AN INVESTIGATION OF THE RISK AND RETURN RELATION AT LONG HORIZONS

Paul Harrison and Harold H. Zhang*

Abstract—This paper examines the relation between expected stock returns and their conditional volatility over different holding periods and across different states of the economy. Seminonparametric density estimation and Monte Carlo integration are used to obtain the expected returns and conditional volatility at various holding intervals. We uncover a significantly positive risk and return relation at long holding intervals, such as one and two years, which is nonexistent at short holding periods such as one month. We also show that the existing finding in the literature of a negative risk and return relation may be attributable to misspecification.

I. Introduction

EXISTING empirical work on the time series risk and return relation has drawn conflicting conclusions. Campbell and Hentschel (1992) and French et al. (1987) find the expected excess return positively related to its conditional variance. On the other hand, Breen et al. (1989), Campbell (1987), Fama and Schwert (1977), Glosten et al. (1993), Nelson (1991), and Pagan and Hong (1991) all report a negative relationship between the expected excess return and conditional volatility.[1] Whitelaw (1994) finds no time-invariant risk and return relation, but that the conditional variance leads the expected return.

We extend the current literature on risk and return in three important directions. First, we examine the risk and return relation at various holding intervals longer than the sampling interval of data. Analyzing the risk and return relation at longer horizons may yield sharper results given the empirical evidence of greater returns predictability at longer horizons. At short horizons, the true long-run risk and return relation could be obscured by short-term noise, which might derive, for instance, from agents trading for portfolio rebalance and/or unexpected immediate consumption need reasons.

Second, we employ a seminonparametric estimation of the joint density of information variables and then use simulation and Monte Carlo integration to form the conditional mean and conditional variance of returns. This technique, which differs from those used in previous studies of returns and volatility, is desirable for a number of reasons.[2] Firstly, our estimation imposes no predetermined functional form on the relation between the conditional mean and volatility of returns nor on the relation between the dependent variable and its conditioning information. This is important when consistent estimates of the expected returns and conditional volatility require correct specification of the underlying data-generating process of the returns. In fact, the mixed findings on the risk and return relation may reflect various misspecifications.[3] Pagan and Hong (1991) and Harvey (1991) both find nonparametric estimation of returns to be preferred to parametric alternatives, while Glosten et al. (1993) and Pagan and Hong (1991) both argue that standard ARCH models are misspecified in this context and can lead to wrong conclusions. Secondly, our simulation technique allows us to examine the behavior of expected returns and volatility at holding periods longer than our monthly data frequency. Consequently, we can get sharper results on the risk and return relation because the high-frequency factors are less likely to exert significant impacts.

Third, we introduce the asset-pricing model of Campbell (1996) to provide some simple analytical examples to explore our empirical results.[4] In Campbell's framework, the conditional covariance between returns and the expected present value of future returns is a priced factor. In many parametric studies, it could be an omitted variable that would obscure a positive risk-return relation. In our study, changes in this covariance over the various holding periods could explain the differences in the long- and short-horizon results.

The cyclicality of expected returns and conditional volatility is first determined individually in regressions with proxies for the state of the economy. We then explore the risk and return relation using scatterplots for each holding interval. Campbell's framework is then used to suggest an explanation for previous findings of a negative risk and return relation. Finally, we explicitly test for cyclical variation in the Sharpe ratio as one alternative to the null hypothesis of a time-invariant risk and return relation.

We find a significant positive relation between the expected returns and conditional volatility at long holding intervals. This result illustrates that different holding periods

Received for publication September 24, 1997. Revision accepted for publication October 9, 1998.

* Brandeis University and Federal Reserve Board, and Carnegie Mellon University, respectively.

We thank John Campbell (the editor), Greg Duffee, Ronald Gallant, Eric Ghysels, Rick Green, Narayana Kocherlakota, Arthur Lewbel, David Marshall, Bryan Routledge, Akhtar Siddique, Norman Swanson, and two anonymous referees for helpful comments, and Huamao Bai for research assistance. We also appreciate the many useful comments we received from presenting this paper at Brandeis University, Carnegie Mellon University, Penn State University, the Federal Reserve Board of Governors, the 1996 North American Econometric Society Summer Meetings in Iowa City, the Third International Conference on Computation in Economics and Finance in Stanford, the 1997 Far Eastern Econometric Society Meetings in Hong Kong, the 1998 Econometric Society Winter Meetings in Chicago, and the 1998 Society for Economic Dynamics Annual Meetings. We, of course, are responsible for remaining errors. Address correspondence to huibingz@andrew.cmu.edu, (412) 268-8836 (Harold H. Zhang) or pharrison@frb.gov (Paul Harrison).

[1] Gallant et al. (1992) and Turner et al. (1989) report mixed results on the risk-return relation. Chan et al. (1992) find no significant variance effect on the returns for the U.S. stock market, but do find one for the world market portfolio.

[2] The seminonparametric method is used in Gallant et al. (1992) in their examination of the dynamics of stock prices and volume.

[3] Most of the studies on the risk and return relation employ the GARCH-M specification and estimate the model using data at sampling frequency such as weekly or monthly.

[4] We thank a referee for suggesting this investigation.

The Review of Economics and Statistics, August 1999, 81(3): 399–408

can lead to significantly different results since, at shorter intervals (such as monthly), no meaningful risk and return relation emerges. This long-horizon positive risk and return relation is nonlinear and strongly resembles the mean-variance efficient frontier. We also show that this risk and return relation itself is time varying, since the Sharpe ratio fluctuates significantly with the business cycle.

The paper continues with section II which describes the formation of the expected holding period returns and conditional volatility. Section III describes the implementation of the seminonparametric method. Section IV presents the empirical results, and section V concludes.

II. Expected Holding Returns and Conditional Volatility

Our investigation of the risk and return relation begins with the computation of the expected market returns and the conditional variance of market returns at various holding intervals assuming that we know the density function for a vector of information variables of which the return is a function. Later, in section III, we discuss the estimation of the density function.

Let $r(t, t + \tau)$ denote the holding returns from time t to $t + \tau$. The time t expected holding return is defined as follows:[5]

$$E_t[r(t, t + \tau)] = E_t\left\{\sum_{i=1}^{\tau} \log\left[1 + r(t-1+i, t+i)\right]\right\}. \tag{1}$$

Let y_{t+1} be a vector of information variables of which the return is a function, so that $r(t, t + 1) = g(y'_{t+1}, \ldots, y'_0)$, and let $x_t = (y'_0, y'_1, \ldots, y'_t)'$. Denote $f(y_{t+1}|x_t)$ as the joint conditional density function. Given $f(\cdot|\cdot)$ and the conditioning set x_t, the expected holding return, $E[r(t, t + \tau)|x_t]$, is given by the following integration:

$$\int \cdots \int \left\{\sum_{i=1}^{\tau} \log\left[1 + g(y'_{t+i}, \ldots, y'_0)\right]\right\} \times \left[\prod_{i=1}^{t+\tau} f(y_i|y_0, \ldots, y_{i-1})\right] dy_{t+1} \ldots dy_{t+\tau}. \tag{2}$$

