The Restrictions on Predictability Implied by Rational Asset Pricing Models

Chris Kirby
University of Texas at Dallas

This article shows how rational asset pricing models restrict the regression-based criteria commonly used to measure return predictability. Specifically it invokes no-arbitrage arguments to show that the intercept, slope coefficients, and R^2 in predictive regressions must take specific values. These restrictions provide a way to directly assess whether the predictability uncovered using regression analysis is consistent with rational pricing. Empirical tests reveal that the returns on the CRSP size deciles are too predictable to be compatible with a number of well-known pricing models. However, the overall pattern of predictability across these portfolios is reasonably consistent with what we would expect under circumstances where predictability is rational.

The use of regression analysis to document return predictability has become commonplace in empirical work. A researcher specifies a model for forecasting returns, fits the model by ordinary least squares, and

This article is a revised version of the second essay in my dissertation at Duke University. I am indebted to my committee members—Cam Harvey, David Hsieh, Tom Smith, George Tauchen, and Bob Whaley—for their support and suggestions. I also appreciate the comments and suggestions of Giorgio DeSantis, Jeff Fleming, Ravi Jagannathan (the editor), Barbara Ostdiek, S. Viswanathan, two anonymous referees, and seminar participants at the Australian Graduate School of Management, the University of Chicago, Duke University, the University of Maryland, the University of Michigan, the University of Texas at Dallas, the University of Utah, and the 1994 meetings of the Western Finance Association. Part of this research was completed while I was visiting Rice University and the University of Maryland. Address correspondence to Chris Kirby, School of Management, University of Texas at Dallas, P.O. Box 830688, Richardson, TX 75083-0688.

The Review of Financial Studies Summer 1998 Vol. 11, No. 2, pp. 343–382

uses the estimates of the slope coefficients and regression R^2 to draw inferences about the ability to predict returns. Although most members of the finance profession seem convinced that the predictability documented in the empirical literature is real, there is considerable disagreement over its source. Some argue that the ability to predict returns arises from market inefficiencies, while others point to rational variation in expected returns as the cause. Prior research, however, has not shown how rational models of asset pricing restrict the regression-based criteria used to measure predictability. This article develops a simple characterization of these restrictions and outlines a general approach for testing whether they are satisfied.

Researchers have long recognized that predictability can exist in efficient capital markets. But to formally test whether the predictability encountered in practice is consistent with market efficiency, we need a model that tells us how assets are priced. Although rational asset pricing models take a variety of different forms, they generally have at least one basic implication in common. In frictionless markets, these models imply that some transformation of the equilibrium price process is a martingale with respect to the information that market participants use to form expectations. This martingale property establishes a direct link between the ability to predict returns and market efficiency. As a consequence, it imposes tight restrictions on the intercept, slope coefficients, and R^2 of the linear models used to predict returns.

The strategy used to characterize these restrictions is similar to that developed by Hansen and Jagannathan (1991). They employ no-arbitrage arguments to construct a lower bound on the variance of the set of admissible stochastic discount factors by showing that every element of this set has the same linear projection onto the asset payoff space. Although the coefficients in this linear projection cannot be estimated directly, they can be inferred by requiring that the projection assign the correct price to all traded assets. The restrictions on predictability are derived in an analogous fashion. Standard no-arbitrage pricing arguments are used to infer the values of the coefficients in a regression of asset returns on a set of instrumental variables. In fact, the variance bound itself can be expressed as a simple quadratic form in these inferred regression coefficients.

Since the restrictions derived in the article apply quite generally, they can aid in understanding the role that predictability plays in the econometric evaluation of asset pricing models. In particular, they make it possible to construct the linear specification for forecasting asset returns that is implied by the asset pricing model under consideration. The forecasts generated in this manner provide a common frame of reference for studying the ability of different pricing models to explain the predictable variation in returns. It might be the case, for example, that the variance of the forecasts associated with one model is too small, while that for another model is too large.

The restrictions can also be used to assess whether the ability to predict returns that has been documented using regression analysis is consistent with market efficiency. This issue is investigated using monthly returns on 10 stock portfolios that are formed on the basis of market capitalization. Harvey (1989) finds that the returns on these portfolios exhibit a significant degree of predictable variation over time. The empirical analysis looks at whether this predictability is consistent with the implications of a number of well-known asset pricing models. They include the utility-based models of Lucas (1978), Abel (1990), and Epstein and Zin (1991), along with simple conditional versions of the Sharpe (1964) and Lintner (1965) CAPM and the Fama and French (1993) three-factor model.

The empirical tests reveal that the size portfolio returns are too predictable to be consistent with these pricing models. Of course, this does not necessarily imply that the return predictability documented in this and previous studies is irrational. Although none of the pricing models fully explains the ability to predict returns, the observed cross-sectional relation between market capitalization and return predictability appears to be reasonably consistent with what we would expect under circumstances where predictability is rational. This finding is significant in light of the potential for data mining that is inherent in the search for instruments that predict returns. Overall the evidence suggests that if we can identify the factors that explain cross-sectional differences in average returns, then we may be able to explain cross-sectional differences in predictability as well.

1. The Pricing Restrictions on Linear Forecasting Models

There is substantial empirical evidence that conditioning information plays an important role in the pricing of many assets. A number of studies, for example, appear to document the existence of significant predictable variation in stock and bond returns.[1] Not surprisingly, these studies have generated a great deal of debate within the finance profession. Some researchers argue that the ability to predict stock and bond returns stems from rational variation in expected returns,

[1] Fama (1991) and Hawawini and Keim (1992) survey the empirical research in this area. The most widely cited studies of the ability to predict stock and bond returns include Keim and Stambaugh (1986), Campbell (1987), and Fama and French (1989).

while others contend it is due to market inefficiencies. There are even those who tend to discount the evidence of predictability altogether.[2]

Since almost all of the empirical work on predicting returns takes the form of regression analysis, it would be easier to interpret the evidence on predictability if we could show how asset pricing theory restricts the parameters of linear forecasting models. Of the existing studies that focus on return predictability, Hansen and Singleton (1983) comes the closest to addressing this issue. They use a representative agent model with constant relative risk aversion to derive explicit restrictions on the joint stochastic process that governs consumption growth and asset returns. Under their model, the predictable component of asset returns is proportional to the predictable component of consumption growth.

Ferson and Harvey (1991) also look at the issue of return predictability and rational pricing in a regression setting. They use a multi-beta CAPM to decompose the variance of the fitted values from a regression of returns on a set of instrumental variables into explained and unexplained components. If the ratio of explained to unexplained variance is large, they argue, then this is evidence that predictability is driven by rational variation in expected returns. Ferson and Harvey estimate this ratio for a number of different portfolios and find that it is reasonably close to one. Consequently, they conclude that the ability to predict returns is, for the most part, a product of rational market forces.

Although these prior studies advance our understanding of the economic basis for predictability, neither develops a general approach for studying predictability within a regression setting. The Hansen and Singleton approach is based on the strong assumptions of constant relative risk aversion and joint lognormality of asset returns and aggregate consumption growth. Ferson and Harvey avoid these assumptions, but their methodology is subject to errors-in-the-variables biases and does not impose the full set of restrictions implied by their pricing model.[3] In contrast, the restrictions on regression-based measures of predictability derived in this article are general in nature. The restrictions analyzed in prior studies of return predictability are special cases of the results presented below.

---

[2] Foster, Smith, and Whaley (1997), for example, take a very skeptical view of the empirical research on return predictability. They contend that a specification search over even a small number of forecasting variables can seriously bias standard procedures for drawing inferences and may produce spurious predictability that is of the same magnitude as that typically reported for stock and bond returns.

[3] The multi-beta CAPM examined by Ferson and Harvey (1991) implies restrictions on the intercept, slope coefficients, and R^2 of their predictive regressions. But their cross-sectional estimation approach leaves the intercept and slope coefficients unrestricted. As a result, their variance-ratio tests may have low power to detect violations of the restrictions implied by the pricing model.

1.1 Predicting returns in an efficient market

Following Hansen and Singleton (1982), consider a general model of a dynamic economy in which an unspecified number of rational investors trade assets in securities markets. Preferences in this economy may vary from one investor to the next and need not conform to any particular class of utility functions. Various nonstandard features, such as time and state nonseparability, are also permitted. It is assumed, however, that each of the investors in the economy solves a dynamic portfolio problem in order to arrive at an optimal holding of securities, and that there are no transactions costs, short sales constraints, or other barriers to trade.[4]

Let R_t denote the date t total return on any one of the traded assets. Under the stated assumptions, each investor's optimization problem yields a first-order condition for this asset of the form

$$\mathrm{E}[m_t R_t \mid \mathcal{F}_{t-1}] = 1, \tag{1}$$

where m_t is the investor's marginal rate of substitution at time t and $\mathcal{F}_{t-1}$ is the information set at time $t-1$. If markets are complete, then marginal rates of substitution are equated across all investors and m_t is uniquely determined. With incomplete markets, marginal rates of substitution differ across individuals, so the set of random variables that satisfy Equation (1) contains multiple elements. Every element of this set is, however, an admissible stochastic discount factor in the sense that it correctly prices all marketed payoffs.

The pricing relation in Equation (1) establishes a necessary condition for markets to be intertemporally efficient. It implies that, in an efficient market, the stochastic process defined by the product of the total return on an asset and the stochastic discount factor must be unpredictable using any of the information contained in $\mathcal{F}_{t-1}$. But this does not require that the returns themselves be unpredictable. In general, the only circumstance in which Equation (1) rules out the ability to predict returns is when the stochastic discount factor, m_t, is equal to a constant. If m_t is constant, then risk-neutral pricing prevails.[5]

Because Equation (1) must hold at each point in time, it can be used to derive explicit restrictions on the covariance between asset returns and variables in $\mathcal{F}_{t-1}$. Although these restrictions could be

[4] Similar assumptions are employed in a number of related studies. See, for example, Harrison and Kreps (1979), Kreps (1981), Hansen and Richard (1987), Hansen and Jagannathan (1991), and Cochrane and Hansen (1992).

[5] Although risk neutrality is sufficient to rule out predictability, risk aversion does not necessarily imply predictability. For example, Hansen and Singleton (1983) show that if asset returns and aggregate consumption growth are jointly lognormal, agents have constant relative risk aversion, and consumption growth is unpredictable, then asset returns will also be unpredictable.

developed in terms of R_t, the analysis is more tractable if we use excess returns. Let R_{ft} denote the total return on a reference asset. This reference asset will generally be a discount bond designed to proxy for the risk-free security, but any traded asset could be used. The important point is that, when combined with Equation (1), the relation $\mathrm{E}[m_t R_{ft} \mid \mathcal{F}_{t-1}] = 1$ implies that

$$\mathrm{E}[m_t r_t \mid \mathcal{F}_{t-1}] = 0, \tag{2}$$

where $r_t \equiv (R_t - R_{ft})$. Since Equation (2) is invariant with respect to scalar transformations of m_t, it is convenient to impose the normalization $\tilde{m}_t = m_t/\mathrm{E}[m_t]$. This yields

$$\mathrm{E}[\tilde{m}_t r_t \mid \mathcal{F}_{t-1}] = 0, \tag{3}$$

with $\mathrm{E}[\tilde{m}_t] = 1$. Given Equation (3), we can easily derive the restrictions on predictability that must be satisfied in an efficient market.

Take any instrumental variable, z_{t-1}, that is an element of the information set, $\mathcal{F}_{t-1}$. Multiply both sides of Equation (3) by z_{t-1} to get

$$\mathrm{E}[\tilde{m}_t r_t z_{t-1} \mid \mathcal{F}_{t-1}] = 0, \tag{4}$$

and apply the law of iterated expectations to obtain

$$\mathrm{E}[\tilde{m}_t r_t z_{t-1}] = 0. \tag{5}$$

Equation (5) places tight restrictions on the ability to predict excess returns. To see this, note that the unconditional analogue of Equation (3) implies the pricing relation

$$\mathrm{E}[r_t] = -\mathrm{cov}(\tilde{m}_t, r_t). \tag{6}$$

Thus we can use the familiar covariance decomposition to write Equation (5) as

$$\mathrm{cov}(r_t, z_{t-1}) = -\mathrm{E}[r_t]\mathrm{E}[z_{t-1}] - \mathrm{cov}(\tilde{m}_t, r_t z_{t-1}), \tag{7}$$

and substitute using Equation (6) for $\mathrm{E}[r_t]$ to obtain

$$\mathrm{cov}(r_t, z_{t-1}) = -\mathrm{cov}(\tilde{m}_t, r_t(z_{t-1} - \mu_z)), \tag{8}$$

where $\mu_z \equiv \mathrm{E}[z_{t-1}]$. This equation provides a general characterization of the restrictions on predictability that must be satisfied in an efficient market. It indicates that the ability to predict r_t, as measured by its unconditional covariance with z_{t-1}, has a well-defined theoretical value.

