NBER WORKING PAPERS SERIES

A CROSS-SECTIONAL TEST OF A PRODUCTION-BASED ASSET PRICING MODEL

John H. Cochrane

Working Paper No. 4025

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1992

I have benefitted greatly from conversations with Lars Hansen on the methods used in this paper. I am grateful to Gene Fama, Ken French and especially Wayne Ferson for many useful comments. Caleb Wong provided excellent research assistance. This research was partially supported by a grant from the National Science Foundation. This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #4025
March 1992

A CROSS-SECTIONAL TEST OF A PRODUCTION-BASED ASSET PRICING MODEL

ABSTRACT

This paper tests a factor pricing model for stock returns. The factors are returns on physical investment, inferred from investment data via a production function. The tests examine the model's ability to explain the variation in expected returns across assets and over time. The model is not rejected. It performs about as well as the CAPM and the Chen, Roll and Ross factor model, and it performs substantially better than a simple consumption-based model. In comparison tests, the investment return factors drive out all the other models.

The paper also provides an easy technique for estimating and testing dynamic, conditional asset pricing models. All one has to do is include factors and returns scaled by instruments in an unconditional estimate. This procedure imposes none of the usual restrictions on conditional moments, and does not require prewhitened or orthogonalized factors.

John H. Cochrane
Department of Economics and
Graduate School of Business
University of Chicago
1101 E. 58th Street
Chicago, IL 60637
and NBER

A Cross-Sectional Test of a Production-Based Asset Pricing Model

1. Introduction

The *investment return* is the marginal rate at which a firm can transfer resources through time, by increasing investment today and decreasing it at a future date, leaving its production plan unchanged at all other dates. This paper examines whether cross-sectional and time-series variation in asset returns can be explained by a factor pricing model with investment returns as the factors. The basic idea is to infer the presence of systematic shocks by watching firms' investment decisions, just as the consumption-based model tries to infer the presence of systematic shocks by watching consumption decisions.

This paper extends the work in Cochrane (1991a), which only tried to explain time-series variation in a single asset return. The techniques I use to estimate and test dynamic, conditional factor models are derived from the work of Hansen (1982), Hansen and Singleton (1982), Hansen and Richard (1987) and Hansen and Jagannathan (1991a). Knez (1991) and Snow (1991) use similar techniques to study factor pricing models; Braun (1991) uses them to investigate consistent pricing of asset and investment returns; DeSantis (1992) uses them to study international capital market integration. Cochrane (1992) contains a more detailed presentation of the methodology, and its connection to more traditional specification and testing of asset pricing models.

The investment return factor model is an example of a production-based asset pricing model. Section 2 explains the idea and philosophy behind production-based asset pricing models. Section 3 shows how to construct investment returns from investment data. Section 4 describes the factor model that uses investment returns as factors. Section 5 describes the empirical methodology, which is a straightforward application of GMM. Section 6 extends the analysis to include conditioning information; it shows how to test a dynamic, conditional factor pricing model by including scaled returns and scaled factors. Section 7 shows how the stochastic discount factor representation used in the paper is equivalent to traditional statements in terms of factor betas and risk premia; it shows how the scaled factor models are equivalent to quite general variation in conditional factor betas and risk premia. Section 8 estimates and tests the investment return model, and section 9 compares it with the consumption-based model, the CAPM and the Chen, Roll Ross model. Section 10 contains some concluding remarks.

2. Motivation for Production-based asset pricing models

Every asset pricing model that precludes arbitrage can be summarized by

$$p = E(m\ x)$$

where

p = price today

x = random payoff at a future date

m = stochastic discount factor,

or, since we usually use returns,

$$1 = E(m\ R)\ . \qquad (4)$$

Asset pricing models differ in how they relate the stochastic discount factor m to observables. For example, the capital asset pricing model (CAPM) is equivalent to

$$m = a + b \text{ market return} \qquad (5)$$

the APT and observable factor pricing models are equivalent to[1]

$$m = \text{linear combination of factors;} \qquad (6)$$

the consumption-based model and general equilibrium models are based on

$$m = \rho\ u'(c_t)/u'(c_{t-1}).$$

Production-based asset pricing models have two distinguishing characteristics: 1) They relate the stochastic discount factor m to a function of *production* data--output, investment, capital stock, inventories, etc. 2) They use as few preference assumptions as possible. They are, as much as possible, based on *firms'* adjustment of investment, output, etc. to changes in asset markets, rather than consumer's saving and asset allocation decisions.

The purpose of a production-based asset pricing model is to explain and model the relation between asset returns and economic fluctuations. There is a great deal of empirical evidence for such a link: the same variables forecast stock returns and GNP, stock returns are associated with contemporaneous and subsequent economic activity, expected returns are related to the covariances of returns with macroeconomic variables. For this purpose, there is reason to hope that production-based models may perform better than other asset pricing models currently popular in finance and macroeconomics.

In traditional asset pricing models, such as the CAPM and APT, expected returns are explained by the behavior of other returns. Though these models may successfully *capture* variation in expected returns, they will never help to *explain* it. To say that the expected return on a given asset varies over the business cycle because (say) the expected return on the market varies leaves unanswered the question, why does the expected return on the market vary? To *explain* variation in prices, it is necessary to examine models that tie prices to quantities.

Asset pricing models that use ad-hoc macroeconomic factors in combination with asset return factors,are also used to capture business-cycle related variation in expected returns (for example, Chen, Roll and Ross (1986), Ferson and Harvey (1991)). The investment return model *derives* the variables that should be the true "macro factors", and uses no "proxies" for factors based on asset returns. It thus has the advantages and disadvantages of any economically-derived model when compared to ad-hoc models.

The consumption-based model ties asset returns to non-durable consumption data (see among many others, Lucas (1978), Hansen and Singleton (1982); Ferson (1992) contains a review). In principle, it is exactly the framework we need to tie asset returns to economic fluctuations. However, despite a tremendous specification search, the consumption-based model has not fared well empirically. In, part, this poor performance is due to the facts that non-durable consumption growth barely moves over the business cycle, and it is very poorly correlated with stock returns.

A hunch motivating production-based models is that aggregate consumption may be de-linked from asset returns at business cycle frequencies due to small transactions costs, as it is at a minute-by-minute level. (Cochrane (1989), Cochrane and Hansen (1992) and Luttmer (1992) contain calculations and literature reviews). If this is true, *no* transformation of nondurable and services consumption growth, by *any* utility function, will provide a useful benchmark for asset returns. Consumption data will simply be uninformative about high frequency movements in asset returns.

Production-based models exploit the *firm's* first order conditions as the consumption-based model exploits the consumer's first order conditions. The model in this paper basically inverts a return version of the q theory of investment to read expected asset returns from investment data, as the consumption-based model inverts the standard theory of consumption and saving to read expected asset returns from consumption data. One can hope for better empirical performance, since production variables do display a substantial cyclical fluctuation (they define it), and they are more highly correlated with stock returns than is non-durable consumption. Furthermore, it is reasonable to hope that the transactions and information costs that seem to de-link consumption from asset returns at business cycle frequencies are less important for firms, due to their larger size.

General equilibrium asset pricing models with production are derived by substituting equilibrium relations between consumption and production variables into the consumption-based model. (A few examples related to the concerns in this paper are Brock (1982), Balvers, Cosimano and MacDonald (1990), Rouwenhorst (1990), Sharathchandra (1991), Detemple and Sundaresan (1991).) They do tie asset returns to production variables. However, they impose preference structures (typically log utility and frictionless markets) that are rejected using consumption data, and it is hard to trace what part of their predictions are due to preference vs. technology assumptions. They both allow and require us to take a stand on the ultimate sources of shocks (Technology? Money? ..), rather than infer the effect of shocks on asset markets through consumption or investment behavior.

Understanding the link between asset returns and economic fluctuations, by identifying and modeling the economic risk factors that determine expected returns, is the central task of asset pricing. There is an enormous literature that documents and characterizes variation in expected returns over time and across assets, but as yet no completely satisfactory class of models for the risk factors that drive this variation in expected returns. As a result, much of empirical finance winds up in a fruitless debate over whether variation in expected returns is due to as-yet-unnamed risk factors, or due to "fads" (see Cochrane (1991b) for an extended discussion). A new class of models for risk factors is obviously useful for finance.

The relation between asset returns and economic fluctuations is also a long-standing concern of macroeconomics. Most obviously, macroeconomists are interested in how changes in the stock market affect investment. A comparison of this paper with the empirical q-theory literature (for example, Abel and Blanchard (1982);

see Hubbard and Kashyap (1992) for a recent example and literature review) suggests that investment responds to changes in *risk premia* that the empirical finance literature has found to dominate changes in expected returns. Most q-theory models specify constant risk premia, and try, without much success, to explain changes in investment from changes in risk-free rates. The large residual in standard q tests is often interpreted as evidence for asymmetric information problems or credit constraints; the relative success of the model presented here may help to rehabilitate the neoclassical view.

More generally, macroeconomists are interested in the links between asset returns and fluctuations for the information they can provide about preferences, technologies and market structures that will be useful in the construction of macroeconomic models. Production-based asset pricing models can provide information about technologies and market structures faced by firms, just as consumption-based asset pricing models are a proving ground for preferences and market structures faced by consumers. For example, one lesson of this paper is that an adjustment cost (or some wedge between the price of installed and uninstalled capital), currently not included in most real business cycle models, is necessary to reconcile investment and asset returns.

3. Investment Returns: Definition and Construction

To construct investment returns from production data, I use adjustment cost technologies of the following form,

$$y_t = f(k_t, l_t) - c(i_t, k_t) \tag{7}$$

$$k_{t+1} = (1-\delta)(k_t + i_t) \tag{8}$$

where

y_t = output

$f(k_t, l_t)$ = production function

k_t = capital stock

l_t = labor input

i_t = investment

δ = depreciation rate

$c(i_t, k_t)$ = adjustment cost function.

The adjustment cost reflects the fact that it's hard to produce in periods of high investment. For example, it's hard to write papers when the painters are in your office.