If the information set evolves according to a Markov process depending on the most recent history with L lags, the expected holding returns conditioning on $x_t = x$ can be rewritten as follows:

$$\int \cdots \int \left\{\sum_{i=1}^{\tau} \log\left[1 + g(y'_i, \ldots, y'_{i-L+1})\right]\right\} \times \left[\prod_{i=1}^{\tau} f(y_i|y_{-L+i}, \ldots, y_{i-1})\right] dy_1 \ldots dy_\tau \tag{3}$$

where $x = (y'_{-L+1}, \ldots, y'_0)$. The above expectation can be approximated using Monte Carlo integration. Let $\{y_j^r\}_{j=1}^{\infty}$, $r = 1, 2, \ldots, R$, denote R simulated realizations of the process starting from x. In other words, y_1^r is a random drawing from $f(y|x)$ with $x = (y'_{-L+1}, \ldots, y'_{-1}, y'_0)'$; y_2^r is a drawing from $f(y|x)$ with $x = (y'_{-L+2}, \ldots, y'_0, y_1^{r\prime})'$; and so forth. The return forecast of holding the stock for τ periods from the initial state x, denoted $\hat{r}_\tau(x)$, is given by

$$\hat{r}_\tau(x) \approx \frac{1}{R}\sum_{r=1}^{R} \{\log\left[1 + g(y'_{-L+1}, \ldots, y'_0)\right] + \cdots + \log\left[1 + g(y^{r\prime}_{\tau-L+1}, \ldots, y^{r\prime}_{\tau-1})\right]\}. \tag{4}$$

The approximation error tends to zero almost surely as $R \rightarrow \infty$, under mild regularity conditions on $f(\cdot|\cdot)$. The conditional variance of the stock returns at various holding periods, denoted $\widehat{\mathrm{Var}}_\tau(x)$, can be similarly obtained by

$$\widehat{\mathrm{Var}}_\tau(x) \approx \frac{1}{R}\sum_{r=1}^{R} \{\log\left[1 + g(y'_{-L+1}, \ldots, y'_0)\right] + \cdots + \log\left[1 + g(y^{r\prime}_{\tau-L+1}, \ldots, y^{r\prime}_{\tau-1})\right]\}^2 - [\hat{r}_\tau(x)]^2. \tag{5}$$

The time-varying expected returns and conditional variance can thus be obtained by setting the conditioning set at each sample point, i.e., $x = x_t$, $t = L, \ldots, T$.

In this application, we examine the nominal excess return, defined as the value-weighted NYSE return minus the one-month Treasury bill rate. Many studies have documented that the dividend yield and risk-free rate have significant predictive power for stock returns (Fama & French, 1988; Hodrick, 1992; Campbell et al., 1997). We thus include the dividend yield and risk-free rate in our information set along with the excess returns. The data for this study is obtained from the CRSP tape and spans the period from January, 1926, to December, 1996, at a monthly frequency. The dividend yield is adjusted for seasonality by averaging the dividends of the previous twelve months as in Fama and French (1988) and Campbell et al. (1997). Considering the possible existence of a unit root in the one-month Treasury bill rate, we follow Campbell (1991), Hodrick (1992), and Campbell et al. (1997) and stochastically detrend the Treasury bill rate by removing a backward one-year moving average of past bill rates from the current bill rate. Thus, $y_{t+1} = (rx_{t+1}, dy_{t+1}, df_{t+1})'$, where rx_{t+1} is the excess return, dy_{t+1} is the dividend yield, and df_{t+1} is the stochastically detrended one-month Treasury bill rate. We can thus obtain the expected excess returns and conditional variance by setting $g(\cdot)$ equal to the first element of vector y.

We create the expected excess returns and conditional variance over four different holding periods: one month, one

[5] This definition follows Campbell et al. (1997).

quarter, one year, and two years (τ = 1, 3, 12, and 24). Of course, since the true conditional density function $f(y|x)$ is unknown, we must estimate it. We choose to apply the seminonparametric (SNP) method developed by Gallant and Nychka (1987) and Gallant and Tauchen (1989) which has since appeared in many finance applications (Gallant et al., 1991, 1997; Gallant et al., 1992, 1993; Tauchen et al., 1996; Harrison, 1998). The SNP method is desirable because it is flexible enough to allow for heterogeneous nonlinear interaction while nesting more standard VAR and ARCH models. The next section describes its implementation.

III. The SNP Density Estimation

The SNP method is based on the notion that a Hermite expansion can be used as a general-purpose approximation to a density function. This basic approach can be adapted to the estimation of the conditional density of a multiple time series $\{y_t\}$ that has a Markovian structure—where the conditional density of the M-vector y_t given the entire history $\{y_{t-1}\}_{l=1}^{\infty}$ depends only on L lags from the most recent past. Collecting these lags together in a single vector gives a $M \times L$-vector denoted as x_{t-1}:

$$x_{t-1} = (y'_{t-L}, \ldots, y'_{t-2}, y'_{t-1})'.$$

The richest SNP approximation of a multivariate density function takes the form of a normal distribution modified by a polynomial:

$$f(y|x, \theta) = \frac{1}{\lambda} [P(z, x)]^2 n_M(y|\mu_x, \Sigma_x) \tag{6}$$

where λ is a scalar that makes the density integrate to one, $z = R_x^{-1}(y - \mu_x)$ is a vector of innovations, $n_M(y|\mu_x, \Sigma_x)$ is a normal distribution with mean μ_x (the location function) and variance-covariance $\Sigma_x = R_x R'_x$, R_x (the scale function) is an upper triangular matrix, and $P(z, x)$ denotes a polynomial in z of degree K_z whose coefficients are polynomials of degree K_x in x.

The constant term of the polynomial is put to one to obtain a unique representation. This normalization means that the leading term of the entire expansion is $n_M(y|\mu_x, \Sigma_x)$—a Gaussian ARCH.

The location function μ_x is given by a vector autoregression

$$\mu_{x_{t-1}} = b_0 + Bx_{t-1}. \tag{7}$$

It is assumed to depend on $L_\mu \le L$ lags. The scale function R_x is given by

$$\text{vech}\,(R_{x_{t-1}}) = \rho_0 + P|e^*_{t-1}| \tag{8}$$

where vech (R) denotes a vector of length $M(M + 1)/2$ containing the elements of the upper triangle of R, $e^*_{t-1} = [(y_{t-L_r} - \mu_{x_{t-}L_r-1}), \ldots, (y_{t-1} - \mu_{x_{t-2}})]$, and $|\cdot|$ denotes elementwise absolute value.

The variance function depends on L_r lagged innovations $(y_t - \mu_{x_{t-1}})$. This is an ARCH-type process with the variance a linear function of the absolute lagged residuals rather than squared lagged residuals and in that way is similar to the specification of Davidian and Carroll (1987) and Nelson (1991).

