

NBER WORKING PAPER SERIES

A SKEPTICAL APPRAISAL OF ASSET-PRICING TESTS

Jonathan Lewellen
Stefan Nagel
Jay Shanken

Working Paper 12360
<http://www.nber.org/papers/w12360>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2006

We are grateful to Christopher Polk, Motohiro Yogo, and workshop participants at the universities of Alberta, Arizona State, Calgary, Georgia, Indiana, Oregon, Stanford, and Washington for helpful comments. We thank Ken French, Sydney Ludvigson, Stijn Van Nieuwerburgh, and Motohiro Yogo for providing data via their websites. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Jonathan Lewellen, Stefan Nagel and Jay Shanken. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Skeptical Appraisal of Asset-Pricing Tests
Jonathan Lewellen, Stefan Nagel and Jay Shanken
NBER Working Paper No. 12360
July 2006
JEL No. G12

ABSTRACT

It has become standard practice in the cross-sectional asset-pricing literature to evaluate models based on how well they explain average returns on size- and B/M-sorted portfolios, something many models seem to do remarkably well. In this paper, we review and critique the empirical methods used in the literature. We argue that asset-pricing tests are often highly misleading, in the sense that apparently strong explanatory power (high cross-sectional R2s and small pricing errors) in fact provides quite weak support for a model. We offer a number of suggestions for improving empirical tests and evidence that several proposed models don't work as well as originally advertised.

Jonathan Lewellen
Tuck School of Business at Dartmouth
305 Tuck Hall
Hanover, NH 03755
and NBER
jon.lewellen@dartmouth.edu

Stefan Nagel
Stanford University
Graduate School of Business
518 Memorial Way
Stanford, CA 94305
and NBER
nagel_stefan@gsb.stanford.edu

Jay A. Shanken
Goizueta Business School
Emory University
1300 Clifton Road
Atlanta, GA 30322
and NBER
jay_shanken@bus.emory.edu

1. Introduction

The finance literature has proposed a wide variety of asset-pricing models in recent years, motivated by evidence that small, high-B/M stocks have positive CAPM-adjusted returns. The models – formal equilibrium theories and reduced-form econometric models – suggest new risk factors to help explain expected returns, including labor income (Jagannathan and Wang, 1996; Heaton and Lucas, 2000), growth in real investment, GDP, and future consumption (Cochrane, 1996; Vassalou, 2003; Li, Vassalou, and Xing, 2005; Parker and Julliard, 2005; Hansen, Heaton, and Li, 2005), housing prices (Kullman 2003), innovations in assorted state variables (Campbell and Vuolteenaho 2004; Brennan, Wang, and Xia, 2004; Petkova, 2006), and liquidity risk (Pastor and Stambaugh, 2003; Acharya and Pedersen, 2005). The literature also proposes a host of new conditioning variables to summarize the state of the economy, including the spread between low- and high-grade debt (Jagannathan and Wang, 1996), the aggregate consumption-to-wealth ratio (Lettau and Ludvigson, 2001), the housing collateral ratio (Lustig and Van Nieuwerburgh, 2004), the expenditure share of housing (Piazzesi, Schneider, and Tuzel, 2006), and the labor income to consumption ratio (Santos and Veronesi, 2005).

Empirically, many of the proposed models seem to do a good job explaining the size and B/M effects, an observation at once comforting and disconcerting: comforting because it suggests that rational explanations for the anomalies are readily available, disconcerting because it provides an embarrassment of riches. Reviewing the literature, one gets the uneasy feeling that it seems a bit too easy to explain the size and B/M effects. This is especially true given the great variety of factor models that seem to work, many of which have very little in common with each other.

Our paper is motivated by that suspicion. Specifically, our goal is to explain why, despite the seemingly strong evidence that many proposed models can explain the size and B/M effects, we remain unconvinced by the evidence. We offer a critique of the empirical methods that have become popular in the asset-pricing literature, a number of prescriptions for improving the tests, and evidence that several of the proposed models don't work as well as originally advertised.

The heart of our critique is that the literature has often given itself a low hurdle to meet in claiming success: high cross-sectional R^2 s (or low pricing errors) when average returns on the Fama-French 25 size-B/M portfolios are regressed on their factor loadings. This hurdle is low because size and B/M portfolios are well-known to have a strong factor structure, i.e., Fama and French's (1993) three factors explain more than 90% of the time-series variation in portfolios' returns and more than 75% of the cross-sectional variation in their average returns. Given those features, obtaining a high cross-sectional R^2 is

easy because almost any proposed factor is likely to produce betas that line up with expected returns; essentially all that's required is for a factor to be (weakly) correlated with SMB or HML but not with the tiny, idiosyncratic three-factor residuals of the size-B/M portfolios.

The problem we highlight is not just a sampling issue, i.e., it is not solved by getting standard errors right. In population, if returns have a covariance structure like that of size-B/M portfolios, loadings on a proposed factor will line up with true expected returns so long as the factor correlates only with the common sources of variation in returns. The problem is also not solved by using an SDF approach. Under the same conditions that give a high cross-sectional R^2 , the true pricing errors in an SDF specification will be small or zero, a result that follows immediately from the close parallel between the regression and SDF approaches (see, e.g., Cochrane, 2001).

This is not to say that sampling issues aren't important. Indeed, the covariance structure of size-B/M portfolios also means that, even if we do find factors that have no ability to explain the cross section of true expected returns, we are still reasonably likely to estimate a high cross-sectional R^2 in sample. As an illustration, we simulate artificial factors that, while correlated with returns, are constructed to have zero true cross-sectional R^2 s for the size-B/M portfolios. We find that a sample adjusted R^2 might need to be as high as 44% to be statistically significant in models with one factor, 62% in models with three factors, and 69% in models with five factors. Further, with three or five factors, the power of the tests is extremely small: the sampling distribution of the adjusted R^2 is almost the same when the true R^2 is zero and when it is as high as 70% or 80%. In short, the high R^2 s reported in the literature aren't nearly as impressive as they might appear.

The obvious question then is: What can be done? How can we improve asset-pricing tests to make them more convincing? We offer four suggestions. First, since the problems are caused by the strong factor structure of size-B/M portfolios, one simple solution is to expand the set of test assets to include other portfolios, for example, industry or beta-sorted portfolios. Second, since the problems are exacerbated by the fact that empirical tests often ignore theoretical restrictions on the cross-sectional slopes, another simple solution is to take the magnitude of the slopes seriously when theory provides appropriate guidance. For example, zero-beta rates should be close to the riskfree rate, the risk premium on a factor portfolio should be close to its average excess return, and the cross-sectional slopes in conditional models should be determined by the volatility of the conditional risk premium (as we explain later; see also Lewellen and Nagel, 2006). Third, we argue that the problems are likely to be less severe for GLS than for OLS cross-sectional regressions, so another (imperfect) solution is to report the GLS R^2 . An added

benefit is that the GLS R^2 has a useful economic interpretation in terms of the relative mean-variance efficiency of a model's factor-mimicking portfolios (this interpretation builds on and generalizes the results of Kandel and Stambaugh, 1995).

Finally, since the problems are exacerbated by sampling issues, a fourth 'solution' is to report confidence intervals for test statistics, not rely just on point estimates and p-values. We describe how to do so for the cross-sectional R^2 and other, more formal statistics based on the weighted sum of squared pricing errors, including Shanken's (1985) cross-sectional T^2 (or asymptotic χ^2) statistic, Gibbons, Ross, and Shanken's (1989) F-statistic, and Hansen and Jagannathan's (1997) HJ-distance. For the latter three statistics, the confidence intervals again have a natural economic interpretation in terms of the relative mean-variance efficiency of a model's factor-mimicking portfolios.

Our suggestion to report confidence intervals has two main benefits. The first is that confidence intervals can reveal the often high sampling error in the statistics – by showing the wide range of true parameters that are consistent with the data – in a way that is more direct and transparent than p-values or standard errors (since the statistics are generally biased and skewed). The second advantage of confidence intervals over p-values is that they avoid the somewhat tricky problem of deciding on a null hypothesis. In economics, researchers typically set up tests with the null hypothesis being that a model doesn't work, or doesn't work better than existing theory, and then look for evidence to reject the null. (In event studies, for example, the null is that stock prices do *not* react to the event.) But asset-pricing tests often reverse the idea: the null is that a model works perfectly – zero pricing errors – which is 'accepted' as long as we don't find evidence to the contrary. This strikes us as a troubling shift in the burden of proof, particularly given the limited power of many tests. Confidence intervals avoid this problem because they simply show the full range of true parameters that are consistent with the data.

We apply these prescriptions to a handful of proposed models from the recent literature. The results are disappointing. None of the five models that we consider performs well in our tests, despite the fact that all seemed quite promising in the original studies.

The paper proceeds as follows. Section 2 formalizes our critique of asset-pricing tests, Section 3 offers suggestions for improving the tests, and Section 4 applies these prescriptions to several recent models. Section 5 concludes.

2. Interpreting asset-pricing tests

Our analysis uses the following notation. Let R be the vector of excess returns on N test assets (in excess of the riskfree rate) and F be a vector of K risk factors that perfectly explain expected returns on the assets, i.e., $\mu \equiv E[R]$ is linear in the $N \times K$ matrix of stocks' loadings on the factors, $B \equiv \text{cov}(R, F) \text{var}^{-1}(F)$. For simplicity, and without loss of generality, we assume the mean of F equals the cross-sectional risk premium on B , implying $\mu = B \mu_F$. Thus, our basic model is

$$R = B F + e, \quad (1)$$

where e are mean-zero residuals with $\text{cov}(e, F) = 0$. We make no assumptions at this point about the covariance matrix of e , so the model is completely general (eq. 1 has no economic content).

We follow the convention that all vectors are column vectors unless otherwise noted. For generic random variables x and y , $\text{cov}(x, y) \equiv E[(x - \mu_x)(y - \mu_y)']$; i.e., the row dimension is determined by x and the column dimension is determined by y . We use $\mathbf{1}$ to denote a conformable vector of ones, $\mathbf{0}$ to denote a conformable vector or matrix of zeros, and I to denote a conformable identity matrix. M denotes the matrix $I - \mathbf{1}\mathbf{1}'/d$ that transforms, through pre-multiplication, the columns of any matrix with row dimension d into deviations from the mean.

The factors in F can be thought of as a 'true' model that is known to price assets; it will serve as a benchmark but we won't be interested in it per se. Instead, we want to test a proposed model P consisting of J factors. The matrix of assets' factor loadings on P is denoted $C \equiv \text{cov}(R, P) \text{var}^{-1}(P)$, and we'll say that P 'explains the cross section of expected returns' if $\mu = C \gamma$ for some risk premium vector γ . Ideally, γ would be determined by theory.

A common way to test whether P is a good model is to estimate a cross-sectional regression of expected returns on factor loadings

$$\mu = z \mathbf{1} + C \lambda + \eta, \quad (2)$$

where λ denotes a $J \times 1$ vector of regression slopes. In principle, we could test three features of eq. (2): (i) z should be roughly zero (that is, the zero-beta rate should be close to the riskfree rate); (ii) λ should be non-zero and may be restricted by theory; and (iii) η should be zero and the cross-sectional R^2 should be one. In practice, empirical tests often focus only on the restrictions that $\lambda \neq 0$ and the cross-sectional R^2 is one (the latter is sometimes treated only informally). The following observations consider the conditions under which P will appear well-specified in such tests.

Observation 1. *Suppose F and P have the same number of factors and P is correlated with R only through the common variation captured by F , by which we mean that $\text{cov}(e, P) = 0$ (e is the residual in eq. 1). Assume, also, that the correlation matrix between F and P is nonsingular. Then expected returns are exactly linear in stocks' loadings on P – even if P has arbitrarily small (non-zero) correlation with F and explains very little of the time-series variation in returns.*

Proof: The assumption that $\text{cov}(e, P) = 0$ implies $\text{cov}(R, P) = B \text{cov}(F, P)$. Thus, stocks' loadings on P are linearly related to their loadings on F : $C \equiv \text{cov}(R, P) \text{var}^{-1}(P) = B Q$, where $Q \equiv \text{cov}(F, P) \text{var}^{-1}(P)$ is the nonsingular matrix of slope coefficients when F is regressed on P . It follows that $\mu = B \mu_F = C \lambda$, where $\lambda = Q^{-1} \mu_F$. \square

Observation 1 says that, if P has the same number of factors as F , testing whether expected returns are linear in betas with respect to P is essentially the same as testing whether P is uncorrelated with e – a test that doesn't seem to have much economic meaning in recent empirical applications. For example, in tests with size and B/M portfolios, we know that R_M , SMB, and HML (the 'true' model F in our notation) capture nearly all (more than 92%) of the time-series variation in returns, so the residual in $R = B F + e$ is both small and largely idiosyncratic. In that setting, we don't find it surprising that almost any proposed macroeconomic factor P is correlated with returns primarily through R_M , SMB, and HML – indeed, we would be more surprised if $\text{cov}(e, P)$ wasn't close to zero. In turn, we are not at all surprised that many proposed models seem to 'explain' the cross-section of expected size and B/M returns. The strong factor structure of size and B/M portfolios makes it likely that stocks' betas on almost any proposed factor will line up with their expected returns.¹

Put differently, Observation 1 provides a skeptical interpretation of recent asset-pricing tests, in which unrestricted cross-sectional regressions (or equivalently SDF tests, as we explain below) have become the norm. In our view, the empirical tests say little more than that a number of proposed factors are correlated with SMB and HML, a fact that might have some economic content but seems like a pretty low hurdle to meet in claiming that a proposed model explains the size and B/M effects. We offer a number of suggestions for improving the tests below.

Observation 2. *Suppose returns have a strict factor structure with respect to F , i.e., $\text{var}(e)$ is a diagonal*

¹ This argument works cleanly if a proposed model has (at least) three factors. It should also apply when P has two factors since size–B/M portfolios all have multiple-regression market betas close to one. In essence, the two-factor model of SMB and HML explains most of the cross-sectional variation of expected returns, so a proposed model really needs only two factors (as long as we ignore restrictions on the intercept).

matrix. Then any randomly chosen set of K assets perfectly explains the cross section of expected returns so long as the K assets aren't asked to price themselves (that is, the K assets aren't included as test assets on the left-hand side of the cross-sectional regression and the cross-sectional risk premia aren't required to equal the expected returns on the K assets). The only restriction is that R_K , the return on the K assets, must be correlated with F , i.e., $\text{cov}(F, R_K)$ must be nonsingular.

