WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

CONTINUOUS TIME PRODUCTION ECONOMIES UNDER INCOMPLETE INFORMATION
I: A SEPARATION THEOREM

by

Gerard Gennotte*

November 1984

Sloan School of Management Working Paper No. 1612-84

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

CONTINUOUS TIME PRODUCTION ECONOMIES UNDER INCOMPLETE INFORMATION
I: A SEPARATION THEOREM

by

Gerard Gennotte*

November 1984

Sloan School of Management Working Paper No. 1612-84

I am grateful to Douglas T. Breeden, John C. Cox, Stanley Fischer, Chi-fu Huang, Terry A. Marsh, Robert C. Merton and Paul R. Samuelson for helpful comments and substantive suggestions. I would also like to thank the members of the finance group at the Sloan School of Management for their contribution. Support from the College Interuniversitaire pour les Sciences du Management is gratefully acknowledged.

*Sloan School of Management, M. I. T.

ABSTRACT

Models of asset pricing generally assume that the variables which characterize the state of the economy are observable. For example, the Capital Asset Pricing Model (CAPM) of Sharpe, Lintner and Mossin, Merton's Intertemporal Asset Pricing Model and Ross's Arbitrage Pricing Theory establish relations between distributional parameters of the rates of return on assets and rates of return on "well diversified" portfolios.

However, the distributional properties of asset prices that are relevant for portfolio decisions are in general not observable, and therefore must be estimated. As discussed, for example, in Merton (1980), the problem of estimating expected returns is particularly difficult and estimation errors are likely to be substantial. In this light, it is reasonable to question whether the assumption of observability of expected returns and other relevant state variables causes significant mis-specification in equilibrium models of asset prices.

This paper has two objectives. The first is the derivation of optimal estimators for the unobservable expected instantaneous returns using observations of past realized returns. The second is the analysis of the optimal portfolio decisions and the characteristics of general equilibrium in which these optimal estimators are used. The main result establishes the separability of the estimation and investment decision problems for general utility functions. Agents derive the conditional distribution of future returns and then select their optimal portfolio using the derived expected returns. The estimators of expected returns are in general not consistent, i.e. the estimation error does not tend to disappear asymptotically (in contrast with the case considered by Merton (1980)). The effects of the estimation error, therefore, cannot be ignored even if realized returns are observed continuously over an infinite time period.

INTRODUCTION

In his seminal papers, Merton (1971), (1973) characterizes the equilibrium expected return on assets and the optimal portfolio demand in an intertemporal capital asset pricing model. Cox, Ingersoll and Ross (1978) extend the analysis to a general equilibrium framework where financial claims prices are derived endogeneously. These analyses, as does the original Sharpe-Lintner-Mossin CAPM, make the assumption that the expected return on investment is known with certainty.

Merton (1980) argues that the expected return on the aggregate market portfolio is unlikely to be constant over time and that even the optimal estimator of this expected return does not yield precise inferences when the observation interval is finite. It is thus unrealistic to assume that the mean return on investment is a known function of observable state variables. Instead, the mean return can only be estimated with error using the observable instantaneous realized returns on investment[1].

The estimation of expected returns on investment is a prerequisite for tests of the CAPM and practical applications of the theory. It also raises important theoretical issues. When the true expected returns are not observable, the usual dynamic programming rules for portfolio optimization are not applicable. When selecting their optimal portfolio, investors have to take into account the fact that they have only estimates of the true expected returns and that variations in the estimators will be caused by the variations in the underlying state variables as well as by fluctuations in the estimation error. Since investors form their portfolios using estimators, the characteristics of equilibrium itself depend on the properties of these estimators. Equilibrium

[1] In real markets, other price series reflect information not contained in past realized returns. In the stylized models of Merton (1971), (1973), Cox, Ingersoll and Ross (1978) and in the analysis here, it is assumed that past realized returns contain all available information.

levels of investments, the term structure of interest rates and contingent claims prices will all be functions of the "perceived" (or estimated) state variables.

The analysis here is intended to improve upon the modest state of our current knowledge regarding portfolio choice and general equilibrium under incomplete information. With the exception of Klein and Bawa (1976), (1977), Williams (1977), Feldman (1983) and De Temple (1984), the heuristic approach in previous research has been to assume implicitly that the portfolio choice problem can be solved in two steps: parameters are first estimated and then portfolios are chosen conditional on these parameters estimates. This separation of the estimation and optimization steps is optimal when a property which Simon (1956) and Theil (1964) termed "certainty equivalence" applies[2]. In this case, the state variables can be replaced in the optimization problem by their certainty equivalent values (i.e. the least square estimators). In a discrete time framework, the certainty equivalence applies if the objective function is quadratic and the process is a linear function of unobservable state variables. In this case, the state variables can be replaced in the optimization problem by their certainty equivalent values (i.e. the least square estimators).

The separation theorem, which is the main result of this paper, establishes that the properties of the certainty equivalence principle can be obtained in continuous time without restrictions on the utility function. Moreover, it is shown that the conditional estimates of expected returns are governed by a system of differential equations similar to the one which characterizes the variations of the true expected returns. This result implies that the dynamic programming methodology of Cox, Ingersoll and Ross (1978) can be generalized to analyze the properties of general equilibrium when the state variables are not observable.

[2] It has been widely used in macroeconomic theory, Lucas and Sargent (1981) provide a survey and several examples.

Except in the log utility case, it is shown that the level of investment and the term structure of interest rates depend upon future variations of the derived investment opportunity set, which arise from randomness in both the true opportunity set and the estimation error. This contrasts with De Temple's (1984) analysis within a general equilibrium framework in which the existence and form of asset prices as functions of estimated state variables are assumed rather than derived, and in which agents ignore information in realized returns when making investment decisions. My results are more general than those of Feldman (1983) who studies a general equilibrium with one unobservable state variable and the effect of estimation error on the term structure in that general equilibrium,but does not develop the implications for the estimation of expected returns. Feldman assumes logarithmic preferences; unfortunately, this is precisely the case where the non-observability of state variables won't affect real investments, portfolio decisions and the term structure of interest rates.

In the analysis to follow, I relax the fundamental assumption of observability of the variables that characterize the state of the economy. In the model, the investment opportunity set is determined by a finite number of random variables. Instantaneous expected returns on physical investment vary stochastically over time. I first consider the problem of inferring expected returns from observed instantaneous returns on investment. Tools of non-linear filtering theory are introduced to solve this difficult estimation problem. The optimal estimators for the expected returns are generalizations of continuous time Kalman filters. Unlike the case of lognormal returns, I show that the precision of the estimators does not necessarily increase with time. Economic intuition supports this result: since expected returns vary randomly over time, past observations will not necessarily contain enough information to assess perfectly the path of the expected returns. Imperfect information is thus shown to be not just a transitory problem. The optimal inference process and the portfolio choice decision are shown to be separable: that is, agents can optimally derive the conditional distribution of future returns and then select their portfolio by replacing the unobservable expected returns with their conditional expectations.

The paper is organized as follows:

Section I describes the continuous time economy and the structure of the uncertainty. Section II states the inference problem and describes the logical steps of the reasoning which leads to its solution. I then consider the simple case of one technology and one state variable. The solution is closely related to standard economic results and provides a useful intuition. It is rigorously proved using the general results of section III. Section III states the non-linear filtering results and the solution to the general inference problem. The mathematical problem was solved by Liptser and Shiryayev (1978) and the proofs are therefore not presented. This introduction to the use of non-linear filtering theory for economic analysis is brief but self-contained. The power of the technique and its applicability to a broad range of economic problems warrant its extensive development in that section. However, the reader is alerted that section III is not necessary for the understanding of the substantive economic argument.

Section IV contains the main result of separability and establishes the general equilibrium conditions. The results are contrasted with the conclusions of prior research.

Finally, the conclusion summarizes the implications for general equilibrium and further research currently underway. A later companion paper[3] will develop these issues.