The one-period investment return is the amount of extra output the firm can sell at t+1 if it invests an additional unit at t, leaving sales at t+2 t+3, .. unchanged. Section 1 of the Appendix goes through the algebra to show that the one-period investment return for the technology specified in (7)-(8) is given by

$$R^I_{t+1} = (1-\delta)\frac{1 + f_k(t+1) + c_i(t+1) - c_k(t+1)}{1 + c_i(t)}. \tag{9}$$

where

$$c_i(t) \equiv \frac{\partial c(i_t,k_t)}{\partial i_t},$$

etc.

The denominator $1+c_i(t)$ in (9) reflects the fact that some output is lost to adjustment costs when increasing investment at time t. The extra time t investment gives rise to extra capital stock at t+1; $f_k(t+1)$ is the extra output that results from the extra capital stock. $c_k(t+1)$ represents the change in adjustment costs at t+1, since the capital stock at t+1 is higher. At t+1, the firm must *lower* investment, to restore the capital stock at t+2 to its original value. The lowered investment means that more can be sold. $1+c_i(t+1)$ represents these lowered investment expenditures.

I use the following parametric specification of technology:[2]

$$y_t = \text{mpk } k_t + \text{mpl } l_t - \frac{\alpha}{2}\left(\frac{i_t}{k_t}\right) i_t. \tag{10}$$

In this case, the investment return (9) becomes

$$R^I_{t+1} = (1-\delta)\frac{1+\text{mpk}+\alpha(i_{t+1}/k_{t+1})+\alpha/2(i_{t+1}/k_{t+1})^2}{1+\alpha(i_t\,/\,k_t)}. \tag{11}$$

Though this function is not pretty, the investment return is approximately proportional to growth in the investment/capital ratio, or, since capital does not vary much, growth in investment. δ and mpk control the mean investment return; α affects the mean and the variance. Ad-hoc models using either investment growth or growth in the investment/capital ratio in place of the investment return (11) perform similarly, so the precise functional form is not crucial to the results. Since this technology is linear in capital, the (marginal) investment return defined in (11) is also equal to the return from holding a claim to the capital stock of the firm over the period (see Cochrane (1991a)).

For given values of the parameters $\{\alpha,\delta,\text{mpk}\}$, I form investment/capital ratios by accumulating capital according to equation (8) starting from the steady-state investment/capital ratio. That ratio is the solution to (8) with constant investment growth equal to the mean investment growth,

$$\frac{i}{k} = E\left(\frac{i_t}{i_{t-1}}\right)/(1-\delta) - 1 \ ;$$

Then, given α and mpk, I construct the investment returns from their definition (11).

4. Factor model.

As mentioned above, absence of arbitrage implies that there exists a strictly positive stochastic discount factor m, such that any asset return R obeys

$$1 = E(m\ R) \tag{12}$$

E can be interpreted as a conditional or unconditional expectation. I will be specific about conditioning information below.

Absence of arbitrage between asset markets and real investment opportunities (the firm's first order condition for profit maximization) implies that the investment return R^I of any production technology (a function of investment, output etc. data) must also obey

$$1 = E(m\ R^I), \tag{13}$$

or, more precisely, that there is a discount factor m that satisfies both (12) and (13). (Section 1 of the Appendix presents a derivation of (13).)

This observation leads immediately to two sets of testable implications. First, one can expand the space of returns on which one tests any asset pricing model (model for m) to include investment returns. Second, one can test for absence of arbitrage or consistent pricing between the set of asset and investment returns, by trying to construct positive m's that satisfy equations (12) and (13). Braun (1991) follows this approach.

This paper concentrates on *asset* pricing. What can we learn about asset returns R from investment returns R^I? The restriction I study is a factor pricing model, namely: *The investment returns are factors for the asset returns.*

Stated more formally, the law of one price (two identical payoffs must have the same price) implies that there is always a discount factor m that is a linear combination of the investment and asset returns and that prices both:[3] there is always an

$$m = \sum_i b_i R_i + \sum_j b_j R_j^I$$

such that

$$1 = E(m\ R_i) \text{ and } 1 = E(mR_j^I)$$

for all asset returns R_i, investment returns R_j^I.

The factor pricing model[4] is the restriction that the discount factor is only a function of the investment returns:

$$m = \sum_j b_j R_j^I \qquad (14)$$

(This is equivalent to the traditional statement in terms of factor betas and risk premia; see section 7). It is also equivalent to the statement that the investment returns span the mean-variance frontier of investment and asset returns.

Why should investment returns be factors for asset returns? There are two ways to derive any factor pricing model. First, one can assume that the space of returns (payoffs) under study has a factor covariance structure, and then appeal to arbitrage arguments. For example, we could assume that the firms on the NYSE are claims to different combinations of N production technologies, plus idiosyncratic components that have small prices. Second, we could invoke assumptions on preferences under which the returns on the N active production processes, which are the only non-diversifiable payoffs in the economy and add up to aggregate wealth, are the factors that price all other assets (for example, see Cox, Ingersoll and Ross (1985)).[5]

Motivated either way, the factor model studied in this paper is not a pure production-based asset pricing model. A pure production-based model would use *no* assumptions on preferences or restrictions on the space of asset returns, and read any asset return off producer's first order conditions, just as the consumption-based model uses no technology assumptions (i.e. is valid for any production technology) and reads asset prices off consumer's first-order conditions. Such models are possible, but the model described here does not quite reach this ideal.

The number and nature of the intertemporal technologies that drive asset returns, or, equivalently, the appropriate level of aggregation of the capital stock, is a modeling choice. "My car" and "your car" are both ways of getting consumption services from today to tomorrow, but hopefully their behavior across states of nature that affect asset returns is sufficiently similar that we can aggregate them into "cars". However, there is no reason to believe *a priori* that all the intertemporal investment opportunities in the economy will be summarized by one or two aggregated production functions. This paper follows the "spirit of the

APT" that there are only a few factors, but this is an additional modeling assumption, not a prediction of the theory; models with highly disaggregated investment opportunities may turn out to be more useful for some purposes.

5. Empirical Methodology

The above statement of the factor pricing model maps naturally into the GMM framework for estimation and testing. I use excess returns (return differences, not necessarily returns less a risk-free rate proxy) to focus on variation in risk premia. Therefore, the moment conditions are

$$0 = E(mR^e) \quad R^e = \text{vector of excess returns.} \tag{15}$$

Using excess returns, the mean discount factor is not identified: if $0 = E(mR^e)$, then $0 = E((\text{constant} \bullet m)R^e)$. It is convenient to normalize so the sample mean discount factor is 1. Thus, we can reexpress any factor model

$$m = f'b \qquad f = \text{vector of factors}$$

in terms of mean zero factors as

$$m = 1 + \tilde{f}'b$$
$$\tilde{f} = f\text{-}E(f) = \text{vector of de-meaned factors} \tag{16}$$

Following the standard GMM procedure (Hansen (1982), Hansen and Singleton (1982)), we choose the parameters b to minimize a weighted combination of the sample moments (15). Using Hansen's notation,

$$\min_{\{b\}} J_T = g_T' W g_T \tag{17}$$

where

$g_T = E_T(mR^e) = E_T(R^e (1 + \tilde{f}'b));$

E_T = sample mean, $\frac{1}{T}\sum_{t=1}^{T}$

W = weighting matrix.

Since the parameters b enter linearly, we can find their estimates analytically[6] rather than by search,

$$\hat{b} = - (C'WC)^{-1}C' W E_T(R^e) \tag{18}$$

with

$C \equiv E_T(R^e\tilde{f}')$.

This estimate has a natural interpretation. The moment condition

$$0 = E(mR^e) = E(R^e) + E(R^e\tilde{f}')b \tag{19}$$

states that mean returns should be a linear function of the covariances of returns with factors. The estimate of b in equation (18) is the coefficient in the GLS regression of expected returns $E_T(R^e)$ on covariances C -- a natural way of making expected returns "as close to" linear in covariances as possible.

The GMM distribution theory (Hansen (1982)) gives an asymptotic joint normal distribution for $\hat{b}$.[7] Hence, t tests on individual b's or χ^2 tests on groups of b's can be used to test whether a factor or group of factors is priced. If a factor is not priced, it does not affect m, so its b should be zero. The GMM distribution theory also provides a χ^2 test whether the minimized value of the objective (17) is significantly different from zero, i.e. of the null that the moments $E(mR^e)$ are equal to zero 0. This test is known as the J_T test or the overidentifying restrictions test. It is the basic test whether we can statistically reject a given observable factor model against a non-specific alternative.

It is also interesting to test a model against specific alternatives, i.e. to ask "given factors a,b,c.., is factor x (or are factors x, y and z) priced?" There are two ways to perform such tests, corresponding to Wald and Likelihood ratio philosophies. Start with a general model that includes both sets of factors. First, we can use the t or χ^2 tests for b=0 to test if a given factor or group of factors is not priced in the presence of the other factors. Second, we can compare the overidentifying restrictions of a restricted system that excludes a given set of factors to the overidentifying restrictions of the unrestricted system that includes all factors. If the excluded factors are not priced, the J_T should not rise much. Precisely, if we use the same weighting matrix to estimate both systems (I use the weighting matrix from the unrestricted system) the difference in J_T statistics has a χ^2 distribution, with degrees of freedom equal to the number of omitted parameters (Newey and West (1987b)).

When the factors are investment returns, I additionally estimate the production function parameters. Since these parameters enter nonlinearly, a search is required. The programming is harder, but the GMM methodology extends trivially.

6. Conditional estimates and conditional factor models

So far, I have considered *unconditional* factor models, and estimates of unconditional moments. The effects of conditioning information are easily included by scaling the returns and/or the factors by instruments. Specifically, to test the conditional moments

$$0 = E(m_{t+1} R^e_{t+1} \mid I_t), \qquad (20)$$

we expand the set of returns to include returns scaled by instruments, and then proceed as before; i.e. we use the moment conditions

$$0 = E(m_{t+1} (R^e_{t+1} \otimes z_t)). \qquad z_t \in I_t \qquad (21)$$

where $\otimes$ denotes the Kronecker product (multiply every asset return by every instrument). To test a model in which the factors conditionally price assets, we expand the set of factors to include factors scaled by instruments,

$$m_{t+1} = b'(f_{t+1} \otimes z_t).$$

To show how scaling returns works, or that (21) tests (20), multiply both sides of (20) by z_t and take unconditional expectations.[8] Conversely, if (21) holds for *all* variables z_t in an information set I_t, then (20) holds.[9] Thus, expanding the payoff space to include scaled returns as in (21) can test *all* of the implications of (20), so that no generality is lost in principle. Of course, the usual instrument selection problem remains, since we cannot in practice test (21) with every variable observed at time t.

