Financial Engineering: New Approaches to Managing Risk Exposure

Stewart D Hodges

Director
Financial Options Research Centre
University of Warwick

Chapter 6 in *New Issues in Financial Services*
Blackwell, 1992, Ed: Ray Kinsella

March 1989

Financial Options Research Centre
School of Industrial and Busines Studies
University of Warwick
Coventry
CV4 7AL
Phone: 0203 523606

FORC Preprint: 89/6

6

Financial Engineering: New Approaches to Managing Risk Exposure

Stewart D. Hodges

We are in the middle of a revolution. A quiet revolution, but a revolution nevertheless. It concerns how financial risk is managed. How it is managed by companies, by investment funds and particularly by financial institutions. I shall describe a number of features of this revolution. This will include a brief introduction to the markets and analytical methods involved, and a rather more detailed description of some key ideas in managing money through time, in hedging interest rate risks, and in coping with market imperfections such as transactions costs and thin markets. By way of introduction, let us first take a look at some recent advertising by the industry.

Figure 6.1 is an advertisement published by Refco: 'Risk is everywhere. That's why you need Refco.' This kind of advertising would have been almost inconceivable in the UK only five or ten years ago. Refco is advertising itself as a world leader in 'financial risk management through the use of futures and options'. These few words describe the essence of the new approach to managing financial risk. An important component of this revolution concerns the development of markets in new securities called financial options, and in other derivative instruments, and in the use of these securities. This is what this chapter is about. Futures,

