

USING CONDITIONAL MOMENTS OF ASSET PAYOFFS TO INFER THE VOLATILITY
OF INTERTEMPORAL MARGINAL RATES OF SUBSTITUTION

A. Ronald Gallant
North Carolina State University

Lars Peter Hansen
The University of Chicago
and NORC

George Tauchen
Duke University

March 1989
Revised June 1989
Revised September 1989

USING CONDITIONAL MOMENTS OF ASSET PAYOFFS TO INFER THE VOLATILITY
OF INTERTEMPORAL MARGINAL RATES OF SUBSTITUTION*

A. Ronald Gallant
North Carolina State
University

Lars Peter Hansen
The University of Chicago
and NORC

George Tauchen
Duke University

ABSTRACT

Previously Hansen and Jagannathan (1989) derived and computed mean-standard deviation frontiers for intertemporal marginal rates of substitution (IMRS) implied by asset market data. These frontiers give the lower bounds on the standard deviations as a function of the mean. In this paper we develop a strategy for utilizing conditioning information efficiently, and hence improve on the standard deviation bounds computed by Hansen and Jagannathan. We implement this strategy empirically by using the seminonparametric (SNP) methodology suggested by Gallant and Tauchen (1989) to estimate the conditional distribution of a vector of monthly asset payoffs. We use the fitted conditional distributions to calculate both conditional and unconditional standard deviation bounds for the IMRS. The unconditional bounds are as sharp as possible subject to robustness considerations. We also use the fitted distributions to compute the moments of various candidate marginal rates of substitution suggested by economic theory, and in particular the time nonseparable preferences of Dunn and Singleton (1986) and Eichenbaum and Hansen (1989). For these preferences, our findings suggest that habit persistence will put the moments of the IMRS inside the frontier at reasonable values of the curvature parameter. At the same time we uncover evidence that the implied IMRS fails to satisfy all of restrictions inherent in the Euler equation. The findings help explain why Euler-equation estimation methods typically find evidence in favor of local durability instead of habit persistence for monthly data.

* Each author acknowledges financial support from the National Science Foundation. We are grateful to Angelo Melino and Whitney Newey for many helpful comments.

Send proofs to: Professor George Tauchen, Department of Economics,
Room 305 Social Science Building, Duke University,
Durham, NC 27706

1. Introduction and basic setup

1.1 Volatility bounds

The goal of this paper is to use conditional moments of asset payoffs to deduce volatility bounds on the intertemporal marginal rates of substitution of consumers. Previously, Hansen and Jagannathan (1989) derived and computed mean-standard deviation frontiers for the intertemporal marginal rates of substitution. These frontiers give the lower bounds on the standard deviations as a function of the mean and are of interest for a variety of reasons. First, they can be used to assess which asset market data sets have the most startling implications for a broad class of asset pricing theories. Second, they can be used as diagnostics for helping to discriminate among various candidate intertemporal asset pricing models. Finally, they can be used to assist in isolating the source of failure of particular asset pricing models that are diagnosed as being implausible using formal statistical methods.

Following Hansen and Richard (1987) and Hansen and Singleton (1982), one tractable way to incorporate conditioning information is to form additional portfolios of asset payoffs using information available to economic agents when securities are traded. Hansen and Jagannathan (1989) and Hansen and Singleton (1982) used this conditioning information in an *ad hoc* manner. There is typically a great degree of flexibility in the manner in which additional portfolio payoffs can be formed using conditioning information. For instance, it is typically possible to take a finite set of primitive securities and form an infinite-dimensional set of portfolio payoffs using conditioning information. In this paper we show how to use the conditioning information as efficiently as possible, and hence improve on the standard deviation bounds computed by Hansen and Jagannathan (1989).

Even though the volatility bounds deduced by Hansen and Jagannathan (1989) apply to *unconditional* moments, we show that efficient use of conditioning information requires knowledge of the first two *conditional* moments of the asset payoffs. There are a variety of nonparametric and seminonparametric methods that can be used to estimate these conditional moments. In this paper we use the seminonparametric (SNP) methodology suggested by Gallant and Tauchen (1989) to first estimate the conditional distributions of a vector of asset payoffs. The conditional moments of the asset payoffs are then inferred from the conditional distributions.

We also show how to compute the standard deviation bounds on intertemporal marginal rates of substitution *conditional* on information available to economic agents. A potential advantage to looking at conditional frontiers is that the conditional distributions of the asset payoffs may have *thinner tails* for most realizations of the conditioning information. Hence, by conditioning, it is often possible to control better for the impact of outlier events on the moments of the asset payoffs. Among other things, we view conditioning as a substitute to the commonly-used practice of splitting the sample into many smaller subsamples.

1.2 Basic setup

We follow Hansen and Richard (1987) in modeling asset prices. For the time being, consider an economy in which asset trades occur at some initial date and the payoffs on these assets occur at some future date. Information is available to the consumers at the trading date and reflected in the equilibrium asset prices. We focus on only two time periods, the trading date and the payoff date, for notation convenience. We have in mind that the admissible

asset payoffs and the corresponding prices are replicated over time in a manner that is stationary, at least asymptotically, and ergodic. This replication is important for our empirical analysis because it permits us to estimate consistently conditional and unconditional moments of the asset payoffs. Hence, while time subscripts are initially suppressed, they become important when we describe methods for estimation and inference in later sections of the paper.

We specify in turn the information, the space of admissible portfolio payoffs and the representation of equilibrium prices.

Information: Let I be a conditioning information set available to economic agents and an econometrician at a particular point in time. Agents are presumed to use information in I to form portfolios of asset payoffs.

Portfolio payoffs: We let P_I be a space of payoffs at some future date on portfolios of assets. For convenience, we impose the restriction that $E(p^2|I)$ is finite with probability one for all p in P_I . In other words, we focus our attention on payoffs that have finite conditional second moments. Given this restriction we can define a conditional inner product

$$\langle p_1 | p_2 \rangle_I = E(p_1 p_2 | I) , \quad (1)$$

and a conditional norm

$$\|p\|_I = [\langle p | p \rangle_I]^{1/2} . \quad (2)$$

We impose two additional restrictions on P_I . The first restriction is

Conditional linearity: For any w_1 and w_2 in I and any p_1 and p_2 in P_I , $w_1 p_1 + w_2 p_2$ is in P_I .

To state the second restriction, we need to define a notion of a convergent sequence and a Cauchy sequence in P_I . The sequence $\{p_j\}$ is conditionally Cauchy if for any $\epsilon > 0$,

$$\lim_{j, k \rightarrow \infty} \Pr\{\|p_j - p_k\|_I > \epsilon\} = 0. \quad (3)$$

The sequence $\{p_j\}$ converges conditionally to p_0 if $\{\|p_j - p_0\|_I\}$ converges in probability to zero. The second restriction on P_I is

Conditional completeness: Every sequence in P_I that is conditionally Cauchy converges conditionally to a point in P_I .

The payoff space P_I is constructed so as to be the conditional counterpart of a Hilbert space. As indicated in Hansen and Richard (1987), the space of all random variables with conditional second moments that are finite almost surely is conditionally linear and complete. For the empirical analysis in this paper, a smaller space is also of interest. Let x denote a vector of variables observed by both the econometrician and economic agents in the current time period, and let y_1 denote an M_1 -dimensional vector of asset payoffs at some future date. Both of these vectors might be constructed from an underlying M -dimensional stochastic process $\{y_t\}$ observed by the econometrician. The vector y_1 is taken to be a subvector of the underlying process at some future date and the vector x contains the current and a finite number of lags of this same process.

$$P_I = \{p: p = v \cdot y_1 \text{ for some } M_1\text{-dimensional vector } v \text{ of random variables in } I\}. \quad (4)$$

It is straightforward to verify that P_I as given by (4) is conditionally linear and complete.

1.3 Asset pricing

We assume that all payoffs in P_I have asset prices that are unambiguously finite and in the information set I . Hence we model asset prices as a function π mapping P_I into I . Hansen and Richard (1987) imposed the following restrictions on π :

Conditional linearity: For any w_1 and w_2 in I and any p_1 and p_2 in P_I , $\pi(w_1 p_1 + w_2 p_2) = w_1 \pi(p_1) + w_2 \pi(p_2)$.

Conditional continuity: For any sequence $\{p_j\}$ for which $\{\|p_j\|_I\}$ converges in probability to zero, $\{\pi(p_j)\}$ converges in probability to zero.

Nondegenerate pricing: There exists a payoff p_0 in P_I for which $\Pr\{\pi(p_0)=0\} = 0$.

Hansen and Richard showed that when π satisfies these restrictions, there exists a payoff p^* in P_I such that:

$$\pi(p) = \langle p^* | p \rangle_I \text{ for all } p \text{ in } P_I. \quad (5)$$

Furthermore, since the pricing function is not degenerate $\|p^*\|_I$ is strictly positive with probability one. Result (5) is just the conditional counterpart to the familiar Riesz Representation Theorem.

