

TEACHING NOTE 00-04:
GIRSANOV'S THEOREM IN DERIVATIVE PRICING

Version date: July 18, 2008

C:\Classes\Teaching Notes\TN00-04.doc

The purpose of this teaching note is to present Girsanov's Theorem as it applies to derivative pricing and to provide some conjunction of the mathematical theory with the financial economic theory. Girsanov's Theorem appears to be a somewhat complex rule that, at best, provides limited use if one already understands derivative pricing. Indeed it is quite possible to obtain a solid understanding of derivative pricing without encountering Girsanov's Theorem. Nowhere did it appear in the literature on derivatives for many years, in spite of many exceptionally powerful advances developed by financial economists. Yet Girsanov's Theorem provides much of the mathematical rigor that underlies derivative pricing. Unfortunately Girsanov's Theorem provides virtually no intuition and is rarely, if ever, integrated with any form of economic discussion of what is really going on. In short, mathematicians approach the derivative pricing problem in their own ways, one being the application of Girsanov's Theorem, and financial economists approach it other ways. That these two groups arrive at the same answer is not surprising, but they speak different languages and travel different routes. Here we shall try to bring them together a little. The discussion is admittedly more mathematical than financial, but that is because the readership is likely to be more financial than mathematical and is more in need of seeing the mathematics and being reminded of when the mathematicians are saying something that a financial economist recognizes.

To price a derivative instrument such as an option, we can resort to any of several methods. One of the more formal mathematical approaches is the application of Girsanov's Theorem. In short, what we do is alter the probability distribution of the stock return such that it follows a stochastic process known as a martingale. A martingale, among other things, is a process without a drift.¹ In that case, we can obtain the value today of the derivative by determining the expected future value of the derivative, where the expectation is arrived at by using the altered probability distribution. In doing it this way, we avoid the problem of having to

¹In addition a martingale must be finite and each realization must be independent of the previous one.

solve a differential equation, though we can still arrive at the solution in this manner by solving the differential equation.

First, however, let us identify why we are doing this. One of the results of mathematical finance is that any two stochastic processes that are both martingales when the same probabilities are applied can be related to each other via a simple transformation. For example, given one martingale, X_t , and another Y_t , the relationship

$$Y_t = Y_0 + \int_0^t \lambda_j dX_j,$$

represents one martingale in terms of the other. Accordingly, this is referred to as the *martingale representation theorem*. It is important to note, however, and often overlooked in the mathematical finance literature that these two processes cannot simply be any two martingales. They must be uniquely linked. Let me explain what this means.

Suppose we consider X and Y , which fluctuate randomly and have expected returns of zero. Nonetheless, X and Y are quite independent. X may be the price of a gas company in Warsaw, Poland while Y might be the price of corn in Madison County, Iowa. There is no amount of mathematical wizardry that can convert the price of a gas company in Warsaw to the price of corn in Iowa, even though both may be martingales given the probability distribution of states of the global economy. The martingale representation theorem seems to say that this can be done. But what is missing can be easily seen by remembering the binomial model.

Suppose X can go up to X^+ or down to X^- . Now we want to model Y . Intuitively we might specify that Y can go up to Y^+ or down to Y^- . If we allow only two states of nature and jointly consider X and Y in the same model, then whenever X is X^+ , then Y is Y^+ and whenever X is X^- , then Y is Y^- . X and Y are perfectly related to each other. Knowledge of the value of one reveals the value of the other. In a very loose sense X and Y are thought of as perfectly correlated, though correlation is a specific type of relationship that does not have to hold to make this point. Obviously when the price of corn is up, the price of the Polish gas company can be up or down. There is probably little relation between our X and Y . Furthermore, we do not require such extreme examples to make this point. X could be the price of Microsoft while Y is the price of Exxon. Though both share a common relationship, as determined by general stock market movements, there is no way that we can completely determine the value of one from the other.

So if Microsoft and Exxon were martingales, we could not relate one to the other as the martingale representation theorem seems to suggest.

But if X is the price of Microsoft and Y is the price of a derivative, such as a call option, on Microsoft, then we may be able to connect the two. Indeed we can, since the price of the latter is completely determined by the price of the former. In the binomial sense, we can easily see this point. The option payoff at expiration for a given outcome is completely driven by the stock price in that outcome.² Thus, when we move to a continuous time framework and attempt to employ the martingale representation theorem, we require that one martingale be completely determined by the other.