[3] Continuous Time Production Economies under Incomplete Information II: Contingent Claims Prices.(forthcoming)

I The Economy

The economy described here is similar to the one considered by Cox, Ingersoll and Ross (1978) with the fundamental difference that the expected returns on the risky technologies are not observable. There is a finite number of physical production technologies. The instantaneous expected returns on investments are stochastic and not observable. In order to establish the statistical inference results, the structure of uncertainty is formally defined in mathematical terms (section I-1). However, the assumptions on the investment opportunity set and on agents' preferences are standard in continuous time general equilibrium analysis.

I-0 Notations

I generally adopt the following notations: capital letters for matrices, $B(t)$ and $\Sigma(t)$ for example, and lower case characters for scalars and vectors, σ and r for example.There are, however, some exceptions motivated by custom: W and J denote wealth and the indirect utility function (both scalars).

Partial derivatives are denoted by subscripts: J_W and the transpose of a matrix by the letter t: $B^t(t)$.

I-1 Structure of the uncertainty

The time span of the continuous time economy is [0,T] (i.e. agents maximize their expected utility function over the time interval [0,T]). The uncertainty in the economy is generated by the Brownian motion X_t. At each point in time, the state of the economy is characterized by the path of X previous to that instant: $\{X_\tau,\ 0\leq\tau\leq t\}$. All the possible paths of X constitute the set of possible events and investors have a common probability measure on that set.

Let (Ω,F,P) be a complete probability space, i.e. each element ω of Ω denotes a complete description of the economic environment (a "state of nature"). Agents are endowed with a common probability measure P. We assume that there is defined on the probability space (Ω,F,P) a s+n-dimensional Brownian motion $X=\{X_\tau, 0\leq\tau\leq T\}$, its components, without loss of generality, assumed to be independent.

Let F_t^X be the sigma algebra generated by $\{X_\tau,\ 0\leq\tau\leq t\}$.

Assume that $F_T^X = F$ and that F_t^X is augmented by all P-negligible sets for all t in [0,T].

I-2 Investment opportunities

There is a single physical good, the numeraire, which may be allocated to consumption or investment. Physical production consists of s technologies. The transformation of an investment of a vector q of the good in the s technologies follows the system of stochastic differential equations:

$$dq = I_q \mu_t \, dt + I_q \Sigma(t) dZ \qquad \text{(I-2.1)}$$

The vector of instantaneous expected rates of return, μ_t, on the production process is governed by the process:

$$d\mu = [\, a(t) + A(t)\mu \,]\, dt + B_z(t)\, dZ + B_y(t)\, dY \qquad \text{(I-2.2)}$$

Where I_q is an s*s diagonal matrix whose ith element is the ith component of q(t). $\Sigma(t)$ is a bounded , non stochastic, s*s matrix function of t. $\Sigma(t)\Sigma^t(t)$, the variance-covariance matrix of physical rates of return is positive definite for all t in [0,T]. a(t), A(t), $B_z(t)$ and $B_y(t)$ are respectively a bounded s vector, two s*s matrices and an s*n matrix, all of which are functions of time. The vector of independent Brownian motions Z_t has as components the first s components of X_t. The last n components of X_t form a vector denoted by Y_t.

Physical production is subject to two different sources of uncertainty. The first one, Z, affects current output dq and also future expected returns μ_t. The second one, Y, affects only future production through its effect on future expected returns. For example, in the context of a wheat economy a new virus affecting wheat reduces current output but also future expected output. In contrast, the discovery of a fertilizer does not modify current output but increases expectations of future output.

Agents can observe instantaneous returns on investments dq and are assumed to know the deterministic functions of time $\sum(t)$, a(t), A(t), $B_z(t)$ and $B_y(t)$. However they are not able to observe the random process μ_t. Consequently, the Brownian motion vectors Z and Y are not observable.

In addition to direct investment in the risky technologies, agents can trade continuously in contingent claims (which are not formally introduced here) and borrow or lend risklessly at the instantaneous rate r. The instantaneous riskless rate r is not specified at this point because it will be endogeneously determined in equilibrium. Agents can continuously invest any amount in the s technologies (i.e. there are no problems of indivisibilities), there are no restrictions on borrowing and lending.[4]

Markets are competitive and agents behave as price takers. Trading takes place continuously at equilibrium prices only, with no adjustment or transaction costs.

[4]The physical production technologies have constant returns to scale. There is thus no incentive to create firms to pool resources. If these existed, however, agents' investment decisions would be unaffected.

I-3 Agents

There is a fixed number of individuals, identical in their preferences and endowments. Agents are endowed with a common probability measure and agree on the characterization of the stochastic processes governing the economy.

Agents are characterized by their initial wealth W_0 and their preferences u. They seek to maximize their expected lifetime time-additive utility of consumption conditional on all available information as of date t (i.e. the rates of return on investment observed in the past):

$$E \{ \int_t^T u[c(s),s] \, ds \mid F_t^q \}$$

subject to their budget constraint, where u is an increasing, concave, twice differentiable von Neuman-Morgenstern utility function and c(t) is the consumption rate selected at date t.[5]

[5]See section IV for the technical conditions imposed on c(t).

II The inference process

Introduction

This section presents the derivation of the conditional distribution of future returns and establishes its important properties. An example is used to illustrate the results and the similarities with standard least square estimation problems.

II-1 The conditional distribution

Agents seek to extract (or "filter") information on future expected instantaneous returns from their observation of past returns. At time 0, they view the distribution of μ_0 as a Gaussian distribution with mean vector m_0 and variance-covariance matrix V_0. As time evolves agents continuously update their beliefs[6]. I will assume that agents are rational, i.e. they derive the optimal estimator using all available information. The theorems of section III, due to Liptser and Shiryayev (1978) are continuous time results. They are essentially a generalization of the theory of filters due to Kalman and Bucy (1961). I present here the important steps of the derivation implicitly assuming discrete infinitesimal observation intervals. The results hold only in continuous time but the reasoning in discrete time is perhaps more intuitive.

Over the infinitesimal interval [0,dt], agents observe the realized returns on investments which are correlated with the change in the expected returns, $d\mu$. These two vectors are (instantly) Gaussian. The distribution of a

[6] The estimation process is thus a continuous Bayesian updating of beliefs. Agents' prior belief is that the distribution of the expected return at time zero, μ_0, is Gaussian. As is shown below, this assumption leads to a Gaussian conditional distribution for μ_t at any time t, $0 \leq t \leq T$.

Gaussian vector conditional on a correlated Gaussian vector is Gaussian. The derivation of the conditional distribution is, therefore, a classic statistics problem and standard econometric analysis can be used to interpret the results in the one dimension case treated in II-2.

Although the random variables q and μ are not Gaussian, the distribution of μ conditional on q is Gaussian (Theorem 1)[7]. Hence the expectations vector and the variance-covariance matrix characterize the distribution. The conditional moments at date t+dt are equal to the conditional moments at date t plus the instantaneous change (or revision) occasioned by the observation of q over the (infinitesimal) period [t,t+dt]. The instantaneous changes in the estimated expected returns vector, dm, and in the variance-covariance matrix, dV, are given by the following equations (Theorem 2):

$$dm = [a(t) + A(t)m_t]\,dt + [B(t)\Sigma^t(t) + V(t)]\,[\Sigma(t)\Sigma^t(t)]^{-1}\,[I_q^{-1}dq - m_t dt] \qquad \text{(II-1.1)}$$

$$\begin{aligned} dV = {} & \{A(t)V_t + V_t A^t(t) + B_z(t)B_z^t(t) + B_y(t)B_z^t(t)\}\,dt \\ & - \{[B(t)\Sigma^t(t) + V_t] * [\Sigma(t)\Sigma^t(t)]^{-1}\,[B(t)\Sigma^t(t) + V_t]\}\,dt \end{aligned} \qquad \text{(II-1.2)}$$

μ_t is a diffusion process and V_t is a deterministic function of time. Equation II-1.2 is a multidimensional differential equation of the Ricatti type.

These expressions are analogous to the classic statistics case where the first terms are the unconditional moments and the second terms result from the inference process. The analogy is discussed in detail in section II-2.

[7] This result stems from the structural form of $d\mu$ and dq. It holds for the wider class of processes considered in section III. Basically, the coefficients Σ, a, A, B_z and B_y of equations I-2.1 and I-2.2 may depend on q as well as on time. The crucial assumption is the linear dependence of dq and $d\mu$ on μ and the fact that the variance-covariance matrices do not depend on μ.