To see the intuition behind Equation (8), consider a simple exchange economy where r_t represents the excess payoff at time t on an investment in the asset of one unit of the consumption good at time $t-1$. In this setting, we can interpret the quantity $r_t z_{t-1}$ as the time t excess payoff on an investment in the asset of z_{t-1} units of the consumption good at time $t-1$. Now, keeping this interpretation in mind, consider the formula for the covariance between r_t and z_{t-1},

$$\mathrm{cov}(r_t, z_{t-1}) = \mathrm{E}[r_t(z_{t-1} - \mu_z)]. \tag{9}$$

The right side of Equation (9) measures the expected excess payoff on a trading strategy that invests $(z_{t-1} - \mu_z)$ in the asset at time $t-1$. In other words, the covariance between r_t and z_{t-1} is the expected excess payoff on a dynamic trading strategy that exploits the information conveyed by the realization of $(z_{t-1} - \mu_z)$.[6]

Because $r_t(z_{t-1} - \mu_z)$ is the excess payoff from a dynamic trading strategy, it has to be priced in the same manner as the excess payoff on any other asset in the economy. We can see this from Equation (8). The right side of this equation is simply the covariance between the excess payoff on the trading strategy and the normalized stochastic discount factor. This covariance measures the systematic risk of the trading rule defined by $(z_{t-1} - \mu_z)$. If investors are rational, then the ability to predict r_t must be consistent with the exposure to systematic risk that an investor takes on by following this trading rule.

1.2 Imposing the restrictions in a regression setting

The restrictions shown in Equation (8) have a simple, intuitive interpretation. But they are not directly applicable to most of the research on predicting asset returns. In practice, the most common approach used to document the ability to predict returns is to fit a linear regression model to the data [see, e.g., Keim and Stambaugh (1986), Campbell (1987), Campbell and Shiller (1988), Fama and French (1989), and Ferson and Harvey (1991)]. Therefore the predictable variation in excess returns is typically measured in terms of the slope coefficients and R^2 from multiple regression analysis. Although the focus thus far has been on the covariance between excess returns and variables in the information set, we can easily generalize the analysis so that it applies to regression-based measures of return predictability.

Let $\mathbf{z}_{t-1}$ denote a $(K-1) \times 1$ vector of instrumental variables. Most researchers assess whether excess returns are predictable by estimating a multiple regression model of the form

$$r_t = \mathbf{x}'_{t-1}\boldsymbol{b} + e_t, \tag{10}$$

where $\mathbf{x}_{t-1} \equiv [1,\ \mathbf{z}'_{t-1}]'$. We know from least-squares theory that the

[6] Hansen and Richard (1987) provide a useful characterization of dynamic trading strategies in their analysis of the role of conditioning information in tests of dynamic asset pricing models.

vector of regression coefficients in Equation (10) is given by

$$\boldsymbol{b} = \mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}]^{-1}\mathrm{E}[r_t\boldsymbol{x}_{t-1}], \tag{11}$$

and from Equation (5) we have the pricing relation

$$\mathrm{E}[r_t\boldsymbol{x}_{t-1}] = -\mathrm{cov}(\tilde{m}_t, r_t\boldsymbol{x}_{t-1}). \tag{12}$$

Substituting Equation (12) into the right side of Equation (11) yields

$$\boldsymbol{b} = -\mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}]^{-1}\mathrm{cov}(\tilde{m}_t, r_t\boldsymbol{x}_{t-1}). \tag{13}$$

If predictability is rational, then the regression coefficients must satisfy the restrictions shown in Equation (13).

A similar approach can be used to derive restrictions on the value of the regression R^2. Recall that the population R^2 measures the proportion of the variance of the dependent variable explained by the regressors. If we partition $\boldsymbol{b}$ as $[b_0, \boldsymbol{b}'_z]'$ where b_0 is the intercept and $\boldsymbol{b}_z$ is the vector of slope coefficients, then the population R^2 for the model can be written as

$$R^2 = \left(\frac{\boldsymbol{b}'_z\boldsymbol{\Sigma}_{zz}\boldsymbol{b}_z}{\sigma_r^2}\right), \tag{14}$$

where $\boldsymbol{\Sigma}_{zz}$ and σ_r^2 denote the variance-covariance matrix of $\boldsymbol{z}_{t-1}$ and the variance of r_t. Substituting for $\boldsymbol{b}_z$ using the formula implied by Equation (13) yields the restriction

$$R^2 = \left(\frac{\boldsymbol{\sigma}'_{m,rz}\boldsymbol{\Sigma}_{zz}^{-1}\boldsymbol{\sigma}_{m,rz}}{\sigma_r^2}\right). \tag{15}$$

where $\boldsymbol{\sigma}_{m,rz} \equiv \mathrm{cov}(\tilde{m}_t, r_t(\boldsymbol{z} - \boldsymbol{\mu}_z))$. This is the only value of the population R^2 that is consistent with rational pricing.

Since the restrictions in Equations (13) and (15) summarize the implications of market efficiency for predictive regressions, they provide a convenient frame of reference for researchers who use regression analysis to investigate return predictability. In particular, they make it easy to assess whether the evidence of predictability uncovered using regression analysis is consistent with any given specification of $\tilde{m}_t$. If it is, then the view that predictability is driven by rational variation in expected returns gains credibility. If not, then we may want to consider whether alternative mechanisms, like market inefficiencies or collective data mining, explain the observed ability to predict returns. Regardless of the outcome, however, examining the regression evidence from an asset pricing perspective can only enhance our understanding of the basis for predictability.

Given a specification for $\tilde{m}_t$, all of the parameters in Equations (13) and (15) can be estimated from financial market data. Consequently it is straightforward to construct direct tests of the restrictions. We can also use the restrictions to identify the linear specification for predicting excess returns that is implied by a particular stochastic discount factor. This follows by noting that, for a given set of instrumental variables, each specification of $\tilde{m}_t$ identifies the regression coefficients needed to construct a linear forecasting model. Studying the time-series behavior of forecasts constructed in this manner may yield valuable insights into the role of predictability in financial markets. We might, for example, want to compare the accuracy of the forecasts implied by different pricing models during various stages of the business cycle.

Before turning to a discussion of the econometric methodology, it is useful to note the link between the foregoing analysis and the work of Hansen and Jagannathan (1991). They construct the linear projection of the true stochastic discount factor onto the set of traded asset payoffs and prove that it has minimum variance in the set of all admissible stochastic discount factors. Although the coefficients in this projection cannot be estimated directly without time-series data on the true stochastic discount factor, we can infer them by invoking standard no-arbitrage arguments. Similarly, in the foregoing analysis, we infer the coefficients in the linear projection of returns onto the instrumental variables using the no-arbitrage pricing results for admissible stochastic discount factors.

As shown in Appendix A, the restrictions in Equations (13) and (15) are closely connected to the Hansen and Jagannathan variance bound and the square of the related distance measure [see Hansen and Jagannathan (1997)]. We can express both as quadratic forms in the unrestricted and restricted vectors of regression coefficients. This connection is convenient because it allows us to use the Hansen and Jagannathan distance measure as a diagnostic in the subsequent empirical tests. Unlike traditional chi-square statistics, the distance measure does not reward variability in the candidate stochastic discount factor. This makes it an appealing metric for evaluating how well a candidate stochastic discount factor explains the observed ability to predict returns.

2. Data and Econometric Methods

This section outlines an approach for testing the restrictions derived in Section 1. The procedure is straightforward. First, select a candidate stochastic discount factor, m_t^c, and estimate a set of predictive regressions under the restrictions it implies. Then, test whether the intercept and slope coefficients from these restricted regressions are significantly different from those obtained via ordinary least-squares (OLS) estimation. A number of candidate stochastic discount factors are employed in the tests. They include both consumption-based and linear factor specifications. The estimation and tests are performed using Hansen's (1982) generalized method of moments (GMM).

2.1 The data

The empirical tests center on the ability to predict monthly excess returns on a set of 10 stock portfolios. These portfolios are formed on the basis of market capitalization using the firms listed on the New York Stock Exchange (NYSE). The portfolio returns begin in July 1963 and end in December 1991 (342 observations). The starting date for the sample is dictated by the availability of the data needed to construct some of the candidate stochastic discount factors. In particular, returns on three factor-mimicking portfolios are required. The first—a proxy for the excess return on the market portfolio—is the monthly return on the value-weighted NYSE index minus the return on the Treasury bill that is closest to 30 days to maturity. The other two, which proxy for size and book-to-market factors, are constructed following the procedures of Fama and French (1993).

The instrumental variables are selected based on the findings of prior research. A total of five instruments are used: the excess return on the equally weighted NYSE index (*xrew*), a dummy variable for the month of January (*jan*), the 1-month return from holding a 90-day Treasury bill less the return on a 30-day bill (*term*), the yield on Moody's Baa rated bonds less the yield on Moody's Aaa rated bonds (*prem*), and the dividend yield on the S&P 500 stock index less the return on a 30-day Treasury bill (*xdiv*). Additional details on the data are provided in Appendix B.

Table 1 reports descriptive statistics for the dataset. Panel A documents the properties of the nominal excess returns on the size portfolios. The smallest decile has the highest mean excess return and the largest return standard deviation. Both the mean and standard deviation of the returns steadily decrease as the market capitalization of the firms in the portfolio becomes larger. A similar pattern is evident in the first-order and twelfth-order autocorrelations for these portfolios. Panel B summarizes the properties of the consumption growth series, the factors, and the instrumental variables. Two of the instruments, *prem* and *xdiv*, exhibit a high degree of persistence, but the autocorrelations decay at a rate that is consistent with the assumption that each time series is stationary.

Table 1
Summary statistics for the dataset

Panel A: Size portfolio returns

Decile	Mean	Std. Dev.	Autocorrelations $\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_6$	$\hat{\rho}_{12}$
1	0.00900	0.0739	0.128	−0.006	−0.063	−0.025	0.215
2	0.00807	0.0655	0.145	−0.015	−0.042	−0.035	0.176
3	0.00715	0.0614	0.143	−0.024	−0.034	−0.015	0.135
4	0.00770	0.0585	0.162	−0.035	−0.031	−0.027	0.093
5	0.00642	0.0561	0.146	−0.041	−0.019	−0.035	0.090
6	0.00684	0.0545	0.145	−0.023	−0.033	−0.048	0.057
7	0.00582	0.0529	0.127	−0.025	−0.010	−0.035	0.020
8	0.00601	0.0512	0.091	−0.053	−0.031	−0.036	0.028
9	0.00469	0.0487	0.083	−0.047	−0.027	−0.046	0.001
10	0.00348	0.0429	−0.002	−0.026	−0.006	−0.047	0.036

Panel B: Consumption growth, factors, and instruments

Variable	Mean	Std. Dev.	Autocorrelations $\hat{\rho}_1$	$\hat{\rho}_2$	$\hat{\rho}_3$	$\hat{\rho}_6$	$\hat{\rho}_{12}$
C_t/C_{t-1}	1.00096	0.0073	−0.377	0.061	0.170	0.031	−0.085
r_{mt}	0.00419	0.0452	0.052	−0.040	−0.005	−0.050	0.025
r_{st}	0.00267	0.0289	0.192	0.067	−0.002	0.107	0.225
r_{ot}	0.00395	0.0256	0.183	0.063	−0.001	0.106	0.072
$xrew_{t-1}$	0.00623	0.0549	0.153	−0.035	−0.028	−0.034	0.082
$term_{t-1}$	0.00066	0.0011	0.296	0.059	0.026	−0.024	−0.035
$prem_{t-1}$	0.00092	0.0004	0.966	0.920	0.886	0.788	0.607
$xdiv_{t-1}$	−0.00212	0.0017	0.912	0.844	0.778	0.679	0.480

In panel B, C_t/C_{t-1} is consumption growth, r_{mt} is the excess return on the value-weighted market, r_{st} is the return on the mimicking portfolio for size, r_{ot} is the return on the mimicking portfolio for book-to-market equity, $xrew_{t-1}$ is the lagged excess return on the equally weighted market, $term_{t-1}$ is the lagged excess return on a 90-day Treasury bill, $prem_{t-1}$ is the lagged yield spread between Baa and Aaa rated bonds, and $xdiv_{t-1}$ is the lagged excess dividend yield on the S&P 500 stock index. See Appendix B for details on the construction of the dataset.

2.2 The consumption-based specifications

Consumption smoothing plays a key role in intertemporal asset pricing theory and is one potential explanation for return predictability. It is natural, therefore, to include consumption-based models in the empirical tests. The first model examined is the standard power utility specification,

$$m_t^c = \frac{1}{1+\delta}\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma}, \tag{16}$$

where C_t denotes per-capita aggregate consumption at time t, δ is the rate of time preference, and γ measures relative risk aversion. A number of authors, beginning with Hansen and Singleton (1982), have studied this specification using monthly data on per-capita consumption expenditures. These studies generally find that the model is rejected by the data. Nonetheless, its performance provides a useful baseline for comparison.

The second consumption-based model allows for habit persistence in consumption expenditures. This class of preferences is of interest because of the evidence that habit persistence may provide a partial solution to the equity premium puzzle. Abel (1990), for example, develops a model where each agent's utility depends on his level of consumption relative to some time-varying benchmark. In equilibrium, the model implies the stochastic discount factor

$$m_t^c = \frac{1}{1+\delta}\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \Bigg/ \left(\frac{C_{t-1}}{C_{t-2}}\right)^{1-\gamma}. \tag{17}$$

Abel finds that this model can not only generate large equity risk premiums, but also substantial time-series variation in conditional expected rates of return. It is the latter finding that motivates its inclusion here.