Proof: Let $P = R_K$ in Observation 1 and re-define R as the vector of returns for the remaining $N - K$ assets and e as the residuals for these assets. The strict factor structure implies that $\text{cov}(e, R_K) = \text{cov}(e, B_K F + e_K) = 0$. The result then follows immediately from Observation 1. \square

Observation 2 is useful for a couple of reasons. First, it provides a simple illustration of our argument that, in some situations, it is easy to find factors that explain the cross section of expected returns: under the fairly common assumption (in the APT literature) of a strict factor structure, any collection of K assets will work. Obtaining a high cross-sectional R^2 just isn't very difficult when returns have a strong factor structure, as they do in most empirical applications.

Second, Observation 2 illustrates that it can be important to take seriously restrictions on the cross-sectional slopes. In particular, Observation 2 hinges on the fact that the K asset factors aren't asked to price themselves, i.e., that the cross-sectional risk premia aren't restricted to equal the vector of expected returns on the K assets, as asset-pricing theory would predict. To see this, Observation 1 (proof) shows that the cross-sectional slopes on C are $\lambda = Q^{-1} \mu_F$, where Q is the matrix of slope coefficients when F is regressed on R_K . In the simplest case with one factor, λ simplifies to μ_K / ρ^2 , where ρ is the correlation between R_K and F .² The slope λ is clearly greater than μ_K unless R_K is perfectly correlated with F . The implication is that the problem highlighted by Observations 1 and 2 – that ‘too many’ proposed factors explain the cross section of expected returns – would be less severe if the restriction on λ was taken seriously (e.g., R_K would then price the cross section only if $\rho = 1$).

Observations 1 and 2 are rather special since, in order to get clean predictions, we've assumed that a proposed model P has the same number of factors as the known model F . The intuition goes through when $J < K$ because, even in that case, we would expect the loadings on proposed factors to line up (imperfectly) with expected returns if the assets have a strong factor structure. The next observation generalizes our results, at the cost of changing the definitive conclusion in Observations 1 and 2 into a

² This follows from the fact that $Q^{-1} = \text{var}(R_K) / \text{cov}(R_K, F)$ and $\mu_K = B_K \mu_F = \mu_F \text{cov}(R_K, F) / \text{var}(F)$.

probabilistic statement.

Observation 3. *Suppose F has K factors and P has J factors, with $J \leq K$. Assume, as before, that P is correlated with R only through the factor F [$\text{cov}(e, P) = 0$], and that P and F are correlated, so that $\text{cov}(F, P)$ has rank J . In a generic sense, made precise below, the cross-sectional R^2 in a regression of μ on C is expected to be J/K .*

Proof: By a ‘generic sense,’ we mean that we don’t have any information about the contribution of each of the factors in F in explaining the cross-section of expected returns and so treat the contributions as random. More specifically, suppose the factor loadings on F satisfy $V_B = B'MB / N = I_K$, i.e., they are cross-sectionally uncorrelated and have unit variances; this assumption is without loss of generality since F can always be transformed to make the assumption hold. A ‘generic sense’ means that we view the risk premia on the transformed factors as being random draws from a normal distribution with mean zero and variance σ_γ^2 . The proof then proceeds as follows: In a regression of μ on C , the risk premia are $\lambda = (C'MC)^{-1}C'M\mu$ and the R^2 is $\lambda'C'M\lambda / \mu'M\mu$. By assumption $\mu = B\mu_F$ and Observation 1 (proof) shows that $C = BQ$, where $Q \equiv \text{cov}(F, P) \text{var}^{-1}(P)$. Substituting for λ , μ , and C , and using the assumption that $V_B = I_N$, the R^2 simplifies to $\mu_F'Q(Q'Q)^{-1}Q'\mu_F / \mu_F'\mu_F$, where $Q(Q'Q)^{-1}Q'$ is a symmetric, idempotent matrix of rank J . The risk premia, μ_F , are assumed $MVN[0, \sigma_\gamma^2 I_K]$, as explained above, from which it follows that the R^2 has a Beta distribution with mean J/K .³ \square

Observation 3 generalizes Observations 1 and 2. Our earlier results show that, if a K -factor model explains both the cross section of expected returns and much of the time-series variation in returns, then it should be easy to find other K -factor models that also explain the cross section of expected returns. The issue is a bit messier with $J < K$. Intuitively, the more factors that are in the proposed model, the easier it should be to find a high cross-sectional R^2 as long as the proposed factors are correlated with the ‘true’ factors. Thus, we aren’t surprised at all if a proposed three-factor model explains the size and B/M effects about as well as the Fama-French factors, nor are we surprised if a one- or two-factor model has some explanatory power. We *are* impressed if a one-factor model works as well as the Fama-French factors, since this requires that a single factor captures the pricing information in both SMB and HML. [We note again that size–B/M portfolios all have Fama-French three-factor market betas close to one, so the model can be thought of as a two-factor model (SMB and HML) for the purposes of explaining cross sectional variation in expected returns.]

³ The distribution follows from the fact that R^2 can be expressed as $z_1 / (z_1 + z_2)$, where z_1 and z_2 are independent, chi-squared variables with degrees of freedom J and $K-J$; see Muirhead, 1982, Thm. 1.5.7.

Figure 1 (on the next page) illustrates these results using Fama and French's 25 size-B/M portfolios, getting away from the specific assumptions underlying Observations 1 – 3. We calculate quarterly excess returns on the 25 portfolios from 1963–2004 and explore, in several simple ways, how easy it is to find factors that explain the cross section of average returns. The figure treats the average returns and sample covariance matrix as population parameters; thus, like Observations 1 – 3, it focuses on explaining expected returns in population, not on sampling issues (which we consider later).

Each of the panels reports simulations using artificial factors to explain expected returns. In Panel A, the factors are constructed to produce a random vector of return betas: a 25×1 vector of loadings (for the 25 size-B/M portfolios) is randomly drawn from a MVN distribution with mean zero and covariance matrix proportional to the return covariance matrix. Thus, although the artificial factors aren't designed to explain expected returns, the loadings will tend to line up with expected returns (positively or negatively) simply because their cross-sectional pattern is determined by the covariance structure of returns. This procedure matches the spirit of Observations 1 – 3 but doesn't impose the requirement that the artificial factors covary only with common components in size-B/M portfolios ($\text{cov}(e, P)$ doesn't have to be zero), though the common components will tend to dominate simply because they are so important.

[An alternative interpretation of these simulations is to note that if we generate a time series of artificial factors uncorrelated with returns, the covariance matrix of *estimated* betas (in a multivariate regression of portfolio returns on the factor) is proportional to the covariance matrix of returns. Thus, the loadings in the simulations can be interpreted as sample betas for random ('useless') factors, and the population R^2 can be interpreted as a sample R^2 when size-B/M portfolios' average returns are regressed on these sample betas. The simulations show how often we expect to find high R^2 's if researchers simply come up with factors that have nothing to do with returns.]

Panel A shows that it is easy to find factors that help explain expected returns on the size-B/M portfolios. With one factor, half of our factors produce an R^2 greater than 0.12 and 25% produce an R^2 greater than 0.30 (the latter isn't reported in the figure). With three factors, the median R^2 is 0.51 and the 75th percentile is 0.64, and with five factors, the median and 75th percentiles are a remarkable 0.68 and 0.76, respectively. Roughly half of our artificial three-factor models and 86% of our artificial five-factor models explain more than half of the cross-sectional variation in expected returns.


Panel B performs a similar exercise but, rather than randomly generate loadings, we randomly generate factors that are zero-investment combinations of the size-B/M portfolios: a 25×1 vector of weights is

Figure 1. Population R²s for artificial factors.

This figure explores how easy it is to find factors that explain, in population, the cross section of expected returns on Fama and French’s 25 size-B/M portfolios. We randomly generate factors – either factor loadings directly or zero-investment factor portfolios, as described in the figure – and estimate the population R² when the size-B/M portfolios’ expected returns are regressed on their factor loadings. The average returns and covariance matrix of the portfolios, quarterly from 1963 – 2004, are treated as population parameters in the simulations. The plots are based on 5,000 draws of 1 to 5 factors.


Panel A: Random draws of factor loadings.

Loadings for the 25 size-B/M portfolios are drawn from a MVN distribution with mean zero and covariance matrix proportional to the return covariance matrix.




Panel B: Random draws of factor portfolios.

Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution.



Panel C: Random draws of zero-mean factor portfolios.

Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution, but only factors with roughly zero expected returns are kept.



constructed by independently drawing from a normal distribution with mean zero and variance one (the weights are shifted and re-scaled to have a cross-sectional mean that is exactly zero and to have one dollar long and one dollar short). These simulations show how easy it is to stumble across factors that help explain the cross section of expected returns. As in Panel A, betas on the artificial factors will tend to line up with expected returns simply because of the covariance structure of returns, even though the factors aren't chosen to have any explanatory power. In fact, Panel B shows that the artificial factors here are even better at explaining expected returns: with one, three, and five factors, the median R^2 s are 0.14, 0.77, and 0.84, while the 75th percentiles are 0.40, 0.81, and 0.87, respectively.

Finally, Panel C repeats the simulations in Panel B with a small twist: we keep only those artificial factors that have roughly zero expected returns [the factors in Panel B are expected to have zero expected returns ($E[\mu'x] = 0$ across draws of x) but don't because of random variation in x]. These simulations illustrate that it can be very important to impose restrictions on the cross-sectional slopes when possible; in particular, theory says that the risk premia on our artificial factors should be zero, equal to their expected returns, but Panel C ignores this restriction and just searches for the best possible fit in the cross-sectional regression. Thus, the actual R^2 differs from zero simply because we ignore the theoretical restrictions on the cross-sectional slopes and intercept. The additional degrees of freedom turn out to be very important, especially with multiple factors: with one, three, and five factors in Panel C, the median R^2 s are 0.03, 0.51, and 0.64, while the 75th percentiles are 0.12, 0.60, and 0.68, respectively (again, properly restricted R^2 s would all be close to zero).

The results above illustrate that the covariance structure of size-B/M portfolios makes it easy to find factors that produce high *population* cross-sectional R^2 s. Our final two observations show that the problem is similar in SDF tests and exacerbated by sampling issues.

Observation 4. *Suppose F has K factors and P has J factors, with $J \leq K$. Assume, as before, that P is correlated with R only through the factor F [$\text{cov}(e, P) = 0$], and that P and F are correlated, so that $\text{cov}(F, P)$ has rank J . In a generic sense, made precise below, the sum of squared pricing errors in an SDF framework, $\varepsilon = E[mR]$, are expected to be $q(K - J)$, where q is defined below. The pricing errors are exactly zero when $J = K$.*

Proof: By a 'generic sense,' we again mean that we don't have any information about the contribution of each factor in F in explaining the cross section of expected returns, but we operationalize the idea slightly differently here (a similar, somewhat messier, result holds if we use the earlier definition). Specifically,

suppose the factor loadings on F satisfy $B'B/N = I_K$; this assumption is without loss of generality since F can always be transformed to make the assumption hold. A ‘generic sense’ means that we view the risk premia on the transformed B as being random draws from a normal distribution with mean α and variance σ_γ^2 . The proof proceeds as follows: Define the SDF as $m = a - b'P$ and the pricing errors as $\varepsilon = E[mR]$. Using excess returns, the SDF is defined up to a constant of proportionality (see Cochrane, 2001); we therefore fix a and find the corresponding b in the SDF. In first-stage GMM, $\min_b \varepsilon'\varepsilon$, the solution is $b = a(D'D)^{-1}D'\mu$ and $\varepsilon = a[I_N - D(D'D)^{-1}D']\mu$, where $D \equiv \text{cov}(R, P)$. The sum of squared pricing errors is $\varepsilon'\varepsilon = a^2\mu'[I_N - D(D'D)^{-1}D']\mu$. By assumption $\mu = B\mu_F$ and $B'B/N = I_K$, so $\varepsilon'\varepsilon$ can be rewritten as $a^2N\mu_F'H\mu_F$, where $H = I_K - B'D(D'D)^{-1}D'B/N$. The matrix H is symmetric and idempotent with rank $K - J$ and the risk premia are assumed to be $MVN[\alpha, \sigma_\gamma^2 I_K]$, as explained above, from which it follows that $\mu_F'H\mu_F$ is proportional to a noncentral chi-squared variate, with mean $(\alpha^2 + \sigma_\gamma^2)(K - J)$ (ε and $\mu_F'H\mu_F$ are exactly zero when $J = K$). The sum of squared errors has expectation $E[\varepsilon'\varepsilon] = a^2 N (\alpha^2 + \sigma_\gamma^2)(K - J)$. \square

