PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Dynamic Asset Allocation: Insights from Theory [and Discussion]

Stewart D. Hodges, R. S. Clarkson and M. A. H. Dempster

Phil. Trans. R. Soc. Lond. A 1994 **347**, 587-598
doi: 10.1098/rsta.1994.0068

Dynamic asset allocation: insights from theory

By Stewart D. Hodges

Financial Options Research Centre, Warwick Business School, University of Warwick, Coventry CV4 7AL, U.K.

This paper provides a survey of the now considerable academic theory relating to the practice of dynamic asset allocation. This work is scattered through the literature and many of the key ideas are not as accessible or well known as they deserve to be.

The paper begins by providing a definition of what is meant by dynamic asset allocation and a description of its most significant features. Next it develops the concept of path independence and its relationship to efficient diversification through time. It is shown that this principle also applies to funds whose performance is appraised relative to an index benchmark. The final sections of the paper describe the implications of recent work on market equilibrium and on performance measurement.

1. Introduction

There is now a substantial body of academic theory relating to the issues involved in dynamic asset allocation. However, this work is rather scattered through the literature and not very easy to find. This paper provides a survey of key ideas in the area. A number of these are not as accessible or as well known as they deserve to be.

The paper begins by providing a definition of what we mean by dynamic asset allocation and some of the most significant features of it. Section 3 discusses the concept of diversification through time. This is a very much less familiar concept than that of diversification across securities. Many funds measure performance relative to benchmarks, so §4 discusses problems related to setting objectives in terms of index benchmarks. Section 5 considers the role of market regularities and equilibrium including the use of the role of forecasts and tactical asset allocation. Finally in §6 we look at issues related to performance measurement.

2. Horizon distributions and contingent pay-offs

It is convenient to start from the definition of dynamic asset allocation given by Trippi & Harrif (1991). They define dynamic asset allocation as 'a class of investment strategies that shifts the content of portfolios between two or more asset classes in response either to changes in the value of the portfolio and/or external economic states, on a more or less continual basis'. The motivation is two-fold: first, to tailor the distribution of fund return at some future date so that it can be an entirely different shape from that of the market index. It can be skewed to the right or skewed to the left or tailored in other more ingenious ways. Second, it may be to exploit predictable regularities, which include market timing and other tactical allocation strategies.

In 1973, Black & Scholes discovered that options on assets can be replicated, and hence valued by shifting the content of a portfolio dynamically between two or more

Phil. Trans. R. Soc. Lond. A (1994) **347**, 587–598
Printed in Great Britain

asset classes. A large body of work on option theory has resulted and it is to this that we must turn in order to understand what dynamic asset allocation strategies are all about. Option theory tells us how to value and synthesize a given contingent profile. The Black–Scholes formula tells us the value of the simplest of contingent profiles, the pay-offs from a call option, for example on a market index. Black–Scholes theory tells us not only how much money we need if we want to construct one of those, but also how to construct it. Provided the assumptions of Black & Scholes hold, the delta hedging strategy will exactly synthesize the contingent pay-offs of the option.

If we have a more complicated contingent profile we can always think of it as being built up in terms of a series of call options. Suppose we can approximate the profile by a series of straight lines. Then corresponding to the intercept we need an investment in that amount of zero coupon bonds. We will then achieve this intercept value even if the index falls to zero. To get the correct slope at this point, we need a corresponding amount in stocks. Finally, every time we have a change in the slope, the amount by which the slope changes tells us how many extra call options we have to have at that particular strike price. Thus any piecewise linear profile is very easy to build up in terms of call options which are easy to value and we know how to replicate using the Black–Scholes theory.

For any fixed investment horizon, the profile of the fund value contingent on possible different values of the market index is exactly equivalent to the probability distribution of the fund value. We can always think in terms of mapping backwards or forwards between a probability distribution and an equivalent contingent profile. To do this we first need to know what the probability distribution of the index is. Provided we know the probability distribution of the index, then given the fund value contingent on the index, it is straightforward to obtain the distribution of the fund value. For example, suppose we want to find what value we are going to exceed, say 20% of the time. We would work out the index value which will be exceeded 20% of the time, and then find the fund value corresponding to it. This gives the fund value that will be exceeded 20% of the time. (This is obviously true, provided the higher fund values correspond to higher index values. We shall see shortly this is usually necessary for efficiency.) So it is easy to map from contingent fund values to the distribution for the fund value.

