IMPLICATIONS OF SECURITY MARKET DATA FOR MODELS OF DYNAMIC ECONOMIES

Lars Peter Hansen

Ravi Jagannathan

IMPLICATIONS OF SECURITY MARKET DATA FOR MODELS OF DYNAMIC ECONOMIES

ABSTRACT

We show how to use security market data to restrict the admissible region
for means and standard deviations of intertemporal marginal rates of
substitution (IMRS's) of consumers. Our approach is (i) nonparametric and
applies to a rich class of models of dynamic economies; (ii) characterizes the
duality between the mean-standard deviation frontier for IMRS's and the
familiar mean-standard deviation frontier for asset returns; and (iii)
exploits the restriction that IMRS's are positive random variables. The region
provides a convenient summary of the sense in which asset market data are
anomalous from the vantage point of intertemporal asset pricing theory.

Lars Peter Hansen   Ravi Jagannathan
University of Chicago   University of Minnesota
Department of Economics   271 19th Avenue South
1126 East 59th Street   Management and Economics Building
Chicago, IL 60637   Minneapolis, MN 55455

INTRODUCTION

In this paper we investigate the implications of asset market data for a rich class of models of dynamic economies. The models within this class differ with respect to the heterogeneity of consumers' preferences, the span of the payoffs on tradeable securities and the role of money in the acquisition of consumption goods. In spite of these differences, a common implication of these models is that the equilibrium price of a future payoff on any traded security can be represented as the expectation (conditioned on current information) of the product of the payoff and an appropriately interpreted intertemporal marginal rate of substitution ($IMRS$) of any consumer [e.g. see LeRoy (1973), Rubinstein (1976), Lucas (1978), Breeden (1979), Harrison and Kreps (1979) and Hansen and Richard (1987)]. This representation is a generalization of the familiar tenet from price theory that prices should equal marginal rates of substitution. To apply this principle to models of asset pricing, securities are viewed as claims to a numeraire good indexed by future states of the world.

If price data were available from a complete set of security markets, the $IMRS$'s of all consumers could be inferred from Arrow-Debreu prices. However, economic agents may not trade in a complete set of contingent claims markets. Furthermore, it may be practical for an econometrician only to use data on a small array of securities. Due to these limitations, asset market data alone is typically not sufficient to identify $IMRS$'s.

One approach that has been used extensively is to identify $IMRS$'s by restricting them to be parametric functions of data observed by an econometrician, [e.g. see Hansen and Singleton (1982), Brown and Gibbons (1985) and Epstein and Zin (1989)]. This approach imposes potentially stringent limits on the class of admissible asset pricing models, and then

1

tests whether the particular parameterizations are consistent with the observed asset market data.

While this parametric approach has yielded interesting insights into the empirical plausibility of particular families of models, the approach proposed in this paper goes to another extreme. We purposely enlarge the class of asset pricing models under investigation by imposing as little structure as possible on the admissible class of models. In so doing we eliminate most of the testable implications except possibly for the *Law of One Price* (portfolios with the same payoffs have the same price) and the absence of arbitrage opportunities (nonnegative payoffs that are positive with positive probability have positive prices). Although we are not able to identify the *IMRS*'s fully, we can extract information about them. When *IMRS*'s are constant, portfolio payoffs with the same price must also have the same mean. Thus the existence of portfolios of securities with the same price but distinct expected payoffs implies that *IMRS*'s must vary. We exploit this observation to derive greatest lower bounds on the standard deviations of *IMRS*'s, i.e. volatility bounds. These bounds are expressed most conveniently as regions of admissible mean–standard deviation pairs for the *IMRS*'s.

The existence of volatility bounds on *IMRS*'s was originally noted by Shiller (1982) [see also Hansen (1982)]. His goal was to construct a diagnostic for a particular family of asset pricing models that is insensitive to the alignment of the data. The volatility implications he deduced for *IMRS*'s used only two asset returns and, even for the two-asset case, are weaker than those reported here.

Our reasons for examining volatility bounds are somewhat different from Shiller's. First, our nonparametric approach can serve as a useful

2

complement to the parametric approach that is prevalent in the literature. In particular, it can assist in understanding better why particular models are rejected on the basis of statistical tests: does the parameterization admit too little variability in the *IMRS*'s? Second, they provide a common set of diagnostics for a potentially large class of asset pricing models. These diagnostics can also be used to evaluate models in which *IMRS*'s are parameterized as functions of observables as well as models for which moments can be computed from characterizations of the stochastic equilibria. Third, they allow us to determine which asset market data sets present the most stringent restrictions for *IMRS*'s and consequently the most startling implications for dynamic economic models. Our approach allows us to make these comparisons without having to focus on a parametric family of such models.

To illustrate these points, we provide an alternative characterization of the so called *equity premium puzzle* [e.g. see Mehra and Prescott (1985)]. In contrast to other characterizations, ours does not depend either on a Markov chain approximation with a small number of states or on a narrow class of asset valuation models. Figure 1 reports a restricted region for the means and standard deviations of *IMRS*'s implied by the annual (1891-1985) time series data on stocks and bonds used by Campbell and Shiller (1988). The shaded region gives the admissible pairs of means and standard deviations for *IMRS*'s. As benchmarks, we also report time series sample means and standard deviations for *IMRS*'s implied by a representative consumer model with commonly-used period utility functions of the form:

$$U(c) = \frac{c^{\gamma+1} - 1}{\gamma+1}$$

3

STANDARD DEVIATION

FIG. 1 - IMRS FRONTIER COMPUTED USING ANNUAL DATA

MEAN

4

for negative values of γ. For this specification of preferences, the *IMRS* can be measured by forming a consumption ratio for two different points in time, raising it to the power γ and discounting. For illustrative purposes the annual subjective discount factor is taken to be .95. The "∎" symbols represent mean-standard deviation pairs for alternative values of γ ranging from 0 to −30. As $|\gamma|$ increases, the volatility of the *IMRS* increases but the effect on the mean of the *IMRS* is not uniform. Initially the mean decreases but subsequently increases so that for large $|\gamma|$ the "∎"'s are in the admissible shaded region.

Our strategy for constructing regions such as that reported in Figure 1 is to construct minimum variance random variables with prespecified means that are related to asset payoffs and prices in the same manner as the *IMRS*'s. We refer to such random variables as being on the *mean-standard deviation frontier* for *IMRS*'s. In section II we construct these frontier random variables ignoring the fact that *IMRS*'s must be positive. In this case the minimum variance random variables are simply linear combinations of the asset payoffs translated by a constant. As a by-product of this construction, we relate our analysis to two commonly used empirical paradigms in finance: mean-variance analysis and linear factor pricing. More precisely we characterize the *duality* between the mean-standard deviation frontier for *IMRS*'s and the familiar mean-standard deviation frontier for asset payoffs. This analysis reveals that asset payoffs on the mean-standard deviation frontier are sufficient to generate the mean-standard deviation frontier for *IMRS*'s. Hence the dimensionality-reduction techniques used in linear factor pricing models can be exploited to derive a region like that reported in Figure 1.

In section III we modify the analysis of section two by incorporating

5

the restriction that *IMRS*'s are positive random variables. For prespecified means, we construct nonnegative random variables that behave like *IMRS*'s and have minimum variances. These random variables are not necessarily linear functions of the payoffs but instead can be interpreted as European *call* and *put* options on portfolios of these payoffs. In contrast to the analysis in section two, for some prespecified means there may not be any nonnegative random variables with finite second moments that behave like an *IMRS*'s. While the approach of this section yields more restrictive (and therefore more informative) volatility bounds, these sharper bounds are harder to compute.

In section IV we illustrate the results in sections II and III by displaying volatility bounds computed using alternative data sets, and generating mean-standard deviation pairs for alternative parametric models of *IMRS*'s. Among other things, we use these bounds to help assess the plausibility of some parametric models of asset prices.

6

I. A General Model of Asset Pricing

In this section we present a general model of asset valuation. Consider an environment in which multiple consumers trade in securities markets. The preferences and information sets of these consumers may be heterogeneous. We fix both the trading period (say time zero) and the time period for the receipts of the asset payoffs (say time $\tau > 0$). Let I^j denote the information set of consumer j at time zero, and $I = \bigcap I^j$ where the intersection is taken over the consumers in the economy who trade securities. The prices of securities traded at date zero are presumed to be in the individual information set I^j of individual j for each j and hence in I. Let P denote a set of portfolio payoffs of the numeraire good at time τ that are traded at time zero. Since the prices of the portfolio payoffs are in I, we represent these prices as a function π_I mapping P into I. Hence $\pi_I(p)$ is the price at time zero of a portfolio which will pay p units of a numeraire good at a future date τ.

Consumers are presumed to solve optimal portfolio problems in determining their asset holdings. This imposes restrictions relating marginal rates of substitution to asset payoffs and prices. To see this let mu_0^j and mu_τ^j denote the equilibrium marginal utilities of consumer j in terms of the numeraire consumption good at date zero and τ respectively. In equilibrium the marginal utility-scaled price must equal the expected marginal utility-scaled payoff conditioned on I^j:

(1.1) $mu_0^j \pi_I(p) = E(mu_\tau^j p | I^j)$ for all p in P.