How we use the martingale representation theorem is that if the price of a stock, as indicated by the X variable above, is a martingale, then by the transformation above, we can turn it into the Y variable, provided that Y is completely determined by X . If Y is the option on X , then we have replicated Y with X . It should be apparent that λ is likely to represent a certain number of shares of X held to replicate Y .

Those who have studied option pricing theory from a financial economics perspective, however, will find this result a bit disconcerting, because you know that one must hold bonds as well as shares to replicate an option. The martingale representation theorem gives us only the condition under which the uncertainty in Y can be captured by the uncertainty in X . More formally, if we differentiate Y_t with respect to X_t , we obtain

$$dY_t = \lambda_t dX_t.$$

This is the result we need and the one that clearly indicates that the uncertainty in y is driven by the uncertainty in X . As we know, we must hold bonds as well as stock to replicate an option.

But the martingale representation theorem does tell us that we need to find a stochastic process λ_t such that we are holding just the right amount of X . In addition we must hold a certain number of bonds, such that we replicate Y . Formally,

$$Y_t = \lambda_t X_t + \theta_t B_t,$$

where B_t is the value of the bonds, and which accrues value by the factor e^{rdt} . Note that λ and θ are stochastic processes. They are indexed by t and are determined as the stock evolves along its

²Of course the exercise price also determines the option payoff, but it is not random so it causes us no problems.

stochastic paths. But they need be known only at t .³ They change as we move through time but are completely determined once we know the new value of the stock price. Financial economists know this as nothing more than the option's delta changing. Indeed delta is what we are currently designating as λ .

Mathematical finance formally tells us that if the stock can be transformed into a martingale, we can find a stochastic process such that the stock can be transformed into another martingale, which replicates the option. Once we have replicated the option, we can price it using the stock price, its number of holdings, the bondholdings, the interest rate, etc., or in other words, known values. We can obtain this by finding the expectation of the future value of this process.

Before we proceed with the mathematics, let us return to one final point. As noted above, mathematicians seem to suggest that we can relate one martingale to the other without considering that these two martingales might be totally unrelated. We cannot replicate the price of corn in Iowa with the price of a gas company in Warsaw. The problem lies in the advantage that financial economists have over financial mathematicians. Financial economists study general equilibrium, which is the study of how markets reach equilibrium where outcomes are determined by the interaction of preferences and expectations to form demand, supply and, ultimately, prices and holdings of assets. General equilibrium models like the Capital Asset Pricing Model, as flawed as they may be, nonetheless, bring us an appreciation for the fact that markets are complex combinations of individuals and assets. Assets do not exist in isolation and, as such, stochastic processes of single assets, which attempt to model outcomes, make us forget everything else that is going on in the market at that time. Derivative pricing theory is considered partial equilibrium. The general equilibrium process is taken as determined externally. Without recognizing the general equilibrium process, mathematicians might have us believe that the price of corn in Iowa could be hedged with the price of a gas company in Warsaw.

Now the one case where the mathematicians would be correct is when there is a single source of uncertainty driving the two variables. That is why, as we shall see, that one random variable can be expressed in terms of another. In financial applications, one random variable is

³Mathematicians call this property *previsibility*.
D.M. Chance, TN00-04

often a derivative of the other. But it is misleading to state the martingale representation theorem as generally as most mathematicians do. They are often thinking of a world in which an option and its underlying, perhaps with a risk-free bond, are the only assets that trade.

Now let us proceed to learn the mathematics of how derivative prices are found. Our first step is to learn how to change the drift, i.e., expected return of a random variable.

Introducing the Radon-Nikodym Derivative by Changing the Drift for a Single Random Variable

Let us first begin by examining the process of changing a probability distribution for a general random variable. We are given a random variable x , which is simply a single unknown outcome and not a stochastic process. We shall take x as distributed normally with mean μ and variance σ^2 . The probability density of x is

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(\frac{x-\mu}{\sigma}\right)^2/2} = \frac{dP(x)}{dx},$$

where the probability distribution function is $P(x)$. Now suppose that we wanted to change the location of this probability distribution. Specifically, we wish to shift its mean by an amount γ . Then the density we want is

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(\frac{x-\mu-\gamma}{\sigma}\right)^2/2}.$$

In other words, we change the mean but not the variance. We need not specify if γ is positive or negative; we could shift the mean upward or downward. Let us call this new density $f^Q(x)$ and the new distribution function $Q(x)$, so $f^Q(x) = dQ(x)/dx$. Let us see how we can make this change whereby $P(x)$ becomes $Q(x)$.