To motivate scaling factors, note that we have supposed so far that the discount factor m is a *fixed* linear combination of a given set of factors. However, the discount factor m might be a linear combination of factors with weights that vary as a vector of instruments z varies across different information sets,

$$m_{t+1} = b(z_t)'f_{t+1}.$$

Again, it is sufficient to consider b's that vary linearly with the instruments, since nonlinear functions can be expressed as linear functions of additional instruments. Thus, with one instrument z, and dropping the time subscripts, the conditional factor model is

$$m = (b_0 + zb_1)'f,$$

But scaling the factors f by the instruments z achieves the same result. The last equation is equivalent to

$$m = b_0 f + b_1(fz).$$
$$= b'(f \otimes z).$$

Therefore, given the choice of instruments, performing the GMM estimation and testing with scaled factors is in principle a completely general test of a dynamic, conditional factor pricing model based on the instruments. Again, the only complaint one can make is that more or other instruments (or functions of instruments) should have been included.

One can form any combination of conditional and unconditional estimates (including scaled returns or not) and conditional and unconditional factor model (including scaled factors or not). For example, scaled factors without scaled returns tests the unconditional implications of a conditional factor model. These are different than the unconditional implications of an unconditional factor model, since the latter does not include scaled factors. Henceforth, I will refer to "scaled factor models" rather than "conditional factor models" to distinguish the two cases.

The models with scaled factors do not eliminate the problem that consumers may observe finer information sets than we do. A conditional factor pricing model with respect to a fine information set does not imply a conditional factor pricing model with respect to a coarser information set, or an unconditional factor model. Equivalently, conditional mean-variance efficiency does not imply unconditional mean-variance efficiency, though the converse is true (Hansen and Richard (1987)). Thus a rejection of any factor model that is derived as a conditional factor model with respect to consumer's information may still be attributed to an insufficiently rich set of instruments. However, scaling factors does provide a very easy method for estimating and testing generally specified conditional factor pricing models *given* an information set.

7. Relation to traditional statement of factor models

The statement that the discount factor m is a linear function of factors is equivalent to the conventional statements of factor pricing models in terms of betas and factor risk premia. Precisely, the scaled factor model

$$m = (f \otimes z)'b, \tag{22}$$

together with the conditional pricing relation

$$1 = E(mR \mid I) \tag{23}$$

for an information set I such that $z \in I$, implies the traditional statement of a factor pricing model.

$$E(R^e \mid I) = \quad cov(R^e, f' \mid I)\ E(ff' \mid I)^{-1} \quad (-1)\ E(fm \mid I)\ /\ E(m \mid I).$$
$$E(R^e \mid I) = \quad \text{-------}\ \beta(I)'\ \text{-------} \qquad \text{-------}\ \lambda(I)\ \text{-------}.$$

i.e., $\beta(I)$ = multiple regression coefficients of R^e on f, conditional on I, and $\lambda(I)$ = vector of conditional factor risk premia. The proof[10] just consists of recognizing b as the (conditional) regression coefficient of m on f, substituting the formula for the regression coefficient in the pricing relation (23), and rearranging.

Conversely, if expected returns are linear in the conditional regression coefficients of returns on some factors, then there exists an m of the form (22) that prices assets, (23). (It is not unique, since one can add any random variable conditionally orthogonal to returns to m.)

In this way, the inclusion of scaled factors can model arbitrary variation in conditional betas and factor risk premia λ, subject only to the choice of instruments. Most tests of factor pricing models include auxiliary assumptions, such as constant conditional betas, constant conditional factor risk premia, constant conditional covariances, etc. Furthermore, the factors do *not* have to be conditionally mean zero (white noise), conditionally or unconditionally orthogonal, or conditionally or unconditionally homoskedastic, as is often assumed.

8. Estimation and testing of the investment return factor model.

8.1 *Set-up*

I use a simple specification of the investment model. There are two investment technologies, corresponding to gross private domestic nonresidential and residential investment (CITIBASE series GIN82 and GIR82. (Section 2 of the Appendix details the sources and transformations used for all data series.) I assume that each investment series corresponds to a technology of the form (10), so that its investment returns are given by (11).

For asset returns, I use the 10 portfolios of NYSE stocks sorted by market value (size) maintained by CRSP. There is a large spread in the mean returns of these portfolios: the small firm decile's mean excess return is almost twice that of the large firm decile. Any asset pricing model must explain this spread in mean returns by spread in assets' covariance with risk factors.

Since the investment returns are based on quarterly average investment, I transformed the asset returns to quarterly average returns rather than use end-of-quarter to end-of-quarter returns. I include moment conditions for investment returns along with the moment conditions generated by asset returns, since both sets of returns should be correctly priced. I created excess returns by subtracting the three month t-bill rate in each case.

I use two instruments: the default premium (yield on BAA corporate bonds - yield on AAA corporate bonds), and the equally-weighted dividend/price ratio. These instruments are popular forecasters of stock returns. In the first-stage estimation, the moments corresponding to scaled returns are treated equally with the non-scaled returns, so it is convenient that the scale of the two is roughly comparable. To this end, I transformed the instruments to have a mean of 1 and a standard deviation of 0.2. To avoid overlap with the averaged return series, I lagged the instruments twice.

If we allow all of the production function parameters $\{\alpha, \delta, \text{mpk}\}$ to vary, the system is over-parameterized. Examining the definition of the investment return (11), the parameters δ and mpk basically affect the mean of the investment return, while α affects the mean and standard deviation. None of the parameters substantially affects the cross-correlation of investment returns with other variables; these are basically given by the cross-correlation of investment growth with the other variables. Furthermore, the mean and standard deviation of the factors are not separately identified,[11] so three parameters control one moment. As a result, the minimization surface has a valley in it, and the program soon crashes with a singular gradient matrix $\partial g_T([\alpha,\delta,\text{mpk}])/\partial[\alpha,\delta,\text{mpk}]$. Therefore, I present results in which α and δ are held fixed, minimizing only over mpk. I tried choosing each of the parameters, and the pricing results are very similar, though the actual parameter values are obviously different. Also, choosing the parameters sequentially (first α, then mpk, etc.) leads to the same general results.

The tables present results using 4 Newey-West (1987a) lags (k=4 in the notation of footnote 7) to construct standard errors. The standard errors are generally a little smaller with k=0, but the results overall are not much changed, indicating little autocorrelation of the residuals. The tables present only the iterated GMM estimates and tests. In most cases the first-stage and second-stage estimates and tests yield similar results.

8.2 *Estimates and tests of the investment model*

Table 1 presents estimates and tests of the investment return factor model. Start with the simple unconditional estimate of the non-scaled factor model, panel 1A. The marginal product of capital parameters mpk are plausible and highly significant. They have about the same value (0.05 - 0.06) and are highly significant in all the following estimates. The estimates and tests of the b's measure whether the investment return factors are priced.[12] The residential factor is significantly priced (t = -2.78), while the nonresidential factor is not (t = 1.10). They are jointly significant (p-value for joint b=0 is 2.08%). Finally, the J_T test of overidentifying restrictions does not reject the model (p-value 54.3%)

The conditional estimates (panel B) are formed by adding scaled returns. Since there are two instruments and a constant, this triples the number of moment conditions, which should sharpen the estimates. It also asks the interesting question whether the model can account for variation in expected returns over time as well as across assets. In this estimate, both investment return factors are now individually significant (t on b 3.47 and -9.78), and jointly highly significant (p value < 0.000%). However, the J_T statistic now convincingly rejects the model (p-value 0.006%).

The natural solution is to include scaled factors, to test a conditional version of the investment return factor model,[13] panel 2. Now there is a factor and a coefficient b corresponding to each factor multiplied by each instrument. These scaled factors are individually and jointly highly significant. Also, the residential and nonresidential factors are significant as sub-groups, as are the scaled and non-scaled factors. Now, the J_T test does not reject (p-value 24.9%).

Figure 1 presents a graphical measure of the non-scaled factor model's fit. It plots the model's predictions for expected returns vs. the sample expected returns. The solid dots are the predictions of the unconditional estimates (panel 1A table 1), so each dot corresponds to a size decile portfolio or an investment return. The triangles correspond to the predictions of the conditional estimates (panel 1B table 1), so each triangle represents a size decile return or investment return scaled by an instrument (including one). The figure shows how the model does a pretty good job of predicting the cross-sectional variation of expected returns by variation in the covariances of returns with the investment returns (the dots lie pretty close to the 45^o line), and how the model does a less good job of explaining the variation in expected scaled returns.[14]

The returns on the bottom left of Figure 1 are the investment returns. Their placement is not an essential feature of the model. It is easy to produce investment returns that lie farther apart or at different places along the line in Figure 1, yet price about as well, by different choices of the fixed parameters α and δ.

Figure 2 presents the mean excess returns vs. model predictions for the scaled factor model. Comparing Figure 2 with Figure 1, the scaled factor model looks much better, especially at pricing the scaled returns, i.e. in its implications for variation in expected returns over time.

9. Comparison with other models

The overidentifying restrictions (J_T) test the investment return model against no specific alternative. But *all* currently available non-trivial models (including the investment model) can undoubtedly be rejected if one uses a sufficiently rich set of assets and instruments. Therefore, in evaluating a model and learning how one might improve it, it may be more interesting to compare a given model to plausible competitors, rather than simply reject or fail to reject it.

In this section, I compare the investment return model to the CAPM, the Chen, Roll and Ross factor model, and the consumption-based model. In each case, I estimate and test the competing model, in the style of table 1 and Figures 1 and 2. Then, I estimate models that include *both* investment return and the other factors, to see which set of factors is priced in the presence of the other.

9.1 *CAPM*

The CAPM is a single factor model with the market return R_m as factor,[15]

$$m = \text{constant} + R_m b.$$

Thus, it trivially maps into the factor pricing-GMM framework outlined above.

