AN INTRODUCTION TO OPTION PRICING AND THE MATHEMATICAL THEORY OF RISK

Marco Avellaneda

Courant Institute-NYU
251 Mercer St.
New York, NY 10012

Morgan Stanley & Co.
1585 Broadway
New York, NY 10036

This review paper discusses the topic of option pricing with emphasis on modeling financial risk. The Black-Scholes formula is derived using the classical dynamic hedging argument. Dynamic hedging justifies the valuation of contingent claims based on the use of risk-neutral, as opposed to "frequential", probabilities. This still leaves open – even in the simplest case of stock option contracts – the issue of specifying the volatility parameter or other charateristics of the model describing the evolution of market prices. This "specification problem" leads us to the issue of economic uncertainty, or risk, the *raison d'etre* of derivaties markets and financial intermediation. Thus, the valuation of contingent claims under uncertainty goes far beyond the exercise of computing expected values of cash-flows. After a discussion of the classical principles of option risk-management using differential sensitivities ("Greeks"), I review some more recent proposals for modeling uncertainty. The idea is to consider, as a starting point, a spectrum of risk-neutral probability measures spanning a set of beliefs and to construct option spreads to reduce uncertainty. This last part of the paper draws on work with my collaborators (Avellaneda, Levy and Parás (1995), Avellaneda and Parás (1996) and Avellaneda, Friedman, Holmes and Samperi (1997)).

Mathematical Finance has produced a true convergence of ideas between different intellectual and applied fields. Presently, we see a strong collaboration between mathematicians, economists and financial professionals in academia and the financial industry. University mathematicians contribute as Wall Street consultants and publish in finance journals. Finance academics use highly quantitative tools and sophisticated econometric analysis. It is has become almost commonplace to find traders with advanced degrees in Mathematics or Physics. As financial markets become increasingly competitive, the demand for sophisticated ideas and creative solutions increases. Markets have thus benefited from the input of mathematics and this trend shows no signs of abating.

In this paper, I present a scientific perspective of one of the cornerstones of Mathematical Finance, the theory of options pricing. This theory was initiated by Fisher Black and Myron Scholes in their seminal 1973 paper (Black and Scholes, 1973) and has grown tremendously since. It would be impossible for me to do justice to this subject in only one lecture. After giving a a general introduction, I will discuss some aspects that have interested me the most. Therefore, paper covers a lot of standard material as well as more advanced research ideas which use non-linear partial differential and ideas from optimal control theory to model risk. For a general introduction to financial markets and financial mathematics, I recommend Hull (1994).

1. Investments and probability. Perhaps a good starting point would be to analyze the situation of an individual which faces the decision of investing money. The most common investments are in stocks, bonds and cash, or short-term deposits, often through mutual funds and pension funds. Even a "no-investment" decision, such as keeping the money in the bank or spending it, is a form of asset allocation. The typical parameters used to evaluate investment decisions are **yield**, or expected return on investment, and **risk** (which is less well defined). In a very schematic form, an investment in a particular asset over a time-period Δt starting at time t gives rise to a **return**

$$\frac{X(t+\Delta t) \;-\; X(t)}{X(t)} \;=\; \frac{\Delta X(t)}{X(t)}$$

where $X(t)$ is the amount initially invested and $X(t+\Delta t)$ the value of the investment at time $t+\Delta t$. A very rough measure of the quality of an investment is obtained if we study a historical time-series of returns and calculate the empirical mean and variance for this investment:

$$\mu \;=\; \frac{1}{N\,\Delta t}\sum_{n=1}^{N}\frac{X(n\,\Delta t) \;-\; X((n-1)\,\Delta t)}{X((n-1)\,\Delta t)}\;,$$

and

$$\sigma^2 \;=\; \frac{1}{(N-1)\,\Delta t}\sum_{n=1}^{N}\left(\frac{X(n\,\Delta t)\;-\;X((n-1)\,\Delta t)}{X((n-1)\,\Delta t)}\;-\;\mu\right)^2 \;.$$

Historical (long-term) numbers for these quantities are, more or less, $\mu = 16\%$, $\sigma \approx 15\%$ for stocks, $\mu = 10\%$, and $\sigma = 12\%$ for bonds, in annual terms. However, these numbers are strongly dependent on the "time-window" or sample size.

Although investment is an art as much as a science, rational investment decisions are often based on evaluating the risk and return of different strategies. Risk-averse individuals will keep their wealth mostly in bonds or in cash (not under the mattress, though!). Investors willing to bear some market risk in exchange of higher returns might invest in stocks or long-term bonds. Long-term bonds are, in principle, riskless instruments, since they have a well defined return if held to maturity. However, the value of a bond changes because there is an "opportunity cost" in holding a bond that yields less than other instruments, *e.g.* cash or shorter-term notes.

It is important to take into consideration the fact that investment strategies can change across time. For example, investment advisors recommend "aggressive" stock portfolios to their younger clients and more "stable" portfolios consisting mostly of bonds, as these approach retirement age. Another form of time-dependent investing is "timing the market" – attempting to buy at the low and sell at the high. This is the goal of all investors, but is certainly not easy to do! The concept of **dynamic investment decision** or **dynamic asset allocation** is extremely important in Finance and Financial Mathematics. This principle applies even more so to firms and corporations in their management of capital and business decisions.

Nobel prize winner Harry Markowitz was one of the the first to propose a coherent theory of investment based on the use of probabilities (Markowitz, 1991). In his approach, the investor considers the **mean** and the **covariance matrix** of the returns of different investment lines. For instance, in the universe of a stock index, a bond index and cash, the investor would consider the 3-vector $(\mu_1,\ \mu_2,\ \mu_3)$ and the **covariance matrix**

$$\begin{matrix} \sigma_1^2 & \sigma_1\,\sigma_2\,\rho_{1\,2} & \sigma_1\,\sigma_3\,\rho_{1\,3} \\ \sigma_1\,\sigma_2\,\rho_{1\,2} & \sigma_2^2 & \sigma_2\,\sigma_3\,\rho_{2\,3} \\ \sigma_1\,\sigma_3\,\rho_{1\,3} & \sigma_2\,\sigma_3\,\rho_{2\,3} & \sigma_3^2 \end{matrix}$$

By allocating his resources among the three different assets in different proportions, the investor can construct a "portfolio" with yield

$$\sum_{i=1}^{3} w_i\,\mu_i \tag{1}$$

and variance

$$\sum_{i=1}^{3} w_i^2\,\sigma_i^2 \;+\; 2\sum_{i<j} w_i\,w_j\,\sigma_i\,\sigma_j\,\rho_{i\,j}\;. \tag{2}$$

The investor can choose the "portfolio weights" w_1, w_2, w_3 so as to maximize returns (1) holding the variance (2) fixed at a desired level. Here, the variance of the portflio is identified with the risk of the investment. This is a classical quadratic optimization problem that can be solved with elementary linear algebra. The procedure, called **mean-variance optimization**, gives a rationale for targeting the maximum mean return for a given risk level. Mean-variance optimization and it many generalizations are the most widely used tools in modern asset allocation and money management.

Implicit in Markowitz's portfolio theory is the idea that the returns are governed by probabilities. We thus (i) regard the outcome of investing as a random variable, and (ii) assume that these probabilities can be inferred from historical data. This raises the fundamental question of to what extent can historical, or **frequential**[1], probabilities predict future returns. Can statistical analysis applied to financial markets predict the future? It is clear that the answer is ultimately no. Unlike physical (mechanical) systems, markets are not "closed systems" determined completely by their initial conditions. The future behavior of the market may depend on information not available at present or by future events that we cannot control and even less model. Furthermore, the market's dependence on a set initial conditions is often murky. The point is that there is a distinction to be made between probability (the calculation of outocmes based on known odds) and risk-analysis (the estimation of outcomes in the presence of odds that are not known with certainty). John Maynard Keynes (Keynes, 1936) put it like this:

> By "uncertain" knowledge...I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty. The sense in which

[1]Based on the observed frequency of past events.