As long as consumer j is not satiated at time zero, $mu_0^j > 0$ and we can divide both sides of (1.1) by mu_0^j which yields:

7

$$\pi_I(p) = E(pm^j \mid I^j) \quad \text{for all } p \text{ in } P$$

where $m^j = mu_\tau^j / mu_0^j$ is the *intertemporal marginal rate of substitution* (*IMRS*) of consumer j. Since asset prices are presumed to be observed by all consumers, it follows from the Law of Iterated Expectations that

$$(1.2) \qquad \pi_I(p) = E(pm^j \mid I) \quad \text{for all } p \text{ in } P.$$

In a world with common information sets and complete markets, marginal rates of substitution are equated across consumers ($m^j = m$ for all j). In such a world, P can be chosen to be sufficiently large so that the common *IMRS* is uniquely determined by (1.2). In general, (1.2) does not uniquely determine m^j. As we will see, however, (1.2) does restrict the unconditional moments of m^j even when markets are incomplete.[1] Since the restrictions we derive apply to all of the individual *IMRS*'s, to simplify notation we drop the j superscript on m.

We now give a more complete description of P and the associated asset pricing function π_I. We do not require that P contain all of the portfolio payoffs that are traded by consumers. Omitting payoffs will, however, weaken the implications for m. As a matter of convenience, we consider the case in which there is an n-dimensional vector x of asset payoffs at date τ. The time zero prices of these assets can also be represented as an n-dimensional vector, say q, and pricing relation (1.2) can be expressed as

$$(1.3) \qquad q = E(xm \mid I) \quad .$$

8

We are interested in the implications of (1.3) for the intertemporal marginal rate of substitution m. To investigate this relation empirically, we must have some way to replicate observations on payoffs, prices and information over time. As in Hansen and Richard (1987), we imagine an environment in which relation (1.3) is replicated over time. In other words there is a composite process $\{(m_t, x_t, q_t)\}$ and a sequence of information sets $\{I_t\}$ that satisfies a version of (1.3) for all t. Econometricians seeking to study this economy are presumed to have data on a finite record (x_t, q_t), for t=1,2,...,T, and the composite process $\{(m_t, x_t, q_t)\}$ is presumed to be sufficiently regular so that a time series version of a law of large numbers applies. Thus sample moments formed from the finite records of data converge to population counterparts as the sample size T becomes large. Even though asset prices are determined τ periods prior to the realization of the asset payoffs, from the vantage point of econometricians, we model $\{q_t\}$ as a stochastic process to accommodate possible variation over time in the asset prices. In what follows we use the unconditional expectation operator E to represent the limit points of the time-series averages of the sample moments.[2]

We now impose restrictions on m, x and q which are expressed in terms of unconditional expectations.

Assumption A1: $E|m|^2 < \infty$, $E|x|^2 < \infty$, Exx' is nonsingular and $E|q| < \infty$.

The restriction that the second moment matrix of x is nonsingular is made as a matter of convenience to rule out cases in which the entries of x are linearly dependent. Among other things, this guarantees that the Law of One Price holds trivially for linear combinations of x. If the moment

9

restrictions imposed on x and q are not satisfied for an original vector of assets, then it is often possible to scale the payoffs and prices so that these restrictions are satisfied. A special case of such scaling is when all of the payoffs are constructed to have a unit price as in the case of measured returns to holding securities between time zero and time τ.

Applying the Law of Iterated Expectations to the pricing relation (1.3) results in the following restriction:

Restriction R1: $Eq = Exm$.

We focus on the unconditional moment *R1* instead of the conditional moment restriction (1.3) because it is typically easier to estimate unconditional moments rather than conditional moments. Restriction *R1*, however, is in general weaker than (1.3). Gallant, Hansen and Tauchen (1989) show how to extend some of the analysis in this paper by exploiting characterizations of the moments of x conditioned on (possibly a subset of) I.

As long as consumers are not satiated in the numeraire consumption good at time τ, the *IMRS* should be strictly positive:

Restriction R2: $m > 0$.

Restriction *R2* is sufficient to imply the absence of arbitrage opportunities. That is, *R2* guarantees that nonnegative payoffs that are strictly positive with positive probability conditioned on I have positive prices. In the next two sections we explore the implications that *R1* and *R2* have for the mean and standard deviation of m.

So far, we have treated the case in which only a finite vector of asset

10

payoffs and prices are investigated. In our subsequent analysis, it will be convenient to extend the pricing function and its unconditional expectation to the linear span of x. Define[3]

$$P \equiv \{c \cdot x \; : \; c \text{ in } \mathbb{R}^n\}.$$

In section three we will also consider derivative claims formed by taking particular nonlinear functions of payoffs in P. In light of $A1$ each portfolio payoff in P has a finite second moment. With this in mind we define a norm on P to be

$$\|p\| \equiv [E(p^2)]^{1/2}.$$

Notice that the standard deviation of a portfolio payoff p, denoted $\sigma(p)$, is given by $\|p - Ep\|$.

Since the portfolio payoffs in x are linearly independent, for each p in P there is a unique c in $\mathbb{R}^n$ for which p is equal to $c \cdot x$. We extend the pricing function so that the prices of these payoffs are given by the corresponding linear combinations of q:

$$\pi_I(c \cdot x) \equiv c \cdot q$$

As required, π_I maps P into I. Notice π_I is constructed so that (1.3) extends to the linear span x:

$$\pi_I(p) = E(pm|I) \quad \text{for all } p \text{ in } P.$$

11

It is also of interest to define a functional π mapping portfolio prices into the *expected value* of the prices:

$$\pi(p) \equiv E\pi_I(p) \quad .$$

Hence π maps P linearly into the real line $\mathbb{R}$. Again the Law of Iterated Expectations implies that

(1.4) $\pi(p) = E(mp)$ for all p in P.

It is straightforward to show that restriction (1.4) is equivalent to *R1*.

II. IMPLICATIONS OF RESTRICTION 1

In this section we characterize the volatility restrictions for m as implied by Restriction $R1$. In subsection II.A we suppose there is a unit payoff in P, while in subsection II.B we consider the more common case in which such a payoff is not included in P. Finally, in subsection II.C we describe how existing empirical methodologies in finance can be used to characterize these volatility restrictions.

II.A Riskless Payoff

Suppose P contains a payoff that is equal to one with probability one. In deriving implications for the volatility of m, it is first convenient to construct a random variable $m^{\bullet}$ in P that satisfies $R1$. This amounts to finding a vector α_{o} in $\mathbb{R}^{n}$ such that

$$(2.1) \qquad Exx'\alpha_{o} = Eq$$

where $m^{\bullet} = x \cdot \alpha_{o}$. Solving (2.1) for α_{o} gives

$$\alpha_{o} = (Exx')^{-1}Eq.$$

Notice that α_{o} depends on the second moment of x and the first moment of q. Hence $m^{\bullet}$ can be constructed from asset market data.

Consider any other random variable m satisfying $R1$. Since P contains a unit payoff,

$$Em = \pi(1) = Em^{\bullet}.$$

13

Consequently all random variables m that satisfy $R1$ have the same mean, and this mean is equal to the expected price of a unit payoff. Also,

$$E[x(m-m^*)] = 0.$$

because both m and m^* satisfy $R1$. In other words the discrepancy between m and m^* is orthogonal to the random vector x. Since m^* is in P, m^* is the least squares projection of m onto P and

$$\sigma^2(m) = \sigma^2(m^*) + \sigma^2(m-m^*).$$

Therefore, we have the following relations:

(2.2) $\sigma(m) \geq \sigma(m^*)$ and $Em^* = Em$.

The volatility bound in (2.2) is as sharp as possible because m^* satisfies $R1$ by construction.

II.B   No Riskless Payoff

Next we consider the more usual case in which P does not contain a unit payoff. It turns out that much of the previous analysis can be exploited in analyzing this case. Let x^a denote the $(n+1)$-dimensional random vector formed by augmenting x with a unit payoff. Since Exx' is nonsingular, and no linear combination of x is equal to one with probability one, $Ex^a x^{a'}$ is also nonsingular. We build an augmented payoff space P^a containing a unit payoff by using x^a in place of x.

14

To apply the analysis in section II.A, we must assign a number v to $\pi(1)$, which is the expected price of a unit payoff. Such price data may not be available, and for this reason we examine implications for an array of hypothetical expected prices. Let v be any candidate for $\pi(1)$ and π_v the corresponding extension of π from P to P^a. We then replicate the analysis in subsection II.A to construct a random variable m_v in P^a such that

$$(2.3) \qquad Exm_v = Eq, \quad Em_v = v.$$

The counterpart to volatility bound (2.2) is

$$(2.4) \qquad \sigma(m) \geq \sigma(m_v)$$

for any random variable m that satisfies RI and has mean v. This volatility bound is as sharp as possible because, by construction, m_v satisfies RI and has mean v.

We replicate the construction of m_v for all real numbers v and generate an indexed collection $\{m_v : v \text{ in } \mathbb{R}\}$ of random variables, each of which satisfies RI. This collection is of interest because for any m satisfying RI, the ordered pair $[Em, \sigma(m)]$ is in the region:

$$(2.5) \qquad S \equiv \{(v, w) \text{ in } \mathbb{R}^2 : w \geq \sigma(m_v)\} .$$

This region summarizes the volatility implications for m implied by RI. We refer to the boundary of S as being the *mean-standard deviation frontier* for *IMRS*'s, and we refer to members of the set $\{m_v : v \text{ in } \mathbb{R}\}$ as being on this frontier.

