

The origins of risk-neutral pricing and the Black-Scholes formula*

L. C. G. Rogers
School of Mathematical Sciences
University of Bath
Bath BA2 7AY, UK

December 13, 2005

1 Introduction

The theory and practice of finance today requires many skills - computing, applied mathematical, probabilistic, statistical, economic - and it is a sad observation that there are many colleagues working in finance who are expert in their own area, but know little of even the basic ideas from the other areas. This article is addressed to those who want to fill in their background a little, and learn something about the fundamental economic ideas which have inspired the development of finance, even though the subject has now become so refined that we may often study it without being aware of its intellectual pedigree. There is nothing original in this article, in the sense that a good financial economist would read it and say, 'Well, of course'. On the other hand, for those who have not seen this way of thinking of things, there are miraculous revelations ahead; we shall see how martingales, stochastic integration and the notion of equivalent martingale measures leap out of the page, for example - these are not just mathematical irrelevancies foisted on the subject by self-satisfied probabilists, they are inevitable consequences of economic thinking. As befits a pedagogical approach, we feel free to cut corners; what follows is rigorous by the standards of physics or applied mathematics, but not rigorous. Indeed, making some of the arguments into proofs would be onerous if not impossible, but the importance of the ideas is that they guide our thinking.

The fundamental concept is that of an economic *equilibrium*. Imagine a market with many agents, each of whom begins with some assets, and may trade them.¹ The objective of each agent is to maximise his utility, which is some function of the assets

*This article appears in *Risk Management and Analysis*, Vol. 2, ed. C. O. Alexander, Wiley, Chichester, 1998, pp81-94.

¹We might also allow the possibility of other economic activities, such as production; this does not alter the essential features of equilibrium, so we shall omit them and discuss only a *pure exchange* economy.

held at the end of trading. When two or more agents meet, if there were any mutually advantageous trades available to them, then they would trade, and all benefit. If trading has taken place to the point where no more such mutually advantageous trades are possible, then the market is in equilibrium. The essential insight (due to Arrow and Debreu) is that *in equilibrium, there are equilibrium prices for the different assets in terms of each other*, and the allocations held by each agent are *what the agents would optimally choose to hold if they were alone in the world, but could buy the different assets at those equilibrium prices*.

This allows us to understand an equilibrium by firstly studying what an agent would optimally choose to do if faced with certain prices for the assets; and then adjusting those prices so that *the markets clear*, that is, the total amounts of the different assets demanded by the optimally-behaving agents are the total amounts present initially in the market.

We shall follow this recipe in a market evolving randomly in time. To begin with (Section 2), we shall study the possible choices available to an agent, and then we shall see (in Section 4) what are the consequences of optimal behaviour of such an agent. This leads naturally to a notion of equivalent martingale measures, and gives a methodology for pricing contingent claims, even in incomplete markets. Finally in Section 5, we take the simple binomial market and see how these ideas work through in that setting, leading eventually to the Black-Scholes equation when we take a suitable limit. Section 3 provides a digression into the background of martingales and equivalent measures, which hopefully many readers will already know. In the Appendix, we show how risk-neutral pricing follows from two completely different approaches, an axiomatic approach, and the no-arbitrage approach. These two routes to risk-neutral pricing are even shorter than that of Sections 2 and 4, but less illuminating.

2 Portfolio choices.

We shall consider a market developing in time, in which there are n shares, the price at time t of the i th share being denoted S_t^i . There is also a ‘zeroth share’, whose price at time t will be written R_t ; we often think of this rather differently, as the value at time t of a unit of money invested at time 0 in a deposit account, though this interpretation is not essential. For brevity, we write

$$S_t \equiv (S_t^1, \dots, S_t^n)^T, \quad \bar{S}_t = (R_t, S_t^1, \dots, S_t^n)^T,$$

where a superscript T denotes transpose. At time t , the agent holds a portfolio

$$\bar{\theta}_t \equiv (\varphi_t, \theta_t^1, \dots, \theta_t^n)^T$$

where θ_t^i is the number of i -shares held at time t , and φ_t is the number of 0-shares held at time t . The market value of the portfolio at time t is therefore

$$(2.1) \quad V_t = \bar{\theta}_t \cdot \bar{S}_t \equiv \varphi_t R_t + \sum_{i=1}^n \theta_t^i S_t^i$$

What portfolios can the agent choose? Clearly we cannot allow the agent to take $\theta_t^i = \exp(10^6 \cdot t)$ - how would he pay for it?! To understand this, suppose that there are two (deterministic) times $T_1 < T_2$, and the agent holds a portfolio process

$$(2.2) \quad \bar{\theta}_t = \begin{cases} H, & 0 \leq t \leq T_1 \\ H', & T_1 < t \leq T_2 \end{cases}$$

at time t . Thus the value of the portfolio at any time $t \in [0, T_1)$ is just $H \cdot \bar{S}_t$, and the change of value at time T_1 is

$$(H' - H) \cdot \bar{S}(T_1),$$

where we use the equivalent notations $\bar{S}_t \equiv \bar{S}(t)$ to avoid clumsiness. If this difference is zero, we say that the portfolio is *self-financing*; if the change in value is zero, no money needs to be added or taken away to finance the change of portfolio! Assuming that the portfolio *is* self-financing, for any $t \in (T_1, T_2]$ the value of the portfolio satisfies