Look at the expression in the exponent above, $-((x - \mu - \gamma)/\sigma)^2/2$ and note that it equals $-(1/2\sigma^2)(x^2 - 2x\mu - 2\gamma x + 2\mu\gamma + \mu^2 + \gamma^2)$. If we compare this to what we started with, $-((x - \mu)/\sigma)^2/2 = -(1/2\sigma^2)(x^2 - 2x\mu + \mu^2)$, we see that we need to multiply the original expression by something. That is,

$$\begin{aligned} & \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(\frac{x-\mu}{\sigma}\right)^2/2} e^{-(-2\gamma x + 2\gamma\mu + \gamma^2)/2\sigma^2} \\ &= \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(\frac{x-\mu-\gamma}{\sigma}\right)^2/2}. \end{aligned}$$

Let us now designate this new multiplier as

$$\phi(x) e^{(2\gamma x - 2\gamma\mu - \gamma^2)/2\sigma^2}.$$

Since $f(x) = dP(x)/dx$, then $f(x)dx = dP(x)$. What we want is $f^Q(x)dx = dQ(x)$, where we have a new probability measure $Q(x)$. Since $f(x)\phi(x) = f^Q(x)$, then $f(x)dx\phi(x) = f^Q(x)dx$. Then

$$\frac{f^Q(x)}{f(x)} = \phi(x).$$

It follows that

$$\frac{dQ(x)}{dP(x)} = \phi(x).$$

Our multiplier $\phi(x)$ can be thought of as an adjustment that converts one probability measure, $P(x)$, into another $Q(x)$. We must be careful, however, in that one cannot just arbitrarily multiply one measure by some other factor, for the resulting measure should be of the same type, here the normal distribution, as the one we started with, though now having a different mean.

In some cases our random variable will be standard normal, meaning that $\mu = 0$ and $\sigma = 1$. In that case,

$$\phi(x) = e^{x\gamma - \gamma^2/2}.$$

Note that in any case if $\gamma > 0$, the mean is shifted upward and if $\gamma < 0$ the mean is shifted downward.

This special function $\phi(x)$, which we noted can be expressed as $dQ(x)/dP(x)$, is a derivative itself and is called the Radon-Nikodym derivative. For this derivative to exist, it is necessary that the function $Q(x)$ and the function $P(x)$ be considered *equivalent probability measures*. What this means is that if an event is possible under one measure, then it is possible under the other measure. In other words, events that cannot occur in the first place, cannot be made possible by simply changing the probability measure. Likewise events that can occur in the first place, cannot be made impossible by changing the probability measure.

Changing the Drift for a Continuous Time Stochastic Process

For our applications in finance, the random variable we deal with is often a stochastic process. In many cases, the random variable will be a Brownian motion, W_t , such that

$$W_t = \int_0^t dW_u,$$

where we know that $W_0 = 0$ and the increments are distributed with mean zero and variance dt .⁴ The density of W_t is⁵

$$f(W_t) = \frac{1}{\sqrt{2\pi t}} e^{-\frac{1}{2}\left(\frac{W_t}{t}\right)^2}.$$

We wish to change this Brownian motion into another one that has a new probability measure Q . We shall shift it by an amount γ_t . In later applications we shall see that γ_t will become an extremely simple function of t , but for now let us leave it unspecified.

So what we want is a new Brownian motion that has a mean of γ_t . What will accomplish this trick is to designate:

$$\phi_t = e^{\int_0^t \gamma_u dW_u - \frac{1}{2} \int_0^t \gamma_u^2 du}.$$

In order for this transformation to be possible, we must impose a constraint on the behavior of γ_u . Specifically, we require that

$$\mathbb{E} \left[e^{\int_0^t \gamma_u^2 du} \right] < \infty.$$

This is called the *Novikov condition*. In simple terms it means that the variation in γ_u must be finite. For all of our applications, the Novikov condition will be met.⁶

If these requirements are met, then N_t can be shown to be a martingale. Let us first apply Itô's Lemma on ϕ_t :

$$d\phi_t = \frac{\partial \phi_t}{\partial W_t} dW_t + \frac{1}{2} \frac{\partial^2 \phi_t}{\partial W_t^2} dW_t^2.$$

First we find the partials:

⁴Recall that specifically, $dW_t = \varepsilon_t \sqrt{dt}$ where ε_t is a standard normal random variable.