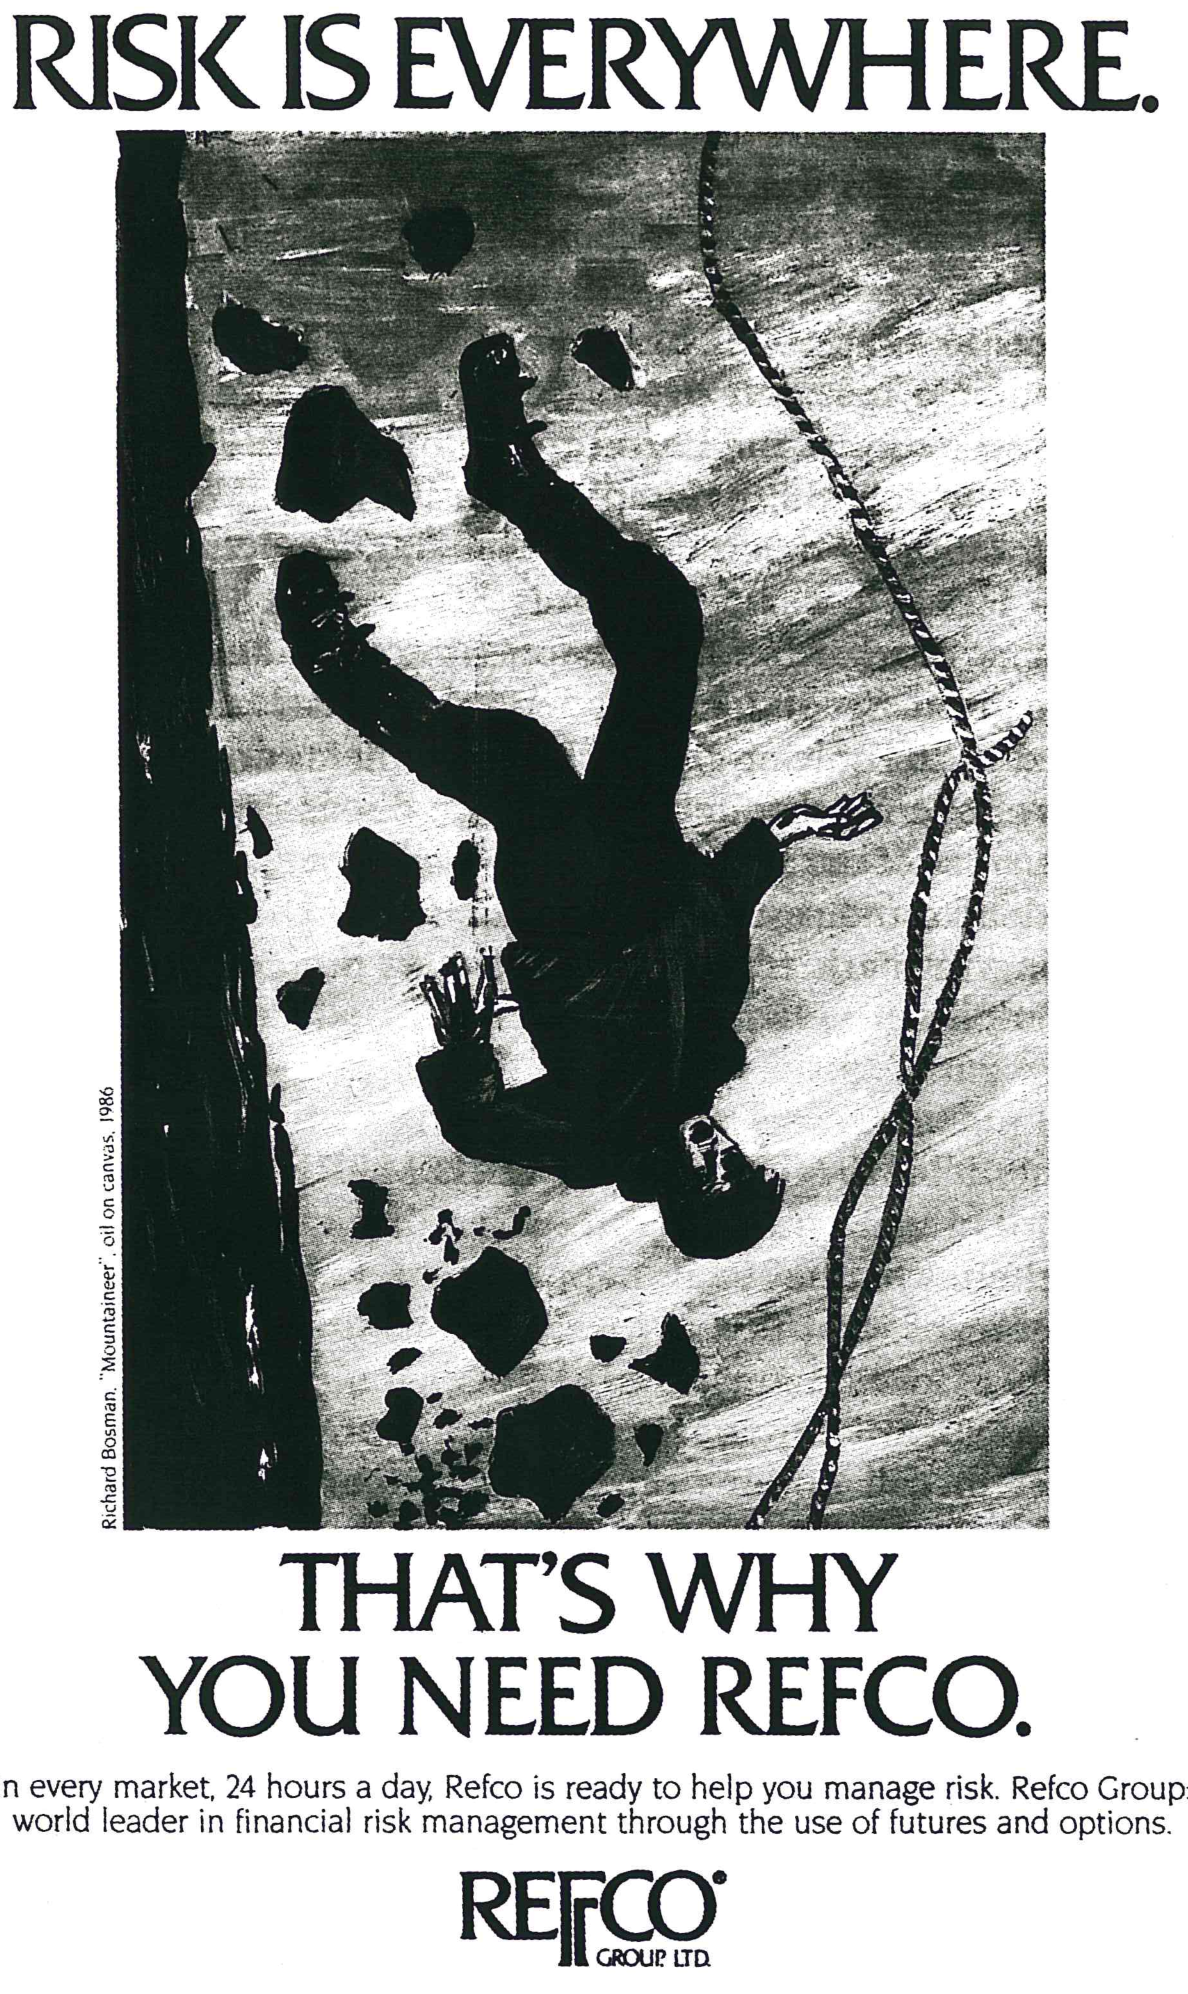

Figure 6.1 Financial risk management advertising

Source: Richard Bosman, *Mountaineer*, oil on canvas, 1986. *The Economist*, 29 May although important, have been highlighted elsewhere (for example, in Merton Miller's (1986) presidential address to the Western Finance Association) and they will not be discussed here and neither will there be much about related markets in swaps.

Academic research has played an important role in facilitating many of the current developments. This is also a field in which there is an increasing amount of direct collaboration between the financial and academic communities. I shall cover three kinds of material. First, a brief summary of what has been happening to financial markets will be given. Second, key academic insights that are directly related to the solving of valuation problems and managing funds efficiently through time will be discussed. Finally, two of the areas of current research at the University of Warwick – models of interest rate movements, and the problems of hedging in markets where transactions are costly – will be covered.

THE NEW FINANCIAL MARKETS

First then, what has been happening out there? Let us find out more of what this innovation consists. The world-wide growth of options markets has been quite dramatic. The Chicago Board Options Exchange (CBOE) established a market in call options on 16 stocks for the first time in 1973. Prior to this, option contracts could be negotiated between parties, but were not standardized or traded in a market place.

The CBOE was an immediate success. Within the year 1.1 million contracts were traded and another 16 stocks were added. Growth continued almost exponentially. In 1977 put options were added and CBOE volume reached 24.8 million contracts. (A put option conveys the right to sell stock at a fixed price whereas a call option gives the right to buy at a fixed price.)

Further exchanges began trading options, including in

1978 the European Options Exchange (in Amsterdam) and the London Stock Exchange. Options on foreign currency, interest rates and bonds first appeared in 1982, and on stock market indices in 1986. These options have particular importance for risk management and 1986 turnover reached 17.3 million contracts on the Chicago Board of Trade treasury bond option, 2.2 million contracts on the Chicago Mercantial Exchange DM option and 1.9 million contracts on the same exchange's option on the Standard and Poor's 500 stock market index. By now options are traded on at least 14 separate exchanges world-wide, and volume continues to increase on virtually all exchanges.