Although p^* is the unique random variable in P_I that satisfies (5), typically there are many random variables not in P_I that also can be used to represent π . Some of these other random variables are easier to interpret and have more direct links to explicit dynamic models. The payoff space P_I is restricted to be tractable for econometric applications and is often smaller than the span of the security market payoffs that economic agents can trade. In the special case in which agents can trade a complete set of contingent claims, the Arrow-Debreu prices can be presented in terms of a strictly positive random variable m that can be used for p^* in (5) except that m may not be in P_I . Since consumers equate marginal rates of substitution to prices, m is also a measure of the common (across consumers) intertemporal marginal rate of substitution (IMRS). More generally, if agents do not face a complete set of securities markets, the individual IMRS's will not necessarily be equated. The IMRS of any consumer can be used, however, to represent π . Hansen and Jagannathan (1989) give several examples of the economic models and the corresponding characterizations of IMRS's.

We focus on properties of random variables m that satisfy

Restriction 1: $\langle m|m \rangle_I < \infty$ and $\langle m|p \rangle_I = \pi(p)$ for all p in P_I .

These random variables are candidates for IMRS's of consumers.¹ Whenever Restriction 1 is satisfied, the pricing function is conditionally linear and complete. In addition, there exists a random variable p^* in P_I that satisfies (5) even if the pricing function is degenerate on P_I . In this paper we analyze implications of Restriction 1 for m given the data on asset prices and payoffs.

2. Conditional analysis

In this section we derive the conditional counterparts to the results reported in Hansen and Jagannathan (1989). First we compute regressions of m onto P_I and a constant term conditioned on I . Then we characterize the *mean-standard deviation frontier* for m conditioned on I .

Since p^* satisfies (5), it follows that $m - p^*$ is orthogonal to P_I conditioned on I :

$$\langle m - p^* | p \rangle_I = 0 \quad \text{for all } p \text{ in } P_I. \quad (6)$$

Recall that p^* is in P_I . As a consequence, p^* is the least squares projection of m onto P_I conditioned on I .

In many cases P_I does not contain a payoff that is riskless conditioned on I . In these circumstances, it is of interest to deduce the least squares projection of m onto a large space that is augmented by a riskless payoff. Let

$$P_I^+ = \{p + w : p \text{ is in } P_I \text{ and } w \text{ is in } I\}. \quad (7)$$

It is straightforward to show that P_I^+ is conditionally linear and complete whenever P_I is conditionally linear and complete. In addition, P_I^+ always contains a payoff o that is identically equal to one. We compute the conditional projection of m onto P_I^+ by adding to p^* the conditional projection of m onto the collection of all random variables in P_I^+ that are orthogonal to P_I conditioned on I . Let p^0 denote the conditional projection of o onto P_I . Then

$$\langle e|p \rangle_I = 0 \quad \text{for all } p \text{ in } P_I, \quad (8)$$

where e is the conditional regression error $o-p^0$. Any random variable in P^+ that is conditionally orthogonal to P_I is the product of a random variable in I and e . Hence, the conditional projection of m onto P_I^+ is given by

$$p^+ = p^* + we \quad (9)$$

where w is the conditional regression coefficient

$$w = \langle e|m \rangle_I / \langle e|e \rangle_I \quad (10)$$

obtained by regressing m onto e conditioned on I . In general this regression coefficient will depend on conditioning information in I . Notice that the two terms on the right side of (9) are conditionally orthogonal [see (8)]. In addition, p^+ satisfies the following two conditional moment restrictions:

$$\langle m-p^+ | p \rangle_I = 0 \quad \text{for all } p \text{ in } P_I \quad (11)$$

and

$$\langle m-p^+ | o \rangle_I = 0. \quad (12)$$

Condition (12) is equivalent to the restriction that m and p^+ have the same mean conditioned on I . It follows from (11) and (12) that the conditional covariance of $m-p^+$ and p is zero for any p in P_I . As a consequence, the standard deviation of m conditioned on I , denoted $\text{std}(m|I)$, is greater than or equal to $\text{std}(p^+|I)$.

It turns out that the conditional regression coefficient w depends on the conditional mean of m . Taking conditional expectations of both sides of (9) yields

$$E(m|I) = E(p^*|I) + wE(e|I) . \quad (13)$$

Also, $\Pr\{e = 0\}$ is zero whenever $\Pr\{\langle e|e \rangle_I = 0\}$ is zero. In this case given knowledge of $E(m|I)$, equation (13) can be solved for w . Conversely, for any choice of w in I there exists a corresponding random variable $E(m|I)$ satisfying (13).

Since $E(m|I)$ is not specified *a priori*, we consider the indexed family of random variables $\{m_w : w \text{ in } I\}$ where

$$m_w = p^* + we. \quad (14)$$

Each member of this family satisfies:

$$\langle m_w | m_w \rangle_I < \infty \text{ and } \langle m_w | p \rangle_I = \pi(p) \text{ for all } p \text{ in } P_I . \quad (15)$$

Hence each member is a valid candidate for m satisfying Restriction 1.

Without knowledge of $E(m|I)$, all we can say is that the conditional projection of m onto P_I^+ is in the set $\{m_w : w \text{ in } I\}$. Each member of this set is on the *conditional mean-standard deviation frontier* for m . That is, even without knowledge of $E(m|I)$, the ordered pair $[E(m|I), \text{std}(m|I)]$ must be in the set

$$S_I = \{(w_1, w_2) \text{ in } I \times I : w_1 = E(m_w|I) \text{ and } w_2 \geq \text{std}(m_w|I) \text{ for some } w \text{ in } I\} \quad (16)$$

It is of interest to determine when the region S_I touches the horizontal axis with probability one. This occurs when $\text{std}(m_w|I)$ is zero with probability one for some w or equivalently m_w is equal to w . In this case

$$\pi(p) = wE(p|I), \quad (17)$$

and therefore we cannot rule out risk-neutral pricing.

We now illustrate the construction of the indexed family $\{m_w: w \text{ in } I\}$ when P_I is given by (4). To construct this set $\{m_w: w \text{ in } I\}$ we must compute p^* and p^0 . Since p^* and p^0 are in P_I , they can be represented as $v^* \cdot y_1$ and $v^0 \cdot y_1$ for some vectors v^* and v^0 of random variables in I . It follows from (5) and (8) that

$$E(y_1 y_1' | I) v^* = \pi(y_1) \text{ and } E(y_1 y_1' | I) v^0 = E(y_1 | I). \quad (18)$$

When $E(y_1 y_1' | I)$ is nonsingular with probability one,

$$v^* = [E(y_1 y_1' | I)]^{-1} \pi(y_1) \text{ and } v^0 = [E(y_1 y_1' | I)]^{-1} E(y_1 | I). \quad (19)$$

Hence the conditional mean-standard deviation frontier for m is constructed using the first two conditional moments of y_1 and the vector of prices of y_1 .

In particular,

$$m_w = y_1' [E(y_1 y_1' | I)]^{-1} [\pi(y_1) - wE(y_1 | I)] + w. \quad (20)$$

3. Unconditional analysis

In this section we deduce implications for the first two unconditional moments of m . The notation $\|\cdot\|$ and $\langle \cdot | \cdot \rangle$ without an I subscript is used to denote the unconditional counterparts to (1) and (2). Let P be the space of all random variables in P_I with finite unconditional second moments. This space can be infinite dimensional even when P is constructed from a finite-dimensional random vector y_1 as in (4). The fact that consumers use conditioning information in I to form portfolios can increase dramatically the dimension of P . Hansen and Richard (1987 Lemma A.4) showed that P is (unconditionally) linear and complete. Hence P is a Hilbert space. For the unconditional analysis to be of interest we must strengthen Restriction 1.

Restriction 2: $\langle m | m \rangle < \infty$ and $\pi(p) = \langle m | p \rangle$ for any p in P .

Our strategy in this section is to replicate the analysis in Section 2. using P in place of P_I and using unconditional projections in place of conditional projections. Since m and o have finite second moments, the unconditional least square projections of m and o onto P are the same as the conditional projections of these random variables onto P_I [see the proof of Theorem A.2 in Hansen and Richard (1987)]. Hence p^* and p^0 are the unconditional projections of m and o onto P . We define

$$P^+ = \{p + c : \text{for some } p \text{ in } P \text{ and some } c \text{ in } \mathbb{R}\} . \quad (21)$$

In this case, the projection of m onto P^+ is given by

$$p^+ = p^* + ce \quad (22)$$

where c is the regression coefficient $\langle m|e \rangle / \langle e|e \rangle$. Finally, the unconditional standard deviation of m , $\text{std}(m)$, is greater than or equal to $\text{std}(p^+)$.