The first term of equation II-1.1 is the instantaneous change in the expected returns vector conditional on realized returns previous to t. The second term represents the amount by which the estimate of the mean is changed by a "surprise" in output. The current dq contains information on the expected return at date t, μ_t, and on the change in the expected return $d\mu$.

The evolution of the variance-covariance matrix through time (II-1.2) is determined by two factors. The first term corresponds to the uncertainty concerning the deterministic part of the variation $d\mu$ resulting from uncertainty about the current value μ_t. The second term reflects the information accrual and actually reduces the variance-covariance matrix[8]. In general, the variance-covariance matrix V_t does not converge to zero as the observation interval tends to $[0,\infty]$. In the particular case of stationary coefficients (section II-1), it is shown that the variance tends to a non-zero limit. The estimation error persists through time and affects equilibrium levels of investment in both the short term and the long term.

The assumption that the variance-covariance matrix of investment returns evolves deterministically while expected investment returns are stochastic might appear to be very specialized. If both expected returns and covariances change stochastically over time, a major added problem is one of identification. However, if enough is specified about joint movements in expected returns and covariances for identification, the analysis here can be extended to take account of the estimation of both parameters. An alternative approach is that taken in Gennotte and Marsh (1984) where we posit variance movements to be a

[8] The uncertainties on past and future expected returns are thus correlated. Consider the case where the same set of state variables Z determines the variations of μ and q. If at any instant expected returns are known for certain, agents can deduce from the realized returns the exact expected returns in the infinitesimal periods following and preceding t. They are thus able to determine expected returns at any point in time. More generally, fluctuations in the variance of the expected returns estimators are intertemporally related.

deterministic function of the investment level (i.e. as in Cox's (1975) constant elasticity of variance model for stock prices). Equilibrium theory is then used to specify the relation between the shifts in expected returns on equity claims in these investments and the variance changes.

The innovation process Z'_t is defined as the normalized deviation of the return from its conditional mean:

$$dZ' = \Sigma^{-1}(t)\ [I_q^{-1} dq - m_t\ dt] \tag{II-1.3}$$

and $Z'_0 = 0$

Substituting dq with its expression in I-2.1, dZ' is given by:

$$dZ' = \Sigma^{-1}(t)\ [\mu_t - m_t]\ dt + dZ$$

Z' is a Brownian vector: dZ' is Gaussian with a zero mean vector and a variance-covariance matrix per unit time equal to the identity matrix; and the increments dZ' are uncorrelated. The Brownian vector Z is not observable, the innovation process Z', however, is derived from observable processes and thus observable.

Substituting dZ' for dZ, the return over the next infinitesimal period becomes:

$$dq = I_q\, m_t\ dt + I_q\ \Sigma(t)\ dZ' \tag{II-1.4}$$

Hence, the paths of m and Z' uniquely determine the path of q.

The path of the conditional expectation m_t is generated by the path of the innovation process Z':

$$dm = [a(t) + A(t)m_t]\ dt + [B(t)\Sigma^t(t) + V(t)]\ \Sigma^{t-1}(t)\ dZ'$$

Intuitively, the information contained in the path of Z' is equivalent to the information contained in the path of q. Formally, the information structure generated by $\{q_0, Z'_s, 0 \leq s \leq t\}$ is equivalent to the information structure generated by $\{q_s, 0 \leq s \leq t\}$ (Theorem 3).

The crucial step now follows in a straightforward fashion. Since Z' is a Brownian motion, it contains no information on the future variations of q and m.

The distribution of the variations dm_t depend on the conditional variance-covariance matrix V_t. V_t is a deterministic function of time (II-1.2), the distribution of the instantaneous change dm_t is thus characterized by the state vector m_t. The distribution of dm_t and the conditional distribution of the instantaneous return $I_q^{-1} dq_t$ are therefore fully characterized by the current value m_t. In other words, the distribution of the change in the opportunity set perceived by investors over the next infinitesimal period is perfectly determined by m_t. Investors do not have to remember past returns to form their expectations of future returns. Hence, the system $[q_t, m_t]$ is Markov[9], i.e. q_t and m_t determine the probability distribution of q and m over the next infinitesimal interval [t,t+dt].

[9] A stochastic process $\{X_s, 0 \leq s \leq T\}$ is a Markov process if and only if the probability of its future evolution depends only on its present value and not on past history.

II-2 An example in dimension one.

To illustrate the results, consider an economy with a single technology which returns are determined by a single state variable z. The instantaneous return on investment and the expected return at date t are given by:

$$dq = q\,\mu\,dt + q\,\sigma\,dz \qquad \text{(II-2.1)}$$

$$d\mu = a\,\mu\,dt + b\,dz \qquad \text{(II-2.2)}$$

where a and b are constant scalars and σ is a strictly positive constant[10]. The expectation and variance of μ_t conditional on observing realized returns are respectively m_t and v_t. The distribution of μ at time zero is assumed to be Gaussian with mean m_0 and variance v_0. The changes in the conditional expectation are given by:

$$dm = a\,m\,dt + \left(b + \frac{v}{\sigma}\right) \frac{1}{\sigma} \left[\frac{dq}{q} - m\,dt \right] dz \qquad \text{(II-2.3)}$$

The innovation process z' is defined by:

$$z'_0 = 0$$

and $$dz' = \frac{1}{\sigma} \left[\frac{dq}{q} - m\,dt\right]$$

Substituting in II-2.1 and II-2.3, dq and dm are rewritten as:

[10]The derivation of the conditional moments is rigourously the same if the coefficients σ, a and b are deterministic functions of q and t. The case of an Ornstein-Uhlenbeck process for μ belongs to that class. But the Ricatti equation which characterizes v_t is solved explicitly only in the constant coefficients case.

$$dq = q\, m\, dt + q\, \sigma\, dz' \qquad \text{(II-2.4)}$$

$$dm = a\, m\, dt + (b + \frac{v}{\sigma})\, dz' \qquad \text{(II-2.5)}$$

Equation II-2.3 can be interpreted as:

$$dm = d[E_t(\mu_t)] + \frac{Cov(dq, \mu_{t+dt})}{Var(dq)} \left[\frac{dq}{q} - m_t dt \right] \qquad \text{(II-2.6)}$$

II-2.6 has the same form as the usual inference equation for Gaussian variables and m_t is also the Least Squares Estimator. The first term is the ex-ante change in the conditional expectation (i.e. before observing the current dq). The second is the innovation in the instantaneous return multiplied by the ratio of the covariance of the two variables and the variance of the signal dq. This ratio is similar to the beta of regression equations[11]. We see that if the innovation is positive and if the correlation between μ and dq is positive, agents will adjust their expectation upward.

The conditional variance v is determined by v_0 and dv:

$$dv = \left[(2\, av + b^2) - [b + \frac{v}{\sigma}]^2 \right] dt \qquad \text{(II-2.7)}$$

v_t is thus a predictable function of time.

Equation II-2.7 can be interpreted as:

$$dv = d[Var_t(\mu_t)] - \frac{[Cov(dq, \mu_{t+dt})]^2}{Var(dq)}$$

This equation is also similar to the usual inference equation for the

[11] It can be graphically interpreted as the projection coefficient of μ on dq.

conditional variance. The first term is the ex-ante increment in the conditional variance corresponding to the unobservable variation of μ from t to t+dt. The second term is always negative and represents the reduction in variance due to the information dq. The variance decreases if the new information is precise enough to outweigh the additional noise due to the variation of μ, $d\mu$.

Equation II-2.7 is of the well-known Riccati type of differential equations. It can be rewritten as:

$$-\frac{dt}{\sigma^2} = \frac{dv}{v(v-2w)}$$

Where w is defined as: $w = a\sigma^2 - b\sigma$.

Integrating with v_0 given yields:

$$v_t = \frac{v_0 \sigma^2}{v_0 t + \sigma^2} \qquad \text{if } w=0$$

and

$$v_t = 2w \left[1 - \frac{1}{1 + \frac{1}{2w-v_0} v_0 \exp(\frac{2wt}{\sigma^2})} \right] \qquad \text{if } w\neq 0$$

If w is negative, the limit of v as t tends to infinity is zero. If w is positive, the limit is 2w. The graphs of figure 1 show that if the initial variance v_0 is larger than 2w, the variance decreases uniformly.

