JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS
November 1975

THEORY OF FINANCE FROM THE PERSPECTIVE OF CONTINUOUS TIME

*Robert C. Merton**

It is not uncommon on occasions such as this to talk about the shortcomings in the theory of Finance, and to emphasize how little progress has been made in answering the basic questions in Finance, despite enormous research efforts. Indeed, it is not uncommon on such occasions to attack our basic "mythodology," particularly the "Ivory Tower" nature of our assumptions, as the major reasons for our lack of progress. Like a Sunday morning sermon, such talks serve many useful functions. For one, they serve to deflate our professional egos. For another, they serve to remind us that the importance of a contribution as judged by our professional peers (the gold we really work for) is often not closely aligned with its operational importance in the outside world. Also, such talks serve to comfort those just entering the field, by letting them know that there is much left to do because so little has been done. While such talks are not uncommon, this is not what my talk is about. Rather, my discussion centers on the positive progress made in the development of a theory of Finance using the continuous-time mode of analysis.

Hearing this in 1975, amidst an economic recession with a baffling new disease called "stagflation" and with our financial markets only beginning to recover from the worst turmoil in almost 40 years, some will say that I am embarked on a fool's errand. Perhaps. At any rate, on this errand, I shall discuss the continuous-time solutions to some of the basic problems of Finance: portfolio selection, capital market equilibrium, and the pricing of capital assets; the derived functions of financial intermediaries and instruments; and the pricing of corporate liabilities. Rather than dwelling on the technical aspects of obtaining these solutions, I will discuss the substantive results and why they seem to differ from those obtained in other modes of analysis. While my main interest is in the substantive results, some methodological discussion is necessary because the continuous-time mode is relatively new. However, I will attempt to keep this discussion to a minimum by emphasizing only those assumptions that differ from its discrete-time counterpart, and leave those assumptions common to both types of analysis as understood.

The natural beginning point for the development of a theory of Finance is the problem of lifetime consumption and portfolio selection for the individual consumer or household under uncertainty. There is a long standing tradition in economic theory to take the existence of households and their tastes as "givens," as exogenous to the theory. But this tradition does *not* extend to economic organizations and institutions: that is, they are regarded as existing primarily because of the functions they serve instead of functioning primarily because they exist. Economic organizations and institutions, unlike households and their tastes, are *endogenous* to the theory. To derive the functions of these financial

**Massachusetts Institute of Technology.* *Aid from the Natural Science Foundation is gratefully acknowledged.*

organizations and institutions, therefore, we must first derive the behavior of individual households.

The basic lifetime consumption-portfolio selection problem can be stated as follows: at each point in time, the consumer must make two types of decisions: (1) How much of his wealth to consume (the "consumption choice"); (2) How to allocate that part of his wealth that he does not consume, across alternative assets (the "portfolio choice"). It is assumed that he makes these decisions so as to maximize the expected value of his utility of lifetime consumption. This function is usually represented as a sum (or in the case of continuous-time, a time integral) of strictly concave von-Neumann-Morgenstern utility functions for consumption at each date plus possibly an end-of-life utility function of wealth representing any bequest motives.

So formulated, a set of stochastic processes for the state variables, typically the returns on assets, are posited, either as objective probability distributions for the future course of returns or as representations of the investor's subjective beliefs, and the problem is solved by stochastic dynamic programming. A complete solution contains the individual's demands for consumption and assets as a function of age, wealth, and the other relevant state variables. To this point, of course, the discussion is consistent with either a discrete-time or continuous-time formulation.

While there are differing assumptions about the continuous-time formulation of the problem depending upon the stochastic processes posited, there are three assumptions common to all such formulations: Namely,

(1) The capital markets are assumed to be open all the time, and therefore economic agents have the opportunity to trade continuously.

(2) Prices of assets traded in speculative markets satisfy the "Efficient Markets Hypothesis" of Fama and Samuelson.[1] Namely, assets are priced so that the stochastic processes describing the *unanticipated* parts of their returns (i.e., the actual returns minus their expected value) are a martingale. This assumption *does not* imply an independent increments process although such a process satisfies this assumption.

(3) The stochastic processes generating the state variables can be described as either:

[1] Diffusion processes with continuous sample paths. Simply described, the state variables generated by these processes are changing all the time but the magnitude of such changes are small over a short time period.

or

[2] Compound Poisson processes with step function type sample paths. Simply described, the state variables generated by these processes will, almost certainly, over a short enough time interval have no change, or with a very small probability, a radical change or "jump" will occur.

or

[3] A mixture of both types.

While all three types have been used in the Finance literature,[2] most of the best-known

---

[1]See Fama [5] and Samuelson [23].

[2]See Merton [10-17]; Black and Scholes [1, 2]; Cox and Ross [4]; Merton and Samuelson [18].

results from continuous-time analysis have come from restricting the state variable dynamics to being diffusion processes with continuous sample paths. Indeed, since virtually any reasonable stochastic process arising in an economics context can be adequately approximated by some mixture of these two types, I would expect any serious disagreement with the stochastic process assumption would be on the use of a special form of these processes, rather than with the processes themselves.