Equally we can go back the other way. If you draw a density function that you like for the fund value, we can just work backwards the other way. Again, we would take some probability, say 20%, locate the fund value we are going to exceed 20% of the time and also the index value we are going to exceed 20% of the time, and then plot that pair of points. Repeating this for all the different probabilities, it builds up the curve of fund values as a function of index ones. So we can always think either in terms of the distribution for the fund value or how the fund value is contingent on the index, and we can map either way. We shall see shortly that this type of framework, at least under simplifying assumptions, is also a condition for the asset allocation rule to be an efficient rule.

3. Efficient diversification through time

This section develops the concept of efficient diversification through time. Some kinds of asset allocation rules essentially throw away money. If you do not want to throw away money, it is important to understand which kinds of rules are efficient and which are not. To get some intuition of how it is possible for a poor dynamic asset allocation rule to throw away money, we will consider a simple mean–variance example that illustrates what can go wrong. The exact relations that need to be followed will be formalized later.

The idea of diversification across securities is now very similar. Here is a very simple numerical example which we will then extend to illustrate time diversification. Suppose we have two securities. Both of them offer an expected risk premium of 10%. Both of them have a standard deviation of 20%, and we will think of this as over a one year period. For simplicity, these securities have returns which are independent. Consider two portfolios, A and B. For Portfolio A we put £100 in security 1 and any remaining wealth is invested in cash. The risk premium we earn is then £10 and the standard deviation is £20. No-one would be silly enough to do that because we know that it is more efficient to diversify. As the two securities are identical the best thing to do is to invest equal amounts in them. A better strategy would be to invest the same amount, £70 say, in each of these two securities, and, again the rest in cash. We shall call this Portfolio B. The variance is now twice 14 squared (from the two securities) which is 392. The standard deviation is £19.8, which is slightly smaller than before, but we have got an extra return premium of £4. We are now earning a £14 risk premium instead of £10 before.

The reason for developing this example in such detail is because we can use exactly the same numbers to tell a different story. Consider an investment in the equity market over two years. In each year the expected risk premium is 10% and the standard deviation is 20%. We know that successive returns are essentially independent so this assumption now looks quite natural. Now see what happens if we have a strategy where we plunge £100 into equity for the first year and then put everything into cash for the second. This corresponds exactly to our Portfolio A: we have got an expected risk premium of £10, and a standard deviation of £20. It would be better to have £70 in equities each year. That way we would still have the same standard deviation for the return at the end of two years but we would have expected to have earned an extra £4 more. Strategies where the manager plunges in and out doing wonderful market timing transactions but actually has no forecasting ability are exactly equivalent to the first very wasteful strategy. You may not realize it is very expensive but we have just shown that it is. Under a market timing strategy of someone who has no forecasting ability, the plunge in and out of equities loses the sort of diversification benefits you would enjoy from a smoother policy.

The magnitudes are significant. If you randomly plunge half the time into cash and half the time into equities, you are giving up 40% of the risk premium that you are earning. You are only getting a £10 risk premium instead of a £14 one. That is an enormous loss.

That is the intuition about why some dynamic allocation strategies are efficient and some throw money away. However, we cannot generally use a mean–variance framework because the pay-offs are too complicated. The mean–variance framework simply helps to develop the intuition. A condition which is often necessary for a dynamic strategy to be efficient, in the sense of not wasting money (the way we wasted £4 on that first strategy) is that the values of the strategy must be path independent. This result, which is not very well known, was first presented by Cox & Leland (1982) in a rather obscure working paper. They show quite rigorously that if we are in a market that satisfies the assumptions for Black–Scholes option pricing (i.e. a constant risk-free rate and constant volatility of the equity market) then a dynamic portfolio strategy must be path independent to be efficient. By efficient we are not referring to mean–variance efficiency, but rather that whatever distribution of outcomes is obtained is purchased as cheaply as possible. (Financial economists refer to this as first order stochastic dominance.) In other words that no money is thrown away.

