defined as having a negative beta. Negative beta denotes an asset, such as gold, that consistently moves counter to market trends—when prices in the market tend to move up, this asset tends to lose value, and vice versa. Investment in such an asset serves to lower the risk of a portfolio without necessarily reducing its expected return.

If the market is in equilibrium, then the chances of finding such a jewel are virtually nil. Market forces will naturally work to drive the price of this kind of asset up, bringing down its future expected rate of return. In other words, to reduce the riskiness of their portfolios, investors will have to sacrifice some reward.

The CAPM has become a key tool of financial economists in understanding the behavior that can be seen in the capital markets every day. But, the beta of a stock is not simply a concern of investors; it's important to the managers of any company who are concerned about share price and the creation of shareholder value. Beta has numerous day-to-day implications for managers. For example, many firms employ a hurdle rate of return to assess whether a new investment is worthwhile in terms of building

shareholder value. This hurdle rate is based on the unique rate of return that the firm thinks investors demand; that is, it is based more or less explicitly on assumptions the firm makes about its beta factor (or about the beta factor of any new project it is considering for an investment). If the firm misunderstands the demands of investors, it is likely to set the wrong hurdle rate. If it sets its target rate of return too high, it will turn down worthwhile investments; if it sets the target too low, it will make investments that offer too low a return. Either way, it will drive down its beta-adjusted returns and make its stock less attractive to investors on a risk-adjusted basis.

As we discuss in Chapter 17 of EoRM2 corporations use a range of related new risk-adjusted measures to better understand the real rate of return they offer to investors. Banks increasingly use a measure called risk-adjusted return on capital (RAROC), and nonfinancial corporations often use a related measure called economic value added (EVA). The implementation of these performance measures necessitates the estimation of the beta factor for a given activity or division of the corporation.

Performance Measures Based on Portfolio Models

In a CAPM world where the market is in equilibrium and is expected to remain in equilibrium, no investor can achieve an abnormal return—that is, an expected return in excess of the return according to the CAPM risk-return relationship. Each stock and each portfolio expects to yield an identical rate of return adjusted for risk. In terms of Figure 5-2, all securities will lie on the capital market line.

This is not the case in the real world. First, we do not know exactly how expected values are determined; we try to estimate the expected values but estimations are always subject to measurement errors. Second, markets are rarely in equilibrium, and once they reach equilibrium, deviation from equilibrium is likely to occur almost instantaneously. In this realistic situation, stocks and portfolios may yield an excess return or a return below the return with fair compensation for the risk exposure. This is why portfolio managers rely on indices to measure the performance of a given stock or a given portfolio relative to the CAPM equilibrium risk-return relationship.

In this section we compare several performance indices and illustrate how they are used in practice. In the following we focus on the three traditional measures of portfolio performance based on the CAPM: (i) the Sharpe reward to volatility ratio, (ii) the Treynor reward to volatility ratio, and (iii) the Jensen performance index. The overall idea is the same—for yielding on average higher returns one must assume a greater amount of risk.

Sharpe Performance Index (SPI)

The slope of the Capital Market Line is the fair equilibrium compensation in term of expected return by unit of volatility of the portfolio, according to the CAPM.

The investment performance index proposed by Sharpe (SPI) is:

$$\text{SPI} = \frac{E(R_I) - r}{\sigma_I}$$

where $E(R_I)$ and σ_I are the expected return and the standard deviation, respectively, for the rates of return on any asset or portfolio. A Sharpe ratio greater than the slope of the Capital Market Line indicates a superior performance to what is expected in equilibrium, while a Sharpe ratio below the slope of the Capital Market Line indicates an inferior performance.

Treynor Performance Index (TPI)

The Treynor performance index (TPI) is:

$$\text{TPI} = \frac{E(R_I) - r}{\beta_I}$$

TPI is similar to SPI in the sense that it measures the risk premium, $E(R_I) - r$, per unit of risk, except for the choice of the risk index. Sharpe uses the standard deviation of the rates of return, while Treynor uses the beta of the portfolio.

In equilibrium, we expect this ratio to be constant across all risky assets and portfolios and equal to the excess expected return on the market portfolio: $E(R_M) - r$. Any TPI greater than $E(R_M) - r$ is considered as a superior performance while it is the contrary for a TPI below $E(R_M) - r$.