The basic assumptions established, let us return to the consumption-portfolio problem, leaving for later the discussion of the merits of the assumptions. In its simplest form,[3] the continuous-time version of the problem assumes that the only source of uncertainty facing the consumer is the rate of return on assets. It is further assumed that these returns are generated by diffusion processes with continuous sample paths and that returns are serially independent and identically distributed through time, i.e., that prices follow a geometric Brownian motion and hence, the prices are lognormally distributed. Also, it is assumed that there is a single consumption good.

In solving this problem, as in more general forms of the continuous time analysis, the first-order conditions for optimal demands for assets are linear, and hence the demand functions can be solved for explicitly by matrix inversion. Moreover, in this problem, the relative demands for risky assets, i.e., the demand for, say, risky asset i divided by the demand for risky asset j, is independent of the investor's preferences of wealth level, and indeed, depend only on the instantaneous means, variances, and covariances of the returns. From this, it is a short step to prove a mutual fund or separation theorem: namely, that all investors agreeing on the distribution of returns will be indifferent between choosing their portfolios from among the basic securities or from just two mutual funds. Further, the separation is complete because the compositions of two funds can be determined solely from the "technological" data about returns without knowledge of investors' preferences, wealth levels, or time horizons. Moreover, if one of the assets is riskless, then one fund can be that asset and the other need only have risky assets, and in this case, the risky fund is called the "optimal combination of risky assets." Although these results are identical in structure to the classic Markowitz-Tobin mean-variance findings, their derivation here comes from a quite different set of assumptions.

The classic mean-variance results are deduced in a static framework by hypothesizing that mean and variance are sufficient statistics for making decisions. Indeed, when it was shown that the only conditions under which this hypothesis is consistent with expected utility maximization are if returns are normally distributed or if utility functions are quadratic, economists (with the exception of those in Finance) lost interest in this approach (except for example purposes) and returned to the more general framework of expected utility maximization.

In contrast, the continuous-time results come from an intertemporal model, are consistent with any concave utility function, and assume a return structure consistent with limited liability: namely, the lognormal distribution which has for a long time been a prototype distribution for security returns.[4]

[3]See Merton [10] and [11, section 5].

[4]It should be emphasized that the lognormality assumption applies only in the simple form of the continuous time formulation, and is not essential to the basic simplifications gained from this mode of analysis.

The natural question to ask is: "Why does it work here, when it didn't in discrete time?" What rabbits went into the hat? Indeed, one *should* be suspicious about any results that obtain in a continuous-time formulation but not in its discrete counterpart. There are two "rabbits": one, the assumption that investors can revise their portfolios frequently, and two, that prices can only change by small amounts in short periods of time. Moreover, the continuous-time result *is* consistent with the discrete-time solution when the appropriate comparison is made.

To see this, I make a short digression: There are three time intervals or horizons involved in the consumption-portfolio problem. (1) *The trading horizon*, is the *minimum* length of time between which successive transactions by economic agents can be made in the market. This is determined by the structure of the markets in the economy. For example, how frequently the market is open, and this time scale is not determined by the individual investor. (2) *The decision horizon*, is the length of time between which the investor makes successive decisions. So, for example, an investor with a fixed decision interval of one month, who makes a consumption decision and portfolio allocation today will under no conditions make any new decisions or take any action prior to one month from now. (3) *The planning horizon*, is the maximum length of time for which the investor gives any weight in his utility function. Typically, this time period would correspond to the balance of the consumer's lifetime. Actually, there is a fourth time interval, not directly related to the problem's solution but related to empirical testing of the derived behavior. I call this the *observation horizon*. It is the length of time between successive observations of the data by the tester, and is typically, either daily, weekly, monthly, quarterly, or annually. It is useful to keep these three or four different time intervals in mind when evaluating the relative merits of alternative formulations for the same problem. Indeed, the choice of time intervals has a significant impact on the derived behavior: a fact too often neglected in many analyses.

The one-period or static approach to portfolio selection implicitly assumes that the decision and planning horizons are the same: "one-period." Moreover, when such models are aggregated to determine market equilibrium relationships, it is further implicitly assumed that these intervals are the *same* for all investors, and therefore, correspond to the trading interval. Further, when such models are tested, the tester usually assumes that these three intervals are in turn equal to the observation interval from which he has chosen his data.

In the typical multiperiod, discrete-time analysis, the trading and decision intervals are assumed equal and the same for all participants. However, the actual time length of these intervals is left unspecified. Hence, while usually not recognized explicitly, every such solution derived has as an implicit argument the length of the time interval, denoted here by "h." Clearly, if one were to vary the "h" in such solutions, the derived behavior of the investors would change, as indeed would any deduced equilibrium relationship. Moreover, for any derived behavior to be true for an arbitrary discrete-time model (i.e., one where "h" is not specified), it would have to be invariant to "h." Think about what this means: such a result would have to obtain whether investors had the opportunity to revise their portfolios every day or were "frozen" into their investments for ten years. In this light, it is not surprising that there are few results from these arbitrary discrete-time models, and those that do obtain are for the most part a qualitative nature such as: "risk-averse investors diversify."