We describe next what is meant by path independence. Suppose starting in 1992 the index was around 2600. Assume that in 1994 the index is up at 3000 but it got there by one of two paths, either going up to 3600 first or a lower path dropping first to 1900. A path independent strategy is one where the portfolio value in two years time will end up at more or less the same value whichever route the market took to get there. This is an extremely useful result because it means that if we have a path independent strategy then we can characterize it in terms of the contingent profile we described earlier. The very complicated multiperiod problem of what market exposure we should have every day collapses down to a contingent pay-off at a single horizon date. We can rely on option pricing theory to tell us what to do in between.

Thus, the rule for efficiency is that our portfolio value at any future date of interest should just depend on where the index gets to and not on how it got there. If you try and time the market without being able to forecast, that will give a dispersion of fund values at any level of the market. We have seen already that this can be very inefficient. In our example we are losing 4% over the two years, or 200 basis points a year.

Other well-known strategies also suffer from the same problem that they are not path independent and therefore they are ways of throwing money away. A simple one is the stop loss rule. Suppose, for example, I start by investing £100 in the equity market and I am going to leave it there unless the index hits 2000, in which case I will move into cash. If the market reaches 3000 without first falling to 2000 we will have gained 15% and end up with £115. On the other hand, if the market falls below 2000 first, we go down to £77 and then lock into cash, so that at the end we still have £77 plus interest. There is a big divergence between the two ending values even though the market index got back to the same level via a different route. So stop loss strategies are intrinsically wasteful. Dybvig (1988) analyses these various strategies and shows that with a stop loss strategy you could easily lose 80 to 90 basis points a year. That may not seem much money but the risk premium you are getting in the first place is less than 8% per annum, so you are throwing away more than 10% of what you are earning.

A lock-in strategy, where we start off in equities and then move into cash if we hit a level we like, is also path dependent and has a similar level of inefficiency. Finally, another example of a dynamic strategy which Dybvig analyses is to look at repeated short term portfolio insurance. What he does is to look at a strategy of rolling over one year portfolio insurance at the end of each year over a five year period. A lot of funds, particularly in the U.S., have done that sort of thing and that is losing probably about 50 basis points a year.

Dybvig's calculations are not difficult. First you simulate the strategy that a fund manager is using to calculate the probability distribution of future fund values. You then translate this into the corresponding values contingent on the market index. Finally, you use option pricing theory to find out the cheapest way to buy that distribution and then see how much cheaper that is than the amount of money you started with and that led you to the distribution in the first place. The cash saved is a tangible measure of the inefficiency of the original strategy.

To summarize, the rule for efficiency (at least under our assumptions) is that the future fund value should be a non-decreasing function of the future index value. Intuitively it is reasonable that the contingent pay-offs go up because it is cheaper to get money when the index is doing well. Contingent claims that pay off in high index states of the world are cheaper than contingent claims that pay off in low index value states of the world and which provide insurance for those states. So you get as much money as you can when money is cheap and you do not buy so much when it is dear.

I have described a world in which the market index is the only uncertain state variable. This whole theory does generalize into richer assumptions where there are many different state variables, but it gets a bit more messy. We will continue to treat only this slightly simpler framework. We will examine next some other aspects of the robustness of this concept. Clearly it is not entirely robust if we have transactions costs in changing the allocations, though work has been done on this issue (see Hodges & Neuberger 1989). We shall now ask whether the concept is robust if we are thinking of tracking error and whether it is robust if the market mean reverts.

4. Benchmark objectives

If we are interested in tracking error and we set our objectives relative to a benchmark portfolio, then similar results still apply. It is still true that when the index is at a high level, money is cheap and when the index is at a low level, contingent pay-offs are more expensive. We therefore need the profile of our surplus or deficit relative to the benchmark to itself be increasing in the benchmark. The path independence result is still true, so the stop-loss strategy is still inefficient but the way you would modify stop-loss will be different from what you would do under normal risk return criteria.

