American Finance Association

Optimal Portfolio Choice Under Incomplete Information
Author(s): Gerard Gennotte
Source: *The Journal of Finance*, Vol. 41, No. 3, Papers and Proceedings of the Forty-Fourth Annual Meeting of the America Finance Association, New York, New York, December 28-30, 1985 (Jul., 1986), pp. 733-746
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/2328506
Accessed: 26/11/2009 08:32

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=black.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Blackwell Publishing and *American Finance Association* are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance.*

http://www.jstor.org

THE JOURNAL OF FINANCE • VOL. XLI, NO. 3 • JULY 1986

Optimal Portfolio Choice Under Incomplete Information

GERARD GENNOTTE*

ABSTRACT

Models of asset pricing generally assume that the variables which characterize the state of the economy are observable. However, the distributional properties of asset prices that are relevant for portfolio decisions are in general not observable, and therefore must be estimated. The estimation of expected returns is a particularly difficult problem and estimation errors are likely to be substantial. In this light, it is reasonable to examine 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 three main objectives: first, to derive optimal estimators for the unobservable expected instantaneous returns using observations of past realized returns; second, to establish that estimation and portfolio choice can be solved in two separate steps; third, to analyze the impact of estimation error on investment choices. The estimators of expected returns are in general not consistent, i.e., the estimation error does not tend to disappear asymptotically. The effects of the estimation error, therefore, cannot be ignored even if realized returns are observed continuously over an infinite time period.

In his seminal papers, Merton [12, 13] characterizes the equilibrium expected return on assets and the optimal portfolio demand in an intertemporal capital asset pricing model. Cox, Ingersoll and Ross [1] (CIR) extend the analysis to a general equilibrium framework where financial claims prices are derived endogeneously. These analyses make the assumption that the state variables which determine the joint distribution of expected returns on investment are known with certainty.

Merton [14] analyzes the properties of the estimation of the expectation and variance of the return on the market portfolio. He argues that the optimal estimator of the expected return does not yield precise inferences when the observation period is finite whereas the estimator of the variance converges to the true value when trading is continuous. It is thus unrealistic to assume that the expected return on investment is a known function of observable state variables. Instead, it can only be estimated with error using the observable instantaneous realized returns on investment.

The estimation of expected returns on investment is a prerequisite for tests of the CAPM and any practical application of the theory. It also raises important

* University of California, Berkeley. 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 and the participants of numerous seminars for their contribution. Support from the College Interuniversitaire pour les Sciences du Management is gratefully acknowledged.

theoretical issues. When selecting their optimal portfolio, investors have to take into account the fact that they have only estimates of "true" expected returns. Since investors form their portfolios using estimators, equilibrium levels of investment, the term structure of interest rates and contingent claims prices will generally 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 on the distribution of asset returns. With the exception of Klein and Bawa [7, 8], Williams [18], Feldman [4], Dothan and Feldman [3] and de Temple [2], 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 parameter estimates. This separation of the estimation and optimization steps is optimal when a property which Simon [16] and Theil [17] termed "certainty equivalence" applies.[1]

In this paper, we establish a separation theorem which extends the certainty equivalence principle to continuous time without restrictions on the utility function. Moreover, the conditional estimates of expected returns are shown to be governed by a system of differential equations similar to the system which characterizes the variations of the "true" expected returns. This result implies that the dynamic programming methodology of CIR can be generalized to analyze the properties of general equilibrium when the state variables are not observable.[2]

Tools of non-linear filtering theory are introduced to derive the optimal estimators of the stochastic expected returns: the expectation of the returns conditional on the observation of past realized returns. They are generalizations of continuously updated Kalman filters. Unlike the case of lognormal returns (Merton [14]), we 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 perfectly assess the path of the expected returns. Imperfect information is thus shown to be not just a transitory problem.

The separation result allows the analysis of the effect of the estimation error on portfolio demand. We show that, except in the case of logarithmic preferences, the levels of investments and the term structure depend upon future variations of the *derived* investment opportunity set, which arise from randomness in both the true opportunity set and in the estimation error. As is well known, in the case of logarithmic preferences and complete information, agents behave myopically and thus derive no utility from hedging against changes in the investment opportunity set. When information is incomplete, logarithmic investors behave

[1] The certainty equivalence principle states that the state variables can be replaced in the optimization problem by their least squares estimators, referred to as the certainty equivalent values. It applies in discrete time when the objective function is quadratic and the process is a linear function of unobservable state variables. It has been widely used in macroeconomic theory. Lucas and Sargent [11] provide a survey and several examples.