$$(2.3) \quad \begin{aligned} V_t &= H' \cdot \bar{S}_t \\ &= H' \cdot (\bar{S}_t - \bar{S}(T_1)) + H \cdot \bar{S}(T_1) \\ &= H' \cdot (\bar{S}_t - \bar{S}(T_1)) + H \cdot (\bar{S}(T_1) - \bar{S}_0) + H \cdot \bar{S}_0 \end{aligned}$$

where we have used the self-financing condition to pass from the first line to the second. The three terms in the final expression have very simple interpretations; the first is the change in value of the portfolio in the time interval $(T_1, t]$, the second is the change in value in the time interval $[0, T_1]$, and the third is the initial value V_0 of the portfolio.

We can easily generalise to a portfolio which gets changed at the times $T_1 < T_2 < \dots$. If the agent chooses to hold H_i throughout the time-interval $(T_{i-1}, T_i]$, and if the portfolio is self-financing (so none of the changes involve any alteration in the value of the portfolio), then the value at time t will be

$$(2.4) \quad \begin{aligned} V_t &= V_0 + H_1 \cdot (\bar{S}(T_1) - \bar{S}(0)) + H_2 \cdot (\bar{S}(T_2) - \bar{S}(T_1)) \\ &\quad + \dots + H_k \cdot (\bar{S}(t) - \bar{S}(T_{k-1})) \end{aligned}$$

if $T_{k-1} < t \leq T_k$. This can be written more concisely as

$$(2.5) \quad V_t = V_0 + \int_{(0,t]} \bar{\theta}_u \cdot d\bar{S}_u.$$

The two alternative expressions can be thought of as the *definition* of the integral notation appearing on the right of (2.5); later, we need to consider what the integral might mean for more complicated portfolio processes $\bar{\theta}$, but for piecewise-constant $\bar{\theta}$ the notational equivalence of (2.4) and (2.5) is unmistakable. Thus we may write

$$V_t - V_0 = \int_{(0,t]} \bar{\theta}_u \cdot d\bar{S}_u.$$

Even though we may not yet know how to define the integral appearing on the right of (2.5) for every $\bar{\theta}$, it is clear that for the piecewise constant $\bar{\theta}$ being used here it *could* only be defined as (2.4). The right-hand side of (2.5) is called the *gains from trade*; it

is the change in value of the portfolio arising from the fluctuations in the prices of the assets. Thus we have the key result that for a self-financing portfolio,

Change in value = Gains from trade

This principle holds in complete² generality, so that (2.5) is true whatever the portfolio process $\bar{\theta}$. To make sense of this, we would need to make a definition of the (stochastic) integral appearing on the right of (2.5), and this is not a minor task as the processes S^i may often have paths of unbounded variation. But perhaps we could get by with using only simple piecewise-constant integrands of the type we have looked at so far? We shall soon have an answer to this question.

Is there any easy way to tell when a portfolio is self-financing? This will obviously be important. To understand this, let us suppose initially that $R_t = 1$ for all t . Then for any self-financing portfolio $\bar{\theta}$

$$(2.6) \quad V_t = \varphi_t + \theta_t \cdot S_t = V_0 + \int_{(0,t]} \theta_u \cdot dS_u.$$

So we see that we can choose *any* θ , and then adjust φ to make the above equation hold. This is intuitively reasonable; we may hold any portfolio of the shares provided we take the money out of our deposit account to pay for it. More generally, if we assume that R grows continuously and define $\tilde{V}_t \equiv V_t/R_t$, $\tilde{S}_t \equiv S_t/R_t$, we have

$$(2.7) \quad \tilde{V}_t = \tilde{V}_0 + \int_{(0,t]} \theta_u \cdot d\tilde{S}_u.$$

In words, the change in the *discounted* value of the portfolio is the integral of the portfolio process with respect to the *discounted* asset price processes; if we realise that working with discounted prices is effectively changing the bank account process to be constant, it is not surprising that we get the same characterisation of self-financing portfolio processes as we had at (2.6). This characterisation of the possible wealth processes from self-financing portfolio choice is the starting point for our understanding of the optimal behaviour of an agent.

3 Some notions and notations from probability.

Everything in this section is quite standard, so please quickly skim it, and if the contents are familiar go immediately to the next section.

As will by now be clear, we are in the business of studying random processes developing in time. The time parameter set, to be denoted \mathcal{T} , will always be *either* the set of non-negative integers, $\{0, 1, 2, \dots\}$ - the ‘discrete-time’ setting - *or* the set of non-negative reals - the ‘continuous-time’ setting. As time passes, an agent gets to know more and more, and his decisions may only be made in the light of information known at the time the decision had to be taken. Thus the choice $\bar{\theta}_t$ of portfolio to be held at time t must

²Of course, there have to be *some* restrictions; $\bar{\theta}$ needs to be non-anticipating in a precise sense, and not to grow too wildly - local boundedness is certainly sufficient.

depend only on the information available at time t - the ‘technical’ way to say this is that $\bar{\theta}_t$ should be \mathcal{F}_t -measurable, where \mathcal{F}_t is the σ -field of events known at time t . It is intuitively clear that if an agent is investing in discrete time in 5 assets, then his choice of portfolio to be held on day n should be a function only of the prices of the 5 assets on earlier days, up to and including the $(n-1)$ th. In this simple setting, this is equivalent to what the ‘technical’ statement says; however, the ‘technical’ statement holds unaltered for much more general situations (continuous time, with an uncountable infinity of available assets, say) in which the notion of ‘a function of earlier prices’ oversteps what the mathematics or the imagination can support!