Table 2 presents GMM estimates and tests of the CAPM. The CAPM behaves about the same way as the investment model. In the unconditional estimate of the non-scaled model, the market return is significantly priced (see t and χ^2 on b) and the unconditional J_T test of overidentifying restrictions does not reject. In the conditional estimate, the market return is more strongly priced, but the J_T test rejects the model. When we include the scaled market return factors, the b's are jointly significant. Unlike the investment model, the overidentifying restrictions are rejected.

Figure 3 and Figure 4 show that the CAPM performs about as well as the investment model, Figure 1 and Figure 2. As is traditional, the worst performance occurs with small firms--the top dot, which is the first decile, is the farthest nonscaled return (dot) away from the 45^{o} line in Figure 3. However, the small firm effect disappears entirely once we include scaled market returns as factors, Figure 4. Thus, the apparent small firm effect may simply be due to inadequate treatment of conditioning information. Most derivations of the CAPM specify that the market is conditionally, but not unconditionally, mean-variance efficient, so this result is not too surprising. (It may also be due to a failure of the CAPM--none of the other models display a small firm effect.)

Do the investment returns drive out the market or vice versa? In a factor model that includes *both* the market and the investment returns, which are significantly priced? The row marked "VW" in Table 3 presents tests based on a model that includes both investment return factors and the market to address this question, as explained above.

In the unrestricted model, the b's corresponding to scaled investment returns are highly significant (p value

0.003 %). However, the scaled market returns are not priced in the presence of the scaled investment return (p value 29%). The "likelihood ratio" test, formed by the rise in χ^2 statistic for constrained vs. unconstrained models, tells the same story. The model that excludes the scaled investment returns is rejected (p-value < 0.000%), while the model that excludes the scaled market returns is not rejected (p value 29%). Thus, the investment returns drive out the market return, and not vice versa.

9.2 *Chen, Roll Ross model*

The Chen, Roll Ross (1986) (CRR) model was explicitly designed to link stock returns to economic fluctuations, and Chen, Roll and Ross claim that their model drives out the market return. Thus, it is an important alternative model to examine. Chen, Roll and Ross advocate a five-factor model, in which the factors are

MP = growth in industrial production

DEI = change in inflation forecast

UI = inflation forecast residual

UPR = return on corporate bonds - return on 10 year government bonds

UTS = return on 10 year government bonds - return on bills.

They advertise these variables as "macroeconomic factors", though in fact all but MP are based on asset returns, just like the CAPM (the inflation forecasts are based on T-bill returns).

Table 4 presents GMM estimates and tests of the CRR model.[16] In the unconditional estimate, the CRR factors are individually and jointly (barely) insignificant. However, in the more powerful conditional estimate, three factors are individually significant, and all the factors together are jointly significant. Both models are not rejected by the J_T test for overidentifying restrictions.

It is not clear whether Chen, Roll and Ross intend their model as a conditional or unconditional factor model. Their test allows some variation in betas, but imposes constant factor risk premia (lambdas), and they only attempt to explain unconditional expected returns. Nonetheless, I include a scaled Chen, Roll and Ross model in table 4, and compare the scaled investment model to a scaled CRR model. This model is suspiciously overparameterized, since it has 15 scaled factors to explain 30 moments. The factor b's are mostly individually insignificant. However, they are jointly significant. The scaled factor b's are significant and the unscaled factor b's are not, suggesting that scaling is important for the CRR model. The overidentifying restrictions test is in the "too good to be true" tail (p-value 96.8%).

Figures 5 and 6 present mean returns vs. model predictions for the Chen, Roll and Ross model, to allow comparison with the investment model and CAPM.[17] The CRR model performance is very similar to the investment model and CAPM.

The row marked "CRR" in Table 3 presents a comparison of the investment model with the CRR model. As with the CAPM, the scaled investment return b's are significant, while the scaled CRR b's are not. Similarly, we reject dropping the investment return factors but not dropping the CRR factors. Thus, the investment model drives out the CRR factors and not vice versa, as with the CAPM.

9.3 *Consumption-based model*

The consumption-based model is perhaps the most appropriate comparison. Like the investment model, the consumption-based model relates asset returns strictly to macroeconomic data rather than other asset returns. It is based on a measure of consumers' intertemporal marginal rate of *substitution* where the investment model is based on a measure of firms' intertemporal marginal rate of *transformation.*

Table 5 presents GMM tests of the basic consumption-based model, and Figure 7 gives the mean return vs. predictions of the consumption based model.[18] The consumption-based model performs much worse than any of the other models studied so far. The expected return scatter is dramatically larger in the figures. While the conditional J_T test is not rejected,[19] the probability value is lower than for the other models. The table shows the large point estimate of the risk aversion coefficient familiar from the equity premium puzzle literature.

The row marked "consumption" in Table 3 presents a comparison of the investment model with the consumption-based model. The unrestricted model here contains the scaled investment returns and consumption growth raised to a risk aversion coefficient. The coefficient of m on the consumption factor is not constrained to one. Again, the "likelihood ratio" test finds that the investment returns drive out the consumption-based model. The b=0 test finds that the investment returns are priced at ridiculous levels of significance (χ^2 = 479 with 6 degrees of freedom), but also finds that the consumption-based discount factor is priced. However, the latter result is not stable: in the first stage estimate the test for consumption b=0 yielded a p-value of 29% rather than 0.11%, and in the second stage it was 53%. Furthermore, the risk aversion coefficient in the iterated GMM estimate was -269, making it harder to take the consumption factor seriously.[20]

10. Concluding remarks

The simple investment return model performs surprisingly well. The investment return factors are significantly priced, the model is not rejected, it is able to explain a wide spread in expected returns. The model performs about as well two standard finance models, the CAPM and the Chen, Roll and Ross factor model. In comparison tests, the investment return factors drive out both the market return and the Chen, Roll and Ross factors. The investment return model performs substantially better than the standard consumption-based model.

Even if the investment returns did not drive the other factors out, the fact that *any* model whose factors are derived from economic theory and are based solely on quantity data is even in a position to challenge the empirical success of traditional finance models may be regarded as an encouraging initial success.

In all cases, the scaled factor models perform substantially better than the non-scaled factor models. This suggests that time-variation in the parameters of asset pricing models, which can be handled by the simple expedient of including scaled factors, is an important ingredient for their empirical success.

Appendix

1. Derivation of investment returns from the production function.

This section derives the investment return from the production technology and shows that the firm's first order conditions direct the firm to remove arbitrage opportunities between investment and asset returns. This derivation follows that of Braun (1991); Cochrane (1991a) presents a derivation of the investment return directly from its definition as the marginal extra sale possible tomorrow from a marginal investment today.

The firm has the production technology given by

$$y_t = f(k_t, l_t) - c(i_t, k_t) \tag{7}$$

$$k_{t+1} = (1-\delta)(k_t + i_t) \tag{8}$$

The firm maximizes its present value,

$$\max_{\{i_t\}} E_t \sum_{j=0}^{\infty} m_{t,t+j} (y_{t+j} - i_{t+j}) \tag{24}$$

Subject to (7) and (8).

In a complete market, m are the contingent claims prices divided by probabilities, so this present value is the firm's time-t contingent claim value. If markets are less than complete, the firm still maximizes (24), but m is now an extension of the stochastic discount factor that prices asset returns, rather than *the* stochastic discount factor for the whole economy.

We derive the first order condition by varying i_t. Note that

$$\partial k_{t+j}/\partial i_t = (1-\delta)^j .$$

Hence,

$$\partial y_{t+j}/\partial i_t = (\partial y_{t+j}/\partial k_{t+j})\,(\partial k_{t+j}/\partial i_t)$$

$$= (1-\delta)^j\,(f_k(t+j) - c_k(t+j)).$$

The notation $f_k(t)$ means "partial derivative with respect to k, evaluated with respect to the appropriate arguments at time t", $f_k(t) \equiv \partial f(k_t, l_t)/\partial k_t$. The first order condition is then

$$1 + c_i(t) = E_t \sum_{j=1}^{\infty} m_{t,t+j}\,(1-\delta)^j\,(f_k(t+j) - c_k(t+j)). \quad (25)$$

Notice the left hand side is the relative price of a unit of installed capital vs. output today; the right hand side is the present value of its benefits.

We desire a model of returns, rather than price and present value. Using $m_{t,t+j} = m_{t,t+1}m_{t+1,t+j}$, break the right hand side of (25) into two pieces,

$$1 + c_i(t) = E_t\Big[\, m_{t,t+1}\,(1-\delta)\Big(\, f_k(t+1) - c_k(t+1) +$$

$$\sum_{j=1}^{\infty} m_{t+1,t+1+j} \, (1-\delta)^j \, (f_k(t+1+j) - c_k(t+1+j))\Big]\Big]$$

Substituting (25) at time t+1 for the sum in the right hand side,

$$1 + c_i(t) =$$

$$E_t\Big[m_{t,t+1} \, (1-\delta)\Big[f_k(t+1) - c_k(t+1) + 1 + c_i(t+1)\Big]\Big]$$

$$1 = E_t\left[m_{t,t+1} \quad (1-\delta) \, \frac{1 + f_k(t+1) + c_i(t+1) - c_k(t+1)}{1 + c_i(t)}\right]$$

or

$$1 = E_t\left[m_{t,t+1} \, R^I_{t+1}\right]$$

with

$$R^I_{t+1} = (1-\delta) \, \frac{1 + f_k(t+1) + c_i(t+1) - c_k(t+1)}{1 + c_i(t)} .$$

For some production technologies it is not possible to summarize the price - vs. present value absence of arbitrage (25) in a single period investment return. For example, if the adjustment cost depends on p lags of investment, then a p-period investment strategy must be considered.

2. Data description

All asset return data are from CRSP. NIPA data and yield data are from CITIBASE. The two investment returns are based on CITIBASE series GIN82 and GIR82. The stock return series are based on CRSP series EWRETD, VWRETD, and the size decile return series DECRET1...DECRET10. The default premium is based on CITIBASE series FYBAAC-FYAAAC. Quarterly data are obtained by using the last month of the quarter. EW d/p is based on CRSP EWRETD and EWRETX, the equally weighted portfolio returns with and without dividends. The returns are cumulated for a year to avoid the seasonal in dividends, then EW d/p = annual EWRETD / annual EWRETX - 1. Again, the last monthly observation in each quarter is the quarterly observation.