⁵The following specification is because W_t has a mean of zero, since it starts off at $W_0 = 0$, and a variance of t . The term in parentheses is the value of the random variable minus its mean divided by its variance.

⁶As we shall ultimately see, γ will be a very simple function that in any rational financial market will have finite variation.

$$\begin{aligned}
\frac{\partial \phi_t}{\partial W_t} &= \frac{\partial}{\partial W_t} \left[e^{\int_0^t \gamma_u dW_u - \frac{1}{2} \int_0^t \gamma_u^2 du} \right] \\
&= \phi_t \frac{\partial}{\partial W_t} \left[\int_0^t \gamma_u dW_u - \frac{1}{2} \int_0^t \gamma_u^2 du \right] \\
&= \phi_t \gamma_t \\
\frac{\partial^2 \phi_t}{\partial W_t^2} &= 0.
\end{aligned}$$

So⁷

$$d\phi_t = \phi_t \gamma_t dW_t = e^{\int_0^t \gamma_u dW_u - \frac{1}{2} \int_0^t \gamma_u^2 du} \gamma_t dW_t.$$

Now let us consider the value at $t = 0$:

$$\phi_0 = e^{\int_0^0 \gamma_u dW_u - \frac{1}{2} \int_0^0 \gamma_u^2 du} \gamma_t dW_t = e^0 = 1.$$

Now we have

$$\begin{aligned}
\int_0^t d\phi_u &= \int_0^t \phi_u \gamma_u dW_u = \phi_t - \phi_0, \text{ so} \\
\phi_t - 1 &= \int_0^t \phi_u \gamma_u dW_u, \text{ and } \phi_t = 1 + \int_0^t \phi_u \gamma_u dW_u.
\end{aligned}$$

The second term on the right-hand side is a known martingale. Its expectation is zero so $E(\phi_t) = 1 = \phi_0$. Consequently, ϕ_t is a martingale.

Thus, we can now be certain that Girsanov's Theorem applies. Our Brownian motion can be transformed as follows:

$$W_t^Q = W_t - \int_0^t \gamma_u du,$$

where W_t^Q is a Wiener process under a new probability measure Q such that

$$\frac{dQ}{dP} = \phi.$$

Remember that N is the Radon-Nikodym derivative.

⁷The result on the third line above is obtained from the result on the second line above by differentiating the stochastic integral at the end point.

We shall ultimately need the Weiner differential, dW_t^Q . It is obtained as follows:

$$W_t^Q = W_t - \int_0^t \gamma_u du,$$

$$W_t^Q - W_t = -\int_0^t \gamma_u du \quad \text{implies}$$

$$dW_t^Q - dW_t = -\gamma_t dt, \quad \text{so}$$

$$dW_t^Q = dW_t - \gamma_t dt.$$

Now let us step back and think about how this is important for our purposes. We shall want to convert our stock price process to a martingale. This will remove the drift and permit us to price the option by evaluating its expected future value under the new probability measure. When we remove the drift, what we are doing is removing the risk premium and the risk-free rate. If we knew the value of the risk premium, it would be no problem: we would simply subtract it out. But we do not know what the risk premium is. We do not know how much of the stock's expected return to remove. We do know, however, that if we remove just enough that the stock return is a martingale, then we require no discounting whatsoever. So the trick is to change the probability distribution so that the stock return is a martingale. Here is where the finance ends and the math takes over. Girsanov's Theorem tells us how to change a probability distribution to leave it the same type of distribution with the same variance but with a different drift. What we have just seen above is that the Brownian motion process, W_t , can be changed such that it is a martingale. Since the stock price process is a simple transformation of the Brownian motion process, it should be easy to transform it as well into a martingale.

We have seen above that we are subtracting a function γ_t . This means that γ can potentially change with t . We are somewhat lucky here, because for our purposes γ_t is a very simple function of t : $\gamma_t = \gamma t$.⁸ Now notice what we obtain for our Radon-Nikodym derivative:

$$\phi_t = e^{\int_0^t \gamma dW_t - \frac{1}{2} \int_0^t \gamma^2 du} = e^{\int_0^t \gamma dW_t - \frac{1}{2} \gamma^2 \int_0^t du} = e^{\gamma W_t - \gamma^2 t/2}.$$