A stochastic process $(X_t)_{t \in \mathcal{T}}$ is a family of random variables, and is said to be *adapted* if X_t is \mathcal{F}_t -measurable for each $t \in \mathcal{T}$. The classic example is an asset price process; assuming perfect information, the price of an asset at time t is always known at time t ! The portfolio $\bar{\theta}_t$ an agent holds at time t is also an adapted process, but it is even a little more; as the discrete-time example above illustrates, the portfolio held on day n had to be decided on day $n-1$, so was known in advance. Such a process is called *previsible*; there is an analogous concept in continuous time, but it is much harder to define, so we shall not even attempt to - suffice it to say that *previsibility is the natural measurability restriction on portfolio processes*.

Probably the most important class of processes is the class of *martingales*. We give here only the briefest summary of the definitions; for the perfect account of martingale theory in discrete time, you have to consult Williams [8]. A martingale is an adapted process $(M_t)_{t \in \mathcal{T}}$ with the properties that $\mathbb{E}|M_t| < \infty$ for every t , and

$$(3.1) \quad \begin{aligned} M_s &= \mathbb{E}_s M_t \\ &\equiv \mathbb{E}[M_t | \mathcal{F}_s] \end{aligned}$$

for all $s \leq t$ in \mathcal{T} . In words, for any $s < t$, the expected value of M_t given \mathcal{F}_s is M_s . The equivalent notations $\mathbb{E}_s Y$ and $\mathbb{E}[Y | \mathcal{F}_s]$ for the *conditional expectation* of a random variable Y given information known by time s are defined by the two properties that $\mathbb{E}_s Y$ is always \mathcal{F}_s -measurable, and

$$(3.2) \quad \mathbb{E}(Y I_F) = \mathbb{E}(\mathbb{E}_s(Y) I_F)$$

for all events $F \in \mathcal{F}_s$. Here, I_F is the random variable which is 1 on the event F , and is otherwise 0; it is an *indicator* random variable, which indicates whether an event has happened or not. An easy and common example from finance would be where the random variable Y is the payoff of a European call option on a share, and the event F is the event that the price of the share has not dropped below 50 before the option expires. In that case, $Y I_F$ is the payoff of a down-and-out call option, with knockout barrier at 50.

To rephrase (3.1) and (3.2) then, a process $(M_t)_{t \in \mathcal{T}}$ is a martingale if for any $s < t$ and any $F \in \mathcal{F}_s$

$$(3.3) \quad \mathbb{E}[(M_t - M_s) I_F] = 0,$$

and in this form we shall presently meet the martingale condition again.

The only other notion we need to introduce for now is that of a *change of measure*. At one level, this is a formal procedure, and very easy to handle as such. Suppose we are given some non-negative random variable Z , for which $0 < c \equiv \mathbb{E}Z < \infty$. Then we can use it to define a new probability \mathbb{P}^* , say, via the recipe

$$(3.4) \quad \mathbb{P}^*(A) = c^{-1} \mathbb{E}[ZI_A].$$

Notice that this *is* a probability - the countable additivity of \mathbb{P}^* is a consequence of properties of the integral, and the fact that $\mathbb{P}^*(\Omega) = 1$ follows from the definition of c . It is clear that \mathbb{P}^* is *absolutely continuous* with respect to \mathbb{P} in the sense that if $\mathbb{P}(A) = 0$ then also $\mathbb{P}^*(A) = 0$ - this is a basic property of integrals. The fact that a measure \mathbb{P}' which is absolutely continuous with respect to \mathbb{P} in this sense must have a representation of the form (3.4) for some random variable Z is a deep and important theorem, the Radon-Nikodym theorem.

As an example of how these things can arise in practice, let's consider a game where you bet on N successive tosses of a coin. This coin lands H with probability p , and lands T with probability $q = 1 - p$. You win 1 each time the coin lands H , you lose 1 each time the coin lands T . Let S_n denote your accumulated gains after the n th toss of the coin, so that $S_0=0$, and $|S_n - S_{n-1}| = 1$ for all $n \geq 1$. The situation is extremely simple; there are just 2^N possible outcomes in the sample space, a typical outcome ω being just a sequence of N symbols, H or T . The probability of a particular sequence ω is simply

$$p^j q^{N-j},$$

where j is the number of H s in the sequence ω . This could equivalently be written as

$$p^{(N+S_N)/2} q^{(N-S_N)/2} \equiv \varphi(p, S_N),$$

say. If we took as our basic probability \mathbb{P} the probability which arises in the situation where $p = 1/2$, we could express the probability \mathbb{P}_p corresponding to a coin with H -probability p via the change-of-measure