By contrast, the continuous-time model is very explicit about the value of "h": namely, h = 0. Of course, this is only a theoretical abstraction since actual continuous trading is not possible. However, with a few technical exceptions the continuous-time solution is the valid continuous limit of the discrete-time solution as the trading interval h tends to zero. I.e., given a delta, I can always find an "h" small enough so that the difference between the continuous-time solution and its discrete counterpart for that h is less than delta, using a reasonable metric for measuring "difference." Thus, the continuous-time solution is a valid approximation to the discrete-time solution, and its accuracy is a function of the actual structure of returns and the length of the "true" discrete time interval. And it is in this sense that the continuous-time and discrete-time results are consistent.

While I thought this correspondence was clear from the derivations in my papers, it has periodically been rediscovered. Indeed, six years after my first paper appeared in print, a new twist in its rediscovery has been to use it to argue for the superiority of discrete-time analysis over the continuous analysis because discrete-time includes continuous time as a limiting case.[5]

While I am on the subject of such comparisons, the major substantive results obtained in discrete-time analysis have required the assumption of a specific utility function or family of functions (for example, the HARA family). Quite aside from the issue of whether assuming a specific family of utility functions is superior to assuming a specific family of stochastic processes, the limited results of sharing rules and separation theorems in the discrete case not only require investors to have the same utility functions but they also must have the same decision interval (what I have called "h"). Indeed, for such results to be empirically testable, the tester will have to settle on what that interval is. In this sense, the discrete analysis is operationally no more general than the continuous analysis. Moreover, having looked at the data on stock and bond returns, if that interval is less than a month, then the discrete and continuous solutions will be virtually the same.

A more serious criticism of the realism of the continuous-model is that with transactions costs investors cannot trade continuously. This is certainly true. Indeed, one reason often given for finite trading-intervals is to give implicit, if not explicit, recognition to these costs. I have argued that this is unsatisfactory because the length of time between trades in the presence of such costs will almost certainly be "action-oriented" and therefore stochastic in nature, and the proper way to handle this problem is to start with the continuous-time model and deduce the optimal intervals. In a recent paper, Magill and Constantinides [9] have deduced such solutions, and the optimal trading intervals are discrete and stochastic. Indeed, the investor could end up trading more than once in a day or not for many months, depending on the *ex-post* time path followed by prices. The derived behavior for investors is to trade when the gain from better diversification offsets the cost of transacting. Their analysis also casts light on an issue raised earlier: namely, there are some cases where the limit of the discrete-time solution is not the continuous solution. For example, if an investor has

[5] See Rubinstein [22, p. 5].

an isoelastic utility function and if security returns are lognormally distributed, then in a standard discrete-time analysis, such an investor will never choose to borrow or short sell.[6] However, in a continuous-time analysis, the same investor may well choose to borrow or short sell. Which is the more "reasonable" description? The Magill-Constantinides analyses demonstrate that even though the trading intervals are discrete with probability one, investors with utility functions in the HARA family may well choose to borrow or short sell. The reason for the difference in derived behavior between the standard discrete-time and the stochastic discrete-time analyses is clear. In the former, the investor is "frozen" into holding his portfolio until the end of the period, and hence by shortselling, or borrowing, risks ruin. In the latter, the investor can revise his portfolio at any time (although he incurs a cost to do so) and hence he will readjust his portfolio, when necessary, to avoid ruin. I submit that the latter behavior is a better description of how investors behave, and it illustrates why using discrete-time analysis as an implicit method of recognizing transactions costs is a poor substitute for its explicit recognition in a continuous-time framework.

In summary, the continuous-time solution is consistent with its discrete-time counterpart when the trading interval is "small," and to my mind, the assumptions required are descriptive of capital markets as they actually function. Moreover, the continuous-time analysis has all the virtues of simplicity and empirical tractability found in the classic mean-variance analysis but without its objectionable assumptions.

Returning to the substantive findings of the basic consumption-portfolio selection problem, once the separation or mutual fund theorem is proven, it is straightforward to derive a continuous-time version of the Capital Asset Pricing Model. Indeed, in any model where all investors hold risky assets in the same relative proportions, it follows immediately that this "optimal combination of risky assets" must be the market portfolio for equilibrium to obtain. For if all investors want to hold risky assets in the same relative proportions, then there is only one way in which this is possible: namely, these relative proportions must be identical to those in the market portfolio. It therefore follows that among all possible investment strategies, the only one that all investors *could* follow is the one that says hold all assets in proportion to their market value.

Let me remark, somewhat parenthetically, that if among the soothsayers and strategists of Wall Street, there were one best investment strategy, and if this "best" strategy became widely known, then whatever the original statement of the strategy, it must lead to simply this imperative: hold all assets in proportion to their market value. For suppose such a strategy required that the investor hold equal amounts of Ford and General Motors. How could all investors following this best strategy do so, unless the total value of each were the same?