Jensen Performance Index (JPI)

Jensen's performance index (JPI) is similar to Treynor performance index as both assume that investors hold well-diversified portfolios. Running a time-series regression of portfolio excess rate of return ($R_{It} - r_t$) on the market portfolio's excess rate of return ($R_{Mt} - r_t$), one can estimate the beta of the portfolio I:

$$(R_{It} - r_t) = \hat{\alpha}_I + \check{\beta}_I (R_{Mt} - r_t) + e_{It}$$

where $\hat{\alpha}_I$ and $\check{\beta}_I$ are the regression coefficients and e_{It} is the deviation of I from the regression line in period t . Taking the mean on both sides, the residual disappears as the average deviation, $\bar{e}_{It}$, is always zero by construction:

$$\bar{R}_I - r = \hat{\alpha}_I + \check{\beta}_I (\bar{R}_M - r)$$

According to the CAPM, one expects $\hat{\alpha}_I$ to be zero in equilibrium. Hence, Jensen concentrates on the Alpha of the regression. If $\hat{\alpha}_I$ is significantly different from zero and positive, then the performance of the portfolio is considered superior, while it is considered inferior if it is negative.

The Arbitrage Pricing Theory

The CAPM is a normative theory of how the expected rate of return on a financial asset is determined. It describes the expected rate of return as a linear function of the market's risk premium, and beta risk is the coefficient, or slope, of the relationship. The arbitrage pricing theory (APT) is an extension of the logic behind the CAPM, explaining the expected rate of return on an asset as a linear function of several market factors. The APT suggests adding more factors that can contribute to the explanation of the expected rate of return, though it does not say which factors to add. It only suggests that there may be factors, such as macroeconomic factors or some stock, bond, or commodity indices, that add to the explanatory power of the relationship. The model is referred to as a multifactor or a multi-index pricing model.

This model was initially proposed in 1976 by Professor Steve Ross. It was later tested by Roll and Ross (1980) and Chen, Roll, and Ross (1986).⁵ Chen, Roll, and Ross found that the following macroeconomic factors were important in explaining the realized average rates of return on the stocks traded on the New York Stock Exchange (NYSE): the surprise in the inflation rate, the unexpected trends in GNP, changes in the default premium of bonds, and drifts in the slopes of the yield curves.

Many empirical works prefer to use the APT approach rather than the CAPM, as the latter is a special case of the former (though the CAPM has more fundamental theoretical foundations). The CAPM is a one-factor model: the market index is the only variable used to explain the expected return for any security. The APT is a multifactor model, whereby a number of different indices can be used to explain the variation in expected rates of return. Therefore, the APT is often used to decompose the factors' contributions to the expected return of stocks, so that we can see the contribution of any fundamental index used to explain the stock's expected rate of return.

How to Price an Option

The next important development in the analysis of risk arrived in 1973 with the publication of two papers on the pricing of options by Fischer Black and Myron Scholes and by Robert Merton.⁶ At the time of the publication of their seminal paper, Black and Scholes were professors at the University of Chicago, while Merton was a professor at MIT. In 1998, Merton and Scholes received the Nobel Prize. (Black had passed away in 1995.)

Options are financial assets that entitle their holders to purchase or sell another asset by or on a predetermined day, at a predetermined price called the striking price, or exercise price. An option to buy the asset is referred to as a *call*, while an option to sell the asset is called a *put*. (See Chapter 6 of EoRM3.)


The price paid up front for the call option is generally a fraction of the current price of the underlying asset on which the option (contract) is written. The remainder, meaning the exercise price, is paid at the time of execution (or exercise) at some point in the future, if the price of the asset is greater than the exercise price. One advantage of purchasing a call option is the ability to buy an asset on credit.

At the time of exercise, the purchaser retains the right to renege on his or her intention to complete the contract. If the asset price at the time of exercise is lower than the price set in the call option contract, the purchaser of the call can opt to let the contract expire. In effect, the call option offers a form of price insurance.


Figure 5-4A describes the cash flow from a call option at maturity (i.e., when it expires). The option has zero cash flow at maturity so long as the price of the underlying instrument is below the striking price (exercise price) K . For prices above the striking price, the owner of the call is entitled to the difference between the price of the underlying instrument and the striking price. This latter cash flow is described by the sloping line, increasing from point K to the right as the price of the underlying asset rises.

5. S. Ross, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory* 13(3), 1976, pp. 341–360; R. Roll and S. Ross, "An Empirical Investigation of the Arbitrage Pricing Theory," *Journal of Finance* 35(5), 1980, pp. 1073–1103; Nai-Fu Chen, R. Roll, and S. Ross, "Economic Forces and the Stock Market," *Journal of Business* 59(3), 1986, pp. 383–403.

