

Foundations of Portfolio Theory

HARRY M. MARKOWITZ*

WHEN I STUDIED MICROECONOMICS forty years ago, I was first taught how optimizing firms and consumers would behave, and then taught the nature of the economic equilibrium which would result from such behavior. Let me refer to this as part one and part two of my microeconomics course. My work on portfolio theory considers how an optimizing investor would behave, whereas the work by Sharpe and Lintner on the Capital Asset Pricing Model (CAPM for short) is concerned with economic equilibrium assuming all investors optimize in the particular manner I proposed. Thus my work on the one hand, and that of Sharpe and Lintner on the other, provide part one and part two of a microeconomics of capital markets.

Professor Sharpe will discuss CAPM, part two of the course. I will confine my remarks to part one, portfolio theory. There are three major ways in which portfolio theory differs from the theory of the firm and the theory of the consumer which I was taught. First, it is concerned with investors rather than manufacturing firms or consumers. Second, it is concerned with economic agents who act under uncertainty. Third, it is a theory which can be used to direct practice, at least by large (usually institutional) investors with sufficient computer and database resources. The fact that it deals with investors rather than producers or consumers needs no further comment. Let me expand on the second and third differences.

In my microeconomics course, the theory of the producer assumed that the competitive firm knows the price at which it will sell the goods it produces. In the real world there is a delay between the decision to produce, the time of production and the time of sale. The price of the product at the time of sale may differ from that which was expected when the production decision was made. This uncertainty of eventual sales price is important in actual production planning but, quite reasonably, was ignored in classical economic models. It was judged not essential to the problem at hand.

Uncertainty cannot be dismissed so easily in the analysis of optimizing investor behavior. An investor who knew future returns with certainty would invest in only one security, namely the one with the highest future return. If several securities had the same, highest, future return then the investor would be indifferent between any of these, or any combination of these. In no

*Marvin Speiser Distinguished Professor of Finance and Economics, Baruch College, CUNY and Director of Research, DAIWA Security Trust Company.

case would the investor actually prefer a diversified portfolio. But diversification is a common and reasonable investment practice. Why? To reduce uncertainty! Clearly, the existence of uncertainty is essential to the analysis of rational investment behavior.

In discussing uncertainty below, I will speak as if investors faced known probability distributions. Of course, none of us know probability distributions of security returns. But, I was convinced by Leonard J. Savage, one of my great teachers at the University of Chicago, that a rational agent acting under uncertainty would act according to “probability beliefs” where no objective probabilities are known; and these probability beliefs or “subjective probabilities” combine exactly as do objective probabilities. This assumed, it is not clear and not relevant whether the probabilities, expected values, etc., I speak of below are for subjective or objective distributions.

The basic principles of portfolio theory came to me one day while I was reading John Burr Williams, *The Theory of Investment Value*. Williams proposed that the value of a stock should equal the present value of its future dividend stream. But clearly dividends are uncertain, so I took William’s recommendation to be to value a stock as the *expected value* of its discounted future dividend stream. But if the investor is concerned only with the expected values of securities, the investor must also be only interested in the expected value of the portfolio. To maximize the expected value of a portfolio, one need only invest in one security—the security with maximum expected return (or one such, if several tie for maximum). Thus action based on expected return only (like action based on certainty of the future) must be rejected as descriptive of actual or rational investment behavior.

It seemed obvious that investors are concerned with risk and return, and that these should be measured for the portfolio as a whole. Variance (or, equivalently, standard deviation), came to mind as a measure of risk of the portfolio. The fact that the variance of the portfolio, that is the variance of a weighted sum, involved all covariance terms added to the plausibility of the approach. Since there were two criteria—expected return and risk—the natural approach for an economics student was to imagine the investor selecting a point from the set of Pareto optimal expected return, variance of return combinations, now known as the efficient frontier. These were the basic elements of portfolio theory which appeared one day while reading Williams.

In subsequent months and years I filled in some details; and then others filled in many more. For example in 1956 I published the “critical line algorithm” for tracing out the efficient frontier given estimates of expected returns, variances and covariances, for any number of securities subject to various kinds of constraints. In my 1959 book I explored the relationship between my mean-variance analysis and the fundamental theories of action under risk and uncertainty of Von Neumann and Morgenstern and L. J. Savage.