Given this dramatic growth in their use, we need to ask what is special about options. A call option conveys the right but not the obligation to buy a financial instrument at a fixed price before (or on) an agreed date. This right without obligation is what creates their characteristic crooked (dog-leg) payoff pattern. This is what makes them unique. Options differ from all other financial instruments in the patterns of risk they produce, and this is what makes them attractive. The 1986 Bank for International Settlements Report, for example, describes how active trading in foreign exchange and interest rate options surged in the early 1980s, 'spurred by growth in customer demand'. Equally though, it seems most unlikely that this level of demand could ever have been met had we not had new analytic techniques for valuing and hedging.

THE NEW VALUATION METHODS

Over several decades, the problem of 'how should options be valued?' was an unsolved one for academics and practitioners alike. The conventional approach to valuing securities with uncertain payoffs consists of the following four steps:

1 work out the probability of future payoffs;
2 find their expected value;

3 decide what rate of return investors require (on average);
4 use this to calculate the value of the security today.

For many years this approach could not be satisfactorily applied to valuing options because we had insufficient knowledge to complete either steps (1) or (3). We did not know how to describe the probable distribution of future payoffs, or how to determine the return expected by investors in options. By 1965 substantial work had been done to enable us to understand much about the probability distributions of security prices in competitive markets (from their serially uncorrelated behaviour) and also about their required rates of return. These ideas can now be found in most textbooks on corporate finance, (see, for example, Brealey and Myers, 1988 or Higson, 1985). Unfortunately this did not help too much, for the risk of an option changes (and so too does its required rate of return), as the underlying share price goes up and down. It was therefore not until 1973 that the breakthrough by Black and Scholes heralded the beginning of a new era, and the use of an entirely new and powerful methodology.

A CASINO EXAMPLE

The concept behind their approach is stunningly simple. By continuously adjusting our holdings of existing securities, we can manufacture entirely new ones. If we know how to manufacture a security, then we know how much it costs to manufacture and how much it is worth. By way of illustration, let us consider an analogy drawn from a betting casino. Suppose I take \$80 to the casino and proceed to make four 50–50 bets (for simplicity, we assume no bias in favour of the casino). One obvious strategy might be to bet \$20 each time. The outcomes from this range between

The concept
By continuously adjusting our holdings of existing securities, we can manufacture entirely new ones.

Example   A casino analogy
Starting with $80, I make four 50–50 bets

(1) Bet $20 each time

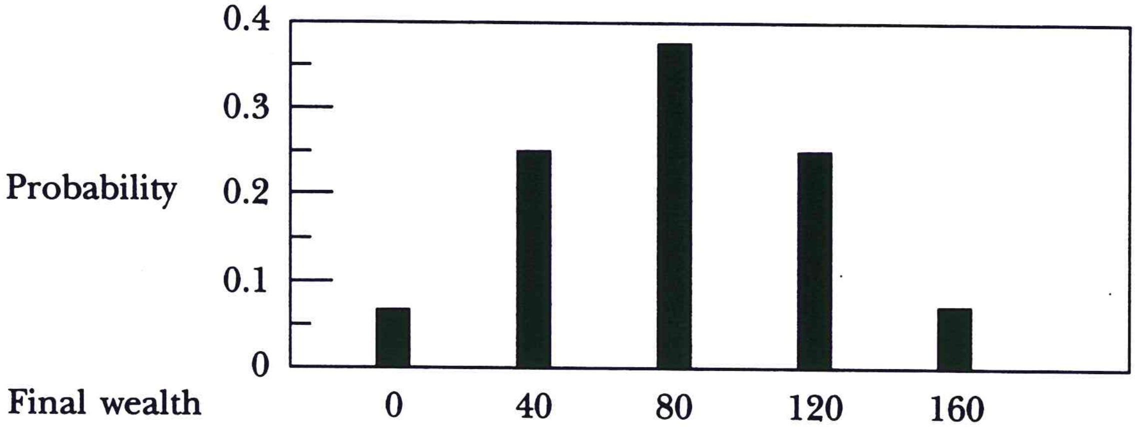

(2) Bet $20 + $5 × (Wins – losses)

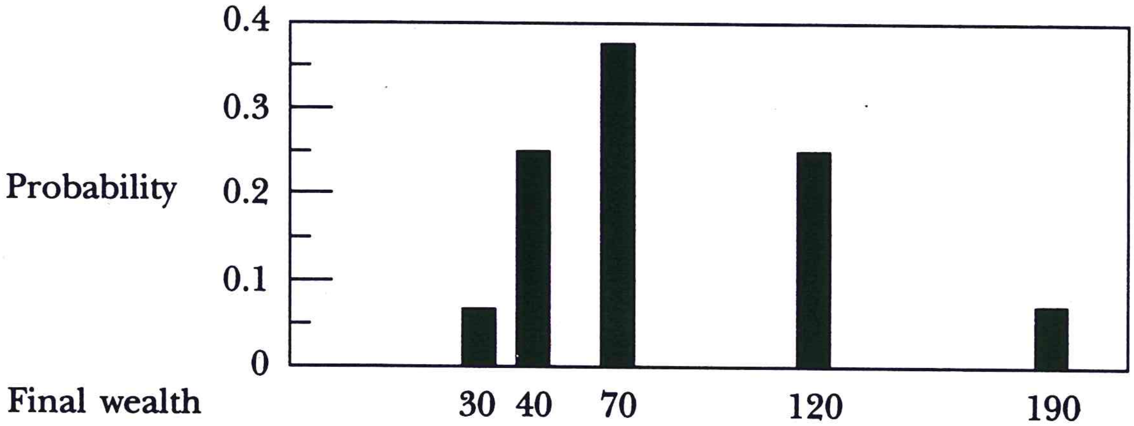

Figure 6.2 Option pricing theory 1983

ending up with $160 and ending up with nothing. The probabilities of different outcomes are illustrated graphically in figure 6.2. This figure also shows that a quite different probability distribution is obtained if I decide to follow a second more subtle strategy of betting an extra $5 every time I win, and betting $5 less every time I lose. This time I end up with a distribution which, instead of being symmetric, is skewed to the right. If I am unlucky, and lose repeatedly, I will only have lost $20 + $15 + $10 + $5 = $50;

so I will still have $30 remaining. If I win repeatedly, I will have won $20 + $25 + $30 + $35 = $110, so I will end up with $190. Nevertheless, since all the bets are 50–50 ones my expected outcome is still to end up with $80.