The third consumption-based model generalizes the conventional expected utility specification by introducing a recursive preference structure. Under this model, current utility depends on both current consumption and the certainty equivalent of future lifetime utility. This generalization is interesting because it allows us to partially separate the effects of intertemporal substitution and risk aversion. Following Epstein and Zin (1991), we can write the stochastic discount factor implied by this model as

$$m_t^c = \left[\frac{1}{1+\delta}\left(\frac{C_t}{C_{t-1}}\right)^{p-1}\right]^{\alpha/p}\left(\frac{1}{R_{mt}}\right)^{1-\alpha/p}, \tag{18}$$

where R_{mt} denotes the total return on the market portfolio, $1-\alpha$ determines relative risk aversion for timeless consumption gambles, and $1/(1-p)$ measures the intertemporal elasticity of substitution for deterministic consumption. For $\alpha/p = 1$, Equation (18) reduces to the standard power utility model. Setting $\alpha = 0$ yields the log utility model of Rubinstein (1976).

2.3 The linear factor specifications

Since linear factor models are the main alternative to consumption-based specifications, they are included in the analysis as well. The first model is a conditional version of the Sharpe (1964) and Lintner (1965) CAPM. At the conditional level, the CAPM implies the pricing relation

$$\mathrm{E}[r_t \,|\, \mathcal{F}_{t-1}] = \beta_{m,t-1}\mathrm{E}[r_{mt} \,|\, \mathcal{F}_{t-1}], \tag{19}$$

where $\beta_{m,t-1} \equiv \mathrm{cov}(r_t, r_{mt} \,|\, \mathcal{F}_{t-1})/\mathrm{Var}(r_{mt} \,|\, \mathcal{F}_{t-1})$ is the conditional beta of the asset. To analyze the conditional CAPM in its most general form, we have to explicitly model conditional expectations. This raises a couple of concerns. First, by modeling conditional expectations, we introduce a flexibility in choosing functional forms that is absent from the consumption-based specifications. Second, it is well known that large-sample chi-square tests have a tendency to reward variability in the stochastic discount factor, and explicitly modeling conditional expectations is likely to increase this variability.

The two specifications of the conditional CAPM examined in the empirical tests are designed to mitigate these concerns by eliminating the need to model conditional expectations. This is accomplished by imposing additional structure on the model. Using the definition of the conditional beta, we have

$$\begin{aligned}(\mathrm{E}[r_{mt}^2 \mid \mathcal{F}_{t-1}] - \mathrm{E}[r_{mt} \mid \mathcal{F}_{t-1}]^2)\beta_{m,t-1} \\ = \mathrm{E}[r_t r_{mt} \mid \mathcal{F}_{t-1}] - \mathrm{E}[r_t \mid \mathcal{F}_{t-1}]\mathrm{E}[r_{mt} \mid \mathcal{F}_{t-1}].\end{aligned} \tag{20}$$

Since $\mathrm{E}[r_t \mid \mathcal{F}_{t-1}] - \beta_{m,t-1}\mathrm{E}[r_{mt} \mid \mathcal{F}_{t-1}] = 0$ under the conditional CAPM, Equation (20) reduces to $\mathrm{E}[r_{mt}^2 \mid \mathcal{F}_{t-1}]\beta_{m,t-1} = \mathrm{E}[r_t r_{mt} \mid \mathcal{F}_{t-1}]$. Substituting this expression for the conditional beta into Equation (19) and rearranging terms reveals that the conditional CAPM implies the stochastic discount factor

$$m_t^c = 1 - \lambda_{m,t-1} r_{mt}, \tag{21}$$

where $\lambda_{m,t-1} \equiv \mathrm{E}[r_{mt} \mid \mathcal{F}_{t-1}]/\mathrm{E}[r_{mt}^2 \mid \mathcal{F}_{t-1}]$ is the price of risk. Thus one way to avoid modeling conditional expectations is to assume that the price of risk is constant. This approach is used in the first set of empirical tests of this model.

The constant price of risk specification implies that the ability to predict the return on an asset is determined by changes in the conditional covariance of its return with the market return. Although this is acceptable from a theoretical perspective, prior research suggests that it may not be empirically plausible. Both Ferson and Harvey (1991) and Ferson and Korajczyk (1995) argue that return predictability is primarily driven by time-series variation in risk premiums.[7] An alternative specification of the conditional CAPM that is consistent with this view is also examined. It assumes that conditional betas are constant. The resulting pricing relation is

$$\mathrm{E}[r_t \mid \mathcal{F}_{t-1}] = \beta_m \mathrm{E}[r_{mt} \mid \mathcal{F}_{t-1}], \tag{22}$$

where $\beta_m \equiv \mathrm{E}[r_t r_{mt}]/\mathrm{E}[r_{mt}^2]$. Note that, since the betas are asset specific, we have to depart from the stochastic discount factor represen-

[7] Ferson and Korajczyk (1995), for example, attribute less than 1% of the predictable variation in returns to changing conditional betas.

tation of the model in order to impose the constant conditional beta assumption. As a result, the distance measure diagnostics discussed at the end of Section 1 are not applicable.

To obtain the restrictions on predictability implied by the constant conditional beta specification, we multiply Equation (22) by $\boldsymbol{x}_{t-1}$, take unconditional expectations, and multiply by $\mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}]^{-1}$. This yields the relation

$$\boldsymbol{b} = \beta_m \boldsymbol{b}_m, \tag{23}$$

where $\boldsymbol{b}_m \equiv \mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}]^{-1}\mathrm{E}[r_{mt}\boldsymbol{x}_{t-1}]$ is the coefficient vector in a regression of the market return on the set of instrumental variables. The constant conditional beta specification of the conditional CAPM restricts the predictable component of asset returns to be proportional to the predictable component of the market return, with the constant of proportionality being the asset's conditional beta with respect to the market.

Given the form of Equation (23), it appears that the constant conditional beta specification may have a built-in advantage in tests for rational predictability. In effect, Equation (23) takes the predictable component of the market return as given. We know, therefore, that if the market return is predictable, then the empirical specification of the model will generate predictability for the test assets. This is not the case for the consumption-based models considered earlier. Of course, just because the constant conditional beta specification generates predictability, there is no guarantee that it will perform well in the empirical tests. It still has to explain cross-sectional differences in the ability to predict returns.

In addition to the conditional CAPM, the tests examine a conditional version of the Fama and French (1993) model. It nests the conditional CAPM as a special case. In its unconditional form, the model says that three factors—the excess return on the market, r_{mt}, the return on a zero-investment portfolio designed to capture risk associated with size (market capitalization), r_{st}, and the return on a zero-investment portfolio designed to capture risk associated with the book-to-market equity, r_{ot}—explain cross-sectional differences in average returns. According to Fama and French, this version of the model fits their sample of return data much better than the unconditional CAPM.

The conditional version of the Fama and French three-factor model implies the stochastic discount factor

$$m_t^c = 1 - \lambda_{m,t-1} r_{mt} - \lambda_{s,t-1} r_{st} - \lambda_{o,t-1} r_{ot}, \tag{24}$$

where $\lambda_{m,t-1}$, $\lambda_{s,t-1}$, and $\lambda_{o,t-1}$ are the prices of market risk, size risk, and book-to-market equity risk, respectively. To obtain the empirical specifications of the model, we employ the same constant price of risk and constant conditional beta assumptions used for the conditional CAPM. The stochastic discount factor and restrictions implied by the constant price of risk assumption are readily apparent. The constant conditional beta assumption yields the relation

$$\boldsymbol{b} = \beta_m \boldsymbol{b}_m + \beta_s \boldsymbol{b}_s + \beta_o \boldsymbol{b}_o, \tag{25}$$

a multifactor generalization of the restriction shown in Equation (23).

In general, fitting the Fama and French model should be viewed as a method of providing diagnostics for the conditional CAPM rather than as a formal test of the model itself. The reason is that, if market participants are rational, then high-risk firms will tend to have lower market values than low-risk firms, regardless of what asset pricing model actually generates returns. Thus, as long as there is no positive relation between the operating size of the firm and its risk, we should expect to find that small capitalization firms are riskier on average than large capitalization firms. As a result, the pricing errors generated by a misspecified pricing model can be expected to exhibit size-related patterns. If we test the three-factor model and find that it explains return predictability better than the CAPM, this suggests that the pricing errors for the CAPM are to some extent explained by risk. But it does not mean that we can interpret r_{st} and r_{ot} as priced factors in the traditional sense.[8]

2.4 Estimation and tests

Let r_{it} denote the excess return on portfolio i. To illustrate the approach used to estimate the parameters and test the restrictions, consider the following example. If we set $\alpha = 0$ in Equation (18), we obtain $1/R_{mt}$ as our candidate stochastic discount factor. Since m_t^c has no unknown parameters, the econometric system estimated for each portfolio is

$$\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i) = \begin{pmatrix} m_t^c - \mu_{m^c} \\ (r_{it} - \boldsymbol{x}'_{t-1}\boldsymbol{b}_{iu})\boldsymbol{x}_{t-1} \\ (-r_{it}(m_t^c - \mu_{m^c}) - \mu_{m^c}\boldsymbol{x}'_{t-1}\boldsymbol{b}_{ir})\boldsymbol{x}_{t-1} \end{pmatrix}, \tag{26}$$

[8] Berk (1995) makes the same point in his critique of the literature on the size effect. He notes that the market capitalization of a firm is a function of at least two things: bigger firms (in an operating sense) tend to have relatively higher market values; and riskier firms tend to have relatively lower market values. This means that, in the cross section, market capitalization will generally be inversely correlated with unmeasured risk. So if we test two pricing models, each of which omits a different priced factor, then size will proxy for a different factor in each of the tests. As a result, it is inappropriate to interpret market-based measures of size as factors in the traditional sense of the arbitrage pricing theory.

where $m_t^c = 1/R_{mt}$, $\boldsymbol{y}_{it} \equiv [R_{mt},\ r_{it},\ \boldsymbol{x}'_{t-1}]'$, and $\boldsymbol{\theta}_i \equiv [\mu_{m^c},\ \boldsymbol{b}'_{iu},\ \boldsymbol{b}'_{ir}]'$.[9]

The first equation in Equation (26) identifies the mean of the candidate stochastic discount factor, the next K equations identify the intercept and slope coefficients in the unrestricted regression of r_{it} on $\boldsymbol{x}_{t-1}$, and the final K equations impose the restrictions in Equation (13) on the regression coefficients. Because the econometric specification in Equation (26) is exactly identified, the GMM estimator of $\boldsymbol{\theta}_i$ is the value that sets $1/T\sum_{t=1}^{T}\boldsymbol{h}(\cdot)$ equal to zero, and the unrestricted estimates of the regression coefficients are the same as the OLS estimates.

Tests of the restrictions are based on the limiting distribution of the GMM estimator. This distribution takes the form

$$\sqrt{T}(\hat{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i) \xrightarrow{d} N(\mathbf{0}, (\boldsymbol{D}_i'\boldsymbol{S}_i^{-1}\boldsymbol{D}_i)^{-1}), \tag{27}$$

where

$$\boldsymbol{D}_i \equiv \mathrm{E}\left[\frac{\partial \boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i'}\right] \quad \text{and} \quad \boldsymbol{S}_i \equiv \sum_{j=-\infty}^{\infty} \mathrm{E}[\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i)\boldsymbol{h}(\boldsymbol{y}_{i,t-j}, \boldsymbol{\theta}_i)']. \tag{28}$$

If we let $\boldsymbol{\Omega}_{ib}$ denote the lower right $2K \times 2K$ submatrix of $\boldsymbol{\Omega}_i \equiv (\boldsymbol{D}_i'\boldsymbol{S}_i^{-1}\boldsymbol{D}_i)^{-1}$, then the Wald statistic for testing $\boldsymbol{b}_{iu} = \boldsymbol{b}_{ir}$ can be written as

$$W_{iT} = T(\hat{\boldsymbol{b}}_{iu} - \hat{\boldsymbol{b}}_{ir})'(\boldsymbol{Q}'\boldsymbol{\Omega}_{ib}\boldsymbol{Q})^{-1}(\hat{\boldsymbol{b}}_{iu} - \hat{\boldsymbol{b}}_{ir}), \tag{29}$$

where $\boldsymbol{Q} \equiv [\boldsymbol{I}_K,\ -\boldsymbol{I}_K]'$ and $\boldsymbol{I}_K$ denotes a $K \times K$ identity matrix. Under the null hypothesis, W_{iT} converges to a central chi-square random variable with K degrees of freedom. Although in general the matrix $\boldsymbol{\Omega}_{ib}$ is unknown, it can be replaced by a consistent estimator without affecting the limiting distribution of the test statistic. Appendix C discusses how this estimator is constructed.