Starting in the 1960s, Sharpe, Blume, King, and Rosenberg greatly clarified the problem of estimating covariances. This past September I attended

the Berkeley Program in Finance at which several analysts reported success in using publicly available accounting figures, perhaps combined with security analysts' earnings estimates, to estimate expected returns. I do not mean that their estimates eliminate uncertainty—only that, on the average, securities with higher estimates outperform those with lower estimates.

So, equipped with database, computer algorithms and methods of estimation, the modern portfolio theorist is able to trace out mean-variance frontiers for large universes of securities. But is this the right thing to do for the investor? In particular, are mean and variance proper and sufficient criteria for portfolio choice?

To help answer this question, let us consider the theory of rational choice under uncertainty. In doing so, let us recall the third way in which portfolio theory is to differ from classical microeconomic theory of the firm or consumer. We seek a set of rules which investors can follow in fact—at least investors with sufficient computational resources. Thus we prefer an approximate method which is computationally feasible to a precise one which cannot be computed. I believe that this is the point at which Kenneth Arrow's work on the economics of uncertainty diverges from mine. He sought a precise and general solution. I sought as good an approximation as could be implemented. I believe that both lines of inquiry are valuable.

The discussion of principles of rational behavior under uncertainty in Part IV of my 1959 book starts with a variant of L. J. Savage's axioms. From such axioms it follows that one should choose a strategy which maximizes expected utility for a many-period game. This, in turn, implies that the investor should act each period so as to maximize the expected value of a single period utility function. This single period utility function may depend on portfolio return and perhaps other state variables. For now, assume that it depends only on portfolio return.

In this case, the crucial question is this: if an investor with a particular single period utility function acted only on the basis of expected return and variance, could the investor achieve almost maximum expected utility? Or, to put it another way, if you know the expected value and variance of a probability distribution of return on a portfolio can you guess fairly closely its expected utility?

A great deal of research has been done on this question, but more is needed. Let me briefly characterize some results, and some open questions. Table I is extracted from Levy and Markowitz. The rows of the table represent various utility functions. For example, the first row reports results for $U(R) = \log(1 + R)$ where R is the rate of return on the portfolio; the second row reports results for $U(R) = (1 + R)^{0.1}$, etc., as indicated in the first column of the table. The second through fifth column of the table represent various sets of historical distributions of returns on portfolios. For example, the second column represents annual returns on 149 investment companies, 1958–67; the third column represents annual returns on 97 stocks.

The calculations associated with the second column in effect assume that an investor must choose one out of 149 portfolios, and his probability beliefs

Table I
Correlation Between EU and $f(E, V)$ for Four Historical Distributions

Utility Function	Annual Returns on 149 Mutual Funds ¹	Annual Returns on 97 Stocks ²	Monthly returns on 97 Stocks ²	Random Portfolios of 5 or 6 Stocks ³
$\text{Log}(1 + R)$	0.997	0.880	0.995	0.998
$(1 + R)^a$				
$a = 0.1$	0.998	0.895	0.996	0.998
$a = 0.3$	0.999	0.932	0.998	0.999
$a = 0.5$	0.999	0.968	0.999	0.999
$a = 0.7$	0.999	0.991	0.999	0.999
$a = 0.9$	0.999	0.999	0.999	0.999
$-e^{b(1+R)}$				
$b = 0.1$	0.999	0.999	0.999	0.999
$b = 0.5$	0.999	0.961	0.999	0.999
$b = 1.0$	0.997	0.850	0.997	0.998
$b = 3.0$	0.949	0.850	0.976	0.958
$b = 5.0$	0.855	0.863	0.961	0.919
$b = 10.$	0.449	0.659	0.899	0.768

¹The annual rate of return of the 149 mutual funds are taken from the various annual issues of A. Wiesenberger and Company. All mutual funds whose rates of return are reported in Wiesenberger for the whole period 1958-67 are included in the analysis.

²This data base of 97 U.S. stocks, available at Hebrew University, had previously been obtained as follows: a sample of 100 stocks was randomly drawn from the CRSP (Center for Research in Security Prices, University of Chicago) tape, subject to the constraint that all had reported rates of return for the whole period 1948-68. Some mechanical problems reduced the usable sample size from 100 to 97. The inclusion only of stocks which had reported rates of return during the whole period may have introduced survival bias into the sample. This did not appear harmful for the purpose at hand.

³We randomly drew 5 stocks to constitute the first portfolio; 5 different stocks to constitute the second portfolio, etc. Since we have 97 stocks in our sample, the eighteenth and nineteenth portfolios include 6 stocks each. Repetition of this experiment with new random variables produced negligible variations in the numbers reported, except for the case of $U = -e^{-10(1+R)}$. A median figure is reported in the table for this case.

concerning returns on these portfolios are the same as historical returns. It is not that we recommend this as a way of forming beliefs; rather, we use this as an example of distributions of returns which occur in fact.

