MODELING EXPECTED STOCK RETURNS FOR LONG AND SHORT HORIZONS

by

Shmuel Kandel
Robert F. Stambaugh

(42-88)

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

Modeling Expected Stock Returns for Long and Short Horizons

by*

Shmuel Kandel
University of Chicago and
Tel Aviv University

and

Robert F. Stambaugh
The Wharton School
University of Pennsylvania

First Draft: November 1988
This Version: March 1989

Comments Welcome

*Some of the material in this paper appeared in earlier drafts of the authors' working papers titled "Long-Horizon Returns and Short-Horizon Models" and "A Model of Expected Returns and Volatilities for Various Forecast Horizons." The authors are grateful to Andrew Abel, John Campbell, George Constantinides, Zvi Eckstein, Eugene Fama, Kenneth French, Michael Gibbons, Joel Hasbrouck, Narayana Kocherlakota, James Poterba, Michael Rozeff, and participants in workshops at the University of Chicago, the University of Pennsylvania, Tel Aviv University, and the 1988 Johnson Symposium for helpful discussions and comments. The first author gratefully acknowledges support from the Batterymarch Fellowship and the Center for Research in Security Prices, and both authors are grateful for support from the Instutute for Quantitative Research in Finance.

ABSTRACT

Expected returns over long and short horizons are modeled using two approaches: an equilibrium asset pricing model and a vector autoregression (VAR). Empirical properties of returns that are consistent with the equilibrium model's implications include (i) an annual "equity premium" of about six percent (ii) a U-shaped pattern of autocorrelations of returns with respect to investment horizons and (iii) a humped pattern with respect to investment horizon for the R-squared in projections of stock returns on predetermined financial variables. Parameters estimated in a monthly VAR for returns and these financial variables also imply autocorrelations, R-squared values, and conditional expected returns that are close to those computed with actual long-horizon returns. Simulations indicate that such a VAR is a reasonable approximation to the equilibrium model for representing the properties of expected short- and long-horizon returns.

1. Introduction

Empirical evidence indicates that expected returns on stocks and bonds vary through time. Much of this evidence is characterized either by autocorrelations of returns or by regressions of returns on various predetermined variables.[1] The length of the holding period over which a return is computed, or the return "horizon," seems to affect the nature of this evidence in significant ways. For example, the sample autocorrelations of returns on indexes of NYSE stocks are positive and in the range of 0.1 to 0.2 for horizons of one month [e.g., Fama and Schwert (1977)], but sample autocorrelations for horizons of five years are negative and in the range of -0.2 to -0.5 [e.g., Stambaugh (1986) and Fama and French (1988)]. In essence, the pattern of autocorrelations is U-shaped with respect to return horizon.[2] Regressions of one-month stock returns on predetermined variables often produce R-squared values less than 0.02 [e.g., Keim and Stambaugh (1986)], whereas regressions of longer horizon returns (several years) on similar predetermined variables often produce R-squared values in excess of 0.30 [e.g., Fama and French (1987)].

Given the apparent sensitivity of evidence about time-varying expected returns to the length of the return horizon, a parsimonious framework capable of integrating the existing empirical evidence would be useful. Such a framework could focus the efforts to pursue economic explanations for the behavior of expected returns. This study investigates two approaches to modeling the behavior of expected returns over both short and long horizons.

---

[1] A partial list of the studies reporting such evidence includes Fama and Schwert (1977), Hall (1981), Huizinga and Mishkin (1984), Fama (1984), Rozeff (1984), Keim and Stambaugh (1986), Campbell (1987), Fama and Bliss (1987), Fama and French (1987, 1988), Poterba and Summers (1989), Lo and MacKinlay (1988), and Huberman and Kandel (1988).

[2] This U-shaped pattern is also discussed by Poterba and Summers (1989). Lo and MacKinlay (1988) find positive autocorrelation in one-week index returns.

The first approach is based on an equilibrium asset-pricing model containing features similar to those in models developed by Lucas (1978), Mehra and Prescott (1985), and Abel (1988). We propose a model with time-additive utility and time-varying conditional moments of the monthly growth rate of output. Implications about the behavior of asset returns over various investment horizons are derived, and numerical examples of these implications are provided.

The examples are constructed for a set of parameters chosen to mimic roughly some of the empirical characteristics of asset returns. Our objective here is not to investigate whether various sample estimates, such as autocorrelations of long-horizon returns, are significantly different from zero or another value implied by a given choice of the model's parameters.[3] Rather, we simply view such sample estimates as empirical benchmarks for selecting the model's parameters. This approach is similar to that of Mehra and Prescott (1985), who attempt to calibrate a pricing model using sample estimates without addressing the statistical precision of those estimates.

An important reason for investigating a specific equilibrium model in this context is to gain additional insights about the extent to which the negative autocorrelations of long-horizon stock returns can be accommodated by explanations devoid of "fads" or other deviations from fundamental value. Poterba and Summers (1989) and Cochrane (1988), for example, discuss some of the properties that equilibrium expected returns would most likely possess in such explanations. The analysis here allows us to characterize other features of the equilibrium as well as the properties of expected stock returns.

The second approach uses a vector autoregression (VAR) to model the

[3]For example, the statistical precision and reliability of autocorrelations of long-horizon returns is considered by Fama and French (1988), Richardson (1988), and Cecchetti, Lam, and Mark (1988).

monthly time-series behavior of returns and other financial variables.[4] Various properties of long-horizon returns can be derived as functions of the parameters of the monthly VAR. We estimate a first-order VAR and then use the parameter estimates to illustrate several of the model's implications for long-horizon returns.

Both modeling approaches appear to be capable of capturing much of the observed behavior of expected returns over various investment horizons. For example, Fama and French (1988) observe a U-shaped pattern of negative first-order autocorrelations of equity returns with respect to return horizons beyond one year. We use the equilibrium model to demonstrate that this pattern is consistent with positively autocorrelated conditional means and variances of output growth. Using the same parameters for the equilibrium model, we also derive the (theoretical) R-squared values in regressions of equity returns for various horizons on several financial variables: a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread. The patterns of the R-squared values across investment horizons are similar to those obtained empirically using similarly defined financial variables [e.g., Fama and French (1987)].

The VAR model is estimated using monthly data on returns and three financial variables similar to those just described. Autocorrelations for long-horizon returns implied by these models are similar to the sample autocorrelations. That is, the same U-shaped pattern is obtained, with small positive autocorrelations at the shortest horizons and large negative autocorrelations at horizons of several years. The R-squared values in regressions of returns on the financial variables implied by the VAR

---

[4] Campbell and Shiller (1988) use a vector autoregression to investigate stock-price volatility and its relation to the forecastibility of long-horizon returns.

estimates are similar to the R-squared values produced in actual regressions--small for short horizons but larger for longer horizons.

The apparent adequacy of a simple VAR representation for stock returns is consistent with the conclusion of Lo (1988) that stock returns do not exhibit "long-term dependence" but instead can be accommodated by a class of models that include finite-order ARMA processes. A VAR can provide a parsimonious representation for a set of more complex univariate ARMA processes.[5]

The VAR approach to modeling, advanced notably by Sims (1980, 1981), does not attempt to specify the underlying equilibrium structure of the economy but instead pursues a reduced-form econometric representation. A question often debated among macroeconomists is whether a given VAR representation can represent the behavior of the economy across various states as well as a more structural model.[6] A similar question is addressed here in the context of our two approaches to modeling expected returns over short and long horizons. We examine the implications about short- and long-horizon returns produced by a large-sample VAR estimation using data simulated from the equilibrium model. We find that, for the specific equilibrium model entertained here, the implications produced by the VAR estimation are close to the true implications of the equilibrium model.

The remainder of the paper is organized as follows. In order to provide a context in which to consider the implications of the models to be developed, section 2 provides a brief summary of some empirically observed properties of returns over short and long horizons. Section 3 presents the equilibrium

---

[5] See, for example, Granger and Newbold (1977) for a summary of the relations between the multivariate and univariate representations.

[6] Kocherlakota (1988a) provides a summary of these issues and addresses some inference problems inherent in both the VAR approach and in more structural approaches.

model and its implications for the behavior of expected equity returns for short and long horizons. These implications are then illustrated for a specific choice of parameter values. Section 4 estimates the VAR model and then examines the implied behavior of returns over various horizons. Section 5 reports results of simulations that address the adequacy of a first-order VAR representation of the equilibrium model. Section 6 presents conclusions and suggests directions for future research.

2. Sample Estimates

Before proceeding to the models, we first provide a brief review of some of the previously observed empirical properties of stock returns over short and long horizons. This evidence, we believe, provides a useful context in which to view many of the results of our modeling efforts. Let $r_{t,N}$ denote the continuously compounded real return on the equally weighted NYSE index for the N-month period starting at the beginning of month t.[7] Figure 1 displays sample estimates of $\text{corr}(r_{t,N}, r_{t-N,N})$, that is, the first order autocorrelation of N-month returns. The estimates are based on monthly data for the period from December 1926 through December 1985 (709 observations) and use overlapping observations in the same manner as Fama and French (1988).[8] Figure 1 displays the same U-shaped pattern reported in previous studies. The sample autocorrelations are positive for horizons of one, two, and six to ten months, negative and decreasing up to the horizon of 54 months, and increasing toward zero at longer horizons.

---

[7] Month t is the month beginning at time t-1 and ending at time t. The monthly change in the natural logarithm of the Consumer Price Index is subtracted from the continuously compounded nominal return.

[8] The estimates are not bias adjusted. Fama and French (1988) report simulation evidence suggesting that the bias in the estimated autocorrelations is, in general, not severe when the true autocorrelations are similar to those displayed in the figure.

We turn next to regressions of $r_{t,N}$, the N-month return, on a set of predetermined variables. One striking characteristic of the results reported by previous studies is the behavior of the R-squared in these regressions. Although the R-squared statistic is generally viewed as, at best, a useful descriptive statistic, the magnitudes obtained in recent studies using long-horizon returns have been objects of significant attention by researchers in finance. Let x_t denote a vector of variables observed at the end of month t. For this exercise, we construct x_t to contain essentially the same three variables used in the regressions of Fama and French (1987):[9]

$(y_{Baa} - y_{Aaa})_t$: the difference at the end of month t between Moody's average yield on bonds rated Baa and bonds rated Aaa.

$(y_{Aaa} - y_{TB})_t$: the difference at the end of month t between the Aaa yield and the yield on a U.S. Treasury Bill with maturity closest to one month.

$(D/P)_t$: for the equally weighted portfolio of NYSE stocks, the ratio of dividends paid for the twelve months ending at t to the price at the end of month t.

Figure 2 displays, for various return horizons (N), the sample R-squared

[9] Similar variables have also been used by other researchers to predict asset returns. For example, Rozeff (1984) finds that dividend-price ratios predict stock returns, and Keim and Stambaugh (1986) find that (among other variables) the difference in yields between low-grade bonds and Treasury Bills predicts stock and bond returns. Contemporaneous changes in similar variables have also been used as common risk factors in empirical investigations of multifactor pricing models. In the latter context, Chen, Roll, and Ross (1986) use return spreads between (i) low-grade and high-grade bonds and (ii) long-term high-grade bonds and Treasury Bills.

value in regressions of $r_{t,N}$ on x_{t-1}.[10] As observed in previous studies, the R-squared is small in regressions using monthly returns, approximately 2%. As the return horizon increases, however, the R-squared increases steadily, to approximately 40% at a four-year horizon. [The values in figure 2 correspond closely to the results reported by Fama and French (1987).] Keeping in mind this glimpse at some of the existing evidence, we now turn to the task of modeling.

3. An Equilibrium-Based Analysis of Returns over Short and Long Horizons

The primary objective of this section is to use an equilibrium model to obtain implications about the behavior of expected returns over horizons of different lengths. This objective plays a role in our choice of some specific features of the model, but the general framework is quite standard and incorporates many features of similar models used previously by researchers to examine expected returns for single-period horizons.

3.1 The Model and Its Implications

We employ a representative-consumer, endowment model in the framework of Lucas (1978), where the physical stock of capital is fixed and aggregate consumption c_t equals aggregate output h_t in each period t. The consumer maximizes expected utility over an infinite horizon,

$$E \{ \sum_{\tau=t}^{\infty} \beta^{\tau-t} U(c_\tau) \} \quad , \tag{1}$$

where β (>0) is a rate of time preference. To this framework we add the

---

[10] Like the autocorrelations, the R-squared values are not adjusted for any finite sample bias and, therefore, provide upward biased estimates of the true R-squared, due primarily to the autocorrelation in the residuals caused by the use of overlapping observations.

following assumptions, which are similar to those contained in models developed by Mehra and Prescott (1985) and Abel (1988).

(A1) The utility function exhibits constant relative risk aversion,

$$U(c) = \frac{c^{1-\alpha} - 1}{1 - \alpha} \qquad 0 < \alpha < \infty \quad . \tag{2}$$

(A2) Let λ_{t+1} denote the one-period growth rate in output, i.e.,

$$h_{t+1} = \lambda_{t+1} h_t \tag{3}$$

The quantity $\ln(\lambda_{t+1})$ is, conditional on information at time t, distributed normally with mean μ_t and variance σ_t^2.

(A3) The pair $(\mu_t\ \sigma_t^2)$ follows a joint stationary Markov process with a finite number of states, S. Let s_t denote the state for $(\mu_t\ \sigma_t^2)$ at time t, where s_t can take values 1, ..., S. Let Φ denote the transition matrix with (i,j) element

$$\phi_{ij} = \text{Prob}(s_{t+1} = j \mid s_t = i) \tag{4}$$

Let π denote the vector of steady-state probabilities, i.e., π_i = Prob(s_t = i). Let $\lambda(i)$ denote the random growth rate drawn from the conditional distribution in which s = i.

(A4) Given s_t, the distribution of s_{t+1} is independent of λ_{t+1}, λ_t, λ_{t-1}, ...

Given the assumptions above, the state of the economy follows a Markov process. There are an infinite number of states, each represented as (c, s), where, at time t, $c_t = c$ and $s_t = s$. There are an infinite number of values for consumption ($0 \le c \le \infty$), but the number of values for s, which represents the conditional moments of consumption growth, is finite.