This approach to GMM estimation and testing is slightly different from what is typically seen in the empirical asset pricing literature. The usual strategy is to explicitly impose the restriction $\boldsymbol{b}_{iu} = \boldsymbol{b}_{ir}$ in the econometric specification and construct a test based on the set of overidentifying restrictions. Note, however, that the two approaches are asymptotically equivalent, so the choice of which to use is primarily a matter of convenience. The advantage of the Wald test is that (i) we can construct it without performing a two-stage nonlinear mini-

---

[9] The subsequent analysis assumes that $\boldsymbol{y}_t$ is generated by a stationary and ergodic process, that the second moment matrix of $\boldsymbol{h}(\cdot)$ exists and is finite, and that the regularity conditions in Hansen (1982) are satisfied. These assumptions are sufficient for the asymptotic distribution theory to hold. Technically the stationarity assumption rules out the use of fixed instruments like the January dummy. But the asymptotic theory for GMM remains valid under other assumptions, such as strong mixing, that are consistent with the use of dummy variables.

mization; and (ii) it facilitates a direct comparison of the unrestricted (OLS) and restricted regression coefficients. Of course, the distributional properties of the two tests could differ in finite samples. But this is not a concern for the dataset used in this study.[10]

In addition to the Wald test, an estimate of the Hansen and Jagannathan distance is computed for each portfolio. This estimate is given by

$$\hat{d}_{iT} = [(\hat{\boldsymbol{b}}_{ir} - \hat{\boldsymbol{b}}_{iu})' \hat{\boldsymbol{G}}_i \hat{\boldsymbol{\Sigma}}_{ix}^{-1} \hat{\boldsymbol{G}}_i (\hat{\boldsymbol{b}}_{ir} - \hat{\boldsymbol{b}}_{iu})]^{\frac{1}{2}}, \tag{30}$$

where $\hat{\boldsymbol{\Sigma}}_{ix}$ and $\hat{\boldsymbol{G}}$ are the sample analogues of $\boldsymbol{\Sigma}_{ix} \equiv \mathrm{E}[r_{it}^2(\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1})]$ and $\boldsymbol{G} \equiv \mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}]$, respectively (see Appendix A). Unlike the Wald test in Equation (29), the distance measure uses a metric that is independent of the functional form of the stochastic discount factor. This makes it useful for comparing the performance of different choices of $\tilde{m}_t^c$.

The econometric analysis for the case where m_t^c contains unknown parameters is similar. The only real difference is that the econometric system in Equation (26) is augmented with the moment equations necessary to identify these parameters. In general, the number and complexity of these additional equations varies depending upon the functional form of $\tilde{m}_t^c$. Although some additional changes are required to accommodate the constant conditional beta specifications, these are straightforward as well. For instance, the constant conditional beta version of the conditional CAPM is estimated using the system

$$\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i) \quad = \quad \begin{pmatrix} (r_{it} - \beta_{im} r_{mt}) r_{mt} \\ (r_{it} - \boldsymbol{x}'_{t-1}\boldsymbol{b}_{iu})\boldsymbol{x}_{t-1} \\ (\beta_{im} r_{mt} - \boldsymbol{x}'_{t-1}\boldsymbol{b}_{ir})\boldsymbol{x}_{t-1} \end{pmatrix}, \tag{31}$$

where $\boldsymbol{y}_{it} = [r_{mt},\ r_{it},\ \boldsymbol{x}'_{t-1}]'$ and $\boldsymbol{\theta}_i = [\beta_{im},\ \boldsymbol{b}'_{iu},\ \boldsymbol{b}'_{ir}]'$. The first equation identifies the portfolio beta, the next K equations identify the unrestricted (OLS) coefficients in the predictive regression, and the final K equations identify the restricted regression coefficients.

3. The Empirical Evidence

The first issue addressed by the empirical tests is whether the ability to predict real excess returns on the size portfolios is consistent with

[10] A modified version of the bootstrap [see Efron (1979)] was used to examine the finite sample properties of the tests for data satisfying the constant conditional beta version of the conditional CAPM. The results suggest that the Wald test generally exhibits better finite sample properties than the test for overidentifying restrictions. Moreover, when used in conjunction with the asymptotic critical values, the Wald test leads to mildly conservative inferences. A separate appendix that describes the details of the bootstrap experiments is available on request.

the restrictions implied by any of the consumption-based asset pricing models. Since the candidate stochastic discount factors implied by these models contain unknown parameters, it is worthwhile to examine the effect of changing these parameters on the empirical results. This is accomplished by choosing a number of different settings for the parameters, fitting the predictive regressions using OLS, and comparing the OLS estimates to those that result from explicitly imposing the pricing restrictions implied by the candidate stochastic discount factor. After documenting the performance of the consumption-based models, the analysis shifts to the linear factor specifications. Their performance is evaluated using the Wald test and the Hansen and Jagannathan distance measure.

3.1 Results for the consumption-based specifications

The first model examined is the power utility specification. Figure 1 presents the results for this model. The graph in the upper left corner shows the OLS estimates of the slope coefficients and R^2 for the predictive regressions. As expected, these estimates indicate that the real excess returns on the size portfolios exhibit a significant degree of predictable variation over time. Although the t-ratios of the slope estimates are not shown, the majority of them are greater than two in magnitude. As for the regression R^2, it ranges from a low of 9.4% for the portfolio with the largest stocks to a high of 20.3% for the portfolio with the smallest stocks. These findings are consistent with Harvey (1989).

The three remaining graphs in Figure 1 illustrate how the estimates of the slope coefficients and regression R^2 change when the restrictions implied by the power utility specification are imposed on the regressions. Each graph corresponds to a different level of relative risk aversion. With a relative risk aversion of five, the restricted estimates of the slopes are close to zero. This indicates that the power utility model implies almost no predictable variation in returns. There is an increase in the magnitude of the slope estimates as the relative risk aversion gets larger. But even a relative risk aversion of 50 does little to alter the results. In every case the R^2 for the restricted regression is below 1%. There is nothing in the results to suggest that the power utility model can explain the observed return predictability.

Abel's (1990) habit persistence model is examined next. The results for this model are shown in Figure 2. Again, the first graph shows the OLS estimates of the regression coefficients. The remaining three graphs show the restricted estimates for different settings of the curvature parameter γ. On balance, the performance of this model is comparable to that of the power utility specification. The restricted estimates of the slopes and R^2 are far removed from the unrestricted

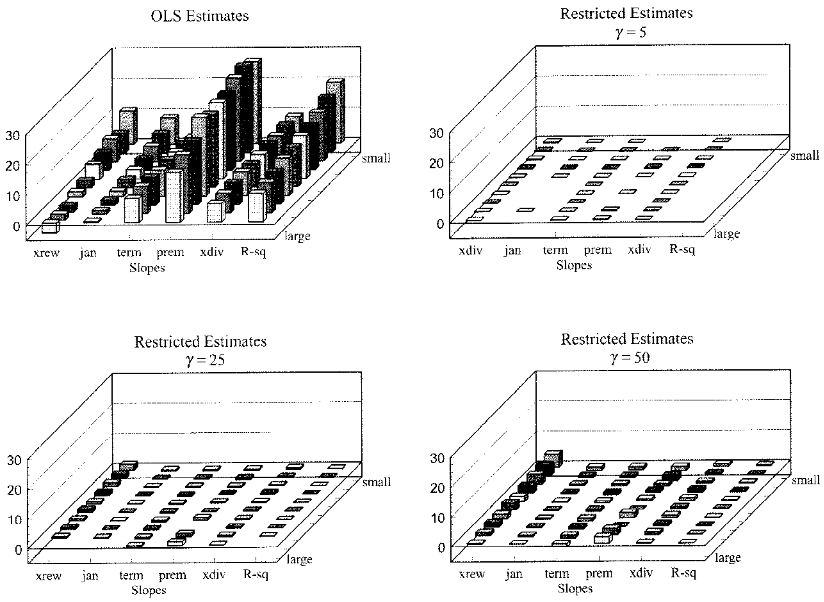

Figure 1
Regression parameters implied by the power utility specification
This figure illustrates the effect of imposing the restrictions implied by the power utility specification on the predictive regressions. The first graph shows the OLS estimates of the slopes and regression R^2. The next three graphs show the restricted estimates of these quantities for three different levels of relative risk aversion (γ = 5, 25, and 50). In all cases, the portfolio market capitalization decreases from the front to back of the graph. The instrumental variables are the lagged excess return on the equally weighted NYSE index ($xrew$), a January dummy (jan), the lagged excess return on a 90-day T-bill ($term$), the lagged yield spread between Baa and Aaa rated bonds ($prem$), and the lagged excess dividend yield on the S&P 500 stock index ($xdiv$). See Table 1 for a description of the data.

estimates. Like the power utility specification, the habit persistence model fails to generate the return predictability that shows up in the unrestricted regressions.

The last consumption-based model examined is the recursive utility specification of Epstein and Zin (1991). Figure 3 presents the results for this model. Since the model has two unknown parameters (α and p), the effect of changing each is examined separately. Panel A shows the OLS and restricted estimates of the regression coefficients for high intertemporal elasticity of substitution ($p = -0.25$; $\alpha = -0.125,\ -0.5,$ and -0.75). The restricted estimates for low intertemporal elasticity of substitution ($p = -4$; $\alpha = -2,\ -8,$ and -12) are shown in panel B. The settings of α are chosen to yield α/p ratios of 0.5, 2, and 3. Values of this ratio less than (greater than) one imply a preference for early (late) resolution of uncertainty.

Despite the increased flexibility afforded by the Epstein and Zin

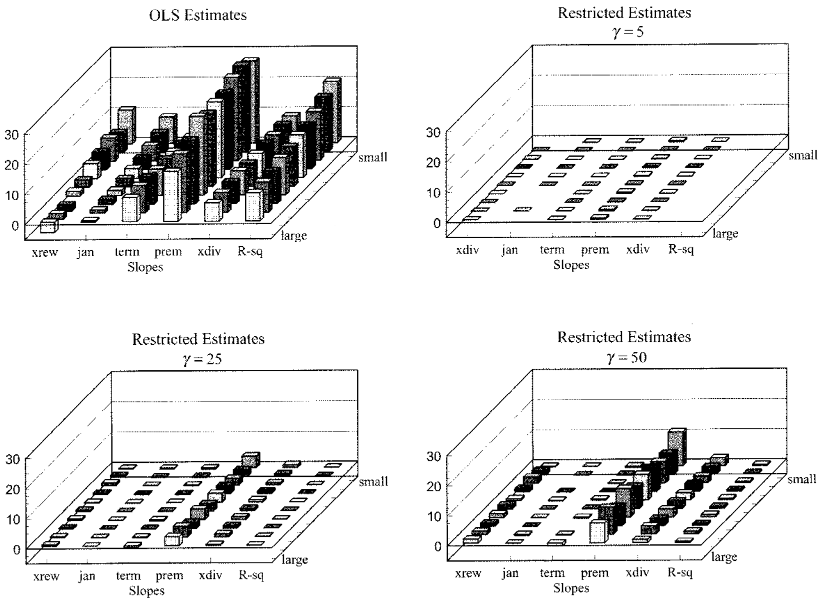

Figure 2
Regression parameters implied by the habit persistence specification
This figure illustrates the effect of imposing the restrictions implied by the habit persistence specification on the predictive regressions. The first graph shows the OLS estimates of the slopes and regression R^2. The next three graphs show the restricted estimates of these quantities for three different levels of the curvature parameter (γ = 5, 25, and 50). In all cases, the portfolio market capitalization decreases from the front to back of the graph. The instrumental variables are the excess return on the equally weighted NYSE index (*xrew*), a January dummy (*jan*), the excess return on a 90-day T-bill (*term*), the yield spread between Baa and Aaa rated bonds (*prem*), and the excess dividend yield on the S&P 500 stock index (*xdiv*). See Table 1 for a description of the data.

model, its performance is no better than that of the two previous consumption-based specifications. Once again, the graphs of the restricted estimates remain flat, regardless of the values chosen for the model parameters. It seems that breaking the link between risk aversion and intertemporal elasticity of substitution does little to improve the poor fit of the consumption-based models. The evidence suggests that we need to focus on other aspects of these models if we hope to explain the predictability displayed by the size portfolio returns.

Although the findings for the consumption-based models may seem unremarkable in light of the statistical rejections of these models reported in the asset pricing literature, they do serve to illustrate that predictability almost certainly plays a key role in these rejections. Most empirical tests of equilibrium asset pricing models are based on overidentifying restrictions that require the disturbance term for the model to be uncorrelated with variables in the information set of investors.

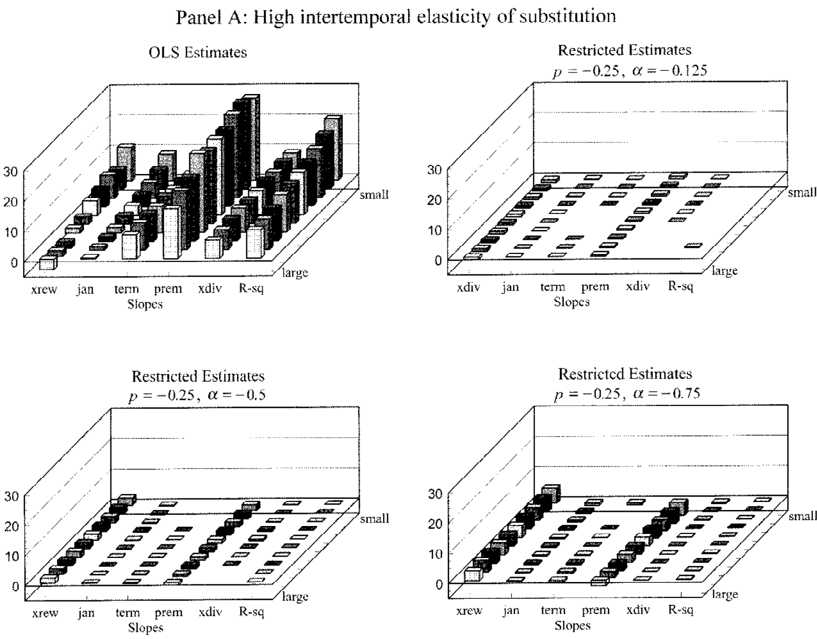

Figure 3
Regression parameters implied by the recursive utility specification
This figure illustrates the effect of imposing the restrictions implied by the recursive utility specification on the predictive regressions. Panel A shows the OLS estimates of the slopes and regression R^2 along with the restricted estimates for high intertemporal elasticity of substitution ($p = -0.25$) and three levels of risk aversion ($\alpha = -0.125, -0.5$, and -0.75). Continued next page.