The precision, defined as the inverse of the variance, does not necessarily increase with time and the estimator m_t is not consistent in general. If μ_0 is perfectly known (i.e. V_0=0), the uncertainty on μ_t disappears completely.

In a more general case where the coefficients depend on time, if μ becomes a constant from time t on (i.e. a=b=0) equations II.2.3 and II-2.7 show that the estimator μ_t still fluctuates but its variance decreases uniformely. There is no

FIGURE 1

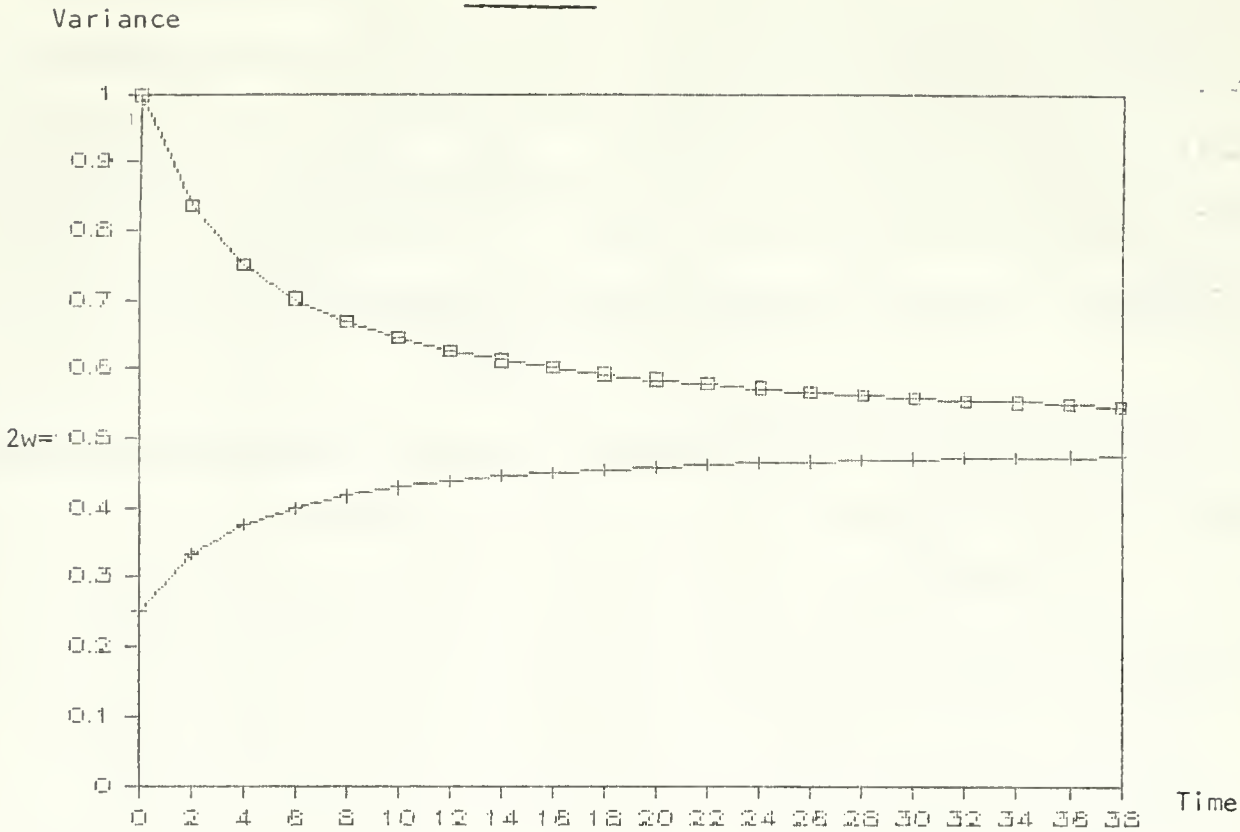

Case w≠0

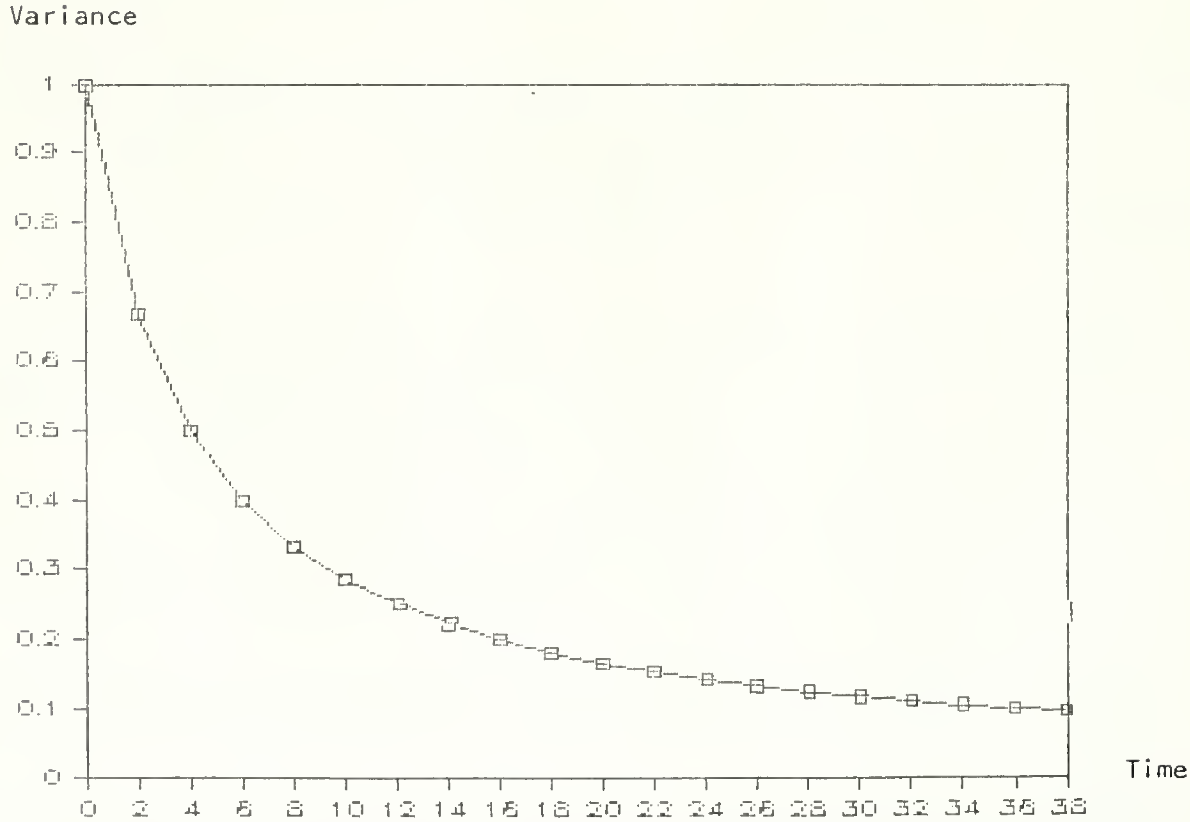

Case w=0

uncertainty on $d\mu$ and the information accruals on μ_t allow agents to refine their estimator of expected returns. Indeed, this is the case considered by Merton (1980), his discrete time estimator and the continuous time one derived here are consistent , i.e. they converge to the true value as the observation interval [0,T] tends to infinity.

III Concepts of non-linear filtering

III-0 Introduction

This section presents some fundamental results of non-linear filtering theory applicable to economic analysis. It is not, strictly speaking, necessary to the understanding of the analysis. This introduction to filtering theory has not appeared before in the substantive context of the financial economics literature (Krishnan (1984) provides an introduction illustrated by examples from physics). The power of its results make it a useful tool for economic analysis. Section II presents the logical steps of the reasoning in a more intuitive fashion.