> I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention... About these matters, there is no scientific basis on which to form any calculable probability whatever. We simply do not know!

At first glance, this profound statement and its implications might lead us to the erroneous conclusion that quantitative tools are only of marginal use in finance and and economics. A deeper analysis – and reality – show that this is not so. The existence of financial risk leads to the need for hedging or diversifying it, hence to more sophisticated investment vehicles. In particular, the appearance and phenomenal growth of **derivatives** has prompted completely new applications of mathematics and probability to finance.

Derivatives are contracts that derive their value from other instruments (stocks, bonds, etc). They include options, which are the main topic of this paper. Derivatives exist because there is volatility, or risk. The effect of derivatives is what economists call "financial intermediation" : the transfer of financial risk from some individuals to others. There are people who, for a price, are willing to assume the investment risks on behalf of other investors who are risk-averse, more or less like an insurance company insures your home against casualty for a fee. The trading of derivative contracts implies that risk itself can be priced and transferred among investors in the marketplace. It is precisely in the area of modeling risk and risk-management that mathematics has proven to be extremely effective.

2. Options. An option is a contract that allows the holder to buy or sell a financial asset at a fixed price in the future. Unlike a forward contract, which consists in a commitment by two counterparties to enter into a transaction at a future date, an option needs not be **exercised** – the holder of the option will use it only if this is convenient. A **call** is an option to buy an asset and a **put** is an option to sell it. An option contract specifies the **exercise price** and the **expiration date** of the contract. For example, the 145 IBM Call of March 1997 gives the holder to buy 100 IBM shares at \$ 145 anytime between now and the third Friday in March. This option is called **American** because it can be exercised anytime before its expiration date. Options that can be exercised *only* at the expiration date are called **European**.

Options increase the spectrum of investments. For instance, an investor shares of a given stock can use options in several ways. Suppose that he thinks that the market is "overvalued" and is due for a correction. He could choose, on the one hand, to "take profits" by selling all or part of his stock holdings and perhaps buying bonds or keeping the proceeds in cash. If he does this, however, he might loose the opportunity of a further rally. With this in mind, he could choose instead to buy a put option, which gives him the right to sell the stock for a period of time at a predetermined exercise price, and maintain his stock position. In this way, he preserves the investment opportunity while insuring himself against a drop in price. This strategy is usually called a **protective put**. Note however that the purchase of the put implies paying up-front for this insurance (perhaps by selling some stock) and that this protection is valid only for a period of time. The investor who implements a protective put strategy reduces a fraction of his potential gains by purchasing the option. Another common strategy is the so called **buy-write** strategy. This is done by investors that hold stock and and believe that the market will not experience much volatility and wish to derive some "income" from the position they hold. In this case, the investor sells ("writes") calls on the stock and receives the option premium. He then becomes obliged to deliver the stock if the calls are exercised. Since he owns stock, he can meet the potential obligation implied by the option contract. There are infinitely many strategies for investment using options; the two mentioned above being the simplest. Puts and calls are traded esseantially like any other financial asset.

A longstanding problem in Finance was the valuation of option contracts. Is there a relationship between the price of the underlying asset, on the one hand, and an option contract written on this asset? This problem was solved by Fisher Black and Myron Scholes in 1973. Let us assume a **stochastic model** for the evolution of the price of the underlying asset:

$$\frac{dS_t}{S_t} \;=\; \sigma\, dZ_t \;+\; \mu\, dt, \tag{3}$$

where Z_t is a Brownian motion and σ, μ represent respectively the volatility and mean of the returns for investing in the stock. This model is just the "continuous-time" version of the "stochastic returns" model of the previous section. We shall be purposely vague about how σ and μ are determined for now. The prevailing short-term interest rate will be denoted by r.

We shall make the initial guess that the value V_t of a call on the stock is given by

$$V_t \;=\; C(S_t,\, t)\ , \tag{4}$$

where $C(S,\, t)$ is a smooth function of S and t.

Suppose that an investor sells one call option and buys Δ shares of the underlying asset at time t. The change in the value of his holdings over the interval $(t,\, t+dt)$ is

$$(\,-V_{t+dt} + \Delta\, S_{t+dt}\,) \;-\; (\,-V_t + \Delta\, S_t\,) \;=\; -\,dV_t + \Delta\, dS_t\ . \tag{5}$$

Using equations (3) and (4) and applying **Itô's formula**, we can express the variation of the portfolio in terms of the variation of call price in terms of the variation of the price of the underlying asset, *viz.*,

$$dV_t \;=\; C_S(S_t,\, t)\, dS_t \;+\; C_t(S_t,\, t)\, dt \;+\; \frac{1}{2}\,\sigma^2\, S^2\, C_{SS}(S_t,\, t)\, dt\ , \tag{6}$$

to leading order in dt. Substituting this expression into (5), we arrive at the following expression for the change in the portfolio value

$$(- C_S(S_t,\, t) \; + \; \Delta\,)\; dS_t \; - \; \left(C_t(S_t,\, t) \; + \; \frac{1}{2}\sigma^2\, S_t^2\, C_{SS}(S_t,\, t) \right)\; dt \; . \tag{7}$$

If the number of shares held in the portfolio was

$$\Delta \; = \; C_S(S_t,\, t)$$

then the dS_t term would vanish in equation (7), rendering the the return of the portfolio non-volatile over the period of time $(t,\, t + dt)$ (to leading order in dt). I claim that the rate of return of this portfolio should be exactly equal to the short-term interest rate. Indeed, if this were not so, there would be an opportunity for making money at no risk - an **arbitrage opportunity**. In fact, by investing in this portfolio (or its "mirror image"), an investor would effectively be able to borrow money at the cheapest rate and lend it out at the more expensive rate. Professionals would take advantage of this and, through the forces of supply and demand, would eventually drive the option price to a level where the return on the option-stock portfolio would be equal to the interest rate for cash.[2] Since the value of the option-stock portfolio at time t is $-C(S_t,\, t) \; + \; \Delta\, S_t \; = \; -C(S_t,\, t) \; + \; S_t\, C_S(S_t,\, t)$, the absence of arbitrage implies that

$$C_t \; + \; \frac{\sigma^2}{2}\, S^2\, C_{SS} \; = \; r\, (C \; - \; S\, C_S \;)$$

or

$$C_t \; + \; \frac{\sigma^2}{2}\, S^2\, C_{SS} \; + \; r\, S\, C_S \; - r\, C \; = \; 0 \; . \tag{8}$$

This is the **Black-Scholes partial differential equation**. To determine the function $C(S,\, T)$ we must specify boundary conditions. In the case of a call with expiration date T we have

$$C(S,\, T) \; = \; (S \; - \; K\,)^+ \; = \; max\, (S \; - \; K \;\; 0) \;\; , \tag{9}$$

where K is the strike price (X^+ represents the positive part of X). Indeed, if $S_T \; \leq \; K$, the option is worthless, and if $S_T > K$, the holder of the call can buy the underlying asset for K dollars and sell it at market price, making a profit of $S_T \; - \; K$.