Computing the regression coefficient requires knowledge of the unconditional mean of m . When this mean is not known *a priori*, we consider the family of random variables $\{m_c : c \text{ in } \mathbb{R}\}$ where

$$m_c = p^* + ce \quad . \quad (23)$$

Notice that each random variable in $\{m_c : c \text{ in } \mathbb{R}\}$ is also in $\{m_w : w \text{ in } I\}$ because random variables that are constant almost surely are in I . Hence m_c satisfies Restriction 1. Since m_c is in P^+ , $\langle m_c | m_c \rangle$ is finite and m_c satisfies Restriction 2 as well.

By changing c we vary the unconditional mean because

$$E(m_c) = E(p^*) + cE(e) \quad (24)$$

and $E(e)$ is different from zero as long as o is not in P . Therefore, for any m satisfying Restriction 2, the ordered pair $[E(m), \text{std}(m)]$ is in the region:

$$S = \{(c_1, c_2) \in \mathbb{R}^2 : c_1 = E(m_c) \text{ and } c_2 \geq \text{std}(m_c) \text{ for some } c \text{ in } \mathbb{R}\} \quad (25)$$

In this sense random variables in $\{m_c : c \in \mathbb{R}\}$ are on the unconditional mean-standard deviation frontier for IMRS's.

If S touches the horizontal axis, then without restricting $E m$ there are no restrictions on $\text{std}(m)$ other than the trivial restriction that it be nonnegative. The region S touches the horizontal axis if, and only if there is a real number c such that $\text{std}(m_c) = 0$ or equivalently m_c is equal to c . In this case,

$$\pi(p) = cE(p|I) . \quad (26)$$

This is a considerably stronger restriction than (17). It requires that the prices on all securities be proportional to their conditional means where the proportionality factor cannot depend on conditioning information.

It is of interest to compare results in this section to related results in Hansen and Richard (1987) and Hansen and Jagannathan (1989). Our conclusion that $\{m_c: c \text{ in } \mathbb{R}\}$ is a subset of $\{m_w: w \text{ in } I\}$ can be thought of as the dual to the result in Hansen and Richard (1987) that asset returns that are on the unconditional mean-standard deviation frontier are also on the conditional mean-standard deviation frontier. Hansen and Jagannathan (1989) derived a region similar to S given in (25). Their region, however, was deduced by regressing m and o onto a space that can be much smaller than P . More precisely, they form a finite-dimensional subspace of P and project m and o onto this space. As a consequence, the region they derive contains the region S given in (25).

4. Strategies for estimating mean-standard deviation frontiers

The M-dimensional process $\{y_t\}$ of observables is assumed to be strictly stationary and to possess a one-step-ahead conditional density $h(y_{t+1}|y'_{t-L+1}, y'_{t-L+2}, \dots, y'_t)$ that depends on at most L lags. For brevity we write $x_t = (y'_{t-L+1} \ y'_{t-L+2} \ \dots \ y'_t)'$, which is $ML \times 1$, and we denote the one-step density as $h(y_{t+1}|x_t)$ or $h(y|x)$. In the empirical work, y_t consists of asset payoffs and other variables that contain information about returns. Let $y_{1,t+1}$ denote the M_1 -dimensional subvector of y_{t+1} containing asset payoffs. All payoffs are assumed to have unit prices; that is, the payoffs are gross returns (1 + net returns). Hence the vector u_1 of the time t prices of $y_{1,t+1}$ consists of M_1 ones.²

For a given x, the conditional mean-standard deviation frontier for m is constructed using the first and second moments of y_1 under $h(y|x)$. Put

$$q(h,x) = \int y_1 h(y|x) dy,$$

$$Q(h,x) = \int y_1 y_1' h(y|x) dy,$$

$$V(h,x) = Q(h,x) - q(h,x)q(h,x)',$$

which are the conditional mean, second moment matrix, and covariance matrix of y_1 given x. Let \hat{h} denote an estimate of h. As can be checked using elementary least squares formulas, the conditional frontier described in Section 2 can be obtained by tracing out the parabolic-shaped region

$$\sigma(m_c|x) = \{[u_1 - cq(\hat{h},x)]'V(h,x)^{-1}[u_1 - cq(\hat{h},x)]\}^{1/2}$$

$$E(m_c|x) = c$$

as c varies over \mathbb{R} . The minimum value of $\sigma(m_c|x)$ is

$$\left\{ u_1' V(\hat{h}, x)^{-1} u_1 - \frac{[u_1' V(\hat{h}, x)^{-1} q(h, x)]^2}{q(\hat{h}, x)' V(\hat{h}, x)^{-1} q(\hat{h}, x)} \right\}^{\frac{1}{2}}$$

which is the length of the residual vector from an oblique projection of u_1 onto the space $\{cq(\hat{h}|x)\}_{c \in \mathbb{R}}$ using $V(\hat{h}|x)^{-1}$ as the weighting matrix. The value of $E(m_c|x)$ at the minimum of the frontier is $u_1' V(\hat{h}, x)^{-1} q(\hat{h}, x) / u_1' V(\hat{h}, x)^{-1} u_1$ which is the coefficient of the projection.

The estimate of the unconditional frontier is more involved as it depends upon the conditional moments at each data point. Put

$$p_t^*(h) = u_1' Q(h, x_{t-1})^{-1} y_{1t}$$

$$e_t(h) = 1 - q(h, x_{t-1})' Q(h, x_{t-1})^{-1} y_{1t}.$$

Then the estimate of the unconditional mean-standard deviation frontier is obtained by tracing out

$$E(m_c) = \frac{1}{n} \sum_{t=1}^n ce_t(\hat{h}) + p_t^*(\hat{h})$$

$$\sigma(m_c) = \left\{ \frac{1}{n} \sum_{t=1}^n [ce_t(\hat{h}) + p_t^*(\hat{h}) - E(m_c)]^2 \right\}^{1/2}$$

as c varies over \mathbb{R} , where n denotes sample size. An equivalent way to do the calculation is, for each scalar $E(m)$, to determine the c such that the first equality holds and then calculate the corresponding σ from the second equation.

Below, we report estimates of conditional and unconditional frontiers using SNP methods to obtain the estimate \hat{h} of h .

5. Data

Two data sets are used in this paper. The first is a long monthly time series on the ex-post real returns on stocks and T-bills, 1926-87, and is described more fully below. The second is used mainly for some auxiliary calculations. It consists of a shorter monthly time series, 1959-84, on ex-post real returns for the same two assets together with consumption data, which become available monthly beginning in 1959. This data set is described more fully in Gallant and Tauchen (1989).

The long time series consists of 744 monthly observations, 1926:01-1987:12, on the ex-post real return on the value weighted NYSE and the ex-post real return on a one-month T-bill. (Throughout, the format for referencing dates is yyyy:mm where yyyy denotes the year and mm denotes the month.) The length of this series appears to be sufficient for reliable estimation of the conditional density. Furthermore, the sample period covers several sub-periods of intense activity on financial markets, including among other things, two crashes.

A deflator must be used to convert nominal returns to real returns. Common practice in testing intertemporal asset pricing models is to use the consumption deflator. This index has some appeal on theoretical grounds, but is only available monthly from 1959 forward. To our knowledge, there are only two monthly price indexes available back as far as 1926. One is the wholesale price index and the other is the consumer price index. Muoio (1988) examines how closely year-to-year percentage changes in annual aggregates of each of these price indexes correspond to those of the annual consumption deflator over the period 1913-1983. His work (Chapter 2, Figures 1 and 2, p. 33) indicates that percentage movements in the annual consumer price index are quite close to those of the consumption deflator, and the agreement is better than for

movements in the wholesale price index. On this basis, we elected to use the consumer price index.

The data are plotted in Figures 1 through 2. Table 1 displays sample means, standard deviations, etc. in the top part of the table and the extreme points in the data expressed in units of standard deviation in the bottom part. Looking at the figures, the secular movements in the volatility of the two series appear to be related, and in particular, the volatilities are higher prior to 1947. From the table, the data are seen to be fat-tailed as evidenced by the large (raw) sample kurtosis and the magnitudes of the extremes. A sample of size 744 from the t-distribution with seven degrees freedom would have about 1% of the sample exceeding ± 3.5 which is in rough agreement with the number of extremes in our data.

Some of these extremes can be associated with distinct events. For instance, the 1946:07 drop in the T-bill return occurred because the nominal T-bill rate was essentially frozen at that time while the price index took a sharp rise in July, 1946, when price controls were lifted. We have not made any special adjustments for this exceptional drop in the T-bill return or any other extreme observations. There are other periods in the data set when prices were controlled during or after wars and nominal rates were slowly changing, though not to the same extent as in 1946. Taken together, these special events are, in our view, reasonably modeled as being the outcome of a stationary process, and should be included in the data set. The remarks of Leamer (1978, pp. 278) are an interesting commentary on the practice of selecting endpoints for the purpose of excluding influential observations.