This illustrates the concept of how new payoff profiles can be constructed by allowing the investments made (bets placed) to vary through time in a way which depends on what has already occurred. What Black and Scholes showed in 1973 was that, by borrowing an amount B and buying H shares (with B and H managed continuously through time) we can precisely manufacture the outcomes of a call option. The value of a call is therefore the cost of the net investment: H times the share price, minus the amount of borrowing, B. The number of shares, H, to be bought is called either the hedge ratio of the option, or its delta. While the concept behind this is simple, the mathematics involved is not. Some quite complicated partial differential equations have to be solved of a kind which also occur in mathematical physics in the context of heat conduction.

SUBSEQUENT DEVELOPMENTS

The approach adopted by Black and Scholes has subsequently spurred a rapid growth in related theoretical work, the analysis of new applications of the technique, and empirical studies of security pricing. Cox and Rubinstein's (1983) textbook lists nearly 300 references, most of them subsequent to Black and Scholes and many of them both important and difficult. Many financial institutions employ scholars who can understand this work and indeed also contribute to it. The institutions' motives are not usually philanthropic! Considerable work has been done to study the traded options markets, and to deal with the problems posed by dividends and the possibility of options being exercised prior to their expiry date. A range of other securities have been studied using the 'continuous-time' methodology of option valuation. These include warrants, convertible bonds, bonds carrying early redemption provisions, corporate debt with a risk of default (and with or without protective covenants), and finally even straightforward government bonds. Equally important are over-the-counter types of risk management products and services such as interest rate caps, floors and collars, and portfolio insurance. Financial institutions have successfully devised, valued and sold new kinds of products, such as 'down and out' options and 'look-back' options. The exploitation of physical assets (such as a gold mine or oil well), the value of underwriting and loan guarantees have also been investigated within this new framework: there are people within the oil companies who specialize in these techniques. We will now look at an example of financial engineering in action.

Chase Manhattan Guaranteed Market Investment Funds

Figure 6.3 shows part of the prospectus for a 'Guaranteed Market Investment' (GMI) fund offered by Chase Manhattan Bank. The investor may choose a particular investment horizon and specify the minimum percentage return required over that period. The fund will then quote an 'index-multiplier' and it will guarantee the investor which ever is the larger of either the minimum percentage return, or the 'index-multiplier' times the return on the market index over the same period. For example, suppose the investor requires at least 6 per cent over a two-year period. The 'index-multiplier' is 70 per cent, and at the end of two years the fund will pay the greater of 6 per cent or 70 per cent of the index return on top of the original investment.

How do they do it? If we draw a diagram of the payoff provided against the future value of the index, we find that from the investor's point of view it amounts to that of a two-year bond plus (70 per cent of) a (two-year European) call option on the market. We can figure out the exercise

The Standard GMI

The Standard GMI pays interest based on a multiple of the S&P 500's percentage *appreciation* from the first to last day of the deposit term. But even if the market falls, Chase preserves your principle in full.

The Standard GMI appeals to investors who seek growth potential without the risk associated with equity investments. It also appeals to equity managers willing to relinquish a degree of upside potential in exchange for full downside protection. Viewed from this perspective, the Standard GMI exhibits characteristics similar to those of a hedged equity portfolio. But it adds one significant breakthrough – the GMI performance parameters are pre-set and assured as a Chase obligation.

The Reverse GMI

The Reverse GMI differs from the Standard GMI by tying interest payments to stock market *depreciation.*

The following tables illustrate Index Multipliers as of July 6, 1987 for a selection of Guaranteed Minimum Interest rates and the standard maturities offered. Your actual Index Multiplier will be determined on the day your account is opened.

Index Multipliers

The standard GMI account

	Guaranteed Minimum Interest							
Maturity	0%	2%	4%	6%	8%	10%	12%	14%
36 mos	101	97	94	90	86	82	77	72
30 mos	94	90	86	81	76	71	64	57
24 mos	86	81	76	70	63	55	45	31
18 mos	77	71	63	55	44	28	–	–
12 mos	62	52	40	21	–	–	–	–
6 mos	47	30	–	–	–	–	–	–

The reverse GMI account

	Guaranteed Minimum Interest							
Maturity	0%	2%	4%	6%	8%	10%	12%	14%
36 mos	303	290	275	259	243	227	210	193
30 mos	249	234	218	202	180	169	152	133
24 mos	197	182	166	149	132	113	92	63
18 mos	153	136	119	101	80	52	–	–
12 mos	100	83	62	33	–	–	–	–
6 mos	65	40	–	–	–	–	–	–

Figure 6.3 Chase GMIs

Source: Chase Manhattan Bank.

price of the option without too much trouble. Now two-year call options cannot be bought in any market place, but Chase can fairly easily value them and hedge their exposure using the modern techniques. Potential problems come from the difficulty of forecasting the future variability of the index, and the transactions costs involved in adjusting hedge positions.

The prospectus also describes another related product, the 'Reverse Guaranteed Market Investment' fund. The reader may enjoy puzzling out how that one works!