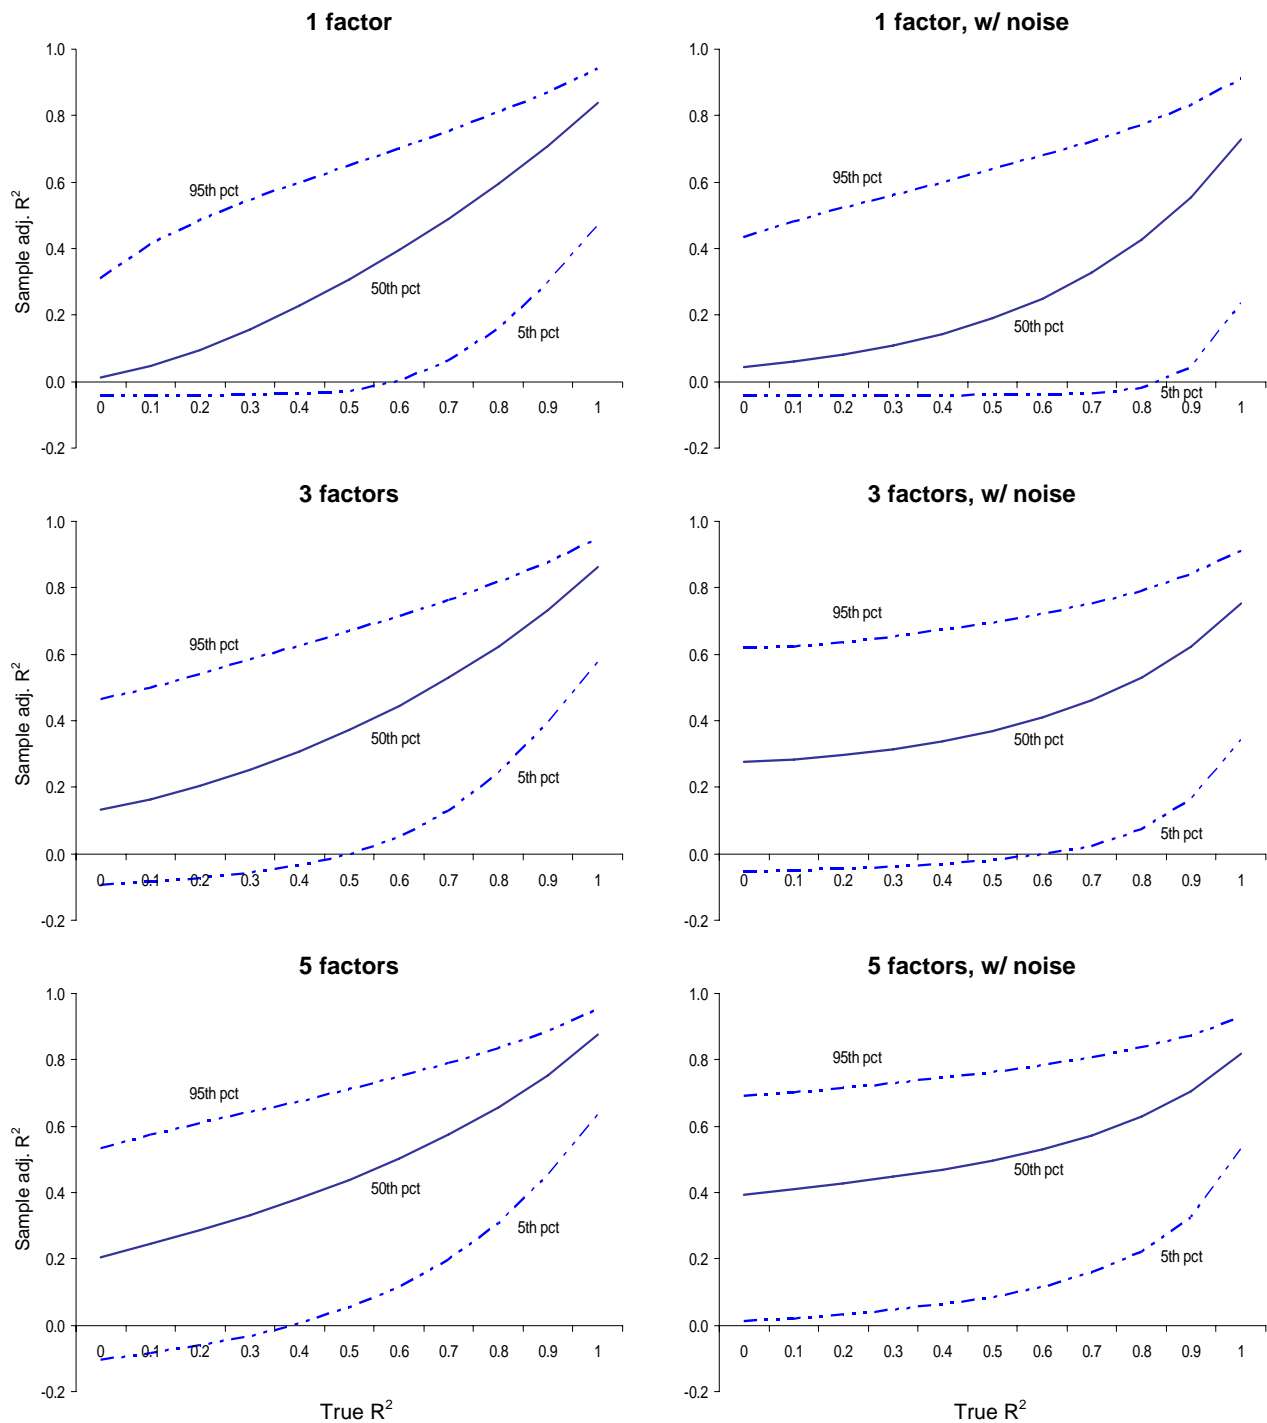
Observation 4 is the SDF equivalent of our earlier cross-sectional R^2 results. It says that, as long as proposed factors P covary with returns only through the factors F – an assumption that seems likely to hold for just about any proposed factor when the test assets are size-B/M portfolios – the model will help reduce SDF pricing errors, $\varepsilon = E[mR]$. The errors are expected to be smaller the more factors that are in P and, in the limit, drop to zero when P has the same number of factors as F. The magnitude of the errors when $J < K$ can be interpreted by noting that with no factor, $J = 0$, the sum of squared pricing errors is expected to be $a^2 N (\alpha^2 + \sigma_\gamma^2) K$ given our assumptions. Thus, every factor reduces $E[\varepsilon'\varepsilon]$ by a fraction $1/K$. The decline in pricing errors is nearly mechanical in tests with size-B/M portfolios, because of their strong covariance structure, and has little economic meaning.

Observation 5. *The problems are exacerbated by sampling issues: If returns have a strong factor structure, it can be easy to find a high sample cross-sectional R^2 even in the unlikely scenario that the population R^2 is small or zero.*

Observation 5 is intentionally informal and, in lieu of a proof, we offer simulations using Fama and French’s 25 size-B/M portfolios to illustrate the point. The simulations differ from those in Figure 1 because, rather than study the population cross-sectional R^2 for artificial factors, we now focus on sampling variation in estimated R^2 s conditional on a given population R^2 . The simulations have two steps: First, we fix a true cross-sectional R^2 that we want a model to have and randomly generate a matrix of factor loadings C which produces that R^2 . Factor portfolios, $P = w'R$, are constructed to have those

Figure 2: Sample distribution of the cross-sectional adj. R^2 .

This figure shows the sample distribution of the cross-sectional adj. R^2 (average returns regressed on estimated factor loadings) for Fama and French's 25 size-B/M portfolios from 1963 – 2004 (quarterly returns). The plots use one to five randomly generated factors that together have the true R^2 reported on the x-axis. In the left-hand panels, the factors are combinations of the size-B/M portfolios (the weights are randomly drawn to produce the given R^2 , as described in the text). In the right-hand panels, noise is added to the factors equal to 3/4 of a factor's total variance, to simulate factors that are not perfectly spanned by returns. The plots are based on 40,000 bootstrap simulations (10 sets of random factors; 4,000 simulations with each).



factor loadings, i.e., we find portfolio weights, w , such that $\text{cov}(R, P) = \text{var}(R) w$ is linear in C .⁴ Second, we bootstrap artificial time series of returns and factors by sampling, with replacement, from the historical time series of size-B/M returns (quarterly, 1963–2004). We then estimate the sample cross-sectional adj. R^2 for the artificial data by regressing average returns on estimated factor loadings. The second step is repeated 4,000 times to construct a sampling distribution of the adj. R^2 . In addition, to make sure the particular matrix of loadings generated in step 1 isn't important, we repeat that step 10 times, giving us a total sample of 40,000 adj. R^2 's corresponding to an assumed true R^2 .

Figure 2 shows results for models with 1, 3, and 5 factors. The left-hand column plots the distribution of the sample adjusted R^2 (5th, 50th, and 95th percentiles) corresponding to true R^2 's of 0.0 to 1.0 for models in which the factors are portfolio returns, as described above. The right-hand column repeats the exercise but uses factors that are imperfectly correlated with returns, as they are in most empirical applications; we start with the portfolio factors used in the left-hand panels and add noise equal to 3/4 of their total variance. Thus, for the right-hand plots, a maximally correlated combination of the size-B/M portfolios would have a time-series R^2 of 0.25 with each factor.

The figure shows that a sample R^2 needs to be quite high to be statistically significant, especially for models with several factors. Focusing on the right-hand column, the 95th percentile of the sampling distribution using one factor is 44%, using three factors is 62%, and using five factors is 69% – when the true cross-sectional R^2 is zero! Thus, even if we could find factors that have no true explanatory power (something that seems unlikely given our population results above), it still wouldn't be terribly surprising to find fairly high R^2 's in sample. Further, with either three or five factors, the ability of the sample R^2 to discriminate between good and bad models is quite small, since the distribution of the sample R^2 is similar across a wide range of true R^2 's. For example, with five factors, a sample R^2 greater than 73% is needed to reject that the true R^2 is 30% or less, at a 5% one-sided significance level, but that outcome is unlikely even if the true R^2 is 70% (probability of 0.17) or 80% (probability of 0.26). The bottom line is that, in both population and sample, a high cross-sectional R^2 seems to provide little information about whether a proposed model is good or bad.

⁴ More specifically, for a model with J factors, we randomly generate J vectors, g_j , that are uncorrelated with each other and which individually have explanatory power of $c^2 = R^2/J$ (and, thus, the correct combined R^2). Each vector is generated as $g_j = c \mu_s + (1-c^2)^{1/2} e_j$, where μ_s is the vector of expected returns on the size-B/M portfolios, shifted and re-scaled to have mean zero and standard deviation of one, and e_j is generated by randomly drawing from a standard normal distribution (e_j is transformed to have exactly mean zero and standard deviation of one, to be uncorrelated with μ , and to make $\text{cov}(g_i, g_j) = 0$ for $i \neq j$). The factor portfolios in the simulations have covariance with returns (a 25×1 vector) given by the g_j .

Related research

Our appraisal of asset-pricing tests overlaps with a number of studies. Roll and Ross (1994) and Kandel and Stambaugh (1995) argue that the cross-sectional R^2 in simple CAPM tests isn't very meaningful because, as a theoretical matter, it tells us little about the location of the market proxy in mean-variance space (see also Kimmel, 2003). We reach a similarly skeptical conclusion about the R^2 but emphasize different issues. The closest overlap comes from our simulations in Panel C of Figure 1, which show that factor portfolios with zero mean returns might still produce high R^2 in unrestricted cross-sectional regressions. These portfolios are far from the mean-variance frontier by construction – they have zero Sharpe ratios – yet often have high explanatory power, consistent with the results of Roll and Ross and Kandel and Stambaugh.

Kan and Zhang (1999) study cross-sectional tests with 'useless' factors, defined as factors that are uncorrelated in population with returns. They show that the usual asymptotics break down because the cross-sectional spread in estimated loadings goes to zero as T gets big (since all the loadings go to zero). Our simulations in Panel A of Figure 1 have some overlap since, as pointed out earlier, they can be interpreted as showing the sample R^2 when randomly generated useless factors are used to explain returns. The issues are different since our simulations generate random factors but hold the time series of returns constant (thus, they don't really consider the sampling issues discussed by Kan and Zhang). More broadly, our results are different because we focus on population R^2 's and, when we do look at sampling distributions in Fig. 2, the factors are not 'useless.'

Some of our results are reminiscent of the literature on testing the APT and multifactor models (see, e.g., Shanken 1987, Reisman, 1992; Shanken, 1992a). Most closely, Nawalkha (1997) derives results like Observations 1 and 2 above, though his focus is different. In particular, he emphasizes that, in the APT, 'well-diversified' variables (those uncorrelated with idiosyncratic risks) can be used in place of the 'true' factors without any loss of pricing accuracy. We generalize his theoretical results to models with $J < K$ proposed factors, consider sampling issues, and emphasize the empirical implications for recent tests using size-B/M portfolios.

Finally, our critique is similar in spirit to a contemporaneous paper by Daniel and Titman (2005). They show that, even if characteristics determine expected returns (e.g., expected returns are linear in B/M), a proposed factor can appear to price characteristic-sorted portfolios simply because, in the underlying population of stocks, factor loadings and characteristics are correlated (forming portfolios tends to inflate that correlation). Our ultimate conclusions about using characteristic-sorted portfolios are similar but we

highlight different concerns, emphasizing the importance of the factor structure of size-B/M portfolios, the impact of using many factors and not imposing restrictions on the cross-sectional slopes, and the role of both population and sampling issues.

3. How can we improve empirical tests?

The theme of Observations 1 – 5 is that, in situations like those encountered in practice, it may be easy to find factors that explain the cross section of expected returns. Finding a high cross-sectional R^2 or small pricing errors often has little economic meaning and, in our view, should not be taken as providing much support for a proposed model. The problem is not just a sampling issue – it cannot be solved by getting standard errors right – though sampling issues exacerbate the problem. Here, we offer a few suggestions for improving empirical tests.

Prescription 1. Expand the set of test portfolios beyond size-B/M portfolios.


Due to the importance of the size and value anomalies, empirical tests often focus on size-B/M portfolios. This practice is understandable but problematic, since the concerns highlighted above are most severe when a couple of factors explain nearly all of the time-series variation in returns, as is true for size-B/M portfolios. One simple solution, then, is to include portfolios that don't correlate as strongly with SMB and HML. Reasonable choices include industry-, beta-, volatility-, or factor-loading-sorted portfolios (the last being loadings on a proposed factor; an alternative would be to use individual stocks in the regression, though errors-in-variables problems could make this impractical). Bond portfolios might also be used. The idea is to price the portfolios all at the same time, not in separate cross-sectional regressions. Also, the additional portfolios don't need to offer a big spread in expected returns; the goal is simply to relax the tight factor structure of size-B/M portfolios.

Figure 3 illustrates this idea. We replicate the simulations in Figure 1 but, rather than use size-B/M portfolios alone, we augment them with Fama and French's 30 industry portfolios. As before, we generate artificial factors and explore how well they explain, in population, the cross section of expected returns (average returns and covariances from 1963 – 2004 are treated as population parameters). The artificial factors are generated in three ways. In Panel A, the factors are constructed to produce a randomly chosen 55×1 vector of factor loadings, drawing from a MVN distribution with mean zero and covariance matrix proportional to the return covariance matrix. In Panel B, the factors are constructed by randomly drawing a 55×1 vector of portfolio weights from a standard normal distribution. And in Panel C, we repeat the simulations of Panel B but keep only the factor portfolios that have (roughly) zero


Figure 3. Population R^2 's for artificial factors: Size-B/M and industry portfolios.