In other words, if we multiply the density function of W_t by N_t as specified above, we should obtain the density function for a new Brownian motion, which we shall call W_t^Q , in which the mean has been shifted by γt . Let us see. Given,

⁸If this were not the case, we would see that our method would not work later on.
D.M. Chance, TN00-04

$$f(W_t) = \frac{1}{\sqrt{2\pi t}} e^{-(w_t^2)}, \quad \text{and}$$

$$\frac{dQ}{dP} = e^{\gamma W_t - \gamma^2 t/2},$$

we obtain by multiplication:

$$f(W_t) \frac{dQ}{dP} = \frac{1}{\sqrt{2\pi t}} e^{-(w_t^2/2t + \gamma W_t - \gamma^2 t/2)}$$

$$= \frac{1}{\sqrt{2\pi t}} e^{-(1/2) \left(\frac{W_t - \gamma t}{t} \right)^2}.$$

This is the density for a Brownian motion with its zero mean shifted by $-\gamma t$. So

$$W_t^Q = W_t - \gamma t$$

To recap, we have that W_t is a Brownian motion under the probability measure P , such that

$$E^P(W_t) = W_0 = 0$$

$$E^P(W_t^Q) = E^P(W_t - \gamma t) = E^P(W_t) - \gamma t$$

$$= 0 - \gamma t.$$

The first statement defines that W_t is a Brownian motion under P . The second statement says that under P , W_t^Q is not a Brownian motion. Its expectation, $-\gamma t$, is not zero, except at $t = 0$, and varies with t . But W_t^Q is a Brownian motion under Q :

$$E^Q(W_t^Q) = W_0^Q = 0.$$

This statement follows since W_t and W_t^Q both start at a value of zero. Under Q , W_t is not a Brownian motion because:

$$E^Q(W_t) = E^Q(W_t^Q + \gamma t) = E^Q(W_t^Q) + \gamma t = \gamma t.$$

When we say that some random process, such as W_t and W_t^Q , is a Brownian motion under a given measure, we are saying that the probabilities of its possible paths are assigned such that its central property, a constant expectation of zero is preserved. When the probabilities are changed such that the process no longer has a zero expected value, it is no longer the same thing. But another process, can and in this case, does have the property of a Brownian motion under the new probability measure.

Changing the Drift of a Stock Price Process

In TN96-04 and TN00-03 we obtained the familiar stochastic process for a stock:

$$dS_t/S_t = \alpha dt + \sigma dW_t.$$

If we change W_t such that now $W_t = W_t^Q + \gamma t$, then we substitute its differential, $dW_t = dW_t^Q + \gamma dt$, into the above stochastic differential equation to obtain

$$\begin{aligned} dS_t^Q/S_t^Q &= \alpha dt + \sigma(dW_t^Q + \gamma dt) \\ &= (\alpha + \sigma\gamma)dt + \sigma dW_t^Q. \end{aligned}$$

So if we change the probability measure for W_t , which is the probability that drives S_t , we are now working with the above stochastic process and a new set of probabilities. But have we converted S_t^Q to a martingale? Not yet. If, however, we specify that $\gamma = -\alpha/\sigma$, the drift becomes zero, leaving us with

$$dS_t^Q/S_t^Q = \sigma dW_t^Q,$$

which is clearly a martingale. From this result, we can assign an obvious interpretation to γ , an interpretation we have already arrived at. First, ignoring the minus sign, the expression α in the numerator is the expected return. The denominator is clearly the risk. Thus, γ is the return over the risk, a kind of risk-return tradeoff. It is somewhat more natural, however, to specify γ in a slightly different manner:

$$\gamma = -\frac{\alpha - r}{\sigma},$$

which, when substituted back into the stochastic differential equation, gives us

$$dS_t^Q/S_t^Q = r dt + \sigma dW_t^Q,$$

Now we have a more natural interpretation of γ . Again, ignoring the minus sign, the numerator is the expected return minus the risk-free rate, or the risk premium. The denominator is the risk.⁹ In financial economics, this ratio is the risk premium per unit of risk and is sometimes called the *market price of risk*. It reflects the risk-return tradeoff, i.e., the additional expected return necessary to induce investors to assume risk.¹⁰

⁹ It is even more natural that instead of defining γ as $-(\mu - r)/\sigma$, we define it as $(\mu - r)/\sigma$ and then change the sign such that instead of adding γ to α , we are subtracting it. In that way, we appear to be subtracting a risk premium, a more sensible way to describe what is happening.