The investment data are quarterly averages, while the asset return data are point-to-point. To correct for this difference, I averaged monthly asset returns over the quarter to correspond with the investment returns.[21] Thus the second quarter return is an average of returns from the last day in December to last day in March, last day in January to last day in April, and last day in February to last day in May. Instruments for the second quarter return are all observed at the end of December (i.e., all instruments are lagged twice). Figure A.1 summarizes the timing relations among the variables.

I constructed Chen Roll and Ross factors as follows:

MP: the growth rate of industrial production. CRR lead this variable by one month to take account of the fact that IP is a monthly average and returns are end-of-month to end-of month. To make the same adjustment for quarterly data, I average IP growth in a similar way to returns. Thus,

$$Q2\ MP = \ln \left[\frac{IP(apr)\ IP(may)\ IP(jun)}{IP(jan)\ IP(feb)\ IP(mar)} \right]$$

UI, unexpected inflation and DEI, change in expected inflation. These variables require a expected inflation series. CRR take their values from Fama and Gibbons (1982). Therefore, I replicated the Fama and Gibbons procedure to extend the data set. Fama and Gibbons start with the Fisher equation

$$E_{t-1}(I_t) = TB_{t-1} - E_{t-1}(R_t)$$

I = inflation

TB = T-Bill rate

R = ex-post real rate, $R_t = TB_{t-1} - I_t$

They add a univariate time-series model for ex-post real rates

$$R_t - R_{t-1} = u_t + \theta u_{t-1}$$

Substituting,

$$E_{t-1}(I_t) = TB_{t-1} - R_{t-1} - \theta\, u_t.$$

To construct this series, I take Fama and Gibbon's value of θ, 0.9923. I start with $u_1 = 0$, then I construct u_t by

$$R_2 - R_1 = u_2$$

$$R_3 - R_2 = u_3 - \theta u_2$$

...

where TB_{t-1} is the one month treasury bill rate and $E_{t-1}R_t$ is the expected real return on T-bills, given by

$$E_{t-1}R_t = E_{t-2}R_{t-1} + (1\text{-}0.9223)u_{t-1}$$

References

Abel, Andrew B., and Olivier J. Blanchard, 1986, The Present Value of Profits and Cyclical Movements in Investment, *Econometrica* 54, 249-273.

Balvers, Ronald J., Thomas F. Cosimano, and Bill McDonald, 1990, Predicting Stock Returns in an Efficient Market, *Journal of Finance*, Forthcoming.

Braun, Phillip, 1991, "Asset Pricing and Capital Investment: Theory and Evidence", Manuscript, Northwestern University

Brock, William A., 1982, Asset Prices in a Production Economy, in Jon J. McCall Ed.: *The Economics of Information and Uncertainty* (University of Chicago Press, Chicago, IL).

Chen, Nai Fu, Richard Roll, and Steven Ross, 1986, "Economic Forces and the Stock Market", *Journal of Business* 59, 383-403.

Cochrane, John H., 1989, "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near Rational Alternatives", *American Economic Review* 79 319-337.

Cochrane, John H., 1991a, "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations", *Journal of Finance* 46 207-234.

Cochrane, John H., 1991b, "Volatility Tests and Efficient Markets, A Review Essay", *Journal of Monetary Economics* 27, 463-485.

Cochrane, John H., 1992, "Simple GMM Tests of Traditional Asset Pricing Models", Manuscript, University of Chicago.

Cochrane, John H. and Lars Peter Hansen, 1992, "Asset Pricing Lessons for Macroeconomics", 1992 NBER Macroeconomics Annual.

Cox, John C., Jonathan E. Ingersoll Jr. and Steven A. Ross, 1985, "An Intertemporal General Equilibrium Model of Asset Prices", *Econometrica* 53, 363-384.

De Santis, Giorgio, 1992, "Volatility Bounds for Stochastic Discount Factors: Tests and Implications From International Stock Returns", Manuscript, University of Chicago.

Detemple, Jerome B. and Suresh Sundaresan, 1991, "Asset Prices with Production Externalities and Habit Formation", Working Paper, Columbia University.

Fama, Eugene F. and Gibbons, Michael, 1982, Inflation, Real

Returns, and Capital Investment, *Journal of Monetary Economics* 13, 327-48.

Ferson, Wayne E., 1992, "Theory and Empirical Testing of Asset Pricing Models", in *The Finance Handbook*, Robert A. Jarrow, William T. Ziemba, and Vojislav Makisimovic, Editors, North-Holland (forthcoming).

Ferson, Wayne E. and Campbell R. Harvey, 1991 "The Variation of Economic Risk Premiums" *Journal of Political Economy* 99, 385-415.

Ferson, Wayne E. and Stephen R. Foerster, 1990, "Finite Sample Properties of Methods of Moments in Latent Variable Tests of Asset Pricing Models" Manuscript, University of Chicago.

Hansen, Lars Peter, 1982, Large Sample Properties of Generalized Method of Moments Estimators, *Econometrica* 50, 1029-1054.

Hansen, Lars Peter, and Scott F. Richard, 1987, The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models, *Econometrica* 55, 587 - 613.

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

Hansen, Lars Peter and Ravi Jagannathan, 1991a, "Implications of Security Market Data for Models of Dynamic Economies", *Journal of Political Economy* 91, 225-262.

Hansen, Lars P. and Ravi Jagannathan, 1991b, "Assessing Specification Errors in Stochastic Discount Factor Models", Manuscript, University of Chicago and University of Minnesota.

Hubbard, R. Glenn and Anil K Kashyap, 1992, "Internal Net Worth and the Investment Process: an Application to U.S. Agriculture", forthcoming, *Journal of Political Economy*

Ingersoll, Jonathan E. Jr., 1988, *Theory of Financial Decision Making* (Rowman and Littlefield, Totowa N.J.).

Knez, Peter J., 1991, "Pricing Money Market Securities With Stochastic Discount Factors", Manuscript, University of Wisconsin.

Lucas, Robert E. Jr., 1978, "Asset Prices in an Exchange Economy", *Econometrica* 46, 1426-1446.

Luttmer, Erzo, 1992, "Implications of Asset Market Data for Economies with Transaction Costs", Manuscript, University of Chicago.

Newey, Whitney K. and Kenneth D. West, 1987a, A Simple, Positive- Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-708.

Newey, Whitney K. and Kenneth D. West, 1987b, "Hypothesis Testing with Efficient Method of Moments Estimation", *International Economic Review* 28, 777-787.

Rouwenhorst, K. Geert, 1990, Asset Returns and Business Cycles: a General Equilibrium Approach, Working paper, Graduate School of Business, University of Rochester.

Sharathchandra, Gopalakrishnan, 1991, Asset Pricing and Production: Theory and Empirical Tests, Manuscript, Finance Department, Cox School of Business, Southern Methodist University.

Snow, Karl N., 1991, "Diagnosing Asset Pricing Models Using the Distribution of Asset Returns", *Journal of Finance* 46, 955-984.

Footnotes

[1]Substituting (5) or (6) in (4), and using E(ab) = E(a)E(b) + cov(a,b) one finds that expected returns are linear in their covariance with factors. See section 7 for details.

[2]This technology may seem unduly simple, in that it ignores imperfect substitutability between capital and labor, declining marginal product of capital, productivity shocks, taxes, and many other features of a plausible production technology.

As a partial defense for simplicity, many of these features have second-order effects on the results. Braun (1991) found this insensitivity in detailed experiments; Sharathchandra (1991) models a concave technology with production function shocks and no adjustment costs, and obtains an essentially constant investment return. Intuitively, production shocks, changes in the productivity of capital induced by changes in the capital stock or employment and tax changes are low - frequency changes to the "dividend" component of the investment returns. The adjustment cost is the "price change" component, (price of installed vs. uninstalled capital) which swamps the other components. Q theory tests based on price vs. present value rather than returns are much more sensitive to these low frequency corrections.

However, this is only a partial defense. Other features such as gestation lags, changes in the investment tax credit or depreciation allowances may have large effects on the results, and using more realistic if more complicated technologies seems an important area for future research.

It is perfectly possible to include production shocks, since they (unlike taste shocks) can be measured. One does not have to believe that all or any shocks to the economy are not productivity shocks in order to use a production-based model.

[3]See, among others, Hansen and Richard (1987). If one imposes the stricter no-arbitrage condition, m must be positive, and m in the next equation is not necessarily a linear function of returns; one may have to include options as well. See Hansen and Jagannathan (1991a).

[4]Equation 14 is a factor *pricing* model. A factor structure on the covariance matrix of returns is sometimes used to derive factor pricing, but factor pricing does not imply or require a factor structure. If we write

$$R = E(R) + \beta' R^{I} + \text{error},$$

factor *pricing* occurs if the errors have zero prices. A factor *structure* occurs if the errors are uncorrelated. Uncorrelated errors may imply zero price errors, but the converse is not true.

[5]The central assumption here is that preferences are time-separable; if they are not, then past investment returns could affect current asset returns. One could, of course, account for potential non-separabilities by including past investment returns as additional factors. With general preferences in discrete time, nonlinear functions of investment returns might also enter m; one can regard linearity either as an assumption on preferences or as a first-order approximation.

[6]The first order conditions to the minimization are

$$\frac{\partial g_T'}{\partial b} W\ g_T = E_T(R^e \tilde{f}')'\ W\ g_T = C'\ W\ (E_T(R^e) + Cb) = 0;$$

Solving, we get (18).

[7] The first-stage GMM estimates use an identity weighting matrix. Thus,

$$\hat{b}_1 = - (C'C)^{-1}C' \, E_T(R^e).$$

Next, construct an estimate of the spectral density at zero of mR^e

$$S_T = \sum_{j=-k}^{k} \frac{k-|j|}{k} \; E_T[(R_t^e(1 + f_t'\hat{b}_1)) \, (R_{t-j}^e(1 + f_{t-j}'\hat{b}_1))']$$

The second-stage GMM estimates use this weighting matrix,

$$\hat{b}_2 = - (C'W^*C)^{-1}C'W^* \, E_T(R^e); \qquad W^* = S_T^{-1}.$$

The asymptotic covariance matrix of either estimate is computed as

$$\text{var}(b) = \frac{1}{T} \left[\frac{\partial g_T'}{\partial b} W^* \frac{\partial g_T}{\partial b} \right]^{-1} = \frac{1}{T} (C' \, S^{-1} \, C)^{-1}.$$

The test for overidentifying restrictions is based on

$$T \, J_T = T \, g_T' \, W^* \, g_T \sim \chi^2 \text{ (DF = \#asset returns - \#factors)}$$

where g_T is formed using the second stage estimates,

$$g_T = E_T(R^e(1+f'\hat{b}_2)).$$

I iterated the GMM procedure, by forming a W* using second-stage estimates, finding third-stage estimates and so forth. Ferson and Forester (1991) find that this procedure gives better small sample performance. It also produced results that were more stable across small variations in the model set-up.

