Macroeconomic Dynamics, **1**, 1997, 551–567. Printed in the United States of America.

ARTICLES

TESTING THE CONSUMPTION CAPM WITH HEAVY-TAILED PRICING ERRORS

NARAYANA R. KOCHERLAKOTA
Federal Reserve Bank of Minneapolis

Many tests have rejected the implications of the consumption CAPM for data on U.S. asset returns. All of the tests, though, assume that the pricing errors satisfy the Central Limit Theorem. I provide empirical evidence that the marginal distributions of the pricing errors are so heavy-tailed that they do not satisfy the Central Limit Theorem. Using recent work on jackknifing, I construct a method of testing asset pricing models with heavy-tailed errors. Using this procedure, I find that the consumption CAPM is *not* rejected by annual U.S. data.

Keywords: Heavy Tails, Jackknife, Asset Pricing

1. INTRODUCTION

There have been many tests of the consumption Capital Asset Pricing Model (consumption CAPM) since its original development by Lucas (1978) and Breeden (1979). Most of the tests reject the model. For example, Mehra and Prescott (1985) argue that the sample mean returns on stocks and bonds are inconsistent with "reasonable" parameterizations of the model [see Kocherlakota (1996) for a survey of the related literature]. Hansen and Singleton (1982, 1983) present evidence that the conditional mean returns on stocks and bonds are inconsistent with any parameterization of the model.

Should we conclude from these studies that the consumption CAPM is false? The above rejections of the consumption CAPM are based, either implicitly or explicitly, on an application of the Central Limit Theorem. I provide empirical evidence that the pricing errors associated with the consumption CAPM are so heavy-tailed that they do not satisfy the Central Limit Theorem. In such a stochastic environment, standard tests of the model can lead to spurious rejections even asymptotically. I use recent work by Politis and Romano (1994) on jackknifing to

I presented earlier versions of this paper at the Wharton School, the University of Iowa, and the University of Houston; I thank the participants of those seminars for their comments. I also thank Krishna Athreya, Barbara McCutcheon, Joe Romano, Gene Savin, and Chuck Whiteman for their help. All remaining errors are my own. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Address correspondence to: Narayana R. Kocherlakota, Research Department, Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, Minneapolis, MN 55480, USA; e-mail: nrk@res.mpls.frb.fed.us.

© 1997 Cambridge University Press 1365-1005/97 $9.00 + .10

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

devise a method of testing the consumption CAPM that is valid for a wide class of pricing errors that do not satisfy the standard Central Limit Theorem. I apply this method to annual U.S. data from the period 1889–1987. I find that once one allows for the heavy tails, the model is not rejected by the data on consumption growth, stock returns, and Treasury-bill returns.

To better understand my argument, it is useful to consider what Weil (1989) terms the risk-free rate puzzle. Let R^f_{t+1} equal the gross real return to a one-year bond. According to the consumption CAPM, this return must satisfy the expression

$$1 = \beta E_t \left\{ (c_{t+1}/c_t)^{-\gamma} R^f_{t+1} \right\}, \tag{1}$$

where β is the discount factor of the representative agent, γ is the agent's coefficient of relative risk aversion, and c_t is period t consumption. Suppose, for example, that $\beta = 0.99$ and $\gamma = 2$. In U.S. annual data, the average value of the right-hand side is equal to 0.971. The fact that this number is so much lower than 1 generally is regarded as strong evidence against the consumption CAPM.

Of course, the difference in the sample means of the two sides of (1) might be due solely to sampling error and not to actual differences in the expectations of the two sides. The standard approach to checking this possibility is to calculate the standard deviation of the pricing error,

$$u_t = \beta \left\{ (c_{t+1}/c_t)^{-\gamma} R^f_{t+1} \right\} - 1,$$

and use a t-test to evaluate the hypothesis that u_t has expectation equal to zero. In annual U.S. data, the t-statistic[1] associated with the sample mean of u_t equals about -3.11. Thus, the left-hand and right-hand sides of (1) appear to be *significantly* different from zero (at least in a statistical sense); it appears that bonds are mispriced by the consumption CAPM.