Having established that the optimal combination of risky assets is the market portfolio, the Security Market Line relationship follows directly. However, in the continuous-

[6] The reason is that the lognormal distribution has the full range of nonnegative outcomes and the expected value of the marginal utility of an isoelastic utility function is infinite if there is a positive probability of zero wealth, and undefined for negative wealth.

time model, it will only hold for short observation intervals, while in the original static version, the observation interval is never specified.

This simple version of the continuous-time model has been attacked on the grounds that it is not consistent with intertemporal equilibrium:[7] namely, it is claimed that if all risky assets are held in the same proportions throughout time, then the only way this is possible is if the *ex-post* returns on all assets are the same. Of course, this is nonsense. For this criticism to follow, one must make some rather absurd assumptions about the supplies of assets. Namely, it must be assumed that firms cannot raise additional investment capital except through internally-generated profits and that firms must reinvest all such profits in their own technology. In other words, a firm cannot distribute profits through dividends or share repurchase and it cannot invest in the technologies of other firms. So, for example, in the early part of this century, buggy whip manufacturers would have had to reinvest all their profits in further production of buggy whips, while automobile manufacturers could not have raised new capital to produce automobiles. Given that a purported major function of the capital markets is the efficient allocation of resources to the most productive investments, it is not surprising that strange results would follow from such a set of restrictions.

I would point out the positive results that the posited structure of returns in this simple version of the continuous-time model is consistent with intertemporal equilibrium for a simple wheat economy where the different risky assets correspond to alternative (uncertain) harvest technologies which remain the same through time. Indeed, in this economy, equilibrium is achieved through time by a pure quantity adjustment in the amount of wheat allocated to each technology.

I would also point out that the behavior implied by this model is consistent with a changing investment opportunity set if such changes are purely random. That is, even if the expected returns and covariances among those returns do change over time, the behavior at each point in time will be as if they are fixed at the current levels provided that such changes are completely random. The proof follows along lines used by Fama in his 1970 *AER* article [6].

In a closely related interpretation, the simple version of the model is also consistent with an investor who does not know the period-by-period, transitory expected return and covariance structure, but who does know the long-run or steady-state equilibrium structure of returns. For example, an investor may not know the *ex-ante* expected return on the market for the next six months, but he may know that historically, the risk premium on the market has averaged 7 percent. If he believes that this structure is constant (at least in real terms), then his optimal behavior will be generated by using this model. Of course, he will still have reason to adjust his portfolio over time to achieve the appropriate pattern of consumption, to maintain diversification, and to maintain the optimal risk-return mix.

This simple version of the continuous-time model forms a point of connection between the models based on the maximization of single-period expected utility of terminal wealth and the more general, multiperiod consumption-portfolio models, and therefore is

---

[7] See Rosenberg and Ohlson [21].

important. However, I view it more as a beginning rather than a final model.

In a recent paper,[8] I extended this basic continuous-time model by allowing for multiple sources of uncertainty in addition to end-of-period return uncertainty. While the concrete examples in that paper focus on the impact of a changing investment opportunity set (i.e., the case where the per period expected returns and covariance structure are changing stochastically over time), I also indicated in a general derivation that similar results would obtain for other sources of uncertainty, for example, multiple consumption goods with uncertain relative prices or when the investor has uncertain wage income. The changing opportunity set is a particularly important type of uncertainty because it only affects intertemporal investors and therefore, its impact would never show itself in a one-period analysis. For example, an investor with a one-period planning horizon facing a specified one-period rate of return structure which includes an interest rate of 8 percent will not change his optimal portfolio if informed that next period's return structure might have an interest rate of either 5 percent or 11 percent, instead of the same 8 percent as this period. However, a multiperiod maximizer facing the same specified one-period rate of return structure will change his *current* portfolio holdings upon being informed of the changed beliefs about future investment rates. And this is so, even though both investors are making an investment decision for one period at a time.

In analyzing this model, it was found that, in contrast to the basic continuous-time model, all investors will not hold the same relative proportions of risky assets, and therefore the standard separation or mutual fund theorem of Markowitz and Tobin will not obtain. However, the first-order conditions are still linear in the demand functions for risky assets, and can, therefore, be solved by matrix inversion. Inspection of these demand functions reveals a rather interesting structure. Namely, the demand for each asset can be written as a sum of terms: the first term is identical to the demand derived in the basic model where the only source of uncertainty is end-of-period returns. Each other term can be identified with a specific additional source of uncertainty. Moreover, differential demand terms have the interpretation of being "hedges" by the investor against these other sources of uncertainty. To see why "hedge" is an appropriate terminology: consider the following.

The actual path of optimal consumption will be a stochastic process because optimal consumption is a function of wealth and the other state variables (e.g., interest rates, prices, and time) which themselves follow stochastic processes. This consumption stochastic process will have an expected time path and a variance around that path. In essence, the derived differential demands for the risky assets are such as to minimize the variance of consumption for a given expected time path. So, for example, if an unanticipated decline in interest rates produces an unanticipated decline in consumption for a given level of wealth, then the investor will tend to hold more of those securities that will produce higher realized returns in the event interest rates do decline. So, in this case, he may hold long maturity bonds. For by doing so, if an unanticipated decline in interest rates does occur, then he will also find himself with an unanticipated higher wealth which will tend to offset the negative impact of the interest rate decline on consumption.