We use this model to derive prices of various types of financial claims as wełl as conditional moments of returns on various assets. For the financial claims considered here, either their prices depend only on s and not on c (as in the case of riskless bonds) or their prices are homogeneous of degree one in c (as in the cases of aggregate wealth, risky debt, and levered equity). The latter property is similar to that exploited by Mehra and Prescott (1985).[11]

Consider a riskless bond representing a claim on one unit of output to be received in N periods. It is easily shown that the price of this bond depends only on s. For example, let $p^{(F2)}(c, i)$ denote the price of a two-period bond at time t when the current state is (c, i). The states in the next two periods are represented as $(c\lambda(i), j)$ and $(c\lambda(i)\lambda(j), k)$. Applying the well known relation that the riskless rate (plus unity) equals the expected ratio of marginal utilities, discounted by the pure rate of time preference,

$$p^{(F2)}(c, i) = \beta^2 E\{ \frac{U'[c\lambda(i)\lambda(j)]}{U'[c]} \mid c, i \} \tag{5}$$

$$= \beta^2 E\{ \lambda(i)^{-\alpha}\lambda(j)^{-\alpha} \mid i \} \tag{6}$$

In computing the price of the bond, it is convenient to represent the

[11]The condition for equilibrium, analogous to that given by Mehra and Prescott (1985), is that the limit of A^m is zero as m becomes infinite, where the (i, j) element of A is $\beta\phi_{ij}E\{\lambda(i)^{1-\alpha}|i\}$.

expectation in (6) as an iterated expectation and then make use of the independence assumption (A4):

$$p^{(F2)}(c,\ i) = \beta^2 E\{\ E\{\ \lambda(i)^{-\alpha}\lambda(j)^{-\alpha}\ |\ i,\ j\ \}\ |\ i\ \} \tag{7}$$

$$= \beta^2 E\{\ E\{\lambda(i)^{-\alpha}\}\cdot E\{\lambda(j)^{-\alpha}\}\ |\ i\ \} \tag{8}$$

$$= \beta^2 E\{\lambda(i)^{-\alpha}\} \sum_{j=1}^{S} \phi_{ij} E\{\lambda(j)^{-\alpha}\} \tag{9}$$

$$= \beta^2 \sum_{j=1}^{S} \sum_{k=1}^{S} \phi_{ij}\phi_{jk} E\{\lambda(i)^{-\alpha}\}\cdot E\{\lambda(j)^{-\alpha}\} \tag{10}$$

The form of the expression in (10) extends easily to additional periods, as summarized by the following proposition.

<u>Proposition 1.</u> Let $p^{(FN)}$ denote the S-vector of prices of the N-period riskless claims in each of the S states. Then

$$p^{(FN)} = \Psi^N \iota_S \quad , \tag{11}$$

where ι_S is an S-vector of ones, Ψ is the $S \times S$ matrix with (i,j) element

$$\psi_{ij} = \beta\ \phi_{ij}\ E\{\lambda(i)^{-\alpha}\} \quad . \tag{12}$$

and Ψ^N denotes the Nth power of the matrix Ψ.[12]

[12]Note that, given (A2), for any m, $E\{\lambda(i)^m\} = \exp(\mu_i m + 0.5\sigma_i^2 m^2)$, where μ_i and σ_i^2 denote the conditional mean and variance of $\ln[\lambda(i)]$.

Next consider the value of aggregate wealth when the state of the economy is (c, i). Let $p^{(A)}(c,i)$ denote this price. It can be shown that a result similar to that in Mehra and Prescott (1985) also obtains in this case. That is, there exists w_i such that $p^{(A)}(c,i) = w_i c$.

Proposition 2. $p^{(A)}(c,i) = w_i c$, where w_i is the ith element of the S-vector **w**, given by

$$\mathbf{w} = (I - H)^{-1} H \iota_S \quad , \tag{13}$$

and H is the S × S matrix with (i,j) element

$$h_{ij} = \beta \, \phi_{ij} \, E\{\lambda(i)^{1-\alpha}\} \quad . \tag{14}$$

We next consider a risky one-period bond that promises to pay, at the end of one period, a fraction θ of current aggregate wealth. That is, if the current state is (c, i) and the state in one period is $(\lambda(i)c, j)$, the payoff on the risky bond will be

$$\min[\, p^{(A)}(\lambda(i)c,\ j) + \lambda(i)c, \quad \theta p^{(A)}(c,\ i)\,] \quad . \tag{15}$$

Let $p^{(B)}(c, i)$ denote the price of this bond.

Proposition 3. $p^{(B)}(c, i) = c \cdot g_i$, and g_i is the i^{th} element of the S-vector $\mathbf{g} = Y\iota_S$, where Y is an S×S matrix with (i, j) element

$$y_{ij} = \beta \, \phi_{ij} E\{\min[\lambda(i)^{1-\alpha}(1 + w_j),\ \lambda(i)^{-\alpha}\theta w_i]\} \quad . \tag{16}$$

Levered equity is defined, for a given θ, as the residual claim on output and capital net of the risky bond defined above. That is, if $p^{(L)}(c, i)$ denotes the value of the levered equity, then $p^{(L)}(c, i) = p^{(A)}(c, i) - p^{(B)}(c, i)$. From propositions 2 and 3

$$p^{(L)}(c, i) = c\cdot(w_i - g_i) \tag{17}$$

The prices of aggregate wealth and levered equity [proposition 2 and equation (17)] are both expressed as the product of current consumption and a quantity depending only on s, the state for (μ_t, σ_t^2). This multiplicative form suggests a decomposition of the natural logarithm of the price as the sum of two components, one stationary and the other nonstationary. With levered equity, for example,

$$\ln[p^{(L)}(c, i)] = \ln(c) + \ln(w_i - g_i) \quad , \tag{18}$$

and it is clear from the assumptions and the previous discussion of the model that ln(c) is nonstationary and $\ln(w_i - g_i)$ is stationary. The same type of decomposition holds for the price of aggregate wealth.[13] Fama and French (1988) and Poterba and Summers (1987) also consider decompositions into stationary and nonstationary components, but their analyses assume that the nonstationary component is a random walk with increments that are independent of innovations in the stationary component. In (18), however, the first-differences of the nonstationary component ln(c) are autocorrelated,

---

[13]The model of Mehra and Prescott (1985) also admits this sort of decomposition. In their model, both components depend on realized consumption growth.

conditionally heteroskedastic, and correlated with innovations in the stationary component, $\ln(w_i - g_i)$.[14]

The Appendix gives, for both aggregate wealth and levered equity, the conditional means and conditional variances of returns in each state of the economy for investment horizons of various lengths. As shown in the Appendix, the conditional moments of returns depend only on s, the state for (μ_t, σ_t^2), and not on c. Also given in the Appendix are unconditional means, variances, and first-order autocorrelations of returns for various investment horizons. The calculation of these unconditional moments is facilitated by the fact that the conditional moments depend only on s.

The model also provides implications about the goodness of fit, or R-squared, of linear projections of return on various financial variables. For this analysis we define three financial variables that also depend only on the state s. These variables correspond roughly to the three financial variables used in the empirical work in the paper. The dividend-price ratio in state (c,i), dp(i), is defined as the ratio of conditional expected consumption to the price of levered equity:

$$dp(i) = \frac{c \cdot E\{\lambda(i)|i\}}{p^{(L)}(c,\ i)} = \frac{E\{\lambda(i)\}}{w_i - g_i} \quad . \tag{19}$$

The default spread in state (c, i), ydef(i), is defined as the yield to maturity on the risky bond in proposition 3 minus the one-period riskless rate:

[14] This does not imply that our model cannot yield a decomposition more similar to those entertained by these studies.

$$\begin{aligned} \text{ydef}(i) &= \frac{\theta p^{(A)}(c,\ i)}{p^{(B)}(c,\ i)} - \frac{1}{p^{(F1)}(c,\ i)} \\ &= \theta \frac{w_i}{g_i} - \frac{1}{\beta E\{\lambda(i)^{-\alpha}\}} \end{aligned} \tag{20}$$

An N-period term spread in state (c, i), yterm(i), is defined as the yield spread between the one-period riskless bond and an N-period riskless bond:

$$\text{yterm}(i) = \left[\frac{1}{p^{(FN)}(c,\ i)} \right]^{1/N} - \frac{1}{p^{(F1)}(c,\ i)} \quad . \tag{21}$$

(Recall from proposition 1 that the prices in (21) do not depend on c.) In the example given below, we set N = 240 (months).

Since the above three financial variables depend only on the state s, it is straightforward to compute, $\bar{\rho}_N$, the multiple correlation coefficient between the expected levered-equity return for an N-period horizon, $R_{t,N}^{(L)}$, and the three financial variables. The implied R-squared in a projection of $R_{t,N}^{(L)}$ on these three variables is then equal to $(\bar{\rho}_N)^2$ times the ratio of the variance of the conditional expected return to the variance of the return.

3.2 A Numerical Example

In this section we illustrate numerically several of the model's implications about the behavior of returns for various horizons. We define a single period as 1 month, and we compute implications for return horizons up to 120 months. We choose parameters to yield implications that coincide reasonably well with characteristics of the data, but we do not claim (and in fact we doubt) that these choices necessarily calibrate the model to give the best "fit" across a number of dimensions. In principle, one could specify a list of moment conditions and estimate the model, but this task is beyond the scope of the present study. Our objective here is simply to sketch out one type of equilibrium model that seems capable of capturing much of the behavior of returns over short and long horizons.

The example is constructed by specifying (i) the Markov process for the conditional moments of consumption growth, (ii) the parameters of the utility function, α and β, and (iii) the value of θ for defining levered equity. The assumed process for consumption growth rates has the same unconditional mean and variance, when stated on an annual basis, as do the annual values used in Mehra and Prescott (1985). The example is also constructed so that the riskless rate and the premium on levered equity implied by the model are equal to those used as the empirical benchmarks in Mehra and Prescott (1985). In addition, we attempt to mimic roughly the overall empirically observed patterns with respect to investment horizon of autocorrelations and R-squared values in regressions of returns on the three financial variables.

A nine-state Markov process is specified for the conditional moments of consumption, (μ_t, σ_t^2). In addition to the unconditional mean and variance of consumption growth, the other parameters specified in constructing the Markov process are (i) the variance of the conditional mean consumption growth, (ii) the variance of the conditional variance of consumption growth, (iii) the autocorrelation of the conditional mean, and (iv) the autocorrelation of the conditional variance. Thus, six parameters are specified in constructing the transition matrix.

We use a procedure that allows these six parameters to imply a unique transition matrix for any given number of states for s. We assume that μ_t and σ_t^2 are independent of each other. A joint Markov process for μ_t and σ_t^2 is constructed by first constructing Markov processes for each parameter. These processes are formed as discrete approximations to first-order autoregressions for μ_t and $\ln(\sigma_t^2)$ with normally distributed errors. The transition matrix can then be calculated easily, given the assumed independence. This technique is similar to that of Tauchen (1986). The assumption that μ_t and σ_t^2 are independent is made for tractability in constructing the example and is not a necessary feature of the model.

Table 1 reports the values and probabilities specified for the Markov process for the conditional moments of consumption growth, and panel A of table 2 reports the parameters summarizing the consumption process.[15] One value that does not match its analogue in Mehra and Prescott (1985) is the autocorrelation of consumption growth rates. They use an annual autocorrelation of -0.14, whereas our assumed process has a monthly autocorrelation of 0.001. The conditional mean growth rate has an autocorrelation of 0.9178, but its variance accounts for a small fraction of the total variance in consumption growth (accounting for the low autocorrelation in the actual growth rate). The conditional variance of the growth rate has a monthly autocorrelation of 0.432.[16]

The parameters of the utility function are specified as $\alpha = 28.55$ and $\beta = 0.99975$, and θ is set equal to 0.413 (so the equity is approximately 60%

---

[15]The "annualized" values (table 2) for the means and variances are simply the monthly values multiplied by 12. Thus, given that the monthly growth rates are neither independently nor identically distributed, this gives only an approximation to the true annual values. Mehra and Prescott's numbers describe simple percentage growth rates, whereas ours describe continuously compounded growth rates.

[16]Six parameters are specified in constructing the Markov process, as discussed earlier, but panel A of table 2 displays nine values. Three of these values are determined by the other six. Specifically, (i) the mean conditional growth rate, (ii) the mean conditional standard deviation, and (iii) the autocorrelation of actual growth rates are determined by the six parameters listed earlier.

of aggregate wealth). This value of the curvature parameter α is high by traditional standards. Indeed, Mehra and Prescott recognize that higher values of α can produce the benchmark interest rate and equity premium, and they restrict α to be less than 10.

A central argument for a lower value of α, cited by Mehra and Prescott and many others, is that sample estimates of the ratio of expected excess return to variance of return, the "price of risk" computed by Friend and Blume (1975), are generally less than 2.0. Friend and Blume interpret the price of risk, divided by a proportion of risky assets demanded between 0.5 and 0.8, as a measure of relative risk aversion. While this calculation may be useful in the case of independently and identically distributed rates of return, such a relation does not necessarily hold in other settings.[17] We note that the price of risk for equity in our example is low--somewhat lower, in fact, than typical sample estimates. Kocherlota (1988b) obtains a similar result using $\alpha = 13.7$ and a value of β greater than unity ($\beta = 1.14$ for annual periods). Higher values of α are also entertained by Black (1988) in a continuous-time model with a time-additive utility and a time-varying price of risk. Black concludes that the model allows both the equity premium and the average price of risk to be calibrated to sample estimates. Although he does not propose a numerical value for α, Black observes that it "must be much larger than one."

Numerous other studies report estimates of relative risk aversion that vary significantly, depending on (i) the specification of the asset pricing model's implications, (ii) the sample period, (iii) the frequency of the data (monthly, quarterly, or annual), (iv) the use of real or nominal quantities, and (v) adjustments for temporal aggregation. Studies using monthly

---

[17] Brown and Gibbons (1985) assume identically and independently distributed consumption growth rates in their approach to estimating the coefficient of relative risk aversion.

consumption data often obtain estimates of relative risk aversion below unity [e.g., Hansen and Singleton (1982, 1983)], while other studies that use quarterly or annual data and attempt to account for temporal aggregation have obtained estimates of 100 or more [e.g., Grossman, Melino, and Shiller (1985), Naik and Ronn (1988).][18]

If the reader has strong priors that α is low, then our inability (and others') to produce reasonable empirical implications in models with time-additive utility and low values of α could be viewed as evidence against such equilibrium explanations of observed return behavior.[19] Given the imprecision associated with estimates of α, however, we are reluctant to exclude equilibrium explanations solely because these explanations include a higher value of α than traditionally believed to be "reasonable."

The observed empirical properties of stock returns can also be captured by the model using a lower value of α if the variance of aggregate consumption is made higher. For example, Cecchetti, Lam, and Mark (1989) investigate a model with time-additive utility and constant relative risk aversion, and they find that the equity premium implied by the model matches the Mehra-Prescott empirical benchmark for $\alpha = 1.40$. In their specification, however, the marginal rates of substitution are computed using real dividends on the

---

[18] Grossman, Melino, and Shiller describe their estimates over 20 as "too large to be plausible," but they do observe that their highest estimates of risk aversion are accompanied by their weakest evidence against the pricing model's overidentifying restrictions. Naik and Ronn argue that the higher values of relative risk aversion are consistent with reasonable real interest rates and equity premiums. Hall (1988) concludes, based on an empirical examination of consumption and asset returns, that the intertemporal elasticity of substitution is probably quite low. If utility is additively separable, this result also implies a high degree of relative risk aversion (although Hall does not argue for such a conclusion).

[19] For an approach to resolving the equity-premium puzzle with utility that is not time additive but instead incorporates "habit formation" see Constantinides (1988).

Standard and Poor's Composite Index. They report for this series an annual standard deviation of the growth rate equal to 0.1359, about four times higher than the standard deviation of annual consumption growth (0.0357) used here and in Mehra and Prescott (1985).

Panel B of table 2 reports some implied properties of returns in our numerical example. The riskless real rate of interest has an (annualized) mean of about 0.8% and a standard deviation of 4.1%. There is no true sample counterpart in the U.S., since one-month index bonds do not exist, but Mehra and Prescott (1985) report a real return on nominally riskless securities with the same mean and a standard deviation of 5.7%. As compared to the real return on the S&P 500 reported by Mehra and Prescott, the model's real return on levered equity has the same mean (by construction) but a higher standard deviation (27.7% versus 16.5%). The ratio of expected excess return to variance (price of risk) equals 1.31 for aggregate wealth and 0.80 for levered equity. Note that, for the monthly return on levered equity, both the conditional mean and the conditional standard deviation have autocorrelations of approximately 0.77.