$$Z = \frac{\varphi(p, S_N)}{\varphi(1/2, S_N)}.$$

In this example, the gains process under \mathbb{P} is actually a martingale, as you are as likely to win 1 as to lose 1 at each toss. If p was different from $1/2$, that is, if the coin were not fair, the gains process would not be a martingale, because you would tend on average to win if $p > 1/2$ or to lose if $p < 1/2$. What we see in this example then is that *by changing measure from \mathbb{P}_p to \mathbb{P} , we convert the process S into a martingale!* This is an important technique in finance, which we are going to justify by economic reasoning in the next section, and in two other ways. It is also important to understand that it is *only* a technique; the transformed probability has no 'real' status, in that it does not describe how assets behave in the real world³, it merely allows us to compute prices and (in some cases) to calculate hedges. The transformed probability can be given various interpretations, as we shall see, and these are illuminating.

³The paper Rogers & Satchell [7] gives a stark warning about the difference between real-world and so-called 'risk-neutral' probabilities.

4 Optimal investment.

There are many different optimisation problems we could pose for the agent, which all lead to broadly similar conclusions. Here is one which is easy to work with. We assume that the agent has a utility function U , which is strictly increasing and strictly concave⁴ on $(0, \infty)$, a fixed investment horizon $T > 0$, and aims to

$$\max \mathbb{E} U(V_T)$$

starting from initial wealth V_0 . Suppose that the optimal wealth is V_T^* , achieved by using the self-financing portfolio θ^* . If we now consider changing slightly from this optimum θ^* to $\theta^* + \varepsilon\eta$, where ε is very small, the discounted terminal wealth $R_T^{-1}V_T^*$ gets changed to

$$R_T^{-1}V_T = R_T^{-1}V_T^* + \varepsilon \int_{(0,T]} \eta_u d\tilde{S}_u \equiv R_T^{-1}(V_T^* + \Delta V),$$

from (2.7). Using the first two terms $U(V_T) = U(V_T^*) + U'(V_T^*) \Delta V + \dots$ of the Taylor expansion, we find that to first order in ε

$$\mathbb{E} U(V_t) = \mathbb{E} \left[U(V_T^*) + U'(V_T^*) \cdot \varepsilon \int_{(0,T]} \eta_u d\tilde{S}_u \cdot R_T \right] + o(\varepsilon) \leq \mathbb{E} U(V_T^*),$$

the inequality coming from the fact that V^* is optimal. From this, we learn that

$$\mathbb{E} \left[U'(V_T^*) \cdot R_T \cdot \int_0^T \eta_u d\tilde{S}_u \right] \leq 0,$$

and by arguing similarly for the perturbation $-\eta$ the conclusion is that for any (admissible) perturbation η

$$\mathbb{E} \left[U'(V_T^*) \cdot R_T \cdot \int_0^T \eta_u d\tilde{S}_u \right] = 0,$$

which we may re-express as

$$(4.1) \quad \mathbb{E}^* \left[\int_0^T \eta_u d\tilde{S}_u \right] = 0$$

when we define the probability \mathbb{P}^* by

$$\frac{d\mathbb{P}^*}{d\mathbb{P}} = c \cdot U'(V_T^*) \cdot R_T,$$

the constant c being chosen for correct normalisation; see (3.4). By taking especially simple perturbations η , we have a remarkable conclusion. Fixing $t \in (0, T)$, taking some event $F \in \mathcal{F}_t$ and using the perturbation

$$\eta_u = I_{\{t < u \leq T\}} I_F,$$

the little change in discounted wealth $\varepsilon \int_{(0,T]} \eta_u d\tilde{S}_u$ becomes $(\tilde{S}_T - \tilde{S}_t) I_F$, statement (4.1) becomes

$$\mathbb{E}^* \left[(\tilde{S}_T - \tilde{S}_t) I_F \right] = 0$$

⁴The strict increase of U is natural, and the strict concavity reflects the agent's risk aversion

and we recognise (see (3.3)) that

under \mathbb{P}^* , discounted price processes are martingales.

So without any contrivance, we have seen that the existence of a probability \mathbb{P}^* equivalent to the original \mathbb{P} under which all discounted asset price processes are martingales *follows from simple ideas of equilibrium!*

Remarks. (i) In recent years, the mathematical finance literature has been full of papers proving various forms of the so-called ‘Fundamental Theorem of Asset Pricing’⁵ which broadly says that in a market there are no arbitrage opportunities if and only if there is an equivalent measure under which all asset price processes are martingales (an EMM). This literature sprang from the original papers of Harrison & Kreps [3] and Harrison & Pliska [4], but perhaps the most important issue is why one should think of formulating such a statement in the first place!! David Kreps, being a good economist, would undoubtedly have been familiar with the kind of arguments we have just seen, and in the light of this, the formulation of the Fundamental Theorem of Asset Pricing is explained.

It is clear that in equilibrium there can be no arbitrage, and we have just seen that in equilibrium there is an EMM. The fact that these two consequences of equilibrium are equivalent is remarkable and very important (or, depending on your point of view, very unimportant!)