[2] Feldman [4] and Dothan and Feldman [3] analyze the term structure of interest rates in a similar framework.

as if their estimators of the expected returns were the true expected returns when selecting their optimal portfolio. In contrast, for more general utility functions, it is shown that risk-averse investors reduce their investments in technologies whose expected returns are more uncertain. This result confirms economic intuition and is consistent with the conclusions of Klein and Bawa [7, 8] in discrete time.

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. We 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. The general inference problem was solved by Liptser and Shiryayev [9, 10] and the proofs are therefore not presented. The power of the technique and its applicability to a broad range of economic problems, however, warrant extensive development. Section III contains the separability result and establishes the general equilibrium conditions. Finally, the conclusion summarizes the implications for general equilibrium.

I. The Economy

The economy described here is similar to the one considered by CIR 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 the uncertainty is formally defined in mathematical terms. However, the assumptions on the investment opportunity set and on agents' preferences are standard in continuous time general equilibrium analysis.

A. 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_τ. At each point in time, the state of the economy is characterized by the path of X previous to that instant: $\{X_\tau, 0 \le \tau \le t\}$. 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 α 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, $X = \{X_\tau, 0 \le \tau \le T\}$, its components, without loss of generality, assumed to be independent.

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

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

B. 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:

$$dq = I_q \mu_t \, dt + I_q \Sigma(t) \, dZ \tag{1}$$

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

$$d\mu = [a(t) + A(t)\mu_t] \, dt + B_z(t) \, dZ + B_y(t) \, dY \tag{2}$$

where I_q is an $s \times s$ diagonal matrix whose i'th element is the i'th component of $q(t)s$ $\Sigma(t)$ is a bounded, non stochastic, $s \times s$ 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 $s \times s$ matrices and an $s \times n$ 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.[3]

Agents observe instantaneous returns on investments, dq, and are assumed to know the deterministic functions of time, $\Sigma(t)$, $a(t)$, $A(t)$, $B_z(t)$ 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), and 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.

C. 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.

[3] 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.

[4] The physical production technologies have constant returns to scale. There is thus no incentive to create firms to pool resources. Constant return to scale on investment is a necessary condition for the separation theorem.

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):

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

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 as date t.[5]

II. The Inference Process

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.

A. 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. Liptser and Shiryayev[6] [10] establish that the conditional distribution of μ is Gaussian and derive the differential equations for the changes in its conditional moments as a function of observed returns. These results are essentially a continuous-time generalization of the theory of filters due to Kalman and Bucy [6]. We 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 Gaussian vector conditional on a correlated Gaussian vector is Gaussian. The derivation of the conditional distribution is therefore intuitively similar to the classical statistics problem and standard econometric analysis can be used to interpret the results in the one dimension case treated in Part B. Although the random variables, q and μ, are not Gaussian, the distribution of μ conditional on q is Gaussian.[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)

[5] See CIR for the technical conditions imposed on $c(t)$.

[6] The interested reader should refer to Liptser and Shiryayev [10] for the rigourous proofs. The technical conditions on the uncertainty structure are not repeated here. They require only mild restrictions on the coefficients once the signal q_s has been redefined in terms of $\text{Log}(q_s)$.

[7] This result stems from the structural form of $d\mu$ and dq. It holds for a wide class of processes: the coefficients Σ, a, A, B_Z and B_Y of equations (1) and (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.

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:

$$dm = [a(t) + A(t)m_t]\,dt + [B_z(t)\Sigma^t(t) + V(t)] \cdot [\Sigma(t)\Sigma^t(t)]^{-1}[I_q^{-1}\,dq - m_t dt] \tag{3}$$

$$dV = \{A(t)V_t + V_tA^t(t) + B_z(t)B_z^t(t) + B_y(t)B_y^t(t) - [B_z(t)\Sigma^t(t) + V_t][\Sigma(t)\Sigma^t(t)]^{-1}[\Sigma(t)B_z^t(t) + V_t]\}\,dt \tag{4}$$

m_t is a diffusion process and V_t is a deterministic function of time. Equation (4) is a multidimensional differential equation of the Ricatti type.

The first term of equation (3) 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 (4) 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 of 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 (Part B), 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 and 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, this analysis can be extended to take the estimation of both parameters into account. An alternative approach is that taken in Gennotte and Marsh [5], where equilibrium theory is used to specify relationships between shifts in expected returns on assets and variance changes.