The low values of the regression R^2 generated by the consumption-based specifications point to a lack of predictability as the driving force behind these rejections. The results suggest, in other words, that excess returns are too predictable to be consistent with the pricing implications of these models.

3.2 The log utility specification

The inability of the power utility, habit persistence, and recursive utility specifications to fit the data suggests that explaining predictability within a consumption-based pricing framework poses a significant challenge. But this finding is not unanticipated. We might suspect that the linear factor models, which can be estimated without consumption data, have a better chance of explaining the observed ability to predict returns. Before turning to the factor models, however, it is useful to examine the log utility specification. Since it can also be es-

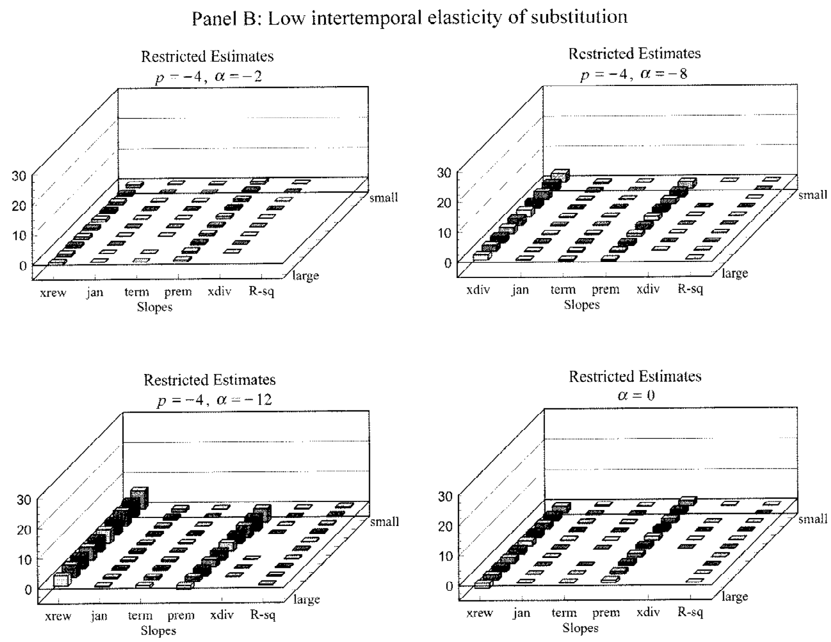

Figure 3 (continued)
Regression parameters implied by the recursive utility specification
Panel B shows the restricted estimates for low intertemporal elasticity of substitution ($p = -4$) and three levels of risk aversion ($\alpha = -2, -8$, and -12). The final graph in panel B shows the restricted estimates for the log utility specification ($\alpha = 0$). In all cases, the portfolio market capitalization decreases from the front to back of the graph. The instrumental variables are the excess return on the equally weighted NYSE index ($xrew$), a January dummy (jan), the excess return on a 90-day T-bill ($term$), the yield spread between Baa and Aaa rated bonds ($prem$), and the excess dividend yield on the S&P 500 stock index ($xdiv$). See Table 1 for a description of the data.

timated without consumption data, its performance can be used as a benchmark for the remaining tests.

The log utility specification is obtained by setting $\alpha = 0$ in the recursive utility model. The performance of this specification is evaluated using the Wald test and the Hansen and Jagannathan distance measure. Nominal rather than real excess returns are used to facilitate comparisons with the linear factor specifications. Table 2 presents the results. The tests in panel A treat each of the portfolios separately. The test in panel B is for a joint econometric specification that combines all 10 portfolios.

The first column in panel A reports the values of the regression R^2 for the unrestricted regressions. These values, which range from a low of 9.3% for decile 10 to a high of 20.3% for decile 1, are virtually identical to those obtained earlier, so the switch from real to nomi-

Table 2
Tests of the restrictions implied by the log utility model

Panel A: Individual portfolio tests

	R^2 (%)		Unrestricted minus restricted coefficient						t-ratio of difference						Wald test and HJ distance		
Decile	$\hat{R}_u^2$	$\hat{R}_r^2$	$\Delta\hat{b}_0$	$\Delta\hat{b}_1$	$\Delta\hat{b}_2$	$\Delta\hat{b}_3$	$\Delta\hat{b}_4$	$\Delta\hat{b}_5$	$t_{\Delta\hat{b}_0}$	$t_{\Delta\hat{b}_1}$	$t_{\Delta\hat{b}_2}$	$t_{\Delta\hat{b}_3}$	$t_{\Delta\hat{b}_4}$	$t_{\Delta\hat{b}_5}$	W_T	p-value	$\hat{d}_T$
1	20.3	0.0	−1.19	13.55	8.46	8.86	25.59	8.89	−1.39	1.45	3.98	3.02	2.28	3.73	65.1	0.000	0.318
2	18.1	0.0	−1.23	8.74	6.22	9.09	27.09	8.87	−1.57	1.30	3.63	3.27	2.78	3.97	68.5	0.000	0.314
3	16.1	0.0	−1.19	9.78	4.92	8.16	26.29	8.33	−1.60	1.67	3.19	3.04	2.84	4.15	68.9	0.000	0.321
4	14.3	0.0	−0.82	7.84	3.68	8.69	23.98	8.44	−1.13	1.37	2.55	3.44	2.70	4.32	75.0	0.000	0.328
5	14.0	0.0	−1.03	7.05	3.12	9.01	24.08	7.90	−1.48	1.31	2.37	3.52	2.91	4.27	67.2	0.000	0.311
6	13.0	0.0	−0.91	4.47	2.65	9.26	23.01	7.63	−1.38	0.86	2.06	3.70	2.94	4.38	63.0	0.000	0.301
7	12.3	0.0	−1.02	3.73	1.63	8.78	25.08	7.93	−1.55	0.74	1.35	3.24	3.10	4.63	57.7	0.000	0.292
8	10.7	0.1	−0.67	−0.21	1.54	9.12	20.11	7.24	−1.06	−0.04	1.28	3.60	2.60	4.31	51.1	0.000	0.272
9	10.1	0.1	−0.75	−0.12	0.93	9.01	18.61	6.55	−1.26	−0.02	0.85	3.22	2.52	4.16	43.0	0.000	0.255
10	9.3	0.1	−0.53	−1.74	0.31	7.83	15.65	6.08	−1.09	−0.32	0.32	2.56	2.38	4.52	35.2	0.000	0.242

Panel B: Joint test using all portfolios

W_T	p-value	$\hat{d}_T$
149.7	0.000	0.484

The regression model is

$$r_{it} = b_{i0} + b_{i1}xrew_{t-1} + b_{i2}jan_t + b_{i3}term_{t-1} + b_{i4}prem_{t-1} + b_{i5}xdiv_{t-1} + e_{it}.$$

Panel A reports the sample R^2 for the unrestricted regression ($\hat{R}_u^2$), the sample R^2 for the restricted regression ($\hat{R}_r^2$), the difference between the unrestricted and restricted coefficient estimates ($\Delta\hat{b}_j,\ j = 0, 1, \ldots, 5$), the associated t-ratios ($t_{\Delta\hat{b}_j},\ j = 0, 1, \ldots, 5$), the Wald statistic and p-value for a test that the unrestricted and restricted coefficient vectors are equal (W_T and p-value), and the Hansen and Jagannathan distance between the unrestricted and restricted coefficient vectors ($\hat{d}_T$). The values of $\Delta\hat{b}_0$, $\Delta\hat{b}_1$, and $\Delta\hat{b}_2$ are scaled by 100. Panel B reports the Wald statistic, p-value, and Hansen and Jagannathan distance for a joint specification that combines the 10 regressions. Individual portfolio tests have six degrees of freedom. The joint test has 60 degrees of freedom. Standard errors are corrected for serial correlation and conditional heteroscedasticity using the procedures of Andrews (1991). See Table 1 for a description of the data.

nal returns does nothing to alter the basic findings on predictability. Column two of the table reports the R^2 values for the restricted regressions. Since all of these values are below 1%, the initial indications are that performance of the log utility model is no better than that of the specifications that use consumption data.

The remaining columns of panel A are used to assess the statistical significance of the difference in the restricted and unrestricted values of the regression R^2. Columns 3 through 8 report the difference between the unrestricted and restricted estimates of the regression coefficients; columns 9 through 14 report the corresponding t-ratios. The last three columns report the Wald statistic associated with the hypothesis that the unrestricted and restricted coefficient vectors are equal, the probability value for the test, and an estimate of the Hansen and Jagannathan distance between the unrestricted and restricted coefficient vectors. The t-ratios and Wald tests are corrected for serial correlation and conditional heteroscedasticity using the approach of Andrews (1991). Under the null, the Wald statistic converges to a chi-square random variable with six degrees of freedom.

The t-ratios and Wald tests indicate that, from a statistical standpoint, the difference between the unrestricted and restricted coefficients is highly significant. Most of the t-ratios are greater than two in magnitude, and the Wald tests strongly reject the null hypothesis that the model explains the return predictability. The results of the joint Wald test reported in panel B are similar. Even when we consider all 10 of the size portfolios together, the p-value for the test is essentially equal to zero. Given these results, it does not appear that we can achieve a meaningful improvement in the fit of the consumption-based models simply by eliminating the use of consumption data.

The Hansen and Jagannathan distance estimates support this conclusion. They range from a low of 0.242 for decile 10 to a high of 0.328 for decile 4. Although not shown, the distance measure estimates for the specifications that use consumption data are similar.[11] In addition, there is a definite cross-sectional pattern in the distance measure estimates. This is evidence that the pricing errors generated by the log utility model are related to the market capitalization of the portfolio. A similar pattern is evident for the unrestricted R^2. It is about twice as big for decile 1 as it is for decile 10. The restricted R^2, on the other hand, is never far from zero. Regardless of how we evaluate the performance of the model, we draw the same conclusion: returns are too predictable to be compatible with log utility.

[11] With a relative risk aversion of five, for example, the corresponding values for the power utility model are 0.251 and 0.337.

3.3 Results for the constant price of risk specifications

The constant price of risk specification of the conditional CAPM is considered next. The results for this specification are summarized in Table 3. As before, the tests in panel A treat each of the portfolios separately. In general, the results look similar to those for the log utility model. The R^2 values for the restricted regressions are in the range of 0.2% to 0.3%, a majority of the t-ratios are greater than two in magnitude, and both the individual portfolio and joint Wald statistics strongly reject the hypothesis that the constant price of risk specification explains the predictability documented in the regressions. Moreover, the estimates of the Hansen and Jagannathan distance range from a low of 0.239 for decile 10 to a high of 0.320 for decile 4, which is nearly identical to the range of values in Table 2.

The cross-sectional pattern in these distance measure estimates is again clearly evident. This pattern may be an important clue to the source of return predictability. To see why, consider a scenario where predictability is due to market inefficiencies. In this case, there is no reason to expect cross-sectional differences in the pricing errors to be a function of cross-sectional differences in risk. If, however, return predictability is rational and we are simply using the wrong pricing model, then cross-sectional differences in the pricing errors are likely to be related to cross-sectional differences in risk.

We can investigate whether this is a plausible explanation for the observed pattern in the distance measure estimates by fitting the Fama and French three-factor model. In an economy where pricing is rational, the market capitalization of a firm partially reflects its risk, so the size and book-to-market portfolios should have some power to capture cross-sectional differences in risk. Since the constant price of risk specification of the conditional three-factor model nests the constant price of risk specification of the conditional CAPM, any reduction in the distance measure estimates obtained by fitting the three-factor model can be directly attributed to the size and book-to-market factors. Therefore the three-factor model provides a way to assess whether the pricing errors for the conditional CAPM exhibit the cross-sectional pattern that we would expect in an economy where the return predictability is rational.