(ii) The optimal portfolio in the general continuous-time setting will not typically be a simple piecewise-constant portfolio, so we really do need a general theory of stochastic integration to underpin the sort of analysis carried out above.

In general, there will be no uniqueness of EMMs; any two agents in the same market would generate an EMM, and these will in general be different. In a *complete* market, which is a market where every contingent claim can be perfectly replicated by suitable trading in the underlying securities, there is of course a unique price for any contingent claim, namely the initial wealth needed to finance the replicating portfolio. We shall study an example of such a market in the next section. The well-established theory of complete markets has been extensively developed in the last twenty years, but incomplete markets (which are much more difficult to handle) have been comparatively neglected. Various mathematicians have made peculiar attempts to define a price of a non-marketed contingent claim in an incomplete market, demonstrating thereby their unfamiliarity with the kinds of arguments sketched above. And indeed, the ideas sketched above need only minor extension to price contingent claims in an incomplete market. Here’s how.

If Y is a bounded non-negative contingent claim whose value will be known by time T (for example, the payoff of a European put option), what would be a fair price p_ε for our agent to pay at time 0 in order to receive εY at time T ? If he paid out p at time

⁵The definitive version of this result appears now to have been proved by Delbaen & Schachermayer [2]. We shall sketch the main ideas of the result in the Appendix.

0, his deposit account would be short by an amount pR_T at time T . To offset this, he would be free to go into the market and invest in the assets traded there; the best he could do by trading in the market, his maximised expected utility of terminal wealth, would then be just

$$\sup \mathbb{E} U(V_T + \varepsilon Y - pR_T).$$

As p increases, this expression decreases, and if p were zero it would certainly be better than $\mathbb{E} U(V_T^*)$, which is the best he could do if he did not enter into any deal involving Y . So there will be (under mild assumptions) a unique p_ε at which he would be just indifferent to this deal, where

$$\mathbb{E} U(V_T^*) = \sup \mathbb{E} U(V_T + \varepsilon Y - p_\varepsilon R_T).$$

This value of p is the agent's maximal buying price for Y ; he would always want to buy Y if it were offered at less than p_ε , and would never pay more than p_ε for Y . Notice that the sup is at least what he would achieve if he used the portfolio θ^* and obtained wealth V_T^* , so

$$\mathbb{E} U(V_T^*) \geq \mathbb{E} U(V_T^* + \varepsilon Y - p_\varepsilon R_T).$$

Expanding the right-hand side to order ε we get

$$\mathbb{E} U(V_T^*) \geq \mathbb{E} [U(V_T^*) + U'(V_T^*)(\varepsilon Y - p_\varepsilon R_T)] + o(\varepsilon).$$

Rearranging, dividing by ε and letting ε drop to zero, we obtain

$$\lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} p_\varepsilon \geq \mathbb{E}^* [R_T^{-1} Y].$$

If we now consider the analogous argument for buying $-\varepsilon Y$, that is, selling εY , we obtain

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} p_\varepsilon = \mathbb{E}^* [R_T^{-1} Y]$$

that is, *the fair time-0 marginal price of a contingent claim is the \mathbb{P}^* -expectation of its value discounted back to time 0*. Thus different agents would have different notions of a fair price for a given contingent claim. There's much more lurking here,⁶ but let's leave it now, and look at a simple example.

5 The binomial market, and the Black-Scholes formula.

In general, the kind of optimisation problem needed to find an EMM by the route described above will not be easy to solve. Here is a very simple example where we can get round the problem altogether.

To start with, assume that we are in discrete time, with a deposit account which starts the period worth 1, and a share that starts the period worth 1. At the end of the period

⁶For example, the result of Jacka [5] that a contingent claim is attainable if and only if its price is the same with respect to all equivalent martingale measures.

the share will be worth u if the period was good, and $d < u$ if the period was bad; the deposit account will be worth $\rho \in (d, u)$ whatever the type of period. Suppose we have an agent who is trying to maximise his expected utility of wealth at the end of the period, in the manner described in the last section. He will end up with a measure \mathbb{P}^* on the set of possible outcomes, which in this easy example contains just two points, ‘good’ and ‘bad’. If p denotes the probability of ‘good’ under \mathbb{P}^* , we know that the discounted share price process becomes a martingale under \mathbb{P}^* , so the expectation of its discounted value at the end of the period must be its value at the beginning of the period:

$$p \frac{u}{\rho} + (1 - p) \frac{d}{\rho} = 1$$

From this, simple algebra leads to the conclusion

$$(5.1) \quad p = \frac{\rho - d}{u - d}$$

Now we see that *the value of p is unique*; it does not depend on the preferences of the agent. ⁷

Now we extend the model, assuming that $u = 1/d$,⁸ and consider one period after another, all independent of each other, and all behaving as in the one-period model above; each period, the share’s value gets multiplied by u (on a good period) or by d (on a bad period), and the deposit account’s value always gets multiplied by ρ . As before, there is a unique EMM, and under this the log-price process X becomes a random walk with ‘up’ probability given by (5.1).