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

$$dZ' = \Sigma^{-1}(t)[I_q^{-1}\,dq - m_t\,dt] \tag{5}$$

and $Z_0' = 0$

Substituting for dq its expression from (1), dZ' is given by:

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

[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.

and $Z'_0 = 0$

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; the increments dZ' being uncorrelated. The Brownian vector, Z, is not observable, the innovation process, Z', however, is derived from observable processes and is thus observable. Substituting dZ' for dZ, the return over the next infinitesimal period becomes:

$$dq = I_q m_t \, dt + I_q \Sigma(t) \, dZ' \tag{6}$$

Hence, the paths of m and Z' uniquely determine the path of q. The path of the conditional expectation, m, 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 $\{q_s, 0 \leq s \leq t\}$.

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. V_t is a deterministic function of time (4), so the distribution of the instantaneous change, dm_t, is 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. 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 $[t, t + dt]$.

B. An Example in Dimension One

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

$$dq = q \, \mu \, dt + q \, \sigma \, dz \tag{7}$$

$$d\mu = a \, \mu \, dt + b \, dz \tag{8}$$

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. Changes in the conditional expectation,

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

[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.

m_t, are given by:

$$dm = a\, m_t\, dt + \left(b + \frac{v}{\sigma}\right)\frac{1}{\sigma}\left[\frac{dq}{q} - m_t\, dt\right] \tag{9}$$

The innovation process z' is defined by:

$$dz' = \frac{1}{\sigma}\left[\frac{dq}{q} - m_t dt\right] \quad \text{and} \quad z_0' = 0$$

Substituting in (7) and (9), dq and dm are rewritten as:

$$dq = q\, m_t\, dt + q\, \sigma\, dz' \tag{10}$$

$$dm = a\, m_t\, dt + \left(b + \frac{v}{\sigma}\right) dz' \tag{11}$$

Equation (9) can be interpreted as:

$$dm = d[E_t(\mu_t)] + \frac{\mathrm{Cov}(dq, \mu_{t+dt})}{\mathrm{Var}(dq)}\left[\frac{dq}{q} - m_t\, dt\right] \tag{12}$$

Equation (12) 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 to the variance of the signal, dq. This ratio is similar to the beta of regression equations.[11] We see that if the innovation and the correlation between μ and dq are positive, agents will adjust their expectation upward.

The conditional variance, v, is determined by v_0 and dv:

$$dv = \left[(2av + b^2) - \left(b + \frac{v}{\sigma}\right)^2\right] dt \tag{13}$$

v_t is thus a predictable function of time.

Equation (13) can be interpreted as:

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

This equation is also similar to the usual inference equation for the 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 (13) is of the well-known Riccati type of differential equations, and

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

its solution is:

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

and

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

where w is defined as: $w = a\sigma^2 - b\sigma$. If w is negative, the limit of v as t tends to infinity is zero. If w is positive, the limit is $2w$; if the initial variance, v_0, is smaller than $2w$, the variance (i.e., the estimation error) increases uniformly. Hence, the estimator m_t is not consistent in general. If μ_0 is perfectly known (i.e., $v_0 = 0$), the uncertainty of μ_t disappears completely. Finally, if μ is an unknown constant (i.e., $v_0 \neq 0$, and $a = b = 0$), m_t still fluctuates but converges to the true μ_t. There is no 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 [14]; 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. Separability

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. However, we have shown that the path of the innovation process, Z_t' contains all the relevant information. The path $(Z_s', 0 \leq s \leq t\}$ determines the conditional moments, m_t and V_t, but contains no additional information on future realizations of q. Rational investors will thus base their investment decisions solely on $[q_t, m_t, V_t]$.

This result is closely related to the discrete time "certainty equivalence" principle due to Simon [16] and Theil [17]. 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 Part B. 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 to 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. If contingent claims are introduced, they can be priced by arbitrage.[12]

A. *The Investment Decision Problem*

Agents seek to maximize their expected lifetime utility of consumption subject to their budget constraint. The investment opportunity set consists of the riskless asset and the s physical technologies as specified in section I by equations (1) and (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 investment, m_t, and the variance-covariance matrix, V_t. 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 (15), (17) and (18). This leads to the following result:

> *Separation Theorem*
> *Agents solve the investment decision problem in two stages: derivation of the vector of (conditional) expected returns, and choice of an optimal portfolio of assets using estimated expected returns.*

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

$$dW = W_t[\alpha^t I_q^{-1}\, dq + (1 - \alpha^t \mathbf{1})r] - c(t)\, dt \tag{14}$$

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

The instantaneous return on investment is given by:

$$dq = I_q m_t\, dt + I_q \Sigma(t)\, dZ' \tag{15}$$

The instantaneous changes in the expected return, m_t, and in the conditional variance, V_t, are given by (3) and (4), and the innovation process is defined by (5).