For a European-style call (which can be exercised only at the date T), $C(S,\, t)$ is determined by solving the Cauchy problem for the Back-Scholes PDE with final condition (8). The explicit solution is known as the **Black-Scholes formula**:

[2] An arbitrage is defined as a transaction in which one buys an asset and immediately sells it realizing a riskless profit. One of the consequences of an (ideal) equilibrium economy is the absence of arbitrage opportunities, or , at least, of "obvious" arbitrage opportunities.

$$C(S,\, t;\, K,\, T;\, r,\, \sigma\,) \;=\; S\, N(d_1) \;-\; K\, e^{-r(T-t)}\, N(d_2)\ , \tag{10}$$

where

$$d_1 \;=\; \frac{1}{\sigma\,\sqrt{T-t}}\, ln\left(\frac{S}{K\, e^{r(T-t)}}\right) \;+\; \frac{1}{2}\,\sigma\,\sqrt{T-t}\ \ ; \qquad d_2 \;=\; d_1 \;-\; \sigma\,\sqrt{T-t}\ ,$$

and

$$N(z) \;=\; \frac{1}{\sqrt{2\,\pi}} \int\limits_{-\infty}^{z} e^{-\frac{y^2}{2}}\, dy\ .$$

A similar formula exists for European puts, replacing the final condition by $(K \;-\; S)^+$.

American-style options, which allow for exercise anytime before the maturity date, can be evaluated in a similar way. However, the possibility of **early exercise** makes the contract more valuable, in principle. The "early exercise premium", or survalue of American options with respect to Europeans, depends on the income stream that can be derived by holding the underlying asset (e.g. stock dividends) and cash (interest rate) instead of waiting to exercise the option. The rule of thumb here is that *there is an early exercise premium whenever the asset that is "bought" has non-zero income stream.* So far, we have not assumed that the stock pays dividends. Therefore, call options o such assets would have no exercise premium. Puts do have early exercise premium since selling the stock is like "buying cash" and the interest rate (income for holding cash) is non-zero. The income earned by exercising the put option and investing the proceeds in cash may exceed the "speculative" value of holding a put expecting the underlying price to drop further. Therein lies the value of early exercise for puts.[3]

In the case of **stock index** options such as OEX options traded at the Chicago Mercantile Exchange, we must take into account the fact that the stocks composing the Standard & Poor 500 Index pay dividends (currently slightly below 2% on an annualized basis); therefore both puts and call have eary exercise premium. It is well-known the value of an American S&P500 call (OEX contract of the CBOE) exceeds the value of the corresponding European contract (SPX).

To value American options, the idea is that we should look for a function $C(S,\, t)$ that satisfies the Black-Scholes equation in regions of the $(S,\, t)$-plane where the option should not be exercised and provide additional boundary conditions along the region corresponding to price levels where the option should be exercised. One way to arrive at this region is to impose the additional conditions on option prices that should hold in the case of American-style options:

[3]The Black-Scholes equation for assets that pay "continuous" dividends at a rate d, is analogous to (8). The only difference is that the drift term $r\,S\,C_S$ is replaced by $(r-d)\,S\,C_S$.

$$C(S,\,t)\ \geq (S-K)^{+}\ \ \text{(calls)}\quad ,P(S,\,t)\ \geq\ (K-S)^{+}\ \ \text{(puts)}\ , \tag{11}$$

since the option is worth as least as much as what you would get by exercising it immediately. These constraints give rise to an **obstacle problem**, or differential inequality, for the Black-Scholes equation which can be solved numerically. The free boundary arising in this problem corresponds to the boundary of the "optimal exercise" region for the holder – the option should be exercised whenever equality holds in (11).

The free-boundary conditions for American options at the boundary of the exercise region are

$$C(S_t,\,t)\ =\ S_t\ -\ K\quad ,\quad \frac{\partial C(S_t,\,t)}{\partial t}\ =\ +1\ .\ \ \text{(calls)}$$

$$P(S_t,\,t)\ =\ K\ -\ S_t\quad ,\quad \frac{\partial P(S_t,\,t)}{\partial t}\ =\ -1\ ,\quad \text{(puts)}\ . \tag{12}$$

These are classical first-order contact free-boundary conditions for obstacle problems encountered in PDE texts.[4]

The story only begins here. The Black-Scholes PDE and its variants are used to value more general contracts which have payoffs depending on the value of another traded asset. Such contracts are generically called **contingent claims** in the Mathematical Finance literature. They include the class of **exotic options**, traded in the inter-bank market, which can have practically any conceivable payoff structure. The end-users of exotic options are banks, corporations and sophisticated investors that with to invest in "tailor-made" derivatives with special cash-flow structures especially suited to their investment needs. "Plain vanilla" stock options, on the other hand, are mostly traded in exchanges such as CBOE, CME, AMEX, etc., and over-the-counter as well.

3. Risk-Neutral Probabilities Suppose that a contract gives rise to a series of cash-flows at different dates $T_1\ <\ T_2\ <\ ...\ <\ T_N$ and that these cash-flows are represented by functions $F_i(S_{T_i}),\ i=1,2...,N$. An argument similar to the one of the previous paragraph shows that the value of the corresponding contingent claim is $V_t\ =\ V(S_t,\,t)$, where $V(S,\,t)$ satisfies the PDE

$$\frac{\partial V}{\partial t}\ +\ \frac{1}{2}\sigma^2\,S^2\,\frac{\partial^2 V}{\partial S^2}\ +\ r\,S\,\frac{\partial V}{\partial S}\ -\ r\,V\ =$$

$$-\sum_{i:\,t<\,T_i} F_i(S)\,\delta(t-T_i)\ , \tag{13}$$

[4] These boundary conditions can be derived also from purely financial considerations (Avellaneda, NYU course notes).

with $V(S,\, T_N + 0) \;=\; 0$. This formula also applies to option portfolios, i.e. to bundles of options held (long or short) by an investor.

The Black-Scholes PDE has a fundamental probabilistic interpretation. The correspondence between PDEs and probabilities via the Fokker-Plank formalism yields

$$V(S_t, t) \;=\; \mathbf{E}\left\{ \sum_{i:\, t< T_i} e^{-t(T_i - t)}\, F(S_{T_i}) \;\middle|\, I_t \right\} \;, \tag{14}$$

where S_t is the diffusion process governed by the stochastic differential equation

$$\frac{dS_t}{S_t} \;=\; \sigma\, dZ_t \;+\; r\, dt \tag{15}$$

and $\mathbf{E}\left\{ \bullet \mid I_t \right\}$ represents the conditional expectation with respect to the σ-algebra generated by $\{Z_s,\, s \le t\}$.

If you compare equation (15) with the proposed statistics for the returns of the index (3), you will notice that the mean returns (drifts) are different. Indeed, in the latter equation, the mean return over the period $(t,\, t+\Delta t)$ is $r\,\Delta t$ instead of $\mu\,\Delta t$ (the "subjective" annual return on investment for the stock). Thus, the Black-Scholes theory tells us that the value of an option (or a more general contingent claim) is equal to the **expected future cash-flows**, calculated under a certain probability measure assigned to the future paths for the price of the underlying asset. This measure, however, is not the "subjective probability" that we started with!

Let us give a concrete example. Suppose that the interest rate is 5%, that the volatility of a stock XYZ is 16% and that it is expected to appreciate in price by 40% annually. The Black-Scholes value of a European-style option to buy this stock at today's price in 180 days is 5.75% of the price of the stock. This result is totally independent of the rate of return on the price of the stock (we assume that no dividends are paid out). On the other hand, the expected value of the cash-flows, $max\,(S_T \;-\; K)^+$, using 40% returns (equation (3)) is a whopping 21.46%. From an investor's point of view, this result may seem paradoxical, since the higher the expected returns, the higher the probability of profiting from holding the call. Therefore, he should be willing to pay more than 5.75% for the option. This argument, based on "frequential probabilities", is nevertheless wrong.