The risky bond used for the default yield spread, ydef(i) in (20), is defined for $\theta = 0.95$, i.e., the bond promises to pay 95% of the current total wealth. Implied properties of the three financial variables are given in panel C of table 2. Comparisons to sample estimates must be crude, since these variables correspond only roughly to those used here and elsewhere in empirical investigations. The implied first-order monthly autocorrelations of these variables range from 0.74 (term spread) to 0.92 (dividend-price ratio), and these are similar to autocorrelations observed by others for similar variables.

The implied first-order autocorrelation of the monthly return on levered equity is -0.04. Therefore, this specification of the model does not imply the positive autocorrelation of monthly stock index returns found in the data. As displayed in figure 3, however, the autocorrelations implied by this specification (solid curve) decline up to horizons of about two years, to almost -0.3, and then they increase gradually toward zero for longer horizons. Thus, for investment horizons beyond several months, the implied pattern of autocorrelations coincides at least roughly to that found in the data (cf. figure 1).

We were unable to find parameter specifications for the model that result in positive autocorrelations at short horizons but negative autocorrelations at longer horizons. The model entertained here is restrictive in a number of respects, however. One feature of the model that could be important in this regard is the assumption (A4) that unexpected consumption growth does not impact the change in the conditional moments of consumption growth. [A similar assumption appears in the model of Abel (1988).][20]

Figure 3 also displays the implied autocorrelations obtained for two lower values of relative risk aversion, $\alpha = 10$ and $\alpha = 2$. As relative risk aversion declines, ceteris paribus, the pattern of autocorrelations flattens toward zero. For a sufficiently low α, between 1 and 2 in this case, the implied autocorrelations for all return horizons become positive. (Of course, the riskless rate, equity premium, etc. also change with α.)

Figure 4 displays the R-squared that the model implies for a projection of $R_{t,N}^{(L)}$ on the three financial variables defined above. The R-squared implied for this specification of the model (solid curve) is low for the one-month horizon (0.02), increases to 0.13 at approximately a two-year horizon, and then declines toward zero. Although the implied R-squared does not reach

[20] We confined our search to cases where μ_t and σ_t^2 are independent.

values as high as the sample estimates, the pattern of low values for short horizons and higher values for longer horizons is similar to that found in the sample estimates (cf. figure 2). Figure 4 also displays the implied R-squared for the two lower values of α considered in figure 3. As α declines, the implied R-squared values move toward zero for all investment horizons.

In addition to high risk aversion, another feature of the model that plays a key role in determining the above patterns is the monthly autocorrelation of the conditional expected consumption growth rate, which we denote as ρ_μ for this discussion. As reported earlier in table 2, $\rho_\mu = 0.92$ in the example. Figure 5 displays autocorrelations of returns for various horizons implied by that value as well as three other values of ρ_μ: 0.5, 0.0, and -0.5 (maintaining all other original specifications). For $\rho_\mu = 0.5$, the pattern of return autocorrelations is U-shaped but reaches its minimum at a shorter horizon (four months). For lower but non-negative values of ρ_μ the U-shape disappears and the negative return autocorrelations increase monotonically toward zero (e.g., $\rho_\mu = 0.0$). Negative values of ρ_μ (e.g., -0.5) also produce the largest negative return autocorrelations at one month but they approach zero with some oscillation.

Similar effects of ρ_μ are found in the implied R-squared values, which are displayed in figure 6 for the same four values of ρ_μ used in figure 5. For $\rho_\mu = 0.5$, the hump occurs at a shorter horizon (four months). Then as ρ_μ declines the hump disappears and R-squared declines monotonically (e.g., $\rho_\mu = 0.0$). Finally, the R-squared starts at its highest value for one-month but moves toward zero with oscillation (e.g., $\rho_\mu = -0.5$).

The autocorrelation of the conditional standard deviation of the consumption growth rate, which we denote ρ_σ, appears to play an interesting secondary role. When conditional mean growth rates vary, as they do in the example, then changing ρ_σ produces little if any effect on either the return autocorrelations or the R-squared values. When the conditional expected growth rate is constant, however, then changing ρ_σ produces patterns and magnitudes for the return autocorrelations and R-squared values that are virtually identical to those produced by the corresponding values of ρ_μ. In fact, we can construct an alternative example that produces very similar implications to those reported for the original example by specifying a constant conditional expected growth rate and $\rho_\sigma = 0.9$.

4. A VAR Approach to Modeling Returns for Various Horizons

In this section we estimate a first-order vector autoregression using monthly time series of four variables--stock index returns and three other financial variables. The estimates obtained are then used to derive a number of implications about the behavior of returns over various horizons.

4.1 Definition and Estimation of the VAR

Let $r_{t,N}$ and x_t be the same quantities defined earlier in section 2. That is, $r_{t,N}$ is the N-period continuously compounded real return on stocks, and x_t contains three financial variables--a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread.

Let

$$y_t \equiv \begin{bmatrix} x_t \\ r_{t,1} \end{bmatrix} .$$

Table 3 reports sample means, standard deviations, and pairwise correlations for the four variables in y_t. (The yield-related variables are stated as percent per month.) .

Define the first-order vector autoregression (VAR),

$$y_t = g_0 + G\, y_{t-1} + \nu_t \qquad , \tag{22}$$

where $E\{\nu_t | y_{t-1}\} = 0$, $E\{\nu_t \nu_{t-s}'\} = 0$ for $s \neq 0$, and the eigenvalues of G are assumed to be less than unity in absolute value.

Table 4 reports ordinary least-squares estimates of the parameters in the VAR model. The estimation is based on monthly data for the period from December 1926 through December 1985 (709 observations). Table 4 also shows the first six autocorrelations of the residuals from each of the four equations, and they are, in general, close to zero. We do not present extensive diagnostic tests of the adequacy of this first-order VAR for representing the behavior of the monthly time series used here. Our primary objective in this study is to investigate the extent to which simple models can capture the observed behavior of returns for various investment horizons. If this basic approach appears to be useful, then future research should consider alternative time-series models.

4.2 Implications of the Estimated VAR

Although the VAR is defined and estimated in terms of one-month-ahead forecasts, the parameters of the model can be used to obtain implications about the behavior of expected returns for longer horizons. To illustrate this point, consider the expected return for a five-year horizon. The estimated parameters reported in table 4 can be used to obtain implied values for the coefficients in a regression of the five-year return, $r_{t,60}$, on y_{t-1}, the lagged values of the four variables in the VAR. Figure 7 plots the implied expected returns for five-year horizons (stated on a per-month basis) obtained from the VAR. Also shown are the fitted expected returns obtained by regressing directly the five-year return $r_{t,60}$ on y_{t-1} (using monthly observations of overlapping five-year returns). The expected returns implied by the VAR(1) model coincide closely with the expected returns estimated directly in the regression.

We illustrate below several additional ways in which the VAR gives implications about the behavior of expected long-horizon returns. In order that the reader not be burdened unnecessarily with algebraic manipulation, we simply present the results of our analyses graphically and omit the underlying formulas. All of the computations follow directly, however, from the estimated parameters reported in table 4.

4.2.1 Autocorrelations of Returns

We consider next the first-order autocorrelation for returns over N-month horizons, corr($r_{t,N}$, $r_{t-N,N}$). Figure 8 plots, for horizons of one month through ten years, the autocorrelations implied by the VAR. Also displayed for comparison are the actual (unadjusted) sample autocorrelations shown earlier in figure 1. Note that the VAR implies properties for corr($r_{t,N}$, $r_{t-N,N}$) that correspond to properties of the sample estimates reported in previous studies. The implied autocorrelation begins at 0.15 for 1-month returns, becomes negative at a 4 month horizon, declines to -0.30 at a 42-month horizon, and then moves back toward zero for longer horizons. Although the implied long-horizon autocorrelation does not reach as low a value as the actual sample autocorrelation, the overall pattern of the implied autocorrelations is quite similar to that of the sample autocorrelations. For short horizons (small N), the implied autocorrelations are positive, and this implication is consistent with previous evidence indicating positive autocorrelation in short-horizon stock-market returns [e.g., Fama and Schwert

(1977) and Lo and MacKinlay (1987)].

None of the values shown are adjusted for finite-sample bias, which is present in the sample autocorrelations as well as in the estimates of the parameters reported in tables 3 and 4. It may be the case that such biases have different effects on the implied autocorrelations and the sample autocorrelations, so precise comparisons based on figure 9 should be made cautiously.

Given that the VAR model includes the lagged one-month return as a predictive variable, it may not be surprising that the model yields an implied autocorrelation for one-month returns that closely resembles the sample estimates. We have also examined a restricted VAR in which (i) the change in the Baa-bond yield appears as an additional variable and (ii) the coefficient on the lagged monthly stock return is constrained to be zero in the stock-return equation. In other words, the lagged stock return is not allowed to contribute directly to predicting the following month's stock return. The estimated parameters of this restricted version of the VAR model also imply positive autocorrelations for short horizons.[21]

4.2.2 Regression R-squared

Figure 9 displays the value of the R-squared implied by the VAR model for regressions of returns on the three predetermined variables in x_{t-1}. That is, in the regression

$$r_{t,N} = \alpha_{0N} + \alpha_N' x_{t-1} + u_{t,N} \quad , \tag{23}$$

the R-squared is the value of $\mathrm{var}(\alpha_N' x_{t-1})/\mathrm{var}(r_{t,N})$ implied by the VAR. This

[21] If the zero restriction is imposed but the additional yield-change variable is not introduced, then the implied one-month autocorrelation is slightly negative.

R-squared value is computed for return horizons (N) ranging from one month to four years. Also shown in figure 9 is the actual (unadjusted) sample R-squared value obtained by regressing $r_{t,N}$ on x_{t-1} using overlapping observations, displayed earlier in figure 2. The R-squared values implied by the VAR exhibit properties similar to those of the sample values. For example, the values begin at 0.02 for a 1-month horizon and increase to 0.23 for a 34-month horizon.[22]

4.2.3 Impulse Response Functions

One potential benefit of the VAR model is that it permits a deeper analysis of the manner in which changes in the predetermined variables impact expected future returns for various horizons. One framework that appears to be especially useful in this analysis is that proposed by Sims (1980, 1981). This approach computes the response of a given variable, in this case $r_{t+n,1}$, to a set of orthogonal shocks in each variable in the system.

The VAR model allows $r_{t+n,1}$ (the return in month t+n) to be written as

$$r_{t+n,1} = \mu + \sum_{j=0}^{\infty} \theta_j' \xi_{t+n-j} \quad , \tag{24}$$

where ξ_s is a vector containing four elements, with $E(\xi_s) = 0$, $\mathrm{var}(\xi_s) = I$, and $\mathrm{cov}(\xi_s, \xi_{s-j}) = 0$ for all $j \neq 0$. The elements of ξ_t are constructed as follows. The first element ξ_{1t} is the shock in period t to the first variable in x_t. The second element ξ_{2t} is the shock to the second variable in x_t that is uncorrelated with ξ_{1t}; the third element ξ_{3t} is the shock to the third variable that is uncorrelated with ξ_{1t} and ξ_{2t}, etc. Thus, the construction

[22] As in the analysis of autocorrelations, the R-squared value implied by the VAR is not adjusted for any finite sample bias.

of ξ_t depends on the ordering of variables. We order the variables as in the vector y_t defined above, so that the fourth element of ξ_t is the shock to $r_{t,1}$ that is uncorrelated with shocks to any of the variables in x_t.

The response of $r_{t+n,1}$ to the shocks (or "impulses") in the variables in period t is represented by the parameter vector θ_n. That is, θ_{1n} is the response of $r_{t+n,1}$ to a one-standard-deviation shock at period t in the first variable, θ_{2n} is the response to the orthogonalized one-standard-deviation shock in the second variable, etc. The values of θ_{in} for various n represent the "impulse response function," the responses of returns in various future months to a one-standard-deviation (orthogonalized) shock to variable i at time t.

Figure 10 plots the responses of monthly returns to shocks in each of the four variables in y_t. Responses of $r_{t+n,1}$ to shocks in the first variable, $(y_{Baa} - y_{Aaa})_t$, are negative for the first two months, positive and increasing for the next three months, and then monotonically decreasing, converging to zero in about four years. The responses to shocks in the variable $(y_{Aaa} - y_{TB})_t$ that are uncorrelated with the shocks in the first variable are positive for almost two years and then decline monotonically toward zero. The responses to shocks in the variable $(D/P)_t$ that are uncorrelated with the shocks in the first two variables are very similar to those of the first variable, but the decline to zero takes almost six years. The responses to shocks in the fourth variable, the shocks in r_t that are uncorrelated with the shocks in the other variables, are positive and decreasing for the first two months, negative and decreasing for two more months, and then increasing monotonically, reaching zero in about two years.

4.2.4 Decomposing the Autocorrelation

The infinite moving average representation in (24) can also be used to analyze the contribution of each of the predictive variables to the autocorrelation of $r_{t,N}$, the return for an N-month horizon. First note that, since

$$r_{t,N} = \sum_{j=0}^{N-1} r_{t+j,1} \quad , \tag{25}$$

the N-horizon return can be written as

$$r_{t,N} = N\mu + \sum_{j=1}^{\infty} \gamma_j' \xi_{t+N-j} \quad . \tag{26}$$

Equation (26) can be rewritten as

$$r_{t,N} = N\mu + \sum_{i=1}^{4} \zeta_{i,t+N-1} \tag{27}$$

where

$$\zeta_{i,t+N-1} = \sum_{j=1}^{\infty} \gamma_{i,j} \xi_{i,t+N-j} \quad . \tag{28}$$

Since the elements of ξ_s are mutually uncorrelated by construction, equation (28) expresses the N-month return as a sum of four orthogonal components. By decomposing the return in this fashion, the autocovariance of $r_{t,N}$ can be similarly decomposed:

$$\mathrm{cov}(r_{t,N}, r_{t-N,N}) = \sum_{i=1}^{4} \mathrm{cov}(r_{t,N}, \zeta_{i,t-1}) \quad . \tag{29}$$

Each term on the right-hand side of equation (29) represents the portion of the autocovariance of $r_{t,N}$ that can be attributed to the covariance between $r_{t,N}$ and past (orthogonalized) shocks to a given variable in the system. Dividing each side of (29) by var($r_{t,N}$) gives the autocorrelation as a sum of four components.

Figure 11 displays, for various return horizons (N), the four components of the autocorrelation in $r_{t,N}$. The solid line represents the autocorrelation of $r_{t,N}$ implied by the VAR, which was displayed previously in figure 8. The other four curves in figure 11 sum to this total autocorrelation. Perhaps the most striking result of this exercise is that shocks to the dividend-price ratio appear to make the most important contribution to both the positive autocorrelation in short-horizon returns as well as to the negative autocorrelation for the longer horizons.

5. <u>Evaluating a VAR Approximation to the Equilibrium Model</u>

This section uses simulations to study (i) the ability of the VAR discussed in section 4 to approximate the equilibrium model in section 3 and (ii) the small-sample bias in the estimates presented. As in previous sections, we focus on the autocorrelations of returns and the regression R-squared values for various investment horizons.