[12] See section IV.

[13] The filtering results hold for a wider class of uncertainty structures. However, the theorem does not follow in the general case. If the distribution of returns depends on the level of investments q_t, the quality of 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. Hence, the estimation problem and the portfolio choice are not separable. To avoid such problems in the present analysis, we will make the assumption that agents are able to draw the same information whether all technologies are effectively used or not.

Equation (15) is similar to equation (1):

$$dq = I_q \mu_t \, dt + I_q \Sigma(t) \, dZ \tag{1}$$

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, is $\Sigma\Sigma^t$, the same as the conditional variance–covariance matrix in (15). Rewriting (15) as follows:

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

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, $I_q \Sigma(t) \, dZ$.

This is why the variance-covariance matrix of the *instantaneous* returns on investment 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 *discrete* time interval. Investors concerned only with returns on investment over the next infinitesimal interval can neglect the estimation error. Indeed, it is shown in Part B that myopic investors optimize their portfolios as if the estimated value of the expected return were the true value.

B. *Equilibrium Conditions*

Let us proceed with the second stage of the optimization; the analysis is similar to that of CIR. Agents choose controls α and c which maximize their expected lifetime utility of consumption subject to the budget constraint (14). We derive here the necessary conditions for the indirect utility function, J, and the optimal set of controls, $[c, \alpha]$. The indirect utility of consumption is a function of time, the expected return, m_t, and wealth, W_t, since V is nonstochastic.

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

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

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

$$E_c = U_c - J_c \le 0 \tag{16}$$

$$c \, E_c = 0 \tag{17}$$

$$E_\alpha = W \, J_W(m_t - \mathbf{1} \, r) + W^2 J_{WW}(\Sigma \, \Sigma^t \alpha) + W(\Sigma \, \theta^t J_{Wm}) \le 0 \tag{18}$$

$$\alpha^t \, E_\alpha = 0 \tag{19}$$

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

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

[15] See Merton [12], for example.

instantaneous returns, m_t:

$$\theta = [B_z(t)\ \Sigma^t(t) + V(t)](\Sigma^t)^{-1} \tag{20}$$

In addition, J verifies the equation:

$$\Phi_{[c,\alpha]}\{J(W, m, t)\} + U(c, t) = 0 \tag{21}$$

The system of equations, (18) and (19) 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 $\alpha_i > 0$), the optimal investment rule is given by:

$$\alpha = -J_W(J_{WW}W\ \Sigma\ \Sigma^t)^{-1}(m_t - r\mathbf{1}) - (J_{WW}W\ \Sigma\ \Sigma^t)^{-1}\Sigma\ \theta^t J_{Wm} \tag{22}$$

where α, m_t, Σ and θ 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, α_i, invested in the risky technologies sum to one and the riskless rate is given by:

$$r = (\mathbf{1}^t(\Sigma\Sigma^t)^{-1}\mathbf{1}\ J)^{-1}\{J_{WW}W + J_W\mathbf{1}^t(\Sigma\Sigma^t)^{-1}m_t + \mathbf{1}(\Sigma^t)^{-1}\theta^t J_{Wm}\} \tag{23}$$

The interest rate is thus a stochastic process, determined by the expected instantaneous returns, wealth and time.[16] It was shown in Part A that the estimation error does not affect the present physical investment opportunity set. 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 (22) 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 θ, the instantaneous "standard error" of the conditional expected returns, that second term can be rewritten as:

$$(\Sigma\Sigma^t)^{-1}\Sigma\theta^t(-J_{WW}W)^{-1}J_{Wm} = (\Sigma\Sigma^t)^{-1}\Sigma B_z^t(-J_{WW}W)^{-1}J_{Wm} + (\Sigma\Sigma^t)^{-1}V(t)(-J_{WW}W)^{-1}J_{Wm}$$

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 [13] in his Intertemporal Capital Asset Pricing Model, assuming perfectly known expected returns. If investors are "myopic", as in the case of logarithmic preferences, they have no interest in hedging against future shifts of the investment opportunity set (i.e., $J_{Wm} = 0$). Consequently, their investment policy, given their estimate of the expected returns, does not depend on $V(t)$.[17] They thus behave exactly as if m_t were the true expected return vector. The separation theorem in this specific case states that agents ignore the uncertainty on the state variables when using their estimate of expected returns to select optimal portfolios.