The Hermite polynomial $P(z, x)$ is given by

$$P(z,x) = \sum_{\alpha=0}^{K_z} \left(\sum_{\beta=0}^{K_x} \alpha_{\alpha\beta} x^\beta \right) z^\alpha \tag{9}$$

where α and β are nonnegative integers, and $z^\alpha = z_1^{\alpha_1} \ldots z_M^{\alpha_M}$ with $|\alpha| = \Sigma_{k=1}^M \alpha_k$; similarly for x^β. It is assumed that the polynomial depends on $L_p \le L$ lags from x. The hierarchical SNP structure can thus be represented by L_μ, L_r, L_p, K_z and K_x. For notational convenience, hereafter, we use SNP(L_μ, L_r, L_p, K_z, K_x).[6]

When K_z is positive, the resulting density function is a modification of the Gaussian due to the multiplication by the polynomial $\{P(z)\}^2$. When K_x is positive, the shape of the density will depend on x. Thus, all moments can depend on x, and the density can approximate any form of conditional heterogeneity (Gallant & Tauchen, 1989). The shape modifications are rich enough to accurately approximate densities from a large class that includes densities that have fat tails, thin tails, or are skewed. The relationship between the parameter settings and the characteristics of the density of y_t is summarized in table (Gallant et al., 1993).

Parameter Setting	Characterization of y_t
$L_\mu = 0, L_r = 0, L_p \ge 0, K_z = 0, K_x = 0$	i.i.d. Gaussian
$L_\mu > 0, L_r = 0, L_p \ge 0, K_z = 0, K_x = 0$	Gaussian VAR
$L_\mu > 0, L_r = 0, L_p \ge 0, K_z > 0, K_x = 0$	non-Gaussian VAR, homogeneous innovations
$L_\mu \ge 0, L_r > 0, L_p \ge 0, K_z = 0, K_x = 0$	Gaussian ARCH
$L_\mu \ge 0, L_r > 0, L_p \ge 0, K_z > 0, K_x = 0$	non-Gaussian ARCH, homogeneous innovations
$L_\mu \ge 0, L_r \ge 0, L_p > 0, K_z > 0, K_x > 0$	general nonlinear process, heterogeneous innovations

The parameter vector θ of $f(y|x, \theta)$ thus consists of the coefficients of the polynomial plus μ_x and R_x and is estimated by maximum likelihood. If the number of parameters p_θ grows with the sample size n, then the true density, its derivatives, and moments are estimated consistently as shown in Gallant and Nychka (1987).

[6] Thus, the model "SNP(17121)" has $L_\mu = 1$, $L_r = 7$, $L_p = 1$, $K_z = 2$, and $K_x = 1$. When the number of lags is greater than nine, we use letters a, b, c, . . . ; for example, a = 10, b = 11, and c = 12.

TABLE 1.—SNP ESTIMATION—MAXIMUM LIKELIHOOD SURFACE

Specification	p_θ	s_n	Schwarz	H-Q	Akaike
SNP (11100)	21	0.9306	1.0148	0.9783	0.9556
SNP (1h100)	69	0.2456	0.5221	0.4022	0.3277
SNP (1c121)	93	0.0579	0.4307	0.2691	0.1687

Notes: This table compares the goodness-of-fit of three key SNP models. SNP(1c121) is the preferred model, SNP(11100) is a simple ARCH model, and SNP(1h100) is a long-lagged ARCH model meant to capture GARCH effects. p_θ is the number of parameters; s_n is the value of the objective function; and the selection criteria are the Schwarz criterion (s_n + $[p_\theta/2n]$ ln (n)), the Hannan-Quinn criterion (s_n + $[p_\theta/n]$ ln [ln (n)]), and the Akaike criterion (s_n + $[p_\theta/n]$).

To select the optimal SNP model, the following fitting strategy is adopted. We start with a simple Gaussian VAR process and gradually add lags in the location function until the marginal decrease in the objective value (s_n) is very small and some model selection criteria reach a minimum.[7] We then incrementally introduce ARCH by adding lags to the scale function, then non-Gaussian ARCH, and finally general nonlinear processes by allowing the polynomial to modify the underlying Gaussian density.

We select the models preferred by the selection criteria. The preferred models are then subjected to a battery of diagnostic tests to determine the goodness of fit. First, simulated series from the models are compared to the originals over their first four moments.[8] Second, the residuals from each SNP model are examined for both short-term and long-term predictability of the mean (residual levels) and variance (squared residuals). For the long-term tests, the residuals and their squares are projected onto ten-year dummy variables. For the short-term tests, the residuals and their squares are projected onto a space formed by the linear, quadratic, and cubic terms of past variables (five lags were used). If an SNP model is the true density, then the residuals should be orthogonal to the regressors. Thus, the null hypothesis that all regressor coefficients equal zero should not be rejected by an F-test and the R^2 will equal zero. Thus, for a given SNP specification, the smaller the R^2 of the regressions, the better the SNP model approximates the true density.

In table 1, we present the maximum-likelihood surface for three key models:

(1) the basic ARCH(1) model, which is SNP(11100);
(2) the ARCH model with many lags in the variance (given by SNP(1h100), which has seventeen lags in the variance) to approximate the GARCH(1,1) specification; and
(3) the preferred model from our selection procedure, which is SNP(1c121).

The preferred model is a general nonlinear process with heterogeneous innovations. It has one lag in the mean of the distribution and twelve lags in the variance-covariance matrix. The polynomial is quadratic in innovations with coefficients being linear in x with one lag. It has 93 parameters that are estimated with 816 observations for each of three series implying a saturation ratio of 26.3.[9] Table 1 indicates that the preferred model performs substantially better than the other two models according to all three model-selection criteria.

The superior performance of the preferred model in matching the data is also revealed by the diagnostic tests summarized in table 2. While all three models do well in the long-term test (the adjusted R^2 for all the long-term tests are no more than 5%), the preferred model outperforms the other two in the short-term tests. The highest R^2 for the preferred model in the short-term tests is 14% compared with 35% for the simple ARCH model and 27% for the long-lagged ARCH model. This evidence indicates that standard ARCH and GARCH models are likely to be misspecified and that an SNP model fits the data better. Examining the normality of the residuals also illustrates the superiority of the preferred model since its excess return residuals exhibit more normality and less excess kurtosis than the ARCH-type models.[10]

There is some evidence of lingering higher-order predictability in the residuals. But, with low adjusted R^2 values, we are optimistic that, while we have not made the perfect fit, it is adequate. If we were to expand the SNP model we could eliminate much of the remaining predictability—that is, the penalty term in our criteria function stops our expansion too soon in this regard. But we prefer a more parsimonious model to facilitate interpretation and to enhance the precision of the estimate.

One concern with any estimation is the possibility of omitted variables through an incorrectly specified information set. Our instruments are the variable of interest, the dividend yield, and the Treasury bill rate. However, other variables have been found to predict stock returns, such as the default premium and the term premium (Fama & French, 1989). Including these in our estimation involves a tradeoff with the power of the estimation because of the large number of additional parameters to be estimated with a finite sample. We thus exclude them.