15

It is of interest to derive an expression for $\sigma(m_\mathbf{v})$ that is both easy to compute and interpret. The moment conditions in (2.3) can be rewritten in terms of the covariance of m and x:

$$(2.6) \qquad E[(x-Ex)(m_\mathbf{v}-\mathbf{v})] = Eq - \mathbf{v}Ex.$$

Now

$$(2.7) \qquad m_\mathbf{v} = (x-Ex)'\beta_\mathbf{v} + \mathbf{v}$$

for some $\beta_\mathbf{v}$ in $\mathbb{R}^n$ because $m_\mathbf{v}$ is a linear combination of a unit payoff and the entries of x and $Em_\mathbf{v}$ is $\mathbf{v}$. Substituting (2.7) into (2.6) and solving for $\beta_\mathbf{v}$ gives

$$\beta_\mathbf{v} = \Sigma^{-1}(Eq - \mathbf{v}Ex)$$

where Σ is the covariance matrix of x. It follows that

$$(2.8) \qquad \sigma(m_\mathbf{v}) = [(Eq - \mathbf{v}Ex)'\Sigma^{-1}(Eq - \mathbf{v}Ex)]^{1/2} \quad .$$

Notice that for a given $\mathbf{v}$, $\sigma(m_\mathbf{v})$ depends only on the means of q and x and the covariance matrix of x.

The standard deviation bound given in (2.8) has the following interpretation. Consider a risk-neutral valuation of the asset payoffs in which m is set to a constant value $\mathbf{v}$ for all states of the world. In this case the means of the prices should be proportional to the means of the asset payoffs with proportionality factor $\mathbf{v}$. The bound in (2.8) is the square root

16

of a quadratic form in the vector of deviations of the observed average prices from the average risk-neutral prices. For a fixed Σ larger deviations from risk neutral pricing imply larger bounds on the volatility of m. Shanken (1987) derived a related bound on the pricing error induced by using error-ridden proxies in computing the valuation of asset payoffs. When a constant v is used as a proxy for m, the bound in (2.8) can be viewed as special case of Shanken's bound [see Proposition 1, page 93-4].

II.C Relation to Empirical Models of Asset Prices

In this subsection we derive the relation between the mean-standard deviation frontier for m and the mean-variance frontier for asset returns. This latter frontier is the focal point of the static capital asset pricing model. The link we deduce between the two frontiers provides an alternative interpretation of the volatility bounds for m. We then describe how linear factor restrictions as imposed in Ross's (1976a) arbitrage-pricing model [see also Chamberlain (1983), Chamberlain and Rothschild (1983) and Connor (1984)] can be used to characterize the mean-standard deviation frontier for m.

Define:

$$(2.9) \qquad R \ast \{p \text{ in } P : \pi(p) = 1\}.$$

When the vector q is not random, R is the collection of (gross) returns on portfolios in P. More generally, R contains all the payoffs in P with *expected prices* that are equal to one.

Consider, first the case in which P contains a unit payoff and $\pi(1)$ is different from zero. Then $1/\pi(1)$ is in R. A second payoff in R is $r^* \ast m^*/\pi(m^*)$. Note that

$$\pi(m^{\bullet}) \;=\; E[(m^{\bullet})^2]$$

and hence

$$(2.10) \qquad \| r^{\bullet} \| \;=\; \| m^{\bullet} \| / \| m^{\bullet} \|^2 \;=\; 1 / \| m^{\bullet} \|.$$

Furthermore, Hansen and Richard (1987) established that $r^{\bullet}$ is the payoff in R that has the smallest norm (second moment). Consequently, $r^{\bullet}$ is the solution to the following optimization problem:

> minimize $\sigma(r)$ subject to $Er = \mu$
> r in R

when μ is set equal to $Er^{\bullet}$. Therefore, $m^{\bullet}$ is proportional to a particular payoff on the mean-standard deviation frontier for R.

To relate the bound for $\sigma(m)$ given in (2.4) to the slope of the *mean-standard deviation frontier* for R, note that

$$(2.11) \qquad \sigma(m)/Em \;\geq\; \sigma(m^{\bullet})/Em^{\bullet} \;=\; \sigma(r^{\bullet}) \| m^{\bullet} \|^2 / Em^{\bullet} \;=\; \sigma(r^{\bullet})/Er^{\bullet}.$$

Recall that the second moment of a random variable r satisfies:

$$E(r^2) \;=\; \sigma(r)^2 + (Er)^2 \quad.$$

Since P contains a unit payoff, the mean-standard deviation frontier for R is a cone with apex at $[0,1/\pi(1)]$ and axis parallel to the horizontal axis. In order that $r^{\bullet}$ be the minimum second moment payoff in R, the ordered pair

18

$[\sigma(r^\bullet), Er^\bullet]$ must occur at the tangency of a circle with center $(0,0)$ and the lower (inefficient) portion of the mean-standard deviation frontier for R. This tangency point is depicted in Figure 2. Since the lower portion of the frontier is a ray from $[0, 1/\pi(1)]$ through $[\sigma(r^\bullet), Er^\bullet]$, the slope of this ray is the *Sharpe Ratio* of the payoff $r^\bullet$, $[Er^\bullet - 1/\pi(1)]/\sigma(r^\bullet)$, and the slope of the circle with center $(0,0)$ that passes through $[\sigma(r^\bullet), Er^\bullet]$ is $-\sigma(r^\bullet)/Er^\bullet$. Therefore,

$$(2.12) \qquad \frac{\sigma(r^\bullet)}{Er^\bullet} = \frac{[1/\pi(1) - Er^\bullet]}{\sigma(r^\bullet)}$$

In light of (2.11) and (2.12), the bound on the ratio $\sigma(m)/Em$ is given by the absolute value of the slope of the mean-standard deviation frontier for R. These relations demonstrate the precise sense in which a steep slope of a mean-standard deviation frontier for asset payoffs can imply a potentially dramatic bound on the volatility of m.

Next we consider the case in which P does not contain a unit payoff and hence R does not contain an (unconditionally) riskless payoff. We follow the strategy used in section II.B by augmenting x with a unit payoff and assigning this payoff an expected price v. This results in an expansion of R to R_v where $1/v$ is now in R_v. Let $r_v^\bullet$ denote the payoff in R_v with the smallest second moment. Since $r_v^\bullet$ is on the mean-standard deviation frontier for R_v, it is well known from static capital asset pricing theory that $r_v^\bullet$ is a linear combination (with coefficients that sum to one) of $1/v$ and any other distinct return on the mean-standard deviation frontier for R_v. As long as $1/v$ is not equal to the mean of the minimum variance payoff in R, we can find a payoff r_v that is on the mean-standard deviation frontier for both R and R_v. Also, for each v the variable m_v is proportional to $r_v^\bullet$. Therefore, with

19

Fig. 2 — Minimum Second Moment Payoff in R

$[std(r^*), E(r^*)]$

one exception, for each random variable m_v on the mean–standard deviation for IMRS's there is a corresponding payoff r_v on the mean–standard deviation frontier for R such that m_v is a linear combination of r_v and a unit payoff. In this sense the mean–standard deviation frontier for IMRS's can be thought of as the *dual* of the mean–standard deviation frontier for R. The exceptional case occurs when $1/v$ is the mean of the minimum variance payoff in R. In this case m_v is a linear combination of a unit payoff and a payoff that is on the mean–standard deviation frontier for the space of payoffs with expected prices equal to zero.

The impact of augmenting R with $1/v$ can be seen graphically by passing a ray from the point $(0,1/v)$ through a tangent point on the mean–standard deviation frontier for R. One side of the mean–standard deviation frontier for the augmented set R_v is given by this tangent ray and the other is a reflection about a horizontal ray from $(0,1/v)$. This construction is displayed graphically in Figure 3. In the special case in which $1/v$ is the mean of the minimum variance payoff in R, it is not possible to draw a tangent line to the mean–standard deviation frontier of R from the point $(0,1/v)$. Instead the frontier for R_v is given by the two asymptotes.

Once the frontier for the augmented set R_v is obtained, the construction illustrated in Figure 2 can be mimicked using R_v in place of R. Thus for any m with mean v that satisfies RI,

$$(2.13) \qquad \frac{\sigma(m)}{Em} \; \geq \; \frac{\sigma(m_v)}{v} \; = \; \frac{\sigma(r_v^\bullet)}{Er_v^\bullet} \; = \; \frac{[1/\pi(1) - Er_v^\bullet]}{\sigma(r_v^\bullet)} \quad .$$

The relations in (2.13) show the connection between the volatility bound on m's with mean v to the slope of the mean–standard deviation frontier for R_v. A steeper slope of the frontier for R_v implies a correspondingly sharper

21

MEAN

0.980   0.988   0.996   1.004   1.012   1.020

STANDARD DEVIATION

0.005 0.010 0.015 0.020 0.025 0.030 0.035 0.040

FIG. 3 –   MEAN-STANDARD DEVIATION FRONTIERS FOR R AND R_v

22

volatility bound for m.[4]

Since the mean-standard deviation frontier for R is known to have a two-fund characterization, the preceding results show that the mean-standard deviation frontier for m can be represented using two distinct frontier payoffs in R. For the general class of asset pricing models considered in this paper, there is no prediction that particular payoffs in R, say the returns on the wealth portfolios of consumers, are mean-variance efficient. Thus without additional restrictions, there is no guidance on how to reduce a *priori* a potentially large collection of portfolio payoffs into a small collection used in a time series analysis.