[16] In order to obtain closed form solutions for equilibrium prices and the term structure of interest rates, Feldman [4] and Dothan and Feldman [3] restrict their model to the case of logarithmic preferences. They show that the interest rate is a linear function of the best estimator of the state variable.

[17] Only future updates of the conditional expected return $m(t)$ depend on the estimation error.

Suppose now that only one technology, say the i'th, has an uncertain expected return. The larger the uncertainty on μ_i, the less agents invest in the i'th security.[18] Hence, non-myopic, risk-averse investors reduce their investment in technologies with more uncertain expected return. Klein and Bawa [8] 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 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 estimation risk.

IV. Conclusions and Future Research

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 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.

In a partially observable economy, no contract with payoffs dependent on unobservable shifts in the investment opportunity set is enforceable (Radner [15]). Hence, we can restrict our attention to the pricing of contingent claims whose payoffs depend only on observable states of the economy. In the framework considered here, markets are not complete; furthermore, agents are not able to identify the "true" state of the world. By contrast, it can be shown that markets are complete with respect to the observable states of the world (as characterized by the set of derived state variables), i.e., agents are able to manufacture all the claims contingent on the observable states of the world by using dynamic portfolio

[18] Assuming that an increase in the conditional expected return corresponds to an increase in future consumption, all other things equal.

strategies of physical investments. It follows that, in contrast to the original CIR model where state variables are observable, financial claim prices are determined by arbitrage.

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, we will investigate the impact of this learning process on equilibrium prices and investment levels.

REFERENCES

1. J. C. Cox, J. E. Ingersoll and S. A. Ross. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica* 53 (March 1985), 363–384.
2. J. De Temple. "Information and Pricing of Assets." Ph.D. dissertation, University of Pennsylvania.
3. U. Dothan and D. Feldman. "Equilibrium Interest Rates and Multiperiod Bonds in a Partially Observable Economy." WP No. 84-133, Owen Graduate School of Management, Vanderbilt University.
4. D. Feldman. "A Theory of Asset Prices and the Term Structure of Interest Rates in a Partially Observable Economy", Ph.D. Dissertation, Northwestern University.
5. G. Gennotte and T. A. Marsh. "Variations in ex ante Risk Premiums on Capital Assets.". Working paper, School of Business Administration, University of California, Berkeley.
6. R. E. Kalman and R. S. Bucy. "New Results in Linear Filtering and Prediction Theory." Trans. ASME, Journal of Basic Engineering Series D, 83, 95–108.
7. R. W. Klein and V. S. Bawa. "The Effect of Estimation Risk on Optimal Portfolio Choice.' *Journal of Financial Economics* 3 (June 1976), 215–232.
8. R. W. Klein and V. S. Bawa. "The Effect of Limited Information and Estimation Risk on Optimal Portfolio Diversification." *Journal of Financial Economics* 5 (August, 1977), 89–111.
9. R. S. Liptser and A. N. Shiryayev. *Statistics of Random Processes* I: *General Theory.* New York: Springer-Verlag, 1977.
10. R. S. Liptser and A. N. Shiryayev. *Statistics of Random Processes* II: *Applications.* New York: Springer-Verlag, 1978.
11. R. E. Lucas and T. J. Sargent. *Rational Expectations and Econometric Practice.* Minneapolis: University of Minnesota Press, 1981.
12. R. C. Merton. "Optimum Consumption and Portfolio Rules in a Continuous Time Model." *Journal of Economic Theory* 3 (December 1971), 373–413.
13. R. C. Merton. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (September 1973), 867–887.
14. R. C. Merton. "On Estimating the Expected Return on the Market: an Exploratory Investigation." *Journal of Financial Economics* 8 (December 1980), 323–362.
15. R. Radner. "The Role of Private Information in Markets and Other Organizations," Bell Laboratories, Working Paper No. 21.
16. H. A. Simon. "Dynamic Programming Under Uncertainty with a Quadratic Objective Function." *Econometrica* 24 (January 1956), 74–81.
17. H. Theil. *Optimal Decision Rules for Government and Industry.* Amsterdam: North Holland, 1964.
18. J. T. Williams. "Capital Asset Prices with Heterogeneous Beliefs," *Journal of Financial Economics* 5 (November 1977), 219–240.