6. F. Black and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 1973, pp. 637–654; R. C. Merton, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science* 4(1), 1973, pp. 141–183.

FIGURE 5-4A The Payoff of a Call Option at Maturity

For a put option, the reverse is true. The buyer pays for the right to sell a stock in the future at a pre-set price. That exercise price constitutes a guaranteed minimum price. On the other hand, if the market price of the asset is higher than the price the put purchaser would receive from exercising the put, he or she can choose to ignore the option contract and opt instead to sell the asset on the open market for its full price. Figure 5-4B shows the payoff of a put option at maturity, which is positive when the underlying asset price falls below the exercise price. So a put option, held with the underlying asset, provides insurance against a reduction in value of the underlying asset to below the exercise price.

FIGURE 5-4B The Payoff of a Put Option at Maturity

In their publications in 1973, Black and Scholes (frequently referred to as B&S) developed a classic model for pricing options. Merton, who collaborated with B&S, also published an important paper in 1973, offering an alternative way to prove the valuation model and many additional extensions. In addition to calculating equilibrium market prices for publicly traded options, this model specified the various components of an option and their interrelationships.

For example, as we've discussed, a call option can be characterized as a "package deal" that includes:

- Buying a stock (or other underlying asset)
- Taking out a loan
- Buying insurance against a decline in the future share price

It can also be shown that over a very short time interval, a call option can be decomposed into buying a certain ratio of the underlying instrument (this ratio is referred to as the "delta" of the option, or

the “hedge ratio”) and taking out a loan (with the amount of the loan being proportional to the probability that the option will be exercised). Options theory has proved invaluable to portfolio and risk management. Using these ideas, portfolio managers can dynamically tailor investment positions to reflect changing expectations, market conditions, and client needs. Purchasing puts against the assets held in a portfolio is synonymous with taking out insurance on those assets. Purchasing and selling combinations of calls and puts can help investors maneuver in volatile or uncertain markets. (See Chapter 6 of EoRM3.)

The detailed mathematics that lie behind the B&S model are somewhat complex and not readily calculated without the aid of computer technology. But the functions that govern an option’s price are quite intuitive. Simply stated, the price of an option is a function of:

- The price of the underlying asset
- The exercise (or striking) price set in the contract
- The prevailing risk-free rate of interest
- The volatility of the underlying asset
- The time remaining until the predetermined exercise date

If the stock price increases and we hold all the other parameters constant, then the value of the call option increases. Similarly, if the option is sold with a lower exercise price or a longer maturity, its value will be higher.

As the underlying asset becomes more volatile, the value of the call increases. This is because a call option has no downside risk—that is, no matter how far the call is out of the money at expiration, its value is still zero—while increasing volatility increases the probability that the option will end up in the money at maturity (i.e., the stock price is more likely to reach a higher value). Furthermore, as the interest rate increases, the value of the call increases (because the present value of the exercise payment in the event of exercise declines as interest rates increase). Similar arguments hold for put options, although the sensitivity of the put options to some of the factors is reversed; it is a declining function of the price of the underlying asset and the risk-free interest rate and an increasing function of the exercise price and volatility.

Of all these factors, the most crucial to the valuation and risk management of an option is the volatility of the underlying asset. It’s often said that options are “risk-friendly”: an increase in the volatility (i.e., an increase in the risk) of the underlying asset, assuming that all other parameters remain constant, leads to an increase in the price of the option, both for calls and for puts. As you may recall, volatility is measured in terms of the standard deviation (the square root of the variance) of the rate of return for the underlying asset during some selected historical period.

In applying historical data to the calculation of future volatilities, we are making a problematic assumption: that volatility remains constant over time. Where there are liquid options markets, however, the B&S model offers one way of working around this problem. Remember that the B&S model offers a way to price options provided that we have access to the inputs, listed earlier. On the other hand, if we already have the price of an option from a liquid options market, then we can use this “output” as one of the inputs. The formula can then be used to calculate a missing input, such as volatility. Using the B&S formula in this way is computing the volatilities *implied* by the prices of options in the market; that’s why this number is often called “implied volatility.”