[8] In the more general case of returns rather than excess returns, one must modify the moment condition to include expected prices rather than 1 on the left hand side: $1 = E(mR \mid I)$ implies $E(1 \otimes z) = E(m(R \otimes z))$.

[9]Note that if $z_t \in I_t$, then $z_t^2 \in I_t$. Thus, "every variable" in I_t means every variable, and every measurable function of every variable.

[10]*Proof:* Recall that the scaled factor model, $m = (f \otimes z)'b$, is equivalent to a model in which the factor b's vary linearly with z, $m = f'b(z)$. This implies that b(z) is the conditional regression coefficient of m on f,

$$b(z) = E(ff' \mid I)^{-1} E(fm \mid I). \tag{1}$$

where I is any information set containing z. Thus, we can write the factor model for m as

$$m = f'E(ff' \mid I)^{-1} E(fm \mid I). \tag{2}$$

The pricing equation (23) implies

$$E(R^e \mid I) = \frac{1}{E(m \mid I)} - \frac{cov(m, R^e \mid I)}{E(m \mid I)}.$$

Substituting from (2), we obtain

$$E(R \mid I) - 1/E(m \mid I) =$$

$$\underbrace{cov(R,f' \mid I)\, E(ff' \mid I)^{-1}}_{\beta(I)} \quad \underbrace{(-1)\, E(fm \mid I)/E(m \mid I)}_{\lambda(I)}$$

If there is a conditionally risk free rate, it is $R^f(I) = 1/E(m \mid I)$. Otherwise, this is the expected conditional zero-beta rate.

Differencing the last equation for two assets, we can similarly write a traditional model for expected returns,

$$E(R^e \mid I) = cov(R^e,f' \mid I)\, E(ff' \mid I)^{-1} \quad (-1)\, E(fm \mid I)/E(m \mid I)$$

$$E(R^e \mid I) = \underbrace{\qquad}_{\beta(I)'} \quad \underbrace{\qquad}_{\lambda(I)}.$$

[11] One can rescale the factors arbitrarily. If factors f satisfy the model, i.e. if they satisfy

$$E(R^e) = -\operatorname{cov}(R^e f')E(ff')^{-1}E(mf)/E(m)$$

then factors Af, where A is a diagonal matrix, also satisfy the same restriction:

$$\operatorname{cov}(R^e f'A')E(Aff'A')^{-1}E(mAf)/E(m) =$$
$$\operatorname{cov}(R^e f')AA^{-1}E(ff')^{-1}A^{-1}AE(mf)/E(m).$$

[12] It may seem initially surprising that the b's come in pairs with one strongly positive and the other strongly negative, but this is the expected pattern. The discount factor m is proportional to the minimum *second-moment* return, which is on the lower portion of the mean-variance frontier. If the investment returns are on the upper portion of the mean-variance frontier, the discount factor m is expected to be very short one investment return and very long the other.

[13] The unconditional estimate of the scaled factor model is suppressed due to low degrees of freedom. Once each factor is scaled by each instrument, there are ten parameters and twelve moments. In all such estimates, the parameters were insignificant and the overidentifying restrictions not rejected.

[14]The figure shows results using first stage estimates because the objective of first stage estimation is precisely to make the points as close to the 45^{o} line as possible. The iterated GMM estimate has a different objective: it wants to minimize statistically informative *linear combinations* of the moments $0 = E(m\ R)$. Since the returns are highly correlated, the iterated estimation values differences and differences of differences (etc.) of moment conditions more than levels. Thus, it is happy to let the points drift away from the 45^{o} line in order to minimize these linear combinations of moments (to make the line joining them straighter). As mentioned in the text, the first-stage and iterated estimates and the pricing tests are usually very similar.

[15]One can equivalently specify the model with the excess market return as the factor, or with the market return and risk free rate or zero beta rate as factors. Note again that we expect $b<0$: the market is on the upper portion of the frontier, and m is proportional to a return on the lower portion. Alternatively, the requirement $1 = E(mR_m)$ implies $b<0$.

[16]Since two of the factors are bond returns, one should in general also check that the model prices these bond returns as well as the asset returns. Though I included the investment returns and value-weighted return in tests of the investment model and CAPM respectively, I follow Chen, Roll and Ross in ignoring this implication of their model.

[17]The unconditional estimate of the scaled CRR model is not identified, so the dots in Figure 6 refer to the non-scaled returns in the conditional estimate, unlike the other figures.

[18] I limit my comparisons to the standard time-separable CRRA formulation, which is about the same level of simplicity as the investment model. It is possible that one of the many variations on the consumption-based model, such as habit persistence, durability, etc. performs better. But it is also possible that one of the many possible variations on the investment model, such as production shocks, gestation lags, etc. performs better still.

The consumption-based model predicts that

$$m = \rho(c_t/c_{t-1})^{\gamma},$$

regardless of conditioning information, so there is no table for "scaled consumption-based model", and no scaled consumption factors in the comparison tables that follow.

[19] Hansen and Singleton (1982) rejected this model, and obtained much smaller estimates of the risk aversion coefficient. Hansen and Singleton used monthly data, fewer assets, and more instruments. When the term premium is added as an instrument, the consumption model quickly rejects with very small p-values. If one drops the wide cross-section of assets, a smaller estimate of γ is obtained.

[20] Large risk aversion coefficients mean that the resulting moments of m are driven by one or two data points, so the distribution theory may be badly approximated.

[21] I thank Campbell Harvey for suggesting this transformation.

Table 1. GMM estimates and tests of investment model

1. Non-scaled factor model $m = b_{NR}R^I_{NR} + b_R R^I_R$

	Parameter estimates				Tests		
	mpk_{NR}	mpk_R	b_{NR}	b_R		mpk=0	b=0	J_T
	A. Unconditional estimates; $0 = E(mR^e)$							
coef:	5.81	6.33	63.1	-52.7	χ^2	2580	7.74	6.94
t:	34.4	41.6	1.10	-2.78	DF	2	2	8
					%p	0.00	2.08	54.3
	B. Conditional estimates; $0 = E(m\, R^e \otimes z)$							
coef:	5.93	6.27	77.0	-66.7	χ^2	10364	142	72.2
t:	97	89	3.47	-9.78	DF	2	2	32
					%p	0.00	0.00	0.006

2. Scaled factor model, conditional estimates; $0 = E(m\, R^e \otimes z)$, $m = b'(R^I \otimes z)$

Parameter Estimates

	mpk_{NR}	mpk_R	b_{NR}	b_R	$b_{NR \cdot def}$	$b_{R \cdot def}$	$b_{NR \cdot d/p}$	$b_{R \cdot d/p}$
coef:	5.82	6.34	358	-349	-92.1	90.2	-200	197
t:	121	74.3	4.64	-4.72	-3.60	3.57	-4.88	4.82

Tests

	Joint b=0					
	all	non-scaled	scaled	NR	R	J_T
χ^2:	626	22.3	301	24.0	25.8	32.6
DF:	6	2	4	3	3	28
%p:	0.000	0.001	0.000	0.002	0.001	24.9

Returns R^e: 10 CRSP size deciles and 2 investment returns, less T-Bill rate

Investment returns R^I: functions of nonresidential (NR) and residential (R) gross fixed investment. (See equation (11).) Search over production function parameters mpk, α=3.0, δ= 0.05 throughout. 100xmpk is reported in the table, i.e. 5.61 for 0.0561

Instruments: constant, default premium (def), equally weighted dividend/price ratio (d/p).

%p gives the percent probability value, i.e. the percent probability that a statistic this high or higher is observed, given the null. (Rejection is a number less than 5% or 1%.). 4 Newey-West lags are used throughout.

Table 2

GMM test of CAPM

1. Non-scaled model; $m = b_{VW} R_{VW}$

Unconditional estimate; $0 = E(m\, R^e)$

Parameter estimate		b=0 test	J_T test
b: -5.54	χ^2:	15.1	5.30
t: -3.88	DF:	1	10
	%p:	0.010	87.0

Conditional estimate; $0 = E(m\, R^e \otimes z)$

Parameter estimate		b=0 test	J_T test
b: -7.70	χ^2:	73.9	48.5
t: -8.59	DF:	1	32
	%p:	0.00	3.11

2. Scaled model, conditional estimate;

$m = b'(R_{VW} \otimes z)$, $0 = E(m\, R^e \otimes z)$

Parameter estimates

	VW	VW•def	VW•d/p
b:	0.47	3.18	-11.7
t:	0.07	0.97	-2.56

Tests

	Joint b=0	J_T
χ^2:	124	45.7
DF:	3	30
%p:	0.000	3.29

Assets R^e: 10 CRSP size deciles plus value-weighted return; excess returns using t-bill rate.

Instruments: default premium, equally weighted dividend/price ratio.

4 Newey-West lags are used throughout.

Table 3.

Model comparison tests

Other Model		Joint b=0: Investment	Joint b=0: Others	"Likelihood ratio": Exclude investment?	"Likelihood ratio": Exclude others?
VW	χ^2	31	3.7	314	3.7
	DF	6	3	8	3
	%p	0.003	29	0.0	29
CRR	χ^2	17	10	37	10
	DF	6	15	8	15
	%p	0.8	79	0.001	79
consum-ption	χ^2	479	11	553	0.25
	DF	6	1	8	2
	%p	0.0	0.11	0.0	88

In each case, I fit an unrestricted model that includes both the investment factors and the other factors. All models except the consumption model include scaled factors, and are tested including returns scaled by instruments.

Joint b=0 tests whether each group of factors is significantly priced in the unrestricted estimate.