We’ll now begin to think of the model with N periods as representing the movement of the asset prices during some fixed time interval $[0, T]$, each period corresponding to $\Delta t \equiv T/N$ units of time. If the compound interest rate of the deposit account is r , we should therefore express the one-period return of the deposit account as $\rho = \exp(r\Delta t)$, and we shall similarly write $u = \exp(\Delta x)$, where we think of Δx as small, and we have in mind ultimately to let both Δx and Δt tend to zero appropriately. In these terms, we shall translate (5.1) to

$$(5.2) \quad p = \frac{e^{r\Delta t} - e^{-\Delta x}}{e^{\Delta x} - e^{-\Delta x}},$$

and in keeping with this way of thinking we shall write $X^{(N)}$ for the log-price process; $X^{(N)}$ is a simple random walk, with steps which take values Δx and $-\Delta x$ only, with probabilities p and $1 - p$ respectively, given by (5.2).

Now consider a contingent claim

$$Y = f(X_{N\Delta t}^{(N)}) \equiv f(X_T^{(N)})$$

⁷This is because the market is complete, and every contingent claim can be replicated perfectly. See any introductory text on finance for the story.

⁸With this assumption, the set of possible values for the share at any time is contained in $\{u^n : n \in \mathbb{Z}\}$.

At time $t = j\Delta t < T$, when $X_t^{(N)} = m\Delta x$, what is the 'fair' price to pay for the contingent claim Y to be delivered at T ? It has to be the expected discounted value of the contingent claim:

$$(5.3) \quad V_N(j\Delta t, m\Delta x) = \mathbb{E}^* [\exp(-(N-j)r\Delta t) \cdot Y]$$

Because we know that the log-price process is a simple random walk, we could rewrite this in terms of an expectation with respect to the binomial distribution: writing $q = 1 - p$ we have

$$V_N(j\Delta t, m\Delta x) = e^{-(N-j)r\Delta t} \sum_{i=0}^{N-j} p^i q^{N-j-i} \binom{N-j}{i} f((m+2i-N+j)\Delta x)$$

But this can be expressed equivalently (and more usefully) as

$$(5.4) \quad V_N(j\Delta t, m\Delta x) = e^{-r\Delta t} \{pV_N((j+1)\Delta t, (m+1)\Delta x) + qV_N((j+1)\Delta t, (m-1)\Delta x)\}$$

with the boundary conditions $V_N(N\Delta t, m\Delta x) = f(m\Delta x)$. We can verify (5.4) directly using Pascal's triangle, or we can preferably interpret it as the Bellman equation of dynamic programming. From this, we obtain

$$(5.5) \quad \begin{aligned} V_N(j\Delta t, m\Delta x) - e^{-r\Delta t} V_N((j+1)\Delta t, m\Delta x) \\ = e^{-r\Delta t} \left(\frac{1}{2} [V_N((j+1)\Delta t, (m+1)\Delta x) - 2V_N((j+1)\Delta t, m\Delta x) \right. \\ \quad \left. + V_N((j+1)\Delta t, (m-1)\Delta x)] \right. \\ \quad \left. + (p - \frac{1}{2}) V_N((j+1)\Delta t, (m+1)\Delta x) \right. \\ \quad \left. + (q - \frac{1}{2}) V_N((j+1)\Delta t, (m-1)\Delta x) \right) \end{aligned}$$

which looks quite a bit more complicated than (5.4). But the grouping of the terms is appropriate for what we intend to do next, which is let N go to infinity, while preserving the relation

$$(\Delta x)^2 = \sigma^2 \Delta t.$$

If we do this, then from (5.2) we see that

$$p - \frac{1}{2} \sim \frac{(r - \frac{1}{2}\sigma^2)\Delta t}{2\Delta x}$$

so dividing (5.5) by Δt and letting Δt go to zero, the differencing operations having been grouped appropriately converge to corresponding differential operators, and we get in the limit

$$(5.6) \quad \frac{\partial V}{\partial t} + rV = \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2} + (r - \frac{\sigma^2}{2}) \frac{\partial V}{\partial x}$$

This formal calculation has led us to the celebrated *Black-Scholes PDE*. We are being rather cavalier in assuming that the discretely-defined value functions V_N have a smooth

limit for which the analogue (5.6) of (5.5) holds, but this is indeed true. The best way to see this is as a consequence of the weak convergence of the random walks $X^{(N)}$ to the process $X_t = \sigma W_t + (r - \sigma^2/2)t$, a Brownian motion with constant drift. As a consequence, the value functions V_N expressed in the form (5.3) converge pointwise to the limit function

$$V(t, x) = E[f(X_T)|X_t = x],$$

at least if f is bounded and continuous (as would be the case for a European put option, for example). But now it is an easy matter to deduce that V solves the Black-Scholes PDE (5.6), since we can write

$$V(T - s, x) = \int \exp(-y^2/2\sigma^2s) f(x + y + (r - \sigma^2/2)s) \frac{dy}{\sqrt{2\pi\sigma^2s}},$$

and directly verify from this that V solves (5.6). Solving the PDE (5.6) with the boundary condition

$$V(T, x) = (K - e^x)^+$$

gives the Black-Scholes formula for the price of a European put option with strike K .