Benchmarks are hard to justify, except possibly when they represent genuine liabilities to be met. Hodges (1976) compared two criteria: mean–variance tracking error with mean–variance return efficiency using a simple portfolio selection model. As you would expect, each criterion is demonstrably inefficient viewed from the perspective of the other. So if you are pursuing a policy that is mean–variance efficient in tracking error, you are giving up a lot in terms of conventional mean–variance efficiency. This suggests a conflict of interest between the fund manager and the ultimate beneficiary.

Roll (1992) gives a much more formal analysis of optimal tracking error betas. Among other things he shows that optimal tracking error portfolios have betas greater than one, whereas conventional mean–variance efficient portfolios have betas less than one. He suggested constraining beta to a value less than one to enhance mean–variance performance.

This may increase the efficiency somewhat but the result remains sub-optimal. So the bottom line is that tracking error criteria are sub-optimal unless you have a very good reason for saying that that really is what the objective ought to be. If there is a clear liability and you are using the liability as a benchmark, we can say the surplus after we have paid off our liabilities may be a reasonable number on which to use a mean–variance criterion. If you are doing index arbitrage then clearly you are concerned with tracking error. However, it is not a good idea to worry about tracking error just because that is how someone else tots up the points at the end of the year to arrive at a bonus!

5. Market regularities and equilibrium

Most investors seem to like positive skewness. Strategies with pay-offs that are convex from below produce distributions with positive skewness. This is fairly easy to see because a straight line contingent pay-off would essentially be cash plus a static holding in the index and would give a log normal distribution (plus a shift to take account for the cash part). As soon as you add convexity to the contingent pay-off you get a longer right tail and more positive skewness. Conversely, for contingent pay-offs with a sufficient concavity we will get a long left-hand tail and negative skewness. Note also that the slope of the contingent value gives some idea of what the fund exposure is. A portfolio insurance type (convex) strategy gives less exposure at low market levels, whereas a contrarian (concave) strategy gives more exposure at low market levels.

Now the market has to clear, and a paper by Leland (1980) considers who should buy and who should sell portfolio insurance in equilibrium. Suppose we can agree that, other things being equal, most investors prefer positive skewness. In equilibrium we would therefore expect that although people are happy to buy stocks when the market is high, they will need more encouragement to hold them when the market is low. Thus, we would expect higher risk premia at historically low market levels, and low risk premia at historically high ones. Empirical work tends to confirm this view. For example, the variance-ratio tests reported by Lo & McKinley (1988) and Poterba & Summers (1988) suggest mean reversion. Fama & French (1988) have also looked at dividend yields as a way of predicting expected returns and again confirm the same kind of effect. Not all the work is terribly significant statistically. There is a suggestion of possible mean reversion, at or about a five year period, but we do not have long enough data sets to get very good statistical significance over this kind of horizon. However, there is now a large number of studies, most of which tend to confirm this view of the world.

It is interesting how little attention has been paid until recently to the way the market risk premium evolves through time. Even the better finance texts seem not to question whether the market risk premium is likely to be constant through time or how it might change. This has now become a topic of theoretical study. He & Leland (1991) and Hodges & Carverhill (1993) have done work which characterizes the evolution of the equilibrium risk premium through time.

This characterization (which is based on a single representative investor assumption) implies that under quite general utility functions, the longer the horizon the less the risk premium will respond to changes in market value. Figure 1 illustrates this by showing some numerical results that we have calculated using this evolution. The solid line shows a hypothesized relationship between the price of risk, defined as the Sharpe ratio measure of the annual risk premium divided by the standard deviation. The risk premium might be 9% and the standard deviation might be 15% to give a figure of 0.6 for this ratio. We are postulating that it depends on the level of the market index, with a higher risk premium when the index is low and a lower risk premium when the index is high. We start off by assuming that when it is close to the horizon it is going to take the shape of the solid line, and our theory tells us how that will evolve when the horizon is a longer way off. What happens is it flattens out. Whereas it is fairly steep at the horizon, it is getting rather flatter when we are six years away. What that seems to imply is that in a real market, the investors with very long horizons will not require such a big risk premium after the