"Likelihood ratio" presents the increase in minimized JT from a model that excludes a given set of factors over the unrestricted model, using the weighting matrix from the unrestricted model.

%p gives the percent probability values: numbers less than 5.0 indicate rejection. "0.0" means "less than 0.000".

Table 4

GMM tests of Chen, Roll Ross factor model

1. Non-scaled factors; m = b'(CRR factors)

	Parameter estimates					Tests		
	MP	DEI	UI	UPR	UTS		Joint b=0	J_T
A. Unconditional estimates; $0 = E(m\ R^e)$								
b:	-1.71	-17.6	34.3	-50.5	-13.8	χ^2	8.79	1.01
t:	-0.29	-0.25	0.87	-1.59	-1.58	DF	5	5
						%p	11.8	96.2
B. Conditional estimates; $0 = E(m\ R^e \otimes z)$								
b:	-8.21	-49.9	11.3	-6.27	-11.2	χ^2	37.4	24.7
t:	-4.28	-2.17	1.29	-0.48	-3.66	DF	5	25
						%p	0.000	47.7

2. Scaled factors, Conditional Estimates; $m = b'(\text{factors} \otimes z)$; $0 = E(m\ R^e \otimes z)$

Parameter Estimates

	Not scaled					Scaled by default					Scaled by EW d/p				
	MP	DEI	UI	UPR	UTS	MP	DEI	UI	UPR	UTS	MP	DEI	UI	UPR	UTS
b:	-14	946	-779	371	113	12	-695	594	-59	-22	-3.7	-317	194	-360	-106
t:	-0.3	1.8	-2.6	1.9	1.4	0.4	-2.1	2.9	0.6	0.4	0.2	-1.2	1.6	-2.2	-1.5

Tests

	Joint b=0			J_T
	all	scaled	nonscaled	
χ^2:	54.2	32.4	10.1	6.62
DF:	15	10	5	15
%p:	0.00	0.03	7.2	96.8

<u>Assets</u> R^e: 10 CRSP size deciles

<u>Instruments</u> z: default premium, equally weighted dividend/price ratio.

4 Newey-West lags are used throughout.

Table 5.

GMM tests of consumption-based model $m = (c_{t+1}/c_t)^{\gamma}$

A. Unconditional estimate;

$0 = E(mR^e)$

Parameter Estimate		J_T test	
γ:	-102	χ^2:	5.87
t:	-2.92	DF:	9
		%p:	75.3

B. Conditional estimate;

$0 = E(m\ R^e \otimes z)$

Parameter Estimate		J_T test	
γ:	-154	χ^2:	37.8
t:	-6.53	DF:	29
		%p:	12.7

Assets R^e: 10 CRSP size deciles less T-bill rate.

Instruments z: default premium, equally weighted dividend/price ratio.

4 Newey-West lags are used throughout.

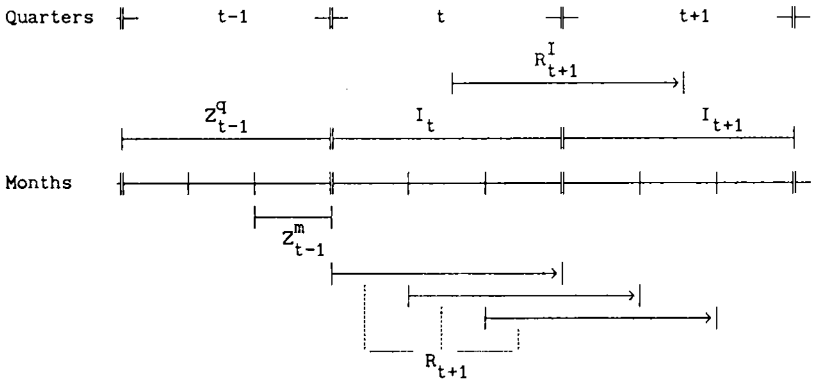

t = time in quarters

R^I = investment return

I = quarterly investment data

R = asset return data used in GMM tests

Z^m = monthly instruments

Z^q = quarterly instruments

Fig. A1. Timing relations among variables.

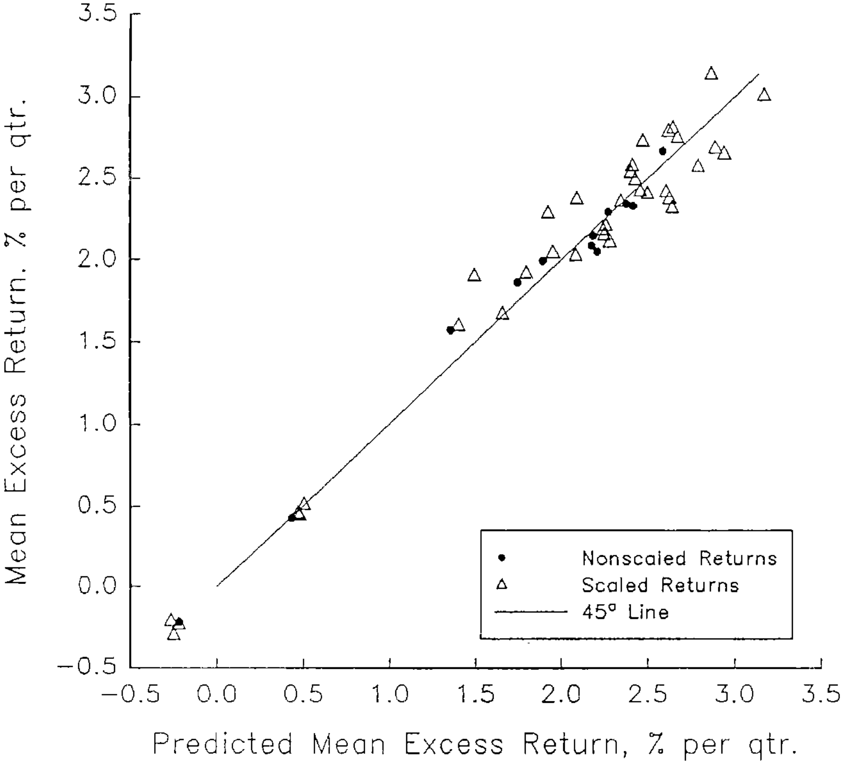

Figure 1. Non-scaled investment model.

Mean returns and scaled returns vs. predictions of two-factor investment model. The solid dots represent the non-scaled returns, estimated by themselves (first stage estimates, corresponding to panel 1A, Table 1). Each triangle corresponds to a size decile portfolio return or investment return, multiplied by a constant, default premium, or equally weighted dividend/price ratio (first stage conditional estimates, corresponding to panel 1B Table 1). The returns are the CRSP size decile portfolios plus two investment returns.

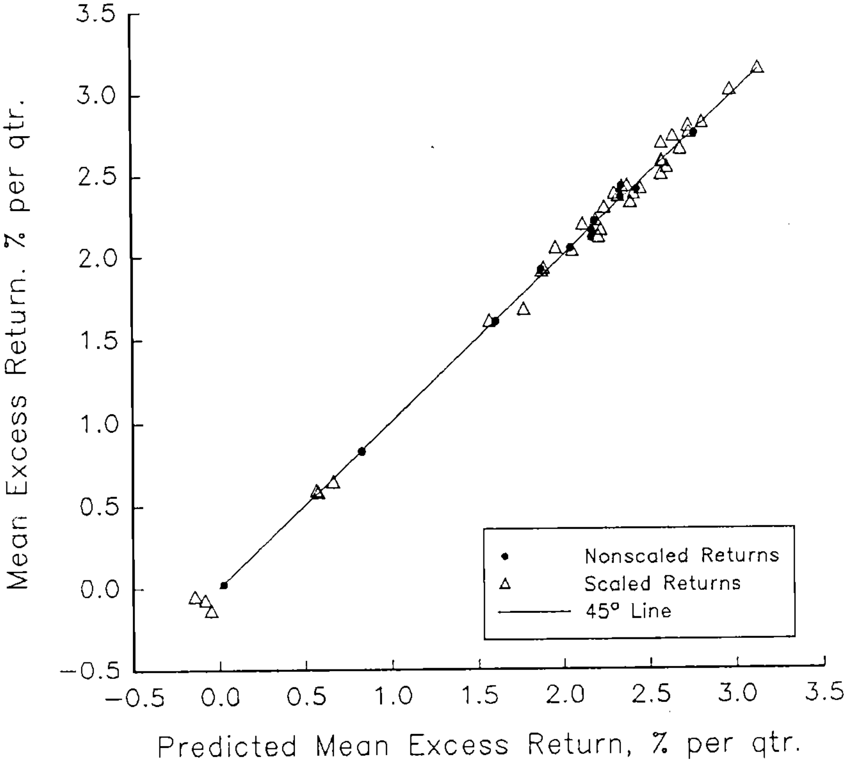

Figure 2. Scaled investment model

Mean returns and scaled returns vs. predictions of the scaled two factor investment model. CRSP size deciles and two investment returns.

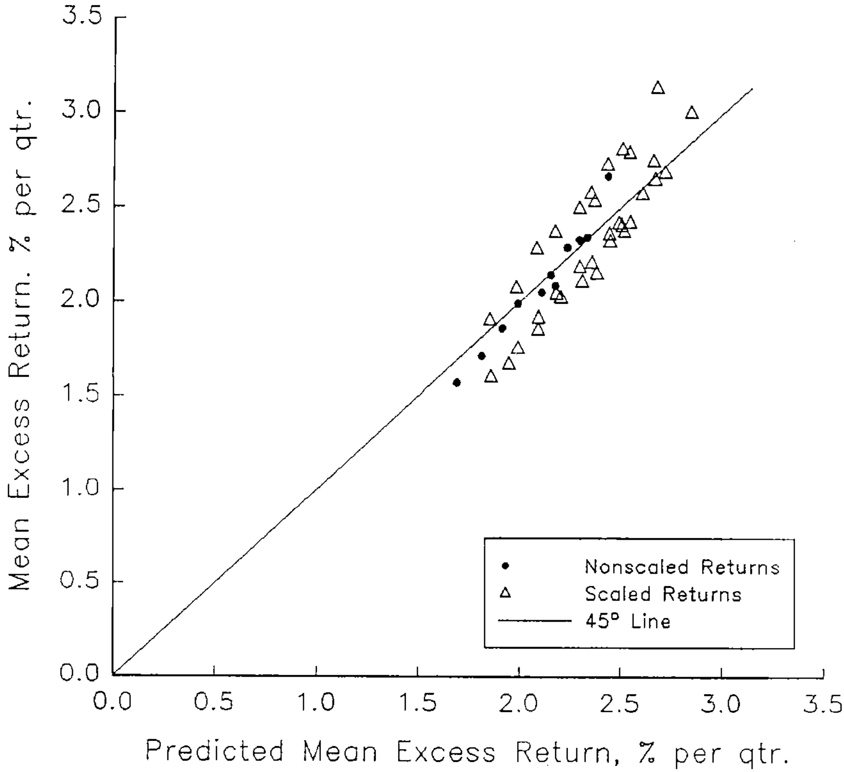

Figure 3. Non-scaled CAPM.