¹⁰ In the study of general market equilibrium, such as the Capital Asset Pricing Model, the appropriate measure of risk is not the standard deviation, but the systematic risk, also known as β . The derivatives pricing framework, however, takes the general D.M. Chance, TN00-04

Now the adjustment is, more or less, just a subtraction of the risk premium (per unit of risk). But the beauty of it all is that we never had to obtain the risk premium. By converting the process to a martingale, we removed the risk premium automatically.

For the log process, recall that its drift, μ , is equal to $\alpha - \sigma^2/2$, so then¹¹

$$\gamma = -\frac{\mu + \sigma^2/2 - r}{\sigma}.$$

But if you have been paying attention, you should note that it appears we no longer have S_t in the form of a martingale. After all, its new expected return is

$$dS_t^Q/S_t = rdt,$$

which is certainly non-zero. Clearly the stock now drifts upward at the risk-free rate. Now we seem to have a problem, but a slightly different spin on things saves the day.

First we should be comforted in knowing that by removing the risk premium, we have taken out the most difficult part of the problem. We ought to be able to solve the option pricing problem from what we now know. Indeed that is the case. Financial economists have long known that if we change the stock's expected return to the risk-free rate, we can then evaluate the expected option payoff under the assumption that the stock price is randomly generated by the standard stochastic differential equation with a drift set at the risk-free rate. Economists then go on to explain that everything we need to know about how investors feel about risk is impounded into the stock price. It is not necessary to reflect any effect of investors' risk preferences in the option price. Consequently, we can proceed to evaluate the option as if the expected return on the stock is the risk-free rate. This approach is often called *risk neutral pricing*. What we have done is equivalent to the well-known procedure of taking a short position in an option, hedging it with a long position in units of the underlying asset, thereby eliminating the risk, followed by setting the return on this hedged portfolio to the risk-free rate. From there we obtain a partial differential equation whose solution is the option pricing model.

But how do we salvage our approach, which now leaves us with S_t no longer a martingale?

The trick lies in recognizing that we can work with the discounted value of S_t . In other words, say we start off at time 0 with a value of S_0 . Then at time t , we have S_t . But suppose we

equilibrium as given and as fully reflected in the price of the underlying asset. Risk is measured in isolation and σ is the only measure of risk recognized within option pricing theory.

¹¹Let us emphasize that this is not a new definition for γ . It simply expresses γ in terms of the log return.

transform our stock price into its discounted value, $S_t e^{-rt}$. Now let us look at some of our previous results. Recall from TN00-03 that the solution to the stochastic differential equation, giving us S_t in terms of S_0 is

$$S_t = S_0 e^{\mu t + \sigma W_t}.$$

Suppose that we adjust S_t to its discounted value, $S_t e^{-rt}$. Then substituting into the above equation:

$$S_t e^{-rt} = S_0 e^{(\mu-r)t + \sigma W_t}.$$

Right now, however, we are under the original probability measure. Substituting $W_t^Q + \gamma t$ for W_t , we obtain

$$S_t e^{-rt} = S_0 e^{(\mu-r)t + \sigma(W_t^Q + \gamma t)}.$$

Noting that we defined γ as $-(\mu + \sigma^2/2 - r)/\sigma$ and substituting this result, we obtain

$$S_t e^{-rt} = S_0 e^{-\sigma^2 t/2 + \sigma W_t^Q}.$$

You may wish to look back in TN00-03 where we used this result for S_t , along with the density function for a normally distributed W_t to obtain the expected future stock price:

$$E[S_t] = S_0 e^{(\mu + \sigma^2/2)t}.$$

If we follow that same procedure here for $S_t e^{-rt}$, we obtain

$$E[S_t e^{-rt}] = S_0 e^{-\sigma^2 t/2} e^{\sigma^2 t/2} = S_0.$$

The absence of a positive expected return shows that the discounted price is a martingale.