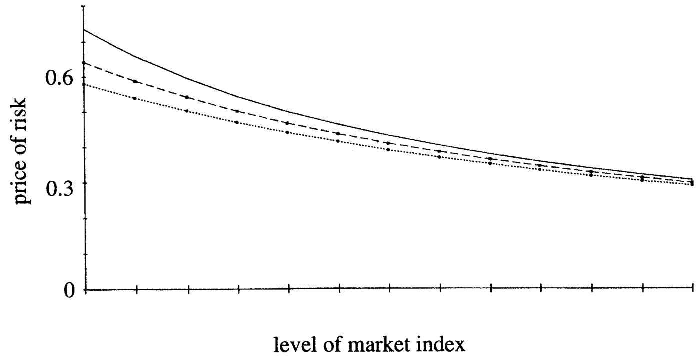

Figure 1. Evolution of the risk premium solution.

market has fallen as compared with the short horizon ones. The market will clear with the investors with long horizons tending to be contrarian, whereas investors with short horizons may feel a need to portfolio insure. We can never say that either a contrarian policy or one of portfolio insurance is the only sensible thing to do. It depends on what your situation is and what your objectives are. We cannot even say that all longer horizon investors should be contrarian, or all short horizon ones should portfolio insure. This is just the balance between the two on average.

The interesting thing about this analysis is that if we are in this kind of equilibrium, where the relationship of the risk premia to the level of the market changes systematically through time, our previous results concerning path independence still hold. On the other hand, there are also more complex equilibrium settings where life is a bit more complicated. We may still be able to use the options framework to think about dynamic asset allocation but the strict form of path independence we had earlier will no longer quite apply.

Tactical allocation

Fama (1991) provides the most recent academic survey on market anomalies, confirming the usual things we all know about: price earnings ratio effects; dividend yield effects; January effects; small firm effects. We know there are studies of in-house analysts that reveal some degree of forecasting skills, and tactical asset allocation policies would aim to be capitalizing on all of these types of things.

Now if we can make forecasts, then again we are making a major change to the Black–Scholes assumptions that Cox & Leland and Dybvig used in their analyses. We can always think in terms of contingent pay-offs as a function of the index. We can also think in terms of option theory to tell us what the exposure should be. However, the Black–Scholes delta is now telling you about our risk-return trade-off, rather than directly about market exposure. Thus, if you forecast, say, that the expected return on the market is 2% higher than normal, you then want to increase exposure by a corresponding amount compared to the calculated delta. This will destroy path independence, but you still have a framework for taking a consistent risk return stance in the market place.

The role of forecasting is particularly interesting. The studies by Hodges & Brealey (1973) and Treynor & Black (1973) look at the relationship between fund performance and the forecasting ability as measured in terms of the correlation between forecasts and outcomes. What they show is that very significant returns can be obtained with remarkably low levels of forecasting ability as measured by the correlation coefficient (R^2). It turns out that if you can make forecasts with an R^2 of 0.01 or 0.02, and use them properly, you can still make quite respectable returns of 2 or 3%. This is a major reason why performance measurement is difficult, because what is significant economically may well not be significant statistically. If you run a regression and find an R^2 of 0.03 you are liable to throw the thing away and say it is no good. That is all you should probably be expecting if you are looking at how good your analysts are anyway. If they are better than that, they will be making a fortune on their own account.

6. Performance measurement

Finally we turn to some issues concerning performance measurement. Most U.K. current practice is rather unsophisticated and the problems are very difficult. The National Association of Pension Funds Committee of Enquiry Report into Investment Performance Measurement (1990) considered, but decided not to pursue, risk adjusted measures. The LIFFE/LTOM (1992) document suggests a sensible treatment for futures but really does not address the issues for options. It treats options as equity substitutes and puts in some sensitivity analysis, but does not really tackle the problems raised by options or by dynamic asset allocation strategies.