For each utility function, and for each of the 149 probability distributions of the second column, we computed its "expected" (that is, its mean) utility

$$EU = \sum_{t=1}^T U(R_t) / T \quad (1)$$

where T is the number of periods in the sample, and R_t the rate of return in period t . We also computed various approximations to EU where the approximation depends only on the mean value E and the variance V of the distribution. Of the various approximations tried in Levy-Markowitz the one

which did best, almost without exception, was essentially that suggested in Markowitz (1959), namely

$$f(E, V) = U(E) + 0.5U''(E)V \tag{2}$$

For example, if $U(R) = \log(1 + R)$,

$$f(E, V) = \log(1 + R) - 0.5V/(1 + E)^2. \tag{3}$$

Equation (2) may be thought of as a rule by which, if you know the E and V of a distribution, you can guess at its expected utility. The figures in Table I are for the Levy-Markowitz approximation which is essentially (2). The entry in the second column, first row reports that, over the 149 probability distributions, the correlation between EU and $f(E, V)$ was 0.997 for $U = \log(1 + r)$. The remaining entries in the second column similarly show the correlation, over the 149 probability distributions, of EU and $f(E, V)$ for the utility functions tested. In most cases the correlation was extremely high, usually exceeding 0.99. We will discuss an exceptional case shortly.

The third column shows the correlation between EU and $f(E, V)$ for a sample of annual return on one-stock "portfolios". The correlations are clearly less than for the diversified investment company portfolios of the second column. The fourth column again considers undiversified, single stock portfolios, but this time for monthly holding period returns. The correlations are much higher than those of column three, usually as high or higher than those in column two. Thus, for the investor who revises his or her portfolio monthly, even for portfolios whose returns were as variable as those of individual stocks, $f(E, V)$ would be highly correlated with EU for the utility functions considered.

The fifth column shows annual holding period returns, now for randomly selected portfolios with 5 or 6 securities each. The correlations are generally quite high again—comparable to those in the second column. Thus, at least, for these probability distributions and most of these utility functions, $f(E, V)$ approximates EU quite well for diversified portfolios, even "slightly" diversified portfolios of size 5 and 6.

Not all expected maximizers are equally served by mean-variance approximations. For example, the investor with $U = -e^{-10(1+R)}$ will find mean-variance much less satisfactory than others presented in Table I. Levy and Markowitz have two observations concerning an expected utility maximizer with $U = -e^{-10(1+R)}$.

The first observation is that an investor who had $-e^{-10(1+R)}$ as his or her utility function would have some very strange preferences among probabilities of return. Reasonably enough, he or she would not insist on certainty of return. For example, the investor would prefer (a) a 50-50 chance of a 5 percent gain vs. a 25 percent gain rather than have (b) a 10 percent gain with certainty. On the other hand there is no R which would induce the investor to take (a) a 50-50 chance of zero return (no gain, no loss) vs. a gain of R rather than have (b) a 10 percent return with certainty. Thus a 50-50

chance of breaking even vs. a 100,000 percent return, would be considered less desirable than a 10 percent return with certainty. We believed that few if any investors had preferences anything like these.

A second observation was that even if some unusual investor did have the utility function in question, such an investor could determine in advance that $f(E, V)$ was not a good approximation for this EU . Table II shows the difference between $U(R)$ and the Taylor approximation upon which (2) is based, namely,

$$Q(R) = U(E) + U'(E)(R - E) + 0.5U''(E)(R - E)^2 \quad (4)$$

for $U = \log(1 + R)$ and $U = -1000e^{-10(1+R)}$, for $E = 0.10$. For the various R listed in the first column, the second through fourth columns show $U(R)$, $Q(R)$ and $\Delta(R) = U(R) - Q(R)$ for $\log(1 + R)$; the following three columns show the same for $-1000e^{-10(1+R)}$. Since the choices implied by a utility function are unaffected by multiplying it by a positive constant, it is not the magnitude of the $\Delta(R)$ s which are important. Rather it is the variation in $\Delta(R)$ as compared to that in $U(R)$. For example, Levy and Markowitz present a lower bound on the correlation between $U(R)$ and $f(E, V)$ as a function of the standard deviations of U and Δ . As we see in the table, as $\log(1 + R)$ goes from -0.357 at $R = -0.30$ to 0.470 at $R = 0.60$, $|\Delta|$ never exceeds 0.024 . In contrast, as $-1000e^{-10(1+R)}$ goes from -0.912 to -0.0001 , $|\Delta|$ often exceeds 0.03 and has a maximum of -0.695 .¹ Thus, if an investor had $U = -e^{-10(1+R)}$ as a utility function, a comparison of $U(R)$, $Q(R)$, and $\Delta(R)$ would provide ample warning that mean-variance is not suitable.