This figure compares how easy it is to find factors that explain, in population, the cross section of expected returns on Fama and French's 25 size-B/M portfolios (dotted lines) vs. 55 portfolios consisting of the 25 size-B/M portfolios and Fama and French's 30 industry portfolios (solid lines). We randomly generate factors – either factor loadings directly or zero-investment factor portfolios, as described in the figure – and estimate the population R^2 when the portfolios' expected returns are regressed on their factor loadings. The average returns and covariance matrix of the portfolios, quarterly from 1963 – 2004, are treated as population parameters in the simulations. The plots are based on 5,000 draws of 1 to 5 factors.

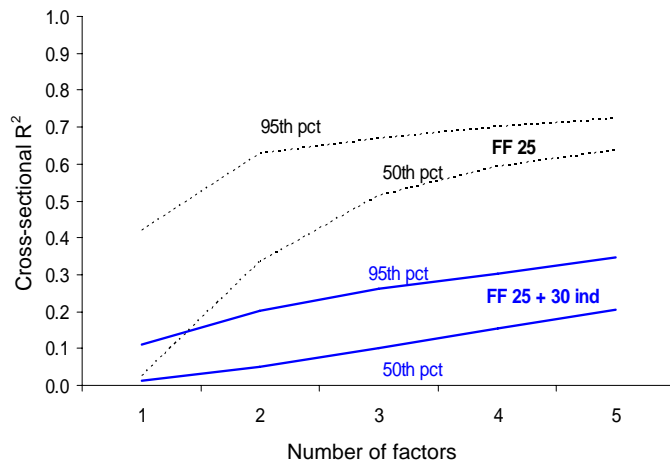
Panel A: Random draws of factor loadings.
 Loadings for the 25 size-B/M portfolios are drawn from a MVN distribution with mean zero and covariance matrix proportional to the return covariance matrix.



Panel B: Random draws of factor portfolios.
 Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution.



Panel C: Random draws of zero-mean factor portfolios.
 Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution, but only factors with roughly zero expected returns are kept.



expected returns. The point in each case is to explore how easy it is to find factors that produce a high cross-sectional R^2 (in population). We refer the reader to the discussion of Figure 1 for the logic and interpretation of each set of simulations.

Figure 3 shows that it is much ‘harder’ to explain expected returns, using artificial factors, on the 55 portfolios than on the 25 size-B/M portfolios (the median and 95th percentiles for the latter are repeated from Fig. 1 for comparison). For example, with three factors, the median R^2 for the full set of 55 portfolios is 15% in Panel A, 20% in Panel B, and 10% in Panel C, compared with median R^2 s for the 25 size-B/M portfolios of 50%, 77%, and 51%, respectively. The difference between the 25 size-B/M portfolios and the full set of portfolios is largest for models with at least three factors, consistent with the three-factor structure of size-B/M portfolios being important. In short, the full set of portfolios seems to provide a more rigorous test of a proposed model.

Prescription 2. Take the magnitude of the cross-sectional slopes seriously.

The recent literature sometimes emphasizes a model’s high cross-sectional R^2 but doesn’t consider whether the estimated slopes and zero-beta rates are reasonable. Yet theory often provides guidance for both that should be taken seriously, i.e., the theoretical restrictions should be imposed ex ante or tested ex post. Most clearly, theory says the zero-beta rate should equal the riskfree rate. The standard retort is that Brennan’s (1971) model relaxes this constraint if borrowing and lending rates differ, but this argument isn’t convincing in our view: (riskless) borrowing and lending rates just aren’t sufficiently different – perhaps 1% annually – to justify the extremely high zero-beta estimates in many papers. An alternative argument is that the equity premium is anomalously high, à la Mehra and Prescott (1985), so it’s unreasonable to ask a consumption-based model to explain it. But it isn’t clear why we should accept a model that doesn’t explain the level of expected returns.

A related restriction, mentioned earlier, is that the risk premium for any factor portfolio should be the portfolio’s expected excess return. For example, the cross-sectional price of market-beta risk should be the market equity premium; the price of yield-spread risk, captured by movements in long-term Tbond returns, should be the expected Tbond return over the riskfree rate. In practice, this type of restriction could be tested in cross-sectional regressions or, better yet, imposed ex ante by focusing on time-series regression intercepts (Jensen’s alphas). Below, we discuss ways to incorporate the constraint into cross-sectional regressions (see, also, Shanken, 1992b).

As a third example, conditional models generally imply concrete restrictions on cross-sectional slopes, a point emphasized by Lewellen and Nagel (2006). For example, Jagannathan and Wang (1996) show that a one-factor conditional CAPM implies a two-factor unconditional model: $E_{t-1}[R_t] = \beta_t \gamma_t \rightarrow E[R] = \beta \gamma + \text{cov}(\beta_t, \gamma_t)$, where β_t and γ_t are the conditional beta and equity premium, respectively, and β and γ are their unconditional means. The cross-sectional slope on $\varphi_i = \text{cov}(\beta_{it}, \gamma_t)$, in the unconditional regression, should clearly be one but that constraint is often ignored in the literature. Lewellen and Nagel discuss this issue in detail and provide empirical examples from recent tests of both the simple and consumption CAPMs. For tests of the simple CAPM, the constraint is easily imposed using the conditional time-series regressions of Shanken (1990), if the relevant state variables are all known, or the short-window approach of Lewellen and Nagel, if they are not.

Prescription 3. Report GLS cross-sectional R^2 s.

The literature typically favors OLS over GLS cross-sectional regressions. The rationale for neglecting GLS regressions appears to reflect concerns with (i) the statistical properties of (feasible) GLS and (ii) the apparent difficulty with interpreting the GLS R^2 , which, on the surface, simply tells us about the model's ability to explain expected returns on 're-packaged' portfolios, not the basic portfolios that are of direct interest (e.g., if μ and B are expected returns and factors loadings for size-B/M portfolios, OLS regresses μ on $X \equiv [1 \ B]$ while GLS regresses $Q \mu$ on $Q X$, where Q is an $N \times N$ matrix such that $\text{var}^{-1}(R) = Q'Q$). We believe these concerns are misplaced (or at least overstated) and that the GLS R^2 actually has a number of advantages over the OLS R^2 .

The statistical concerns with GLS are real but not prohibitive. The main issue is that, since the covariance matrix of returns must be estimated, the exact finite-sample properties of GLS are generally unknown and textbook econometrics emphasizes GLS's asymptotic properties, which can be a poor approximation when the number of assets is large relative to the length of the time series (Gibbons, Ross, and Shanken, 1989, provide examples in a closely related context; see also Shanken and Zhou, 2006). But we see no reason this problem can't be overcome using standard simulation methods or, in special cases, using the finite-sample results of Shanken (1985) or Gibbons et al.

The second concern – that the GLS R^2 is hard to interpret – also seems misplaced. In fact, Kandel and Stambaugh (1995) show that the GLS R^2 is in many ways a more meaningful statistic than the OLS R^2 : when expected returns are regressed on betas with respect to a factor portfolio, the GLS R^2 is completely determined by the factor's proximity to the minimum-variance boundary while the OLS R^2 has little

connection, in general, to the factor's location in mean-variance space (see also Roll and Ross, 1994; this result assumes the factor is spanned by the test assets). Thus, if a market proxy is nearly mean-variance efficient, the GLS R^2 is nearly one but the OLS R^2 can, in principle, be anything. A factor's proximity to the minimum-variance boundary may not be the only metric for evaluating a model, but it does seem to be both economically reasonable and easy to understand.


The same idea applies to models with non-return factors. In this case, Appendix A shows that a GLS regression is equivalent to using maximally-correlated mimicking portfolios in place of the actual factors and imposing the constraint that the risk premia on the portfolios equal their excess returns (in excess of the zero-beta rate if an intercept is included). The GLS R^2 is determined by the mimicking portfolios' proximity to the minimum-variance boundary (more precisely, the distance from the minimum-variance boundary to the 'best' combination of the mimicking portfolios). Again, this distance seems like a natural metric by which to evaluate a model, since any asset-pricing theory boils down to a prediction that the factor-mimicking portfolios span the mean-variance frontier.

One implication of these facts is that obtaining a high GLS R^2 would seem to be a more rigorous hurdle than obtaining a high OLS R^2 : a model can produce a high OLS R^2 even though the factor mimicking portfolios are far from mean-variance efficient, while the GLS R^2 is high only if a model can explain the high Sharpe ratios available on the test portfolios.


The implicit restrictions imposed by GLS aren't a full solution to the problems discussed in Section 2. Indeed, Observations 1 and 2 apply equally to OLS and GLS regressions (i.e., both R^2 s are one given the stated assumptions). But our simulations with artificial factors, which relax the strong assumptions of the formal propositions, suggest that, in practice, finding a high (population) GLS R^2 s is much less likely than finding a high OLS R^2 . Figure 4, on the next page, illustrates this result. The figure shows GLS R^2 s for the same simulations as Figure 1, using artificial factors to explain expected returns on Fama and French's (1993) size-B/M portfolios (treating their sample moments as population parameters; the OLS plots are repeated for comparison). The plots show that, while artificial factors have some explanatory power in GLS regressions, the GLS R^2 s are dramatically lower than OLS R^2 s. The biggest difference is in Panel C, which constructs artificial factors that are random, zero-cost combinations of the 25 size-B/M portfolios, imposing the restriction that the factors' Sharpe ratios are zero. The GLS R^2 s are appropriately zero, since the risk premia on the factors match their zero expected returns, while the OLS R^2 are often 50% or more in models with multiple factors. The simulations suggest that obtaining a high GLS R^2 represents a more stringent hurdle.

Figure 4. Population OLS and GLS R^2 's for artificial factors.


This figure explores how easy it is to find factors that explain, in population, the cross section of expected returns on Fama and French's 25 size-B/M portfolios. We randomly generate factors – either factor loadings directly or zero-investment factor portfolios, as noted – and estimate the population OLS and GLS R^2 's when expected returns are regressed on factor loadings. The average returns and covariance matrix of the size-B/M portfolios, quarterly from 1963 – 2004, are treated as population parameters in the simulations. The plots are based on 5,000 draws of 1 to 5 factors.



Panel A: Random draws of factor loadings. Loadings for the 25 size-B/M portfolios are drawn from a MVN distribution with mean zero and covariance matrix proportional to the return covariance matrix.



Panel B: Random draws of factor portfolios. Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution.



Panel C: Random draws of zero-mean factor portfolios. Zero-investment factors, formed from the size-B/M portfolios, are generated by randomly drawing a 25×1 vector of weights from a standard normal distribution, but only factors with roughly zero expected returns are kept.

Prescription 4. If a proposed factor is a traded portfolio, include the factor as one of the test assets on the left-hand side of the cross-sectional regression.

Prescription 4 builds on Prescription 2, in particular, the idea that the cross-sectional price of risk for a factor portfolio should be the factor's expected excess return. One simple way to build this restriction into a cross-sectional regression is to ask the factor to price itself; that is, to test whether the factor portfolio itself lies on the estimated cross-sectional regression line.

Prescription 4 is most important when the cross-sectional regression is estimated with GLS rather than OLS. As mentioned above, when a factor portfolio is included as a left-hand side asset, GLS forces the regression to price the asset perfectly: the estimated slope on the factor's loading exactly equals the factor's average return in excess of the estimated zero-beta rate (in essence, the asset is given infinite weight in the regression). Thus, a GLS cross-sectional regression, when a traded factor is included as a test portfolio, is similar to the time-series approach of Black, Jensen, and Scholes (1972) and Gibbons, Ross, and Shanken (1989).

Prescription 5. Report confidence intervals for the cross-sectional R^2 .


Prescription 5 is less a solution to the problems highlighted above – indeed, it does nothing to address the concern that it may be easy to find factors that produce high population R^2 s – than a way to make the sampling issues more transparent. We suspect researchers would put less weight on the cross-sectional R^2 if the extremely high sampling error in it was clear (extremely high when using size-B/M portfolios, though not necessarily with other assets). More generally, we find it odd that papers often emphasize this statistic without regard to its sampling properties.

The distribution of the sample R^2 can be derived analytically in special cases, but we're not aware of a general formula or one that incorporates first-stage estimation error in factor loadings. An alternative is to use simulations like those in Figure 2, one panel of which is repeated in Figure 5. The simulations indicate that the sample R^2 (OLS) is often significantly biased and skewed by an amount that depends on the true cross-sectional R^2 . These properties suggest that reporting a confidence interval for R^2 is more meaningful than reporting just a standard error.

The easiest way to get confidence intervals is to 'invert' Figure 5, an approach suggested by Stock (1991) in a different context. In the figure, the sample distribution of the estimated R^2 , for a given true R^2 , is

Figure 5. Sample distribution of the cross-sectional adj. R^2 .

This figure repeats the ‘1 factor’ panel of Fig. 2. It shows the sample distribution of the cross-sectional adj. R^2 , as a function of the true cross-sectional R^2 , for a model with one factor using Fama and French’s 25 size-B/M portfolios from 1963 – 2004 (quarterly returns). The simulated factor is a combination of the size-B/M portfolios (the weights are randomly drawn to produce the given R^2 , as described in Section 2). The plot is based on 40,000 bootstrap simulations (10 sets of simulated factors; 4,000 simulations with each).



found by slicing the picture along the x-axis (fixing x , then scanning up and down). Conversely, a confidence interval for the true R^2 , given a sample R^2 , is found by slicing the picture along the y-axis (fixing y , then scanning across). For example, a sample R^2 of 0.50 implies a 90% confidence interval for the population R^2 of roughly [0.25, 1.00], depicted by the dark dotted line in the graph. The confidence interval represents all values of the true R^2 for which the estimated R^2 falls within the 5th and 95th percentiles of the sample distribution. The extremely wide interval in this example illustrates just how uninformative the sample R^2 can be.

Prescription 6. Report confidence intervals for the (weighted) sum of squared pricing errors.

Prescription 6 has the same goals as Prescription 5: to provide a better summary measure of how well a model performs and to make sampling issues more transparent. Again, Prescription 6 doesn’t address our concern that it may be easy to find factors that produce small population pricing errors for size-B/M portfolios. But reporting confidence intervals should at least make clear when a test has low power: we may not reject that a model performs perfectly (the null of zero pricing errors), but we also won’t reject that the pricing errors are quite large. And the opposite can be true as well: confidence intervals can reveal if a model is rejected not because the pricing errors are economically large but because the tests simply have strong power. In short, confidence intervals allow us to better assess the economic significance of the results.