Let us do one more thing. We shall apply Itô's Lemma to the discounted stock price. First for simplicity of notation, let us designate S_t^* as $S_t e^{-rt}$. Now applying Itô's Lemma to S_t^* , we obtain:

$$dS_t^* = \frac{\partial S_t^*}{\partial W_t} dW_t + \frac{\partial S_t^*}{\partial t} dt + \frac{1}{2} \frac{\partial^2 S_t^*}{\partial W_t^2} dW_t^2.$$

Using the above result,

$$S_t e^{-rt} = S_0 e^{(\mu-r)t + \sigma W_t},$$

we can obtain the partial derivatives:

$$\frac{\partial S_t^*}{\partial W_t} = S_t^* \sigma, \quad \frac{\partial^2 S_t^*}{\partial W_t^2} = S_t^* \sigma^2, \quad \frac{\partial S_t^*}{\partial t} = S_t^* (\mu - r).$$

Substituting and noting that $dW_t^2 = dt$, we obtain the following stochastic differential equation for S_t^* :

$$\begin{aligned}\frac{dS_t^*}{S_t^*} &= (\mu - r + \sigma^2/2)dt + \sigma dW_t \\ &= (\alpha + r)dt + \sigma dW_t.\end{aligned}$$

We see that once we have taken the risk-free rate into effect in defining the underlying variable, we no longer can account for the risk-free rate in the drift. But still we do not have a martingale. Recall that to obtain a martingale, we substitute $dW_t^Q + \gamma t$ where $\gamma = -(\alpha - r)/\sigma$ into the above, giving us

$$\frac{dS_t^*}{S_t^*} = \sigma dW_t^Q,$$

which is clearly a martingale.

To summarize, we adjust the drift of the stock price process by changing the probabilities such that we obtain a martingale. What we ultimately have done is gone ahead and discounted the stock price and then worked with the discounted stock price to change its probability measure, leaving us with a martingale. Then we can easily evaluate the option by applying the probability distribution of the discounted stock price to the option payoff. In this way the option price is its expected payoff at expiration without discounting.

If we do the one, discounting the stock price, without the other, changing the measure, we have technically not completed the process. But as it turns out, we can get away with changing the measure without discounting the stock price. Recall that $dS_t/S_t = \alpha dt + \sigma dW_t$. Substitute $dW_t^Q + \gamma dt$ and $\gamma = -(\alpha - r)/\sigma$ and we obtain:

$$\frac{dS_t}{S_t} = rdt + \sigma dW_t^Q,$$

We can work with this model and do the discounting later, i.e., after we have evaluated the expected option payoff at expiration. That is because whether we discount before we have performed the expectation or after, we have not altered the fundamental process or results of taking expectations other than by the simple linear adjustment, e^{-rt} . While mathematicians would prefer that we convert the stock price to a martingale, requiring that we do the discounting

beforehand, financial economists would prefer to do the discounting afterwards. That is because the latter approach is more in line with the intuition provided by economic theory: the price of any asset is its expected future value, discounted to the present at an appropriate rate. That in this case the appropriate rate is the risk-free rate is quite intuitive. The risk has been removed via the risk-free hedge, or alternatively, can be viewed as fully imbedded into the price of the underlying asset and, therefore, cannot legitimately be incorporated again. Moreover, if the risk is either not present or removed, investors' risk preferences are irrelevant to the valuation process. In that case, one might just as well use the simplest form of risk preferences: risk neutrality, wherein investors discount future values at the risk-free rate.

Finally, we might just simply say that if the price of the underlying asset is given, any two investors regardless of their feelings about risk will value the option in the same manner. Consequently, we can treat both investors as though they had the simplest risk preferences of all, risk neutrality.

References

- Baxter, M. and A. Rennie. *Financial Calculus*. Cambridge: Cambridge University Press (1996), Ch. 3.
- Dothan, M. U. *Prices in Financial Markets*. New York: Oxford University Press (1990), Chs. 5, 6.
- Duffie, D. *Dynamic Asset Pricing Theory*, 2nd. ed. Princeton: Princeton University Press (1996), Ch. 6, Appendix D.
- Duffie, D. *Security Markets: Stochastic Models*. Boston: Academic Press (1988), Chs. 21.
- Epps, T. W. *Pricing Derivative Securities*. Singapore: World Scientific (2000), Ch. 3.
- Karatzas, I. and S. E. Shreve. *Brownian Motion and Stochastic Calculus*, 2nd. ed. Berlin: Springer-Verlag (1991), Chs. 3, 5.
- Musiela, M. and M. Rutkowski. *Martingale Methods in Financial Modelling*. Berlin: Springer-Verlag (1997), Appendix B.
- Neftci, S. *An Introduction to the Mathematics of Financial Derivatives*, 2nd. ed. San Diego: Academic Press (2000), Chs. 14, 15.

Nielsen, L. T. *Pricing and Hedging Derivative Securities*. Oxford, U.K.: Oxford University Press (1999), Ch. 2.