Bookstaber & Clarke (1984) have provided analysis which shows clearly that 'methods which depend on mean and variance measures cannot be applied because options strategies mould the return distributions bringing the higher moments into play'. For the distribution of return on an all equity portfolio, there is no problem in using mean–variance analysis. However, suppose the fund manager has written a lot of out-of-the-money covered call options, it pushes the whole distribution to the right and then truncates the right tail. We are left with a distribution which also has a big spike where the call options start to be exercised against the manager, and just by looking at it, you can see that a mean–variance framework cannot be used to compare performance. That epitomizes the problem of performance measurement and it does not matter of course whether the firm actually wrote covered calls or whether it pursued a dynamic strategy which had the same effect. Either way, the probability distribution is quite distorted and we cannot use mean–variance analysis.

We therefore have to relinquish all of the Sharpe (1966), Treynor (1966) and Jensen (1968) philosophy of performance measurement. However, we can still use some of the philosophy of Fama's (1972) decomposition of returns. Fama suggests using benchmark portfolios to attribute performance to various sources. We can have a benchmark that has a moving beta equal to the actual beta of the fund, and the difference between that benchmark and what the fund does you attribute to stock selection. You then compare the moving beta benchmark against something with a fixed beta at the average level, and you attribute that to market timing, and so on. The philosophy can be extended in a variety of ways. One extension is described by Sharpe (1992). He suggests that we calculate the returns on benchmark indices for various asset classes and then run regressions between the fund return and the returns on these indices. That enables us to identify the effective allocation across classes. We can do this in a moving window, so we can track how the allocation across classes is changed. In his paper he contrasts the Trustees Commingled Fund which had a very static mix of asset classes, mostly in small stocks, with Fidelity Magellan

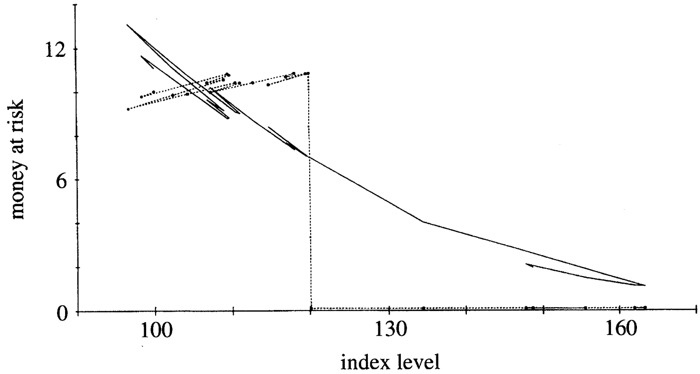

Figure 2. Lock-in strategy.

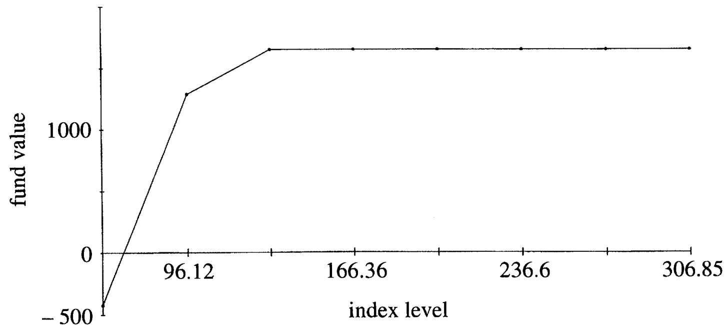

Figure 3. Estimated objective lock-in strategy.

which has an increasing proportion through time in growth stocks and a decreasing one in small stocks.

Finally, Hodges (1991) suggests that we can regress the portfolio exposure on the deltas of options spreads in order to understand horizon objectives and also to understand how close the fund comes to path independence. We create bull option spreads at various market levels and then we explain how close the observed market exposure is to a linear combination of the deltas from each of these bull option spreads. If the residual is zero then the manager is pursuing a classic path independent strategy. If there is a high residual, it is not at all path independent. The analysis gives us directly the inferred future fund value as a function of the index value (i.e. its contingent pay-off). Surprisingly (at least running that with simulated data) this procedure seems to work fairly well. Figure 2 shows a path independent strategy (solid) fitted to an actual simulated lock-in strategy (dotted). They are quite a long way apart, which shows that the lock-in is fairly inefficient. The inferred objectives are shown in figure 3 and you can see the lock-in aspect is revealed from the analysis.