6 Appendix: two other approaches

In this Appendix, we present two very different approaches, which both lead to risk-neutral pricing. Each is quite direct, indeed, arguably simpler than the route we took earlier in Sections 2 and 4, but does not yield the insight of that (economic) approach. To emphasise this, note that each of the approaches here yields only the *existence* of a risk-neutral pricing measure, but gives no guidance on which to choose when there is more than one.

AXIOMATIC APPROACH. We put ourselves in a filtered probability space $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ such that the σ -field \mathcal{F}_0 is trivial, and suppose that we have pricing operators $(\pi_{tT})_{0 \leq t \leq T}$ for contingent claims; if Y is some bounded random variable which is \mathcal{F}_T measurable, the time- t ‘market’ price of Y is

$$\pi_{tT}(Y).$$

This will be a bounded \mathcal{F}_t -measurable random variable; any sensible definition of ‘the market price’ at time t would have to be a random variable, as the information contained in \mathcal{F}_t would inevitably affect what the market was willing to pay for Y .

We shall assume that the pricing operators $(\pi_{tT})_{0 \leq t \leq T}$ satisfy certain axioms:

- (A1) Each π_{tT} is a bounded positive linear operator from $L^\infty(\mathcal{F}_T)$ to $L^\infty(\mathcal{F}_t)$;
- (A2) If $Y \in L^\infty(\mathcal{F}_T)$ is almost surely 0, then $\pi_{0T}(Y)$ is 0, and if $Y \in L^\infty(\mathcal{F}_T)$ is non-negative and not almost surely 0, then $\pi_{0T}(Y) > 0$;
- (A3) For $0 \leq s \leq t \leq T$ and each $X \in L^\infty(\mathcal{F}_t)$ we have

$$\pi_{st}(X\pi_{tT}(Y)) = \pi_{sT}(XY);$$

(A4) For each $t \geq 0$ the operator π_{0t} is bounded monotone-continuous - which is to say that if $Y_n \in L^\infty(\mathcal{F}_t)$, $|Y_n| \leq 1$ for all n , and $Y_n \uparrow Y$ as $n \rightarrow \infty$, then $\pi_{0t}(Y_n) \uparrow \pi_{0t}(Y)$ as $n \rightarrow \infty$.

Axiom (A1) says that the price of a non-negative contingent claim will be non-negative, and the price of a linear combination of contingent claims will be the linear combination of their prices - which are reasonable properties for a market price. Axiom (A2) says that a contingent claim that is almost surely worthless when paid, will be almost surely worthless at all earlier times (and conversely) - again reasonable. The third axiom, (A3), is a ‘consistency’ statement; the market prices at time s for XY at time T , or for X times the time- t market price for Y at time t , should be the same, for any X which is known at time t . The final axiom is a natural ‘continuity’ condition which is needed for technical reasons.

Let’s see where these axioms lead us. Firstly, for any $T > 0$ we have that the map

$$A \mapsto \pi_{0T}(I_A)$$

defines a non-negative measure on the σ -field \mathcal{F}_T , from the linearity and positivity (A1) and the continuity property (A4). Moreover, this measure is absolutely continuous with respect to \mathbb{P} , in view of (A2). Hence there is a non-negative \mathcal{F}_T -measurable random variable ζ_T such that

$$\pi_{0T}(Y) = \mathbb{E}[\zeta_T Y]$$

for all $Y \in L^\infty(\mathcal{F}_T)$. Moreover, $\mathbb{P}[\zeta_T > 0] > 0$, because of (A2) again. Now we exploit the consistency condition (A3); we have

$$\pi_{0t}(X\pi_{tT}(Y)) = \mathbb{E}[X\zeta_t\pi_{tT}(Y)] = \pi_{0T}(XY) = \mathbb{E}[XY\zeta_T].$$

Since $X \in L^\infty(\mathcal{F}_t)$ is arbitrary, we deduce that

$$\pi_{tT}(Y) = \mathbb{E}_t[Y\zeta_T]/\zeta_t,$$

which shows that the pricing operators π_{st} are actually given by a risk-neutral pricing recipe, with the state-price density process ζ . In practice, the state-price density process is often decomposed as the product of the discount factor $\exp(-\int_0^t r_s ds)$ and the change-of-measure martingale.

NO-ARBITRAGE APPROACH. The proof that the absence of arbitrage implies that there exists an equivalent measure under which all discounted asset prices are martingales was first given in discrete time by Dalang, Morton & Willinger [1]. The full story for continuous time is much harder; a non-trivial part of the difficulty lies in the fact that the obvious definitions of ‘no arbitrage’ don’t work for various reasons, and the framing of the correct definition is a major part of the work of Delbaen & Schachermayer [2]. We shall here just give the flavour of the proof of the discrete-time result from Rogers [6], which you will see has a lot in common with the ideas of Section 4 above.