One model that is suggested by a theory of speculative markets and that can account for data with these characteristics is an ARCH-type model with a fat-

tailed innovation distribution (Gallant, Hsieh, and Tauchen, 1989). Our estimation strategy, discussed below, can track data that follow this model and can accommodate (nonlinear) departures from it, if present.

Another characteristic that one might note in Table 1 is the extent to which the mean of the stock return exceeds that of the T-bill return. This reflects the so-called equity premium anomaly, which we discuss further in Sections 7 and 8 below.

Table 2a contains a summary of linear VAR estimation. There is evidence for autocorrelation in both variables, though the R^2 's indicate that the linear predictability of the stock return is very slight and substantially below that of the T-bill series. The Granger tests reveal some evidence of linear feedback from the stock return to the T-bill return at the lower order lags. The p-values for the tests shown in the table were computed using the conventional formulas for linear models, and should thus be interpreted with some caution. We do not regard them as precise tools for inference but rather as familiar statistics that give a qualitative feel for the underlying characteristics of the data.

Table 2b contains a summary of the results of fitting linear VAR models to squared residuals from VAR models in Table 2a as a crude indicator of conditional heteroskedasticity. The results suggest that relative to the information set comprised of past VAR squared residuals, there is some predictability in the magnitude of the stock residual but very little predictability in the magnitude of the T-bill residual. The same caveat regarding the p-values applies.

6. Semiparametric estimation of the conditional density $h(y|x)$

6.1 SNP models

We utilize semiparametric (SNP) methods to estimate the conditional density of the observed data. In subsequent sections, the fitted density is used to derive estimates of conditional moments for the frontier calculations. It is also used to compute expectations of certain nonlinear functions that are important for the interpretation of the frontiers.

SNP methods (Gallant and Tauchen, 1989) are a nonparametric approach for time series density estimation. The SNP density, denoted by $h_K(y|x)$, is the K^{th} term in sequence of approximations to the underlying density $h(y|x)$. The leading term of the approximation is a linear vector autoregression with Gaussian errors. The higher order terms accommodate departures from Gaussianity and possible conditional nonlinear dependence in moments, to the extent these higher order effects need to be introduced.

An SNP approximation takes the form

$$h_K(y|x) = f_K(R^{-1}[y - b_0 - Bx]|x)|R|^{-1}$$

where b_0 is $M \times 1$, B is $M \times 1$, R is an $M \times M$ upper triangular matrix, and $f_K(z|x)$, $z \in \mathbb{R}^M$, is a modified Hermite density to be defined presently. The vector b_0 and the matrix B are the parameters of the conditional mean of a VAR model and R is the upper triangular square root of the error covariance matrix. Hence, $f_K(z|x)$ is an approximation to the conditional density of the standardized VAR error. Its form is

$$f_K(z|x) = \frac{\left[\sum_{|\beta|=0}^K \sum_{|\alpha|=0}^K a(\beta, \alpha) z^\alpha x^\beta \right]^2 \varphi(z)}{I(a, x)}$$

where $\varphi(z)$ is the M dimensional standardized multivariate normal density, $I(a,x)$ is the integral of the numerator over $z \in \mathbb{R}^M$, $K=(K_z, K_x)$, and the rest of the notation is as follows. The vector α is a multi-index of length M , that is, an $M \times 1$ vector whose elements are nonnegative integers, and

$$z^\alpha = (z_1)^{\alpha_1} \cdot (z_2)^{\alpha_2} \cdot \dots \cdot (z_M)^{\alpha_M},$$

$$|\alpha| = \alpha_1 + \alpha_2 + \dots + \alpha_M.$$

Likewise, β is a multi-index of length ML and

$$x^\beta = (x_1)^{\beta_1} \cdot (x_2)^{\beta_2} \cdot \dots \cdot (x_{ML})^{\beta_{ML}},$$

$$|\beta| = \beta_1 + \beta_2 + \dots + \beta_{ML}.$$

The approximation $f_K(z|x)$ takes the form of a squared polynomial in z and x times the standard normal density. The polynomial is of degree K_z in z and K_x in x , and the coefficients are $\{a_{\alpha,\beta}\}$, with $a_{0,0}$ normalized to equal unity.

There are three tuning parameters, L , K_z , and K_x , of the SNP approximation. In the discussion below, we shall use the notation $\text{SNP}(L, K_z, K_x)$ to identify the corresponding density $h_K(y|x)$. When it is important, we will write $h(y|x; \theta)$, where θ is a vector containing the polynomial parameters and the VAR parameters of the model. The parameter vector θ is estimated using standard maximum likelihood methods. In order to obtain consistency using this representation, both K_z and K_x must grow with sample size, either deterministically or adaptively (Gallant and Nychka, 1987; Gallant and Tauchen, 1989).

SNPRX models were developed in Gallant, Hsieh, and Tauchen (1989). An SNPRX is similar to an SNP model, except that the leading term is an ARCH-type model with Gaussian errors, a linear conditional mean, and a conditional

covariance matrix that depends upon the past of the process. The higher order terms thus accommodate deviations from this model. The motivation for introducing this class of models is the extensive set of empirical results presented by Rob Engle and his collaborators that document the presence of strong conditional heteroskedasticity in financial market data (see Engle and Bollerslev, 1986, for a review). To the extent that a good parameterization can be found for the conditional covariance matrix, then the SNPRX models permit one to subsume into the leading term of the model aspects of the data that, on a priori grounds, are likely to be important.

The structure of an SNPRX model is similar to that of an SNP model, except that the matrix R is allowed to depend upon x . An SNPRX(L, K_z, K_x) model has as its leading term a linear vector autoregression with mean, $b_0 + Bx$, and Gaussian errors with conditional covariance matrix $R(x)R(x)'$. The higher order terms capture deviations from that model. In the estimations reported below, the parameterization

$$\text{vech}(R(x)) = p_0 + \sum_{j=1}^{ML} p_j \text{abs}(x_j)$$

is used, where p_j is a parameter vector of length $M(M+1)/2$, $j=0,1,\dots,ML$. In the estimation, the series $\{y_t\}$ is linearly transformed to have mean zero and covariance matrix I_M . Thus, the upper triangular square root of the covariance matrix of the leading term depends linearly on the absolute values of L lags of the process, after centering and rescaling. The parameterization makes it straightforward to impose positivity on the implied conditional variance matrix, at the expense of creating multiple peaks in the likelihood surface.

In a multivariate context like this, the optimization algorithm can use an extreme value from one series as an explanatory variable for another series. This allows it to fit an observation nearly exactly, reduce the corresponding conditional variance to near zero, and inflate the likelihood. This problem is endemic to procedures that adjust variance on the basis of observed explanatory variables. We have compensated for this effect by an additional transformation

$$\tilde{x}_i = (4/c)\exp(cx_i)/[1+\exp(cx_i)] - 2/c \quad i = 1, \dots, ML$$

with $c = 1/2$. This is a one-to-one (logistic) transformation that has a negligible effect on values of x_i between -3.5 and 3.5 but progressively compresses values that exceed ± 3.5 so they are bounded by ± 4 . The inverse transformation is $x = (1/c)\ln[(2+c\tilde{x})/(2-c\tilde{x})]$. This transformation is roughly equivalent to variable bandwidth selection in kernel density estimation. Because it affects only x , and not y , the asymptotic properties of SNP estimators discussed above are unaltered.

6.2 SNP estimation

Previous experience (Gallant, Hsieh, and Tauchen, 1989) suggests a reasonable strategy for choosing the tuning parameters L , K_z , and K_x . The first step is to use the Schwarz criterion (Schwarz, 1978) to select an initial tentative model. Since the Schwarz criterion was developed for different estimation contexts and is known to be conservative, a second step undertakes a battery of specification tests to check for whether the initial model has missed any important conditional variation in the first two moments. The initial model is then expanded, if need be, to the point where the specification tests are passed at conventional significance levels. This strategy gives sensible results, though there is certainly a need for further

theoretical work on strategies for optimally selecting the tuning parameters of these models.

Table 3a contains a summary of estimated (logistic) VAR and SNPRX models fitted on 738 observations on a bivariate process comprised of the stock return and the T-bill return, 1926:07-1987:12. The first six observations of the 744 available, are not used so that values of the objective function across models with different lag lengths are comparable. The table shows the tuning parameters, the number of parameters in each model, the value of the objective function, and the Schwarz criterion. Because the objective function is proportional to minus log likelihood,

$$s_n(\theta) = -\frac{1}{n} \sum_{t=1}^n \ln[h_K(y_t | x_{t-1}), \theta],$$

the values shown in the table decrease when more parameters are included in the model. Due to the sign change, models with smaller values in the column labeled 'Schwarz' perform better under that criterion.