Data for this section's analyses are generated by simulating the equilibrium model with the parameter values specified in the numerical example of section 3.2. Two types of simulations are conducted, small-sample and large-sample. In a small-sample simulation, 100 independent samples are generated. Each sample includes 709 observations of the continuously compounded monthly return on levered equity and the three financial variables defined in the model: a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread. The estimated statistics of the 100 samples are averaged to obtain "average small sample" values. We also generate a data set of 70,900 observations and denote its statistics as "large sample" values.

The simulation results for the autocorrelations are presented graphically in figure 12, and figure 13 displays corresponding results for the regression R-squared. The solid line in each figure represents the true values implied by the equilibrium model, shown previously in figures 3 and 4, for autocorrelations of simple returns and R-squared values in regressions with simple returns. The sample estimates in section 2 and the VAR estimation in section 4 use continuously compounded returns, however, and we have not obtained analytical expressions for the equilibrium model's autocorrelations and R-squared values with continuously compounded returns. Therefore, we simulate 70,900 monthly observations and calculate the autocorrelations and the R-squared values in the same manner used to compute the statistics reported in section 2 using actual data. As can be seen in figures 12 and 13, these values are very close to the exact values for simple returns.

A different large-sample simulation is used to assess the VAR as an approximation to the equilibrium model. We estimate the VAR using the simulated 70,900 monthly observations and then calculate the autocorrelations and R-squared values implied by the estimated VAR parameters. The implied autocorrelations are very close to the large-sample estimates obtained above and to the true values for simple returns obtained from the equilibrium model. However, the R-squared values for long horizons implied by the large-sample estimates of the VAR are larger than the R-squared values calculated in the large sample above.

We conduct another simulation to assess the small-sample bias in the statistics. For each of the 100 generated data sets of 709 observations, we calculate the autocorrelations and the R-squared values in the same way used with the actual data in section 2. The dashed line in each of figures 12 and 13 represents the average over the 100 data sets of these small-sample statistics. Note that, for long horizons, the average small-sample autocorrelations are much lower (larger negative) than those in the large-sample simulation, and the average small-sample R-squared values are much higher than those in the large-sample simulation.

A final simulation is conducted to examine the behavior of the VAR-implied autocorrelations and R-squared values computed from small-sample estimates of the VAR parameters. We estimate a VAR model for each of the 100 simulated data sets and compute the implied autocorrelations and R-squared values. Figures 12 and 13 plot the averages of these statistics (over the 100 data sets). The average implied autocorrelations are very close to those implied by the large-sample estimation of the VAR, to the true autocorrelations of the simple returns, and to the autocorrelations of the continuously compounded returns obtained in the large sample. The R-squared results are somewhat unexpected. The average implied R-squared values are smaller than those implied by the large-sample VAR estimates, but they are almost identical to those for the continuously compounded returns calculated in the large sample and to the true values for simple returns.

6. Conclusions

A number of the empirically observed properties of expected stock returns for various investment horizons appear to be consistent with an equilibrium pricing model that includes (i) positively autocorrelated conditional means and variances of consumption growth and (ii) risk aversion that is high by traditional modeling standards. The high risk aversion in this model does not, however, imply a "price of risk" above what is typically estimated.

Other properties of returns that are consistent with the model include (i) an annual "equity premium" of about six percent (ii) a U-shaped pattern of negative autocorrelations of returns with respect to investment horizons beyond several months, and (iii) a humped pattern with respect to investment horizon for the R-squared in projections of stock returns on predetermined financial variables--a dividend price ratio, a default yield spread, and a term yield spread.

The equilibrium model analyzed here also implies values in each state of the economy for the conditional expected return and the conditional variance of return for various investment horizons. The states of the economy are characterized by the conditional moments of consumption growth. Kandel and Stambaugh (1989) examine the model's implications across states of the economy for the same parameter values used here in the numerical example. They find that the tendency for estimates of the price of risk to be higher during recessions, when coupled with business-cycle variation in estimates of the moments of consumption growth, appears to be consistent with the equilibrium model.

Empirical properties of expected long- and short-horizon stock returns also appear to be represented well by a first-order vector autoregression (VAR) for monthly returns and these three financial variables. Parameters estimated for the monthly VAR can be used to obtain implications about (i) expected long-horizon returns, (ii) the autocorrelations of returns for various horizons, and (iii) the R-squared values in regressions of short- and long-horizon returns on the financial variables. The values for these statistics implied by the estimated VAR parameters are close to the sample estimates computed with actual long-horizon returns. Simulation evidence suggests that such a VAR is also a reasonable approximation to the equilibrium model in terms of capturing the properties of expected short- and long-horizon returns.

APPENDIX

This appendix gives moments of returns, as implied by the equilibrium model in section 3, for investment horizons of arbitrary length. We first give the conditional means and variances of returns on aggregate wealth and levered equity, and then we use these conditional moments to obtain unconditional means, variances, and autocorrelations. Proofs of the propositions are omitted but are available upon request to the authors.

Using proposition 2, the one-period simple rate of return on aggregate wealth, where the economy goes from state (c, i) to state (cλ(i), j), is given by

$$R^{(A1)} = \frac{c\lambda(i) + w_j c\lambda(i)}{cw_i} - 1 \tag{1a}$$

$$= \frac{\lambda(i)(1 + w_j)}{w_i} - 1 \quad . \tag{2a}$$

Note that the conditional distribution of $R^{(A1)}$ depends only on i. Following a line of reasoning similar to that invoked in (7) through (10) leads to conditional means and variances for rates of returns over various horizons.

<u>Proposition A1.</u> Let $E^{(AN)}$ denote the S-vector of conditional expected N-period (simple) rates of return on aggregate wealth in each of the S states, $1 \le s \le S$.

$$E^{(AN)} = M^N \iota_S - \iota_S \quad , \tag{3a}$$

where M is an S×S matrix with (i, j) element

$$m_{ij} = \phi_{ij} E\{\lambda(i)\} \frac{(w_j + 1)}{w_i} \quad . \tag{4a}$$

<u>Proposition A2.</u> Let $V^{(AN)}$ denote the vector of conditional variances of N-period (simple) rates of return on aggregate wealth in each of the S states. Then[24]

$$V^{(AN)} = Q^N \iota_S - [(M^N \iota_S)*(M^N \iota_S)] \quad , \tag{5a}$$

where Q is an S×S matrix with (i, j) element

$$q_{ij} = \phi_{ij} E\{\lambda(i)^2\} \frac{(w_j + 1)^2}{w_i^2} \quad . \tag{6a}$$

Given the definition of levered equity and the relations in (15) and (17), the one-period rate of return on the levered equity, where the economy goes from state (c, i) to state (cλ(i), j), is given by

$$R^{(L1)} = \frac{\max[0,\ \lambda(i)(1 + w_j) - \theta w_i]}{w_i - g_i} - 1 \quad . \tag{7a}$$

As with the returns on aggregate wealth analyzed earlier, the conditional distribution of $R^{(L1)}$ depends only on i. Using an approach similar to that followed in developing the previous propositions, we obtain the conditional means and variances of returns on the levered equity for various investment

[24] The symbol "*" denotes a Hadamard matrix product. If $[a_{ij}]$ and $[b_{ij}]$ denote the elements of m×n matrices, then $[a_{ij}]*[b_{ij}] = [a_{ij} \cdot b_{ij}]$.

horizons.

Proposition A3. Let $E^{(LN)}$ denote the S-vector of conditional expected N-period (simple) rates of return on levered equity in each of the S states, $1 \leq s \leq S$.

$$E^{(LN)} = \Gamma^N \iota_S - \iota_S \quad , \tag{8a}$$

where Γ is an S×S matrix with (i, j) element

$$\gamma_{ij} = \frac{E\{\max[0, \lambda(i)(1 + w_j) - \theta w_i]\}}{w_i - g_i} \quad . \tag{9a}$$

Proposition A4. Let $V^{(LN)}$ denote the vector of conditional variances of N-period (simple) rates of return on levered equity in each of the S states. Then

$$V^{(LN)} = \Xi^N \iota_S - [(\Gamma^N \iota_S)*(\Gamma^N \iota_S)] \quad , \tag{10a}$$

where Ξ is an S×S matrix with (i, j) element

$$\xi_{ij} = \phi_{ij} \frac{E\{(\max[0, \lambda(i)(1 + w_j) - \theta w_i])^2\}}{(w_i - g_i)^2} \quad . \tag{11a}$$

We also examine the unconditional moments of returns on aggregate wealth and on levered equity. Let $\overline{E}^{(AN)}$, and $\overline{V}^{(AN)}$ denote the unconditional mean and variance of the N-period simple rate of return on aggregate wealth. Given the conditional moments provided in propositions A1 and A2, unconditional means and variances are obtained directly from the conditional moments using π, the vector of steady-state probabilities:

$$\bar{E}^{(AN)} = \pi' E^{(AN)} \tag{12a}$$

$$\bar{V}^{(AN)} = \pi' V^{(AN)} + \pi' [(E^{(AN)} - \bar{E}^{(AN)}\iota_S)*(E^{(AN)} - \bar{E}^{(AN)}\iota_S)] \tag{13a}$$

Following the notation defined earlier, let $R^{(A)}_{t,N}$ denote the N-period return on aggregate wealth starting at the beginning of period t. In obtaining the first-order autocorrelation of $R^{(A)}_{t,N}$, $\mathrm{corr}(R^{(A)}_{t,N}, R^{(A)}_{t-N,N})$, we note that

$$\mathrm{corr}(R^{(A)}_{t,N}, R^{(A)}_{t-N,N}) = \frac{E\{1 + R^{(A)}_{t,2N}\} - [E\{1 + R^{(A)}_{t,N}\}]^2}{\mathrm{var}\{R^{(A)}_{t,N}\}} ,$$

$$= \frac{\bar{E}^{(A2N)} - 2\bar{E}^{(AN)} - [\bar{E}^{(AN)}]^2}{\bar{V}^{(AN)}} \tag{14a}$$

The expressions for the unconditional moments of returns on levered equity are of precisely the same forms [simply replace "A" with "L" in (12a)-(14a)].

REFERENCES

Abel, Andrew B., 1988, "Stock Prices under Time-Varying Dividend Risk: An Exact Solution in an Infinite-Horizon General Equilibrium Model," Journal of Monetary Economics 22, 375-394.

Black, Fischer, 1988, "Mean Reversion and Consumption Smoothing," Working paper, Goldman, Sachs & Co.

Brown, David P. and Michael R. Gibbons, 1985, "A Simple Econometric Approach for Utility-Based Asset Pricing Models," Journal of Finance 40, 359-381.

Campbell, John Y., 1987, "Stock Returns and the Term Structure," Journal of Financial Economics 18, 373-399.

Campbell, John Y. and Robert J. Shiller, 1988, "Stock Prices, Earnings, And Expected Dividends," Working Paper, National Bureau of Economic Research.

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1988, "Mean Reversion in Equilibrium Asset Prices," Working paper, Ohio State University.

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1989, "The Equity Premium and the Risk Free Rate: Matching the Moments, Working paper, The Ohio State University.

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, "Economic Forces and the Stock Market," Journal of Business 59, 383-403.

Cochrane, John H., 1988, "Is Long Horizon Predictability of Stock Returns Due to Fads or Time-Varying Expected Returns?," Working Paper, Department of Economics, University of Chicago.

Constantinides, George M., 1988, "Habit Formation: A Resolution of the Equity Premium Puzzle," Working Paper, University of Chicago.

Fama, Eugene F., 1984, "The Information in the Term Structure," Journal of Financial Economics 13, 509-528.

Fama, Eugene F. and Robert R. Bliss, 1987, "The Information in Long-Maturity Forward Rates," American Economic Review 77, 680-692.

Fama, Eugene F. and Kenneth R. French, 1987, "Forecasting Returns on Corporate Bonds and Common Stocks," Working Paper, University of Chicago.

Fama, Eugene F. and Kenneth R. French, 1988, "Permanent and Temporary Components of Stock Prices," Journal of Political Economy 96, 246-273.

Fama, Eugene F. and G. William Schwert, 1977, "Asset Returns and Inflation," Journal of Financial Economics 5, 115-146.

Friend, Irwin and Marshall E. Blume, 1975, "The Demand for Risky Assets," American Economic Review 65, 900-922.

Grossman, S.J., A. Melino, R.J. Shiller, 1985, "Estimating the Continuous Time Consumption Based Asset Pricing Model," NBER Working Paper.

Granger, C. W. J. and Paul Newbold, 1977, Forecasting Economic Time Series, Academic Press, New York.

Hall, Robert E., 1981, "Intertemporal Substitution in Consumption," Working Paper, Stanford University.

Hall, Robert E., 1988, "Intertemporal Substitution in Consumption," Journal of Political Economy 96, 339-357.

Hansen, Lars Peter and Kenneth J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," Econometrica 50, 1269-1286 (with corrections in Econometrica 52, 267-268).

Hansen, Lars Peter and Kenneth J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," Journal of Political Economy 91, 249-265.

Hsieh, David A., 1983, "A Heteroscedasticity-Consistent Covariance Matrix Estimator for Time Series Regressions," Journal of Econometrics 22, 281-290.

Huberman, Gur and Shmuel Kandel, 1988, "Market Efficiency and Value Line's Record," Working Paper, Graduate School of Business, University of Chicago.

Huizinga, John and Frederick S. Mishkin, 1984, "Inflation and Real Interest Rates on Assets with Different Risk Characteristics," Journal of Finance 39, 699-712.

Kandel, Shmuel, and Robert F. Stambaugh, 1989, "Expectations and Volatility of Long-Horizon Stock Returns," Working Paper, University of Chicago and University of Pennsylvania.

Keim, Donald B. and Robert F. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," Journal of Financial Economics 17, 357-390.

Kocherlakota, Narayana R., 1988a, "Evaluating Policy Innovations Using Vector Autoregressions and Structural Methods," Working Paper, Northwestern University.

Kocherlakota, Narayana R., 1988b, "In Defense of the Time and State Separable Utility-Based Asset Pricing Model," Working Paper, Northwestern University.

Lo, Andrew W., 1988, "Long-Term Memory in Stock Market Prices," Working Paper, Massachusetts Institute of Technology.

Lo, Andrew W. and A. Craig Mackinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," Review of Financial Studies 1, 137-158.

Lucas, Robert E. Jr., 1978, "Asset Prices in an Exchange Economy," Econometrica 46, 1429-1445.

Mehra, Rajnish and Edward C. Prescott, 1985, "The Equity Premium: A Puzzle," Journal of Monetary Economics 15, 145-162.

Naik, Vasanttilak T. and Ehud I. Ronn, 1988, "The Impact of Time Aggregation and Sampling Interval on the Estimation of Relative Risk Aversion and the Ex Ante Real Interest Rate," Working Paper, University of Texas and University of British Columbia.

Poterba, James A. and Lawrence H. Summers, 1989, "Mean Reversion in Stock Prices: Evidence and Implications," Journal of Financial Economics 22, 27-59.

Rozeff, Michael S., 1984, "Dividend Yields are Equity Risk Premiums," Journal of Portfolio Management 10 (Fall), 68-75.

Richardson, Matthew, 1988, "Temporary Components of Stock Prices: A Skeptic's View," Working paper, Stanford University.

Sims, Christopher A., 1980, "Macroeconomics and Reality," Econometrica 48, 1-48.