Implied volatilities are of tremendous practical importance to those who regularly trade options in the market, and they are often re-input into the B&S model to calculate the price of a slightly different option series with different exercise prices or maturities. Because implied volatility cannot be

directly observed and is dependent on the model, it has in the past been a weak link in the operational risk management of option desks. A trader who prices option positions for risk management purposes using an implied volatility number that he or she computes from the market can be faced with some severe temptations. In the past, traders at certain investment banks have deliberately input wrong implied volatility numbers to transform the apparent value of their under-the-water options portfolio. Implied volatility is a particular worry because it's the one input into the B&S model that cannot be checked by an auditor without some degree of specialized knowledge. This is one example of how issues surrounding the principles of risk modeling affect the operational practice of risk management, a theme we'll return to in Chapter 15 of EoRM2.

Since 1993, the Chicago Board Options Exchange (CBOE) has calculated an implied volatility index for the S&P 500 Index. Known as the VIX, this volatility index is calculated using the prices of synthetic 30-day, at-the-money traded options on the stock index, an approach initially proposed in 1986 by Professors Menachem Brenner and Dan Galai of the Hebrew University.⁷ Futures on the VIX started trading in 2004, and options started trading in 2006. Since then, many other volatility indices have been published, and options and futures contracts are often traded on them. Figure 5-5 shows the evolution of the VIX from July 1997 until June 2022.

The VIX is often referred to as a fear indicator because it spikes during crises and market crashes and in reaction to unanticipated events. It is evident in the major spike to 80 percent in September 2008 with Lehman Brothers collapse, followed by spikes in 2011 and 2012 with the Sovereign Debt crises in Europe, and again with the eruption of the COVID-19 pandemic in March 2020. As such, VIX-based strategies can be used to protect equity portfolios against “gap risk”—that is, a significant unexpected drop in the equity markets. The short-term VIX responds quickly and strongly to negative events, jumping upward, sometimes quite dramatically, as shown in Figure 5-5. So, one way to buy protection against “gap risk” is to purchase futures contracts on the VIX.

7. M. Brenner and D. Galai, “New Financial Instruments for Hedging Changes in Volatility,” *Financial Analysts Journal*, July/August 1989, pp. 61–65; M. Brenner and D. Galai, “Options on Volatility,” *Option Embedded Bonds*, I. Nelken (ed.), Irwin Professional Publishing, 1997, pp. 273–286.

FIGURE 5-5 Evolution of the VIX During the Period July 2002 to July 2022

Source: Bloomberg

The B&S model also provides insights into how options introduced into a portfolio of financial assets interact with those assets and affect the overall risk of the portfolio. The systematic risk (beta) of a call or put option is a function of the beta of the underlying asset multiplied by the *elasticity* of the call or put. By elasticity, we mean the percentage change in the value of the option for a 1 percent change in the price of the underlying security. The model determines that the elasticity of a call is positive and greater or equal to 1, while the elasticity of a put is less than or equal to -1 . Hence, adding call options to a portfolio will often increase the overall risk of the portfolio (assuming a positive beta), while adding puts will have a mitigating effect on a portfolio's risk. Shorting calls—that is, writing call options—can also have a mitigating effect on the portfolio risk, since such a position has a negative beta.

The B&S model can also be used to compute the *hedge ratio* of an option position, also known as the *delta*. This ratio describes the change in value of an option resulting from a small change in the price of the underlying asset. The hedge ratio indicates how the risk of a financial asset can be hedged with options. The price of both the underlying asset and the option changes over time, so the hedge ratio is in fact dynamic, requiring that adjustments to the portfolio be made in order to maintain a target level of hedging. The hedge ratio of a call is between 0 and 1, and the delta of a put option is between -1 and 0.

For example, imagine that the delta of a call option that is slightly out of the money is 0.5, meaning that if the price of the underlying stock increases (decreases) by \$1, the value of the call increases (decreases) by \$0.50. This implies that if we want to neutralize, over a short time horizon, the risk of a long position in a call contract on 100 shares, we should sell short 50 shares.

The insights of the model introduced by Black and Scholes in the 1970s have led to further applied research in finance, particularly in relation to volatility. For example, in the last two decades, in reaction to evidence that volatility in financial markets may undergo major shifts in its behavior over time (more technically, that it is “nonstationary”), researchers have begun to make use of a more dynamic

approach to financial asset valuation. In particular, Robert Engle, a statistics professor⁸ and a leading researcher in this area, introduced the ARCH (autoregressive conditional heteroskedasticity) volatility model during the 1980s in an attempt to estimate volatility as an “auto-regressive” process. The key feature of the model is that it treats today’s volatility as a function of the previous day’s volatility and also introduces a correction factor for deviations from an expected volatility. (Robert Engle was the recipient of the Nobel Prize in Economics in 2003 for his work on volatility modeling.)⁹

Modigliani and Miller (M&M)

In order to complete this brief introduction to the theoretical basis of modern risk management, we must turn to the work published by Franco Modigliani and Merton Miller in 1958 (for which they were each awarded the Nobel Prize in Economics, Modigliani in 1985 and Miller in 1990).¹⁰ Their work does not directly involve financial markets, but rather focuses on corporate finance. Modigliani and Miller showed that in a perfect capital market, with no corporate or income taxes, the capital structure of a firm (i.e., the relative balance between equity and debt capital) has no effect on the value of the firm.