The explanation lies in the concept of dynamic hedging. Under the Black-Scholes assumptions, the holder of $\Delta_t \;=\; C_S(S_t,\, t)$ shares of stock and $C(S_t,\, t) \;-\; C_S(S_t,\, t) \cdot S_t$ dollars in a money-market account has a portfolio worth $C(S_t,\, t)$ dollars at time t. In the period $(t,\, t + dt)$, the change in the value of the shares plus the interest accrued in the cash account add up to the change in the function $C(S_t,\, t)$, because

$$dC(S_t,\, t) \;=\; C_S(S_t,\, t)\, dS_t \;+\; \left(C_t(S_t,\, t) \;+\; \frac{\sigma^2\, S^2}{2}\, C_{SS}(S_t,\, t) \right) dt$$

$$= \Delta_t \, dS_t \; + \; r \, (\, C(S_t, \, t) \; - \; C_S(S_t, \, t) \cdot S_t \,) \; dt \; ,$$

by virtue of the Black-Scholes PDE. Therefore, by successively adjusting the number of shares, Delta after each trading period, it is possible to maintain a net portfolio value equal to $C(S_t, \, t)$. At the expiration date, the value of the portfolio is exactly equal to $C(S_T, \, T) \; = \; (S_T \; - \; K \,)^+$ which is the market value of the option.

The conclusion is that if you have an initial reserve at time t of $C(S_t, \, t)$ dollars, you can implement a dynamic trading strategy that generates a return identical to the one of the option at the expiration date. This strategy is called **option replication**, or **dynamic hedging**.

Referring back to the example, the point is that an individual that can engage in dynamic hedging, does not care if investing 21.46% of price of the underlying stock in the call will make him break even "in the long run". The option can be "manufactured" with only 5.75%, so why pay more? Conversely, if the expected return on the stock was less than 5% and the option was priced as a "statistical bet" instead of using Black-Scholes, the investor would probably buy but definitely not sell at that price! The probability measure on price movements implied by (15) is called a **risk-neutral probability** because it has the property that the option premium corresponds to the value of a "statistical bet" (expectation) under this modified probability. is

The Black-Scholes formula represents the *cost of of replicating the option*, rather than the expected value of the payoff under a subjective probability (as in a game of chance). It is a consequence of the absence of arbitrage opportunities: if the assumptions of the model are correct (volatility, interest rate) then if the market traded at another price, this would give rise to a profit at no risk. One would simply buy or sell the option and offset the risk by dynamic hedging.

An important caveat at this point is the ability of investors to engage in dynamic hedging in practice, which involves actively trading in the underlying security over the lifetime of the option. As some readers might know, the possibility of dynamic hedging is only available to professional option dealers, due to the large costs of execution, transaction costs, etc. However, option dealers, who compete for customers in the marketplace, estimate the cost of managing an option inventory using Black-Scholes and make prices accordingly. Since prices must be competitive with other dealers, they reflect the business costs of dynamic hedgers (risk-neutral valuation) rather than the expected value of the option payoff under subjective probabilities.

Let me also mention here that the use of risk-neutral probabilities for pricing contingent claims goes far beyond Black-Scholes theory. Harrison and Kreps (1979) formulated a general theory of **no-arbitrage pricing**, often called **Arbitrage Pricing Theory** (APT) which can be summarized as follows:

Suppose that a market has no arbitrage opportunities. Then, there exists a (risk-neutral) probability measure P defined on the paths of prices of traded assets such that

$$P_t \;=\; \mathbf{E}^P \left\{ e^{-\int\limits_t^T r_s\, ds} P_T \;+\; \sum_{t\,<\,T_i\,<\,T} e^{-\int\limits_i^T -t\, r_s\, ds} C_{T_i} \;\middle|\; I_t \right\} \,. \tag{15}$$

Here, $P_\bullet$ represents the price of a traded asset, C_{T_i} represent cash-flows which are paid to the holder of the asset at different dates and r_s is a short-term interest rate. The measure P is called a **martingale measure** in the Mathematical Finance literature.

The significance of APT is that, *some* probability measure with the property (15), i.e a risk-neural probability, *must* exist in the absence of arbitrage. In other words, if such a measure would not exist, we would be able to find, in principle, a dynamic trading strategy that generates profits with no risk. Equation (15) can also be viewed as a formula for evaluating the price of securities, and it applies to all classes of financial assets. A complete discussion of the implications of APT is beyond the scope of this lecture. I recommend, for instance Duffie (1992) for an in-depth discussion of this subject.

3. Risk-management using the "Greeks". This section discusses practical uses of the Black-Scholes formula or, more generally, the Black-Scholes PDE, as a tool for for hedging an options portfolio. For simplicity, we consider a "pure" situation of a stock which pays no dividends between now and the expiration dates of the options.

Recall that the parameters that enter the Black-Scholes formula are (i) the exercise price, or strike price, K, (ii) the expiration date, T, (iii) the price of the underlying asset, S, (iv) the interest rate, r, and (v) the volatility, σ. Of these five parameters, the first four are observable at any given time (r is known for short expiration dates). In contrast, the volatility of the underlying asset is not directly observable. For each value of the volatility parameter we obtain a different theoretical option value. Conversely, it is easy to show that to each possible option value (in the range of the Black-Scholes formula) corresponds a unique volatility parameter. This is a consequence of the fact that the Black-Scholes option premium is a strictly increasing function of σ . The **implied volatility** of a traded call is, by definition, the value of σ that solves the equation

$$C\,(S,\, t;\, K,\, T;\, r,\, \sigma\,) \;=\; \text{market price of the call}\ ,$$

where the left-hand side represents the BS theoretical value, with the same definition applying to puts.

The market price of the option defines its implied volatility, which is the volatility that "makes Black-Scholes true". We will assume henceforth that σ is the implied volatility. The main use of the Black-Scholes formula is to control the exposure of an options portfolio to different types of market risk. We already introduced the parameter Delta, which is the derivative of the portfolio value with respect to changes in S:

$$\Delta \;=\; \frac{\partial V(S,\,t)}{\partial S}\;.$$

Here, V represents the total value of the portfolio, i.e. the sum of the values of all the options in the portfolio). The portfolio Delta is the algebraic sum of the Deltas corresponding to each option computed with its own implied volatility. The fundamental result of BS is that a position in $-\Delta$ shares of the underlying asset renders the position "market-neutral" — the value of the portfolio will vary less than the stock price by an order of magnitude (dt rather than $dt^{1/2}$). Moreover, the Black-Scholes theory implies that the dt term is equal to the cost of funding the position at the riskless rate. (If funding is taken into account the variation of the total portfolio value is therefore $O(t^{3/2})$ - negligible even after adding all the variations over the life of the option).

For example, a 180-day European at-the-money call ($K = S$) with $r = 5\%$ and $\sigma = 16\%$ has a delta of 0.6086. This means that to hedge a short position in 100 calls (or, equivalently, one option to buy 100 shares), one should buy 60.86 shares of the underlying stock (say, 60 shares). Similarly, the holder of an option to buy 100 shares will be hedged by short-selling 60 shares. The same philosophy applies to option portfolios. At least in theory, Delta-hedging is a way of protecting an option portfolio against moves in the price of the underlying asset.

An important observation concerning **delta-hedged** option positions is that the distinction between owning a call and owning a put disappears completely! Indeed,the identity

$$(S\;-\;K)^{+}\;-\;(K\;-\;S)^{+}\;=\;S\;-\;K$$

implies that

$$C(S,t;K,T)\;-\;P(S,t;K,T)\;=\;S\;-\;K\,e^{-r(T-t)}\;,$$

where the functions in the left-hand side of the equation represent the value of a call and a put with same strike and expiration date. The right-hand side can be interpreted as the fair value of a contract to buy the stock at K dollars at date T. This result is called **put-call parity**. It states that a put can be "synthesized" with a portfolio consisting of a call, one share held short, and a note with face value K. Differentiating this equation, we obtain obtain a relation between the Deltas of puts and calls:

$$\Delta_{call}\;-\;\Delta_{put}\;=\;1\;.$$

The point is that , while unhedged option positions represent leveraged bets on the direction of the market, hedged positions are "non-directional". Delta-hedged option positions stand to profit or lose according to the behavior of volatility or interest rates.