Three important results of non-linear filtering are used to reformulate the investment decision problem in terms of observable variables. The extensive proofs can be found in Liptser and Shiryayev (1978). However the technical conditions of application of the theorems are given in the appendix.
Results are presented here for general forms of the stochastic processes.
Application to the specific case considered in the analysis is done in sections II and IV.
I will consider the general class of stochastic processes:

$$dq = [d(t,q) + D(t,q)\ \mu_t]\ dt + \sum(t,q)\ dZ \qquad \text{(III-0.1)}$$

$$d\mu = [a(t,q) + A(t,q)\ \mu_t]\ dt + B_z(t,q)\ dZ + B_y(t,q)\ dY \qquad \text{(III-0.2)}$$

Dependence on t and q will be omitted in the statement of the theorems unless specifically needed.

III-1 Conditional distribution of the expected returns

The fundamental theorem of non-linear filtering (theorem 8.1 in Liptser and Shiryayev) characterizes the conditional distribution of unobservables for a wide class of processes. However, the equations yield the conditional moments as functions of conditional moments of an higher order. In particular, one fundamental difficulty is that the conditional expectation depends on the second and third conditional moments. We thus need additional relations between the moments to obtain a closed system.

If the random process $[\mu, q]$ is Gaussian, higher order moments are functions of the expectation and variance. We have the classic relation:

$$E(\mu^3 \mid F_t^q) = 3\, E(\mu \mid F_t^q)\, E(\mu^2 \mid F_t^q) - 2\, [E(\mu \mid F_t^q)]^3 \qquad \text{(III-1.1)}$$

which closes the system of equations giving $E(\mu \mid F_t^q)$ and $E(\mu^2 \mid F_t^q)$.

This case was first analyzed by Kalman and Bucy and later extended by Liptser and Shiryayev to the conditionally Gaussian case. Here, $(\mu \mid F_t^q)$ is not a Gaussian system but:

$$f_t^q\,(x) = P(\, \mu_t \leq x \mid F_t^q),$$

the conditional probability distribution, is Gaussian: as is proved by Theorem 1. Consequently III-1.1 closes the system of equations on the conditional moments. The Gaussian property also implies that the distribution is characterized by its first two moments.

Theorem 1

Let conditions A-1 to A-7 of the appendix be verified and, with probability one, let the conditional distribution:

$f_0^q(x) = P\{\mu_0 \leq x | q_0\}$ be (P a.s.) Gaussian, $N(m_0, V_0)$, where $m_0 = E(\mu_0 | F_0^q)$ and the matrix $V_0 = E\{(\mu_0 - m_0)(\mu_0 - m_0)^t \mid F_0^q\}$ is such that Tr $V_0 < \infty$ (P a.s.). Then a random process $[\mu, q]$ verifying (III-0.1) and (III-0.2) is conditionally Gaussian, i.e. for any t_j, $0 < t_0 < t_1 < \ldots < t_n < t$, the conditional distribution:

$f_t^q(x_0, x_1, \ldots, x_n) = P(\mu_0 \leq x_0, \ldots \mu_n \leq x_n \mid F_t^q\}$ is (P a.s.) Gaussian.

III-2 Conditional moments

Theorem 2 gives the conditional expectations vector and variance-covariance matrix of the expected return μ. Since the conditional distribution is Gaussian, it is fully characterized by the first two moments.

Theorem 2

Let conditions A-1 to A-10 be verified, then the vector $m_t = E(\mu_t | F_t^q)$ and the matrix $V(t) = E\{(\mu_t - m_t)(\mu_t - m_t)^t | F_t^q\}$ are unique, continuous, F_t^q-measurable for any t, solutions of the system of equations:

$$dm = [a + A\, m_t]\, dt + [B \Sigma^t + V D^t]\, [\Sigma \Sigma^t]^{-1}\, [dq - (d + D\, m_t)\, dt]$$

and

$$dV = \{A V + V A^t + B_y B_y^t + B_z B_z^t - [B_z \Sigma^t + V D^t]\, [\Sigma \Sigma^t]^{-1}\, [B_z \Sigma^t + V D^t]\}\, dt$$

with initial conditions $m_0 = E(\mu_0 | F_0^q)$ and

$$V_0 = E\,\{(\mu_0 - m_0)(\mu_0 - m_0)^t \mid F_0^q\}.$$

If in this case the matrix V_0 is positive definite, then the matrices V_t, $0 \leq t \leq T$, will have the same property.

The random vector μ_t is thus a diffusion process which variations reflect the fluctuations of the signal q_t. V_t is a deterministic matrix valued function of time and of the investment levels q.

The filtering is called "non-linear" because the estimator is a non-linear function of the observations q_t.

III-3 The innovation process

Theorem 2 gives the conditional moments of μ which characterize the Gaussian distribution. In this section, I reformulate the stochastic processes of interest in terms of the innovation process Z'_t. Z'_t is an observable standard Brownian motion constructed from dq and m_t.

Theorem 3 states that the sigma algebra $F_t^{q_o,Z'}$, generated by Z'_t and q_0, is equivalent to F_t^q , the sigma algebra generated by q_s $\{0\leq s\leq t\}$. That is, all the information observed by agents can be summarized into the observation of the innovation process Z'_t and the initial value of q, q_0. Since Z'_t is a Brownian motion, its path does not contain information on the future except for the value it implies for the state variables μ_t, V_t and q_t. Define Z'_t as follows:

$$dZ' = \Sigma^{-1} [dq - (d + D m) dt] \quad \text{and} \quad Z'_0=0$$

dZ' is the normalized deviation of dq from its conditional expectation at date t. The stochastic processes of interest q_t, m_t and V_t are easily expressed in terms of Z'_t:

$$dq = [d + D \mu] dt + \Sigma dZ'$$

$$dm = [a + A m] dt + [B \Sigma^t + V D^t] (\Sigma^t)^{-1} dZ'$$

$$dV = \{A V + V A^t + B_z B_z^t + B_y B_y^t - [B \Sigma^t + V D^t] [\Sigma \Sigma^t]^{-1} [B \Sigma^t + V D^t]\} dt$$

(system III-3.1)

Theorem 3

Let functionals a(t,q), A(t,q), $B_z(t,q)$, $B_y(t,q)$, d(t,q), D(t,q), $\Sigma(t,q)$ satisfy A-4.

Let also V_0, $a(t,q)$, $A(t,q)$, $B_z(t,q)$, $B_y(t,q)$, $d(t,q)$, $D(t,q)$ and $(\sum\sum^t)^{-1}(t,q)$ be uniformely bounded.

Then the system of equations III-3.1 has a unique strong (i.e. ${}_tF^{q_0,Z'}$ measurable for each t) solution. In this case,

$$F_t^q = {}_tF^{q0,Z'}, \quad 0 \leq t \leq T$$

The knowledge of the past history of instantaneous returns dq is equivalent to the knowledge of the path of Z'_t. Since Z'_t is a Brownian motion, the joint process $[q_t, m_t, V_t]$ is Markov.

<u>Applicability of the filtering results to the structure of uncertainty of sections II and IV</u>

Inspection of equation I-2.1: $dq = I_q\, \mu_t\, dt + I_q\, \sum(t)\, dZ$ shows that the coefficients $D(t,q)=I_q$ and $(\sum\sum^t)^{-1}(t,q)=I_q^{-1}\ (\sum\sum^t)^{-1}(t)\ I_q^{-1}$ are not uniformly bounded. In particular, when one of the investment levels q_i tends to zero, the ith element of I_q^{-1} tends to infinity. Intuition suggests to reformulate the problem in terms of returns, i.e. to redefine the "signal" as the log-transform $Log(q_t)$. The expectation and variance-covariance matrix of μ_t conditional on knowing $\{dLog(q_s),\ 0 \leq s \leq t\}$ are equal to m_t and V_t respectively. In this formulation, the conditions of Theorems 1, 2 and 3 only imply that the coefficients: $a(t)$, $A(t)$, $B_z(t)$, $B_y(t)$ and $(\sum\sum^t)^{-1}(t)$ are bounded. Consequently Theorems 1, 2 and 3 hold for the structure of uncertainty considered in sections II and IV provided that the coefficients satisfy mild regularity assumptions.