Suppose we are in a one-period world, and that there are n assets. Also assume for simplicity that there is no discounting: $R \equiv 1$ in the notation of Section 2. The changes

in the prices of the n assets from the start to the end of the period is denoted by the random vector $X \equiv (X^1, \dots, X^n)^T$, so that if the agent chooses to hold the portfolio $\theta \in \mathbb{R}^n$ during the period, his gain by the end of the period will just be $\eta \equiv \theta \cdot X$. The no-arbitrage assumption is that there is no θ for which $\mathbb{P}[\eta \geq 0] = 1$, and $\mathbb{P}[\eta > 0] > 0$; in words, you cannot make a gain without also facing some risk of losing. Without loss of generality, we can assume that there is no $a \in \mathbb{R}^n$ for which $\mathbb{P}[a \cdot X = 0] = 1$, for then X would lie in a smaller-dimensional subspace, and we could drop down to that subspace and work there.

The theorem says that if the no-arbitrage assumption holds, then there is a measure \mathbb{P}^* equivalent to \mathbb{P} under which

$$\mathbb{E}^*[X] = 0;$$

the random vector X is a vector of martingale differences.

The idea is to consider the moment-generating function

$$a \mapsto \varphi(a) \equiv \mathbb{E} \exp(a \cdot X) \quad (a \in \mathbb{R}^n)$$

which we assume without loss of generality is everywhere finite.⁹

Now consider $\alpha \equiv \inf\{\varphi(a) : a \in \mathbb{R}^n\}$. There are two cases to deal with: *either* this infimum is attained, *or* this infimum is not attained.

In the first case, suppose that the infimum is attained at a^* . Then for any non-zero $\theta \in \mathbb{R}^n$ we shall have

$$\begin{aligned} 0 &\leq \varepsilon^{-1}[\varphi(a^* + \varepsilon\theta) - \varphi(a^*)] \\ &= \varepsilon^{-1}\mathbb{E}[\exp(a^* \cdot X) (\exp(\varepsilon\theta \cdot X) - 1)] \\ &\rightarrow \mathbb{E}[\exp(a^* \cdot X) \theta \cdot X] \end{aligned}$$

and by applying the same argument to $-\theta$ we obtain the conclusion

$$\mathbb{E}[\exp(a^* \cdot X) \theta \cdot X] = 0$$

whatever θ , which is what we were after; the measure whose density with respect to \mathbb{P} is proportional to $\exp(a^* \cdot X)$ will be an equivalent martingale measure.

In the second case, the infimum is not attained, is at least zero (since $\varphi > 0$) and is less than $1 = \varphi(0)$. So there is some sequence of points a_n such that $\varphi(a_n) < \alpha + n^{-1}$. The sequence must be unbounded, otherwise there would be an accumulation point where, by Fatou's lemma, φ would have to be equal to α and the infimum would be attained. By passing to a subsequence, we may suppose that $a_n/|a_n|$ converges to some point γ on the unit sphere. The claim is that

$$\mathbb{P}[\gamma \cdot X > 0] = 0,$$

⁹If not, we would replace \mathbb{P} by the equivalent measure $\tilde{\mathbb{P}}$ defined by $\tilde{\mathbb{P}} = c \cdot \exp(-|X|^2)\mathbb{P}$, for a suitable normalising constant c . For this measure, the moment generating function of X is everywhere finite, and the argument proceeds as given.

for, if not, there would be some neighbourhood U of γ such that $\mathbb{P}[\theta \cdot X > 0] > 0$ for all $\theta \in U$. Since $a_n/|a_n| \in U$ for large enough n , it would follow that $\varphi(a_n) \rightarrow \infty$ as $n \rightarrow \infty$, a contradiction. Thus $\mathbb{P}[\gamma \cdot X > 0] = 0$, and so (since X does not lie in a proper subspace) it has to be that $\mathbb{P}[\gamma \cdot X < 0] > 0$, and investing in the portfolio $-\gamma$ gives us an arbitrage. So under the assumption of no arbitrage, the second possibility cannot in fact occur.

References

- [1] DALANG, R.C., MORTON, A., and WILLINGER, W. Equivalent martingale measures and no-arbitrage in stochastic securities market models. *Stochastics and Stochastics Reports* **29**, 185–201, 1990.
- [2] DELBAEN, F., and SCHACHERMAYER, W. The fundamental theorem of asset pricing for unbounded stochastic processes. Preprint, 1997.
- [3] HARRISON, J.M., and KREPS, D.M. Martingales and arbitrage in multiperiod securities markets. *J. Econ. Theory* **20**, 381-408 (1979).
- [4] HARRISON, J.M., and PLISKA, S.R. Martingales and stochastic integrals in the theory of continuous trading. *Stoch. Proc. Applications* **11**, 215-260, 1981.
- [5] JACKA, S.D. A martingale representation result and an application to incomplete financial markets. *Math. Finance* **2**, 239–250, 1992.
- [6] ROGERS, L.C.G. Equivalent martingale measures and no-arbitrage. *Stochastics and Stochastics Reports* **51**, 41–49, 1994.
- [7] ROGERS, L.C.G., and SATCHELL, S.E. Does the behaviour of the asset tell us anything about the option price formula? A cautionary tale. Preprint, 1996.
- [8] WILLIAMS, D. *Probability with Martingales*, Cambridge Univ. Press, 1991.