[8] Merton [13].

Indeed, this general pattern repeats itself for each additional source of uncertainty for which securities can be used to hedge. Moreover, the two-mutual fund or separation theorem of the basic model generalizes to a multifund theorem where in addition to the two funds, there is a fund for each additional source of uncertainty. Of course, the term "mutual fund" is used broadly since simple financial instruments like bonds may serve the function of some of the funds.

Thus, in this more general model securities have, in addition to their manifest function of providing an "efficient" risk-return tradeoff for end-of-period wealth, a latent function of allowing consumers to "hedge" against other uncertainties.

These findings suggest what I call a "Consumer Services" model of asset pricing[9] which falls somewhere in between a model that considers only the uncertainty associated with end-of-period wealth (for example, the Capital Asset Pricing Model) and the complete markets model of Arrow and Debreu. The model is rich enough to explain the functions of financial intermediaries and financial instruments while maintaining sufficient specification to be empirically testable. The basic ideas behind the model go as follows below.

If we start with the Arrow-Debreu model with complete markets where there are more securities than states of nature n, then Cass and Stiglitz [3] have proven a "mutual fund" or "separation" theorem which states that there can be constructed n mutual funds or (composite) securities made up of linear combinations of all the securities such that (1) all consumer-investors would be indifferent between having available just these n mutual funds or all the securities; (2) the construction of these composite securities requires no knowledge of consumer preferences, wealth allocations, or their subjective probabilities for each of the states of nature.

While the theorem states "indifferent" when there are no transactions costs and information is freely available to everyone, it is reasonable to presume strict preference for the mutual funds if the number of securities greatly exceeds the number of states. Economies of scale in transactions costs and information gathering and processing make it more sensible to have a centralized compilation of the distributions for each of the primary securities rather than have each investor do it for himself. For the same reason, it would make sense for each mutual fund to have the property of paying a positive amount in one state of nature and zero otherwise (i.e., a set of basic contingent claims) rather than some other more complicated combination which in theory would be equivalent.

Since the number of possible states of nature, n, is very large, such a complete set of markets is economically infeasible. There are three basic reasons: (1) there are the direct costs of operation of so many separate mutual funds; (2) that despite the reduction in the number of securities to be analyzed, the large size of n would make the consumer's information processing costs very large; (3) the occurrence of certain states may be controllable by some consumers (i.e., the moral hazard problem). Hence, some "compromise" is obviously required. To do so, we retain the notion that the mutual fund approach is preferred when there are large numbers of securities and large numbers of relatively small economic units (e.g., consumers), but we restrict the number of funds. Consumers have access to limited amounts of information and have limited abilities to process the information that they have. Further, because of the costs of information

[9] This section was originally presented as a part of "A Reexamination of the Capital Asset Pricing Model," (mimeo) July 1973.

gathering and processing, one would expect the consumer to center his attentions on the major sources of uncertainty which affect his consumption plan. Hence, it is reasonable to assume that these sources of uncertainty can be represented by a finite and not very large number of state variables. Further, one would expect the number and type of mutual funds that would be created would correspond roughly to the number and type of major uncertainties which consumers face. The primary prerequisites for such a fund to be created are: (1) the source of uncertainty must be important to a sufficient number of consumers; (2) it must be possible to have a standardized contract with payoffs in contingencies which are easily recognizable; (3) the source of uncertainty must not be controllable by the consumer(s). Thus, for some of the major uncertainties, it is virtually impossible to construct a financial security which would allow the consumer to hedge against them. Broadly, we would expect to find two types of securities traded: (1) "natural" or primary securities, such as common stocks, which are issued by firms to finance production of real output; (2) financial securities or financial intermediaries created to serve the purposes of the "mutual funds" to be derived [namely, to aid the consumer in achieving a higher (expected) consumption level (through a return to capital), a better intertemporal allocation of resources (by not requiring that savings equals investment at each point in time), a lower level of risk (by providing hedges against the major (common) sources of uncertainty faced by the consumer)].

The individual consumer's demand for assets can be partitioned as follows: for those important sources of uncertainty which he faces in common with other consumers, he will take positions in the mutual funds or financial securities created for that purpose. For those important sources of uncertainty for which no mutual fund exists (either because it is specialized to him or there is inherent moral hazard), he will take positions in those primary securities, if they exist, to hedge. For those sources of uncertainty for which no security is a hedge and for those sources which he simply neglects in his analysis, no security position can help so his (differential) demand for securities will be unaffected.

As an example, if the consumers are one-period maximizers of the utility of mean and variance of end-of-period wealth, then the well-known separation or two-fund theorem obtains. Namely, each investor will be indifferent between selecting a portfolio from all the primary securities and from two funds: (1) the market portfolio of risky assets and (2) a riskless asset. Presumably, such funds or financial instruments would be created since there is an obvious common demand. However, consider an investor who also has (uncertain) labor income which he cannot sell forward to eliminate its risk because of the moral hazard problem. Suppose further that it is a very highly-specialized form of labor which can only be used by one company (or a small number of similar companies). Then, from risk-aversion, it would be natural to suppose that this investor would want to hold less of this company's stock to hedge against unfavorable changes in his labor income. If, on the other hand, there were no security whose outcome was correlated with this particular source of uncertainty, he can do nothing to hedge, and his optimal portfolio would be generated by the two mutual funds alone.