The second-derivative of the BS price with respect to the spot price

$$\Gamma = \frac{\partial^2 V(S, t)}{\partial S^2},$$

is an important sensitivity for practical Delta-hedging. Gamma measures the change in Delta per change in the value of the underlying asset. This sensitivity does not enter the derivation of the Black-Scholes theory, which assumes continuous adjustments in Delta. However, in real life, adjustments to the hedge portfolio are done at discrete dates usually in response to changes in the price of the underlying asset). Discrete hedging gives rise to "hedge slippage": the option is no longer perfectly replicated an there are gains and losses at each time (relative to the Black-Scholes) fair value. The Gamma of an option (put or call) is given explicitly by

$$\Gamma = \frac{1}{S\sqrt{2\,\pi\,\sigma^2}}\, e^{-\frac{d_1^2}{2}}.$$

It is a positive quantity: the Delta of an option increases as the stock price increases. The graph of a hedged (long) option as a function of S is convex and has a minimum at the current spot price. It is easy to see that the graph at time $t + dt$ is also convex but it is uniformly lower in value for all S. This means that the holder of a hedged long option position will make money if $| S_{t+dt} - S_t |$ is large and will lose if the change in value of the stock is nearly unchanged. If we consider the reverse position, i.e., a short option hedged with the Black-Scholes delta, we see that the graph is now concave. Small stock moves give rise to a profit and large stock moves give rise to a loss, due to the negative convexity. Market professionals refer to the former situation as being "long Gamma" and the latter as being "short Gamma". This concept, like Delta, applies at the portfolio level. The Gamma of the portfolio is just the net Gamma obtained by adding all the option Gammas weighted by their sign (long/short) and by the number of contracts. Needless to say, positions with large Gammas are risky because they are difficult to hedge without introducing too much error with respect to theory. The size of Gamma, measures, in a sense, the risk exposure to missing the hedge by trading discretely. Managing Gamma is achieved by buying or selling options so as to to keep this sensitivity under prescribed limits, therefore limiting the risk of "hedge-slippage".

Gamma also becomes important when we take into account **transaction costs** incurred by hedging dynamically. For instance, traders that are short Gamma have to buy stock when its price goes up and sell when the price drops (always a painful proposition). On the other hand, hedging a long-Gamma position involves selling stock as the market rallies and buying when it drops. This creates a noticeable asymmetry in the expected cost of replicating long and short Gamma positions. In general, dynamic hedging in the presence of transaction costs is expensive and erodes traders' profits. A "compromise" must therefore be made between avoiding transaction costs and limiting the contract's risk-exposure. This subject leads to very interesting mathematics and is the object of numerous studies (Leland(1985), Davis, Panas and Zariphopolou (1993), Avellaneda and Parás (1994) , Taleb(1997)).

Another important sensitivity of the BS equation is

$$\Theta = \frac{\partial V(S,\,t)}{\partial t} \ ,$$

known as the **time-decay**. It is of crucial importance in the risk-management of an options portfolio, because it tells the trader by how much the value of the position will change if the spot price stays the same. Theta and Gamma have opposite signs, as we argued above. We can see this more precisely by making the change of variables

$$\tilde{V} \ = \ e^{-r\,t}\,V \quad , \quad \tilde{S} \ = \ e^{-r\,t}\,S \ ,$$

which expresses values in constant dollars, removing the effect of interest rates. In terms of these variables, the Black-Scholes PDE becomes

$$\frac{\partial \tilde{V}}{\partial t} \ + \ \frac{1}{2}\,\sigma^2\,\tilde{S}^2\,\frac{\partial^2 \tilde{V}}{\partial \tilde{S}^2} \ = \ 0 \ ,$$

or

$$\tilde{\Theta} \ = \ -\,\frac{1}{2}\,\sigma^2\,\tilde{S}^2\,\tilde{\Gamma} \ ,$$

where we use tildes to emphasize that values are computed in constant dollars. This relation can be rewritten as

$$\tilde{V}(t+dt) \ - \ \tilde{V}(t) \ \propto \ -\,\frac{1}{2}\,\tilde{\Gamma}\,\cdot\mathbf{E}\left\{\left(\tilde{S}(t+dt) \ - \ \tilde{S}(t)\right)^2\right\} \ ,$$

where we identified the variance with the square of the spot price increment. Thus, *the BS equation expresses a proportionality relation between time-decay, convexity and volatility.*

Let us illustrate this with a numerical example. The change in value a 180-day at-the-money call with $\sigma \ = \ 16\%$ over one day expressed in constant dollars is approximately $\Delta V \ = \ 0.01\%$ of the value of the underlying asset. This is what the holder of the option will lose over one day if the price of the underlying remains unchanged. On the other hand, the change in Delta for a 1% move in spot is 0.0357. A volatility of 16% represents approximately a $(\Delta\tilde{S})^2$ of approximately $16 \times 16/365 \ = \ 0.71$ in percentage terms. Multiplying this by 0.5×0.0357 yields approximately 0.012%, as claimed. This is what the holder of the option expects to make if the spot price changes by one standard deviation.

As one might expect, the sensitivity of the BS value with respect to the volatility parameter plays a crucial role. This sensitivity is known as **Vega**:[5]

$$Vega \;=\; \frac{\partial V}{\partial \sigma}\;.$$

Traders often use other higher-order sensitivities of the Black-Scholes formula to analyze the risk of an option portfolio, such as

$$\frac{\partial^2 V}{\partial S \partial \sigma} \;=\; \frac{\partial \Delta}{\partial \sigma} \;=\; \frac{\partial\,(Vega)}{\partial S}\;,$$

and

$$\frac{\partial^2 V}{\partial \sigma^2} \;=\; \frac{\partial\,(Vega)}{\partial \sigma}\;.$$

The intuition behind these higher-order sensitivities is probably mastered only by the most seasoned option professionals. For example, $\frac{\partial \Delta}{\partial \sigma}$ represents the change in the Delta, the amount of shares of the underlying asset to be held in the hedge, under a change in the options implied volatility with all the other parameters held constant. Another parameter often used to manage option positions is the sensitivity with respect to the short-term interest rate,

$$\rho \;=\; \frac{\partial V}{\partial r}\;.$$

Although, the effect of ρ is less important in general, considerations about changes in funding rates can be important in countries where interest rates are very high and/or fluctuate considerably over short periods of time.

By and large, option portfolio risk-management consists in neutralizing, or at least keeping with reasonable limits, the above-mentioned sensitivities in a portfolio. Thus, an options dealer will manage his Delta (usually set to zero), as well as Gamma, Theta, Vega and Rho. Except for Delta, which is managed by trading in the underlying asset, risk-management of other "Greeks" involves buying and selling options so as to keep the net sensitivities near target levels. This includes taking positions in these higher-order sensitivities as well. For example, a trader can gain exposure to a rise in implied volatility, for example, by having a positive-Vega portfolio for instance, but at the same time be neutral in Delta, Gamma and Rho. An obvious but important consideration is the opposite

[5]The derivative of the stock price with respect to σ is now sometimes referred to as κ (Kappa), with the advantage that the latter is a Greek letter. Vega appears to be the terminology used to by option risk-managers in the "early days" and is currenlty widely used. The name Vega, on the other hand, corresponds to the name of a Chevrolet model sold in this country in the 1970's, a somewhat amusing/nostalgic coincidence.

sign of Theta and Gamma – the "buyer of convexity" gets exposed naturally to time-decay, while time is "on the side" of the seller of convexity if the market remains quiet.