These ways of managing money come under the label of portfolio insurance schemes. The adjustments in market exposure required to implement them are often accomplished by means of the futures markets. Let us take a quick look at some of the theory about right and wrong ways of managing money through time.

PATH INDEPENDENCE AND OTHER FUNDAMENTAL IDEAS

The analytic methods used in financial engineering have a generality that reaches far beyond the valuation of straightforward European call options or put options. We can understand some of this generality by considering a simplified problem of managing an investment fund through time. We will suppose that our funds must be divided between a single risk free instrument and an equity market index. At each date in time we simply need to decide how much to bet on the index, and the remainder will be invested at the risk free rate. It turns out that (with frictionless continuous trading) any strategy which is rational must satisfy a path independence condition with respect to the level of the index. In other words, the value of the fund (if properly managed) must at any date in the future only depend on the value of the index at that date, and not on what path it took to get there! A derivation of this result is provided in the Appendix to this paper. It is originally due to Cox and Leland (1982). This rather surprising path independence property also implies that the fund values satisfy simple averaging relationships through time (of the kind used in the binomial method of option pricing of Cox, Ross and Rubinstein, 1979). The optimal amount to hold in the index also satisfies these same relationships. It is these averaging relationships which therefore govern all problems in financial engineering and which are exploited to find solutions to a wide variety of valuation and hedging problems. The averaging I have described is exactly equivalent to the partial differential equation that was solved by Black and Scholes. It characterizes not only the behaviour of option values, but also the single method of managing money through time that cannot be shown to be inefficient.

Often, the endeavours of financial economists have tended to suggest that the actions of financial managers are irrelevant. That is certainly not the case here. If there are infinitely many portfolio strategies which satisfy the fundamental equations for path independence, there are even more which do not. Stop-loss and dollar-averaging strategies are well known examples of policies that can be shown to be inefficient. Policies of repeatedly rolling over short-term portfolio insurance can be even more devastating. Work by Dybvig (1988a; 1988b) demonstrates that we can now precisely quantify how much money is wasted by any particular strategy that is inefficient. This is accomplished by calculating what terminal distribution of wealth it leads to, and then how much more cheaply the same distribution could be obtained by a preferred strategy.

CURRENT RESEARCH

Currently research is being pursued within the larger institutions and, more particularly, in a number of centres of excellence in UK universities. Among the latter, the Financial Options Research Centre, based at Warwick University, is playing a leading role. An important strand of its work, supported by the ESRC, is focused on the applications of options-based approaches to the problems associated with financial valuation. More specifically, this has four main elements:

1 A survey and taxonomy of option pricing models according to the nature of each application. Some earlier surveys exist, but there is no up-to-date survey. No clear taxonomy has yet appeared in the literature.
2 The formal mathematical structures of the model, and the way these too may be classified.
3 Except for the simplest models, there is a great deal of confusion about numerical methods. Techniques in use include the binomial method, finite difference methods, and numerical integration. Rather little work has been done to compare the relative efficiency of different approaches for solving different problems. There will be clear guidelines concerning numerical methods.
4 Finally, sufficient empirical work is being undertaken to demonstrate the viability of the modelling approaches to make some comparisons between models.

There has been much interest in this work in City institutions. Most seem relatively comfortable in valuing simple equity of foreign exchange (FX) options. Their expertise varies greatly for the more esoteric instruments. Many institutions are interested in the rapidly developing topic of interest rate models, applied to hedging interest rate risk, or to the valuation and hedging of derivatives based on interest rates or bond prices. The broader field of risk control and hedging generally is also one where access to current research is regarded as important.

Much work is currently being carried out on both the problems of how interest rates move and can be hedged (leading to the valuation of bond options) and of hedging in markets where transactions costs are significant or where there may be illiquidity.

GOVERNMENT BONDS

The notion of using the duration of a bond to measure and hedge interest rate risk goes back at least to Macauley (1938) and Redington (1952). This concept is still an important one, but today we need more refined techniques than this particularly when interest rate derivatives are involved. Curiously, even 'straight' bonds involve elements of options. If we form a portfolio of a one-year and a 20-year bond it will go up a lot if interest rates fall, but only fall slightly if rates increase. Such a portfolio can be constructed so that it appears to dominate a single 5-year bond. Of course, it will not really dominate, for the yield of the 5-year bond will be bid up to compensate for its different risk profile. A rapidly developing literature here is helping us to understand how these effects work and the exact nature of the mechanisms in these markets.

The basics concepts are essentially the same as we have already encountered. Just as we can manufacture options by means of dynamic trading strategies, so too can we manufacture bond type instruments. Any real bond that gets out of line with the cost of a manufactured one can, in principle, be arbitraged. Vasicek (1977) showed that the assumption of a particular process for the short term interest rate (and a term structure entirely dependent on that rate) led to a no-arbitrage characterization of the entire term structure. Cox, Ingersoll and Ross (1985) embedded this within a more general structure, and also developed an equally simple model for a slightly different interest rate process. Although Brown and Dybvig (1986) demonstrated that the Cox, Ingersoll and Ross model could even be used to measure the market's view of the uncertainty of future interest rates, this class of models had a severe limitation: the models could not easily be made to exactly reflect the actual form of the current term structure. A paper by Ho and Lee (1986) illustrated that this problem could be overcome, and researching good financial engineering solutions to this problem is now a 'hot' topic. At Warwick, Andrew Carverhill and Jim Steeley have made both theoretical and empirical contributions to our knowledge in this area. Carverhill (1988) provides a generalization of Ho and Lee, which is a precursor to further extensions with empirical content. Steeley (1989) provides an improved method of estimating the term structure of interest rates. Figure 6.4 reproduces some of his estimates for the UK gilts market.