In Table 3a, the SNPRX(L,0,0) models are seen to perform substantially better under the Schwarz criterion than the VAR(L) models, which are linear Gaussian vector autoregressive models save for the logit transformation to x . Of the SNPRX(L,0,0) models, the preferred model is the SNPRX(3,0,0), which indicates that a lag length of three is needed to capture the conditional heteroskedasticity. Of all of the models shown in the table, the preferred model under the Schwarz criterion is the SNPRX(3,4,0). This model has a homogeneous error density with the polynomial component being a quartic constrained to include at most quadratic interactions. A quartic was chosen because the results of (Gallant, Hsieh, and Tauchen, 1989) indicated that a quartic is needed to make the density assume the shape -- peaked near zero

and thick in the extreme tails -- that is characteristic of the probability density of long time series on movements in financial market data.

Table 3b shows the outcome of the battery of specification tests for each of the models in Table 3a. The tests for homogeneous location are the regression F tests from regressions of the standardized residuals (the residuals divided by the square roots of the predicted conditional variances) on constants and six lags of linear terms, squared terms, and cubic terms of the two series. The tests for homogeneous scale are similar except that the dependent variables are the squared standardized residuals. Because of the "Durbin effects" of prefitting discussed in Newey (1985) and Tauchen (1985), the p-values reported in Table 3a are probably too large and certainly should be interpreted with caution. However, they do suggest that there is additional conditional scale variation in the stock return that is not accounted for by the SNPRX(3,4,0) specification and that a SNPRX(3,4,1) specification is preferred. Because the Schwarz criterion is known to be conservative, we accept this implication and adopt the SNPRX(3,4,1) model as the preferred specification.

7. Conditional frontier estimates

We begin with a discussion of the general properties of the conditional frontiers and then proceed to an analysis of their economic implications.

7.1 Location and shape properties of conditional frontiers

Figure 3 shows mean-standard deviations frontiers conditional on x equaling its unconditional mean over the data set. Conditioning on this particular x value confines the analysis to the central part of the data and thereby helps circumvent the effects of the influential observations. The conditional frontiers shown in Figure 3 (and subsequent figures) were computed using the method described in Section 4 above. The figure shows frontiers calculated from four estimated specifications of the conditional density of the bivariate stock and T-bill returns series:

- (A) $y_t | x_{t-1} \sim \text{SNPRX}(3, 4, 1)$
- (B) $y_t | x_{t-1} \sim \text{SNPRX}(3, 0, 0)$
- (C) $y_t | x_{t-1} \sim \text{SNP}(3, 0, 0)$
- (D) $y_t | x_{t-1} \sim \mu \text{ constant, } \sigma \text{ constant}$

Specification (A) is the preferred specification obtained in Section 6 above. This specification entails conditional heteroskedasticity and a conditionally dependent non-Gaussian error density. Specification (B) entails only conditional heteroskedasticity with a conditionally homogeneous Gaussian error density. Specification (C) is essentially the standard VAR model with a linear mean and constant conditional covariance matrix, except for the logit transformation to x (see Section 6 above); the logit transformation essentially has no effect in the central part of the data. Specification (D) uses no conditioning information at all and is obtained by calculating the

unconditional first two moments of the series. The figure also shows points corresponding to conditional mean and standard deviations of the reciprocals of the (gross) stock and T-bill returns. The expectations were computed from the preferred SNPRX(3,4,1) specification. These points are shown because they help provide a visual sense of scale and because the reciprocal of the stock return is a candidate IMRS under certain assumptions about preferences.

Some interesting features of Figure 3 are the similarities between the frontiers computed under specifications A and B together with the contrasts between these two frontiers and those computed under specifications C and D. Both the A and B specifications account for conditional heteroskedasticity while the C and D specifications do not. The similarity between the A and B frontiers indicates that, in the central part of the data, the estimate of the conditional frontier is quite robust with respect to the specification of the model, once conditional heteroskedasticity is taken into account. Some confirmatory evidence on robustness is the frontier computed under an SNPRX(3,4,0) specification, which is not shown in Figure 3 but is shown in Figure 4 (to be discussed below). The fact that the extra complications entailed in specification A relative to B do not have noticeable effects on the frontiers is due to ignoring tail behavior by conditioning on a point in the central part of the data.

In Figure 3 the minimum of all of the frontiers occur at a value for $E(m|x)$ very near the conditional mean of the reciprocal of the

T-bill return. The minimum of the conditional frontier corresponds to the oblique projection of a vector of ones onto the conditional mean vector, with the weighting matrix being the inverse of the conditional variance matrix. (See Section 4 above.) Because of low variability of the T-bill return relative to the stock return, the T-bill component dominates that projection.

As discussed in Section 2 above, the conditional frontier touches the horizontal axis if the conditional means of the returns are equal. In this case risk neutral asset pricing cannot be ruled out. The simplest way to test the hypothesis of equal unconditional means is to test for equal unconditional means using the paired t-test on the difference between the stock and T-bill returns. There are advantages to this approach: Any errors in converting nominal returns to real returns net out (approximately) and the t-statistic is asymptotically normally distributed even in the presence of heteroskedasticity. (The null hypothesis implies all autocorrelations are zero so serial correlation need not be taken into account in computing standard errors.) The paired t-test rejects the hypothesis of equal conditional returns at a p-value of 0.003. We are relying on the large sample (744 observations) and have not done simulation to determine if the fat-tails or heteroskedasticity are severe enough to affect this conclusion.

The hypothesis can also be viewed as a conditional moment restriction on the density which can be tested using the methods in Gallant and Tauchen (1989). The top part of Table 4 shows the outcome of tests of this restriction when imposed on the bivariate fits. When imposed on the SNPRX(3,0,0) specification, the restriction is rejected at the same p-value as the paired t-test. This specification allows for conditional heteroskedasticity but imposes a normal innovation distribution. When a fat-tailed innovation distribution is

allowed, an SNPRX(3,4,0), the p-value increases to 0.1096 and one would accept the restriction at conventional significance levels. Due to problems with numerical stability the restriction was not fully imposed on the SNPRX(3,4,1) specification. Thus the p-value is overstated and is only suggestive of the consequence of allowing for additional nonlinearities in the fit. The bottom part of the Table 4 shows the results of testing the restriction that the conditional means are equal and constant, which is a stronger form of risk neutrality that implies a constant IMRS. The test outcomes provide substantial evidence against this hypothesis.

Figure 4 shows the conditional frontiers computed from the unrestricted SNPRX(3,4,0) specification and the same specification restricted to have equal conditional means, both fitted using the bivariate series. The SNPRX(3,4,0) is the most liberally parameterized model for which the restriction could be fully imposed. Since the p-value for the restriction when tested against the SNPRX(3,4,0) model is about 0.11, the downward shift in the frontier indicates the width of an approximate 10 percent confidence about the frontier in the vicinity of $E(m|x) = 1$. In view of the other p-values reported in Table 4, a 10 percent confidence band about the frontier computed under the SNPRX(3,0,0) specification would be narrower while a band for the frontier computed under the SNPRX(3,4,1) would probably be somewhat wider.

Each of the conditional frontiers shown in Figures 3 and 4 was computed at an x value set to the unconditional means of the series. Figures 5 and 6 provide some indication as to extent to which the time series variation in x shifts the conditional frontier. Each of the two figures pertain to the SNPRX(3,4,1) specification fitted to the bivariate stock and T-bill returns.

Figure 5 is a plot of the the minimum value of $\sigma(m|x_t)$ along the conditional frontier, while Figure 6 is a plot of the value of $E(m|x_t)$ at which the minimum occurs. Figure 5 indicates that nearly all of the time series variation in the minimal $\sigma(m|x)$ is concentrated in a band between 0.00 and 0.30. Figure 6 indicates the value of $E(m|x)$ where the minimum occurs is substantially more volatile in the period 1926-1951 than in the subsequent period.

7.2 Conditional first and second moment properties of IMRS's based on time separable utility

Under certain aggregation conditions (Eichenbaum and Hansen, 1989) the intertemporal marginal rate of substitution can be related to the preference function of a fictitious representative agent. Under further assumptions about time separability and homotheticity, the IMRS between time t and $t+1$ can be expressed as $m_{t,t+1} = \beta(c_{t+1}/c_t)^{-\gamma}$, where β is a subjective discount factor between zero and unity, $\gamma \geq 0$ is a curvature parameter, and $\{c_t\}$ is the equilibrium consumption process. This is the IMRS based on the CRR or power utility specification, which has been used in empirical work by Grossman and Shiller (1981), Hansen and Singleton (1982), Grossman, Melino, and Shiller (1987), and many others.

In this sub-section we examine the extent to which the conditional first and second moments of this candidate IMRS violate the bounds implicit in estimated conditional mean - standard deviation frontiers. Because the analysis is conditioned at the mean of the data, the moment bounds we examine are, in one sense, weaker than those examined by Hansen and Jagannathan (1989). The reason is that this candidate, like any other, could conceivably satisfy the moment bounds in the central part of the data while not in the extremes.