Sims, Christopher A., 1981, "An Autoregressive Index Model for the U.S., 1948-1975," in Large-Scale Macro-Econometric Models, edited by J. Kmenta and J.B. Ramsey (North-Holland Publishing Co., Amsterdam).

Stambaugh, Robert F., 1986, "Discussion," Journal of Finance 41, 601-602.

Tauchen, George, 1986, "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," Economic Letters 20, 177-181.

White, Halbert, 1980, "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," Econometrica 48, 817-838.

Table 1

Markov Process for the Conditional Mean and Standard Deviation of the Monthly Consumption Growth Rate in the Example of the Equilibrium Model

State	Unconditional Probability	Conditional Mean of the Monthly Growth Rate (%)	Conditional Standard Deviation of the Monthly Growth Rate (%)
1	0.084	0.111	0.915
2	0.088	0.153	0.915
3	0.084	0.194	0.915
4	0.160	0.111	1.023
5	0.168	0.153	1.023
6	0.160	0.194	1.023
7	0.084	0.111	1.144
8	0.088	0.153	1.144
9	0.084	0.194	1.144

Probability of moving from state i to state

state i	1	2	3	4	5	6	7	8	9
1	0.459	0.041	0.000	0.397	0.036	0.000	0.062	0.006	0.000
2	0.039	0.422	0.039	0.034	0.365	0.034	0.005	0.057	0.005
3	0.000	0.041	0.459	0.000	0.036	0.397	0.000	0.006	0.062
4	0.209	0.019	0.000	0.500	0.045	0.000	0.209	0.019	0.000
5	0.018	0.192	0.018	0.042	0.460	0.042	0.018	0.192	0.018
6	0.000	0.019	0.209	0.000	0.045	0.500	0.000	0.019	0.209
7	0.062	0.006	0.000	0.397	0.036	0.000	0.459	0.041	0.000
8	0.005	0.057	0.005	0.034	0.365	0.034	0.039	0.422	0.039
9	0.000	0.006	0.062	0.000	0.036	0.397	0.000	0.041	0.459

Table 2

Unconditional Moments of Various Monthly Series in the Example of the Equilibrium Model

Quantity	Mean[a]	Standard Deviation[b]	First-Order Autocorrelation
A. Consumption Growth Rates (exogenous)			
Actual Growth Rate	1.8300	3.5684	0.0010
Conditional Expected Growth Rate	1.8300	0.1163	0.9178
Conditional Standard Deviation of the Growth Rate	3.5551[c]	0.2852	0.4320
B. Rates of Return (endogenous)			
Riskless Rate	0.7907	4.0964	0.7509
Return on Aggregate Wealth	4.4262	16.6290	-0.0544
Return on Levered Equity	6.9865	27.7465	-0.0439
Conditional Expected Return on Levered Equity	6.9865	3.9415	0.7776
Conditional Standard Deviation of Return on Levered Equity	27.0820[c]	4.5712	0.7764
C. Other Financial Variables (endogenous)			
Dividend-Price Ratio	2.0886	0.0630	0.9155
Short-Term Default Spread	5.3292	1.0810	0.8681
Term-Structure Spread	-1.3336	3.9515	0.7410

[a]All numbers in this column, except those noted otherwise, are multiplied by 1200.

[b]The numbers in this column are multiplied by $100 \times \sqrt{12}$.

[c]Multiplied by $100 \times \sqrt{12}$.

Table 3

Sample Means, Standard Deviations, and Correlations

Variable	Mean	Standard Deviation	Correlations $y_{Aaa}-y_{TB}$	D/P	r
$y_{Baa}-y_{Aaa}$	0.1007	0.0676	0.58	0.08	0.03
$y_{Aaa}-y_{TB}$	0.1839	0.1077		-0.04	0.10
D/P	0.0402	0.0133			-0.19
r	0.0078	0.0754			

Note: The variables are defined as follows.

$(y_{Baa} - y_{Aaa})_t$: the difference at the end of month t between Moody's average yield on bonds rated Baa and bonds rated Aaa.

$(y_{Aaa} - y_{TB})_t$: the difference at the end of month t between the Aaa yield and the yield on a U.S. Treasury Bill with maturity closest to one month.

$(D/P)_t$: for the equally weighted portfolio of NYSE stocks, the ratio of dividends paid for the twelve months ending at t to the price at the end of month t.

r_t : the continuously compounded real return in month t on the equally weighted portfolio of NYSE stocks.

Table 4

Estimates of the VAR's Parameters

Dependent Variable[a]	Intercept	$y_{Baa}-y_{Aaa}$	$y_{Aaa}-y_{TB}$	D/P	r	adj. R^2	$\chi^{2(c)}$	ρ_1	ρ_2	ρ_3	ρ_4	ρ_5	ρ_6
	Independent variables (lagged one month)[b]							Residual Autocorrelations					
$y_{Baa}-y_{Aaa}$	0.0006 (0.0031)	0.9706 (0.0309)	0.0076 (0.0090)	0.0423 (0.0437)	-0.0828 (0.0176)	0.96	3983.0 (0.000)	0.09	-0.09	-0.26	-0.13	0.12	0.10
$y_{Aaa}-y_{TB}$	0.0077 (0.0059)	0.1303 (0.0421)	0.8749 (0.0369)	0.0744 (0.0804)	-0.0592 (0.0225)	0.85	3460.3 (0.000)	-0.11	0.06	-0.06	-0.02	-0.04	0.09
D/P	0.0024 (0.0008)	-0.0070 (0.0055)	-0.0012 (0.0013)	0.9633 (0.0129)	-0.0042 (0.0028)	0.93	9045.3 (0.000)	0.02	-0.00	-0.08	0.03	0.12	0.05
r	-0.0245 (0.0134)	0.0975 (0.1073)	0.0295 (0.0311)	0.3937 (0.2466)	0.1605 (0.0714)	0.04	12.3 (0.016)	-0.00	0.02	-0.08	-0.01	0.07	-0.02

[a]The variables are defined as follows.

$(y_{Baa} - y_{Aaa})_t$: the difference at the end of month t between Moody's average yield on bonds rated Baa and bonds rated Aaa.

$(y_{Aaa} - y_{TB})_t$: the difference at the end of month t between the Aaa yield and the yield on a U.S. Treasury Bill with maturity closest to one month.

$(D/P)_t$: for the equally weighted portfolio of NYSE stocks, the ratio of dividends paid for the twelve months ending at t to the price at the end of month t.

r_t : the real return in month t on the equally weighted portfolio of NYSE stocks.

[b]The coefficients are estimated using ordinary least squares and the standard errors (in parentheses) are based on the heteroskedasticity-consistent estimator of the covariance matrix of White (1980) and Hsieh (1983).

[c]The statistic reported is asymptotically distributed as χ^2 with five degrees of freedom under the null hypothesis that all of the coefficients on the independent variables (excluding the intercept) are equal to zero. The p-value is shown in parentheses.

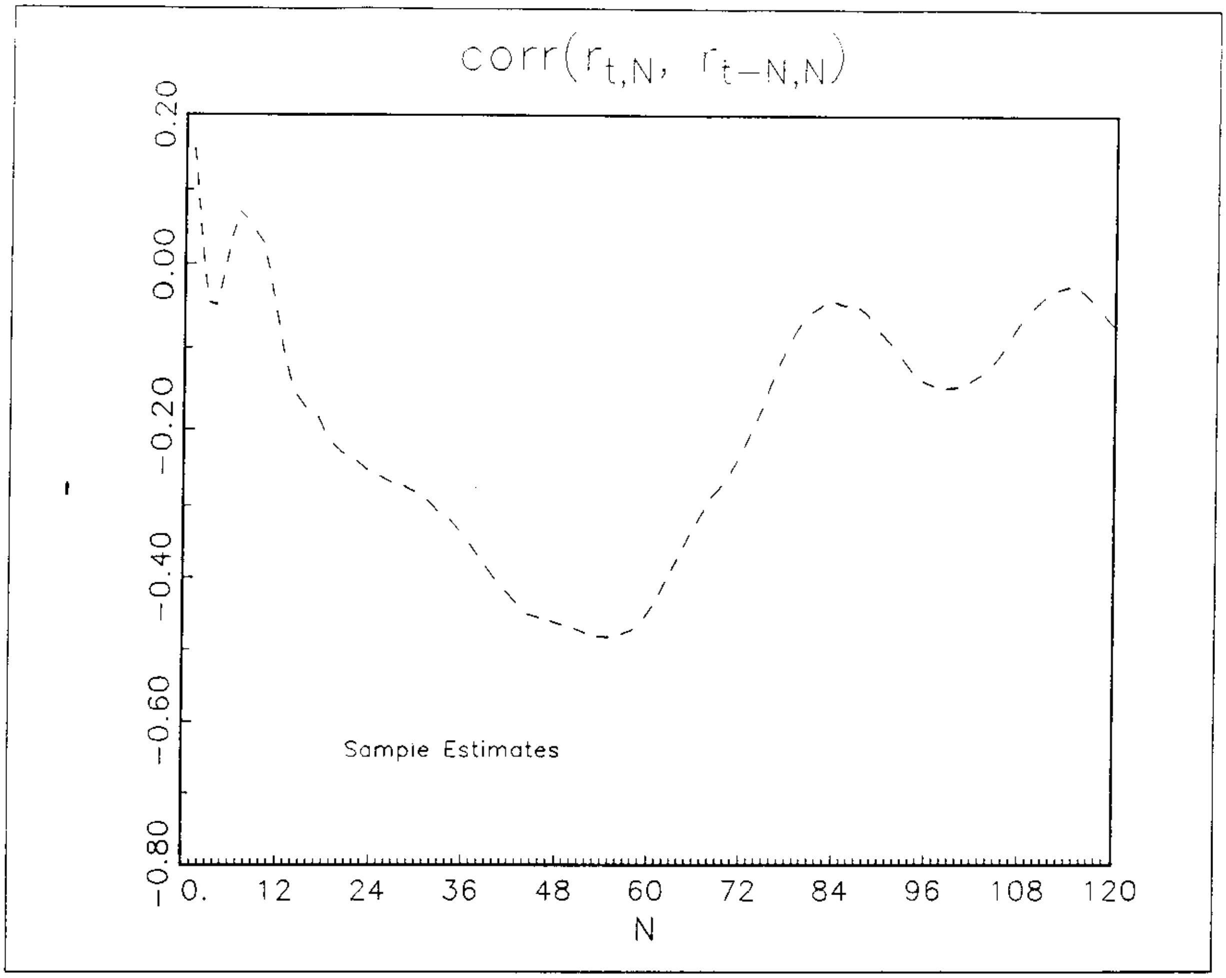

Figure 1. Sample estimates of first-order autocorrelations of N-month real returns (continuously compounded) on the equally weighted NYSE portfolio. The estimates are obtained by regressing the N-month return on its lagged value, using monthly observations with overlapping return horizons.

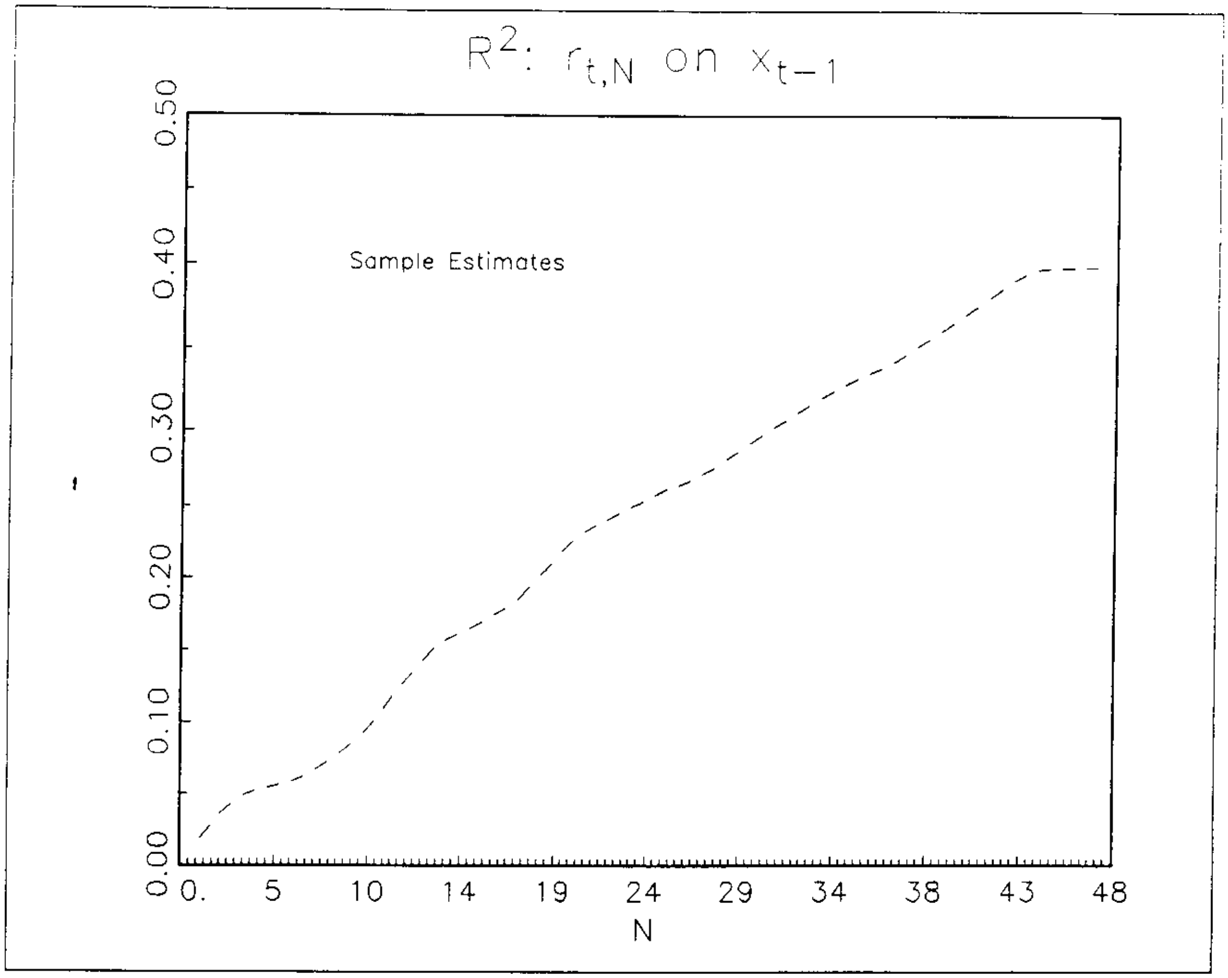

Figure 2. R-squared values in regressions of N-month real returns (continuously compounded) on the equally weighted NYSE portfolio on three predictive variables (the Baa yield minus the Aaa yield, the Aaa yield minus the T-Bill yield, and the dividend-price ratio). For each N, the R-squared value is obtained in a regression using monthly observations with overlapping return horizons.