One *ad hoc* approach that is often used to reduce the dimensionality of the collection of payoffs is factor analysis as employed in empirical arbitrage-pricing models [e.g. see Connor and Korajczyk (1988) and Lehmann and Modest (1988)]. Suppose that P is generated by a sequence $\{p_j\}$ where

$$(2.14) \qquad p_j = \gamma_j \cdot f + e_j$$

and f is a vector of common factors for all of the payoffs.[5] Often, the factors f are in (an appropriately defined) span of $\{p_j\}$. Hence, it follows from the Law of One Price that there exists a unique vector $\pi(f)$ of hypothetical expected prices for the factor payoffs. One possible strategy for deducing volatility bounds on m is to use the extensive collection of payoffs $\{p_j\}$ (or possibly a subset of it) to identify the first two moments of f and the expected price vector $\pi(f)$. A region S then could be constructed from these factor moments and prices using formula (2.8).

In general, information is lost in going from the larger space P to the smaller space F of linear combinations of factors. Tests of factor models of

23

asset pricing examine whether the pricing relation:

$$(2.15) \qquad \pi(p_j) = \gamma_j \cdot \pi(f)$$

holds at least approximately. When (2.15) holds exactly, the regions S generated by P and F coincide. Therefore, if asset payoffs can be priced in terms a small number of factors f, there is no loss to constructing the region S from F instead of the larger space P.[6]

As argued in Hansen and Richard (1987), an unconditional factor decomposition as in (2.14) may not be very appealing when economic agents can use conditioning information in I to make investments. If the factor decomposition (2.14) is conditioned on an information set I and γ_j is a vector of random variables in I, a reduction in payoffs is more complicated but still feasible.

24

III. IMPLICATIONS OF RESTRICTION 2

In section II we showed how to construct minimum variance random variables that satisfy restriction *R1*. These random variables may be negative with positive probability and hence may fail to satisfy *R2*.[7] As long as we limit ourselves to candidate *IMRS*'s that are translations of payoffs in *P*, it may not be possible to ensure that frontier random variables are strictly positive, or for that matter nonnegative.

In this section we initially replace *R2* by a weaker requirement that *m* be nonnegative. We then construct minimum variance candidates for *m* among the class of nonnegative random variables satisfying *R1*. It turns out that these minimum variance random variables can be interpreted as either European call or put options on payoffs in *P*. Recall that when the payoff on the underlying portfolio is *p* and the strike price is *k*, a European call option entitles an investor to the payoff $max\{p-k,0\}$ and a put option to $max\{k-p,0\}$. These payoffs are clearly nonnegative, but they may be nonlinear functions of *x*. The resulting volatility bounds for nonnegative random variables satisfying *R1* also apply when the random variables are restricted to be strictly positive (satisfy *R2*). However, in this case the lower bounds may only be approximated rather than attained.

This section is divided into three subsections. In section III.A we suppose there is a unit payoff in *P* while in section III.B we consider the more common case in which such a payoff is not included in *P*. Finally, in section III.C we discuss the close connection between our analysis and work by Harrison and Kreps (1979) and Kreps (1981) on the viability of equilibrium pricing functions consistent with the absence of arbitrage opportunities.

25

III.A Riskless Payoff

First consider the case in which there is a unit payoff in P. For each p in P, let p^+ denote $max\{p,0\}$. Note that for any p' in P and any nonnegative strike price k that is proportional to the unit payoff, the payoffs $p'-k$ and $k-p'$ are in also in P. Therefore the collection of all random variables p^+ for some p in P includes the payoffs on European call and put options with constant strike prices.

Suppose that we weaken $R2$ to the requirement that m be nonnegative. By construction, all derivative claims of the form p^+ for payoffs p in P are nonnegative.   It turns out that the minimum variance nonnegative random variable $\tilde{m}$ satisfying $R1$ is given by such a derivative claim.   Hence we are led to the problem of finding a vector α_o in $\mathbb{R}^n$ such that

$$(3.1) \qquad E[x(x'\alpha_o)^+] = Eq$$

where $\tilde{m} = (x'\alpha_o)^+$. In what follows we will first show that $\tilde{m}$, when it exists, has the smallest variance among all nonnegative random variables, m, satisfying $R1$. We then discuss the existence and computation of a solution to (3.1).

To show that $\tilde{m}$ has the smallest variance, consider any other nonnegative random variable m satisfying $R1$. Clearly

$$Exm = Ex\tilde{m} \quad .$$

Exploiting the nonnegativity of m, we have that

$$(3.2) \qquad E\tilde{m}m \geq \alpha_o'Exm = \alpha_o'Ex\tilde{m} = E[(\tilde{m})^2] \quad .$$

26

It follows from the Cauchy-Schwarz Inequality that

$$\|m\| \geq \|\tilde{m}\| \quad .$$

Since P contains a unit payoff, both m and $\tilde{m}$ must have the same mean. Therefore, we have the following relations:[8]

$$(3.3) \qquad \sigma(m) \geq \sigma(\tilde{m}) \quad \text{and} \quad Em = E\tilde{m} \quad .$$

Next we ask whether the volatility bound in (3.3) can be sharpened by requiring m be strictly positive instead of nonnegative. If $\tilde{m}$ is strictly positive (with probability one), then clearly the answer is no. This can only occur when $\tilde{m}$ coincides with $m^{\bullet}$ computed in subsection II.A. Consider the case in which $m^{\bullet}$ is not strictly positive with probability one, and let m be any random variable satisfying $R1$ and $R2$. Then $\tilde{m}$ is zero with positive probability, and it follows from (3.2) that

$$0 < \|m - \tilde{m}\|^2 = \|m\|^2 - 2E\tilde{m}m + \|\tilde{m}\|^2 \leq \|m\|^2 - \|\tilde{m}\|^2.$$

Therefore, at the very least the weak inequality ($\geq$) in (3.3) is replaced by the strong inequality ($>$). In fact no further improvements are possible. To see this form

$$(3.4) \qquad m_j = [1-(1/j)]\tilde{m} + (1/j)m \quad .$$

Then m_j is strictly positive and $\{\sigma(m_j)\}$ converges to $\sigma(\tilde{m})$. Therefore, $\sigma(\tilde{m})$

is in fact the greatest lower bound for $\sigma(m)$ when m is restricted to satisfy *R1* and *R2*.

Equation system (3.1) is nonlinear in the parameter vector α_o, and its solution cannot necessarily be represented in terms of matrix manipulations. There is a a closely related optimization problem whose solution may be easier to compute. This problem entails finding a payoff in R whose truncation has the smallest second moment:

$$(3.5) \qquad \min_{r \ in \ R} \ \|r^+\|^2 \quad .$$

In Appendix A we show that (3.5) has a solution, although this solution may not be unique. Furthermore, a necessary and sufficient condition for $\tilde{r}$ to be a solution to (3.5) is

$$(3.6) \qquad E(\tilde{r}^+ z) = 0 \qquad \text{for all } z \text{ in } P \text{ such that } \pi(z) = 0.$$

We can think of (3.5) as being the first-order condition for optimization problem (3.5).

It turns out that we can construct a solution to (3.1) by scaling $\tilde{r}^+$ appropriately. Let

$$(3.7) \qquad \tilde{m} = \tilde{r}^+ / \|\tilde{r}^+\|^2 .$$

This scaling is permissible because $\|\tilde{r}^+\|$ must be strictly positive as long as there exists at least one random variable m satisfying *R1* and *R2*. To see this, suppose to the contrary that $\|\tilde{r}^+\|$ is zero. Then $-\tilde{r}$ is a nonnegative payoff with a strictly negative expected price. Such a payoff is

28

inconsistent with *R2* because it implies that there exists an arbitrage opportunity.

Clearly $\tilde{m}$ as given by (3.7) can be represented as $(\alpha_o{}'x)^+$ for some α_o in $\mathbb{R}^n$. To verify that α_o solves (3.1), we must show that $\tilde{m}$ as given by (3.7) satisfies *R1*. Let p be any payoff in P, and form the payoff:

$$z = p - \pi(p)\tilde{r}.$$

Note that $\pi(z) = 0$ because $\pi(\tilde{r}) = 1$. It follows from first-order condition (3.6) that

$$
\begin{aligned}
0 \;=\; E\tilde{m}z \;&=\; E\tilde{m}p - \pi(p)E\tilde{m}\tilde{r}\\
&=\; E\tilde{m}p - \pi(p)E[\tilde{r}^+\tilde{r}]/\|\tilde{r}^+\|^2\\
&=\; E\tilde{m}p - \pi(p) \quad.
\end{aligned}
$$

Thus $\tilde{m}$ satisfies *R1* as required.