III-4 Summary of non-linear filtering results

Let us put the results in the perspective of our problem. Under rather mild assumptions on the stochastic processes, Theorem 1 shows that the conditional distribution of μ knowing (the path of) q is Gaussian. This leads to a simple solution to the fundamental filtering equations. Theorem 2 characterizes the moments of the distribution m_t and V_t as functions of investments q.

The conditional expectation m_t of μ_t, knowing the past realizations of dq, is a stochastic process. The variance-covariance of μ_t conditional on the past realizations of dq is a deterministic function of time and the investment levels q[12]. The conditional distribution being Gaussian, it is characterized by its first two moments. Theorem 3 shows that the set of all available information, i.e. the past realizations of dq, is equivalent to information about the path of the innovation process, Z', and the initial value of q, q_0. Since Z' is a Brownian motion, its past values contain no information on its future values. Hence, all the available information is equivalent to the knowledge of the current values q_t, m_t and V_t. The probability distribution of the variations of q, m and V over the next infinitesimal instant [t,t+dt] are uniquely determined by the present values q_t, m_t and V_t. That is, the system $[q_t,m_t,V_t]$ is Markov.

[12]The conditional variance matrix V_t is equal to the unconditional mean squared error $E[(\mu_t-m_t)^2]$ (as proved by Liptser and Shiryayev).

IV Separability

IV-0 Introduction

The fundamental results of section II and III are twofold. First, the conditional distribution of the vector of instantaneous rates of return on investment conditional on knowing past realizations is Gaussian. Second, the conditional distribution of future returns on investment is determined by the level of investments q_t and the conditional moments at date t, m_t and V_t; consequently the system $[q_t, m_t, V_t]$ is Markov.

The investment decision problem can now be restated in the usual framework of dynamic programming. A priori, all past observations of the instantaneous returns influence investors' perceived opportunity set and thus define the state of the economy. In the original problem, the number of state variables is thus potentially infinite. The analysis of sections II and III shows that all the information is equivalent to that contained in the path of the innovation process Z'_t. The path $\{Z'_s \; 0 \leq s \leq t\}$ determines the conditional moments m_t and V_t, but contains no information on the future realizations of Z'. Investors thus use available information only to infer the conditional moments and the variables that characterize the investment opportunity set are $[q_t, m_t, V_t]$.

This result is closely related to the discrete time "certainty equivalence" principle due to Simon (1956) and Theil (1964). That principle states that if the utility function (or more generally the objective function) is quadratic and if the controlled process is a linear function of the unobservable state variables, the optimization problem can be solved as if the state variables were known for certain to be equal to their conditional expectation. The intuition is that, with a linear marginal utility function and a linear dependence of the process on the state variables, the expected marginal utility depends only on the expectation of the state variable. In the present continuous time case, the dependence of the instantaneous returns on the state variables is linear and the expected marginal utility is a linear function of the first and second moments of the distributions of q_t and μ_t. Intuitively, one is thus led to expect a separation property where only the first two conditional moments of dq_t and $d\mu_t$ matter, and this is precisely what is shown here.

While the assumption of homogeneous agents is not required for the inference process, it becomes necessary for the equilibrium results of section IV-2. Because of the homogeneity of tastes, endowments and beliefs, all agents will choose the same portfolio of assets. Agents are allowed to issue contingent claims and borrow or lend risklessly. However, since these claims are in zero net supply, no agent will hold them in equilibrium.

Consequently, we can restrict our attention to an economy where no contingent claims are traded. The resulting equilibrium will be defined by the same investment proportions and the same interest rate. If contingent claims are introduced, they can be priced by arbitrage.

In this section, I describe the investor's decision problem and prove that the estimation and consumption-portfolio choice steps of this problem can be separated and solved sequentially. I then show how optimal investment choices differ from that obtained by Cox, Ingersoll and Ross (1978).

IV-1 The investment decision problem

Agents seek to maximize their expected lifetime utility of consumption subject to their budget constraints. The investment opportunity set consists of the riskless asset and the s physical technologies as specified in section I by equations (I-2.1) and (I-2.2). Since agents are not able to observe the true expected returns on investments, μ_t, they derive the expected returns conditional on available information. The investment opportunity set at any point in time is characterized by the conditional expected returns on investments m_t and the variance-covariance matrix V_t. Section III shows that the system $[q_t, m_t, V_t]$ is Markov. Since the instantaneous returns do not depend on the level of investment q_t and the variance-covariance matrix V_t is a deterministic function of time[13], the vector m_t fully characterizes the investment opportunity set perceived by investors at any date t (equations IV-1.2,IV-I.4 and IV-1.5).

This leads to the following result:

[13]The filtering results of section III are derived for a wider class of uncertainty structures. However, Theorem 4 does not follow in the general case. If the distribution of returns depends on the level of investments q_t, the quality of the information extracted from past realizations depends on the past levels of investments. Agents would therefore select their optimal investments taking into account the quality of information they generate. This information will be used to optimize future investments. Hence, the estimation problem and the portfolio choice are not separable. To avoid such problems in the present analysis, I will make the assumption that agents are able to draw the same information whether all technologies are effectively used or not.

Separation Theorem (Theorem 4)

Assuming the structure of the uncertainty of section II and that conditions A1-A14 hold, agents solve the investment decision problem in two stages:

.the derivation of the optimal estimator (the conditional expectation) of expected returns.

.the choice of an optimal portfolio of assets using estimated expected rates of return.

Agents are constrained by their individual wealth W_t determined by its initial value W_0 and the differential equation:

$$dW = W_t \; [\omega^t \; I_q^{-1} \; dq + (1 - w^t \; \underline{1}) \; r] - c(t) \; dt \qquad \text{(IV-1.1)}$$

where ω denotes the vector of proportions of wealth invested in the risky technologies and $\underline{1}$ a vector with components equal to 1.

The instantaneous return on investment is given by:

$$dq = I_q \; m_t \; dt + I_q \; \underline{\Sigma}(t) \; dZ' \qquad \text{(IV-1.2)}$$

The instantaneous change in the expected return m_t is:

$$dm = [a(t) + A(t) m_t] \; dt + [B(t) \underline{\Sigma}^t(t) + V_t] \; [\underline{\Sigma}^t(t)]^{-1} \; dZ' \qquad \text{(IV-1.3)}$$

and the change in the conditional variance V_t is:

$$dV = \{A(t) V_t + V_t A^t(t) + B_z(t) B_z^t(t) + B_y(t) B_y^t(t)\} \; dt \qquad \text{(IV-1.4)}$$

$$- \{[B_z(t) \underline{\Sigma}^t(t) + V_t] \; [\underline{\Sigma}(t) \underline{\Sigma}^t(t)]^{-1} \; [B_z(t) \underline{\Sigma}^t(t) + V_t]\} \; dt$$

where the innovation process Z'_t is a standard Brownian motion defined as follows:

$$dZ' = \underline{\Sigma}^{-1}(t) \; [I_q^{-1} dq - m_t \; dt] \quad \text{and} \quad Z'_0 = 0 \qquad \text{(IV-1.5)}$$

Equation (IV-1.2) is similar to equation (I-2.1):

$$dq = I_q \; \mu_t \; dt + I_q \; \underline{\Sigma}(t) \; dZ \qquad \text{(I-2.1)}$$

The expected return μ_t is replaced by its conditional expectation m_t and the stochastic process Z_t is replaced by the observable innovation process Z'_t. If μ_t is perfectly known the variance-covariance matrix of the returns, μ_t, is $\sum\sum^t$, the same as the conditional variance, $Var(I_q^{-1} dq | F_t^q)$ in IV-1.2. Rewriting IV-1.2 as follows:

$$dq = I_q\ m_t\ dt + I_q\ [\mu_t - m_t]\ dt + I_q\ \sum(t)\ dZ$$

we see that the uncertainty on the true expected return, $I_q\ [\mu_t - m_t]\ dt$, is insignificant when compared to the intrinsic stochastic factor $I_q\ \sum(t)\ dZ$. This is why the variance-covariance matrix of the <u>instantaneous</u> returns on investments is unaffected by the estimation error. This rather surprising result holds independently of the information-uncertainty structure assumed for the expected returns. The estimation error, however, modifies the variance-covariance matrix of returns on investment over any <u>discrete</u> time interval. Investors concerned only with returns on investment over the next infinitesimal interval can neglect the estimation error. Indeed, it is shown in section IV-2 that such myopic investors optimize their portfolios as if the estimated value of the expected return were the true value.