We begin with the special case $\gamma = 1$, which is logarithmic preferences. In this case one can arguably measure the implied IMRS without utilizing consumption data. In particular, one can show (Rubenstein, 1976) that $m = 1/r^*$, where r^* is the return on an aggregate wealth portfolio. This measure of m is also implied by a particular version of preferences that are not state separable (Epstein and Zin, 1989). In this case, though, the implied m is not necessarily representable as a consumption ratio as is the case for state separable preferences. Figures 3 and 4 give an indication of the extent to which this candidate IMRS violates the conditional moment bounds when we take r^* to be the stock return variable of the long time series.

More generally, though, conducting a *conditional* analysis of the first two moment properties of consumption-based IMRS's using monthly data requires that we confine attention to a period beginning in 1959, which is when monthly consumption data first becomes available. Figure 7 shows the conditional frontier implied by the SNP(2,2,1) specification that Gallant and Tauchen (1989) estimated for a trivariate series comprised of stock and T-bill returns together with consumption growth, 1959:05-1984:12. The asterisks are points indicating the conditional mean and standard deviation of $\beta(c_{t+1}/c_t)^{-\gamma}$ implied by the fitted SNP model, with numerical integration used to compute required conditional expectations. The parameter β is set to unity and γ varies from zero to 19 in increments of one. Lower values of β simply translate the set of asterisks leftwards. The asterisk at (0.0,0.0) corresponds to $\gamma=0.0$, and the points move up and to the left as γ increases. The asterisks trace out a parabola laid on its side, and ultimately pass back through the frontier at γ values in excess of 200. The figure suggests that this candidate IMRS can accommodate large conditional risk premiums only with very high values of γ . This finding is the conditional analogue of findings by Mehra and Prescott (1985) and others.

8. Conclusion

We conclude by using unconditional frontier analysis to interpret recent empirical work utilizing time nonseparable preference specifications in asset pricing applications. A discussion of some robustness considerations precedes the analysis.

8.1 Robustness

Estimates of unconditional frontiers that fully utilize the conditioning information in the data turn out to be fairly sensitive to model specification, and are thus not reported. The reason for the lack of robustness is both subtle and interesting. The calculations entail computing sample moments of quantities that depend directly on the data and on the conditional second moment matrix of the asset returns at each data point. That matrix can be written as $\mu_t \mu_t' + \Sigma_t$, where μ_t is the predicted conditional mean vector and Σ_t is the predicted conditional variance matrix. Because returns fluctuate in a relatively narrow band about unity, the matrix $\mu_t \mu_t'$ is very close to the matrix whose elements are unity. Furthermore, the elements of Σ_t are fairly small, being on the order of 10^{-3} to 10^{-5} . Hence, in a few places in the data the conditional second moment matrix can become very close to being a rank one matrix. For example, the 1984:08 observation is an otherwise innocuous looking observation, but the ratio of the largest to the smallest eigenvalue of the conditional second moment matrix of returns is 2319, which is a very large condition number. A large condition number means that the elements of $(\mu_t \mu_t' + \Sigma_t)^{-1}$ will be very sensitive to small changes in the elements of Σ_t , which appears to be the cause of the lack of robustness.

This difficulty is intrinsically multi-dimensional and will not arise when μ_t and Σ_t are scalars. We can therefore circumvent it by estimating separate

mean - standard deviation frontiers, each based on the implied conditional marginal distributions from the long series estimation, and then take their intersections.³ The strategy still exploits conditioning information, but not to its full extent. The individual frontiers are estimated using the method outlined in Section 4. In each case, the $\{y_t\}$ process is the bivariate stock and T-bills series, and the fitted conditional density is the preferred SNPRX(3,4,1) specification. One frontier is obtained by taking the $\{y_{1t}\}$ process to be the stock return alone, while the other is obtained by taking $\{y_{1t}\}$ process to be the T-bill return. Additional work suggested that these individual frontiers are reasonably robust.

Figure 8 shows these separately estimated frontiers along with pointwise two-sigma confidence bands. The one on the left corresponds to stocks while the other corresponds to T-Bills. The confidence bands were computed using the method described in Hansen and Jagannathan (1989). Covariance terms to lag seven and Parzen's weights were employed in the weighted covariance estimation. The confidence bands only give a rough indication of the estimation uncertainty and should be interpreted with caution, as they do not take account of the estimation of the incidental parameters of the fitted conditional density. Doing so in a reliable manner appears to entail some extensive Monte Carlo work, which we defer to subsequent research.

8.2 Time nonseparable preferences

Sims (1980), Novales (1989), and many others have argued that time nonseparable preferences might be needed to reconcile the time series properties of asset returns and consumption. A particularly convenient specification for time nonseparable preferences was introduced by Dunn and Singleton (1986) and Eichenbaum and Hansen (1989). In their specification, the IMRS is

$$m_{t,t+1} = \beta(c_{t+1}/c_t)^{-\gamma}(\psi_{t+1}/\psi_t)$$

$$\psi_t = E\{ [\alpha(c_{t-1}/c_t) + 1]^{-\gamma} + \alpha\beta[1 + \alpha(c_{t+1}/c_t)]^{-\gamma} | I_t \}$$

where c_t is consumption goods acquired in period t , I_t is time t information, β is a subjective discount factor, γ is a curvature parameter, and α is a parameter capturing the intertemporal service flow from previously acquired consumption goods.

The sign of the parameter α reflects the nature of the time non-separability of preferences. When α is positive, then newly acquired consumption goods have a locally durable component. Acquisitions of consumption goods in nearby periods are thus substitutable, which is a characteristic of preferences consistent with the arguments of Huang and Kreps (1987). On the other hand, when α is negative, then consumption goods are complementary across adjacent time periods. This characteristic of preferences is consistent with notions of adjustment costs and habit persistence. Constantinides (1988) argues that IMRS's based on preferences displaying habit persistence can account for the equity premium anomaly.

When this preference specification is estimated using GMM-Euler equation methods, then the estimates of α are generally statistically significant and in the range 0.10 to 0.50 (Dunn and Singleton, 1986; Eichenbaum and Hansen, 1989). This evidence is consistent with local durability of consumption. Gallant and Tauchen (1989) likewise examine the empirical implications of time non-separability, though using a different preference specification and estimation strategy. Their estimation strategy enables them to produce estimates of the means of asset returns subject to the Euler equation restrictions. They find that sets of preferences that best fit the data, in the sense of not failing

tests of overidentifying restrictions, also display local durability but at the same time predict an equity premium that is too small relative to the observed premium.

Figure 9 is useful for interpreting these empirical findings. The figure displays evidence on the impact that various assumptions regarding local durability and habit persistence have on the first two moment properties of the IMRS. The frontier shown in the figure is the intersection of the two separately estimated frontiers presented in Figure 8. The figure also shows the points corresponding to $(E(m), \sigma(m))$ for the IMRS implied by the time nonseparable preferences under various assumptions about γ and α . We calculated the coordinates of these points using the SNP(2,2,1) model fitted by Gallant and Tauchen to the shorter 1959-84 data set. We used numerical integration to compute the conditional expectations required for ψ_t , and we then formed sample moments of $\beta(c_{t+1}/c_t)^{-\gamma}(\psi_{t+1}/\psi_t)$. This calculation presumes that agents' information set I_t is current and lagged values of the two asset returns and consumption growth. The $(E(m), \sigma(m))$ points in Figure 8 are for these candidate IMRS's when α is set to -0.50, 0.00, and 0.50, and are indicated by minus signs, zeros, and plus signs, respectively. The γ values increase from zero in increments of 1. The points for $\alpha = 0.00$ and $\alpha = 0.50$ leave the left edge of the diagram for γ above 12 and then cut back through the frontier at γ values in excess of 100. In all work β is set to unity; other values for β will simply rescale the calculated means and variances equiproportionally.

When α equals 0.00, these preferences are time separable. For modest values of γ , the associated $(E(m), \sigma(m))$ points are seen in Figure 9 to violate the moment restrictions embodied in our frontier estimate based on the long

time series. This is consistent with the findings of Hansen and Jagannathan (1989) and Mehra and Prescott (1985). The same is true in Figure 10, which shows the means and standard deviations for these candidate IMRS's compared to the frontier Hansen and Jagannathan estimated using monthly scaled returns, 1959-86.

When local durability is introduced by setting α equal to 0.50, then the extent of the violation of the moment bounds is exaggerated relative to the case $\alpha = 0$. This is interesting in view of the fact that Euler equation methods generally give positive estimates of α . Apparently, the statistical characteristics of the data are such that Euler-equation methods will place more weight on fitting other moments, in particular the cross serial correlation structure of the IMRS and the observables, and relatively less weight on fitting the first two moment properties of the IMRS.