COPING WITH MARKET IMPERFECTIONS

One final area of research relates to the practical problems in hedging risk. Some kinds of risk can be hedged quite cheaply and easily. Others may be much more difficult to manage. The usual valuation of call options is based on the notion that it is possible to manufacture call options synthetically by trading continuously at no cost. This is certainly an optimistic picture, and particularly so when we realize that the replication strategy would actually involve an infinite amount of turnover. (We can show that turnover increases in proportion to the square root of the number of rebalancing intervals).

Leland (1985) presents and describes properties of a method for hedging call options when, in addition to the usual assumptions of Black and Scholes, there is a proportional transactions cost in trading the risky asset. However, this method is in no sense an optimal one. Other important papers by Davis (1988), Davis and Norman (1988) and Taksar, Klass and Assaf (1988) describe optimal portfolio policies to maximize expected utility over an infinite horizon. These papers generalize earlier seminal work of Merton (1971) to a world which includes transactions costs, but they are not really concerned with the 'financial engineering' type problems of replicating or hedging. However, these papers do contain the key to what is needed.

Many 'financial engineering' problems can be formulated as what are known as stochastic optimal control problems.

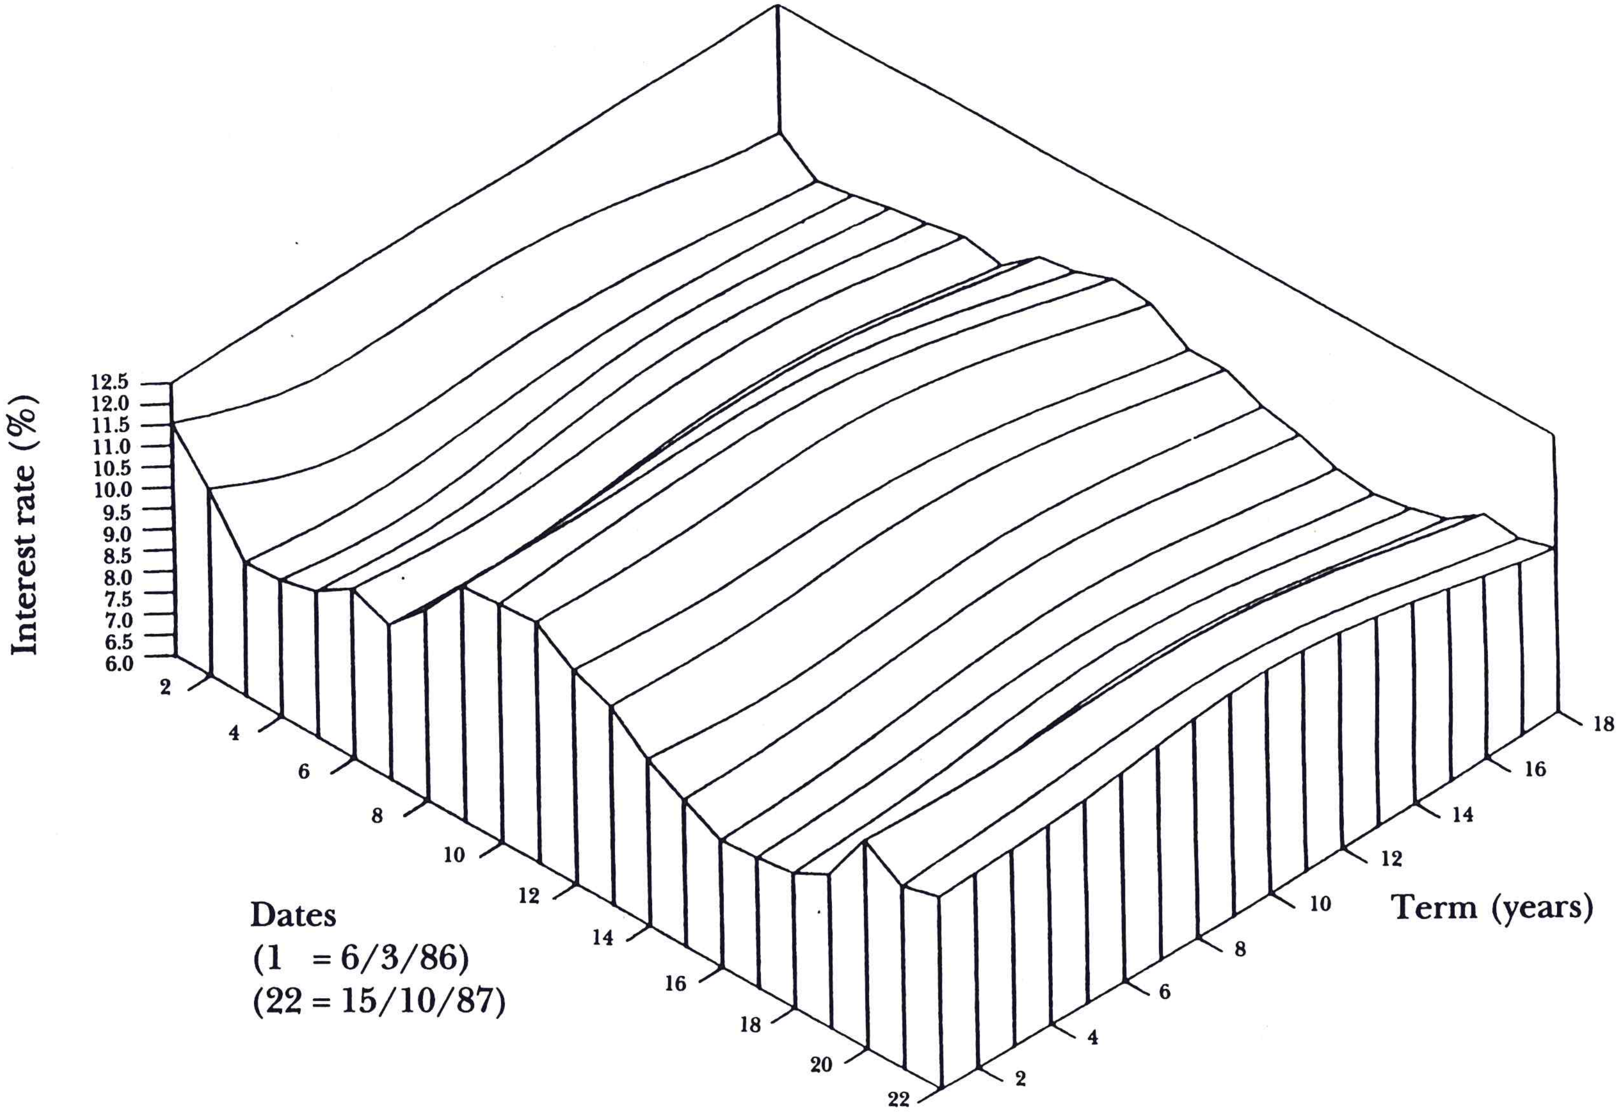

Figure 6.4 Time path of the spot rate curve

Fleming (1969) provides a useful survey article of these techniques. Solutions can be characterized by a dynamic programming (Bellman-Hamilton-Jacobi) equation, and which can be solved recursively from the horizon date backwards (as, for example, we do for the option pricing binomial procedure). This approach is a particularly important one for dealing with problems such as those posed by transactions costs or market incompleteness.

In replicating a contingent claim under transactions costs (or in undertaking almost any other kind of hedging under costs or incompleteness) we can no longer require exact replication. It is therefore necessary to define a loss function, and it is important to choose one that is appropriate. We can make some general statements about the properties of optimal policies: usually it is appropriate to define control bands and transact only to remain within those bands. Leland's (1985) policy is severely suboptimal because he always revises his hedge ratio right back to the 'ideal' one, instead of to within a band.