Mean returns and scaled returns vs. predictions of CAPM. CRSP size deciles plus value-weighted NYSE.

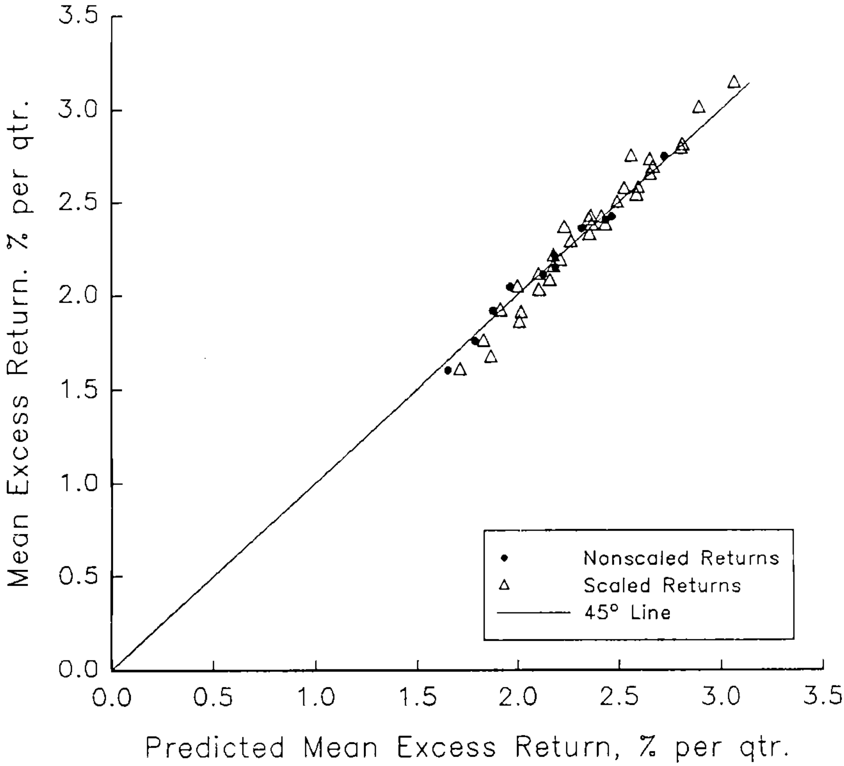

Figure 4. Scaled CAPM.

Mean returns and scaled returns vs. predictions of scaled CAPM. CRSP size deciles plus value-weighted NYSE.

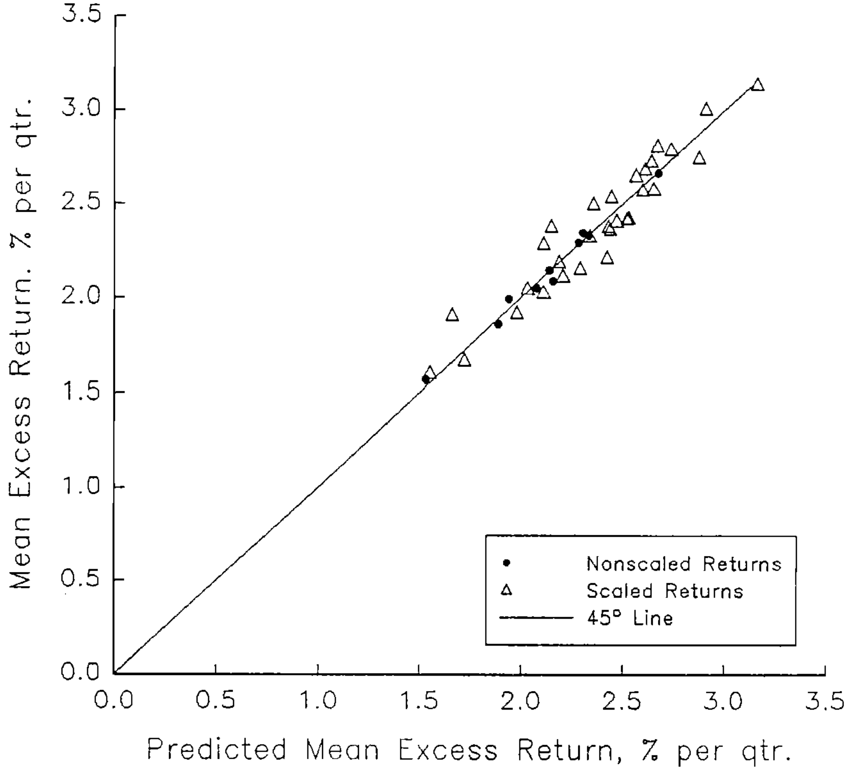

Figure 5. Non-scaled Chen Roll Ross model

Mean returns and scaled returns vs. predictions of Chen, Roll and Ross five-factor model. CRSP size decile portfolios.

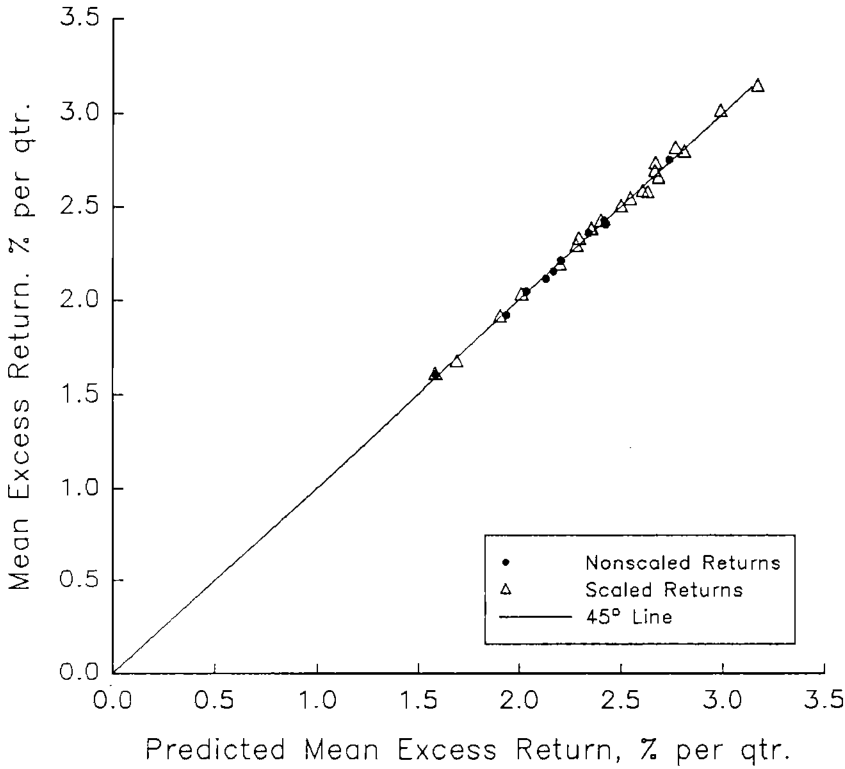

Figure 6. Scaled Chen, Roll Ross model.

Mean returns and scaled returns vs. predictions of scaled Chen, Roll and Ross five-factor model. CRSP size decile portfolios. (Dots are conditional estimates of non-scaled returns.)

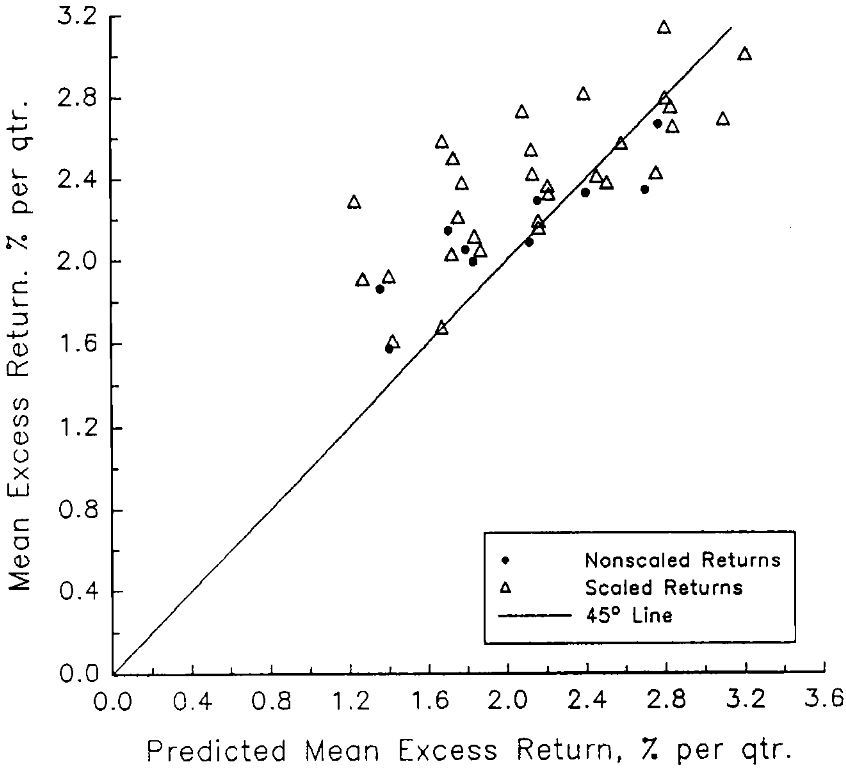

Figure 7. Consumption-based model.

Mean returns and scaled returns vs. predictions of the consumption-based model with CRRA utility. CRSP size decile portfolios.