However, it is important to consider what is lost by this omission. To explore the impact of other variables, we take the residuals from the three models and regress them on multiple lags of the default premium (the yield differential between the BAA and AAA corporate bonds) and the term premium (the yield differential between ten-year maturity

[7] Specifically, three model-selection criteria are calculated for each SNP fit: the Schwarz criterion (s_n + $[p_\theta/2n]$ ln (n)), the Hannan-Quinn criterion ($s_n + p_\theta/n$ ln (n)]), and the Akaike criterion ($s_n + p_\theta/n$).

[8] The results are not reported and are available upon request.

[9] The saturation ratio is defined as the total number of observations divided by the number of parameters estimated. Portnoy (1985) gives the maximum number of parameters in a linear regression as a function of sample size such that asymptotic normality of a linear function of the parameters is preserved. Using the Marron and Wand (1992) test suite, Fenton and Gallant (1996) find that the finite sample performance of the SNP estimator for a univariate series is both qualitatively and asymptotically similar to the kernel estimator, which is optimal. However, we have not yet seen studies on the finite sample performance of the SNP estimator in a multivariate setting.

[10] The results are not reported and are available upon request.

TABLE 2.—DIAGNOSTIC TESTS

	Excess Return				Dividend Yield				Risk-free Rate			
	Mean		Var		Mean		Var		Mean		Var	
Model	p	R^2	p	R^2	p	R^2	p	R^2	p	R^2	p	R^2
short-term												
11100	0.000	0.17	0.000	0.35	0.000	0.18	0.000	0.29	0.550	0.00	0.000	0.05
1h100	0.000	0.09	0.020	0.03	0.000	0.14	0.000	0.27	0.930	0.00	0.870	0.00
1c121	0.000	0.11	0.002	0.04	0.000	0.11	0.000	0.14	0.940	0.00	0.980	0.00
long-term												
11100	0.010	0.01	0.000	0.04	0.000	0.04	0.000	0.05	0.550	0.00	0.000	0.05
1h100	0.050	0.01	0.004	0.02	0.000	0.04	0.000	0.04	0.930	0.00	0.870	0.00
1c121	0.030	0.01	0.010	0.01	0.000	0.01	0.000	0.04	0.940	0.00	0.980	0.00

Notes: This table reports the results from regressions of the SNP residuals for each series (mean) and the residuals squared (variance) as a check for lingering predictability. The reported "p" is the probability from an F-test that all of the regressor coefficients are equal to zero, and the reported R^2 is the adjusted R-square from the same regression. The short-term tests use five lags of the information variables (excess return, dividend yield, and risk-free rate) and their squares and cubes as regressors. The long-term tests use ten-year dummy variables as regressors.

government bonds and three-month Treasury bills).[11] For the excess return residuals from the preferred model, the null hypothesis of no effect from the additional two variables cannot be rejected at the 5% level (the P-value of the F statistic is 0.866) and the adjusted R^2 of the regression is 0.00. Thus, whatever predictability may exist is less than 1% of the total unexplained variation. This is not necessarily an endorsement of our estimation, but it does suggest that the marginal impact, if any, from expanding the information set is likely to be very small. Contrastingly, for the simple ARCH model, the null hypothesis of no effect from the omitted variables is rejected at the 1% level. For the long-lagged ARCH model, the null hypothesis of no effect cannot be rejected at the 5% level, but it can be at 15% (the P-value is 0.146), and the adjusted R^2 is greater than 1%.

Similar results are obtained by regressing the excess return residuals onto the lagged consumption growth rate.[12] The preferred model yields a P-value of 0.508 and an R^2 of zero; the simple ARCH yields a P-value of 0.135 and an R^2 of 0.01; and the long ARCH yields a P-value of 0.487 and an R^2 of zero. Again, we see that the preferred model performs the best and that this particular omitted variable is not likely to be a concern.

IV. Risk and Return

A. *Cyclicality of Expected Returns and Conditional Volatility*

Figure 1 presents the estimated time series of the expected excess returns,[13] and figure 2 is its volatility measured by the conditional variance (using 10,000 replications in the Monte Carlo integration ($R = 10{,}000$)) from the preferred SNP model. Both the expected excess return and conditional volatility series appear to reach a peak during recessionary periods classified by the National Bureau of Economic Research (NBER), which are marked by vertical lines. In other words, both the expected returns and conditional volatility are countercyclical. It also appears that the countercyclical movement of the volatility is more prominent than that of the expected returns.

FIGURE 1.—EXPECTED NOMINAL EXCESS RETURNS AT (A) MONTHLY, (B) QUARTERLY, (C) ONE YEAR, AND (D) TWO YEAR HOLDING INTERVALS. THE VERTICAL GRID LINES ARE NBER BUSINESS-CYCLE PEAKS AND TROUGH.

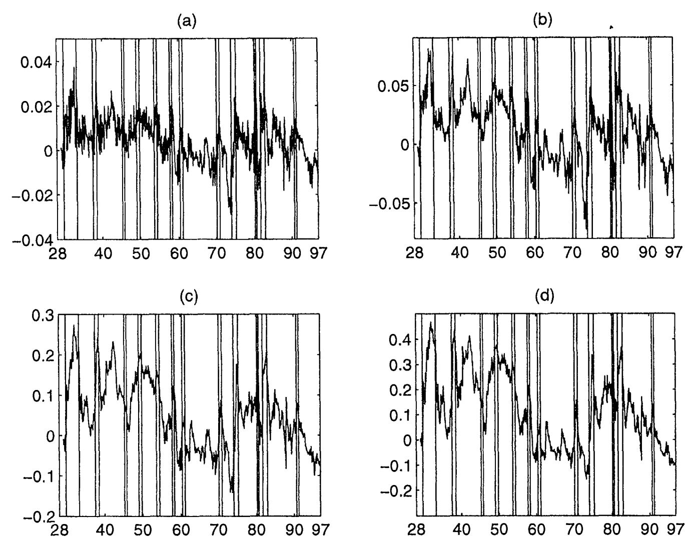

To assess their statistical significance and to quantify the size of this relation, we relate our estimates to actual measures of the business cycle. Let $r(\widehat{t, t} + \tau)$ and Var $(r(\widehat{t, t}, + \tau))$ be the predicted holding returns and condi-

[11] Unfortunately, we do not have this data prior to 1959 so that this test is only over the more recent subsample.

[12] The consumption growth rate is defined as the growth rate of per capita nondurable and service expenditure in constant dollars obtained from CITIBASE.

[13] There is a long period of negative expected excess returns as implied by figure 1, which has also been found in other papers in the related literature. The model of Campbell (1986) generates a simple example of how this could come about. In his framework, the risk premium on a stock can be decomposed into the holding premium of bonds and a factor resulting from the payoff uncertainty on the stock. If the maturity of the stock is different than the holding period, neither of the above two terms is unambiguously positive or negative. Therefore, an asset whose payoff is positively correlated with consumption may have a return that is negatively correlated with consumption (and thus a negative risk premium). Alternatively, as pointed out by a referee, negative expected returns may indicate that market participants were relying on additional variables in their information set not observed by econometricians or displayed some degree of irrationality.

FIGURE 2.—CONDITIONAL VOLATILITY OF EXCESS RETURNS AT (A) MONTHLY, (B) QUARTERLY, (C) ONE-YEAR, AND (D) TWO-YEAR HOLDING INTERVALS. THE VERTICAL GRID LINES ARE NBER BUSINESS-CYCLE PEAKS AND TROUGH.