Exact solutions to these kinds of problems are often rather nasty to compute, as there will usually be at least two state variables in the dynamic programming optimization. However, we can learn from studying the solutions to simple problems, and with modern computing techniques we can cope with increasing degrees of realism and complexity. A more detailed account of this work is contained in Hodges and Neuberger (1989).

CONCLUSIONS

I have tried to convey the essence of some of the key strands which characterize the financial engineering revolution we are witnessing. More than any other single factor, it is undoubtedly the product of the advances in computing and communications. It has created a financial community in which many practitioners are also academics and many academics are also practitioners. Some have criticized the results of these innovations as being destabilizing. Only time will reveal whether they are right. Whatever further reforms may be made to the markets themselves, financial engineering is here to stay.

APPENDIX: FUNDAMENTAL PROPERTIES

METHOD OF ANALYSIS

Maximize expected utility of terminal wealth, $E[U(w_T)]$.

For simplicity we shall consider

1 Additive gambles with payoffs ±1.
2 Zero interest rate.
3 Constant probabilities

p of +1
q of −1

These all generalize, and we can also have

$$U = U(w_T, S_T) \text{ where } S_T = \tilde{r}_1 + \tilde{r}_2 + \ldots \tilde{r}_T$$

(e.g. final value of underlying)

1 Path independence result

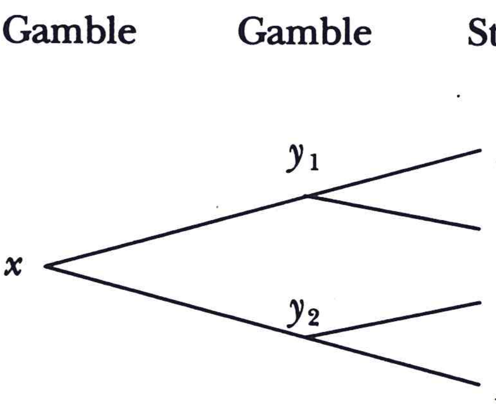

State	Probability	Wealth
A	p^2	$k + x + y_1$
B	pq	$k + x - y_1$
C	pq	$k - x + y_2$
D	q^2	$k - x - y_2$

The maximization problem is equivalent to:

Maximize $E[U] = p^2U(x + y_1) + pqU(x - y_1) + pqU(-x + y_2) + q^2U(-x + y)$

$$\text{This gives } \frac{\partial E}{\partial x} = p^2U'_A + pqU'_B - pqU'_C - q^2U'_D = 0 \qquad (1)$$

$$\frac{\partial E}{\partial y_1} = p^2U'_A - pqU'_B \qquad = 0 \qquad (2)$$

$$\frac{\partial E}{\partial y_2} = \qquad pqU'_C - q^2U'_D = 0 \qquad (3)$$

Hence $U'_A = \frac{q}{p} U'_B, U'_C = \frac{q}{p} U'_D$ and from (1):