IV-2 Equilibrium conditions

Let us proceed with the second stage of the optimization; the analysis is similar to that of Cox, Ingersoll and Ross (1978). Agents choose controls ω and c which maximize their expected lifetime utility of consumption subject to the budget constraint (IV-1.1).

I derive here the necessary conditions for the indirect utility function J and the optimal set of controls $[c,\omega]$. The indirect utility of consumption is a function of time, the expected return m_t and wealth W_t, since V is non-stochastic.

The first order conditions[14] result from differentiating the function :

$E(W,m,t) = \Phi_{[c,\omega]}\{J(W,m,t)\} + U(c,t).$

Where $\Phi_{[c,\omega]}$ is the Dynkin[15] operator associated with the controls $[c,\omega]$.

$$E_c = U_c - J_W \leq 0 \qquad \text{(IV-2.1)}$$

$$c\ E_c = 0 \qquad \text{(IV-2.2)}$$

$$E_\omega = W\ J_W(m_t - \underline{1}r) + W^2\ J_{WW}(\Sigma\ \Sigma^t\omega) + W\ (\Sigma\ \Delta^t J_{Wm}) \leq 0 \qquad \text{(IV-2.3)}$$

$$\omega^t\ E_\omega = 0 \qquad \text{(IV-2.4)}$$

[14] I will assume that agents restrict their attention to the set of non anticipative feedback controls satisfying the growth and Lipschitz conditions. I will also assume that there exists an indirect utility function J and a set of optimal controls $[c,\omega]$. They then verify the first order conditions. (See Fleming and Rishel (1975) for the technical discussion).

[15] Define the function $H^*(x,t)$ by:

$$H^*(x,t) = \lim_{h \to 0} E_t \left[h^{-1} [H(x(t+h),t+h) - H(x,t)] \right]$$

H^* can be interpreted as the expected time rate of change of $H(x,t)$ and is thus a generalization of the ordinary time derivative for deterministic functions. And we have $H^* = \Phi\{h(x,t)\}$. A heuristic method for finding $\Phi\{H(x,t)\}$ is to take the conditional expectation of dH (found by Ito's lemma) and "divide" by dt (see Merton (1971)).

Where J_{Wm} denotes the vector of partial derivatives of J_W with respect to the components of m and $\Delta\Delta^t$ denotes the variance-covariance matrix of the expected instantaneous retuns m_t:

$$\Delta = [\ B(t)\ \Sigma^t(t) + V(t)\ ]\ (\Sigma^t)^{-1} \qquad \text{(IV-2.5)}$$

In addition, J verifies the equation:

$$\phi_{[c,\omega]}\{\ J(W,m,t)\ \} + U(c,t) = 0 \qquad \text{(IV-2.6)}$$

The system of equations IV-2.3 and IV-2.4 yields the optimal equilibrium levels of investments and the interest rate. The complete solution of the system is complex even for simple utility functions. However, if we restrict our attention to the set of technologies actually used in equilibrium (i.e. i's for which $\omega_i > 0$), the optimal investment rule is given by:

$$\omega = -J_W\ (J_{WW}W\ \Sigma\ \Sigma^t)^{-1}\ (m_t - r\underline{1}) - (J_{WW}\ W\ \Sigma\ \Sigma^t)^{-1}\ \Sigma\ \Delta^t\ J_{Wm} \qquad \text{(IV-2.7)}$$

Where ω, m_t, Σ and Δ are to be interpreted as referring to the set of technologies actually used.

Since there is no net supply of the riskless asset, the shares of wealth ω_i invested in the risky technologies sum to one and the riskless rate is given by:

$$r = (\underline{1}\ (\Sigma\ \Sigma^t)^{-1}\ \underline{1}\ J_W\ )^{-1}\ \{\ J_{WW}\ W + J_W\ \underline{1}^t\ (\Sigma\ \Sigma^t)^{-1}\ m_t\ + \underline{1}\ (\Sigma^t)^{-1}\ \Delta^t\ J_{Wm}\ \} \qquad \text{(IV-2.8)}$$

The interest rate is thus a stochastic process, determined by the expected instantaneous returns, wealth and time.

It was shown in section IV-1 that the estimation error does not affect the present physical investment opportunity set. Equation IV-2-8 shows that the interest rate is a function of the conditional variance V_t. Since investors are homogeneous, they will not borrow or lend and the present investment opportunity set is unaffected by the estimation error.

The first term of equation IV-2.7 is the familiar demand for investment by single period optimizers. The second term corresponds to the demand for hedging against changes in the investment opportunity set. Substituting Δ, the instantaneous "standard error" of the conditional expected returns, that second term can be rewritten as:

$$(\sum\sum^{t})^{-1}\sum \Delta^{t}(-J_{WW}\ W)^{-1}J_{Wm} = (\sum\sum^{t})^{-1}\sum B^{t}(-J_{WW}\ W)^{-1}J_{Wm} + (\sum\sum^{t})^{-1}V_{t}(-J_{WW}\ W)^{-1}J_{Wm}$$

The second term of this equation disappears when V_t is equal to zero and the first term is identical to the one obtained by Merton (1973) in his Intertemporal Capital Asset Pricing Model, assuming perfectly known expected returns. If investors are "myopic", as in the case of logarithmic preferences[16], they have no interest in hedging against future shifts of the investment opportunity set (i.e. $J_{Wm}=0$). Consequently, their investment policy does not depend on V_t. They thus behave exactly as if m_t were the true expected return vector. The Separation Theorem in this very specific case states that agents completely ignore the uncertainty on the state variables.

Suppose now that only one technology, say the ith, has an uncertain expected return. The larger the uncertainty on μ_i, the less agents invest in the ith security[17]. Hence, non myopic, risk-averse investors reduce their investment in technologies with more uncertain expected return.

Klein and Bawa (1977) explicitly incorporate estimation risk into the investment decision problem in a one period model. They show that, under some assumptions for the uncertainty structure, the set of efficient portfolios is identical to that given by traditional analysis but the optimal portfolio choice differs. They conclude that risk-averse agents invest less, ceteris paribus, in

[16] Feldman's (1983) derivation of equilibrium investment levels and of the interest rate makes that assumption. The implications for the term structure of interest rates are thus utility specific.

[17] Assuming that an increase in the conditional expected return corresponds to an increase in future consumption, all other things equal. See Merton (1973) for a more detailed discussion of comparative statics of this type.

differs. They conclude that risk-averse agents invest less, ceteris paribus, in assets with more uncertain expected returns. We reach the same conclusion here but the continuous time framework allows us to highlight the mechanism: the demand for hedging causes intertemporal utility maximizers to reduce their investments in lesser known technologies but the uncertainty on production in the next (infinitesimal) period is unaffected by the estimation risk.

Conclusions and future research

A summary of the main results of the analysis is followed by an outline of the research currently underway.

The process of estimating expected returns on investment was explicitly integrated into a general equilibrium model of real investment under uncertainty. Agents were assumed to have rational expectations, i.e. they use all available information in an optimal way to form their expectations. The optimal estimator of expected returns on investment, namely the conditional expectation, was derived using continuous time filtering techniques. The dynamic estimator is a diffusion process governed by a differential equation similar to the one which characterizes the true expected returns. The estimation and the investment decision problems exhibit separability for a wide class of uncertainty structures.

The separation property leads to the derivation of general equilibrium conditions using the customary dynamic programming techniques without restrictive assumptions on agents' preferences. Uncertainty about the expected returns affects the value of the interest rate and the level of physical investments. Risk-averse, intertemporal utility maximizers reduce their investments in technologies with more uncertain expected returns. High levels of uncertainty lead to a reduced diversification of the optimal portfolios. Finally, the effect of estimation risk on equilibrium does not disappear as the observation interval increases without bound and agents gather more information on the technologies.