This construction of $\tilde{m}$ parallels a similar construction reported in Hansen and Richard (1987) and in section II. Ignoring *R2*, one way to construct the random variable $m^\bullet$ which has minimum variance among the class of random variables satisfying *R1* is to compute the minimum second moment payoff, $r^\bullet$, in R and divide it by its second moment, $\|r^\bullet\|^2$. We have just demonstrated that a similar strategy works for constructing a random variable $\tilde{m}$ that attains the volatility bound among the class of nonnegative random variables satisfying *R1*. Instead of computing the minimum second moment payoff in R, we calculate the minimum *truncated* second moment payoff, $\tilde{r}$, in R. To form $\tilde{m}$, the *truncation* of this payoff, $\tilde{r}^+$, is divided by the second moment of its truncation $\|\tilde{r}^+\|^2$. Whereas $\|m^\bullet\|$ is given by $1/\|r^\bullet\|$, $\|\tilde{m}\|$ is

29

given by $1/\|\tilde{r}^+\|$. Since truncating a random variable reduces its norm, as required $\tilde{m}$ has a larger second moment than m^*. The difference in the two norms reflects the incremental contribution of restriction $R2$ for the volatility bound on m.

One advantage to solving optimization problem (3.5) instead of solving directly the nonlinear equation system (3.1) is that optimization problem (3.5) has a convex objective function $\|r^+\|^2$ and a convex constraint set R so that numerical solutions are quite feasible to obtain. Although $\tilde{r}$ is not necessarily unique, its truncation $\tilde{r}^+$ is (see Appendix A). A sufficient condition for $\tilde{r}$ to be unique, which is often satisfied in practice, is that no two payoffs in R have the same truncation.

III.B No Riskless Payoff

Consider the more common case in which P does not contain a unit payoff. As in section II.B augment x with a unit payoff and form an augmented payoff space P^a. Similarly, assign alternative strictly positive numbers v for $\pi(1)$ and extend π from P to P^a. Let R_v be the augmented set of payoffs with expected prices equal to one when $\pi(1)$ is assigned v. The counterpart to equation (3.1) is not guaranteed to have a solution, however. It turns out that there additional limits on the admissible choices of v consistent with $R2$.

To investigate these limits, we study the counterparts to optimization problem (3.5) using the augmented space of payoffs R_v in place of R. Define:

$$(3.8) \qquad \delta_v \; \equiv \; \inf_{r \text{ in } R_v} \|r^+\|^2 \; .$$

When δ_v is positive, the bound on $\|m\|^2$ among the class of nonnegative random

variables satisfying *R1* and *R2* with mean v is $1/\delta_v$.   However, particular choices of v may result in δ_v being zero and hence $1/\delta_v$ being infinite.   For instance, when there is a portfolio payoff *p* in *P* such that *p* is less than or equal to one with probability one and v is strictly less than $\pi(p)$, δ_v is zero.   This is true because the random variable $(1-p)/[v-\pi(p)]$ is in R_v and is less than or equal to zero with probability one.   Consequently, the norm of its truncation is zero.

As noted by Merton (1973), Cox, Ross and Rubinstein (1979), Harrison and Kreps (1979) and Kreps (1981), it possible to obtain arbitrage bounds on the admissible (expected) prices that can be assigned to payoffs not in *P*.   In the case of a unit payoff, the upper and lower bounds are given by

$$\bar{\pi}(1) \equiv \begin{cases} inf\{\pi(p) \,:\, p \geq 1\} & \text{if } \{p \text{ in } P \,:\, p \geq 1\} \text{ is not empty} \\ +\infty & \text{otherwise} \end{cases}$$

$$\underline{\pi}(1) \equiv sup\{\pi(p) \,:\, p \leq 1\}$$

respectively.[9]   Since the zero payoff is in *P*, $\underline{\pi}(1)$ is always and nonnegative.   The arbitrage bounds $\underline{\pi}(1)$ and $\bar{\pi}(1)$ determine the range of admissible values of *Em* that are compatible with *m* being a nonnegative random variable.   Clearly, $\{v \,:\, \delta_v > 0\}$ must be a subset of the interval $[\underline{\pi}(1), \bar{\pi}(1)]$. In fact, the interiors of these sets coincide (see Appendix A).

When δ_v is strictly positive, there exists a minimum variance, nonnegative random variable with mean v that satisfies *R1* (see Lemma A.4 in the Appendix A).   Let this random variable be denoted $\tilde{m}_v$.   The corresponding volatility bound is:

$$(3.9) \qquad \sigma(m) \geq \sigma(\tilde{m}_v),$$

31

and the family of random variables, $\{\widetilde{m}_{\mathbf{v}} : \delta_{\mathbf{v}} > 0\}$, comprise the mean–standard deviation frontier for nonnegative random variables satisfying $R1$. Thus the counterpart to the region S given in (2.5) is

(3.10) $S^+ \equiv \{(\mathbf{v}, \mathbf{w}) : \delta_{\mathbf{v}} > 0 \text{ and } \mathbf{w} \geq \sigma(\widetilde{m}_{\mathbf{v}})\}$.

The set S^+ is convex. To see this consider two values of $\mathbf{v}$ for which $\delta_{\mathbf{v}}$ is strictly positive, say $\mathbf{v}(\ell) \leq \mathbf{v}(u)$. Form convex combinations of the random variables $\widetilde{m}_{\mathbf{v}(\ell)}$ and $\widetilde{m}_{\mathbf{v}(u)}$. These convex combinations are nonnegative random variables that also satisfy $R1$. Recall that the mean of a convex combination of random variables is equal to the convex combination of the means, and by the Triangle Inequality, the standard deviation of a convex combination is less than or equal to the convex combination of the standard deviations. While convex combinations of $\widetilde{m}_{\mathbf{v}(\ell)}$ and $\widetilde{m}_{\mathbf{v}(u)}$ are not necessarily on mean–standard deviation frontier, the ordered pairs of their means and standard deviations must be in S^+. This is sufficient for S^+ to be convex.

Next we consider the incremental contribution of requiring that m be strictly positive as in $R2$. It is shown in Appendix A that Em must be in the open interval $(\underline{\pi}(1), \bar{\pi}(1))$ (see Lemma A.6). Hence one effect of the imposition of strict positivity is that end points of the interval $\{\mathbf{v} : \delta_{\mathbf{v}} > 0\}$ are eliminated.

For $\mathbf{v}$'s in the $(\underline{\pi}(1), \bar{\pi}(1))$, $\widetilde{m}_{\mathbf{v}}$ can be interpreted as either a call or put option on a payoff in P. More precisely, we obtain the counterpart to the result in section III.A that

$$\widetilde{m}_{\mathbf{v}} = (\widetilde{p}_{\mathbf{v}} - k)^+$$

32

where $\tilde{p}_v$ is a portfolio payoff in P and k is in $\mathbb{R}$. When k is nonnegative, $\tilde{m}_v$ is a call option on a portfolio with payoff $\tilde{p}_v$ and strike price k, and when k is negative $\tilde{m}_v$ is a put option on a portfolio with payoff $-\tilde{p}_v$ and strike price $-k$. Therefore for any v in $(\underline{\pi}(1), \bar{\pi}(1))$ the counterpart to equation (3.1) has a solution.

As in section III.A, this gives us a simple check of the incremental impact of positivity on the volatility bounds given in (3.9). For any v in $(\underline{\pi}(1), \bar{\pi}(1))$ such that $\tilde{m}_v$ is strictly positive, the bound in (3.9) cannot be improved by restricting m to be strictly positive. This can only occur when $\tilde{m}_v$ coincides with m_v calculated in section III.B. On other hand, for any v for which m_v is not strictly positive, the weak inequality ($\geq$) in (3.9) is replaced by a strong inequality ($>$).

Even though S^+ may be a proper subset of S, the region S is still of interest for a variety of reasons. First, S is easier to use in practice because a characterization of S^+ may require that a nonquadratic optimization problem be solved for each value of v. Second, the lower boundaries of S^+ and S coincide for values of v for which m_v is nonnegative. Consequently, it is advantageous to characterize S as a first step in characterizing S^+ and then check for nonnegativity of m_v. Finally, even for values of v for which m_v is negative with positive probability, the coefficients on x^a given in representation (2.7), when scaled appropriately, can be used as starting values for a numerical search routine used in computing δ_v.

In Figure 4 we report plots of the regions S and S^+ for the same financial data set as was used to generate Figure 1. The region S^+ is shaded, and the lower boundary of the region S is given by the dashed line below S^+. While the lower boundaries of these regions coincide for points closest to the the horizontal axis, they diverge for other points. The

FIG. 4 – IMRS FRONTIER WITH AND WITHOUT POSITIVITY IMPOSED

divergence between the boundaries is greater when the volatility bounds are more restrictive. Recall that in generating the lower boundary of S, we constructed random variables $m_\mathbf{v}$ with mean $\mathbf{v}$ that satisfy $R1$ and are linear combinations of $x^\bullet$. When these random variables have large standard deviations relative to their means, it is not surprising that they are negative with high probability. As a result, the positivity restriction ($R2$) often has more bite when $\sigma(m_\mathbf{v})/\mathbf{v}$ is larger.

As is true for S, the dimensionality of P can sometimes be reduced prior to the construction of S^+. Suppose that members of $\dot P$ have factor decompositions of the form:

$$p = \gamma \cdot f + e$$

where f is a vector of common factors. Suppose further that the idiosyncratic components of the payoffs satisfy $E(e|f) = 0$ and $\pi(e) = 0$. Hence we have exact factor pricing and each payoff p in P is a mean-preserving spread of a payoff $\gamma \cdot f$ with the same price. Consequently,

$$\|p^+\|^2 \geq \|(\gamma \cdot f)^+\|^2$$

because the function $[(p)^+]^2$ of p is convex. Hence in solving (3.5) or (3.8) it suffices to restrict attention to linear combinations of the factors with expected prices equal to one. While it is evident how to use this reduction when the factors are observed, unobserved factors are problematic because it may be difficult to compute or estimate $\|(\gamma \cdot f)^+\|^2$ for arbitrary vectors γ. Due to the truncation of $(\gamma \cdot f)$, calculating $\|(\gamma \cdot f)^+\|$ requires knowledge of the entire probability distribution of f, whereas typical factor analytic

procedures identify only the first two moments of f.