$U'_B = U'_C$ which implies $w_B = w_C$

□

2 Averaging

$$k + x - y_1 = w_B = w_C = k - x + y_2$$

Hence, $$x = \tfrac{1}{2}(y_1 + y_2)$$

Similarly for w itself

3 The heat equation

At each time step δ, $S \to S \pm \sigma\sqrt{\delta}$ $w = w(S, t)$

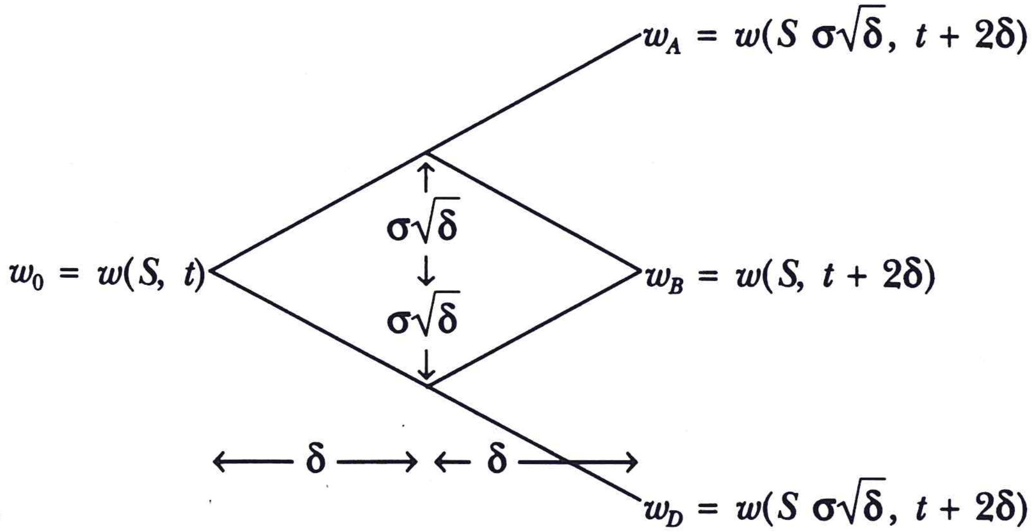

$$\frac{\partial^2 w}{\partial S^2} \simeq \frac{w_A - 2w_B + w_D}{\left(2\sigma\sqrt{\delta}\right)^2}$$

$$= \frac{w_A + 2w_B + w_D - 4w_B}{4\sigma^2\delta} = \frac{w_o - w_B}{\sigma^2\delta}$$

$$\frac{\partial w}{\partial t} \simeq \frac{w_B - w_0}{(2\delta)} \simeq -\frac{\sigma^2}{2}\frac{\partial^2 w}{\delta\sigma^2}$$

References

Bank for International Settlements 1986: *Recent Innovations in International Banking.* (Chairman Mr S. Y. Cross.)

Black, F. and Scholes, M. 1973: The pricing of options and corporate liabilities. *Journal of Political Economy,* 81, 637–55.

Brealey, R. A. and Myers, S. C. 1988: *Principles of Corporate Finance,* New York: McGraw-Hill.

Brown, S. and Dybvig, P. H. 1986: The empirical implications of the Cox, Ingersoll, Ross theory of the term structure of interest rates. *Journal of Finance,* 41, 617–30.

Carverhill, A. P. 1988: The Ho and Lee term structure theory: a continuous time version. Financial Options Research Centre Preprint 88/5, University of Warwick.

Cox, J. C. and Leland, H. 1982: On dynamic investment strategies. *Proceedings, Seminar on the Analysis of Security Prices* 26, No. 2, Centre for Research in Security Prices, University of Chicago, Chicago, IL.

— and Rubinstein, M. 1983: *Options Markets.* Englewood Cliffs, NJ: Prentice Hall.

—, Ross, S. A. and Rubinstein, M. 1979: Option pricing: a simplified approach. *Journal of Financial Economics,* 7, 229–63.

—, Ingersoll, J. E. and Ross, S. A. 1985: A theory of the term structure of interest rates. *Econometrica,* 53, 385–407.

Davis, M. H. A. 1988: Local time on the Stock Exchange. In J. R. Norris (ed.), *Stochastic Calculus in Application,* Pitman Research Notes in Mathematics, London.

— and Norman, A. R. 1988: Portfolio selection with transactions' costs. Working Paper, London: Imperial College.

Dybvig, P. H. 1988a: Inefficient dynamic portfolio strategies, or how to throw away a million dollars in the Stock Market. *Review of Financial Studies*, 1(1), 67–88.

— 1988b: Distributional analysis of portfolio choice. *Journal of Business*, 61, 369–93.

Fleming, W. H. 1969: Optimal continuous-parameter stochastic control. *SIAM Review*, 11, 470–509.

Higson, C. 1985: Business Finance, London: Butterworth.

Ho, T. S. and Lee, S. 1986: Term structure movements and the pricing of interest rate contingent claims. *Journal of Finance*, 41, 1011–29.

Hodges, S. D. and Neuberger, A. 1989: Optimal replication of contingent claims under transactions costs. *Review of Futures Markets*, 8(2), 222–39.

Leland, H. E. 1985: Option pricing and replication with transaction costs. *Journal Of Finance*, 1283–301.

Macaulay, F. R. 1938: *Some theoretical problems suggested by the movements of interest rates, bond yields and stock prices in the US since 1856*. National Bureau of Economic Research.

Merton, R. C. 1971: Optimum consumption and portfolio rules in a continuous time model. *Journal of Economic Theory*, 373–413.

Miller, M. H. 1986: Financial innovation: the last twenty years and the next. *Journal of Financial and Quantitative Analysis*, 21, 459–71.

Redington, F. M. 1952: Review of the principles of life office valuations. *Journal of the Institute of Actuaries*, 78, 286–340.

Steeley, J. 1991: Estimating the term structure of interest rates. Financial Options Research Centre, Preprint 89/1, University of Warwick, *Journal of Business, Finance and Accounting*, 18, 513–29.

Taksar, M., Klass, M. J. and Assaf, D. 1988: A diffusion model for optimal portfolio selection in the presence of brokerage fees. *Mathematics of Operations Research*, 13(2), 277–94.

Vasicek, O. A. 1977: An equilibrium characterisation of the term structure. *Journal of Financial Economics*. 5, 177–88.