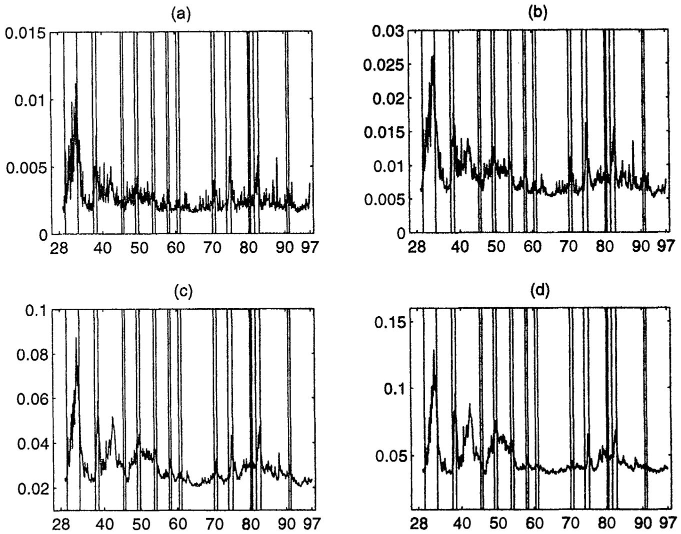

tional volatility, respectively. We then estimate the following linear models to determine their cyclical behavior:[14]

$$r(\widehat{t, t} + \tau) = \delta(\tau) + \gamma(\tau)BC_t + u_{t+\tau} \tag{10}$$

$$\mathrm{Var}\,(r(\widehat{t, t} + \tau)) = \zeta(\tau) + \xi(\tau)BC_t + v_{t+\tau} \tag{11}$$

where BC_t is the proxy for the state of the economy. Three proxies are used in this application: an NBER business-cycle dummy, the growth rate of industrial production, and an index of consumer confidence.[15]

In the first regression, BC_t is a dummy variable representing the NBER classification of business cycles, $BC_t = BCD_t$. It is defined as

$$BCD_t = \begin{cases} 1 & \text{if Peak} < t \leq \text{Trough} \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Because of the limited number of recessions over this time period—as well as the fact that they are determined ex post—we recognize that this might not be the ideal proxy. To ameliorate this concern, other proxies are also used.

In the second regression, BC_t is the growth rate of industrial production, $BC_t = IP_t$. Downturns in industrial production are, of course, associated with recessions. But, as a continuous variable, it may provide a richer comovement with the expected returns and account for the severity of peaks and troughs.

The third regression employs a more forward-looking variable by using an index of consumer expectations as a measure of economic conditions. The Index of Consumer Expectations is part of the Survey of Consumer Sentiment conducted by the University of Michigan which dates back to 1953 and was obtained from CITIBASE. We follow Marcus and Ors (1996) who use the Consumer Confidence Index as a proxy for economic activity. As they point out, it is attractive to have a measure based on then-current economic expectations. We use the level of the index itself as our proxy, so that higher-than-average expectations correlate with peaks and lower-than-average with troughs, so $BC_t = CE_t$.

Table 3 presents the results on the relationship between the expected excess returns (for each holding period) and each of the three different business-cycle proxies.[16] The findings indicate that the expected return is countercyclical. The results are similar across holding period and business-cycle proxy. They are also consistent with the previous literature (Fama & French, 1989; Kandel & Stambaugh, 1990).[17]

In table 4, the same regressions are done with the conditional volatility as the dependent variable instead of the

TABLE 3.—REGRESSIONS OF NOMINAL EXCESS RETURNS ON THE BUSINESS-CYCLE PROXY VARIABLE

Horizon	Variable	Param. Estimate	Standard Error	T for H0: $\beta = 0$	$Prob > \|T\|$
M	BCD	0.0074	0.00155	4.77	0.0001
Q	BCD	0.0251	0.00520	4.82	0.0001
1Y	BCD	0.0870	0.02117	4.11	0.0001
2Y	BCD	0.1450	0.3765	3.85	0.0001
M	IP	−0.1003	0.05456	−1.84	0.0662
Q	IP	−0.3287	0.17542	−1.87	0.0610
1Y	IP	−1.1232	0.60423	−1.86	0.0630
2Y	IP	−1.8943	0.64085	−2.96	0.0031
M	CE	−0.0001	0.00006	−1.79	0.0737
Q	CE	−0.0004	0.00019	−2.13	0.0334
1Y	CE	−0.0018	0.00061	−2.94	0.0001
2Y	CE	−0.0034	0.00116	−2.93	0.0031

Note: The constant is statistically significant at the 1% level in all regressions and so is not reported in the table. In the first regression, the proxy variable is the business-cycle dummy variable ("BCD"); in the second, the proxy is the growth rate of industrial production ("IP"); and, in the third, the proxy is an index of consumer expectations ("CE"). The standard errors are calculated to be robust for heteroskedasticity and serial correlation according to the Newey-West procedure (using 3 lags for the monthly, 6 for quarterly, 12 for yearly, and 24 for the two-year holding period).

[14] We do not include the proxies for the business cycle in our density estimation of the information variables to confine the dimension of the SNP estimation to be relatively parsimonious. Our two-step procedure is a computational compromise. A recent study by Perez-Quiros and Timmermann (1997) uses a Markov switching process to explicitly model the expected excess returns jointly with business cycles.

[15] Because the true expected returns are not observed, in practice we use the estimated expected returns. As a result, the regression is subject to errors in the dependent variable. As long as the estimation errors are uncorrelated with the business-cycle variable, the results should be unbiased. To evaluate this possibility, we create the residuals between the realized returns and the estimated expected returns, and calculate the correlations between the residuals and the business-cycle proxies. These correlations are very small in magnitude, which suggests the regression is unbiased. If any effect, the greater regression error should therefore inflate the standard errors on γ, and thus understate the t-statistics and significance levels.

[16] Because of the availability of data, the regressions with industrial production and consumer expectations are done only with the subsample since 1959.

[17] We find that the selection of information variables is very important. For instance, procyclical returns can be obtained when the risk-free rate is not detrended. This sensitivity could be due to the near unit root in the risk-free rate.