The results for the constant price of risk specification of the three-factor model are shown in Table 4. Adding the size and book-to-market factors to the analysis produces a noticeable improvement in fit over the corresponding specification of the conditional CAPM. The restricted regression R^2 ranges from 0.6% for decile 10 to 4.5% for decile 1, and the maximum Hansen and Jagannathan distance falls to 0.294. But the degree of predictability indicated by the restricted estimates is still too low. Both the individual portfolio and joint tests

Table 3
Tests of the restrictions implied by the constant price of risk version of the conditional CAPM

Panel A: Individual portfolio tests

	R^2 (%)		Unrestricted minus restricted coefficient						t-ratio of difference						Wald test and HJ distance		
Decile	$\hat{R}_u^2$	$\hat{R}_r^2$	$\Delta\hat{b}_0$	$\Delta\hat{b}_1$	$\Delta\hat{b}_2$	$\Delta\hat{b}_3$	$\Delta\hat{b}_4$	$\Delta\hat{b}_5$	$t_{\Delta\hat{b}_0}$	$t_{\Delta\hat{b}_1}$	$t_{\Delta\hat{b}_2}$	$t_{\Delta\hat{b}_3}$	$t_{\Delta\hat{b}_4}$	$t_{\Delta\hat{b}_5}$	W_T	p-value	$\hat{d}_T$
1	20.3	0.2	−1.28	16.10	7.94	8.88	23.24	8.60	−1.58	1.83	4.14	3.02	2.16	3.67	64.1	0.000	0.318
2	18.1	0.2	−1.31	11.02	5.82	9.14	24.89	8.67	−1.77	1.76	3.70	3.30	2.61	3.87	63.5	0.000	0.310
3	16.1	0.2	−1.28	11.96	4.57	8.19	24.33	8.17	−1.82	2.17	3.22	3.04	2.68	4.09	62.2	0.000	0.316
4	14.3	0.2	−0.91	9.95	3.35	8.66	22.21	8.33	−1.34	1.88	2.48	3.41	2.54	4.30	63.2	0.000	0.320
5	14.0	0.2	−1.13	9.18	2.83	8.93	22.43	7.79	−1.72	1.82	2.29	3.45	2.74	4.21	55.7	0.000	0.304
6	13.0	0.2	−1.02	6.49	2.37	9.19	21.42	7.52	−1.64	1.33	1.95	3.64	2.76	4.30	51.7	0.000	0.291
7	12.3	0.2	−1.12	5.86	1.38	8.62	23.52	7.85	−1.80	1.27	1.21	3.21	2.93	4.58	46.9	0.000	0.284
8	10.7	0.2	−0.77	1.89	1.29	8.90	18.65	7.16	−1.29	0.42	1.13	3.56	2.42	4.27	40.7	0.000	0.262
9	10.1	0.3	−0.85	1.97	0.72	8.68	17.37	6.50	−1.53	0.43	0.69	3.16	2.37	4.15	34.0	0.000	0.248
10	9.3	0.3	−0.61	0.07	0.14	7.52	14.49	6.06	−1.33	0.01	0.15	2.54	2.24	4.46	31.0	0.000	0.239

Panel B: Joint test using all portfolios

W_T	p-value	$\hat{d}_T$
173.1	0.000	0.481

The regression model is

$$r_{it} = b_{i0} + b_{i1}xrew_{t-1} + b_{i2}jan_t + b_{i3}term_{t-1} + b_{i4}prem_{t-1} + b_{i5}xdiv_{t-1} + e_{it}.$$

Panel A reports the sample R^2 for the unrestricted regression ($\hat{R}_u^2$), the sample R^2 for the restricted regression ($\hat{R}_r^2$), the difference between the unrestricted and restricted coefficient estimates ($\Delta\hat{b}_j,\ j = 0, 1, \ldots, 5$), the associated t-ratios ($t_{\Delta\hat{b}_j},\ j = 0, 1, \ldots, 5$), the Wald statistic and p-value for a test that the unrestricted and restricted coefficient vectors are equal (W_T and p-value), and the Hansen and Jagannathan distance between the unrestricted and restricted coefficient vectors ($\hat{d}_T$). The values of $\Delta\hat{b}_0$, $\Delta\hat{b}_1$, and $\Delta\hat{b}_2$ are scaled by 100. Panel B reports the Wald statistic, p-value, and Hansen and Jagannathan distance for a joint specification that combines the 10 regressions. Individual portfolio tests have six degrees of freedom. The joint test has 60 degrees of freedom. Standard errors are corrected for serial correlation and conditional heteroscedasticity using the procedures of Andrews (1991). See Table 1 for a description of the data.

strongly reject the model. This indicates that the size and book-to-market factors capture only a modest portion of the predictable variation in returns when we force the prices of risk to be constant.

It might seem surprising that the three-factor model does not perform better in light of the findings of Fama and French (1993). However, Fama and French do not test the stochastic discount factor representation of their model, so direct comparisons could be misleading. Instead, they focus on the statistical significance of the coefficient estimates in multiple regressions of excess returns on the factors. Their methodology has more in common with the tests of the constant conditional beta specifications. The results of these tests are discussed in Section 3.4.

3.4 Results for the constant conditional beta specifications

The results for the constant conditional beta specification of the conditional CAPM are reported in Table 5. This version of the model performs substantially better than the constant price of risk specification. The restricted regression R^2 ranges from a low of 6.2% for decile 1 to a high of 10.3% for deciles 9 and 10. Although most of the individual portfolio tests still reject the model, a majority of the t-ratios are less than two in magnitude and the statistic for the joint test is insignificant at the 5% level. So, the model has at least some degree of success in explaining the predictable variation in returns.

It is important to place these results in proper context, however. Recall that, under the constant conditional beta version of the conditional CAPM, the predictable variation in the portfolio returns is driven by predictable variation in the market return. Consequently tests of this specification treat the predictability of the market return as given. This must be considered when interpreting the test results. Suppose, for example, we are in a situation where the market return is predictable solely as a result of inefficiencies. As long as the returns on individual firms are contemporaneously correlated, portfolio returns will almost certainly exhibit some degree of predictability, regardless of whether there is a structural link between the beta of the portfolio and the predictability of its return. Depending on the values of the portfolio betas, the tests could have low power to reject the restrictions of the conditional CAPM.

Since it is necessary to depart from the stochastic discount factor representation to impose the constant conditional beta assumption, we cannot use the Hansen and Jagannathan distance measure to examine how the pricing errors for this specification behave in the cross section. But we can get some idea of the cross-sectional pattern of the pricing errors by comparing the unrestricted and restricted values of the regression R^2. The unrestricted R^2 starts at about 9% for decile

Table 4
Tests of the restrictions implied by the constant price of risk version of the conditional three-factor model

Panel A: Individual portfolio tests

	R^2 (%)		Unrestricted minus restricted coefficient						t-ratio of difference						Wald test and HJ distance		
Decile	$\hat{R}_u^2$	$\hat{R}_r^2$	$\Delta\hat{b}_0$	$\Delta\hat{b}_1$	$\Delta\hat{b}_2$	$\Delta\hat{b}_3$	$\Delta\hat{b}_4$	$\Delta\hat{b}_5$	$t_{\Delta\hat{b}_0}$	$t_{\Delta\hat{b}_1}$	$t_{\Delta\hat{b}_2}$	$t_{\Delta\hat{b}_3}$	$t_{\Delta\hat{b}_4}$	$t_{\Delta\hat{b}_5}$	W_T	p-value	$\hat{d}_T$
1	20.3	4.5	−1.19	19.11	3.57	10.25	13.26	6.34	−1.49	2.73	1.99	3.01	1.08	2.63	40.3	0.000	0.282
2	18.1	3.5	−1.37	13.63	2.39	10.18	18.71	6.90	−1.86	2.31	1.62	3.13	1.70	3.08	46.6	0.000	0.277
3	16.1	2.9	−1.41	14.51	1.72	9.24	19.68	6.57	−2.06	2.66	1.30	2.92	1.91	3.22	52.8	0.000	0.289
4	14.3	2.3	−1.05	11.66	0.98	9.47	18.98	6.98	−1.58	2.34	0.78	3.08	1.95	3.70	51.5	0.000	0.294
5	14.0	1.9	−1.27	10.60	0.81	9.97	19.70	6.66	−1.95	2.11	0.69	3.23	2.14	3.54	46.2	0.000	0.283
6	13.0	1.7	−1.17	8.06	0.56	10.17	19.15	6.47	−1.92	1.59	0.49	3.31	2.17	3.59	44.6	0.000	0.279
7	12.3	1.2	−1.24	7.09	0.01	9.67	21.42	6.81	−1.99	1.44	0.00	2.85	2.39	3.99	45.5	0.000	0.275
8	10.7	1.3	−0.89	3.02	−0.05	9.97	16.23	5.99	−1.49	0.65	−0.05	3.15	1.89	3.58	35.3	0.000	0.251
9	10.1	0.9	−0.97	2.40	−0.27	10.02	15.36	5.43	−1.72	0.50	−0.27	2.82	1.85	3.53	32.8	0.000	0.243
10	9.3	0.6	−0.70	0.07	−0.35	8.55	13.45	4.96	−1.42	0.01	−0.39	2.16	1.73	3.81	27.7	0.000	0.228

Panel B: Joint test using all portfolios

W_T	p-value	$\hat{d}_T$
224.5	0.000	0.474

The regression model is

$$r_{it} = b_{i0} + b_{i1} xrew_{t-1} + b_{i2} jan_t + b_{i3} term_{t-1} + b_{i4} prem_{t-1} + b_{i5} xdiv_{t-1} + e_{it}.$$

Panel A reports the sample R^2 for the unrestricted regression ($\hat{R}_u^2$), the sample R^2 for the restricted regression ($\hat{R}_r^2$), the difference between the unrestricted and restricted coefficient estimates ($\Delta\hat{b}_j,\ j = 0,\ 1,\ \ldots,\ 5$), the associated t-ratios ($t_{\Delta\hat{b}_j},\ j = 0,\ 1,\ \ldots,\ 5$), the Wald statistic and p-value for a test that the unrestricted and restricted coefficient vectors are equal (W_T and p-value), and the Hansen and Jagannathan distance between the unrestricted and restricted coefficient vectors ($\hat{d}_T$). The values of $\Delta\hat{b}_0$, $\Delta\hat{b}_1$, and $\Delta\hat{b}_2$ are scaled by 100. Panel B reports the Wald statistic, p-value, and Hansen and Jagannathan distance for a joint specification that combines the 10 regressions. Individual portfolio tests have six degrees of freedom. The joint test has 60 degrees of freedom. Standard errors are corrected for serial correlation and conditional heteroscedasticity using the procedures of Andrews (1991). See Table 1 for a description of the data.

Table 5
Tests of the restrictions implied by the constant conditional beta version of the conditional CAPM

Panel A: Individual portfolio tests

	R^2 (%)		Unrestricted minus restricted coefficient						t-ratio of difference						Wald test	
Decile	$\hat{R}_u^2$	$\hat{R}_r^2$	$\Delta\hat{b}_0$	$\Delta\hat{b}_1$	$\Delta\hat{b}_2$	$\Delta\hat{b}_3$	$\Delta\hat{b}_4$	$\Delta\hat{b}_5$	$t_{\Delta\hat{b}_0}$	$t_{\Delta\hat{b}_1}$	$t_{\Delta\hat{b}_2}$	$t_{\Delta\hat{b}_3}$	$t_{\Delta\hat{b}_4}$	$t_{\Delta\hat{b}_5}$	W_T	p-value
1	20.3	6.2	−0.42	13.03	7.52	−1.91	4.65	0.94	−0.80	1.57	4.51	−0.85	0.66	0.64	42.9	0.000
2	18.1	7.6	−0.47	8.39	5.23	−1.50	6.48	1.01	−1.11	1.55	4.35	−0.78	1.15	0.81	35.3	0.000
3	16.1	8.0	−0.45	9.48	3.93	−2.02	6.34	0.74	−1.21	2.02	3.89	−1.15	1.36	0.68	30.9	0.000
4	14.3	8.7	−0.08	7.56	2.69	−1.40	3.93	0.83	−0.26	1.89	3.34	−0.90	0.97	0.86	24.8	0.000
5	14.0	9.0	−0.31	6.72	2.13	−0.77	4.53	0.51	−1.04	1.89	3.15	−0.52	1.21	0.56	22.1	0.001
6	13.0	9.4	−0.19	4.20	1.68	−0.44	3.60	0.31	−0.72	1.39	2.84	−0.32	1.09	0.37	20.4	0.002
7	12.3	9.8	−0.31	3.37	0.64	−0.78	5.78	0.64	−1.38	1.49	1.46	−0.84	2.23	0.95	16.3	0.012
8	10.7	10.0	0.02	−0.58	0.58	−0.16	1.21	0.11	0.11	−0.29	1.46	−0.17	0.53	0.17	5.4	0.499
9	10.1	10.3	−0.07	−0.54	−0.03	0.15	0.24	−0.36	−0.55	−0.41	−0.12	0.26	0.14	−0.82	1.8	0.938
10	9.3	10.3	0.06	−2.10	−0.53	0.02	−0.50	−0.01	0.49	−1.65	−2.47	0.03	−0.39	−0.02	14.8	0.022

Panel B: Joint test using all portfolios

W_T	p-value
76.0	0.080

The regression model is

$$r_{it} = b_{i0} + b_{i1} xrew_{t-1} + b_{i2} jan_t + b_{i3} term_{t-1} + b_{i4} prem_{t-1} + b_{i5} xdiv_{t-1} + e_{it}.$$

Panel A reports the sample R^2 for the unrestricted regression ($\hat{R}_u^2$), the sample R^2 for the restricted regression ($\hat{R}_r^2$), the difference between the unrestricted and restricted coefficient estimates ($\Delta\hat{b}_j,\ j = 0, 1, \ldots, 5$), the associated t-ratios ($t_{\Delta\hat{b}_j},\ j = 0, 1, \ldots, 5$), and the Wald statistic and p-value for a test that the unrestricted and restricted coefficient vectors are equal (W_T and p-value). The values of $\Delta\hat{b}_0$, $\Delta\hat{b}_1$, and $\Delta\hat{b}_2$ are scaled by 100. Panel B reports the Wald statistic and p-value for a joint specification that combines the 10 regressions. Individual portfolio tests have six degrees of freedom. The joint test has 60 degrees of freedom. Standard errors are corrected for serial correlation and conditional heteroscedasticity using the procedures of Andrews (1991).