The management of volatility risk (and, in some cases, interest rate risk) of an options book is a highly non-trivial matter. Among other things, the implied volatilities of options with different maturities and strikes are generally not equal. Managing volatility risk requires therefore monitoring the joint movement of a "matrix" of implied volatilities

$$\sigma(K_i,\,T_j)\ ,\quad i\,j\,=\,1,\,2\,...$$

where K_i and T_i represent different strikes and expiration dates. This implies frequent adjustments of the option portfolio to keep the Greek sensitivities in line and to design profitable option positions according to traders' expectations on the future behavior of volatilities and rates. In other words, professional option traders must evaluate the risk-return characteristics of **option spreads**, or different option combinations, in terms of volatility and interest-rate forecasts, price ranges, dividend pay-outs, liquidity of different contracts, etc. Natanberg (1988) and Taleb(1977) contain lucid analyses of option strategies from a practical point of view.

4. Uncertain volatility models. From a fundamental point of view, the assumption that volatility is constant and Vega-hedging as a risk-management practice, have several shortcomings which are more and more recognized. Specifically:

- The Black-Scholes model, (15), assumes implicitly that variations or the spot price are *homogeneous*, i.e. that price returns have the same statistics at different dates;
- Vega gives the sensitivity of an option only for *small changes* in the implied volatility;
- Vega-hedging is inconsistent with the fact that options with the same maturity and different strikes usually trade at different implied volatilities.

In mathematical parlance, one would say that Greek hedging corresponds to a "linearized" approach to managing risk. In particular, it is not expected to work in the event of a large move, or crash, or to a regime shift, in which market conditions change dramatically after some event.

The constant volatility assumption is actually inconsistent with APT. In fact, if two options with the same expiration date trade at different implied volatilities, which is often the case, then the "spot volatility" of the underlying asset

$$\frac{dS_t}{S_t}\ =\ \sigma_t\,dZ_t\ +\ r\,dt$$

cannot be constant, or even a deterministic function of time. If σ_t was constant or $\sigma_t = \sigma_0(t)$ where $\sigma_0(\cdot)$ is deterministic, at least one of the two options would be mispriced by the market if we used the original Black-Scholes model. However, there is ample evidence that demonstrates that a difference in implied volatilities does not necessarily imply an arbitrage opportunity. Instead, it is currently believed that the implied volatility **smile** and **skew** observed in many markets, reflect traders' "inhomogeneous" volatility expectations. Traders estimate the cost of option replication conditionally on future events such as changes in market conditions, changes in the liquidity of the underlying asset, etc. There is no reason why the constant-volatility/lognormal assumptions of Black-Scholes should hold: only the *consequence* of the theory – the consequences of no-arbitrage – should hold, once a probabilistic model is specified.

Here, mathematics comes to the rescue in a big way. According to the No-Arbitrage Theorem of §2, no-arbitrage implies the existence of a probability measure on paths $\{S_t\}$ such that the prices of traded assets are *simultaneously* reproduced by taking expectations with respect to the same probability. The only restriction, dictated by APT, is that this probability should be such that that prices computed in constant dollars are martingales. By making volatility σ_t a random process, or a function of S_t and t, for example, we can expect to satisfy the APT equation with a single probability measure for $S_\bullet$ and thus obtain therefore a more accurate valuation of option positions.

This prompted researchers to look for specifications of the volatility process that would reconcile Black-Scholes theory with observed behavior of implied volatility of option markets. One proposal has been to make the spot volatility a stochastic process, which may be statistically correlated with price shocks. Hull and White (1987) showed, among other things, that a negative correlation between volatility shocks and price shocks (as the market drops volatility rises) reproduces qualitatively (but, in my own experience, not quantitatively) the "volatility skew" observed in equity options markets. Typically, in these markets, out-of-the-money puts have a higher implied volatility than out-of-the-money calls. Other proposals for volatility modeling (Engle (1984), Noh *et al* (1994)) use **conditionally heteroskedastic models** (the ARCH-GARCH family) to model the behavior of the underlying asset. However, the use of stochastic volatility models or ARCH models raises the important problem of model specification (by econometric analysis or otherwise) and the relevance of historical data for managing future risk. Another critique of stochastic volatility models, with which Keynes would probably not disagree, is that the differential sensitivities with respect to the parameters of these more complicated models may not provide protection against large moves in the market. Making the model more elaborate may reflect better current options prices quoted in the market but still missed the notion of risk, or uncertainty about the model itself.

In an attempt to remedy these shortcomings of parametric models, I present here a new approach for managing volatility risk. This approach is based on the premise that we have little knowledge of the spot volatility process and that it may be preferable to use the concept of "uncertainty" or *lack of information*, rather than an elaborate specification the statistics of the volatility.

Let us assume that we have determined a **confidence interval** for the spot volatility process $\{\sigma_t,\ 0 \le t \le T\}$, without going into details of how this interval, or "cone" was obtained. We postulate therefore that the process of conditional volatility for the price of the underlying asset satisfies the inequalities

$$\sigma_{min}(t) \ \leq \ \sigma_t \ \leq \ \sigma_{max}(t) \ , \tag{16}$$

where $0 \ < \ \sigma_{min} \ < \ \sigma_{max}$ are deterministic functions. We shall consider the collection of all probability specifications on the underlying price process that satisfy the volatility bounds. Given this range of uncertainty, we cannot provide a single price for any given contract whose value is sensitive to volatility. Instead there is continuum of possible prices. We shall focus on the extreme model values, or upper and lower bounds on prices, which correspond to the **worst-case scenario replication costs** for short and long positions, respectively.

In this setting, the problem of calculating extreme prices is isomorphic to a stochastic control problem in which the volatility is the the control variable. Extremal prices can be computed using the Bellman dynamic programming principle. More specifically, the partial differential equation for the upper bound has the form

$$\frac{\partial V}{\partial t} \ + \ \Phi_0 \left[\frac{S^2}{2} \frac{\partial^2 V}{\partial S^2} \right] \ + \ r\,S\,\frac{\partial V}{\partial S} \ - \ r\,V \ = \ - \sum_{i:\, t\, <\, T_i} F_i(S)\,\delta(t \ - \ T_i) \ , \tag{17}$$

where

$$\Phi_0\,[X] \ = \ \begin{cases} \sigma^2_{min}\,X \quad , \quad X \ < \ 0 \ , \\ \\ \sigma^2_{max}\,X \quad , \quad X \ \geq \ 0 \ . \end{cases} \tag{18}$$

This equation constitutes a simple but important modification of the Black-Scholes PDE; it reduces to the latter when there is no uncertainty ($\sigma_{min} \ = \ \sigma_{max}$). The PDE (17) is known as the Uncertain Volatility Model (UVM) (Avellaneda, Levy and Paras, 1995). The Delta of the UVM equation can be used to immunize the portfolio against market risk, in the following sense: if the agent uses this Delta and starts with the reserves calculated with the UVM equation he or she will break even if the worst-case volatility scenario is realized and will otherwise make a profit (use less reserves than budgeted). This idea is known as a **dominating strategy** in Mathematical Finance. If the agent budgets initially less reserves that the UVM value, the strategy may lose money but losses are limited to the difference between the BS and UVM premia. It is intuitively clear that this method of valuation is robust with respect to volatility sepcification. However, this protection does not come for free, since UVM requires more reserves than Black-Scholes using, say, the center of the band as volatility parameter. In particular, UVM can generate option prices that are not competitive with those of a dealers using a constant volatility inside the band.

The difference between upper and lower extremal prices, or between the extremes and Black Scholes using the center of the band (16), is due to the fact that the new valuation provides protection against *all volatility paths in the band.* This gap in prices reflects our uncertainty about future volatility paths. One way to narrow this "uncertainty gap " is is to consider options as a partial alternative to dynamic hedging. An extension of UVM, which I describe now, provides a rational approach for doing this (Avellaneda and Parás, 1996). The method consists in using the UVM equation to price "packages" formed by the contingent claim of interest combined with traded options. Thus, the idea is to perform the UVM analysis on the new package package rather than on the original claim. Since the options can be bought/sold in the market we can concentrate on delta-hedging only the net exposure. Usually, this proposal gives rise to narrower price bands while, at the same time, hedges that perform well under a broad range of volatility scenarios.