TABLE 4.—REGRESSIONS OF CONDITIONAL VOLATILITY OF NOMINAL EXCESS RETURNS ON THE BUSINESS-CYCLE PROXY VARIABLE

Horizon	Variable	Param. Estimate	Standard Error	T for H0: $\beta = 0$	$Prob > \|T\|$
M	BCD	0.0014	0.00031	4.46	0.0001
Q	BCD	0.0039	0.00083	2.77	0.0059
1Y	BCD	0.0107	0.00381	2.80	0.0051
2Y	BCD	0.0143	0.00612	2.34	0.0195
M	IP	−0.0130	0.00491	−2.65	0.0081
Q	IP	−0.0362	0.01642	−2.21	0.0274
1Y	IP	−0.0878	0.04642	−1.89	0.0587
2Y	IP	−0.1032	0.03950	−2.61	0.0090
M	CE	−0.00002	0.000004	−4.68	0.0001
Q	CE	−0.00006	0.000016	−3.76	0.0002
1Y	CE	−0.00015	0.000044	−3.40	0.0007
2Y	CE	−0.00020	0.000064	−3.15	0.0017

Notes: The constant is statistically significant at the 1% level in all regressions and so is not reported in the table. In the first regression, the proxy variable is the business-cycle dummy variable ("BCD"); in the second, the proxy is the growth rate of industrial production ("IP"); and, in the third, the proxy is an index of consumer expectations ("CE"). The standard errors are calculated to be robust for heteroskedasticity and serial correlation according to the Newey-West procedure (using 3 lags for the monthly, 6 for quarterly, 12 for yearly, and 24 for the two-year holding period).

expected returns. The conditional volatility is also found to be countercyclical irrespective of holding period or proxy variable.[18] This, too, is consistent with the prior literature, such as Schwert (1989) and Kandel and Stambaugh (1990).

B. *The Risk and Return Relation*

The systematic time variation in the expected returns and conditional volatility begs the question of how they are related to each other. Of particular interest is the idea that the volatility of returns is taken by investors as a measure of riskiness of financial assets and thus should be priced. The standard intuition is that high volatility will induce investors out of stocks and into bonds. This should raise the expected return to equity and lower the expected return to bonds.

Figure 3 presents the scatterplots of the expected nominal excess returns at the various holding intervals against the corresponding conditional variances at the same intervals for the preferred SNP model. We make the following observations. First, at the longer holding interval, a positive relation appears between the expected returns and the volatility—where, at the monthly holding interval, we observe an undistinguished cloud of points. Second, the risk and return relation at long horizons (such as one and two years) exhibits resemblance to the mean-variance efficient frontier. In particular, this illustrates a strong nonlinear relation between risk and return: the reward to risk-taking increases at a decreasing rate.

A possible explanation for a holding-period effect is that investments over short horizons may sometimes be influenced by portfolio balance and transaction cost considerations or by unexpected immediate consumption needs

[18] Because of serial correlation in the expected returns and conditional volatility, the standard errors in both tables are estimated using the Newey-West procedure (Newey & West, 1987). Also, the use of overlapping observations in the longer holding periods does not appear to be a concern because of the consistency in results across all of the intervals and across the business-cycle proxies (which have varying levels of autocorrelation).

FIGURE 3—.EXPECTED EXCESS RETURNS VERSUS THEIR CONDITIONAL VOLATILITY FOR THE PREFERRED SNP MODEL AT (A) MONTHLY, (B) QUARTERLY, (C) ONE-YEAR, AND (D) TWO-YEAR HOLDING INTERVAL, RESPECTIVELY.

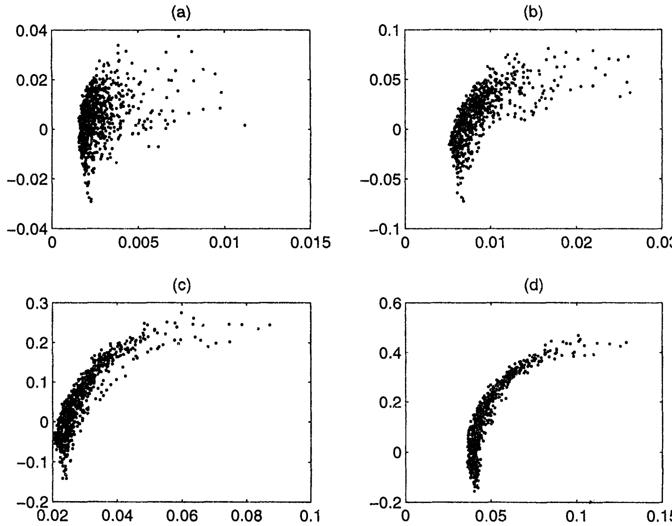

(Daniel & Marshall, 1997). All these factors may obscure the risk and return relation in the short horizon. But, in the long term, these factors are less likely to play an important role compared with the risk factor. Thus, the true risk and return relation is more likely to be revealed at longer holding intervals than in the short horizons. The findings are, to some extent, consistent with the empirical evidence that intertemporal consumption asset pricing (ICAPM) models cannot be rejected at longer holding intervals, even though they can be rejected with monthly data (Lewis, 1991; Handa et al., 1993; Daniel & Marshall, 1997). Our findings are also consistent with the fact that longer holding periods have been associated with a greater predictability of returns (Fama & French, 1988), which may imply that statistically sharper results on the risk and return relation are obtained at longer holding intervals.

To assess the effect of different parametric specifications on the risk and return relation, we also estimate the expected excess return and conditional variance for the simple ARCH model and the long-lagged ARCH model using the same Monte Carlo simulation technique. The results are plotted in figures 4 and 5, respectively. They illustrate no meaningful positive relation between risks and returns. In some cases, we find a negative risk and return relation, but in most cases there appears to be no trade-off. These results support the argument that misspecifications in prior parametric studies could have an important impact on their inferences about the true risk and return relation.

C. *The Negative Risk-Return Relation: A Theoretical Framework and Possible Explanation*

To aid the interpretation of our risk-return findings and to offer additional insight of why they differ from previous studies, we introduce the theoretical framework proposed by

FIGURE 4.—EXPECTED EXCESS RETURNS VERSUS THEIR CONDITIONAL VOLATILITY FOR THE SIMPLE ARCH MODEL AT (A) MONTHLY, (B) QUARTERLY, (C) ONE-YEAR, AND (D) TWO-YEAR HOLDING INTERVAL, RESPECTIVELY.

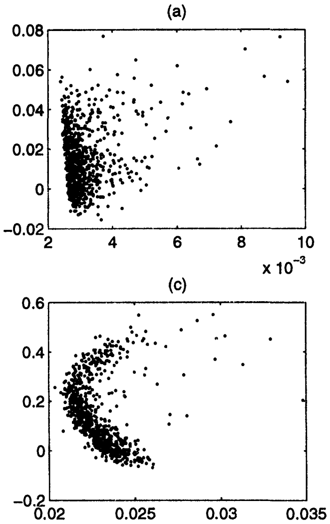

FIGURE 5.—EXPECTED EXCESS RETURNS VERSUS THEIR CONDITIONAL VOLATILITY FOR THE LONG-LAGGED ARCH MODEL AT (A) MONTHLY, (B) QUARTERLY, (C) ONE-YEAR, AND (D) TWO-YEAR HOLDING INTERVAL, RESPECTIVELY.