One line of research motivated by these results is the analysis of the implications for general equilibrium of the non-observability of state variables. It will be further developed in a later companion paper. The results obtained by CIR (1978) will be reexamined in this more general case. In particular they show that, in the case of a single state variable and logarithmic preferences, the interest rate is a linear function of the state variable. In the case of more general utility functions, the interest rate and financial claim prices are functions of the derived moments of the conditional distribution of the state variables. An important question is, therefore, whether or not the conditional moments can be deduced from equilibrium prices. We should not expect markets to become complete by the addition of contingent claims to the investment opportunity set. In fact, no contract with payoffs dependent on unobservable shifts in the investment opportunity set is enforceable (Radner (1981)). There exists, however, a set of contingent claims which allows agents to hedge against changes in their perceived investment opportunity set characterized by the conditional moments of the returns distribution.

The imperfection of the estimation of the state variables adds a factor to the variability of prices, interest rates and equilibrium investments. Consider the case of lognormal prices: the opportunity set is constant. However agents' perception of the investment opportunity set evolves over time as they improve their estimate of mean returns. In this case the learning process is the only determinant of changes in the interest rate and equilibrium investments. Consider now a related situation: the economy goes through a turbulent period followed by a return to stationarity in the investment opportunity set. The characteristics of equilibrium widely vary after the shock responding not to present events but to accruals of information on past events. Using the equations derived here, I will investigate the impact of this learning process on equilibrium. As was shown in this analysis, the conditional distribution of returns has the same instantaneous variance as the distribution of returns conditional on observing the true expected return. However, uncertainty on the state variables adds to the uncertainty on future opportunity sets. In the case of one state variable and logarithmic preferences, the CIR (1978) analysis is easily extended. Logarithmic preferences imply myopic behaviour and precludes demand for hedging against changes in the investment opportunity set. Therefore, further research will focus instead on the more interesting case where investors do not behave myopically and modify their portfolios to hedge against changes in their derived opportunity set. Further research will make the assumption of isoelastic preferences to study the demand for hedging in a tractable case.

APPENDIX

The technical assumptions of the non-linear filtering results are:

Elements of the vector functions:

$$d(t,x) = (d_1(t,x), \ldots d_s(t,x))$$

$$a(t,x) = (a_1(t,x), \ldots a_s(t,x))$$

and matrices:

$$D(t,x) = [D_{i,j}(t,x)]_{s*s}$$

$$\Sigma(t,x) = [\Sigma_{i,j}(t,x)]_{s*s}$$

$$A(t,x) = [A_{i,j}(t,x)]_{s*s}$$

$$B_z(t,x) = [B_{zi,j}(t,x)]_{s*s}$$

$$B_y(t,x) = [B_{yi,j}(t,x)]_{s*n}$$

are assumed to be measurable nonanticipative (i.e B^s_t-measurable where B^s_t is the σ-algebra in the space of C^s_t of the continuous functions $x=\{x_r,\ 0 \le r \le T\}$ generated by the functions x_r, $r \le t$) functionals on

$\{[0,T] * C^s_T,\ B_{[0,T]} * B^s_T\}$, $x=(x_1 \ldots x_s)$ belonging to C^s_t.

and verify (for any x in C^s_T, the space of continuous s-vector valued functions on [0,T], and all the admissible values for i and j):

$$\int_0^T \Big[|a_i(t,x)| + |A_{ij}(t,x)| + (B_{zij}(t,x))^2 + (B_{yij}(t,x))^2 + (\Sigma_{ij}(t,x))^2 \Big] dt < \infty \quad \text{(A-1)}$$

$$\int_0^T \Big[(d_i(t,x))^2 + (D_{ij}(t,x))^2 \Big] < \infty \quad \text{(A-2)}$$

The matrix $\Sigma(t,x) * \Sigma^t(t,x)$ is uniformely non-singular, i.e. the elements of the reciprocal matrix are uniformely bounded. (A-3)

If $g(t,x)$ denotes any element of the matrix $\sum(t,x)$ then for any x and x' in C_T^s,

$$[g(t,x)-g(t,x')]^2 \leq L_1 \int_0^t |x_s-x'_s|^2 \, dK(s) + L_2 \, |x_t-x'_t|^2$$

$$g^2(t,x) \leq L_1 \int_0^t (1+|x_s|^2) \, dK(s) + L_2 \, (1+|x_t|^2) \qquad \text{(A-4)}$$

where $|x_t| = x_1^2(t) + \ldots + x_s^2(t)$ and $K(s)$ is a nondecreasing right continuous function, $0 \leq K(s) \leq 1$.

$$\int_0^T E|D_{ij}(t,q)\mu_j(t)|dt < \infty \qquad \text{(A-5)}$$

$$E|\mu_j(t)| < \infty \qquad 0 \leq t \leq T \qquad \text{(A-6)}$$

$$P\left[\int_0^T [D_{ij}(t,q)m_j(t)]^2 dt < \infty\right] = 1 \qquad \text{(A-7)}$$

$$|A_{i,j}(t,x)| \leq L \qquad |D_{i,j}(t,x)| \leq L \qquad \text{(A-8)}$$

$$\int_0^T E[a_i^4(t,q)+B_{zij}^4(t,q)+B_{yij}^4(t,q)]dt < \infty \qquad \text{(A-9)}$$

$$E \sum \mu_i^4(0) < \infty \qquad \text{(A-10)}$$

where E denotes the expectation and $m_j(t) = E[\mu_j(t) | F_t^q]$.

REFERENCES

F. Black, (1976), "Studies of stock price volatility changes", proceedings of the 1976 meetings of the American Finance Association.

J.C. Cox, (1975), "Notes on option pricing 1: constant elasticity of variance diffusions", Working Paper, Graduate School of Business, Stanford University.

J.C. Cox, J.E. Ingersoll, S.A. Ross, (1978): "A theory of the term structure of interest rates", Research paper No 468, Graduate School of Business, Stanford University.

J. De Temple, (1984): "Information and pricing of assets", Ph.D. dissertation, University of Pennsylvania.

D. Feldman, (1983): "A theory of asset prices and the term structure of interest rates in a partially observable economy", PhD Dissertation, Nortwestern University.

G. Gennotte and T.A. Marsh, (1984): "Determinants of asset prices variations", Working Paper.

C.F. Huang, (1984): "Information structure and equilibrium asset prices" WP 1477-83, A.P. Sloan School of Management, M.I.T.

R.E. Kalman and R.S. Bucy, (1961): "New results in linear filtering and prediction theory", Trans. ASME, Journal of Basic Engineering, ser D, 83, 95-108.

R.W. Klein and V.S. Bawa, (1976): "The effect of estimation risk on optimal portfolio choice", Journal of Financial Economics, Vol 3,3.

R.W. Klein and V.S. Bawa, (1977): "The effect of limited information and estimation risk on optimal portfolio diversification", Journal of Financial Economics, Vol 5,1.

V. Krishnan, (1984): "Non linear filtering and smoothing: an introduction to martingales, stochastic integrals and estimation", John Wiley β sons, New York.

R.S. Liptser and A.N. Shiryayev, (1977): "Statistics of random processes I: general theory", Springer-Verlag, New York.

R.S. Liptser and A.N. Shiryayev, (1978): "Statistics of random processes II: applications", Springer-Verlag, New York.

R.E. Lucas and T.J. Sargent, (1981): "Rational expectations and econometric practice", University of Minnesota Press, Minneapolis.

R.C. Merton, (1971): "Optimum consumption and portfolio rules in a continuous time model", Journal of Economic Theory,Vol 3,4.

R.C. Merton, (1973): "An intertemporal capital asset pricing model", Econometrica, Vol 41,5.

R.C. Merton, (1980): "On estimating the expected return on the market: an exploratory investigation", Journal of Financial Economics Volume 8,4

R. Radner, (1981): "The role of private information in markets and other organisations", Bell Laboratories, Working Paper No.21

H.A. Simon, (1956): "Dynamic programming under uncertainty with a quadratic objective function", Econometrica 24, No 1, Jan 1956, 74-81.

H. Theil, (1964): "Optimal decision rules for government and industry", Amsterdam: North Holland.

J.T. Williams, (1977): "Capital asset prices with heterogeneous beliefs", Journal of Financial Economics Vol 5,2.

Bar Code On Last Page