10 and rises to around 20% for decile 1, while the restricted R^2 starts at about 10% for decile 10 and falls to around 6% for decile 1. This indicates that there is still a relation between the pricing errors generated by the model and the market capitalization of the portfolio. As before, we can explore whether this relation appears to be consistent with rational pricing by fitting the Fama and French model.

The results for the constant conditional beta specification of the Fama and French model are shown in Table 6. Most of the individual portfolio tests reject the model, but it clearly fits the data better than the constant beta specification of the conditional CAPM. In particular, the restricted R^2 increases from around 9% for decile 10 to over 15% for decile 1. Moreover, the relation between market capitalization and the magnitude of the pricing errors largely disappears, and the joint test statistic is insignificant at standard levels.

Again, however, we have to be careful not to read too much into these results. Market capitalization and book-to-market ratios cannot be treated as factors in the conventional sense. Factors represent fundamental sources of undiversifiable risk in the economy. The returns on the Fama and French portfolios simply pick up the differences in discount rates that are embedded in market prices. As long as assets are priced rationally, the market capitalization and book-to-market factors should have some power to capture cross-sectional differences in risk, regardless of whether the true model is the CAPM, the APT, or a consumption-based specification. This makes the Fama and French model better suited to a diagnostic role than to traditional uses like estimating expected returns. At best, we can say the results in Table 6 appear to be reasonably consistent with what we would expect to see in an economy where predictability is rational.

3.5 Comparisons with prior research

Hansen and Singleton (1983) conducted the first comprehensive study of rational pricing and return predictability. They developed explicit restrictions on the joint stochastic process that governs consumption growth and asset returns by assuming constant relative risk aversion and lognormality. In their representative agent framework, the predictable component of returns is proportional to the predictable component of consumption growth. Hansen and Singleton tested this restriction and found that it was strongly rejected by the data. Although the tests of the consumption-based models in this study eliminate the strong assumptions of Hansen and Singleton, they yield similar results. The consumption-based models cannot generate the degree of predictability that is observed in the data.

In more recent studies, Ferson and Harvey (1991) and Ferson and Korajczyk (1995) argue that conditional beta pricing models explain

Table 6
Tests of the restrictions implied by the constant conditional beta version of the conditional three-factor model

Panel A: Individual portfolio tests

	R^2 (%)		Unrestricted minus restricted coefficient						t-ratio of difference						Wald test	
Decile	$\hat{R}^2_u$	$\hat{R}^2_r$	$\Delta\hat{b}_0$	$\Delta\hat{b}_1$	$\Delta\hat{b}_2$	$\Delta\hat{b}_3$	$\Delta\hat{b}_4$	$\Delta\hat{b}_5$	$t_{\Delta\hat{b}_0}$	$t_{\Delta\hat{b}_1}$	$t_{\Delta\hat{b}_2}$	$t_{\Delta\hat{b}_3}$	$t_{\Delta\hat{b}_4}$	$t_{\Delta\hat{b}_5}$	W_T	p-value
1	20.3	15.6	−0.57	−5.21	2.54	−1.02	4.28	0.66	−1.83	−1.14	4.18	−1.00	1.06	0.81	22.8	0.001
2	18.1	15.6	−0.58	−6.35	1.46	−0.72	6.20	0.77	−2.74	−3.16	4.07	−0.98	2.20	1.35	25.5	0.000
3	16.1	15.4	−0.54	−4.13	0.69	−1.25	6.10	0.49	−3.48	−2.62	2.24	−2.10	3.07	1.10	19.1	0.004
4	14.3	14.7	−0.15	−3.88	0.15	−0.71	3.74	0.60	−1.01	−2.95	0.65	−1.02	2.04	1.37	11.6	0.073
5	14.0	14.2	−0.37	−3.30	−0.07	−0.16	4.36	0.31	−2.66	−3.01	−0.29	−0.24	2.43	0.79	19.2	0.004
6	13.0	13.8	−0.24	−4.10	−0.25	0.04	3.45	0.15	−1.76	−3.53	−1.09	0.07	2.01	0.36	17.6	0.007
7	12.3	12.6	−0.34	−2.78	−0.62	−0.38	5.68	0.50	−2.46	−2.31	−2.67	−0.75	3.48	1.24	26.1	0.000
8	10.7	12.2	−0.01	−4.64	−0.53	0.04	1.13	0.05	−0.09	−4.16	−2.32	0.06	0.73	0.11	24.6	0.000
9	10.1	11.1	−0.09	−1.74	−0.53	0.17	0.20	−0.36	−0.76	−1.67	−2.49	0.34	0.13	−0.94	11.4	0.077
10	9.3	9.3	0.08	1.48	0.22	−0.21	−0.44	0.07	1.31	3.43	2.44	−0.79	−0.63	0.34	24.0	0.001

Panel B: Joint test using all portfolios

W_T	p-value
62.5	0.39

The regression model is

$$r_{it} = b_{i0} + b_{i1}xrew_{t-1} + b_{i2}jan_t + b_{i3}term_{t-1} + b_{i4}prem_{t-1} + b_{i5}xdiv_{t-1} + e_{it}.$$

Panel A reports the sample R^2 for the unrestricted regression ($\hat{R}^2_u$), the sample R^2 for the restricted regression ($\hat{R}^2_r$), the difference between the unrestricted and restricted coefficient estimates ($\Delta\hat{b}_j,\ j = 0, 1, \ldots, 5$), the associated t-ratios ($t_{\Delta\hat{b}_j},\ j = 0, 1, \ldots, 5$), the Wald statistic and p-value for a test that the unrestricted and restricted coefficient vectors are equal (W_T and p-value). The values of Δb_0, $\Delta\hat{b}_1$, and $\Delta\hat{b}_2$ are scaled by 100. Panel B reports the Wald statistic and p-value for a joint specification that combines the 10 regressions. Individual portfolio tests have six degrees of freedom. The joint test has 60 degrees of freedom. Standard errors are corrected for serial correlation and conditional heteroscedasticity using the procedures of Andrews (1991).

a large fraction of the predictable variation in portfolio returns. The evidence discussed in Section 3.4 is generally consistent with their results. But it may be premature to conclude that these models truly explain return predictability. In conditional beta models, return predictability is driven by predictable variation in both the betas and the factors. So, if one or more of the factors is the return on a portfolio, we know before we conduct the empirical tests that the model will almost certainly imply some level of predictability. This is an unavoidable consequence of the contemporaneous correlation in returns. The real question, therefore, is not whether conditional beta pricing models capture a substantial fraction of the predictable variation in returns, but whether they explain cross-sectional differences in return predictability.

To see why this is an important distinction, consider the potential effects of data mining. If we have a large number of researchers examining essentially the same data, then it is quite likely that some of them will uncover predictability that appears to be statistically significant but is really due to chance occurrences. Ferson and Korajczyk (1995) note that this is a concern and argue that, since they select factors based on research that predates the predictability literature, it should be difficult to explain the predictability that results from data mining using their pricing model. They contend, in other words, that from the standpoint of their work, data-mining biases should be conservative.

But, given the form of the restrictions implied by their pricing model, it appears that this need not be true. We know that (i) under their model, return predictability is closely linked to the predictability of the factors; and (ii) asset returns are contemporaneously correlated. So, if we specify the market return as a factor and it is predictable due to data mining, then we should expect the pricing model to explain at least some of the predictability displayed by other portfolios. Even though the market return was identified as a factor long before researchers began to report that returns are predictable, this is no guarantee that using it as a factor will lead to conservative tests. A pricing model that uses portfolio returns as factors may appear to explain a substantial fraction of the predictable variation in returns, even when this predictability is due to collective data mining.[12]

The form of the restrictions implied by the constant beta specifica-

[12] Foster, Smith, and Whaley (1997) make a similar point when discussing the potential effects of data mining. They note that some researchers argue that data mining is not a plausible explanation for the ability to predict U.S. returns because recent studies have documented predictability using international returns as well. Their analysis reveals, however, that this is not a valid inference. Since international returns and instruments are correlated with the U.S. returns and instruments, mining the U.S. data will also bias the results of the international studies.

tions may also help explain why Ferson and Harvey (1991) and Ferson and Korajczyk (1995) conclude that time-varying betas explain only a small fraction of the predictable variation in returns. The empirical evidence suggests that, even when the betas are held constant, any linear factor pricing model in which the market return is specified as one of the factors will probably have a reasonable degree of success in explaining return predictability. As a result, the additional flexibility afforded by allowing the betas to change through time is unlikely to have a large impact on the analysis.

Despite such caveats, the results for the constant conditional beta specifications remain interesting. They reveal that the observed pattern of predictability across the size portfolios is reasonably consistent with what we would expect under circumstances where predictability is rational. This finding suggests that if we can identify the factors that lead to cross-sectional differences in average returns, then we have a chance of explaining cross-sectional differences in predictability as well.

4. Conclusions

The ability to predict asset returns is often measured by regressing them on a set of predetermined instrumental variables. This approach is favored over other strategies because it is simple to implement and the resulting estimates of the slope coefficients and regression R^2 have an appealing interpretation. The primary contribution of this article is to show that asset pricing models imply that the parameters in such regressions must take on certain values. With the form of these restrictions established, it is straightforward to determine whether the predictability documented using regression analysis is consistent with the implications of rational asset pricing models. The article also shows how to construct the linear forecasting specifications implied by different asset pricing models. Forecasts generated in this manner are useful for exploring the economic basis for predictability.

The empirical investigation focuses on whether the ability to predict size portfolio returns is consistent with a number of well-known pricing models. These models include the consumption-based specifications of Lucas (1978), Abel (1990), and Epstein and Zin (1991), along with simple conditional versions of the Sharpe (1964) and Lintner (1965) CAPM and the Fama and French (1993) three-factor model. The evidence indicates that, in general, returns are too predictable to be consistent with these models. Of course, rejecting the restrictions for these models does not necessarily mean that predictability is irrational. Any test for rational predictability involves a joint hypothesis that markets are efficient and the pricing model used in the test is correct. Although none of models in this study explain the overall pattern of predictable variation in returns, the results do suggest that the cross-sectional differences in predictability are reasonably consistent with what we would expect under circumstances where predictability is rational.

As for future research, there are several ways to refine the analysis. Recent work by Luttmer (1994) suggests that it is important to take transaction costs into account. Thus one useful extension would be to explicitly incorporate market frictions into the restrictions on regression-based measures of predictability. The work of Hansen, Heaton, and Luttmer (1995) should prove useful in this regard. Another possibility is to pursue a seminonparametric approach like that of Gallant and Tauchen (1989). They argue that it is local durability of nondurable consumption, not habit persistence or adjustment costs, that is the source of intertemporal nonseparability in preferences. Perhaps consumption-based specifications that incorporate this feature would have better success in explaining the observed ability to predict returns than the models considered here.

Appendix A: The Variance Bound and Distance Measure

The relation between the restrictions on predictability developed in Section 1, the Hansen and Jagannathan (1991) variance bound, and the Hansen and Jagannathan (1997) distance measure is illustrated below.

A.1 The Variance Bound

Let $\boldsymbol{r}_t$ denote an $N \times 1$ vector of portfolio returns. Hansen and Jagannathan (1991) show that the lower bound on the variance of the set of admissible stochastic discount factors is given by the variance of the linear projection of any element of this set onto the asset payoff space. Since the assets considered in Section 1 include dynamic trading strategies, this projection takes the form

$$\tilde{m}_t^* = \phi_0^* + (\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})' \boldsymbol{\phi}_{rx}^*, \tag{A1}$$

where ϕ_0^* is the intercept and $\boldsymbol{\phi}_{rx}^*$ is a $KN \times 1$ vector of slope coefficients. So the lower bound on the variance of $\tilde{m}_t$ is given by

$$\sigma_{\tilde{m}^*}^2 = \boldsymbol{\phi}_{rx}^{*\prime} \boldsymbol{\Sigma}_{rx} \boldsymbol{\phi}_{rx}^*, \tag{A2}$$

where $\boldsymbol{\Sigma}_{rx} \equiv E[(\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})(\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})']$ is the variance-covariance matrix of the subset of payoffs used to construct $\tilde{m}_t^*$.

Since $\tilde{m}_t$ is an admissible stochastic discount factor, it must assign the correct price to each of the asset payoffs. As a result, we can rearrange Equation (5) to obtain the relation $\mathrm{E}[\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1}] = -\mathrm{cov}(\tilde{m}_t, \boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})$ and write the vector of slopes as

$$\phi^*_{rx} = -\boldsymbol{\Sigma}^{-1}_{rx}\mathrm{E}[\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1}]. \tag{A3}$$

Substituting Equation (A3) into the expression for the variance bound yields

$$\sigma^2_{\tilde{m}^*} = \boldsymbol{\mu}'_{rx}\boldsymbol{\Sigma}^{-1}_{rx}\boldsymbol{\mu}_{rx}, \tag{A4}$$

where $\boldsymbol{\mu}_{rx} \equiv \mathrm{E}[\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1}]$ is the vector of expected payoffs.