Levy and Markowitz present other empirical results. They also explain the difference between assuming that an investor has a quadratic utility function

Table II
Quadratic Approximation to Two Utility Functions
 $E = 00.1$

R	$\log(1 + R)$	$Q_L(R)$	Δ_L	$-1000e^{10(1+R)}$	$Q_E(R)$	Δ_E
-.30	-.35667	-.33444	-.02223	-.91188	-.21712	-.69476
-.20	-.22314	-.21461	-.00854	-.33546	-.14196	-.14950
-.10	-.10536	-.10304	-.00232	-.12341	-.08351	-.03990
.00	.00000	.00027	-.00027	-.04540	-.04175	-.00365
.10	.09531	.09531	.00000	-.01670	-.01670	.00000
.20	.18232	.18209	.00023	-.00614	-.00835	.00221
.30	.26236	.26060	.00176	-.00226	-.01670	.01444
.40	.33647	.33085	.00563	-.00083	-.04175	.04092
.50	.40546	.39283	.01263	-.00031	-.08351	.08320
.60	.47000	.44655	.02345	-.00011	-.14196	.14185

¹Among the 149 mutual funds, those with E near .10 all had annual returns between a 30% loss and a 60% gain. Specifically, 64 distributions had $0.08 \leq E \leq 0.12$, and all had returns within the range indicated.

versus using a quadratic approximation to a given utility function to develop an $f(E, V)$ approximation, such as that in (2). In particular, they show that $f(E, V)$ in (2) is not subject to the Arrow, Pratt objection to a quadratic utility function, that it has increasing risk aversion. Indeed, Levy and Markowitz show that a large class of $f(E, V)$ approximations, including (2), have the same risk aversion in the small as does the original EU maximizer.

I will not recount here these further Levy and Markowitz results, nor will I go into important results of many others. Chapter 3 of Markowitz (1987) includes a survey of the area up to that time. I will, however, briefly note results in two important unpublished papers.

Levy and Markowitz measure the efficacy of $f(E, V)$ by the correlation between it and EU . Y. Simaan defines the optimization premium to be the percent the investor would be just willing to pay out of the portfolio for the privilege of choosing the true EU maximizing portfolio rather than being confined to the mean-variance "second best". The reason for performing a mean-variance analysis in fact, rather than a theoretically correct expected utility analysis, is convenience, cost or feasibility. It is typically much more expensive to find a utility maximizing portfolio than to trace out an entire mean-variance frontier. The data requirements for an expected utility analysis can substantially exceed those of a mean-variance analysis, since estimates of first and second moments generally are not sufficient for the former. Finally, there is the problem of determining the investor's utility function. Simaan's criteria measures the worth, as a percent of the portfolio, paid out of the portfolio, of incurring the added expenses of finding an EU maximizing portfolio. He solves for this optimization premium analytically under certain assumptions.

L. Ederington evaluates EU approximations using thousands of synthetic time series generated by randomly selecting from actual time series. He evaluates approximations like (2), except that they use the first three or four moments, as well as (2) that uses the first two. It is all very well to point out theoretically that more moments are better than fewer. The practical question is: how much?

Ederington finds, as did Levy and Markowitz, that for some utility functions the mean-variance approximation is so good that there is virtually no room for improvement. Where the mean-variance approximation falters, Ederington finds that typically three moments provides little improvement to the approximation whereas four moments improves the approximation considerably.

Despite noteworthy results reported above, and many more that I have not described here, there is much to be done. Three examples will illustrate the need.

First, all the experimentation and analysis to date give us a rather spotty account of where mean-variance serves well and where it falters. Perhaps it is possible to develop a more systematic characterization of the utility functions and distributions for which the mean-variance approximation is good, bad and marginal.