III.C Viability of Equilibrium Pricing Functions and Arbitrage Pricing

The analysis in this section is intimately connected to general treatments of pricing derivative claims [e.g. see Ross (1978), Harrison and Kreps (1979) and Kreps (1981)].   Among other things, Harrison and Kreps (1978, 1979) and Kreps (1981) consider the following question.   Given a set of payoffs on primitive securities and the prices of those securities, when is it possible to extend the pricing function to a larger collection of payoffs in such a way as to preserve no-arbitrage? As emphasized by Kreps (1981), this experiment should not be construed as introducing new markets in an economy that might alter the resulting competitive equilibrium allocations.   It is merely a hypothetical extension leaving intact the (expected) prices of the payoffs in P.   When such an extension is possible, Harrison and Kreps (1979) and Kreps (1981) refer to the pricing function as being *viable*.

Throughout the analysis in this section, we presumed that the family of m's that satisfying $R1$ and $R2$ is not empty.   Clearly this is sufficient to eliminate arbitrage opportunities on P.   Rather than assuming that this family is not empty, an alternative starting point is to verify that no arbitrage opportunities exist on P, and then to appeal to Theorem 3 in Kreps (1981) to show that π can be extended from P to the collection L^2 of all random variables that are (Borel measurable) functions of x and have finite second moments.[10]   The existence of an m satisfying $R1$ and $R2$ then follows from the Riesz Representation Theorem applied to L^2. [See also Lemma 2.3 in Hansen and Richard (1987)].

IV. Illustrations and Discussion

We now illustrate our analysis with alternative parametric models of m and alternative data sets on asset payoffs and prices. The model of m described in the introduction and used to generate Figure 1 assumed that consumers' preferences are separable over time and states of the world. In section IV.A we investigate the impact on m of relaxing time separability. In subsection IV.B we focus on logarithmic risk preferences but do not require that these preferences be state separable. Finally in section IV.C, we describe the implications of price data on short term Treasury bills for *IMRS*'s and comment briefly on the implications for monetary models.

IV.A: Preferences that are not Time Separable

Consider the following stylized version of a model with time nonseparabilties in preferences. As in the introduction, we use a time- and state-separable specification of preferences for *consumption services* with a power utility function:

$$(4.1) \qquad E \sum_{t=0}^{\infty} \lambda^t \frac{s_t^{\gamma+1} - 1}{\gamma+1}$$

except now s_t depends on measured consumption in the current period and one previous period:

$$(4.2) \qquad s_t = c_t + \theta c_{t-1}.$$

More general versions of this model have been investigated by Dunn and Singleton (1986), Eichenbaum, Hansen and Singleton (1988), Gallant and

37

Tauchen (1989) and Eichenbaum and Hansen (1990). We will proceed as if there is a single representative consumer. As noted by Wilson (1968) and Rubinstein (1974), this assumption can be relaxed when θ is zero, the consumption allocations are consistent with the existence of complete contingent claims markets and all consumers have the same preferences. This aggregation result also applies more generally, say when θ is different from zero, as long as there are, in effect, complete markets in consumption services. [See Eichenbaum, Hansen and Richard (1987)]. When θ is positive, consumption generates positive services in the current as well as in one subsequent time period. In this case there is *intertemporal substitution* in generating consumption services from consumption goods. More precisely, there is a durable component to consumption that depreciates fully after one time period. Alternatively when θ is negative, there is *intertemporal complementarity* in generating consumption services from consumption goods. Put somewhat differently, the term $-\theta c_{t-1}$ is a component of current period consumption which reflects either *committed consumption* from the previous time period or *habit persistence*. Sundarasen (1989), Novales (1990) and Constantinides (1988) have argued that habit persistence may be important in explaining the relation between asset market data and economic aggregates.

For these forms of time nonseparabilities, the marginal utility of consumption is

$$mu_\tau = (c_\tau)^\gamma + \lambda\theta E[(c_{\tau+1})^\gamma | I_\tau] \quad .$$

The *IMRS* between time zero and time τ is the corresponding ratio of marginal utilities scaled by λ^τ. Constructing m requires computation of the conditional expectation $E[(c_{\tau+1})^\gamma | I_\tau]$ except in the special case in which θ

38

is zero.

To illustrate what impact positive and negative values of θ have for the volatility of m, we report calculations from Gallant, Hansen and Tauchen (1989). For these calculations the ratio (c_t/c_{t-1}) is a component of a Markov process with a stochastic law of motion estimated by Gallant and Tauchen (1989) for monthly data on the consumption of nondurables and services [for more details see Gallant and Tauchen (1989) and Gallant, Hansen and Tauchen (1989)]. The estimated law of motion was then used to compute $E[(c_{\tau+1})^{\gamma}|I_{\tau}]$ required in forming a time series for m. Sample means and standard deviations were calculated for m's implied by alternative values of γ and θ.[11]

For this illustration we let $\theta = -.5$, $\theta = 0$ and $\theta = .5$. The results are reported in Figure 5. The "■"'s are used to denote mean-standard deviation pairs for $\theta = 0$, the "Δ"'s for $\theta = .5$ and the "o"'s for $\theta = -.5$. For each choice of θ, we let γ range from 0 to -14 with decrements of minus one. In all cases the subjective discount factor λ is set to one. Smaller values of λ decrease proportionately the mean and standard deviation of m. When $\gamma = 0$, m is one for all choices of θ. In this case $[E(m),\sigma(m)] = (1,0)$.

Consider first the case in which $\theta = 0$. Increasing $|\gamma|$ magnifies the volatility of m but initially reduces its mean. Extrapolated much further, the curve (indexed by γ) does not turn around until $|\gamma|$ is in the vicinity of one hundred, after which increasing $|\gamma|$ enlarges the mean of m. The initial decline in the mean of m reflects the dominant role of positive growth rates in consumption. For extremely large values of $|\gamma|$, observations with negative growth rates in consumption come to dominate the sample mean eventually resulting in a change of slope of the curve. In comparing the curves denoted by "■"'s in Figures 1 and 5, recall that the long annual time

39

FIG. 5 - IMRS FRONTIER COMPUTED USING MONTHLY DATA

MEAN

STANDARD DEVIATION

series used to generate Figure 1 contains negative growth rate observations on consumption during the depression. The absence of *bad* events in the monthly data set is responsible for the fact that the curve (indexed by γ) does not turn until magnitude of γ is substantial.[12]

Consider next the case in which $\theta = .5$. Not surprisingly, introducing this local durability into preferences reduces the volatility of m. The quantitative effect of this smoothing does not appear to be very substantial, however. The curves for $\theta = .5$ and $\theta = 0$ are similar for the range of γ's that are plotted. Hence there is little adverse effect on the volatility of m to introducing durability by setting $\theta = .5$.

Finally, consider the case in which $\theta = -.5$. This intertemporal complementarity has the anticipated impact of increasing the volatility of m for a given value of γ. This effect is quite dramatic as indicated in Figure 5. Furthermore, the value of $|\gamma|$ at which the curve turns is reduced dramatically. For $\theta = -.5$ the turning point for $|\gamma|$ is in the vicinity of -7, and the initial decline in the mean of m is much less dramatic.

We now compare the three curves, which describe alternative mean–standard deviation pairs for parametric models of m, to a region S^* generated using monthly data on asset payoffs and prices. The asset market data are the same as were used by Hansen and Singleton (1982) except that data revisions were incorporated and more recent data points were included. The resulting time period is from 1959:3 – 1986:12. The first two asset payoffs are the one-month real return on Treasury bills and the one-month real value-weighted return on the New York Stock Exchange. Six additional time series of asset payoffs were constructed using these data by scaling the original two payoffs and prices by the one-period lagged returns and the one period lag in the consumption ratio. For the the range of hypothetical means

considered, the region S described in section II was essentially the same as the region S^+ described in section III.

For the specification of preferences with $\theta = 0$, larger values of $|\gamma|$ initially make the mean–standard deviation pair for m further from S^+ region because of the adverse effect on the mean of m. This is consistent with the fact that Hansen and Singleton (1982) found point estimates for γ that were close to zero but substantial evidence against the over-identifying restrictions. As emphasized by Singleton (1988), estimates of the discount factor λ are often greater than one when bond returns are included in the analysis. For a fixed γ, enlarging λ has the desired effect of increasing proportionately the mean and standard deviation of m.

From the vantage point of Figure 5, the case for intertemporal complementarities in preferences is appealing. For a given value of γ, a negative value of θ increases both the mean and the standard deviation of m. However, it is quite possible for m to have a mean and standard deviation in S^+ and not satisfy $R1$. In other words, for a given parametric specification of m, requiring $[E(m),\sigma(m)]$ be in S^+ does not exhaust the testable implication of $R1$. As emphasized by Gallant, Hansen and Tauchen (1989), there is substantial statistical evidence that the resulting m's violate $R1$. In fact empirical studies that use similar data and preference specifications, such as Dunn and Singleton (1986), Eichenbaum, Hansen and Singleton (1988), and Eichenbaum and Hansen (1990), typically find parameter estimates that reflect intertemporal substitution ($\theta > 0$) although they ｀ ｡d statistical evidence against the resulting parametric model of m.