The (weighted) sum of squared pricing errors (SSPE) is an alternative to the cross-sectional R^2 as a measure of performance. Like the R^2 , sample estimates of the SSPE are generally strongly biased and skewed, suggesting that confidence intervals are more meaningful than standard errors or p-values. The literature considers several versions of such statistics, including Shanken's (1985) cross-sectional T^2 statistic, Gibbons, Ross, and Shanken's (GRS 1989) F-statistic, Hansen's (1982) J-statistic, and Hansen and Jagannathan's (1997) HJ-distance. Confidence intervals for any of these can be obtained using an approach like Figure 4, plotting the sample distribution as a function of the true parameter. We describe here how to get confidence intervals for the GRS F-test, the cross-sectional T^2 (or asymptotic χ^2) statistic, and the HJ-distance, all of which have useful economic interpretations and either accommodate or impose restrictions on the zero-beta rate and risk premia.

The GRS F-statistic tests whether the time-series intercepts (pricing errors) are all zero when excess returns are regressed on a set of factor portfolios, $R = a + B R_p + e$. (The F-test can be used only if the factors are all portfolio returns or if the non-return factors are replaced by mimicking portfolios.) Let \hat{a} be the OLS estimates of a and let $\Sigma \equiv \text{var}(e)$. The covariance matrix of \hat{a} , given a sample for T periods, is $\Omega = c \Sigma$, where $c = (1 + s_p^2) / T$ and s_p^2 is the sample maximum squared Sharpe ratio attainable from combinations of P . GRS show that, under standard assumptions, the weighted sum of squared pricing errors, $S = \hat{a}' \Omega_{OLS}^{-1} \hat{a} = c^{-1} \hat{a}' \Sigma_{OLS}^{-1} \hat{a}$ is asymptotically chi-squared and, if returns are multivariate normal, the statistic $F = c^{-1} \hat{a}' \Sigma_{OLS}^{-1} \hat{a} [(T - N - K) / N (T - K - 1)]$ is small-sample F with non-centrality parameter $\lambda = c^{-1} a' \Sigma^{-1} a$ and degrees of freedom N and $T - N - K$. Moreover, the quadratic $\theta_z^2 = a' \Sigma^{-1} a$ is the model's unexplained squared Sharpe ratio, the difference between the population squared Sharpe ratio of the tangency portfolio (θ_τ^2) and that attainable from R_p (θ_p^2). Thus, a confidence interval for θ_z^2 can be found by inverting a graph like Fig. 5 but showing the sample distribution of F as a function of θ_z^2 (or, more formally, finding the set of θ_z^2 for which F is less than, say, the 95th percentile of an F-distribution with non-centrality parameter $c^{-1} \theta_z^2$).

Figure 6 illustrates the confidence-interval approach for testing the unconditional CAPM. The test uses quarterly excess returns on Fama and French's 25 size-B/M portfolios from 1963 – 2004 and our market proxy is the CRSP value-weighted index. The size and B/M effects are quite strong during this sample (the average absolute quarterly alpha is 0.96% across the 25 portfolios), and the GRS F-statistic strongly rejects the CAPM, $F = 3.491$ with a p-value of 0.000. The graph shows, moreover, that we can reject that the squared Sharpe ratio on the market is within 0.21 of the squared Sharpe ratio of the tangency portfolio: a 90% confidence interval for $\theta_z^2 = \theta_\tau^2 - \theta_M^2$ is $[0.21, 0.61]$. Interpreted differently, following

Figure 6. Sample distribution of the GRS F-statistic and confidence interval for θ_z^2 .

This figure provides a test of the CAPM using quarterly excess returns on Fama and French’s 25 size-B/M portfolios from 1963–2004. The sample distribution of the GRS F-statistic, for a given value of the true unexplained squared Sharpe ratio, θ_z^2 , can be found by slicing the graph along the x-axis (fixing θ_z^2 then scanning up to find percentiles of the sample distribution). A confidence interval for θ_z^2 , given the sample F-statistic, is found by slicing along the y-axis (fixing F then scanning across). θ_z^2 is the difference between the squared Sharpe ratio of the tangency portfolio and that of the CRSP value-weighted index. $F = c^{-1} \hat{a}' \Sigma_{OLS}^{-1} \hat{a} (T-N-1)/[N(T-2)]$; it has an F-distribution with noncentrality parameter $c^{-1}\theta_z^2$ and degrees of freedom N and T–N–1, where N=25, T=168, and $c \approx 1/T$. The sample F-statistic is 3.49 and the corresponding 90% confidence interval for θ_z^2 is depicted by the dark dotted line.



MacKinlay (1995), there exists a portfolio z that is uncorrelated with the market and, with 90% confidence, has a quarterly Sharpe ratio between 0.46 ($=0.21^{1/2}$) and 0.78 ($=0.61^{1/2}$). This compares with a quarterly Sharpe ratio for the market portfolio of 0.18 during this period. The confidence interval provides a good summary measure of just how poorly the CAPM works.

Shanken’s (1985) CSR T^2 test is like the GRS F-test but focuses on pricing errors (residuals) in the cross-sectional regression, $\mu = z \iota + B \lambda + \alpha$. (The T^2 test can be used with non-return factors and doesn’t restrict the zero-beta rate to be r_f , unless the intercept is omitted.) The test is based on the traditional two-pass methodology: Let b be the matrix of factor loadings estimated in the first-pass time-series regression and let $x = [\iota \ b]$ be regressors in the second-pass cross-sectional regression with average returns, r, as the dependent variable. The pricing errors, $\hat{a} = [I - x(x'x)^{-1}x']r$, have asymptotic variance $\Sigma_a = (1 + \lambda' \Sigma_F^{-1} \lambda) y \Sigma y / T$, where $y = I - x(x'x)^{-1}x'$ and the term $(1 + \lambda' \Sigma_F^{-1} \lambda)$ accounts for estimation error in b.⁵ The T^2 -statistic is then $\hat{a}' S_a^+ \hat{a}$, where S_a^+ is the pseudoinverse of the estimated Σ_a (based on consistent estimates of λ , Σ_F , y, and Σ ; the pseudoinverse is necessary because Σ_a is singular). Appendix A shows that T^2 is

⁵ Appendix A explains these results in detail. The variance Σ_a assumes that returns are IID over time and that $\alpha = 0$. Also, Shanken analyzes GLS, not OLS, cross-sectional regressions. The appendix shows that the OLS-based T^2 test described here is equivalent to his GLS-based test.

asymptotically χ^2 with degrees of freedom $N - K - 1$ and non-centrality parameter $n = \alpha' \Sigma_a^{-1} \alpha = \alpha' (y \Sigma y)' \alpha [T / (1 + \lambda' \Sigma_F^{-1} \lambda)]$. The quadratic, $q = \alpha' (y \Sigma y)' \alpha$, again has an economic interpretation: it measures how far factor-mimicking portfolios are from the mean-variance frontier. Specifically, let R_p be K portfolios that are maximally correlated with F , and let $\theta(z)$ be what we'll refer to as a generalized Sharpe ratio, using the zero-beta rate, $r_f + z$, in place of the riskfree rate. The appendix shows that $q = \theta_\tau^2(z) - \theta_p^2(z)$, the difference between the maximum generalized Sharpe ratio on any portfolio and that attainable from R_p . (The zero-beta rate in this definition is the z that minimizes q ; it turns out to be the GLS zero-beta rate.) Therefore, as with the GRS F-test, a confidence interval for q can be found by plotting the sample distribution of the T^2 -statistic as a function of q , using either the asymptotic χ^2 distribution or a simulated small-sample distribution.⁶

The final test we consider, the HJ-distance, differs from the prior two because it focuses on SDF pricing errors, $\varepsilon = E[y(1+R) - 1]$, where $y = a + b P$ is a proposed SDF and we now define R as an $N+1$ vector of total (not excess) returns including the riskless asset. Let m be any well-specified SDF. Hansen and Jagannathan (1997) show that the distance between y and the set of true SDFs, $D = \min_m E[(y-m)^2]$, also equals the largest squared pricing error available on any portfolio relative to its second moment, i.e., $D = \max_x (\varepsilon'x)^2 / E[(1+R_x)^2]$. Using the second definition, the distance is easily shown to equal $D = \varepsilon' H^{-1} \varepsilon$, where $H \equiv E[(1+R)(1+R)']$ is the second moment matrix of gross returns. To get a confidence interval for D , Appendix B shows that $D = \theta_z^2 / (1 + r_f)^2$, where θ_z^2 is the model's unexplained squared Sharpe ratio, as defined earlier (see also Kan and Zhou, 2004). Thus, like the GRS F-statistic, the estimate of D is small-sample F up to a constant of proportionality.⁷ A confidence interval is then easily obtained using the approach described above.

4. Empirical examples

The prescriptions above are straightforward to implement and, while not a complete solution to the problems discussed in Section 2, should help to improve the power and rigor of empirical tests. As an illustration, we report tests for several models that have been proposed recently in the literature. Cross-sectional tests in the original studies focused on Fama and French's 25 size-B/M portfolios, precisely the

⁶ A third possibility, applying the logic of Shanken (1985), would be to replace the asymptotic χ^2 distribution with a finite-sample F distribution (details available on request).

⁷ We assume that the parameters a and b are chosen to minimize D and that the factors are portfolio returns, or that all non-return factors are replaced by mimicking portfolios. In the latter scenario, the small-sample F-distribution would not take into account estimation error in the mimicking-portfolio weights; simulations could be used instead to approximate the sampling distribution.

scenario for which our concerns are greatest. Our goal here is not to disparage the papers – indeed, we believe the studies provide economically important insights – nor to provide a full review of the often extensive empirical tests in each paper, but only to show that our prescriptions can dramatically change how well a model seems to work.

We investigate models for which necessary data are readily available. The models include: (i) Lettau and Ludvigson’s (LL 2001) conditional consumption CAPM (CCAPM), in which the conditioning variable is the aggregate consumption-to-wealth ratio *cay* (available on Ludvigson’s website); (ii) Lustig and Van Nieuwerburgh’s (LVN 2004) conditional CCAPM, in which the conditioning variable is the housing collateral ratio *mymo* (we consider only their linear model with separable preferences; *mymo* is available on Van Nieuwerburgh’s website); (iii) Santos and Veronesi’s (SV 2004) conditional CAPM, in which the conditioning variable is the labor income-to-consumption ratio s^0 ; (iv) Li, Vassalou, and Xing’s (LVX 2005) investment model, in which the factors are investment growth rates for households (ΔI_{HH}), non-financial corporations (ΔI_{Corp}), and the non-corporate sector (ΔI_{Ncorp}) (we consider only this version of their model); and (v) Yogo’s (2006) durable–consumption CAPM, in which the factors are the growth in durable and non-durable consumption, Δc_{Dur} and Δc_{Non} , and the market return (R_M) (we consider only his linear model; the consumption series are available on Yogo’s website). For comparison, we also report results for the simple unconditional CAPM, the unconditional consumption CAPM, and Fama and French’s (FF 1993) three-factor model (collectively called the benchmark models).

Table 1 reports cross-sectional regressions of average returns on estimated factor loadings for the eight models. The tests use quarterly excess returns (in %), from 1963 – 2004, and highlight Prescriptions 1, 5, and 6, our suggestions to expand the set of test portfolios beyond size-B/M portfolios and to report confidence intervals for the cross-sectional R^2 and Shanken T^2 (asymptotic χ^2) statistic. Specifically, we compare results using Fama and French’s 25 size-B/M portfolios alone (‘FF25’ in the table) with results for the expanded set of 55 portfolios that includes Fama and French’s 30 industry portfolios (‘FF25 + 30 ind’). Our choice of industry portfolios is based on the notion that they should provide a fair test of the models (in contrast to, say, momentum portfolios whose returns seem to be anomalous relative to any of the models). We report OLS regressions, despite our advocacy of GLS regressions, for two main reasons: (i) to enhance comparison with prior studies; and (ii) Appendix A shows that there is a close link between the GLS R^2 and the cross-sectional T^2 -statistic that do report, since both measure how far a model’s factor-mimicking portfolios are from the minimum-variance boundary. Therefore, to limit redundancy, we just report a p-value for the T^2 statistic and a confidence interval for q , the difference between the

maximum generalizable Sharpe ratio and that attainable from a model's mimicking portfolios (q is zero if the model fully explains the cross section of expected returns).

The confidence intervals for q and the cross-sectional R^2 are obtained using the methods described in the previous section. For R^2 , we simulate the distribution of the sample adjusted R^2 for true R^2 's between 0.0 to 1.0 and invert plots like Figure 5; the simulations are similar to those in Figures 2 and 5, with the actual factors for each model now used in place of the artificial factors.⁸ We also use simulations to get a confidence interval for q , rather than rely on asymptotic theory, because the length of the time series in our tests (168 quarters) is small relative to the number of test assets (25 or 55). The confidence interval for q is based on the T^2 -statistic since, as explained above, q determines the noncentrality parameter of T^2 's (asymptotic) distribution. Thus, we simulate the distribution of the T^2 -statistic for various values of q and invert a plot like Figure 6, with q playing the same role as θ_z^2 in the GRS F-test. The p-value we report for the T^2 -statistic also comes from these simulations, with $q = 0$.

Table 1 shows four key results. First, adding industry portfolios dramatically changes the performance of the models, in terms of slope estimates, cross-sectional R^2 's, and T^2 statistics. Compared with regressions using only size-B/M portfolios, the slope estimates are almost always closer to zero and the cross-sectional R^2 's often drop substantially. The adj. R^2 drops from 58% to 0% for LL's model, from 57% to 9% for LVN's model, from 41% to 3% for SV's model, from 80% to 42% for LVX's model, and from 18% to 3% for Yogo's model. In addition, for these five models, the T^2 statistics are insignificant in tests with size-B/M portfolios but reject, or nearly reject, the models using the expanded set of 55 portfolios. The performance of FF's three-factor model is similar to the other five – it has an R^2 of 78% for the size-B/M portfolios and 31% for all 55 portfolios – while the simple and consumption CAPMs have small adj. R^2 's for both sets of test assets.