7. Conclusions

In conclusion, the real power of derivatives for fund management lies more than anything for their use in modifying return distributions. We must, of course, recognize the liquidity limitations of the markets, and this conclusion may not apply so accurately to the very largest funds. Inconsistent risk exposure can be very expensive. Market regularities matter, so portfolio insurance probably has a hidden cost of lower expected returns. If you are contrarian you have to be able to bear the risk when the market is bad, but there is probably a reward. For performance measurement, mean–variance measures are inadequate and we are now just beginning to see new approaches that work directly with the pattern of risk exposure through time.

I acknowledge many helpful comments to earlier versions of this paper from Dr Les Clewlow, Dr Michael Selby and other colleagues in FORC, and from participants at seminars in London, Mannheim and Rotterdam.

References

Admati, A. R. & Ross, S. A. 1985 Measuring performance in a rational expectations equilibrium model. *J. Business* **58**, 1–26.

Allen, G. C. 1991 Performance attribution for global equity portfolios. *J. Portfolio Management* Fall, 59–65.

Black, F. & Scholes, M. 1973 The pricing of options and corporate liabilities. *J. Political Econ.* **81**, 637–659.

Bookstaber, R. 1990 *Option pricing and investment strategies* (3rd edn). Probus Publishing.

Bookstaber, R. & Clarke, R. 1984 Options portfolio strategies: measurement and evaluation. *J. Business* **57**, 469–492.

Cox, J. C. & Leland, H. E. 1982 On dynamic investment strategies. *Proc. seminar analysis of security prices*, Centre for Research in Security Prices, University of Chicago.

Dybvig, P. H. 1988 Inefficient dynamic portfolio strategies, or how to throw away a million dollars. *Rev. Financial Stud.* **1**, 67–88.

Fama, E. F. 1972 Components of investment performance. *J. Finance* **27**, 551–567.

Fama, E. F. 1991 Efficient capital markets: II. *J. Finance* **46**, 1575–1617.

Fama, E. F. & French, K. R. 1988 Divident yields and expected stock returns. *J. Financial Econ.* **22**, 3–25.

Galai, D. & Geske, R. 1984 Option performance measurement. *J. Portfolio Management.*

He, H. & Leland, H. 1993 On equilibrium asset price processes. *Rev. Financial Stud.* **6**, 593–617.

Hodges, S. D. 1976 Problems in the application of portfolio selection models. *Omega* **4**, 699–709.

Hodges, S. D. 1991 Ex-post evaluation of dynamic portfolio strategies (or how to tell whether a million dollars has been thrown away). Financial Options Research Centre, University of Warwick (Pre-print 91/23).

Hodges, S. D. & Brealey, R. A. 1973 Portfolio selection in a dynamic uncertain world. *Financial Analysts J.* **29**, 50–65.

Hodges, S. D. & Carverhill, A. P. 1993 The characterization of economic equilibria which support Black–Scholes option pricing. *Economic J.* **103**, 395–405.

Hodges, S. D. & Neuberger, A. 1989 Optimal replication of contingent claims under transactions costs. *Rev. Futures Markets* **8**, 223–242.

Jensen, M. C. 1968 The performance of mutual funds in the period 1945–1964. *J. Finance* **23**, 389–416.

Kritzman, M. P. 1990 *Asset allocation for institutional portfolios*. Homewood, Illinois: R. D. Irwin.

Kritzman, M. P. 1992 Asset allocation for individual investors. *Financial Analysts J.* Jan.–Feb., 12–13.

Leland, H. E. 1980 Who should buy a portfolio insurance? *J. Finance* **35**, 581–596.

LIFFE/LTOM 1992 The reporting and performance measurement of financial futures and options in investment portfolios. *LIFFE/LTOM Recommendations*, January.

Lo, A. W. & McKinley, A. C. 1988 Stock market prices do not follow random walks: evidence from a simple specification test. *Rev. Financial Stud.* **1**, 41–66.

Poterba, J. & Summers, L. 1988 Mean reversion in stock prices: evidence and implications. *J. Financial Econ.* **22**, 27–59.