To see the role that predictability plays in Equation (A4), consider a multivariate model for forecasting returns of the form

$$\boldsymbol{r}_t = \boldsymbol{X}'_{t-1}\boldsymbol{B} + \boldsymbol{e}_t, \tag{A5}$$

where $\boldsymbol{B} = \mathrm{E}[\boldsymbol{X}_{t-1}\boldsymbol{X}'_{t-1}]^{-1}\mathrm{E}[\boldsymbol{X}_{t-1}\boldsymbol{r}_t]$ is the $KN \times 1$ vector of regression coefficients, $\boldsymbol{X}_{t-1} \equiv \boldsymbol{I}_N \otimes \boldsymbol{x}_{t-1}$, and $\boldsymbol{I}_N$ is an $N \times N$ identity matrix. Following Kirby (1997), we can show that the limiting distribution of the least-squares estimator of $\boldsymbol{B}$ is

$$\sqrt{T}(\hat{\boldsymbol{B}} - \boldsymbol{B}) \stackrel{d}{\to} N(\boldsymbol{0}, \boldsymbol{G}^{-1}\boldsymbol{H}\boldsymbol{G}^{-1}), \tag{A6}$$

where

$$\begin{aligned} \boldsymbol{G} &\equiv \boldsymbol{I}_N \otimes \mathrm{E}[\boldsymbol{x}_{t-1}\boldsymbol{x}'_{t-1}] \\ \boldsymbol{H} &\equiv \sum_{j=-\infty}^{\infty} \mathrm{E}[(\boldsymbol{e}_t \otimes \boldsymbol{x}_{t-1})(\boldsymbol{e}_{t-j} \otimes \boldsymbol{x}_{t-j-1})']. \end{aligned} \tag{A7}$$

Therefore the Wald statistic

$$W_T = T\hat{\boldsymbol{B}}'(\boldsymbol{G}\boldsymbol{H}^{-1}\boldsymbol{G})\hat{\boldsymbol{B}}, \tag{A8}$$

provides a joint test of whether the regression coefficients in Equation (A5) are significantly different from zero.

The Wald test in Equation (A9) can be viewed as a test of the hypothesis that investors are risk neutral. This follows by noting that $\boldsymbol{B} = \boldsymbol{0}$ corresponds to a scenario where excess returns have a mean of zero and are uncorrelated with the instruments. In general, this will be the case only if investors are risk neutral. Under risk-neutral pricing, $\mathrm{E}[(\boldsymbol{e}_t \otimes \boldsymbol{x}_{t-1})(\boldsymbol{e}_{t-j} \otimes \boldsymbol{x}_{t-j-1})']$ reduces to $\mathrm{E}[(\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})(\boldsymbol{r}_{t-j} \otimes \boldsymbol{x}_{t-j-1})']$ and $\mathrm{E}[(\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})(\boldsymbol{r}_{t-j} \otimes \boldsymbol{x}_{t-j-1})'] = 0$ for all $j \neq 0$. Therefore the matrix $\boldsymbol{H}$ reduces to $\boldsymbol{\Sigma}_{rx}$ and the Wald statistic can be computed as

$$W_T = T\hat{\boldsymbol{\mu}}'_{rx}\hat{\boldsymbol{\Sigma}}^{-1}_{rx}\hat{\boldsymbol{\mu}}_{rx}, \tag{A9}$$

where $\hat{\boldsymbol{\mu}}_{rx}$ and $\hat{\boldsymbol{\Sigma}}_{rx}$ are the sample analogues of $\boldsymbol{\mu}_{rx}$ and $\boldsymbol{\Sigma}_{rx}$.

Comparing Equation (A10) to Equation (A4) reveals that $W_T = T\hat{\sigma}^2_{\tilde{m}^*}$. The intuition behind this relation is simple. A large value of $\hat{\sigma}^2_{\tilde{m}^*}$ suggests that the stochastic discount factor is highly variable, which in turn provides evidence against risk neutrality. Thus the variance bound has a natural interpretation within a linear regression framework for predicting returns. Strong predictability will generally translate into a large lower bound on the variance of $\tilde{m}_t$. So if we find that a specification of $\tilde{m}_t$ satisfies the bound, we can say that the candidate $\tilde{m}_t$ displays enough variability to be consistent with the observed ability to predict returns.

A.2 The Distance Measure

Hansen and Jagannathan (1997) use the least-squares distance between a candidate stochastic discount factor and the minimum-variance stochastic discount factor as a measure of model misspecification. To construct this distance measure, we note that the intercept in Equation (A1) is given by $(1 - \boldsymbol{\mu}'_{rx}\boldsymbol{\phi}^*_{rx})$, so we can write $\tilde{m}^*_t$ as

$$\tilde{m}^*_t = 1 + ((\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1}) - \boldsymbol{\mu}_{rx})'\boldsymbol{\phi}^*_{rx}. \tag{A10}$$

Similarly, the linear projection of the candidate stochastic discount factor onto a constant and the set of payoffs can be written as

$$\tilde{m}^c_t = 1 + ((\boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1}) - \boldsymbol{\mu}_{rx})'\boldsymbol{\phi}^c_{rx}. \tag{A11}$$

The distance between these two projections is

$$d \equiv [(\boldsymbol{\phi}^c_{rx} - \boldsymbol{\phi}^*_{rx})'\boldsymbol{\Sigma}_{rx}(\boldsymbol{\phi}^c_{rx} - \boldsymbol{\phi}^*_{rx})]^{\frac{1}{2}}. \tag{A12}$$

Using the relations shown in Equations (13), (A3), and (A6), Equation (A13) can be written as

$$d = [(\boldsymbol{B}_r - \boldsymbol{B}_u)'\boldsymbol{G}\boldsymbol{\Sigma}^{-1}_{rx}\boldsymbol{G}(\boldsymbol{B}_r - \boldsymbol{B}_u)]^{\frac{1}{2}}, \tag{A13}$$

where $\boldsymbol{B}_u = \boldsymbol{G}^{-1}\boldsymbol{\mu}_{rx}$ and $\boldsymbol{B}_r = -\boldsymbol{G}^{-1}\mathrm{cov}(\tilde{m}^c_t, \boldsymbol{r}_t \otimes \boldsymbol{x}_{t-1})$ are the unrestricted and restricted vectors of regression coefficients. If we square Equation (A14) and multiply by T, the resulting expression is similar to the Wald statistic for testing whether the restrictions on the regression coefficients are satisfied. The only difference is that Equation (A14) uses $\boldsymbol{G}\boldsymbol{\Sigma}^{-1}_{rx}\boldsymbol{G}$ as a distance metric instead of the inverse of the variance-covariance matrix of the difference in the unrestricted and restricted estimators.

Appendix B: The Dataset

The 10 size portfolios are constructed using data provided by the Center for Research in Security Prices (CRSP) at the University of Chicago. First, the NYSE firms are sorted into deciles based on the market value of equity outstanding. Then, monthly portfolio returns are computed by taking a value-weighted average of the monthly returns for the individual firms. The portfolios are rebalanced annually. For tests that require real returns, the nominal returns are deflated by one plus the inflation rate implied by monthly consumption expenditures on nondurable goods. The consumption data are from CITIBASE. Aggregate per-capita consumption is constructed by taking seasonally adjusted, aggregate, real expenditures on nondurable goods and dividing by the total resident U.S. population.

The size and book-to-market factor are constructed as in Fama and French (1993). Each year, firms are assigned to one of six different portfolios depending on their market capitalization and book equity to market equity ratio. These six portfolios are then combined to form two zero-investment portfolios whose returns are designed to capture cross-sectional differences in systematic risk associated with market capitalization and book-to-market equity. The data used to construct the instrumental variables are drawn from three sources: CRSP (the value-weighted NYSE index, 30-day Treasury bill, and 90-day Treasury bill returns); Standard & Poor's *Current Statistics* (the dividend yield on the S&P 500 stock index); and Moody's *Industrial Manual* (the Aaa and Baa bond yields).

Appendix C: The Variance-Covariance Matrix

To construct tests of the restrictions on predictability, we need a consistent estimator for $\boldsymbol{\Omega}_i \equiv (\boldsymbol{D}_i' \boldsymbol{S}_i^{-1} \boldsymbol{D}_i)^{-1}$. In most applications, the elements of the disturbance vector $\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i)$ are serially correlated. Consequently, a suitable weighting scheme is required in order to obtain a consistent, positive semidefinite estimator of $\boldsymbol{S}_i$. All of the results reported in the empirical section of the article are based on estimates of $\boldsymbol{S}_i$ constructed using quadratic spectral weights. The bandwidth used to compute these weights is an estimate of the optimal bandwidth for the dataset employed in the empirical analysis. As suggested by Andrews (1991), this estimate is obtained by fitting an AR(1) model to each element of the sample analogue of $\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i)$.

The complete procedure for constructing the covariance matrix estimator is as follows. Let P denotes the dimension of the disturbance vector $\boldsymbol{h}(\boldsymbol{y}_{it}, \boldsymbol{\theta}_i)$. First, the sample autocorrelation coefficients, $\hat{\rho}_{ij}$, and the innovation variances, $\hat{\sigma}_{ij}^2$, from the AR(1) models are substituted into the formula,

$$\hat{\eta}_i = \sum_{j=1}^{P} \frac{4\hat{\rho}_{ij}^2 \hat{\sigma}_{ij}^4}{(1-\hat{\rho}_{ij})^8} \Bigg/ \sum_{j=1}^{P} \frac{\hat{\sigma}_{ij}^4}{(1-\hat{\rho}_{ij})^4}. \tag{C1}$$

Then, the bandwidth, ω_i, is set equal to $1.3221(\hat{\eta}_i T)^{0.2}$ and $\hat{\mathbf{S}}_i$ is constructed as

$$\begin{aligned} \hat{\mathbf{S}}_i \; = \; & \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{h}(\boldsymbol{y}_{it}, \hat{\boldsymbol{\theta}}_i) \boldsymbol{h}'(\boldsymbol{y}_{it}, \hat{\boldsymbol{\theta}}_i) \\ & + \sum_{j=1}^{T-1} w_{iL} \Bigg[\frac{1}{T} \sum_{t=l+1}^{T} [\boldsymbol{h}(\boldsymbol{y}_{it}, \hat{\boldsymbol{\theta}}_i) \boldsymbol{h}'(\boldsymbol{y}_{i,t-j}, \hat{\boldsymbol{\theta}}_i) \\ & \qquad + \boldsymbol{h}(\boldsymbol{y}_{i,t-j}, \hat{\boldsymbol{\theta}}_i) \boldsymbol{h}'(\boldsymbol{y}_{it}, \hat{\boldsymbol{\theta}}_i)] \Bigg], \end{aligned} \tag{C2}$$

where

$$w_{iL} \equiv \frac{25}{12\pi^2 (L/\omega_i)^2} \left(\frac{\sin(6\pi L/(5\omega_i))}{6\pi L/(5\omega_i)} - \cos(6\pi L/(5\omega_i)) \right) \tag{C3}$$

is the weight at lag L given by the quadratic spectral kernel. Although the bandwidth used to compute the weights will generally deviate from the true optimal bandwidth because of estimation error, the Monte Carlo results of Andrews (1991) suggest that the resulting estimator of the variance-covariance matrix performs well under most circumstances.

References

Abel, A., 1990, "Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review*, 80, 38–42.

Andrews, D., 1991, "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817–858.

Berk, J., 1995, "A Critique of Size-Related Anomalies," *Review of Financial Studies*, 8, 275–286.

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

Campbell, J., and R. Shiller, 1988, "Stock Prices, Earnings and Expected Dividends," *Journal of Finance*, 46, 661–676.

Cochrane, J., and L. Hansen, 1992, "Asset Pricing Explorations for Macroeconomics," *NBER Macroeconomics Annual*, 115–169.

Efron, B., 1979, "Bootstrap Methods: Another Look at the Jacknife," *Annals of Statistics*, 7, 1–26.

Epstein, L., and S. Zin, 1991, "Substitution, Risk Aversion, and Temporal Behavior of Consumption and Asset Returns II: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.

Fama, E., 1991, "Efficient Capital Markets: II," *Journal of Finance*, 46, 1575–1617.

Fama, E., and K. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

Fama, E., and K. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

Ferson, W., and C. Harvey, 1991, "The Variation in Economic Risk Premiums," *Journal of Political Economy*, 99, 385–415.

Ferson, W., and R. Korajczyk, 1995, "Do Arbitrage Pricing Models Explain the Predictability of Stock Returns?," *Journal of Business*, 68, 309–349.

Foster, D., T. Smith, and R. Whaley, 1997, "Assessing Goodness-of-Fit of Asset Pricing Models: The Distribution of the Maximal R^2," *Journal of Finance*, 52, 591–607.

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

Hansen, L., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

Hansen, L., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.

Hansen, L., and R. Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.

Hansen, L., J. Heaton, and E. Luttmer, 1995, "Econometric Evaluation of Asset Pricing Models," *Review of Financial Studies*, 8, 237–274.

Hansen, L., and S. Richard, 1987, "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55, 1029–1054.

Hansen, L., and K. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1285.

Hansen, L., and K. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91, 249–264.

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

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

Hawawini, G., and D. Keim, 1992, "On the Predictability of Common Stock Returns: World-Wide Evidence," working paper, Wharton School, University of Pennsylvania.

Keim, D., and R. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," *Journal of Financial Economics*, 17, 357–390.

Kirby, C., 1997, "Measuring the Predictable Variation in Stock and Bond Returns," *Review of Financial Studies*, 10, 579–630.

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

Lintner, J., 1965, "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," *Review of Economics and Statistics*, 47, 13–37.

Lucas, R., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Luttmer, E., 1994, "Asset Pricing in Economies with Frictions," working paper, Northwestern University.

Rubinstein, M., 1976, "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics*, 7, 407–425.

Sharpe, W., 1964, "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance*, 19, 425–442.