Thus, to a reasonable approximation, most of the aggregate demand for the individual primary securities can be viewed as coming in an "indirect" way through mutual funds, i.e., individual consumers for the most part only purchase a relatively small number of composite financial securities or mutual funds. Mutual fund managers serve the function of purchasing the primary securities to form the portfolio necessary to perform these services. Therefore, the aggregate demand for a primary security will depend on how its return contributes to the formation of these "service" portfolios.

Since the equilibrium expected return on an asset is "determined" in part by the aggregate demand for it, one would expect to find a correspondence between its expected return and the statistical dependence between the asset's return and the various major sources of uncertainty. All risk-averse consumers would prefer to have less uncertainty for the same expected consumption stream, and therefore would "give up" some (expected) return on an asset in return for that asset providing a hedge against some of these uncertainties. Hence, to the extent that any asset's return contributes to (or aggravates) the consumers' attempts to hedge against these uncertainties, one would expect the equilibrium return on that asset to be affected. If, on average, a particular asset's return contributes to consumers' attempts to hedge against a common source of uncertainty, then one would expect the equilibrium expected return on that asset to be differentially lower than on a similar asset which does not provide that "service." This negative differential in expected return can be interpreted as the market "cost" to the consumer for the hedging service provided by this asset. If, on average, an asset's return aggravates consumers' attempts to hedge, then its equilibrium expected return would be differentially higher, and this positive differential in expected return can be interpreted as the market "premium" to the consumer in return for bearing the extra risk caused by holding this asset. A simple illustration of this principle can be found in the CAPM. Since the only source of uncertainty is end-of-period wealth and all investors are assumed to be risk-averse, a given investor would view an asset as providing a ("diversification") service if it lowers the variance of his end-of-period wealth and hence, would accept a lower expected return on this asset than on one which did not provide this service. However, since all investors' optimal portfolios are perfectly correlated, an asset which aids diversification for one investor does so for all investors, and therefore, all investors would accept a lower expected return on this asset. Inspection of the Security Market Line, indeed shows this is the case.

However, with respect to most sources of uncertainty, such unanimity among investors' views of whether an asset contributes to risk or not will be the exception. Thus, one group of consumers may consider a *long* position in an asset as contributing to a reduction of the risks it perceives, while another group may view a *short* position as contributing to a reduction in its risks. Thus, whether the equilibrium expected return on the asset reflects a differential cost or premium will depend on the aggregation of investors' demands, and unless there is a systematic "weak side" to the market, the sign of the differential may fluctuate through time. One example of this type is the Modigliani-Sutch [19] Habitat theory of bond pricing. If a consumer has preferences which induce risk-aversion with respect to wealth, then he will view long-term bonds as risky and would require a market premium over short-term bonds to hold them. However, if a consumer has preferences which induce risk-aversion with respect to income, then short-term bonds are risky to him, and he would require a market premium over long-term bonds to hold them.

Thus, with respect to the uncertainty about future interest rates, the differential expected return between long- and short-term bonds could be of either sign.

To determine the types of securities one would expect to find and the sources of differentials in expected returns, it is necessary to establish what the important uncertainties are for a typical consumer in making his plan. Although not a complete listing, the following seven items would seem to cover most of the important common sources of uncertainty for a consumer:

(S.1) Uncertainty about his own future tastes;

(S.2) Uncertainty about the menu of possible consumption goods that will be available to the future;

(S.3) Uncertainty about relative prices of consumption goods;

(S.4) Uncertainty about his labor income;

(S.5) Uncertainty about future values of nonhuman assets;

(S.6) Uncertainty about the future investment opportunity set; i.e., the future rates of return which can be earned on capital;

(S.7) Uncertainty about the age of death.

While all of these have been considered in one model or another, it is important to note that those models based on the criterion of maximizing the expected utility of end-of-period wealth explicitly take into account only the uncertainty associated with nonhuman wealth. Included in this class of models is the CAPM.

Even though all these uncertainties are important to the consumer, not all will differentially affect security prices or returns. It is difficult to imagine a financial security which could reduce the uncertainties associated with one's own future tastes or the menu of possible consumption goods in the future. While uncertainty about the age of death is an important problem for all consumers and life insurance was created in response to this demand, the event of death is probably reasonably statistically independent across people, and it is unlikely that the returns on securities (other than life insurance policies) would be statistically dependent on the event of an individual's death. Hence, one would not expect this source of uncertainty to have differential effects on security prices. The risks associated with labor income can be completely eliminated if the consumer could sell forward his wage income in the same way shares are issued on nonhuman capital. But, because of the moral hazard problem, it is difficult for the consumer to do so. While some of the individual risk can be eliminated by disability and life insurance and by "investing" in education to make his labor more substitutable across firms, there still will be systematic risk due to (unanticipated) shifts in capital and labor's relative shares (i.e., the wage-rental ratio). This uncertainty could produce a differential demand for shares in labor-intensive versus capital-intensive industries. Similarly, inflation risk may cause differentials in demand between different maturity "money" securities. Although information costs and the uncertainties about tastes and future products prohibit complete future markets for consumption goods, it is reasonable to expect consumers to differentiate broad classes of consumption (e.g., housing, food, transportation, clothing, and recreation) and hence, differentials in demand for shares in different industries could occur as the result of uncertainty about relative consumption good prices. The standard end-of-period wealth uncertainty will induce differential demands for those securities which aid diversification. Finally, if there is uncertainty about the rates of return which will be available in the future, differential demands may occur between long- and short-term bonds and between shares of firms whose returns are sensitive to shifts in capitalization rates versus ones that are not.