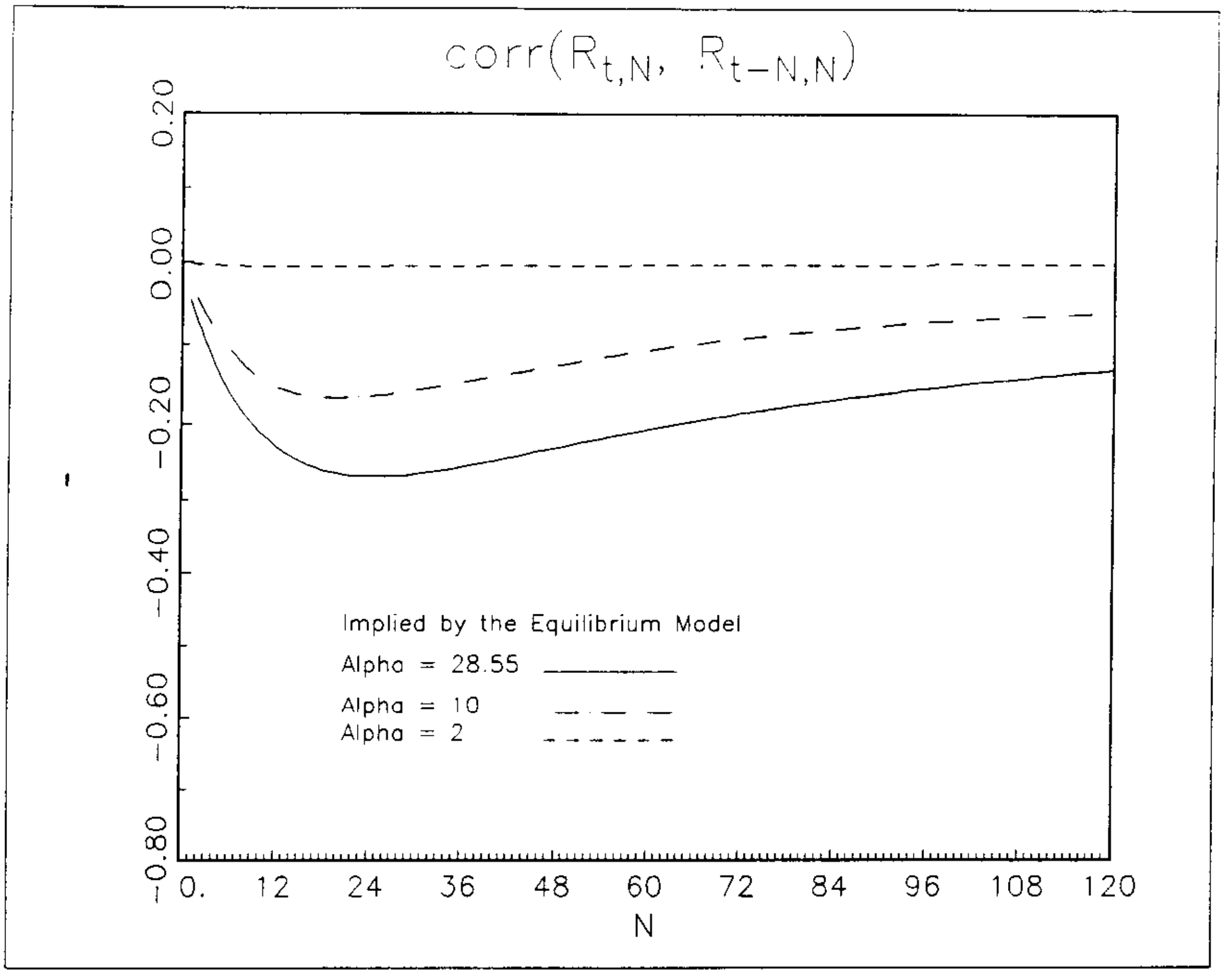

Figure 3. First-order autocorrelations of N-month returns on levered equity implied by the equilibrium model. The solid line represents values implied by the model using the parameters of the numerical example of section 3. The long-dashed and the short-dashed lines display the values for this example where the risk aversion coefficient is changed to 10 and 2, respectively.

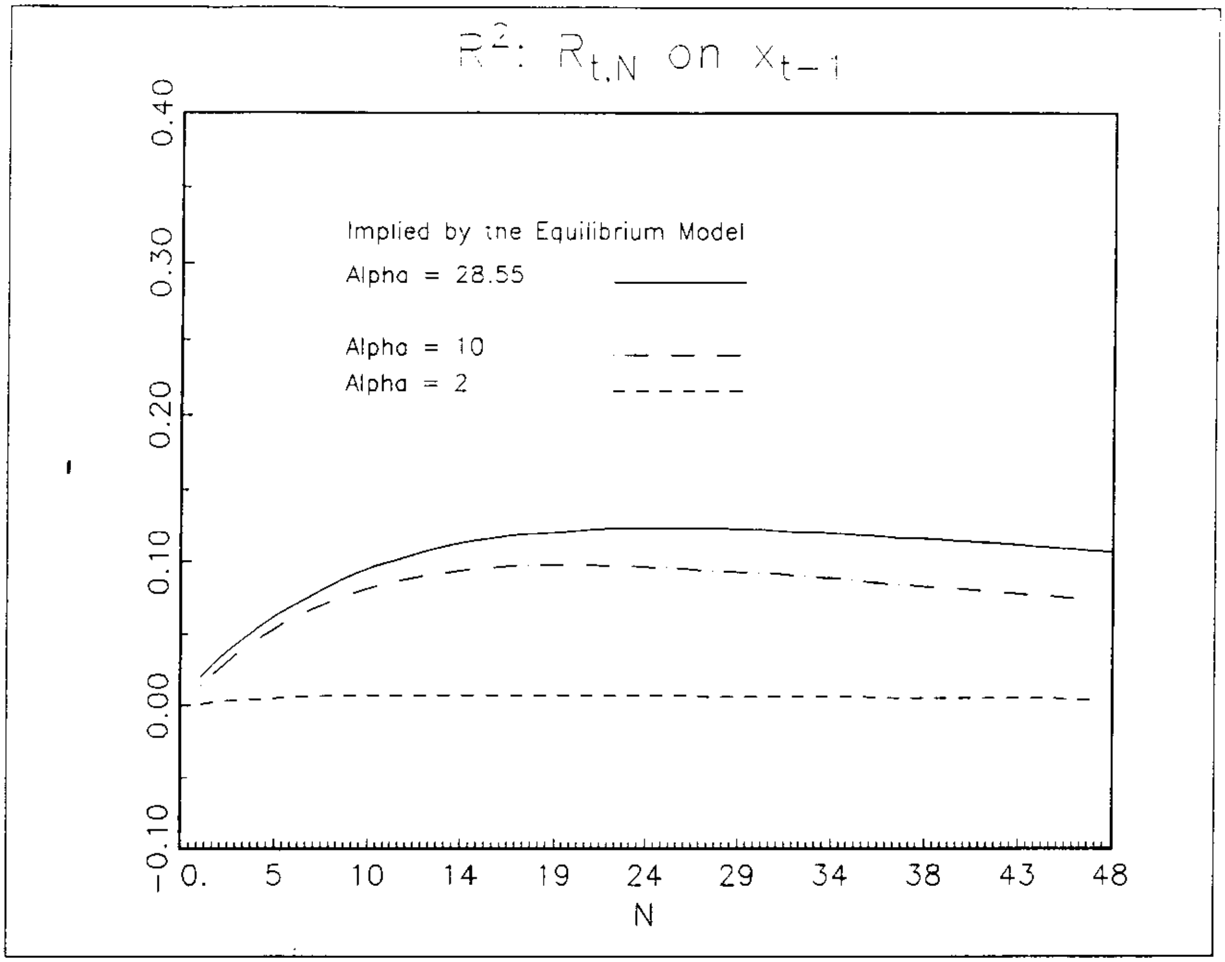

Figure 4. R-squared values in regressions of N-month returns on levered equity on three predictive variables (a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread) implied by the equilibrium model. The solid line displays values implied by the model using the parameters of the numerical example of section 3. The long-dashed and the short-dashed lines display the values for this example where the risk aversion coefficient is changed to 10 and 2, respectively.

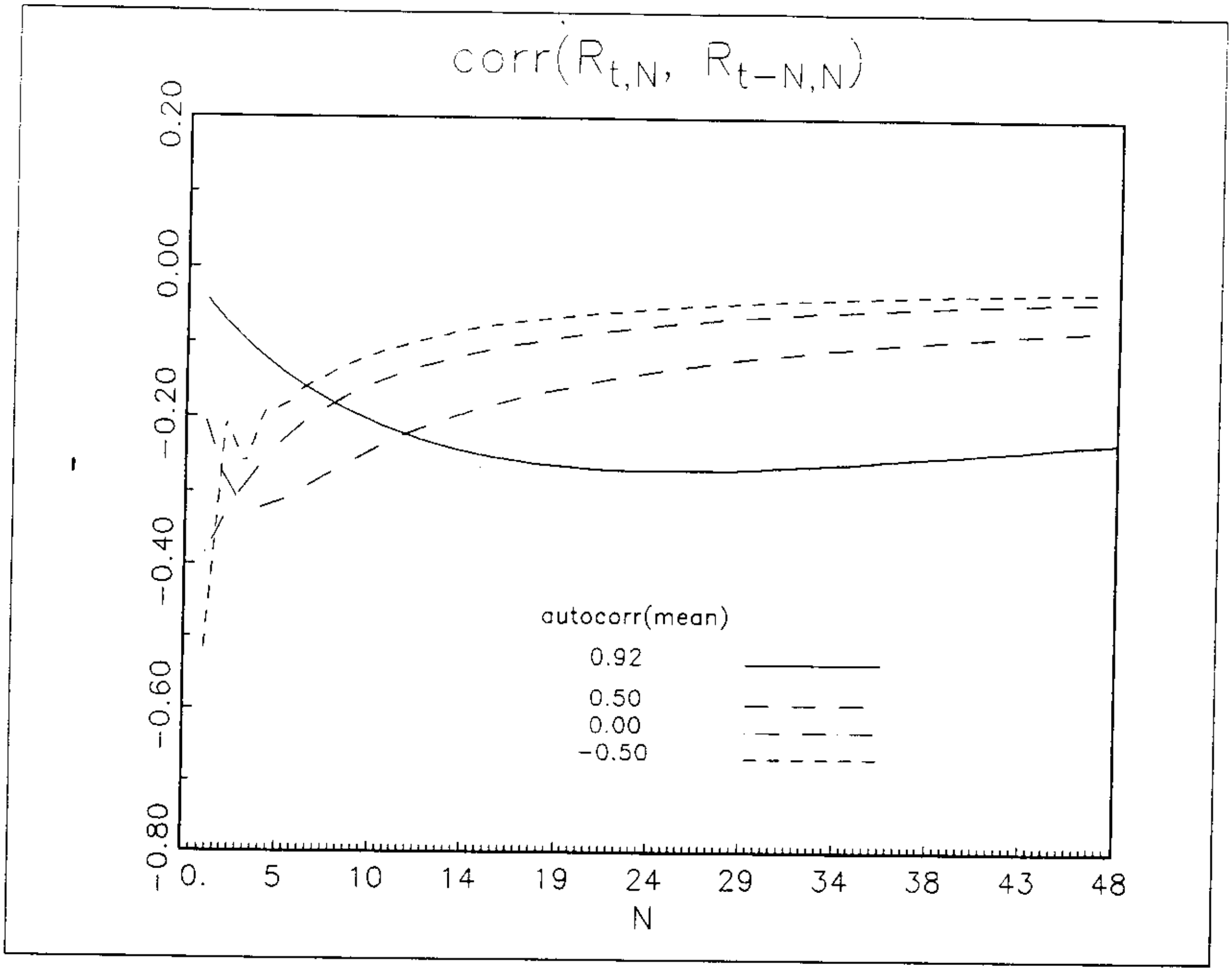

Figure 5. First-order autocorrelations of N-month returns on levered equity implied by the equilibrium model. The solid line represents values implied by the model using the parameters of the numerical example of section 3. The other lines are obtained using alternative specifications of the first-order autocorrelation of the monthly growth rate in consumption.

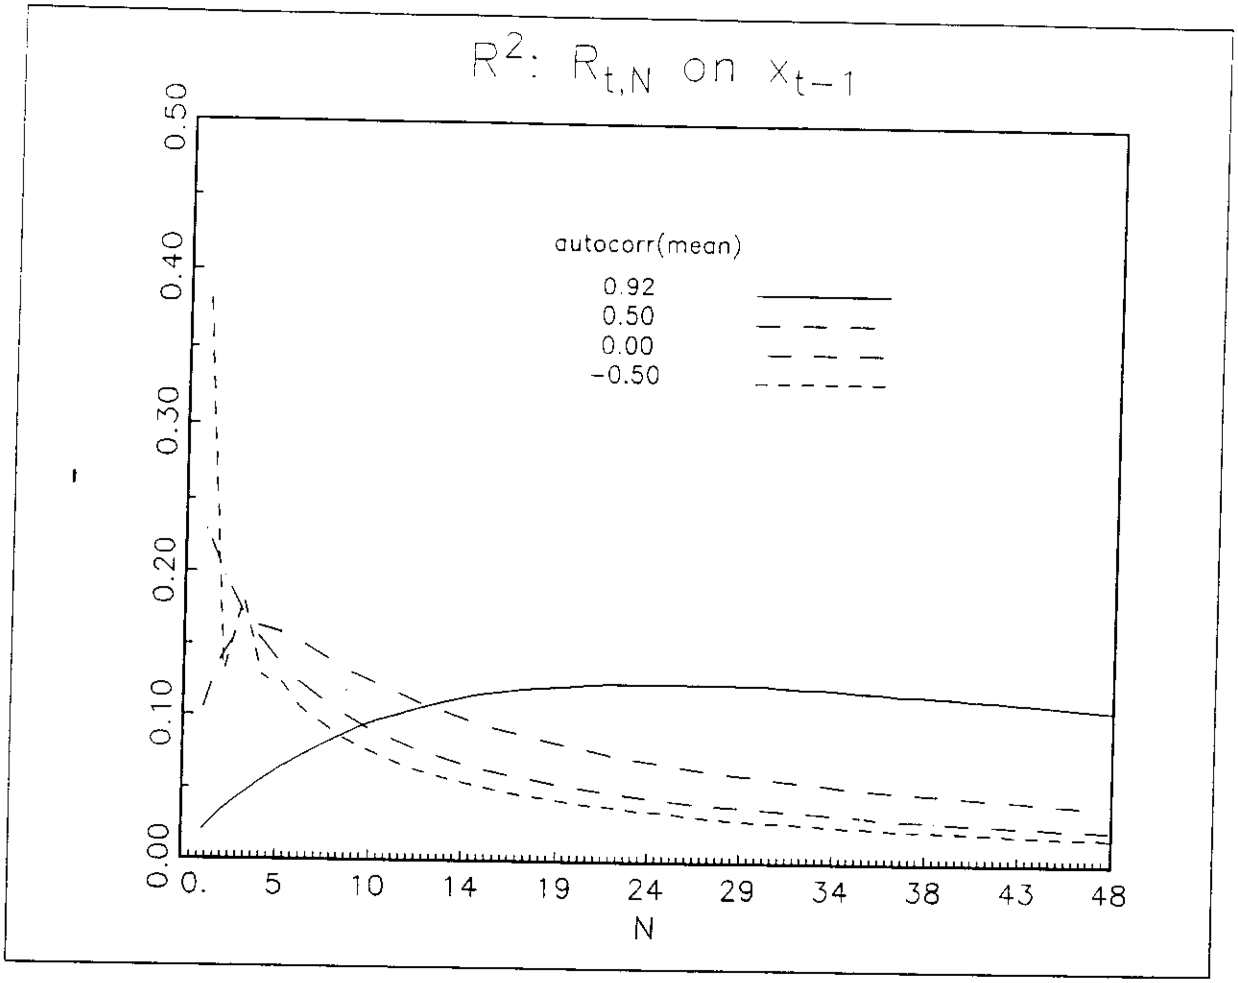

Figure 6. R-squared values in regressions of N-month returns on levered equity on three predictive variables (a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread) implied by the equilibrium model. The solid line displays values implied by the model using the parameters of the numerical example of section 3. The other lines are obtained using alternative specifications of the first-order autocorrelation of the monthly growth rate in consumption.