When habit persistence is introduced by setting α equal to -0.50, then the first two moment properties of these candidate IMRS's are seen to change rather dramatically. In this case the $(E(m), \sigma(m))$ points come closer to satisfying the moment restrictions with much smaller values of γ , and in particular will enter the frontier when γ equals to 14. The pattern is consistent with an argument that habit persistence can account for the apparent anomalies regarding the unconditional first two moments of asset returns. In light of the previously discussed empirical work, however, one might expect that candidate IMRS's with α negative might not do well on Euler equation checks. We computed the R^2 's from regressions of the associated Euler equation errors on three lags each of the stock return, T-bill return, and consumption growth, using the 1959-84 monthly data set. With $\alpha = -0.50$, $\gamma = 14$, the calculation gives $R^2 = 0.37$ for the stock return error and $R^2 = 0.35$ for the T-bill

equation error, and suggests that there is predictability in the Euler equation errors. The p-values of the regression F statistics are below 0.0001. Because of prefitting effects, one should not take the p-values too seriously, and we prefer to view the R^2 's as simple estimates of population quantities. By way of contrast, for the time separable logarithmic preferences, $\gamma = 1$, $\alpha = 0$, the Euler equation check gives much smaller R^2 's: 0.04 (p-value = 0.28) for stocks, and 0.19 (p-value < 0.0001) for T-bills.

Hansen and Jagannathan (1989) also explore the effects that introducing local durability and habit persistence has on the first two moments of this IMRS. Their approach does not employ an estimate of the law of motion of asset returns and consumption, which precludes computing the ψ_t 's and also entails working only with excess returns. They present indirect evidence that would lead one to expect to see the patterns we obtain in Figures 9-10. However, using only excess returns, they could uncover virtually no evidence against the hypothesis that the candidate IMRS with α negative satisfies the Euler equations. If we use only excess returns then we can produce that same finding, as our calculation gives $R^2 = 0.02$ (p-value = 0.67) for the Euler equation check on the excess return using the IMRS generated by the $\gamma = 14$, $\alpha = -0.50$ parameter configuration. Above, however, we uncover evidence against the same hypothesis, which suggests that additional discriminatory information has been incorporated into the analysis.

Appendix

We first describe the construction of the variables in the data set and then give the detailed sources.

Stock Return: This is the real value weighted return on the NYSE computed as

$$\begin{aligned} \text{gross: } & (1 + vw_t)(cpi_{t-1}/cpi_t) \\ \text{net: } & (1 + vw_t)(cpi_{t-1}/cpi_t) - 1 \end{aligned}$$

where vw_t is the net nominal value weighted return from the end of month $t-1$ to the end of month t , and cpi_t is the consumer price index in month t .

T-bill Return: This is the real one-month return on a one-month T-bill computed as

$$\begin{aligned} \text{gross: } & (1 + tb_t)(cpi_{t-1}/cpi_t) \\ \text{net: } & (1 + tb_t)(cpi_{t-1}/cpi_t) - 1 \end{aligned}$$

where tb_t is the net nominal return on a one-month T-bill and cpi_t is the consumer price index in month t .

Sources

Both the nominal stock and T-bill returns were obtained from the Center for Research in Security Prices (CRSP), University of Chicago, Chicago, Illinois. The nominal stock return is the CRSP Value Weighted Index for the New York Stock Exchange. The nominal T-bill return is Ibbotson and Sinquefeld series on the one-month return on one-month Treasury Bills.

The price index is the Consumer Price Index - All Urban Consumers (CPI-U; 1967=100). The specific sources are as follows:

1986:12-1987:12 — U.S. Department of Commerce, Bureau of Economic
Analysis, Survey of Current Business, January
1988, Vol. 68, No. 1, p. S-5.

1925:12-1986:11 — Computer printout supplied by U.S. Department of
Labor, Bureau of Labor Statistics, Washington
D.C., 20212.

Table 1

Basic Statistics, 1926:01-1987:12

	Mean	Std. Dev.	Skewness	Kurtosis	Min	Max
Stocks	1.006751	0.057370	0.3882	10.8331	0.7153	1.3805
T-bills	1.000369	0.005765	-1.6319	17.1875	0.9446	1.0234

Extremes in Units of Standard Deviation

Lowest	Date	Highest	Date
----- Stocks -----			
-5.08394	1931:09	3.50645	1933:05
-4.21136	1938:03	3.97934	1938:06
-3.99084	1940:05	5.58224	1932:07
-3.92780	1987:10	6.48917	1932:08
-3.52938	1929:10	6.51820	1933:04
----- T-Bills -----			
-9.68646	1946:07	2.72965	1930:07
-4.90222	1933:07	2.74393	1931:01
-4.08944	1946:11	2.80872	1932:02
-3.98466	1947:09	3.82211	1927:07
-3.74431	1946:08	3.99518	1932:01

Table 2a

Summary of Bivariate Linear Vector Autoregressive Estimation
 Period: 1926:07-1987:12

Adjusted R²

Dep. Var.	Lag		
	1	3	6
Stocks	.02	.03	.03
T-bills	.30	.34	.35

Granger Causality Tests

Dep. Var.	Lag = 1		Lag = 3		Lag = 6	
	Stocks	T-bills	Stocks	T-bills	Stocks	T-bills
Stocks	**	*	***		***	
T-bills		***	***	***	***	***

*** denotes p-value \leq .001
 ** denotes .001 < p-value \leq .01
 * denotes .01 < p-value \leq .05

Incremental Granger Causality Tests

Dep. Var.	Lag 1 of 1		Lag 2,3 of 3		Lag 4,5,6 of 6	
	Stocks	T-bills	Stocks	T-bills	Stocks	T-bills
Stocks	**	*	**		*	
T-bills		***	***	***		***

*** denotes p-value \leq .001
 ** denotes .001 < p-value \leq .01
 * denotes .01 < p-value \leq .05

Table 2b

Summary of Bivariate Autoregressive Estimation on Squared Residuals
 Period: 1926:07-1987:12

Adjusted R²

Dep. Var.	Lag		
	1	3	6
eeST	.07	.10	.09
eeTB	<.01	.02	.01

Granger Causality Tests

Dep. Var.	Lag = 1		Lag = 3		Lag = 6	
	eeST	eeTB	eeST	eeTB	eeST	eeTB
eeST	***		***		***	
eeTB						

*** denotes p-value \leq .001
 ** denotes .001 < p-value \leq .01
 * denotes .01 < p-value \leq .05

Incremental Granger Causality Tests

Dep. Var.	Lag 1 of 1		Lag 2,3 of 3		Lag 4,5,6 of 6	
	eeST	eeTB	eeST	eeTB	eeST	eeTB
eeST	***		***			
eeTB						

*** denotes p-value \leq .001
 ** denotes .001 < p-value \leq .01
 * denotes .01 < p-value \leq .05

Table 3a

Bivariate Models for Stocks and T-bills
 Period: 1926:07-1987:i2

Fitted Models

Model ^a	L	K _Z	K _X	p _θ	s _n	Schwarz
VAR	0	0	0	5	2.837857	2.860075
VAR	1	0	0	9	2.651005	2.690997
VAR	2	0	0	13	2.636136	2.693902
VAR	3	0	0	17	2.618241	2.693782
VAR	4	0	0	21	2.608032	2.701348
SNPRX	1	0	0	15	2.482177	2.548831
SNPRX	2	0	0	25	2.390109	2.501199
SNPRX	2	4	0	34	2.242368	2.393449
SNPRX	2	4	1	74	2.177562	2.506386
SNPRX	3	0	0	35	2.296991	2.452516
SNPRX	3	4	0	44	2.180962	2.376479
SNPRX	3	4	1	84	2.125640	2.498900
SNPRX	4	0	0	45	2.254970	2.454931

Note: a. In models with $K_Z > 0$ third and higher order interactions have been deleted. In models with $K_X > 0$ the lag length of the polynomial in x is 2.

Table 3b

Bivariate Models for Stocks and T-bills
Period: 1926:07-1987:12

Specification Tests

Model ^b	L	K _Z	K _X	Location				Scale			
				Stocks		T-bills		Stocks		T-bills	
				F	p-val	F	p-val	F	p-val	F	p-val
VAR	0	0	0	2.93	.00	13.39	.00	10.07	.00	3.43	.00
VAR	1	0	0	2.60	.00	3.49	.00	9.70	.00	1.78	.00
VAR	2	0	0	2.59	.00	2.83	.00	9.69	.00	1.67	.01
VAR	3	0	0	2.32	.00	2.30	.00	9.56	.00	1.62	.01
VAR	4	0	0	2.25	.00	1.90	.00	9.70	.00	1.67	.01
SNPRX	1	0	0	1.76	.00	2.41	.00	4.47	.00	1.82	.00
SNPRX	2	0	0	1.44	.05	1.80	.00	2.24	.00	1.94	.00
SNPRX	2	4	0	1.62	.01	1.81	.00	2.88	.00	1.85	.00
SNPRX	2	4	1	1.27	.14	1.72	.01	1.42	.05	1.90	.00
SNPRX	3	0	0	1.18	.22	1.09	.34	1.41	.06	1.04	.41
SNPRX	3	4	0	1.42	.05	1.08	.34	1.90	.00	1.19	.21
SNPRX	3	4	1	1.20	.20	1.11	.31	0.86	.71	1.18	.22
SNPRX	4	0	0	1.04	.41	0.70	.91	1.35	.09	0.74	.87

Note: b. The location diagnostic is a regression of residuals from the fit on 6 lags of linear, quadratic, and cubic ST TB. The scale diagnostic is a regression of squared residuals on same. The F has 36 degrees freedom in the numerator and 701 in the denominator.