If these are the sources of uncertainty common to most investors, then we can identify a set of mutual funds which would be (approximately) sufficient to span the space of consumers' optimal portfolios. Specifically, we might identify these funds to be: (1) the "market" portfolio; (2) a (short-term) riskless asset; hedging portfolios for *unanticipated;* (3) shifts in rates of return; (4) shifts in the wage-rental ratio; (5) changes in prices for basic groups of consumption goods. Further, consumer demands for individual securities can be written as if they came indirectly through the demands for these mutual funds. Hence, the equilibrium expected return on a security will be a function of the expected return on each of these funds and the statistical dependence between the security's return and the return on each of these funds.

Indeed, this model is testable in its continuous-time formulation because, in equilibrium, the expected excess return on an individual security will be equal to a weighted sum of the excess returns on the mutual funds where the weights will equal the instantaneous multiple regression coefficients between the return on the individual security and the returns on the funds. This result is a natural generalization of the standard Security Market Line to a Security Market Hyperplane:

$$\bar{R}_i - r = \Sigma_j^m \beta_i^j (\bar{R}^j - r).$$

For my last topic, I turn to the theory for pricing financial instruments and in particular the pricing of corporate liabilities. It is here where the continuous-time analysis has had its greatest impact. In what may be one of the most important contributions to Finance in this decade, Fischer Black and Myron Scholes[10] used the continuous-time analysis to deduce a formula for common stock options. In essence, they were able to demonstrate that by following a specified dynamic portfolio policy consisting of mixtures of positions in the underlying common stock and riskless borrowing or lending, they could exactly replicate the payoff structure associated with a call option.

Given this demonstration, it is immediately obvious that this portfolio strategy is a perfect substitute for the option, and indeed, any two of the three securities involved can be combined in an appropriate portfolio strategy to exactly replicate the payoff structure of the third. Hence, given the prices of any two of the securities (for example, the stock and riskless bond), the third security's price (in this case the option) is uniquely determined, i.e., their analysis is clearly a relative pricing theory.

Having solved the call option case, it is straightforward to see that the same technique can be used to price other types of options, and from there, it is not a long step to recognize that their approach is equally valid when applied to the firm as a whole to price its entire capital structure. As a result, their analysis has led to a unified

---

[10] See Black and Scholes [1]. For further discussion, see Merton [12].

theory for pricing of virtually any financial claim on the firm.[11]

I view the Black-Scholes contribution in three parts. The first and most important part is the initial insight in setting up the problem. The second part is the development of a quantitative formula based for the most part on observable or reasonably estimated variables. In particular, rather surprisingly, consumer preferences and the expected return on the underlying stock do not enter the formula. Moreover, one need not even assume that the market is in equilibrium. Of course, a quantitative formula is important because it allows for empirical testing and for its direct use in applications. This is one of those rare cases where a piece of analysis conceived entirely in theory has had an immediate and important impact on actual operations. Indeed, the Black-Scholes formula is widely used by most firms trading in the option market. The third part is that the Black-Scholes analysis has led to some important qualitative propositions. For example, it is widely held that the Modigliani-Miller irrelevance-of-financing theorem does not hold when there is a positive probability of bankruptcy because in such cases, personal borrowing and corporate borrowing are no longer perfect substitutes for one another. The Black-Scholes type analysis can be used to demonstrate that this conclusion is false. Namely, while no *fixed* portfolio strategy can replicate the nonlinear payoff structure of such corporate debt, a continuous-time, dynamic portfolio strategy can.

Also, some earlier studies[12] attempted to use warrant price and stock price comparisons to estimate investor's expectations for the stock or investor's risk preferences. But since the Black-Scholes formula requires neither as inputs for pricing warrants, it clearly demonstrates that such attempts are doomed to failure.

Since the pricing formula does not depend upon knowledge of the expected return on the stock, both Bulls and Bears who agree on the other inputs, will agree on the price of the option relative to the stock, or the prices of the individual elements of the capital structure relative to the total value of the firm. This demonstrates that the investor's decision as to what investment action to take in regard to a specific firm can be separated into two parts. First, in an evaluation of the firm as a whole, is it a good or bad investment? Second, given this decision, what financial claim on the firm is the best vehicle for either being long or short in the firm? Since the Black-Scholes formula requires an estimate of the variance on the stock, it has created a demand for further research into estimation techniques for variances.