Roll, R. 1992 A mean–variance analysis of tracking error. *J. Portfolio Management* Summer, 13–23.

Samuelson, P. A. 1969 Lifeline portfolio selection by dynamic stochastic programming. *Rev. Econ. Stat.* **51**, 239–246.

Sharpe, W. F. 1966 Mutual fund performance. *J. Business* **39**, 119–138.

Sharpe, W. F. 1992 Asset allocation: management style and performance measurement. *J. Portfolio Management* Winter, 7–19.

Treynor, J. L. 1966 How to rate management investment funds. *Harvard Business Rev.* **43**, 63–75.

Treynor, J. L. & Black, F. 1973 How to use security analysis to improve portfolio selection. *J. Business* **46**, 66–86.

Trippi, R. R. & Harrif, R. B. 1991 Dynamic asset allocation rules: survey and synthesis. *J. Portfolio Management* Summer, 19–26.

Discussion

R. S. Clarkson (*Dalserf, U.K.*). Experience suggests that an academic analysis of risk is futile in most of the situations. We need to concentrate instead on investor psychology; what Keynes described as 'anticipating the anticipations of others'. As opposed to the second order differential equations we have been discussing, I regard the behaviour of 'good' investors as being at least fourth order in nature, but we simply do not yet have such high order mathematical models. My guess is that in addition to the aggregate uncertainty introduced at the portfolio level by Professor Markowitz more than 40 years ago we need to build in uncertainty at the microscopic level along the lines of the Heisenberg Uncertainty Principle in quantum theory if we are to explain how individual investors behave and thereafter build up a new theory of finance on sounder foundations than those provided by the linear methodologies of the current theory.

There is a story that emphasizes the two crucial components of successful investment: obtaining important information before your competitors and understanding investor psychology. When the Battle of Waterloo was fought in 1815, there were no Reuters or Topic screens to give a real-time account of the engagement. Rothschild, however, had his own private sources of intelligence, and when his messenger brought him the crucial news at about 11 a.m. he went into the gilts market and sold heavily. Others assumed that we had lost at Waterloo, and the market plunged in waves of panic selling. Rothschild bought back at the bottom of the panic and then went out to lunch. When the Government messenger arrived in the afternoon it was realized that we had indeed won the Battle of Waterloo, and prices soared spectacularly. The enormous profits that George Soros made out of currency speculation around the time of sterling's ignominious exit from the Exchange Rate Mechanism suggests that investor psychology has not changed much over the past two or three centuries.

M. A. H. Dempster (*University of Essex, U.K.*). A class of models for dynamic management of portfolios of both assets and liabilities which are capable of capitalizing on even the low levels of accuracy in forecasting returns on securities mentioned – including those for options and futures – are discrete time dynamic stochastic programming models. This type of model was first applied by Bradley & Crane (1972) and Lane & Hutchinson (1980) and modern versions (Dempster & Ireland 1988; Carino *et al.* 1993)) tend to maximize the expected utility of terminal wealth at a finite horizon, taking into account horizon end effects and possibly subject to down-side risk control through explicit (almost sure) constraints. Rate and yield forecasts may be incorporated into the generation of explicit realizations of future rate and yield processes, termed *scenarios*, whose random variations are considered in the optimal evolution of portfolio decisions. More details may be found in Dempster (1994).

Additional references

Bradley, S. P. & Crane, D. B. 1972 A dynamic model for bond portfolio management. *Management Sci.* **19**, 139–151.

Carino, D. R. *et al.* 1993 The Russell–Yasuda Kasai financial planning model. *Management Sci.* (submitted).

Dempster, M. A. H. 1994 Dynamic asset/liability portfolio management. *Mathematical Finance.* (In the press.)

Dempster, M. A. H. & Ireland, A. M. 1988 A financial expert decision support system. In *Mathematical models for decision support* (ed. G. Mitra), pp. 415–440. Berlin: Springer.

Lane, M. & Hutchinson, P. 1980 A model for managing a certificate of deposit portfolio under uncertainty. In *Stochastic programming* (ed. M. A. H. Dempster), pp. 473–495. London: Academic Press.