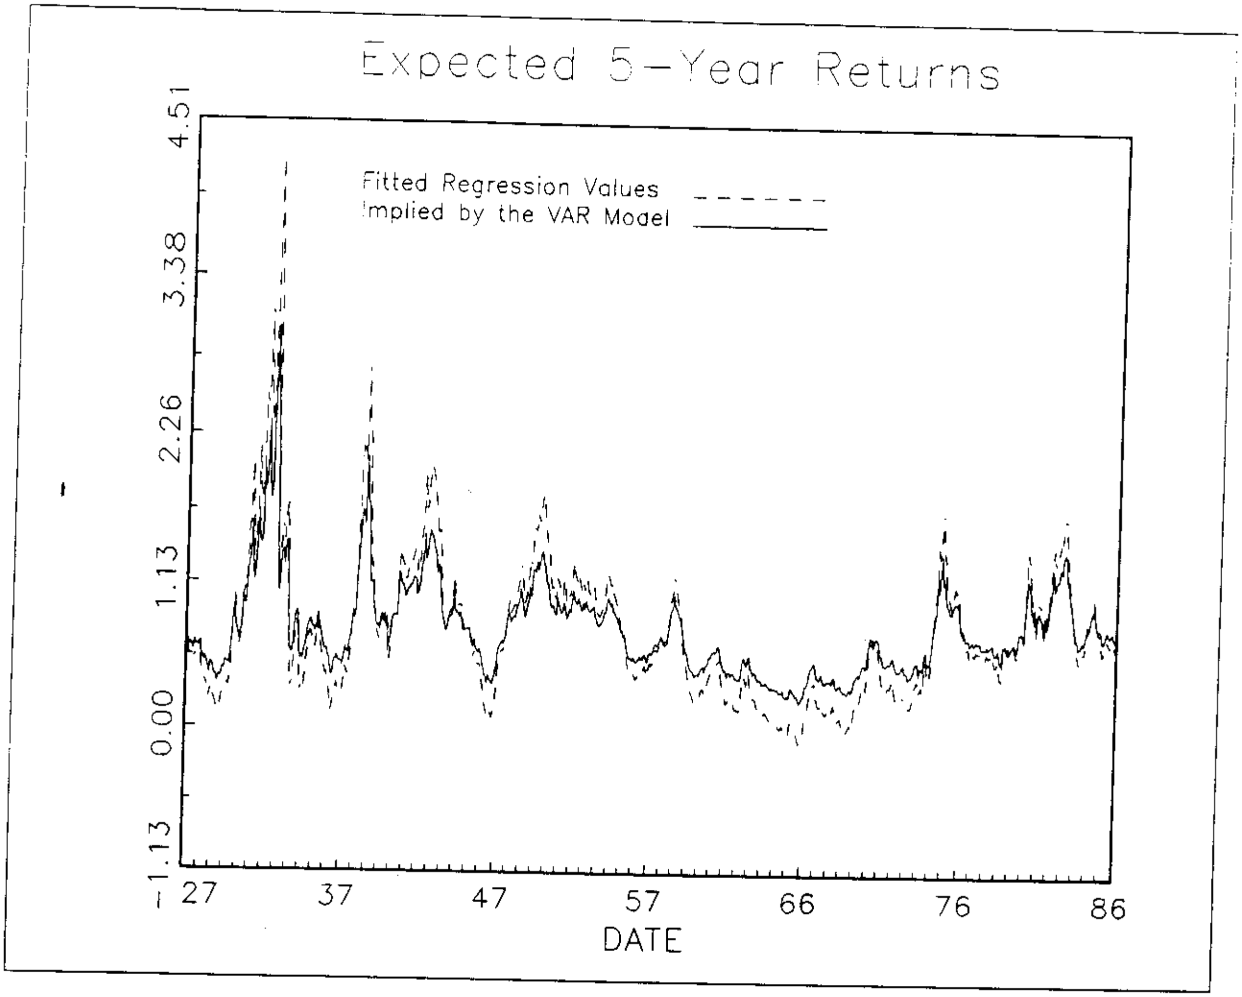

Figure 7. Estimated expected real returns (continuously compounded) on the equally-weighted NYSE portfolio for five-year horizons. The value plotted corresponds to the expected monthly return for the five-year horizon beginning on the given date. The solid line represents expected five-year returns implied by the VAR model, and the dashed line represents expected returns estimated directly in a regression with five-year returns as the dependent variable.

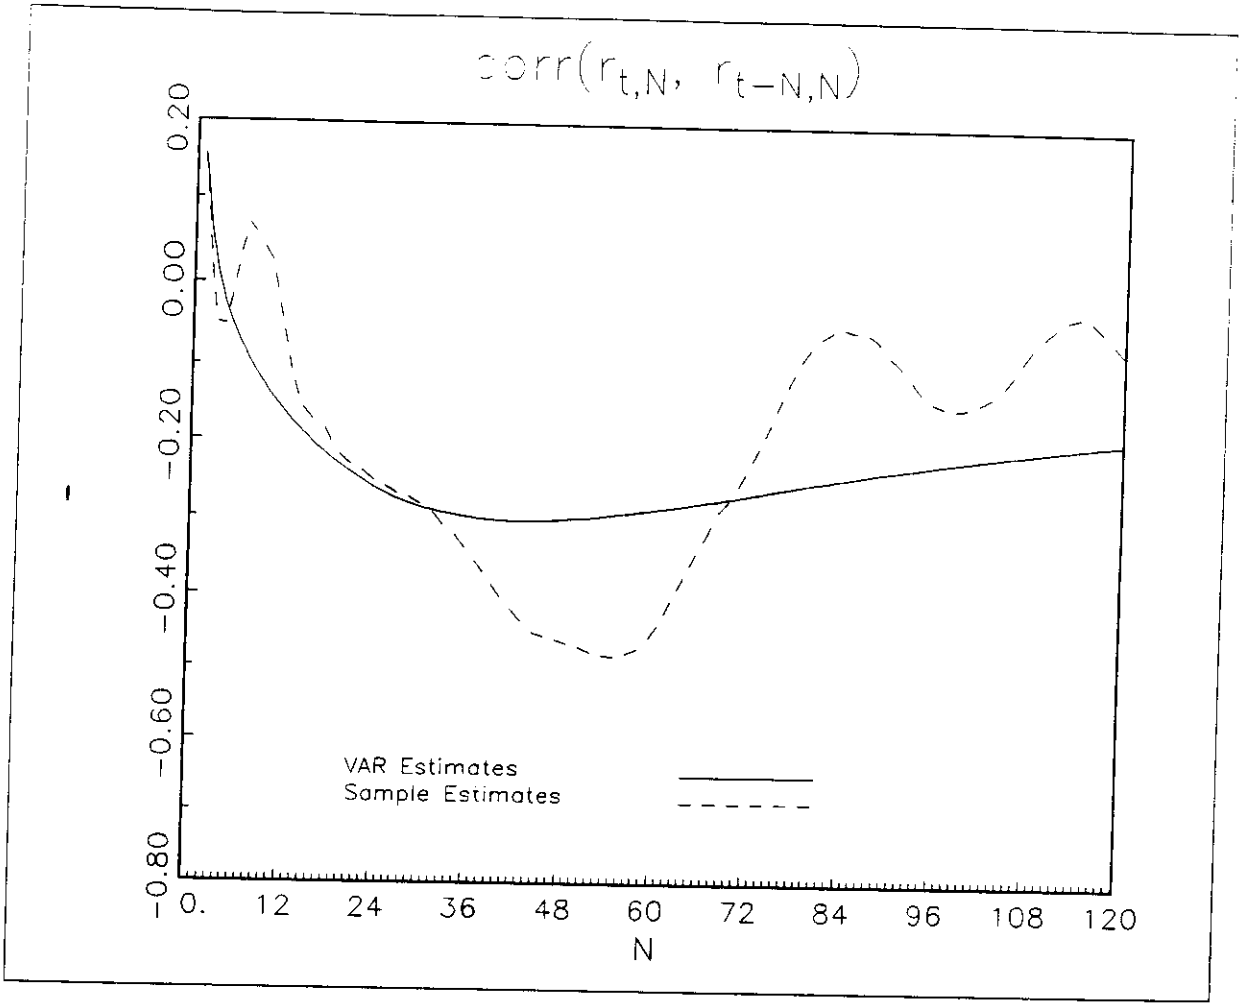

Figure 8. First-order autocorrelations of N-month real returns (continuously compounded) on the equally weighted NYSE portfolio. The solid line displays values implied by the VAR model. The dashed line displays sample estimates obtained by regressing the N-month return on its lagged value, using monthly observations with overlapping return horizons.

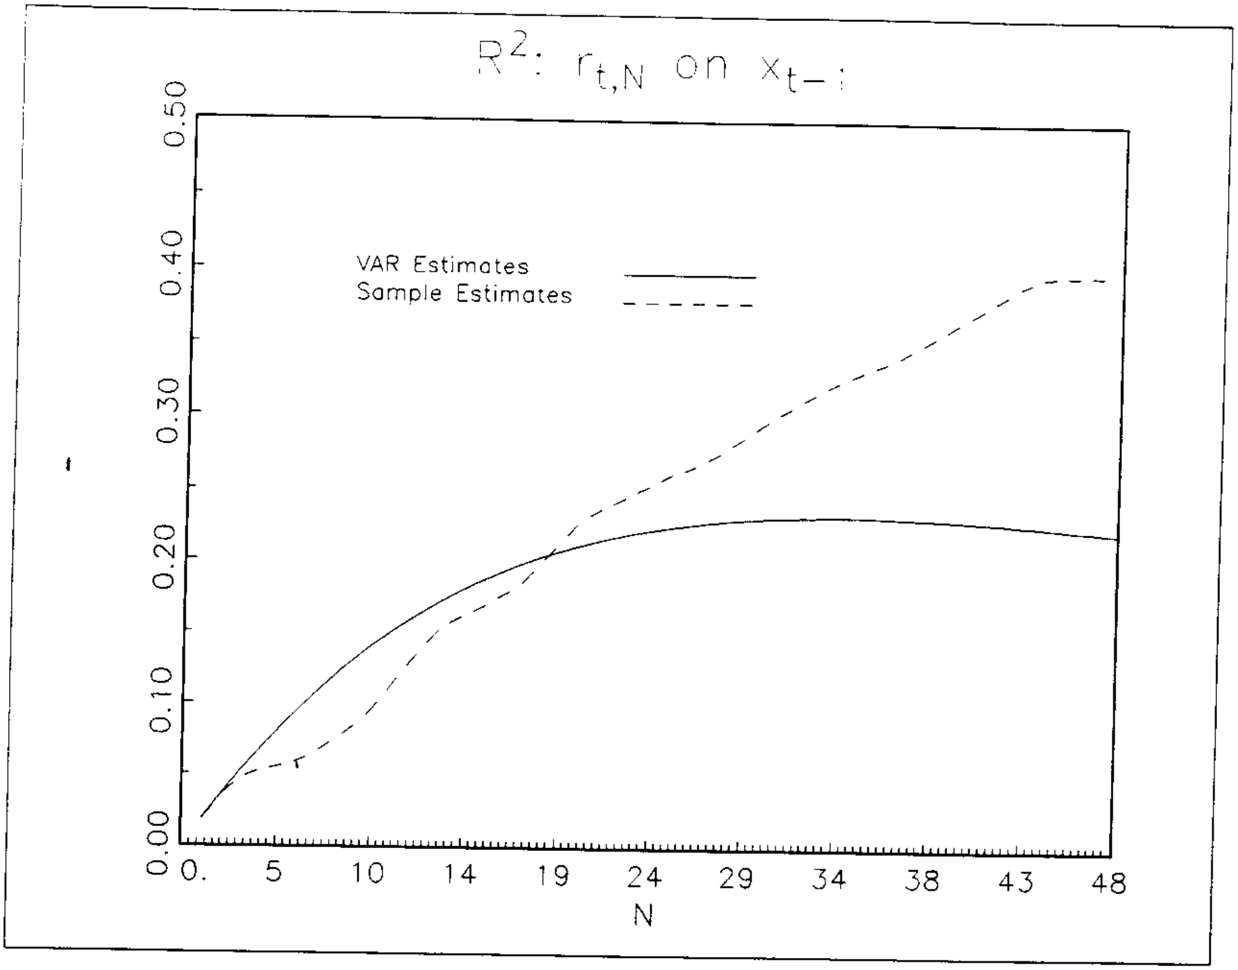

Figure 9. R-squared values in regressions of N-month real returns (continuously compounded) on the equally weighted NYSE portfolio on three predictive variables (the change in the Baa yield, the Baa yield minus the Aaa yield, the Aaa yield minus the T-Bill yield, and the dividend-price ratio). The solid line displays values implied by the VAR model. The R-squared value is equal to the ratio of the implied variance of the expected N-month return to the implied variance of the total N-month return. The dashed line displays sample values obtained in a regression of the N-month return on the three variables, using monthly observations with overlapping return horizons.