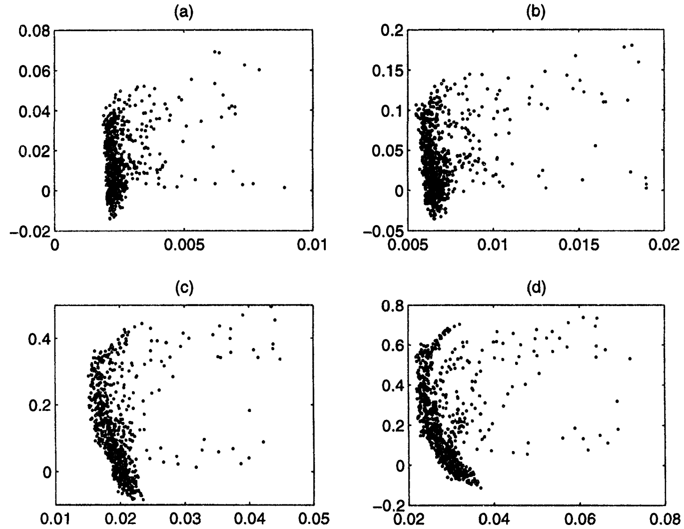

Campbell (1996). In analyzing the cross-sectional pattern of postwar U.S. stock and bond returns, Campbell derives the following asset-pricing formula:[19]

$$E_t r_{i,t+1} - r_{f,t+1} + \frac{V_{ii}(t)}{2} = \gamma V_{im}(t) + (\gamma - 1)V_{ih}(t) \quad (13)$$

where $r_{i,t+1}$ is the return on asset i from time t to $t + 1$, $r_{f,t+1}$ is the return on a riskless asset from time t to $t + 1$, $V_{ii}(t)$ is the conditional variance of the return on asset i at time t, $V_{im}(t)$ is the conditional covariance between the return on asset i and the market return at time t, $V_{ih}(t)$ is the conditional covariance of the unexpected return on asset i with the present value of future market returns, and γ is the relative risk aversion coefficient.

As pointed out by Campbell (1996), the above asset-pricing formula expresses the risk premium, after being adjusted for the Jensen's inequality effect, as a weighted sum of two conditional covariances. The first covariance, with a weight of γ, is the covariance of the asset's return with the market return. The second covariance, with a weight of $\gamma - 1$, is the covariance of the asset return with news about future returns on the market. If we let asset i be the market portfolio, we arrive at the following equation:

$$E_t(r_{m,t+1} - r_{f,t+1}) = (\gamma - \tfrac{1}{2})V_{mm}(t) + (\gamma - 1)V_{mh}(t) \quad (14)$$

where the left-hand side is the expected excess return on the market portfolio from time t to $t + 1$, $V_{mm}(t)$ is the conditional variance of the market return at time t, and $V_{mh}(t)$ is the time t conditional covariance. The first term on the right-hand side summarizes the risk and return trade-off. The second term is the analogue to the hedge portfolios of Merton's (1973) intertemporal model.

The above equation collapses to the conditional CAPM only if the risk aversion coefficient is 1, if $V_{mh}(t)$ is 0, or if the conditional covariance $V_{mh}(t)$ is perfectly correlated with the conditional variance $V_{mm}(t)$. Without one of these three circumstances, the standard CAPM will not hold and Merton's hedge term is an omitted variable. Indeed, if the conditional variance $V_{mm}(t)$ and the conditional covariance $V_{mh}(t)$ are negatively correlated,[20] omission of $V_{mh}(t)$ could account for findings of a negative or weak positive relation between the expected excess return and the conditional variance.

Under this framework, our long-horizon risk and return results could differ from the short horizon if the monthly covariance $V_{mh}(t)$ differs from the two-year covariance.[21] At short horizons, $V_{mh}(t)$ might be more noisy and move independently or opposite of $V_{mm}(t)$. However, as the horizon increases far enough, $V_{mh}(t)$ must eventually equal zero, leaving a linear risk-return relation, since at very long horizons returns will not be forecastable.

D. *Time-Varying Risk and Return Relation*

We also want to test how the relation between expected returns and conditional variance varies with the business cycle. To achieve this goal, we first construct the price of risk measure (the Sharpe ratio) for the nominal excess returns as

[19] In deriving the asset pricing formula, Campbell (1996) uses the objective function proposed by Epstein and Zin (1989, 1991) and Weil (1989) to distinguish the coefficient of relative risk aversion and the elasticity of intertemporal substitution. He also assumes that the log consumption/wealth ratio is stationary. We refer readers to Campbell (1996) for the details.

[20] For instance, in a recession, when $V_{mm}(t)$ is high, $V_{mh}(t)$ might be negative because future returns move opposite current returns, thus giving the two an inverse correlation. Campbell et al. (1997) point out that a negative $V_{mh}(t)$ could exist because of mean reversion in the market.

[21] We thank the editor for pointing out this interpretation.

TABLE 5.—REGRESSIONS OF PRICE OF RISK OF NOMINAL EXCESS RETURNS ON THE BUSINESS-CYCLE PROXY VARIABLE

Horizon	Variable	Param. Estimate	Standard Error	T for H0: $\beta = 0$	$Prob > \|T\|$
M	BCD	0.1209	0.02546	4.75	0.0001
Q	BCD	0.2252	0.04517	4.98	0.0001
1Y	BCD	0.4147	0.09197	4.51	0.0001
2Y	BCD	0.5378	0.12573	4.28	0.0001
M	IP	−1.8166	1.06110	−1.71	0.0869
Q	IP	−3.3464	1.80323	−1.86	0.0634
1Y	IP	−6.1082	3.24510	−1.88	0.0598
2Y	IP	−8.0754	2.77532	−2.91	0.0036
M	CE	−0.0019	0.00118	−1.65	0.0994
Q	CE	−0.0041	0.00209	−1.98	0.0474
1Y	CE	−0.0099	0.00348	−2.85	0.0043
2Y	CE	−0.0148	0.00523	−2.82	0.0048

Notes: The constant is statistically significant above the 1% level in all regressions and so is not reported in the table. In the first regression, the proxy variable is the business-cycle dummy variable ("BCD"); in the second, the proxy is the growth rate of industrial production ("IP"); and, in the third, the proxy is an index of consumer expectations ("CE"). The standard errors are calculated to be robust for heteroskedasticity and serial correlation according to the Newey-West procedure (using 3 lags for the monthly, 6 for quarterly, 12 for yearly, and 24 for the two-year holding period).

the ratio of the expected returns and the square root of the conditional variance, i.e.,

$$r(\widehat{t, t} + \tau)/\sqrt{\mathrm{Var}\,(r(\widehat{t, t} + \tau))}\ .$$

We then regress the price of risk on the three business-cycle proxy variables. Under the null hypothesis that the risk and return relation is time-invariant, the price of risk of the expected returns will be constant, and the coefficient of the business-cycle proxy variables in the above regressions will then be zero.

In table 5, we report the results on the relationship between the price of risk and the business-cycle proxies for the nominal excess returns. We make the following observations. First, the price of risk varies with the business cycle. This is supported by the evidence that most of the coefficient estimates of the business-cycle proxies are statistically significant. Second, the price of risk is countercyclical, which is consistent with previous studies by Fama and French (1989), Harvey (1989, 1991), Kandel and Stambaugh (1990), Campbell and Cochrane (1999), and others. This result is analogous to that found for the cyclical behavior of the expected returns.