Let me show how this works. Assume that there are M traded options in the market, with payoffs $G_j,\ j = 1, 2, \ldots M$ and expiration dates $\tau_1, \tau_2, \ldots \tau_M$. Assume also that each option trades at a price C_j respectively. If an agent sells the derivative security represented by the cash-flows $\{F_j\}$ and buys a portfolio of λ_1 contracts of the first option, λ_2 contracts of the second, and so forth, the amount of cash needed to hedge this position under the worst-case scenario is

$$\sup_P \mathbf{E} \left\{ \sum_{i=1}^{N} e^{-r\,T_i}\, F(S_{T_i}) \ - \ \sum_{j=1}^{M} \lambda_j\, e^{-r\,\tau_j}\, G(S_{\tau_j}) \right\} \ ,$$

where the sup is taken over all probabilities with volatility paths in the band (16). If we add to this the market price of the options portfolio, we obtain

$$V(\lambda_1, \ldots \lambda_M) \ = \ \sup_P \mathbf{E} \left\{ \sum_{i=1}^{N} e^{-r\,T_i}\, F(S_{T_i}) \ - \ \sum_{j=1}^{M} \lambda_j\, e^{-r\,\tau_j}\, G(S_{\tau_j}) \right\} \ + \ \sum_{j=1}^{M} \lambda_j\, C_j \ . \tag{19}$$

$V(\lambda_1, \ldots \lambda_M)$ can be interpreted as the worst-case scenario reserves necessary to hedge the contingent claim using a partial "static" options hedge and Delta-hedging the residual. If we minimize $V(\lambda_1, \ldots \lambda_M)$ over all combinations of the variables, we will have found the cheapest hedge using options that immunizes the portfolio over all possible volatility paths inside the band. This hedge will be represented by a vector of **option hedge-ratios** $(\lambda_1^*, \ldots \lambda_M^*)$. It is noteworthy that the function $V(\lambda_1, \ldots \lambda_M)$ is strictly convex, so that if a minimum exists, it must be unique. Moreover, the probability measure P^* which gives the worst-case scenario with these hedge-ratios satisfies the market price conditions

$$C_j \ = \ \mathbf{E} \left\{ e^{-r\,\tau_j}\, G(S_{\tau_j}) \right\} \ , \quad j = 1, 2, \ldots M. \tag{20}$$

To see this, it suffices to use the first-order conditions for the minimum in (19) taking into account that the partial derivatives of V satisfy

$$\frac{\partial V(\lambda_1, ...\lambda_M)}{\lambda_j} = C_j - \mathbf{E}^{P^*} \left\{ e^{-r\,\tau_j} \, G(S_{\tau_j}) \right\} , \quad j = 1, 2, ... M , \tag{21}$$

where P^* is the measure that realizes the supremum. In particular, the new worst-case scenario probability produces a **calibrated model**, matching the prices of all the traded options.

Let us illustrate the theory with an example. Suppose that the stock price is $S_0 = \$\,100$, and that $r = 5\%$. Suppose that a 180-day option on this stock with strike $K_1 = 100$ is trading with an implied volatility of 16%. Moreover, you expect the" spot" volatility σ_t to vary between the limits $\sigma_{min} = 8\%$ and $\sigma_{max} = 24\%$ How could you use this information to price and hedge an option with a strike $K_2 = 115$ expiring in 160 days?

In this problem, we have $M = N = 1$. The market price of the at-the-money option with $\sigma = 0.16$ is \$5.75. Optimizing the function $V(\lambda_1)$ with $F_1(S) = (S - 115)^+$, $G_1(S) = (S - 100)^+$ and $C_1 = 5.75$ yields, for the short position,

$$\lambda_1 = 0.302 , \quad V(\lambda_1) = 1.97 , \quad \Delta = 0.06 .$$

This means that optimal hedge consists in buying 0.302 at-the-money options for each 115-option sold. The UVM Delta hedge for the residual portfolio is 0.06, i.e., 6% of the notional amount of shares in the contract. The quantity $V(\lambda_1) = 1.97$ represents the cost of this hedge, which is broken down as follows: $5.75 \times 0.302 = \$1.73$ invested in the option hedge and $1.97 - 1.73 = \$0.24$ invested in reserves for dynamic hedging. We can compare this with the Black-Scholes premium with different values of σ. For $\sigma = 0.16$ we have a BS value of \$ 0.75, which is much less than $V(\lambda_1) = 1.97$. On the other hand , if your worst-case fears materialize and the term volatility turned out to be 24% instead of 16%, the correct premium should have been \$2.24! This is greater than the optimized value, 1.97. The price could be narrowed further if there were more options to add to the mix!

This hedging technique, called **Lagrangian Uncertain Volatility Model** ($\lambda{-}UVM$), can be used to systematically construct option hedges for exotic options and other non-standard derivatives which are more reliable with respect to changes in the volatility environment. They provide an alternative to hedging volatility risk using Vega and other σ-related Greeks. From a theoretical point of view, we see how the introduction of multiple martingale measures (through a band of volatility paths) can be used to eliminate volatility risk by selecting a package of options that gives rise to a less risky net position. In my own opinion, this tradeoff between delta-hedging and buying option protection up-front is the key for managing volatility risk.

5. Relative entropy: combining volatility uncertainty with a-priori beliefs.
In this section, I sketch a generalization $\lambda - UVM$ which attempts to narrow the lack of information inherent in assumtion (16) (Avellaneda, Friedman, Holmes, Samperi, 1997).

Let us assume now that we have a definite belief about the behavior of forward volatility, based upon historical analysis, expectations about future option prices, etc. This belief is expressed mathematically as a choice for the "most likely" risk-neutral probability measure, P_0, that will be realized, *e.g.*,

$$\frac{dS_t}{S_t} \;=\; \sigma_0(S_t,\, t)\, dZ_t \;+\; r\, dt \;, \tag{22}$$

where $\sigma_0(S,\, t)$ is a given function of $(S,\, t)$. We now wish to price and hedge contingent claims taking into account also volatility risk – the possibility that the realized volatility will be differ from our a-priori belief σ_0. A natural approach is to compute a worst case-scenario probability measure, as before, introducing a penalization for probabilities which are "far away" from the prior measure (22) in some norm or distance and which we consider therefore unlikely to occur.

It turns out that a convenient notion of distance for this purpose is the **Kullback-Leibler relative entropy distance**, given by

$$\mathcal{E}\,(P;\, P_0) \;:=\; \mathbf{E}^P \;\left\{\; log\left(\frac{dP}{dP_0}\right) \;\right\} \;, \tag{23}$$

where $\frac{dP}{dP_0}$ is the Radon-Nykodym derivative of P with respect to P_0. The KL relative entropy is a tool used in computer science to analyze the complexity of codes and in other fields of science as well, usually in connection with inverse problems and parameter estimation (Jaynes (1996, 1984)). It can be interpreted as a generalized likelihood ratio. To proceed along these lines, we must compute the relative entropy distance between P_0 and a generic probability measure (P) for S_t of the form

$$\frac{dS_t}{S_t} \;=\; \sigma(t)\, dZ_t \;+\; r\, dt \;. \tag{24}$$

As it turns out, the relative entropy of Ito processes with different volatilities is infinite. Nevertheless, this divergence can be removed by passing to a discrete lattice model and analyzing the asymptotic behavior of (23) as the mesh-size tends to zero. This yields the result