The second key result is that the cross-sectional R^2 is often very uninformative about a model's true (population) performance. Our simulations show that, across the five main models in Table 1, a 95% confidence interval for the true R^2 has an average width of 0.69, using either the size-B/M portfolios or the expanded set of 55 portfolios. For regressions with size-B/M portfolios, we cannot reject that all

⁸ The only other difference is that, to simulate data for different true cross-sectional R^2 's, we keep the true factor loadings the same in all simulations, equal to the historical estimates, and change the vector of true expected returns to give the right R^2 . Specifically, expected returns in the simulations equal $\mu = h(C\lambda) + \varepsilon$, where C is the estimated matrix of factor loadings for a model, λ is the estimated vector of cross-sectional slopes, h is a scalar constant, and ε is a random drawn from a $MVN[0, \sigma_\varepsilon^2 I]$ distribution; h and σ_ε are chosen to give the right cross-sectional R^2 and to maintain the historical cross-sectional dispersion in expected returns.

models work perfectly, as expected, but neither can we reject that the true R^2 s are quite small, with an average lower bound for the confidence intervals of 0.31. (Li, Vassalou, and Xing's model is an outlier, with a lower bound of 0.75.) For regressions with all 55 portfolios, four of the five confidence intervals include 0.00 and the fifth includes 0.20 – that is, using just the sample R^2 , we can't reject that the models have essentially no explanatory power. One of the confidence intervals covers the entire range of R^2 s from 0.00 to 1.00. The table suggests that, as a general rule, sampling variation in the R^2 is just too large to use it as a reliable metric of performance.

The third key result is that none of the models provides much improvement over the simple or consumption CAPM when performance is measured by q , the distance a model's mimicking portfolios are from the minimum-variance boundary. (By implication, none shows much improvement when performance is measured by the GLS R^2 , with an average GLS R^2 of only 0.10 for the five models using size-B/M portfolios and 0.04 using the full set of 55 portfolios.) This is true even for tests with size-B/M portfolios, for which OLS R^2 s (point estimates) are quite high, and is consistent with our view that q provides a more rigorous hurdle than the OLS R^2 . The distance q can be interpreted as the maximum generalized squared Sharpe ratio (quarterly, defined relative to the optimal zero-beta rate) on a portfolio that is uncorrelated with the factors, which is zero if the model is well-specified. For the size-B/M portfolios, the sample q is 0.46 for the simple and consumption CAPMs, compared with 0.44 for LL's model, 0.45 for LVN's model, 0.43 for SV's model, 0.34 for LVX's model, and 0.46 for Yogo's model. Adding the 30 industry portfolios, the simple and consumption CAPM q 's are both 1.34, compared with 1.31, 1.32, 1.28, 1.27, and 1.24 for the other models. Just as important, confidence intervals for the true q are generally quite wide, so even when we can't reject that q is zero, we also cannot reject that q is very large. Again, this is true even for the size-B/M portfolios, for which the models seem to perform well if we narrowly focus on the T^2 -statistic's p-value under the null.

Finally, in the spirit of taking seriously the cross-sectional parameters (Prescription 2), the table shows that none of the models explains the level of expected returns: the estimated intercepts are all substantially greater than zero for tests with either the size-B/M portfolios or the expanded set of 55 portfolios. The regressions use excess quarterly returns (in %), so the intercepts can be interpreted as the estimated quarterly zero-beta rates over and above the riskfree rate. Annualized, the zero-beta rates range from 7.8% to 14.3% *above the riskfree rate*. These estimates cannot reasonably be attributed to differences in lending vs. borrowing costs.

In sum, despite the seemingly impressive ability of the models to explain the cross section of average

returns on size-B/M portfolios, none of the models performs very well once we expand the set of test portfolios, consider confidence intervals for the true R^2 and cross-section T^2 statistic, or ask the models to price the riskfree asset.

5. Conclusion

The main point of the paper is easily summarized: Asset-pricing models should not be judged by their success in explaining average returns on size-B/M portfolios (or, more generally, for portfolios in which a couple of factors are known to explain most of the time-series and cross-sectional variation in returns). High cross-sectional explanatory power for size-B/M portfolios, in terms of high R^2 or small pricing errors, is simply not a sufficiently high hurdle by which to evaluate a model. In addition, the sample cross-sectional R^2 , as well as more formal test statistics based on the weighted sum of squared pricing errors, seems to be uninformative about the true (population) performance of a model, at least in our tests with size, B/M, and industry portfolios.

The problems we highlight are not just sampling issues, i.e., they are not solved by getting standard errors right (but sampling issues do make them worse). In population, if returns have a covariance structure like that of size-B/M portfolios, true expected returns will line up with true factor loadings so long as a proposed factor is correlated with returns only through the variation captured by the two or three common components in returns. The problems are also not solved by using an SDF approach, since SDF tests are very similar to traditional cross-sectional regressions.

The paper offers four key suggestions for improving empirical tests. First, since the problems are tied to the strong covariance structure of size-B/M portfolios, one simple suggestion is to expand the set of test assets to include portfolios sorted in other ways, for example, by industry or factor loadings. Second, since the problems are exacerbated by the fact that empirical tests often ignore theoretical restrictions on the cross-sectional intercept and slopes, another suggestion is to take their magnitudes seriously when theory provides appropriate guidance. Third, since the problems we discuss appear to be less severe for GLS regressions, another suggestion is to report the GLS R^2 in addition to, or instead of, the OLS R^2 . Last, since the problems are exacerbated by sampling issues, our fourth suggestion is to report confidence intervals for cross-sectional R^2 s and other test statistics using the techniques described in the paper. Together, these prescriptions should help to improve the power and rigor of empirical tests, though they clearly don't provide a perfect solution.

The paper also contributes to the cross-sectional asset-pricing literature in a number of additional ways: (i) we provide a novel interpretation of the GLS R^2 in terms of the relative mean-variance efficiency of factor-mimicking portfolios, building on the work of Kandel and Stambaugh (1995); (ii) we show that the cross-sectional T^2 statistic based on OLS regressions is equivalent to that from GLS regressions (identical in sample except for the Shanken-correction terms), and we show that both are a transformation of the GLS R^2 ; (iii) we derive the asymptotic properties of the cross-sectional T^2 statistic under both the null and alternative hypotheses, offering an economic interpretation of the non-centrality parameter; and (iv) we describe a way to obtain confidence intervals for the GRS F-statistic, cross-sectional T^2 statistic, and Hansen-Jagannathan distance, in addition to confidence intervals for the cross-sectional R^2 . These results are helpful for understanding cross-sectional asset-pricing tests.

Appendix A

This appendix derives the asymptotic distribution of the cross-sectional T^2 -statistic, under the null and alternatives, and provides an economic interpretation of the non-centrality parameter.

Let R_t be the $N \times 1$ vector of excess returns and F_t be the $K \times 1$ vector of factors in period t . Both are assumed, in this appendix, to be IID over time. The matrix of factor loadings is estimated in the first-pass time-series regression, $R_t = c + B F_t + e_t$, and the relation between expected returns and B is estimated in the second-pass cross-sectional regression, $\mu = z \iota + B \gamma + \alpha = X \lambda + \alpha$, where $\mu \equiv E[R_t]$, $\lambda' \equiv [z \ \gamma']$, $X \equiv [\iota \ B]$, and α is the vector of the true pricing errors. More precisely, the parameters in the cross-sectional equation depend on whether we are interested in an OLS or GLS regression: for OLS, the population slope is $\lambda = (X'X)^{-1}X'\mu$ and the pricing errors are $\alpha \equiv [I - X(X'X)^{-1}X'] \mu \equiv y \mu$; for GLS, the slope is $\lambda^* = (X'V^{-1}X)^{-1}X'V^{-1}\mu$ and the pricing errors are $\alpha^* \equiv [I - X(X'V^{-1}X)^{-1}X'V^{-1}] \mu \equiv y^* \mu$, where $V \equiv \text{var}(R_t)$. In practice, of course, the cross-sectional regression is estimated with average returns substituted for μ and estimates of B substituted for the true loadings.

We begin with a few population results that will be useful for interpreting empirical tests. We omit the time subscript until we turn to sample statistics.

Result 1. The cross-sectional slope and pricing errors in a GLS regression are identical if V is replaced by $\Sigma \equiv \text{var}(e)$. Thus, we will use V and Σ interchangeably in the GLS results below depending on which is more convenient for the issue at hand.

Proof: See Shanken (1985). The result follows from $V = B \Sigma_F B' + \Sigma$ and $V^{-1} = \Sigma^{-1} - \Sigma^{-1} B [\Sigma_F^{-1} + B' \Sigma^{-1} B]^{-1} B' \Sigma^{-1}$, which, along with the definition of X , imply that $(X'V^{-1}X)^{-1}X'V^{-1} = (X'\Sigma^{-1}X)^{-1}X'\Sigma^{-1}$. \square

Recall that $\alpha = y \mu$ and that $\alpha^* = y^* \mu$. The quadratics $q = \alpha' [y \Sigma y]^{+} \alpha$ and $q^* = \alpha^{*'} [y^* \Sigma y^{*'}]^{+} \alpha^*$, where a superscript '+' denotes a pseudoinverse, will be important for interpreting the cross-sectional T^2 test. The analysis below uses the facts, easily confirmed, that y and $\Sigma^{-1}y^*$ are symmetric, y and y^* are idempotent ($y = yy$ and $y^* = y^*y^*$), $y = yy^*$, $y^* = y^*y$, and $yX = y^*X = 0$.

Result 2. The quadratics q and q^* are unchanged if Σ is replaced by V . Together with Result 1, this result will imply that the T^2 statistics, from either OLS or GLS, are the same regardless of which covariance matrix we use.

Proof: $yX = y^*X = 0$ implies that $yB = y^*B = 0$, since $X = [\iota \ B]$. Thus, $yV = yB \Sigma_F B' + y \Sigma = y \Sigma$ and $y^*V = y^*B \Sigma_F B' + y^* \Sigma = y^* \Sigma$. The result follows immediately. \square

Result 3. The OLS and GLS quadratics are identical, i.e., $q = q^*$.

Proof: The quadratics are defined as $q = \alpha' [y \Sigma y]^+ \alpha$ and $q^* = \alpha^{*'} [y^* \Sigma y^{*'}]^+ \alpha^*$. Using the definition of a pseudoinverse, it is easy to confirm that $[y \Sigma y]^+ = \Sigma^{-1} y^*$, implying that $q = \alpha' \Sigma^{-1} y^* \alpha$. Further, $[y^* \Sigma y^{*'}]^+ = \Sigma^{-1} (y^*)^+$, implying that $q^* = \alpha^{*'} \Sigma^{-1} (y^*)^+ \alpha^*$.⁹ Moreover, $\alpha^* = y^* \mu = y^* \alpha$, from which it follows that $q^* = \alpha^{*'} y^{*'} \Sigma^{-1} (y^*)^+ y^* \alpha = \alpha' (\Sigma^{-1} y^*)' (y^*)^+ y^* \alpha = \alpha' \Sigma^{-1} y^* \alpha = q$, where the second-to-last equality uses the fact that $\Sigma^{-1} y^*$ is symmetric and that $y^* (y^*)^+ y^* = y^*$. \square

Result 4. The OLS and GLS quadratics equal $q^* = \alpha^{*'} \Sigma^{-1} \alpha^* = \alpha^{*'} V^{-1} \alpha^*$. (This result implies that our cross-sectional T^2 -statistic matches that of Shanken, 1985.)

Proof: Result 3 shows that $q^* = \alpha' \Sigma^{-1} y^* \alpha$. Recall that y^* is idempotent, $\Sigma^{-1} y^*$ is symmetric, and $\alpha^* = y^* \alpha$. Therefore, $q^* = \alpha' \Sigma^{-1} y^* y^* \alpha = \alpha' y^{*'} \Sigma^{-1} y^* \alpha = \alpha^{*'} \Sigma^{-1} \alpha^*$. \square

Mimicking portfolios, R_p , for the factors are defined as the K portfolios maximally correlated with F . The $N \times K$ matrix of mimicking-portfolio weights are slopes in the regression $F = k + w_p' R + s$, where $\text{cov}(R, s) = 0$ (we ignore the constraint that $w_p' \iota = \iota$ for simplicity; the weights can be scaled up or down to make the constraint hold without changing the substance of any results). Thus, $w_p = V^{-1} \text{cov}(R, F) = V^{-1} B \Sigma_F$ and stocks' loadings on the mimicking portfolios are $C = \text{cov}(R, R_p) \Sigma_p^{-1} = V w_p \Sigma_p^{-1} = B \Sigma_F \Sigma_p^{-1}$.

Result 5. The cross-sectional regression (OLS or GLS) of μ on B is equivalent to the cross-sectional regression of μ on C , with or without an intercept, in the sense that the intercept, R^2 , pricing errors, and quadratics q and q^* are the same in both.

Proof: The first three claims, that the intercept, R^2 , and pricing errors are the same, follow directly from the fact that C is a nonsingular transformation of B . The final claim, that the quadratics are the same regardless of whether we use F or R_p , follows from the fact that the pricing errors are the same and the quadratics can be based on V , i.e., $q^* = \alpha^{*'} V^{-1} \alpha^*$, where V is invariant to the set of factors. \square

⁹ More precisely, $\Sigma^{-1} (y^*)^+$ is a generalized inverse of $y^* \Sigma y^{*'}$, though not necessarily the pseudoinverse (the pseudoinverse of A is such that $AA^+A = A$, $A^+AA^+ = A^+$, and A^+A and AA^+ are symmetric; a generalized inverse ignores the two symmetry conditions). It can be shown that the use of $\Sigma^{-1} (y^*)^+$ in the quadratic q^* is equivalent to using the pseudoinverse, which may or may not be different (we don't know).

Result 6. The GLS regression of μ on B or μ on C, with or without an intercept, prices the mimicking portfolios perfectly, i.e., $\alpha_p^* = w_p' \alpha^* = 0$. It follows that the slopes on C equal $\mu_p - z^* \iota$, the expected return on the mimicking portfolio in excess of the GLS zero-beta rate (for this last result, we assume that w_p is scaled to make $w_p' \iota = 1$).