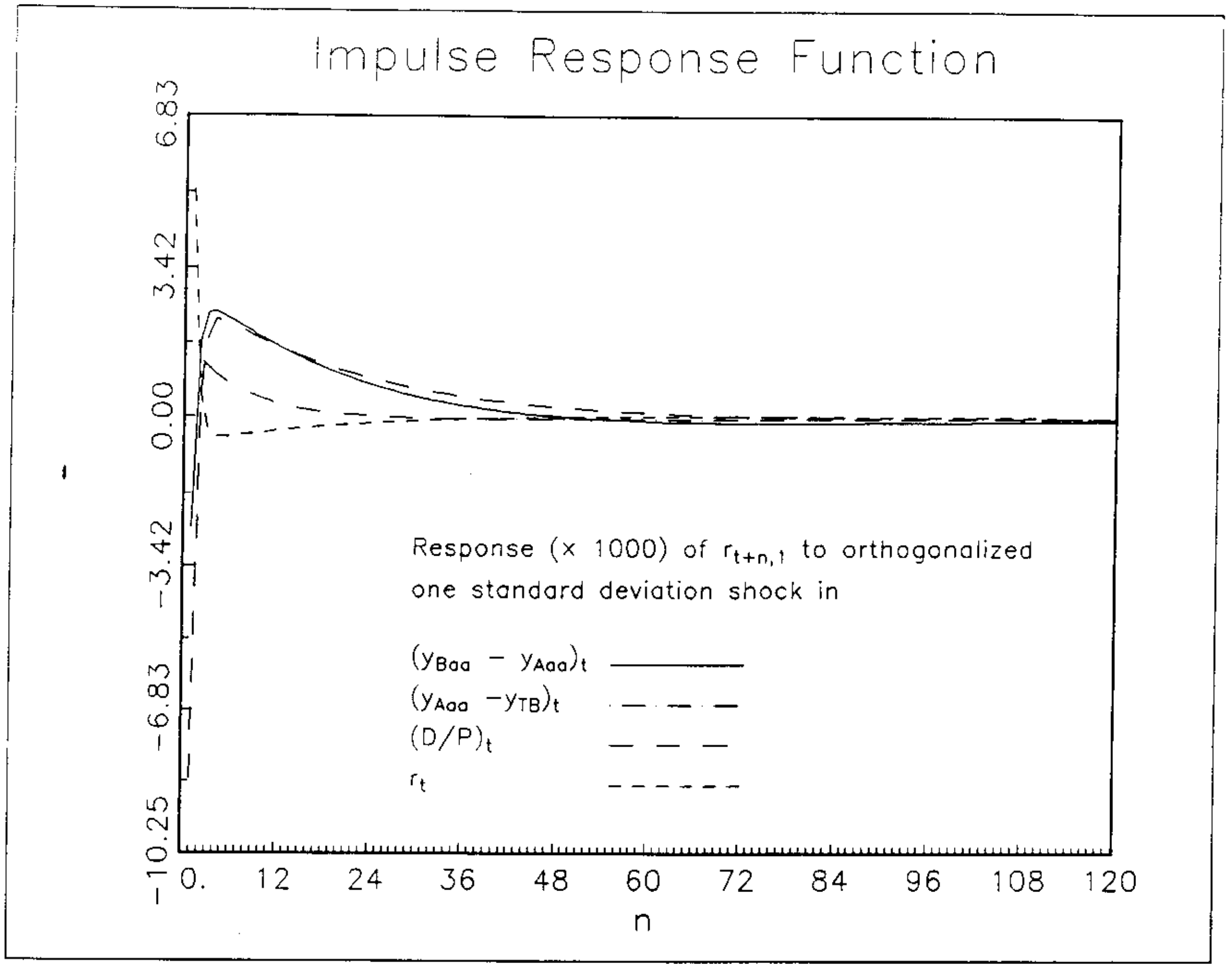

Figure 10. Responses of one-month real returns (continuously compounded) on the equally weighted NYSE portfolio in month t+n to shocks in the predictive variables in month t, as implied by the VAR model.

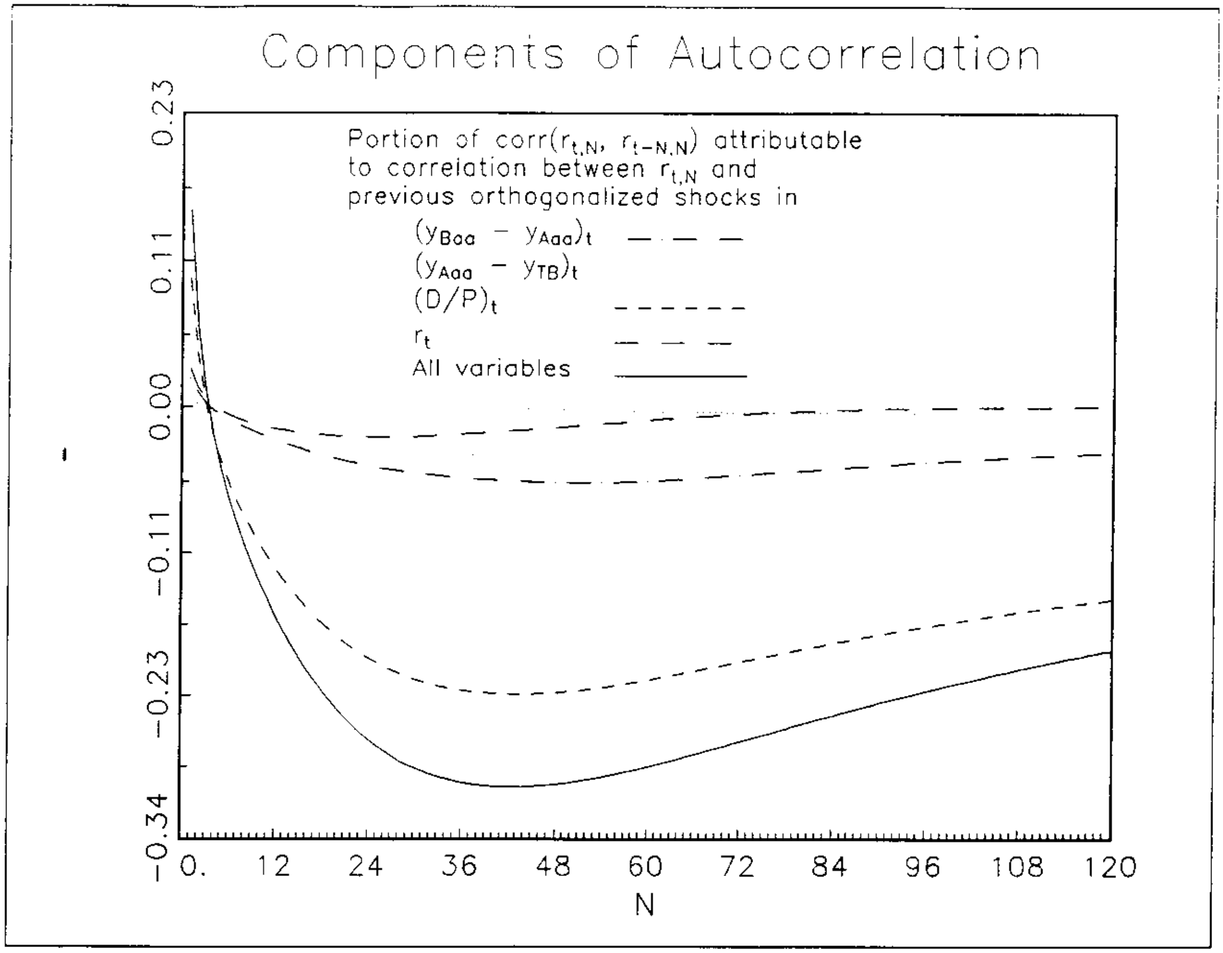

Figure 11. Components of the first-order autocorrelation of N-month real returns (continuously compounded) on the equally weighted NYSE portfolio, as implied by the unrestricted VAR model.

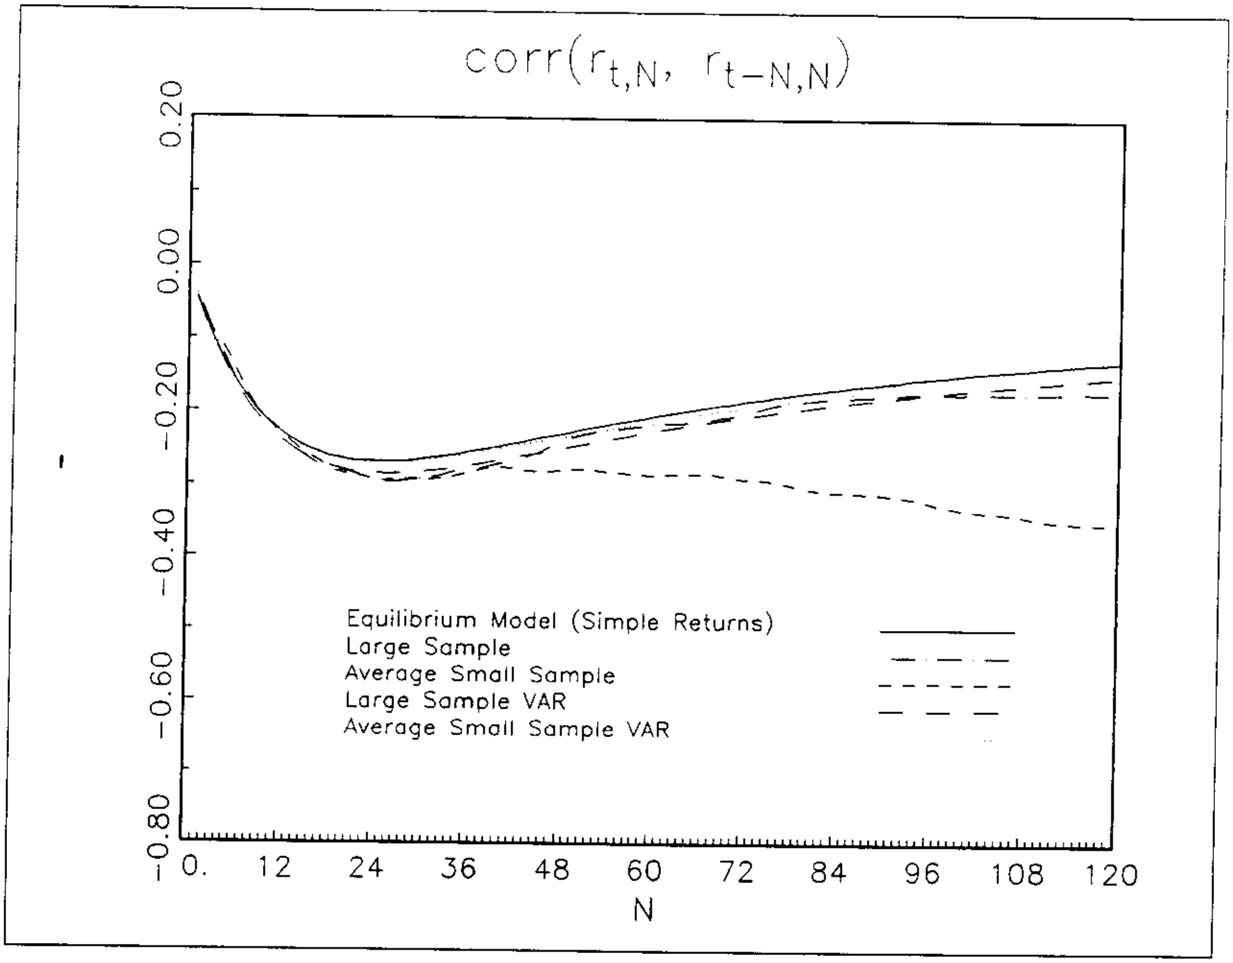

Figure 12. First-order autocorrelations of N-month returns on levered equity. The solid line displays values for simple returns implied by the equilibrium model with the parameters of the example in section 3. All other lines display values for continuously compounded returns based on data generated by simulating the equilibrium model with these parameters. The line with dots and dashes displays estimates obtained from a simulated data set with 70900 observations. The line with short dashs displays the average values of 100 simulated data sets with 709 observations each. For these two lines, each value is obtained by regressing the N-month return on its lagged value, using overlapping return horizons. The dashed line displays values implied by a VAR model estimated from a simulated data set with 70900 observations. The dotted line displays the average (over 100 data sets) values implied by a VAR model estimated from a simulated data set with 709 observations.

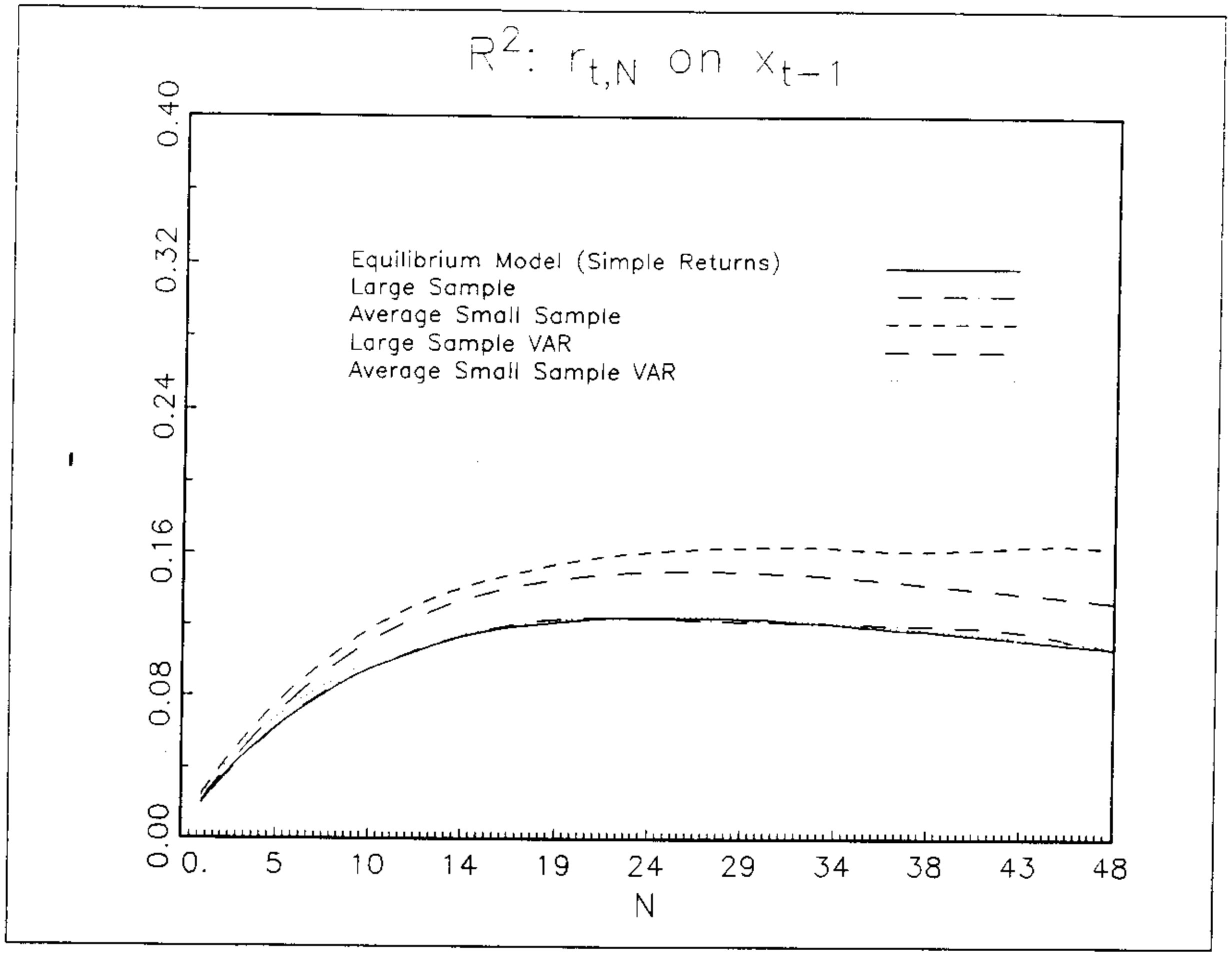

Figure 13. R-squared values in regressions of N-month returns on levered equity on three predictive variables (a dividend-price ratio, a low-grade-versus-high-grade yield spread, and a short-term-versus-long-term yield spread). The solid line displays values implied for simple returns by the equilibrium model with the parameters of the numerical example of section 3. All other lines display values for continuously compounded returns based on data generated by simulating the equilibrium model with these parameters. The line with dots and dashes displays estimates obtained from a simulated data set with 70900 observations. The line with short dashes displays the average values of 100 simulated data sets with 709 observations each. For these two lines, each value is obtained by regressing the N-month return on its lagged value, using overlapping return horizons. The dashed line displays values implied by a VAR model estimated from a simulated data set with 70900 observations. The dotted line displays the average (over 100 data sets) values implied by a VAR model estimated from a simulated data set with 709 observations. The VAR-implied R-squared value is equal to the ratio of the implied variance of the expected N-month return to the implied variance of the total N-month return.