Table 4

Bivariate Tests That Frontiers Touch the Horizontal Axis^a

Conditional Frontiers

Model ^a	L	K _Z	K _X	p _θ	s _n	Schwarz	p-value
Constrained	3	0	0	28	2.311790	2.436210	.0027
SNPRX	3	0	0	35	2.296991	2.452516	
Constrained	3	4	0	31	2.194144	2.331895	.1096
SNPRX	3	4	0	44	2.180962	2.376479	
Constrained	3	4	1	67	2.139290	2.437010	.2668
SNPRX	3	4	1	84	2.125640	2.498900	

Unconditional Frontier

Model ^a	L	K _Z	K _X	p _θ	s _n	Schwarz	p-value
Constrained	3	0	0	22	2.414714	2.512473	.0000
SNPRX	3	0	0	35	2.296991	2.452516	
Constrained	3	4	0	31	2.284002	2.421752	.0000
SNPRX	3	4	0	44	2.180962	2.376479	
Constrained	3	4	1	54	2.199811	2.439764	.0000
SNPRX	3	4	1	84	2.125640	2.498900	

Note: a. In models with $K_Z > 0$ third and higher order interactions have been deleted. In models with $K_X > 0$ the lag length of the polynomial in x is 2.

Footnotes

1. For the purposes of this discussion we are ignoring the restriction that m be strictly positive. Hansen and Jagannathan (1989) showed that additional implications are obtained by imposing positivity.
2. In our empirical analysis we use asset returns, i. e., asset payoffs with unit prices. Scaling payoffs to have unit prices is more than just a normalization in our analysis unless the scale factors are incorporated into the information set used by the econometrician. A computational advantage to using returns is that we do not have to augment the information set with the asset prices because these prices are degenerate by construction.
3. Notice that volatility bounds on the IMRS could also be obtained by examining the univariate law of motion for a single security return. Cochrane (1988) exploited this observation in deducing bounds implied by market-wide returns from the New York Stock Exchange.

REFERENCES

- Cochrane, J. H., 1988, Bounds on the variance of discount rates implied by long horizon predictability of stock returns, Manuscript (University of Chicago, Chicago, IL).
- Constantinides, G. M., 1988, Habit formation: A resolution of the equity premium puzzle, Manuscript (University of Chicago, Chicago IL).
- Dunn, K. and K. J. Singleton, 1986, Modeling the term structure of interest rates under non-separable utility and durability of goods, *Journal of Financial Economics* 17, 27-55.
- Eichenbaum M., and L. P. Hansen, 1989, Estimating models with intertemporal substitution using aggregate time series data, *Journal of Business and Economic Statistics*, forthcoming.
- Epstein, L. G. and S. E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns, *Econometrica* 57, 937-970.
- Engle, R. F., and T. Bollerslev, 1986, Modeling the persistence of conditional variances, *Econometric Reviews* 5, 1-50.
- Gallant, A. R., *Nonlinear statistical models* (Wiley, New York).
- Gallant, A. R., D. Hsieh, and G. E. Tauchen, 1989, On fitting a recalcitrant series: the Pound/Dollar exchange rate, 1974-83, in: W. Barnett, J. Powell, and G. Tauchen, eds., *Nonparametric and semiparametric methods in econometrics and statistics, Proceedings of the fifth international symposium in economic theory and econometrics*, (Cambridge University Press, Cambridge), forthcoming.
- Gallant, A. R. and D. Nychka, 1987, Semiparametric maximum likelihood estimation, *Econometrica* 55, 363-90.
- Gallant, A. R. and G. E. Tauchen, 1989, Semiparametric estimation of conditionally constrained heterogeneous processes: Asset pricing applications, *Econometrica*, forthcoming.
- Grossman, S. and R. Shiller, 1981, The determinants of the variability of stock market prices, *American Economic Review* 71, 222-227.
- Grossman, S., A. Melino, and R. Shiller, 1987, Estimation of the continuous time consumption based asset pricing model, *Journal of Business and Economic Statistics* 5, 315-327.
- Hansen, L. P. and R. Jagannathan, 1989, Restrictions on intertemporal marginal rates of substitution implied by asset returns, Manuscript (Northwestern University, Evanston, IL).

- Hansen, L. P. and S. F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587-613.
- Hansen, L. P. and K. J. Singleton, 1982, Generalized instrumental variables estimators of nonlinear rational expectations models, *Econometrica* 50, 1269-1286.
- Huang, C., and D. Kreps, 1987, On intertemporal preferences with a continuous time dimension, I: The case of certainty, Manuscript (Massachusetts Institute of Technology, Cambridge, MA).
- Leamer, E., 1978, Specification searches: Ad hoc inference with nonexperimental data (Wiley, New York).
- Mehra, R. and E. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145-162.
- Muoio, P., 1988, Empirical tests of the utility-based asset pricing model using temporally aggregated twentieth century data, Duke University Ph.D. thesis.
- Newey, W. K., 1985, Maximum likelihood specification testing and conditional moment tests, *Econometrica* 53, 1047-1071.
- Newey, W. K., and K. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703-708.
- Novalés, A., 1989, Solving nonlinear rational expectations models: A stochastic equilibrium model of interest rates, *Econometrica*, forthcoming.
- Rubenstein, M., 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics* 7, 407-425.
- Schwarz, G., 1978, Estimating the dimension of a model, *Annals of Statistics* 6, 461-464.
- Sims, C. A., 1980, International evidence on monetary factors in macroeconomic fluctuations, Manuscript (University of Minnesota, Minneapolis, MN).
- Tauchen, G., 1985, Diagnostic testing and evaluation of maximum likelihood models, *Journal of Econometrics* 30, 415-443.
- U.S. Department of Commerce, Bureau of Economic Analysis, 1988, Survey of Current Business, January 1988, Vol. 68, No. 1.

Real Value Weighted Return, 1926 - 1987




Fig. 1

Real T - bill Return, 1926 - 1987





Fig. 2


Conditional Frontier at $x =$ Unconditional Mean
 Series = Bivariate: Stocks & T-bills, 1926 - 1987
 Various Specifications for the SNP Model



s: Stocks, b: T-bills
 A: SNPRX(3.4,1), B: SNPRX(3,0,0), C: SNP(3,0,0), D: Constant mu & sigma

Fig. 3

Conditional Frontier at $x =$ Unconditional Mean
Series = Bivariate: Stocks & T-bills, 1926-1987
Constrained and Unconstrained SNPRX(3,4,0)



s: Stocks, b: T-bills

Fig. 4

$\sigma(m|x)$ at the Minimum of the Conditional Frontier
Series = Bivariate: Stocks & T-bills, 1926 - 1987
Law of Motion = SNPRX(3,4,1)




Fig. 5

E(m|x) at the Minimum of the Conditional Frontier
Series = Bivariate: Stocks & T-bills, 1926-1987
Law of Motion = SNPRX(3,4,1)





Fig. 6


Conditional Frontier at $x = \text{Unconditional Mean}$
 Series = Stocks, T-bills, & Consumption Growth, 1959 – 1984
 Fitted Law of Motion = SNP(2,2,1)



s: Stocks, b: T – bills
 *: Time separable (CRR or power) utility, γ varies from 0.00 in increments of 1.0

Fig. 7

Unconditional Single Asset Frontiers with Two Sigma Confidence Bands
Series = Bivariate: Stocks & T-bills, 1926-1987
Law of Motion = SNPRX(3,4,1)



s: Stocks, b: T-bills

Fig. 8

Envelope of Unconditional Single Asset Frontiers
 Series = Bivariate: Stocks & T-bills, 1926 - 1987
 Law of Motion = SNPRX(3,4,1)





Fig. 9

Unconditional Frontier with Two Sigma Confidence Band
 Series = Monthly Scaled Returns, 1959 - 86



- o: time separable utility, $\alpha = 0.0$
- +: time nonseparable utility, $\alpha = 0.50$
- : time nonseparable utility, $\alpha = -0.50$
- γ varies from 0.00 in increments of 1.0

Fig. 10