Proof: From the discussion prior to Result 5, $w_p = V^{-1} B \Sigma_F$, implying that $\alpha_p^* = w_p' \alpha^* = \Sigma_F B' V^{-1} \alpha^* = \Sigma_F B' V^{-1} y^* \mu$. Further, $X' V^{-1} y^* = 0$, from which it follows that $B' V^{-1} y^* = 0$ and, hence, $\alpha_p^* = 0$. This proves the first half of the result. Also, by definition, $\alpha^* = \mu - z^* \iota - C \gamma_p^*$, where γ_p^* are the GLS slopes on C. Therefore, $\alpha_p^* = w_p' \alpha^* = \mu_p - z^* \iota - \gamma_p^* = 0$, where $\mu_p = w_p' \mu$ and $C_p = w_p' C = I_K$. Solving for γ_p^* proves the second half of the result. \square

Result 7. Pricing errors in a GLS cross-sectional regression, of μ on either B or C, are identical to the intercepts in a time-series regression of $R - z^* \iota$ on a constant and $R_p - z^* \iota$. It follows that the quadratics q and q^* equal $\theta_\tau^2(z^*) - \theta_p^2(z^*)$, where $\theta_i(z^*)$ is a generalized Sharpe ratio with respect to $r_f + z^*$, defined as $(\mu_i - z^*) / \sigma_i$, and σ_i is asset i 's standard deviation, τ is the 'tangency' portfolio with respect to $r_f + z^*$, and θ_p is the maximum squared generalized Sharpe ratio attainable from R_p .

Proof: Intercepts in the time-series regression are $\alpha_{TS} = \mu - z^* \iota - C (\mu_p - z^* \iota)$. From Result 6, these equal α^* since $\gamma_p^* = \mu_p - z^* \iota$. The interpretation of the quadratics then follows immediately from the well-known interpretation of $\alpha_{TS}' \Sigma^{-1} \alpha_{TS}$ (Jobson and Korkie, 1982; Gibbons, Ross, Shanken, 1989), with the only change that the Sharpe ratios need to be defined relative to $r_f + z^*$. \square

Result 8. The GLS R^2 equals $1 - q / Q = 1 - q^* / Q$, where $Q = (\mu - \mu_{gmv} \iota)' V^{-1} (\mu - \mu_{gmv} \iota)$ and μ_{gmv} is the expected return on the global minimum variance portfolio (note that Q depends only on the characteristics of asset returns, not the factors being tested). Further, the GLS R^2 is zero if and only if the factors' mimicking portfolios all have expected returns equal to μ_{gmv} (i.e., they lie exactly in the middle of mean-variance space), and GLS R^2 is one if and only if some combination of the mimicking portfolios lies on the minimum-variance boundary.

Proof: The first claim follows directly from the definition of the GLS R^2 , i.e., $GLS R^2 = 1 - \alpha^{*'} V^{-1} \alpha^* / (\mu - z_{nf} \iota)' V^{-1} (\mu - z_{nf} \iota)$, where z_{nf} is the GLS intercept when μ is regressed only a constant. z_{nf} is the same as μ_{gmv} and $\alpha^{*'} V^{-1} \alpha^*$ is the same as q^* (see Result 4). The second claim, which we state without further proof, is a multifactor generalization of the results of Kandel and Stambaugh (1995) with mimicking portfolios substituted for non-return factors (see Result 5). The key fact is that $q^* = Q - (\mu_p -$

$\mu_{\text{gmV}} \mathbf{1}' \Sigma_P^{*-1} (\mu - \mu_{\text{gmV}} \mathbf{1})$, where Σ_P^* is the residual covariance matrix when R_P is regressed on the R_{gmV} . Thus, q^* is zero (the GLS R^2 is one) only if some combination of R_P lies on the minimum-variance boundary, and q^* equals Q (the GLS R^2 is zero) only if $\mu_P = \mu_{\text{gmV}} \mathbf{1}$. \square

Together, Results 1 – 8 describe key properties of GLS cross-sectional regressions, establish the equality between the OLS and GLS quadratics q and q^* , and establish the connections among the location of R_P in mean-variance space, the GLS R^2 , and the quadratics. All of the results have exact parallels in sample, redefining moments as sample statistics rather than population parameters.

Our final results consider the asymptotic properties of the cross-sectional T^2 statistic under the null that pricing errors are zero, $\alpha = 0$ and $\alpha^* = 0$, and generic alternatives that they are not. The T^2 statistic is, roughly speaking, the sample analog of the quadratics q and q^* based on the traditional two-pass methodology (defined precisely below). Let r be average returns, b be the sample (first-pass, time-series regression) estimate of B , $x = [\mathbf{1} \ b]$ be the corresponding estimate of X , v and S be the usual estimates of V and Σ , $\hat{y} = I - x(x'x)^{-1}x'$ be the sample estimate of y , and $\hat{y}^* = I - x(x'v^{-1}x)^{-1}x'v^{-1}$ be the sample estimate of y^* . The estimated OLS cross-sectional regression is $r = x \hat{\lambda} + \hat{a}$, where $\hat{\lambda} = (x'x)^{-1}x'r$ and $\hat{a} = \hat{y} r$ is the sample OLS pricing error. The estimated GLS regression is $r = x \hat{\lambda}^* + \hat{a}^*$, where $\hat{\lambda}^* = (x'v^{-1}x)^{-1}x'v^{-1}r$ and $\hat{a}^* = \hat{y}^* r$. Equivalently, since $r = (1/T) \sum_t R_t$, the estimated slope and pricing errors can be interpreted as time-series averages of period-by-period Fama-MacBeth cross-sectional estimates. We will focus on OLS regressions but, as a consequence of the sample analog of Result 3 above, we show that the T^2 statistic is equivalent from OLS and GLS.

Our analysis below uses the following facts:

- (1) $R_t = \mu + B UF_t + e_t$, where $UF_t = F_t - \mu_F$.
- (2) $\mu = X \lambda + \alpha = x \lambda + (X - x) \lambda + \alpha = x \lambda + (B - b) \gamma + \alpha$.
- (3) $B UF_t = b UF_t + (B - b) UF_t$.
- (4) $\hat{y} x = 0$ and $\hat{y} b = 0$

Combining these facts, the pricing error in period t is $\hat{a}_t \equiv \hat{y} R_t = \hat{y} (B - b) \gamma + \hat{y} \alpha + \hat{y} (B - b) UF_t + \hat{y} e_t$ and the time-series average is $\hat{a} = \hat{y} (B - b) \gamma + \hat{y} \alpha + \hat{y} (B - b) \overline{UF} + \hat{y} \bar{e}_t$, where an upper bar denotes a time-series average. Asymptotically, $b \rightarrow B$, $\hat{y} \rightarrow y$, and \overline{UF} and \bar{e}_t both go to zero. These observations, together with $y \alpha = \alpha$, imply that \hat{a} is a consistent estimator of α . Also, the second-to-last term, $\hat{y} (B - b) \overline{UF}$, converges to zero at a faster rate than the other terms and so, for our purposes, can be dropped: $\hat{a} = \hat{y} (B - b) \gamma + \hat{y} \alpha + \hat{y} \bar{e}_t$.

Result 9. Define $d \equiv \hat{a} - \hat{y} \alpha$. Asymptotically, $T^{1/2} d$ converges in distribution to $N(0, T \Sigma_d)$, where $\Sigma_d = y \Sigma y (1 + \gamma' \Sigma_F^{-1} \gamma) / T$.

Proof: The asymptotic mean is zero since $\hat{a} - \hat{y} \alpha \rightarrow \alpha - y \alpha = 0$. The asymptotic covariance follows from observing that d is the same as \hat{a} under the null that $\alpha = 0$, the scenario considered by Shanken (1985, 1992b), and the term $(1 + \gamma' \Sigma_F^{-1} \gamma)$ is just the Shanken correction for estimation error in b . More precisely, $d = \hat{y} (B - b) \gamma + \hat{y} \bar{e}_i$. The asymptotic distribution is the same substituting y for \hat{y} , and the two terms are uncorrelated with each other under the standard assumptions of OLS regressions (i.e., in a regression, estimation error in the slopes is uncorrelated with the mean of the residuals). Therefore, $\Sigma_d = \text{var}[y (B - b) \gamma] + \text{var}[y \bar{e}_i]$. Let $\text{vec}(b - B)$ be the $NK \times 1$ vector version of $b - B$, stacking the loadings for asset 1, then asset 2, etc, which has asymptotic variance $\Sigma \otimes \Sigma_F^{-1} / T$ from standard regression results. Rearranging and simplifying, the first term is $\text{var}[y(B - b)\gamma] = \gamma' \Sigma_F^{-1} \gamma (y \Sigma y) / T$ and the second term is $\text{var}[y \bar{e}_i] = y \Sigma y / T$. Summing these gives the covariance matrix. \square

A corollary of Result 9 is that, under the null that $\alpha = 0$, $T^{1/2} \hat{a}$ also converges in distribution to $N(0, T \Sigma_d)$. To test whether $\alpha = 0$, the cross-sectional T^2 statistic is then naturally defined as $T^2 = \hat{a}' S_d^+ \hat{a}$, where S_d is the sample estimate of Σ_d substituting the statistics \hat{y} , S , $\hat{\gamma}$, and S_F for the population parameters y , Σ , γ , and Σ_F . Thus, $T^2 = \hat{a}' [\hat{y} S \hat{y}]^+ \hat{a} [T / (1 + \hat{\gamma}' S_F^{-1} \hat{\gamma})]$. The key quadratic here, $\hat{q} = \hat{a}' [\hat{y} S \hat{y}]^+ \hat{a}$, is the sample counterpart of q defined earlier. Result 3 implies that this OLS-based T^2 statistic is identical to a GLS-based T^2 statistic defined using \hat{q}^* , the sample equivalent of the GLS quadratic q^* (the T^2 statistics are identical assuming the same Shanken correction term, $\hat{\gamma}' S_F^{-1} \hat{\gamma}$, is used for both; they are asymptotically equivalent under the null as long as consistent estimates of γ and Σ_F are used for both). Moreover, Result 4 implies that $T^2 = \hat{a}' S^{-1} \hat{a}^* [T / (1 + \hat{\gamma}' S_F^{-1} \hat{\gamma})]$.

Result 10. The cross-sectional T^2 statistic is asymptotically χ^2 with degrees for freedom $N - K - 1$ and non-centrality parameter $n = k q^*$, where $k = T / (1 + \gamma' \Sigma_F^{-1} \gamma)$.¹⁰ Equivalently, from Result 7, the non-centrality parameter is $n = k [\theta_\tau^2(z^*) - \theta_P^2(z^*)]$.

Proof, part 1 (distribution under the null): From Result 9, if $\alpha = 0$, \hat{a} is the same as d and $T^{1/2} \hat{a}$ converges

¹⁰ We use the terminology of a limiting distribution somewhat informally here (notice that, as the result is stated, the noncentrality parameter goes to infinity as T get large unless α and q^* are zero). The asymptotic result can be stated more formally by considering pricing errors that go to zero as T gets large: Suppose that $\alpha^* = T^{-1/2} \delta^*$, for some fixed vector δ^* . For this sequence of α^* , the T^2 statistic converges in distribution to a χ^2 with noncentrality parameter $kq^* = \delta^{*'} \Sigma^{-1} \delta^* / (1 + \gamma' \Sigma_F^{-1} \gamma)$, where $k = T / (1 + \gamma' \Sigma_F^{-1} \gamma)$ and $q^* = \delta^{*'} \Sigma^{-1} \delta^* / T$.

in distribution to $N(0, T \Sigma_d)$. Σ_d is nonsingular with rank $N - K - 1$, so $\hat{a}' \Sigma_d^+ \hat{a}$ is asymptotically χ^2 with degrees of freedom $N - K - 1$. Further, S_d^+ is a consistent estimate of Σ_d^+ , implying that $\hat{a}' S_d^+ \hat{a}$ converges to the same distribution. [S_d^+ converges to Σ_d^+ since $S_d^+ = S^{-1} \hat{y}^* T / (1 + \hat{\gamma}' S_F^{-1} \hat{\gamma})$, from Result 3, which clearly converges to $\Sigma_d^+ = \Sigma^{-1} y^* T / (1 + \gamma' \Sigma_F^{-1} \gamma)$].

Proof, part 2 (distribution under alternatives): In general, $\hat{a} = d + \hat{y} \alpha = \hat{y} (d + \alpha)$, where the first equality follows from the definition of d and the second follows from the fact that \hat{y} is idempotent and $\hat{y} d = d$. Following the proof of Result 3, it is straightforward to show that $\hat{y} [\hat{y} S \hat{y}]^+ \hat{y} = [\hat{y} S \hat{y}]^+$, which implies that $\hat{q} = \hat{a}' [\hat{y} S \hat{y}]^+ \hat{a} = (d + \alpha)' [\hat{y} S \hat{y}]^+ (d + \alpha)$. The matrix $[\hat{y} S \hat{y}]^+$ converges to $[y \Sigma y]^+$ and the T^2 statistic is $T^2 = \hat{q} T / (1 + \hat{\gamma}' S_F^{-1} \hat{\gamma})$, which together imply that T^2 has the same asymptotic distribution as $(d + \alpha)' \Sigma_d^+ (d + \alpha)$, where $\Sigma_d = [y \Sigma y]^+ (1 + \gamma' \Sigma_F^{-1} \gamma) / T$. Recall that $T^{1/2} d$ converges in distribution to $N(0, T \Sigma_d)$. Therefore, $d' \Sigma_d^+ d$ is asymptotically central χ^2 and $(d + \alpha)' \Sigma_d^+ (d + \alpha)$ is noncentral χ^2 with noncentrality parameter $n = \alpha' \Sigma_d^+ \alpha$ (both with degrees of freedom $N - K - 1$). Result 10 then follows from observing that the noncentrality parameter can be rewritten as $n = \alpha' [y \Sigma y]^+ \alpha T / (1 + \gamma' \Sigma_F^{-1} \gamma) = q^*$ k , where $k = T / (1 + \gamma' \Sigma_F^{-1} \gamma)$. \square

Appendix B

This appendix derives the small-sample distribution of the HJ-distance when returns are multivariate normal and the factors in the proposed model are portfolio returns (or have been replaced by maximally correlated mimicking portfolios). R is defined, for the purposes of this appendix, to be the $N+1$ vector of total rates of return on the test assets, including the riskless asset.

Let $y = a + b R_p$. The HJ-distance is defined as $D = \min_m E[(m - y)^2]$, where m represents any well-specified SDF, i.e., any variable for which $E[m(1+R)] = 1$. Hansen and Jagannathan (1997) show that, if y is linear in asset returns (or is the projection of a non-return y onto the space of asset returns), then the m^* which solves the minimization problem is linear in the return on the tangency portfolio, i.e., $m^* = v_0 + v_1 R_\tau$ for some constants v_0 and v_1 , and $D = E[(m^* - y)^2]$.