The implication of their work is that a corporation cannot increase its value by assuming greater debt, even though the expected cost of debt is lower than the expected cost of equity. Increasing the leverage of a firm (i.e., taking on more debt relative to equity) means increasing the financial risk of the firm. Equity holders (whose claims on the firm’s assets are subordinate to those of lenders and bondholders) will demand compensation for this risk and expect higher rates of return. Under M&M assumptions, the weighted average cost of capital (WACC) for a firm will remain constant regardless of the financial leverage. This is a variation on the theme that has run through this chapter: investors look not for higher returns, but for higher risk-adjusted returns.

Modigliani and Miller’s work also has important implications for discussions about capital adequacy in banks and for approaches to performance measurement such as RAROC (see Chapter 17 of EoRM2). If we accept their propositions, ignoring for the moment any corporate tax effects, shareholders should be indifferent as to whether a bank has low levels of equity and high levels of leverage, or vice versa. This should not adversely affect the value of the bank, nor the welfare of the shareholders. The Basel III intention to require much higher capital requirements from banks—that is, more equity and lower leverage—is consistent with M&M, in cases where the tax issue is not a significant factor.

Behavioral Finance

One recent trend in the financial literature is to investigate the subjective attitude of investors to risk and return and how they react to different situations in the financial markets. Various apparent anomalies in the behavior of stock markets are then explained using theories based on psychology and human cognition.

8. Previously at University of California San Diego and currently at New York University.

9. Robert Engle is the founder of the V-Lab at New York University (NYU): <https://vlab.stern.nyu.edu>. The volatility laboratory (V-Lab) provides real time measurement, modeling and forecasting of volatility, correlations for a wide range of financial assets and indices.

10. F. Modigliani and M. H. Miller, “The Cost of Capital, Corporation Finance, and the Theory of Investment,” *American Economic Review* 4, 1958, pp. 261–297.

The initial breakthrough came from the work of two psychologists, Amos Tversky and Daniel Kahneman (the latter received the Noble Prize for his work in 2002; Tversky passed away in 1996). In 1979, Kahneman and Tversky wrote “Prospect Theory: An Analysis of Decision Under Risk,” a paper that used cognitive psychology to explain various divergences of economic decision making from neoclassical theory.¹¹ They showed that investors misjudge opportunities with very high or very low probabilities.¹² The different cognitive biases have become known by nicknames such as “herding behavior,” describing the tendency to mimic the investment behavior of large groups, and “mental accounting,” which explains how investors tend to divide their investments into separate mental accounts based on criteria such as the source of the funds or the use of the funds. The “ostrich effect,” meanwhile, describes how many investors do not seem to want to see risky situations and are willing to accept a lower return for an identical risk if the risky situation is not presented to them.¹³

The major issue with behavioral finance is whether it can help in pricing securities and, especially, in pricing risks. While it can explain some anomalies, and what to theoreticians may seem irrational investment decisions, it is not yet certain that a knowledge of behavioral theories can help organizations to manage their risks more rationally.

Conclusion

The key theoretical models help us to define risk in a consistent way and indicate which measures of risk are relevant to specific situations, highlighting the importance of:

- Elimination of arbitrage opportunities when valuing financial instruments and positions
- The critical difference between idiosyncratic (specific) risk and systematic risk
- The dependence of financial modeling on key parameters and inputs

Above all, perhaps, they help us forge a rational link between the risk management perspective of the corporation and the desires of its shareholders.

11. Daniel Kahneman and Amos Tversky, “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica* 47(2), 1979, pp. 263–291.

12. The interested reader can find basic information on the common biases in Wikipedia (http://en.wikipedia.org/wiki/Behavioral_economics) and Investopedia (http://www.investopedia.com/university/behavioral_finance).

13. D. Galai and O. Sade, “The ‘Ostrich Effect’ and the Relationship Between the Liquidity and the Yields of Financial Assets,” *Journal of Business* 79(5), 2006, pp. 2741–2759.