IV.B: Logarithmic Risk Preferences

In (4.1) and (4.2) suppose that θ is zero and γ is minus one. In this case preferences are logarithmic. As noted by Rubinstein (1976), m is equal to the reciprocal of the return on the wealth portfolio of the representative consumer between time zero and time τ. [See also Brown and Gibbons (1985)]. Epstein and Zin (1989) showed that this same conclusion applies to a parametric class of recursive preferences that are not state separable so long as the risk preferences remain logarithmic. Whereas in the state-separable case, the return on the wealth portfolio is equal to the discounted consumption ratio, this exact relation no longer applies when state separability is relaxed. Nevertheless, the return on the wealth portfolio can still be used as a valid measure of m.

For this reason we have included an "x" in Figures 1 and 5. In the case of Figure 1 this "x" denotes the sample mean-standard deviation pair for the reciprocal of the measured annual return on the Standard and Poors 500 stock price index and in Figure 5 it represents the sample mean-standard deviation pair for the reciprocal of measured monthly value-weighted return on the New York Stock Exchange. In both cases the means are near the points in S^+ that are closest to the horizontal axis, but the mean-standard deviation pair is outside of S^+. However, after taking of sampling error, there is very little evidence against the null hypothesis that this model of m satisfies $R1$.[13]

IV.C: Treasury Bill Data and Monetary Models

We also calculated the regions S and S^+ using monthly data on three-month holding period returns on Treasury bills. The holding period returns were constructed using bond prices on three, six, nine and twelve month discount bonds from 1964:7 - 1986:12. Nominal returns were converted

43

to real returns using the implicit price deflator on nondurables and services. These bond price data (excluding the most recent periods) have been used by Fama (1984), Dunn and Singleton (1986) and Stambaugh (1988), among others, to investigate time variation in risk premia and particular models of bond prices. In Figure 6 we report the resulting regions S and S^*. The region S^* is shaded and the lower boundary of the region S is given by the dashed line below S^*. The resulting standard deviation bounds for m are quite striking. For means of m in the vicinity of one, the bound on the standard deviation is near one. Given the magnitude of these bounds, it is not suprising that $R2$ has an important incremental contribution *vis-a-vis R1*.

The bounds reported in Figure 6 appear to us to pose quite a challenge to a large class of asset valuation models. For instance, the quarterly counterpart to the $\theta = 0$ curve in Figure 5 ranges from $(1,0)$ to $(.90,.08)$ as γ ranges from 0 to -14. Volatility bounds of a similar magnitude were also obtained using monthly data on one-month holding period returns for Treasury bills with maturities from one to six months. These bounds are directly comparable to the three curves plotted in Figure 5. Since these latter bounds apply to *IMRS*'s measured over a shorter time period (one month instead of three months), they are even more startling. However, short term Treasury bills are often held to maturity and trading of these Treasury bills takes place in secondary markets except for the three, six and twelve month bills. The bid-ask spreads for the short-term bills can be quite substantial [see Stambaugh (1988) and Knez, Litterman and Scheinkman (1989)] so that the prices used in our calculations may be less reliable. These concerns should be less problematic for the results in Figure 6 since they were computed using price data from the more richly traded three, six, nine and twelve month Treasury bills.

44

STANDARD DEVIATION

FIG. 6 – IMRS FRONTIER COMPUTED USING QUARTERLY HOLDING-PERIOD RETURNS

MEAN

with positivity
without positivity

As emphasized by Knez, Litterman and Scheinkman (1989), short term Treasury bills often may be held to maturity as cash substitutes for particular transactions. As such, these bills may generate important liquidity services that are not measured appropriately by the implied *ex post* real returns. Hence the measured real returns may understate the value of the assets to the holders of the securities. Recall from section II.C, that large standard deviation bounds for m occur when the slope of the mean-standard deviation frontier for R is steep. For the Treasury bill data, this means that the reason that the volatility bounds on m are large (as reflected in the region S) is that the expected short-term gain associated with holding longer term bills is large relative to the increase in the standard deviation. Abstracting from the liquidity services of the short term bills may distort the magnitude of the resulting volatility bounds on m.

Refinements of real asset pricing models to incorporate money, such as the cash-in-advance models of Lucas and Stokey (1987), Swennson (1985) and Townsend (1987) are designed to accommodate the rate of return dominance between one-period bonds and money. However, in their current form they are not well suited to differentiate among short term Treasury bills with different maturity dates. Although the link between measured *real IMRS*'s and security market data may be confounded in these models, there is an alternative notion of the *indirect IMRS* for money which reflects the fact that the cash-in-advance constraint may not always be binding. Hence for an appropriate interpretation of m, these monetary models are compatible with *R1* and *R2* as long as money is not included among the vector of assets used to generate S^{+}.

Conclusions and Extensions

In this paper we have characterized the implications of security market data for means and standard deviations of *IMRS*'s. This exercise is important in evaluating alternative models of dynamic economies because *IMRS*'s are the channels by which the attributes of these models impinge on asset prices. Abstracting from the restriction that *IMRS*'s are positive, we established the connection between volatility bounds on *IMRS*'s and mean–standard deviation frontiers for asset payoffs. Thus we showed how diagnostics commonly used in empirical finance can be translated into information about *IMRS*'s. We also showed how to extract sharper volatility bounds by taking account of the fact that *IMRS*'s should be positive. These sharper bounds exploit more fully the absence of arbitrage opportunities in the underlying economic environment than, say, linear factor representations of asset prices.

There are three important directions in which the ideas in this paper can be developed further. An earlier version of this paper has already provoked some work along these three lines.

i) In this paper we focused exclusively on deriving implications for *IMRS*'s expressed in terms of population moments of asset payoffs and prices. In practice, these attributes of asset market data will not be known a *priori*, but only can be approximated by using time series averages in place of population moments. This introduces sampling error into the analysis. A major drawback in the discussion in Section 4 is that it abstracted from the presence approximation error introduced by using sample averages from historical time series in place of population moments. Hansen and Jagannathan (1990) shows how to use large sample theory to both assess whether there is sufficient statistical evidence to reject that the bounds are degenerate (equal to zero) and to assess the magnitude of the approximation errors.

ii)   The restrictions on *IMRS*'s derived in this paper all pertain the first and second moments.   More generally, it would be desirable to characterize the admissible family of distributions for *IMRS*'s given asset market data.   An additional step towards such a characterization is taken by Snow (1990), who shows how to extend the analysis in this paper to obtain bounds on other moments of the *IMRS*'s.

iii)   The diagnostics derived in this paper   can be applied to any intertemporal asset pricing model for which moments of m can be computed. While the calculations in section IV were performed by first constructing hypothetical time series on m, such a construction is not necessary.   All that is really essential is the ability to compute the moments of m implied by the model.   As an alternative to constructing a time series on m, these moments can be deduced from the equilibrium stochastic law of motion for the model [e.g. see Heaton (1990)].   Therefore, calculations like those illustrated in section IV can be performed for an extensive array of intertemporal asset pricing models including models that take account of measurement errors in consumption, seasonality and aggregation-over-time biases.

48

1.   In the case of incomplete markets, formula (1.2) abstracts from the possible existence of short sale constraints.

2.   Our use of the unconditional expectation operator in this context is justified formally when the time series converges appropriately to a stochastic steady state and is ergodic.   In this case unconditional expectations are computed using the stationary distribution.   For processes that are asymptotically stationary but not ergodic, the limit points can often be represented as conditional expectations where the conditioning is on the invariant sets for the approximating stationary stochastic process.

3.   Notice that a larger set P of portfolio payoffs could be constructed by following more complex trading strategies in which the vector c is replaced by a vector of random variables in I.   The theoretical analysis in Hansen and Richard (1987) is designed to accommodate this case as is the econometric analysis in Gallant, Hansen and Tauchen (1989).   We focus on the linear span of x for pedagogical convenience and empirical tractability.

4.   Using conditioning information in clever ways can sharpen the volatility bounds on m by increasing the maximum Sharpe Ratio of the payoffs in R_v.   For example Breen, Glosten and Jagannathan (1989) show that information in Treasury bill returns can be used to construct a portfolio which has the same average return as the value-weighted index of New York Stock Exchange securities but is only half as variable.

5.   Although our derivation of the volatility bounds for m assumed that the payoff space P is finite-dimenionsal, this restriction was made for pedagogical convenience.   In fact the duality relation between the mean-standard deviation frontiers for m's that satisfy restriction $R1$ and for payoffs in R extends to environments in which P is generated by an infinite number of payoffs, say by $\{p_j\}$.

6.   In contrast to factor analytic approaches, Huberman and Kandel (1987) test whether the dimensionality of P can be reduced to a prespecified observed subset of security returns, namely three size-based portfolios of New York Stock Exchange Securities.   In this case F can be constructed using these three returns.   Huberman and Kandel find, however, that this construction of F is not adequate to span the mean-standard deviation frontier for the original P constructed using thirty three size-sorted portfolios.   Hence in this case the dimensionality reduction from P to F will result in weaker implications for m.