The constants a and b are generally unknown and chosen to minimize D . Therefore, a and b solve $\min_{a,b} [E(m^* - a - b R_p)^2]$. This problem is simply a standard least-squares projection problem, so D turns out to be nothing more than the residual variance when m^* is regressed on a constant and R_p . Equivalently, D is v_1^2 times the residual variance when R_τ is regressed on a constant and R_p : $D = v_1^2 \text{var}(\omega)$, where ω is

from the regression $R_\tau = a' + b' R_p + \omega$. Kandel and Stambaugh (1987) and Shanken (1987) show that the correlation between any portfolio and the tangency portfolio equals the ratio of their Sharpe measures, $\text{cor}(R_x, R_\tau) = \theta_x/\theta_\tau$. Thus, b' gives the combination of R_p that has the maximum squared Sharpe ratio, denoted θ_p^2 , from which it follows that $\text{var}(\omega) = (1 - \theta_p^2 / \theta_\tau^2) \sigma_\tau^2$. Cochrane (2001) shows that $v_1 = -\mu_\tau / [\sigma_\tau^2 (1+r_f)]$, implying that the HJ distance is $D = v_1^2 \text{var}(\omega) = (\theta_\tau^2 - \theta_p^2) / (1 + r_f)^2$, where $\theta_\tau^2 - \theta_p^2$ can be interpreted as the proposed model's unexplained squared Sharpe ratio.

The analysis above is cast in terms of population parameters, but equivalent results go through in sample, re-defining all quantities as sample moments. Thus, the *estimated* HJ-distance, d , is proportional to the difference between the sample squared Sharpe ratios of the ex post tangency portfolio and the portfolios in R_p . Following the discussion in Section 4, the sample HJ-distance is therefore proportional to the GRS F-statistic: $d = F c / (1 + r_f)^2 [N (T - K - 1) / (T - N - K)]$, where $c = (1 + s_p^2) / T$ and s_p^2 is the sample counterpart to θ_p^2 . It follows immediately that, up to a constant of proportionality, d is non-central F with non-centrality parameter $\lambda = c^{-1} a' \Sigma^{-1} a = c^{-1} (1 + r_f)^2 D$.

References

- Acharya, Viral and Lasse Pedersen, 2005. Asset pricing with liquidity risk. *Journal of Financial Economics* 77 (2), 375-410.
- Black, Fisher, 1972. Capital market equilibrium with restricted borrowing. *Journal of Business* 45, 444-454.
- Black, Fisher, Michael C. Jensen, and Myron Scholes, 1972. The capital asset pricing model: Some empirical findings, in: *Studies in the theory of capital markets*, ed. by Michael C. Jensen. New York: Praeger.
- Brennan, Michael, 1971. Capital market equilibrium with divergent borrowing and lending rates. *Journal of Financial and Quantitative Analysis* 6, 1197-1205.
- Brennan, Michael, Ashley Wang, and Yihong Xia, 2004. Estimation and test of a simple model of intertemporal asset pricing. *Journal of Finance* 59, 1743-1775.
- Campbell, John and Tuomo Vuolteenaho, 2004. Bad beta, good beta. *American Economic Review* 94, 1249-1275.
- Cochrane, John, 1996. A cross-sectional test of an investment-based asset pricing model, *Journal of Political Economy* 104, 572-621.
- Cochrane, John, 2001. *Asset Pricing*. Princeton University Press, Princeton, NJ.
- Daniel, Kent and Sheridan Titman, 2005. Testing factor-model explanations of market anomalies. Working paper (Northwestern University, Evanston, IL).
- Fama, Eugene and Kenneth French, 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3-56.
- Gibbons, Michael, Stephen Ross, and Jay Shanken, 1989. A test of the efficiency of a given portfolio. *Econometrica* 57, 1121-1152.
- Hansen, Lars, 1982. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50, 1029-1054.
- Hansen, Lars, John Heaton, and Nan Li, 2005. Consumption strikes back? Measuring long-run risk. Working paper (University of Chicago, Chicago, IL).
- Hansen, Lars and Ravi Jagannathan, 1997. Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance* 52, 557-590.
- Heaton, John and Deborah Lucas, 2000. Asset pricing and portfolio choice: The role of entrepreneurial risk. *Journal of Finance* 55 (3), 1163-1198.
- Jagannathan, Ravi and Zhenyu Wang, 1996. The conditional CAPM and the cross-section of stock returns. *Journal of Finance* 51, 3-53.
- Jobson, J. D. and Bob Korkie, 1982. Potential performance and tests of portfolio efficiency. *Journal of*

- Financial Economics* 10, 433-466.
- Kan, Raymond and Chou Zhang, 1999. Two-pass tests of asset pricing models with useless factors. *Journal of Finance* 54, 203-235.
- Kan, Raymond and Guofu Zhou, 2004. Hansen-Jagannathan distance: Geometry and exact distribution. Working paper (University of Toronto and Washington University in St. Louis).
- Kandel, Shmuel and Robert Stambaugh, 1987. On correlations and inferences about mean-variance efficiency. *Journal of Financial Economics* 18, 61-90.
- Kandel, Shmuel and Robert Stambaugh, 1995. Portfolio inefficiency and the cross-section of expected returns. *Journal of Finance* 50, 157-184.
- Kimmel, Robert, 2003. Risk premia in linear factor models: Theoretical and econometric issues. Working paper (Princeton University, Princeton, NJ).
- Kullman, Cornelia, 2003. Real estate and its role in asset pricing. Working paper (University of British Columbia, Vancouver, BC).
- Lettau, Martin and Sydney Ludvigson, 2001. Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy* 109, 1238-1287.
- Lewellen, Jonathan and Stefan Nagel, 2006. The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics*, forthcoming.
- Li, Qing, Maria Vassalou, and Yuhang Xing, 2005. Sector investment growth rates and the cross-section of equity returns. *Journal of Business*, forthcoming.
- Lustig, Hanno and Stijn Van Nieuwerburgh, 2004. Housing collateral, consumption insurance, and risk premia. *Journal of Finance* 60 (3), 1167-1221.
- Mehra, Rajnish and Edward Prescott, 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15, 145-161.
- Muirhead, Robb, 1982. *Aspects of Multivariate Statistical Theory*. Wiley, New York, NY.
- Nawalkha, Sanjay, 1997. A multibeta representation theorem for linear asset pricing theories, *Journal of Financial Economics* 46, 357-381.
- Pastor, Lubos and Robert Stambaugh, 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111, 642-685.
- Parker, Jonathan and Christian Julliard, 2005. Consumption risk and the cross section of expected returns. *Journal of Political Economy* 113 (1), 185-222.
- Petkova, Ralitsa, 2006. Do the Fama-French factors proxy for innovations in predictive variables? *Journal of Finance* 61, 581-612.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2006. Housing, consumption, and asset pricing. Forthcoming in *Journal of Financial Economics*.

- Reisman, Haim, 1992. Reference variables, factor structure, and the approximate multibeta representation. *Journal of Finance* 47, 1303-1314.
- Roll, Richard and Stephen Ross, 1994. On the cross-sectional relation between expected returns and betas. *Journal of Finance* 49, 101-121.
- Santos, Tano and Pietro Veronesi, 2006. Labor income and predictable stock returns. *Review of Financial Studies* 19, 1-44.
- Shanken, Jay, 1985. Multivariate tests of the zero-beta CAPM. *Journal of Financial Economics* 14, 327-348.
- Shanken, Jay, 1987. Multivariate proxies and asset pricing relations: Living with the Roll critique. *Journal of Financial Economics* 18, 91-110.
- Shanken, Jay, 1990. Intertemporal asset pricing: An empirical investigation. *Journal of Econometrics* 45, 99-120.
- Shanken, Jay, 1992a. On the current state of the arbitrage pricing theory, *Journal of Finance* 47, 1569-1574.
- Shanken, Jay, 1992b. On the estimation of beta-pricing models. *Review of Financial Studies* 5, 1-34.
- Shanken, Jay and Goufu Zhou, 2006. Estimating and testing beta pricing models: Alternative methods and their performance in simulations. Working paper (Emory University and Washington University in St. Louis).
- Stock, James, 1991. Confidence intervals for the largest autoregressive root in U.S. macroeconomic time series. *Journal of Monetary Economics* 28, 435-459.
- Vassalou, Maria, 2003. News related to future GDP growth as a risk factor in equity returns. *Journal of Financial Economics* 68, 47-73.
- Yogo, Motohiro, 2006. A consumption-based explanation of expected stock returns. *Journal of Finance* 61, 539-580.

Table 1. Empirical tests of asset-pricing models, 1963 – 2004.

The table reports slopes, Shanken (1992b) t-statistics (in parentheses), and adj. R^2 s from cross-sectional regressions of average excess returns on estimated factor loadings for eight models proposed in the literature. Returns are quarterly (%). The test assets are Fama and French's 25 size-B/M portfolios used alone or together with their 30 industry portfolios. The cross-sectional T^2 (asymptotic χ^2) statistic tests whether residuals in the cross-sectional regression are all zero, as described in the text, with simulated p-values in brackets. T^2 is proportional to the distance, q , that a model's true mimicking portfolios are from the minimum-variance boundary, measured as the difference between the maximum generalized Sharpe ratio and that attainable from the mimicking portfolios; the sample estimate of q is reported in the final column. 95% confidence intervals for the true R^2 and q are reported in brackets below the sample values. The models are estimated from 1963 – 2004 except Yogo's, for which we have factor data through 2001.

Model and test assets	Variables				Adj. R^2	T^2	q
Lettau & Ludvigson	const.	cay	Δc	cay$\times\Delta c$			
FF25	3.33 (2.28)	-0.81 (-1.25)	0.25 (0.84)	0.00 (0.42)	0.58 [0.30, 1.00]	33.9 [p=0.08]	0.44 [0.00, 0.72]
FF25 + 30 industry	2.50 (3.29)	-0.48 (-1.23)	0.09 (0.53)	-0.00 (-1.10)	0.00 [0.00, 0.35]	193.8 [p=0.00]	1.31 [0.18, 1.08]
Lustig & V Nieuwerburgh	const.	my	Δc	my$\times\Delta c$			
FF25	3.58 (2.22)	4.23 (0.76)	0.02 (0.04)	0.10 (1.57)	0.57 [0.35, 1.00]	20.8 [p=0.57]	0.45 [0.00, 0.48]
FF25 + 30 industry	2.78 (3.51)	0.37 (0.13)	-0.02 (-0.09)	0.03 (1.40)	0.09 [0.00, 1.00]	157.1 [p=0.04]	1.32 [0.00, 0.96]
Santos & Veronesi	const.	R_M	$s^0 \times R_M$				
FF25	2.45 (1.39)	-0.32 (-0.17)	-0.22 (-2.04)		0.41 [0.15, 1.00]	19.7 [p=0.83]	0.43 [0.00, 0.18]
FF25 + 30 industry	2.29 (2.75)	-0.17 (-0.16)	-0.05 (-1.51)		0.03 [0.00, 0.70]	188.7 [p=0.01]	1.28 [0.06, 0.90]
Li, Vassalou, & Xing	const.	ΔI_{HH}	ΔI_{Corp}	ΔI_{Ncorp}			
FF25	2.47 (2.13)	-0.80 (-0.39)	-2.65 (-1.03)	-8.59 (-1.96)	0.80 [0.75, 1.00]	25.2 [p=0.29]	0.34 [0.00, 0.48]
FF25 + 30 industry	2.22 (3.14)	0.20 (0.19)	-0.93 (-0.58)	-5.11 (-2.32)	0.42 [0.20, 1.00]	141.2 [p=0.11]	1.27 [0.00, 0.84]
Yogo	const.	Δc_{Ndur}	Δc_{Dur}	R_M			
FF25	1.98 (1.36)	0.28 (1.00)	0.67 (2.33)	0.48 (0.29)	0.18 [0.00, 1.00]	24.9 [p=0.69]	0.46 [0.00, 0.30]
FF25 + 30 industry	1.95 (2.27)	0.18 (1.01)	0.19 (1.52)	0.12 (0.11)	0.02 [0.00, 0.60]	159.3 [p=0.06]	1.24 [0.00, 0.78]
CAPM	const.	R_M					
FF25	2.90 (3.18)	-0.44 (-0.39)			-0.03 [0.00, 0.55]	77.5 [p=0.00]	0.46 [0.12, 0.48]
FF25 + 30 industry	2.03 (2.57)	0.10 (0.09)			-0.02 [0.00, 0.35]	225.2 [p=0.00]	1.34 [0.18, 0.96]

Table 1 continues on next page (variables are defined at the end of the table)

Table 1, continued.

Model and test assets	Variables				Adj. R ²	T ²	q
Consumption CAPM	const.	Δc					
FF25	1.70 (2.47)	0.24 (0.89)			0.05 [0.00, 1.00]	60.6 [p=0.01]	0.46 [0.06, 0.66]
FF25 + 30 industry	2.07 (3.51)	0.03 (0.15)			-0.02 [0.00, 0.65]	224.5 [p=0.00]	1.34 [0.18, 1.02]
Fama & French	const.	R_M	SMB	HML			
FF25	2.99 (2.33)	-1.42 (-0.98)	0.80 (1.70)	1.44 (3.11)	0.78 [0.60, 1.00]	56.1 [p=0.00]	0.37 [0.06, 0.42]
FF25 + 30 industry	2.21 (2.14)	-0.49 (-0.41)	0.60 (1.24)	0.87 (1.80)	0.31 [0.00, 0.90]	200.4 [p=0.00]	1.24 [0.12, 0.90]

Variables:

R_M = CRSP value-weighted excess return

Δc = log consumption growth

cay = Lettau and Ludvigson's (2001) consumption-to-wealth ratio

my = Lustig and Van Nieuwerburgh's (2004) housing collateral ratio (based on mortgage data)

s^w = labor income to consumption ratio

ΔI_{HH} , ΔI_{Corp} , ΔI_{Ncorp} = log investment growth for households, non-financial corporations, and the non-corporate sector

Δc_{Ndur} , Δc_{Dur} = Yogo's (2005) log consumption growth for non-durables and durables

SMB, HML = Fama and French's (1993) size and B/M factors