Second, suppose that the investor has a utility function for which mean-variance provides a close approximation, but the investor does not know precisely which one. In this case the investor need not determine his or her utility function to obtain a near optimum portfolio. The investor need only pick carefully from the (one-dimensional) curve of efficient *EV* combinations in the two dimensional *EV* space. To pursue a similar approach when four moments are required, the investor must pick carefully from a three-dimensional surface in a four-dimensional space. This raises serious operational problems in itself, even if we overcome computational problems due to the nonconvexity of sets of portfolios with given third moment or better.

But perhaps there is an alternative. Perhaps some other measure of portfolio risk will serve in a two parameter analysis for some of the utility functions which are a problem to variance. For example, in Chapter 9 of Markowitz (1959) I proposed the "semi-variance" *S* as a measure of risk where

$$S = E(\text{Min}(0, R - c)^2)$$

where $c = E(R)$ or c is a constant independent of choice of portfolio. Semi-variance seems more plausible than variance as a measure of risk, since it is concerned only with adverse deviations. But, as far as I know, to date no one has determined whether there is a substantial class of utility functions for which mean-semi-variance succeeds while mean-variance fails to provide an adequate approximation to *EU*.

Third, in general the derived, single period utility functions can contain state-variables in addition to return (or end of period wealth). Expected utility, in this case, can be estimated from return and state-variable means, variances and covariances provided that utility is approximately quadratic in the relevant region. (Recall the Levy-Markowitz analysis of quadratic utility versus quadratic approximation in the relevant region.) To my knowledge, no one has investigated such quadratic approximation for cases in which state variables other than portfolio value are needed in practice.

In sum, it seems to me that the theory of rational behavior under uncertainty can continue to provide insights as to which practicable procedures provide near optimum results. In particular, it can further help evaluate the adequacy of mean and variance, or alternate practical measures, as criteria.

Finally, I would like to add a comment concerning portfolio theory as a part of the microeconomics of action under uncertainty. It has not always been considered so. For example, when I defended my dissertation as a student in the Economics Department of the University of Chicago, Professor Milton Friedman argued that portfolio theory was not Economics, and that they could not award me a Ph.D. degree in Economics for a dissertation which was not in Economics. I assume that he was only half serious, since they did award me the degree without long debate. As to the merits of his arguments, at this point I am quite willing to concede: at the time I defended my dissertation, portfolio theory was not part of Economics. But now it is.

REFERENCES

- Arrow, K., 1965, *Aspects of the Theory of Risk Bearing*, (Helsinki).
- Blume, M., 1971, On the assessment of risk, *Journal of Finance* 26, 1-10.
- Ederington, L. H. 1986, Mean-variance as an approximation to expected utility maximization, Working paper 86-5, School of Business Administration, Washington University, St. Louis, Missouri.
- King, B. F., 1966, Market and industry factors in stock price behavior, *Journal of Business* 39, 139-190.
- Levy, H. and H. M. Markowitz, 1979, Approximating expected utility by a function of mean and variance, *American Economic Review* 69, 308-317.
- Lintner, J., 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* February.
- Markowitz, H. M., 1952, Portfolio selection, *The Journal of Finance* 7, 77-91.
- , 1956, The optimization of a quadratic function subject to linear constraints, *Naval Research Logistics Quarterly* 3.
- , 1959, *Portfolio Selection: Efficient Diversification of Investments*, (Wiley, Yale University Press, 1970, Basil Blackwell, 1991).
- , 1987, *Mean-Variance Analysis in Portfolio Choice and Capital Markets*, (Basil Blackwell, paperback edition, Basil Blackwell, 1990).
- Pratt, J. W., 1964, Risk aversion in the small and in the large, *Econometrica* 32, 122-136.
- Rosenberg, B., 1974, Extra-market components of covariance in security returns, *Journal of Financial and Quantitative Analysis* 29, 263-273.
- Savage, L. J., 1954, *The Foundations of Statistics*, 2nd ed., (Wiley, Dover, 1972).
- Sharpe, W. F., 1963, A simplified model for portfolio analysis, *Management Science* 9, 277-293.
- , 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *The Journal of Finance* 19, 425-442.
- Simaan, Y., 1987, Portfolio selection and capital asset pricing for a class of non-spherical distributions of assets returns, dissertation, Baruch College, The City University of New York.
- A. Wiesenberger and Company, *Investment Companies*, annual editions, New York.
- Von Neumann, J., and O. Morgenstern, 1944, *Theory of Games and Economic Behavior*, 3rd ed., (Princeton University Press, 1953).
- Williams, J. B., 1938, *The Theory of Investment Value*, (Harvard University Press, Cambridge).