7.   As Dybvig and Ingersoll (1982) have pointed out, naive use of m^* to compute (expected) prices of contingent claims may lead to assignment of negative (expected) prices to some positive payoffs and hence to the appearance of an arbitrage opportunity.

8.   An alternative way to deduce these bounds it to exploit the fact that when payoffs on calls and puts are included in the analysis, the space of admissible payoffs is essentially complete [see Ross (1976b), Breeden and Litzenberger (1978), Arditti and John (1980) and Green and Jarrow (1987)]. If the prices of all such payoffs were available, the counterpart to the random variable m in section II would be strictly positive.   Although this extensive collection of option price data is typically not available, we can follow Merton (1973) and use lower bounds on option prices to obtain a lower bound on the volatility of m.

9.   The characterization reported in Harrison and Kreps (1979) and (1981) is somewhat more complicated because they allow the counterpart to the space P to be infinite dimensional.

10.   In addition to a *no-arbitrage* restriction, Kreps (1981) also imposed a *no-free-lunch* restriction on (P,π).   The extra restriction is required because the counterpart to P in Kreps' analysis is allowed to be infinite dimensional.

11.   Note that the calculations of mean–standard deviation pairs for m when θ = 0 do not exploit the Markov specification estimated by Gallant and Tauchen (1989) and are consequently more robust.

12.   The *sample* volatility of m may be substantially lower than the *population* volatility if consumers anticipate that extremely bad events can occur with small probability when such events do not occur in the sample. Reitz (1988) argued that this phenomenon could explain the *equity premium puzzle*.

13.   In the case of Figure 1, one of the two moment conditions $E(mx-q) = 0$ is satisfied by construction.   The other condition was tested using the method suggested in Hansen and Singleton (1982):   the $\chi^2(1)$ statistic is 1.40 with probability value 0.24.   Similarly, for Figure 5, four of the eight moment conditions are satisfied by construction.   The $\chi^2(4)$ statistic for the four remaining conditions is 4.88 with probability value of 0.30.

REFERENCES

Arditti, F. D., and K. John (1980), "Spanning the State Space with Options," *Journal of Financial and Quantitative Analysis*, 50, 1-9.

Breeden, D. T. (1979), "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265-296

Breeden, D. T., and R. H. Litzenberger (1978), "Prices of State-Contingent Claims Implicit in Options Prices," *Journal of Business*, 51, 621-651.

Breen, W., L. Glosten, and R. Jagannathan (1989), "Economic Significance of Predictable Variations in Stock Index Returns," *Journal of Finance*, 44, 1177-1189.

Brown, D.P., and M.R. Gibbons (1985), "A Simple Econometric Approach for Utility-Based Asset Pricing Models," *Journal of Finance*, 11, 359-381.

Campbell, J. and R. Shiller (1988), "Dividend-Price Ratios and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195-228

Chamberlain, G. (1983), "Funds, Factors, and Diversification in Arbitrage Pricing Models," *Econometrica*, 51, 1305-1323.

Chamberlain, G., and M. Rothschild (1983), "Arbitrage, Factor Structure, and Mean-Variance Analysis of Large Asset Markets, *Econometrica*, 51, 1281-1304.

Connor, G. (1984), "A Unified Beta Pricing Model," *Journal of Economic Theory*, 34, 13-41.

Connor, G., and R. A. Korajczyk (1988), "Risk and Return in an Equilibrium APT: Application of a New Test Methodology," *Journal of Financial Economics*, 21, 255-289.

Constantinides, G. (1988), "Habit Formation: A Resolution of the Equity Premium Puzzle," CRSP Working Paper 225, Graduate School of Business, University of Chicago.

Cox, J.C., S.A. Ross and M. Rubinstein (1979), "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7, 229-263

Dunn, K., and K.J. Singleton (1986), "Modeling the Term Structure of Interest Rates Under Nonseparable Utility and Durability of Goods," *Journal of Financial Economics*, 41, 333-355.

Dybvig, P.H., and J.E. Ingersoll (1982), "Mean Variance Theory in Complete Markets," *Journal of Business*, 55, 233-251.

Eichenbaum, M.S., and L.P. Hansen (1990), Estimating Models with Intertemporal Substitution using Aggregate Time Series Data, *Journal of Business and Economic Statistics*, 8, 53-69.

Eichenbaum, M.S., L.P. Hansen and S.F. Richard (1987), "Aggregation, Durable Goods and Nonseparable Preferences in an Equilibrium Asset Pricing Model," Program in Quantitative Analysis Working Paper 87-9, NORC, Chicago.

Eichenbaum, M.S., L.P. Hansen and K.J. Singleton (1988), "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice Under Uncertainty," *The Quarterly Journal of Economics*, 103, 51-78.

Epstein, L., and S. Zin (1989), "Substitution, Risk Aversion, and Temporal Behavior of Consumption and Asset Returns II: An Empirical Analysis," manuscript, University of Toronto.

Fama, E.F. (1984), "The Information in Term Structure," *Journal of Financial Economics*, 13, 509-528.

Gallant, R., L.P. Hansen, and G. Tauchen (1989), "Using Conditional Moments of Asset Payoffs to infer the Volatility of Intertemporal Marginal Rates of Substitution", forthcoming in *Journal of Econometrics*.

Gallant, R. and G. Tauchen (1989), "Seminonparametric Estimation of Conditionally Constrained Heterogeneous Processes: Asset Pricing Applications," *Econometrica*, 57, 1091-1120.

Green, R.C. and R.A. Jarrow (1987), "Spanning and Completeness in Markets with Contingent Claims," *Journal of Economic Theory*, 202-210.

Hansen, L.P. (1982), "Consumption, Asset Markets and Macroeconomic Fluctuation: A Comment," *Carnegie-Rochester Conference Series on Public Policy* 17, 239-250.

Hansen, L.P. and R. Jagannathan (1990), "Econometric Methods For Analyzing Single Factor Models of Asset Prices," Manuscript, University of Chicago.

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, 1987, 587-613.

Hansen, L.P., and K.J. Singleton (1982), "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269-1286.

Harrison, J.M. and D.M. Kreps (1979), "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381-408.

Heaton, J. (1990), "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," manuscript.

Huberman, G., and S. Kandel (1987), "Mean-Variance Spanning," *Journal of Finance*, 42, 873-888.

Knez, P., R. Litterman and J. Scheinkman (1989), "Explorations into Factors Explaining Money Market Returns," Financial Strategies Group, Goldman, Sachs and Company.

Kreps, D. (1981), "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities," *Journal of Mathematical Economics*, 8, 15-35.

Lehmann, B. N. and D. M. Modest (1988), "The Empirical Foundations of Arbitrage Pricing Theory," *Journal of Financial Economics*, 21, 213-254.

LeRoy, S. F. (1973), "Risk Aversion and the Martingale Property of Stock Prices," *International Economic Review*, 14, 436-446.

Lucas, R.E., Jr. (1978), "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1978, 1429-1445.

Lucas, R.E., and N.L. Stokey (1987), "Money and Interest in a Cash-in-Advance Economy," *Econometrica*, 55, 491-513.

Mehra, R., and E.C. Prescott (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145-161, 1985.

Merton, R.C. (1973), "Theory of Rational Options Pricing," *Bell Journal of Economics and Management Science*, 4, 141-183.

Novales, A. (1990) "Solving Nonlinear Rational Expectations Models: A Stochastic Equilibrium Model of Interest Rates," *Econometrica*, 58, 93-112.

Reitz, T.A. (1988), "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, 117-131.

Ross, S.A. (1976a), "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 13, 341-360.

Ross, S.A. (1976b), "Options and Efficiency," *Quarterly Journal of Economics*, 90, 75-89.

Ross, S.A. (1978), "A Simple Approach to the Valuation of Risky Streams," *Journal of Business*, 3, 453-476.

Rubenstein, M. (1974), "An Aggregation Theorem for Securities Markets," *Journal of Financial Economics*, 1, 225-244.

Rubenstein, M. (1976), "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics*, 7, 407-425.

Shanken, J. (1987), "Multivariate Proxies and Asset Pricing Relations: Living with the Roll Critique," *Journal of Financial Economics*, 18, 91-110.

Shiller, R.J. (1982), "Consumption, Asset Markets, and Macroeconomic Fluctuations," *Carnegie-Rochester Conference Series on Public Policy*, 17, 203-238.

Singleton, K.J. (1988), "Specification and Estimation of Intertemporal Asset Pricing Models, "*Handbook of Monetary Economics*, B. Friedman and F. Hahn, eds., Amsterdam: North-Holland, 1988.

Snow, K. N. (1990), "Diagnosing Asset Pricing Models Using the Distribution of Asset Returns," Manuscript, University of Chicago.

Stambaugh, R.F. (1988), "The Information in Forward Rates: Implications for Models of the Term Structure," *Journal of Financial Economics*, 21, 41-70.

Sundaresan, M. (1989), "Intertemporally Dependent Preferences and the Volatility of Consumption and Wealth," *The Review of Financial Studies*, 2, 73-89.

Swensson L.E.O. (1985), "Money and Asset Prices in a Cash-in-Advance Economy," *Journal of Political Economy*, 93, 919-44.

Townsend, R.M. (1987), "Asset-Return Anomalies in a Monetary Economy," *Journal of Economic Theory*, 41, 219-247.

Wilson, R. (1968), "The Theory of Syndicates," *Econometrica*, 36, 119-132.