$$\mathcal{E}\,(P;\, P_0) \;\approx\; \frac{\text{const.}}{\Delta t} \;\times \mathbf{E}^P \left\{\; \int_0^T \eta\left(\frac{\sigma(t)}{\sigma_0(S_t,\, t)}\right)\, dt \;\right\} \;, \tag{25}$$

where ΔT represents the lattice time-time step and η is a convex function defined on non-negative real numbers. Equation (25) shows that the relative entropy of two measures defined on a lattice (trinomial tree) which approximate the Ito processes P and P_0 diverges at a rate inversely proportional to the lattice step, with a well-defined rate, given by the right-hand side of (25). The function η depends on the discretization scheme, unfortunately. However $\eta(Y)$ is, generically, a convex function of its argument, independently of the scheme used to approximate the diffusions. A possible choice for η, which is obtained from the standard trinomial lattice and is therefore useful for computations, is

$$\eta(Y) \;=\; Y\,log(Y) \;-\; Y \;+1\;. \tag{26}$$

This computation suggests that we consider the following variant of λ-UVM (compare with (20)):

$$V(\lambda_1,\,...\,\lambda_M) \;=$$

$$\sup_P \mathbf{E}^P \left\{ -\,\epsilon \int_0^T \eta\left(\frac{\sigma(t)}{\sigma_0(S_t,\,t)}\right) dt \;+\; \sum_{i=1}^{N} e^{-r\,T_i}\,F(S_{T_i}) \;-\; \sum_{j=1}^{M} \lambda_j\,e^{-r\,\tau_j}\,G(S_{\tau_j}) \right\}$$

$$+\;\sum_{j=1}^{M} \lambda_j\,C_j\;, \tag{27}$$

where the supremum is taken with respect to risk-netural measures such that σ_t lies in the band (16). In this equation, ϵ represents a numerical constant, which can be "tuned" in different ways. Large values of ϵ correspond to a large penalization for deviating from the a-priori measure P_0; small values of ϵ correspond to a weak effect of the relative entropy and hence to a larger volatility uncertainty.

For finite values of ϵ, minimization of (27) over all vectors $(\lambda_1,\,...\,\lambda_M)$ will give the best option hedge taking into account uncertain volatility and, at the same time, information about the "most likely risk-neutral measure" P_0. The minimization of the function in (27), which is convex in its arguments, is done using a non-linear PDE analogous to (17)-(18), but where the function Φ_0 is replaced by Φ_ϵ, given by

$$\Phi_\epsilon(X) \;=\; \cdot\sigma_0^2\cdot \sup_{\frac{\sigma_{min}}{\sigma_0}\leq Y\leq\frac{\sigma_{min}}{\sigma_0}} \left(X\,Y \;-\; \epsilon\cdot\eta(Y)\,\right)\;. \tag{28}$$

A precise correspondence between this theory and λ-UVM can be made if we take $\eta(Y)$ to be a function equal to zero for $\sigma_{min}/\sigma_0 \;\leq\; Y\;\sigma_{max}/\sigma_0$, and to $+\infty$ otherwise (a degenerate convex function). In this case, the entropy term will be finite if and only if σ is inside the volatility band. A smooth, strictly convex function η which is finite in the interval and infinite outside this interval gives rise to a smooth transition, as a function of *Gamma*, between the extreme volatility values σ_{min}, σ_{max} which is equal to σ_0 when the Gamma of the portfolio vanishes (in the absence of volatility risk). For details on this theory and a variety of numerical results, see Avellaneda, Friedman, Holmes and Samperi (1997).

6. Conclusion This paper sketched some of the main themes and results in the area of valuation of derivative securities. We have covered the principles underlying the valuation of options and contingent claims and also discussed the management of the risk of an options portfolio using differential sensitivities. This led us naturally to the original question of model specification and to managing the uncertainty as to which risk-neutral probability will be "realized" by the market in the future. We made a fundamental distinction between pricing contingent claims using a well-defined probabilistic model, and pricing in the presence of uncertainty with respect to model parameters. This uncertainty is ever-present in dervatives valuation, the argument being that if the pricing measure was known with certainty by all market participants, there would be ultimately no incentive for trading options. Option risk-management is beyond the scope of linear models. Managing the "Greeks" is the preferred technique for handling model uncertainty, but we have shown that this method can be inconsistent with APT and, more importantly, that it will not provide protection against large changes in market conditions (*i.e.* large changes in risk-neutral probabilites). Proposals such as $\lambda - UVM$, which are based on the simultaneous consideration of different pricing scenarios, seem to be natural alternative to differential hedging.

I conclude by stating that the focus of this paper has been limited to a single risky asset. In reality, we are usually confronted with an economy with multiple surces of risk and multiple underlying assets, such as the universe of fixed-income instruments with different tenors. In the multidimensional case the complexity of the theory increases considerably, and even the apparently simple problem of evaluating expected cash-flows (under a single probability measure) presents difficulties. The need for exploring uncertainty in a multi-asset economy, such as correlation uncertainty, makes the analysis of risk under uncertain probabilities even more relavant. I hope that some of the methods discussed here, or at least the ideas behind them, may prove useful for solving the challenging problems that exist in this field.

References

Avellaneda., M. and Parás, A., (1994), Pricing and hedging derivative securities in the presence of large transaction costs, *Applied Math. Finance*, 1,

Avellaneda, M., Levy, A. and Parś, A. (1995), Pricing and hedging derivative securities in markets with uncertain volatilities, *Applied Math. Finance*, vol 2, 73-88.

Avellaneda, M. and Parás A. (1996), Pricing and Hedging Derivative Securities in Markets with Uncert ain Volatilities: the Lagrangian Uncertain Voaltility Model, *Applied Mathematical Finance*, 3, 21-52.

Avellaneda, M., Friedman, C., Holmes, R. and Samperi, D. (1997), iCalibrating volatility surfaces via relative entropy minimization, *Applied Math Finance*, 3, March issue.

Black F. and Scholes, M. (1973), The pricing of options and corporate liabilities, *J. Pol. Econ.*, 81, 637-659.

Davis, M. H., Panas, V.G.,and Zariphopolou, T. (1993), European option pricing with transaction costs, *SIAM J. Control and Optimization*, 31, 470-493.

Duffie, D. (1992), *Dynamical Asset Pricing Theory*, Princeton Univ. Press

Engle, R., (1984), Wald likelihood ratio and Langrange multiplier tests in econometrics, it Handbook of Econometrics, vol II, Grilches and Intrilligator eds. North-Holland, Amsterdam; see also Noh, J. Engle, R. and Kane A. (1994), Forecasting volatility and options prices of the S&P 500 Index, *The Journal of Derivatives*, Fall issue.

Harrison, J.M., and Kreps, D. (1979), Martingales and arbitrage in multiperiod securities markets, *J. Econ. Theory*, 381-408.

Hull, J. (1994), *Options, futures and other derivative securities*, 2nd ed. , Prentice Hall, Englewod Hills, NJ.

Hull, J. and White, A., (1987), The pricing of options on assets with stochastic volatilities, *J. of Finance* June, 281-300.

Jaynes, E. T. (1996), *Probability theory: The logic of science.* Unpublished Manuscript, Washington University, St. Louis, MO., see also article in *Inverse Problems*, vol 14 D. McLaughlin, ed., American Math. Soc., Providence, Rhode Island, 1984.

Keynes, J. M. (1936), *The General Theory of Investment Employment and Money*, Hatrcourt Brace and Co. New York.

Leland, H. (1985), Option pricing and replication in the presence of transaction costs, *J. of Finance*, 40, 1283-1301

Markowitz, H. (1991), *Portfolio Selection*, J. Wiley and Sons, New York, 2nd edition.

Natanberg (1988), *Option volatility and pricing strategies*, Probus.

Taleb (1997), *Dynamic Hedging*, John Wiley and Sons, New York.