While much of the research in this area is still in progress,[13] I believe that before long we will have answers to many of the long outstanding questions about corporate financing policy including a rationale for debt targets and debt capacity, and a quantitative analysis of the optimal debt structure when interest payments are deductible.

Moreover, most of the assumptions required in the original Black-Scholes derivation have been substantially weakened in more recent research with no significant changes in their more important conclusions.

While time does not permit me to cover all the continuous-time research going on in

---

[11] See Merton [12, 15] and Ingersoll [8].

[12] See Sprenkle [86].

[13] An excellent survey can be found in Smith [25].

Finance, I will simply mention the Solnik [27] analysis in international capital markets, the Richard [20] analysis of the demand for insurance, the Fischer [7] analysis of index-linked bonds, and Scheffman [24] analysis of the optimal investment decision by firms.

In summary, the continuous-time mode of analysis has proved fruitful in analyzing some of the basic problems in Finance. The costs associated with getting these results are the twin assumptions that markets are open most of the time and that the stochastic processes can be described by either diffusion or compound Poisson processes. The benefits are generally sharper results that are easier to interpret than in the discrete-time case, and an enormous literature on these stochastic processes which allows one to analyze rather complex models and still get quantitative results.

REFERENCES

[1] Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy,* vol. 81 (1973), pp. 637-659.

[2] __________. "The Valuation of Option Contracts and a Test of Market Efficiency." *Journal of Finance,*vol. 27 (1972), pp. 399-417.

[3] Cass, D., and J. Stiglitz. "The Structure of Investor Preferences and Asset Returns, and Separability in Portfolio Allocation: A Contribution to the Pure Theory of Mutual Funds." *Journal of Economic Theory,* vol. 2 (1970), pp. 122-160.

[4] Cox, J., and S. Ross. "The Pricing of Options for Jump Processes." *Journal of Financial Economics,* (forthcoming).

[5] Fama, E. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance,* vol. 25 (1970), pp. 383-417.

[6] __________. "Multiperiod Consumption-Investment Decisions." *American Economic Review,* vol. 60 (1970), pp. 163-174.

[7] Fischer, S. "The Demand for Index Bonds." *Journal of Political Economy,* vol. 83 (1975), pp. 509-534.

[8] Ingersoll, J. "A Theoretical and Empirical Investigation of the Dual Purpose Funds: An Application of Contingent Claims Analysis." *Journal of Financial Economics,* (forthcoming).

[9] Magill, M., and G. Constantinides. "Portfolio Selection with Transactions Costs." *Journal of Economic Theory,* (forthcoming).

[10] Merton, R. C. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *Review of Economics and Statistics,* vol. 51 (1969), pp. 247-257.

[11] __________. "Optimum Consumption and Portfolio Rules in a Continuous Time Model." *Journal of Economic Theory,* vol. 3 (1971), pp. 373-413.

[12] __________. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science,* vol. 4 (1973), pp. 141-183.

[13] __________. "An Intertemporal Capital Asset Pricing Model." *Econometrica,* vol. 41 (1973), pp. 867-887.

[14] __________. "Appendix: Continuous-Time Speculative Processes." In P. A. Samuelson, "Mathematics of Speculative Price," *SIAM Review.*

[15] __________. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance,* vol. 29 (1974), pp. 449-470.

[16] ________. "Option Pricing when Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics*, (forthcoming).

[18] Merton, R. C., and P. A. Samuelson. "Fallacy of the Log-Normal Approximation to Optimal Portfolio Decision-Making over Many Periods." *Journal of Financial Economics*, vol. 1 (1974), pp. 67-94.

[19] Modigliani, F., and R. Sutch. "Innovations in Interest Rate Policy." *American Economic Review*, vol. 56 (1966), pp. 178-197.

[20] Richard, S. "Optimal Consumption, Portfolio and Life Insurance Rules for an Uncertain Lived Individual in a Continuous-Time Model." *Journal of Financial Economics*, vol 2 (1975), pp. 187-204.

[21] Rosenberg, B., and J. Ohlson. "The Stationary Distribution of Returns and Portfolio Separation in Capital Markets: A Fundamental Contradiction." University of California, Berkeley (unpublished, 1973).

[22] Rubinstein, M. "The Strong Case for the Generalized Logarithmic Utility Model as the Premier Model of Financial Markets." W. P. #34, Institute of Business and Economic Research, University of California, Berkeley (1975).

[23] Samuelson, P. A. "Proof that Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review*, vol. 6 (1965), pp. 41-49.

[24] Scheffman, D. "'Optimal' Investment under Uncertainty." University of Western Ontario (unpublished, 1975).

[25] Smith, C. "Option Pricing: A Review." *Journal of Financial Economics*, (forthcoming).

[26] Sprenkle, C. "Warrant Prices as Indicators of Expectations and Preferences." In *The Random Character of Stock Market Prices*, edited by P. Cootner. Cambridge: M.I.T. Press, pp. 412-474.

[27] Solnik, B. *European Capital Markets*. Lexington, Mass: Lexington Books (1973).