V. Conclusion

This paper explores the cyclical behavior of expected stock returns, their volatility, and the risk and return relation in association with various holding intervals. Using the seminonparametric (SNP) method of Gallant and Nychka (1987) and Gallant and Tauchen (1989) and Monte Carlo integration, we obtain the expected returns and the conditional volatility for the excess returns for holding intervals from one month to two years. This method is preferred over more-rigid parametric specifications in that no prior restrictions are imposed on either the conditional mean and conditional variance relation or on the relation of the dependent variables with the predicting variables. Using various proxies for the business conditions of the economy, we find that both the expected excess return and risks measured by the conditional volatility are countercyclical for the excess returns.

We find no obvious risk and return relation at short horizons such as one month. However, a positive risk and return relation appears at longer holding intervals such as a quarter, and one and two years. We find that the existing results of a negative risk and return relation could be due to an omitted variable effect in parametric specifications.

The time-invariant risk and return relation is also explicitly tested, and we find that the Sharpe ratio varies with the business cycle. Our study indicates that correctly specified economic models should be able to generate positive time-varying contemporaneous risk and return relations, at least at long holding intervals.

REFERENCES

Breen, William, Lawrence R. Glosten, and Ravi Jagannathan, "Economic Significance of Predictable Variations in Stock Index Returns," *Journal of Finance* 44 (1989), 1177–1189.

Campbell, John Y., "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics* (November, 1986) 785–803.

——— "Stock Returns and the Term Structure," *Journal of Financial Economics* 18 (1987), 373–399.

——— "A Variance Decomposition for Stock Returns," *Economic Journal* 101 (1991), 157–179.

——— "Understanding Risk and Return," *Journal of Political Economy* 104 (1996), 298–345.

Campbell, John Y., and John H. Cochrane, "By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107 (1999), 205–251.

Campbell, John Y., and Ludger Hentschel, "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns," *Journal of Financial Economics* 31 (1992), 281–318.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, *The Econometrics of Financial Markets* (Princeton: Princeton University Press, 1997).

Chan, K. C., G. Andrew Karolyi, and Rene Stulz, "Global Financial Markets and the Risk Premium on U.S. Equity," *Journal of Financial Economics* 32 (1992), 137–167.

Daniel, Kent, and David Marshall, "Consumption-Based Modeling of Long-Horizon Returns," *Macroeconomic Dynamics* 1 (1997), 452–484.

Davidian, Marie, and R. J. Carroll, "Variance Function Estimation," *Journal of the American Statistical Association* 82 (1987), 1079–1091.

Epstein, Larry G., and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica* 57 (1989), 937–969.

——— "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Journal of Political Economy* 99 (1991), 263–286.

Fama, Eugene F., and Kenneth R. French, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics* 22 (1988), 3–25.

——— "Business Conditions and Expected Stock Returns," *Journal of Financial Economics* 25 (1989), 23–49.

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

Fenton, Victor M., and A. Ronald Gallant, "Qualitative and Asymptotic Performance of SNP Density Estimators," *Journal of Econometrics* 74 (1996), 77–118.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, "Expected Stock Returns and Volatility," *Journal of Financial Economics* 19 (1987), 3–30.

Gallant, A. Ronald, David A. Hsieh, and George Tauchen, "On Fitting a Recalcitrant Series: The Pound/Dollar Exchange Rate, 1974–83," in William A. Barnett, James Powell, and George Tauchen (eds), *Nonparametric and Semiparametric Methods in Econometrics and Statistics* (Cambridge, UK: Cambridge University Press, 1991).

——— "Estimation of Stochastic Volatility Models with Suggestive Diagnostics," *Journal of Econometrics* 81 (1997), 159–192.

Gallant, A. Ronald, and Douglas W. Nychka, "Semiparametric Maximum Likelihood Estimation," *Econometrica* 55 (1987), 363–390.

Gallant, A. Ronald, Peter E. Rossi, and George Tauchen, "Stock Prices and Volume," *Review of Financial Studies* 5 (1992), 199–242.

——— "Nonlinear Dynamic Structure," *Econometrica* 61 (1993), 871–907.

Gallant, A. Ronald, and George Tauchen, "Seminonparametric Estimation of Conditionally Constrained Heterogeneous Processes: Asset Pricing Applications," *Econometrica* 57 (1989), 1091–1120.

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Returns on Stocks," *Journal of Finance* 48 (1993), 1779–1802.

Handa, Puneet, S. P. Kothari, and Charles Wasley, "The Sensitivity of Multivariate Tests of the Capital Asset-Pricing Model to the Return Measurement Interval," *Journal of Finance* 48 (1993), 1543–1553.

Harrison, Paul, "Similarities in the Distribution of Percentage Price Changes Between the 18th and 20th Centuries," *Journal of Business* 71 (1998), 55–79.

Harvey, Campbell R., "Time-Varying Conditional Covariances in Tests of Asset Pricing Models," *Journal of Financial Economics* 24 (1989), 289–317.

——— "The Specification of Conditional Expectations," working paper, Duke University (1991).

Hodrick, Robert J., "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies* 5 (1992), 357–386.

Kandel, Shmuel, and Robert F. Stambaugh, "Expectations and Volatility of Consumption and Asset Returns," *Review of Financial Studies* 3 (1990), 207–232.

Lewis, Karen, "Should the Holding Period Matter for the Intertemporal Consumption-Based CAPM?" NBER working paper series, no. 3583 (1991).

Marcus, Alan J., and Evren Ors, "Hedging Corporate Bond Portfolios Across the Business Cycle," *Journal of Fixed Income* (March, 1996), 56–60.

Marron, J. S., and M. P. Wand, "Exact Mean Integrated Squared Error," *Annals of Statistics* 20 (1992), 712–736.

Merton, Robert C., "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41 (1973), 867–887.

Nelson, Daniel B., "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica* 59 (1991), 347–370.

Newey, Whitney K., and Kenneth D. West, "A Simple Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55 (1987), 703–708.

Pagan, Adrian R., and Y. S. Hong, "Nonparametric Estimation and the Risk Premium," in William Barnett, James Powell, and George Tauchen (eds.), *Nonparametric and Semiparametric Methods in Econometrics and Statistics* (Cambridge, UK: Cambridge University Press, 1991), 51–75.

Perez-Quiros, Gabriel, and Allan Timmermann, "Stock Returns in Recessions and Expansions: Assessing the Economic Value of Asymmetries," Federal Reserve Bank of New York and University of California at San Diego (1997).

Portnoy, Stephen, "Asymptotic Behavior of M Estimators of p Regression Parameters when p^2/n is Large; II. Normal Approximation," *Annals of Statistics* 13 (1985), 1403–1417.

Schwert, G. William, "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance* 44 (1989), 1115–1153.

Tauchen, George, Harold Zhang, and Ming Liu, "Volume, Volatility, and Leverage: A Dynamic Analysis," *Journal of Econometrics* 74 (1996), 177–208.

Turner, Christopher M., Richard Startz, and Charles R. Nelson, "A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market," *Journal of Financial Economics* 25 (1989), 3–22.

Weil, Philippe, "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics* 24 (1989), 401–421.

Whitelaw, Robert F., "Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns," *Journal of Finance* 49 (1994), 515–541.