The purpose of this paper is to argue that the marginal distributions of pricing errors such as u_t have tails so thick that the Central Limit Theorem does not apply to them. Hence, the t-statistic described in the preceding paragraph will not have a normal distribution asymptotically. Standard notions of statistical significance are no longer relevant. Instead, an alternative approach, such as the one described in this paper, is needed to assess the null hypothesis that u_t has mean zero.

Why do the pricing errors associated with the consumption CAPM have such heavy tails? The pricing errors are essentially the product of asset returns and the representative investor's intertemporal marginal rates of substitution (IMRS). Hence, one explanation for the pricing errors' having such heavy tails is that asset returns are generated by a symmetric stable distribution with infinite variance, as argued by Mandelbrot (1963), Fama (1965), and McCulloch (1978), among others. I show that in annual data, there is little evidence to support this explanation. Instead, it appears that the distribution of the representative investor's IMRS, $\beta(c_{t+1}/c_t)^{-\gamma}$, has a right tail heavy enough that it does not satisfy the Central Limit Theorem.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

To see why this might happen, suppose that the distribution of consumption is not bounded away from zero: There exist states of the world with positive probability in which consumption growth is extremely low. If the investor is highly risk averse, the representative investor's IMRS is extremely high in those states [this phenomenon is also important in Rietz's (1988) explanation of the equity premium puzzle]. Thus, if the distribution of consumption growth has a sufficiently thick left tail and/or the investor is sufficiently risk averse, then the representative investor's IMRS is sufficiently heavy-tailed that the Central Limit Theorem will not apply to u_t.

The rest of the paper is organized as follows. Section 2 demonstrates that if it is assumed that the pricing errors associated with the consumption CAPM have finite variance, then the sample means of the pricing errors are significantly different from zero. Section 3 argues that it is theoretically and empirically plausible that the pricing errors associated with the consumption CAPM are heavy-tailed. Section 4 discusses how to construct an asymptotically valid test in this situation. Section 5 applies the test to U.S. annual data and fails to reject the consumption CAPM. Section 6 concludes.

2. THE PROBLEM

Most asset pricing models (such as the Capital Asset Pricing Model or the consumption CAPM) take the form of specifying some observable variable as a pricing kernel (alternatively termed a pricing measure or a benchmark payoff). Cross-sectional or intertemporal differences in expected returns are then explained via covariances with this observable variable. More specifically, a typical asset pricing model specifies a strictly positive stochastic process m_t such that

$$1 = E_t m_{t+1} R_{t+1}, \tag{2}$$

where R_{t+1} is the gross return to any traded asset. The empirical researcher in asset pricing wants to know how to test this null hypothesis, given a finite number of observations on m_t and R_t. A useful approach developed by Hansen and Jagannathan (1991) starts by constructing the *pricing error* u_t by using the formula

$$u_t = m_t R_t - 1.$$

Then, under the null hypothesis (2), u_t is a *martingale difference*: Its expectation, conditional on information available at time $(t-1)$, is zero.

Billingsley (1961) has proved that, as long as u_t is a stationary ergodic martingale difference with finite second moment, the statistic

$$J_T = T^{1/2}\left(\frac{1}{T}\sum u_t\right) \Big/ \left(\frac{1}{T}\sum u_t^2\right)^{1/2} \tag{3}$$

converges in distribution to $N(0, 1)$. This result suggests the use of the method of moments to test the consumption CAPM [as has been done by Hansen and

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

TABLE 1. Method-of-moments tests of expected pricing error equaling zero using stock returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	0.992	0.998	1.000	1.000
1	0.924	0.977	0.995	0.999
2	0.690	0.853	0.944	0.983
3	0.377	0.583	0.764	0.889
4	0.173	0.322	0.504	0.682
5	0.081	0.169	0.298	0.457
6	0.045	0.096	0.180	0.297
7	0.030	0.064	0.121	0.204
8	0.024	0.049	0.091	0.154
9	0.023	0.044	0.078	0.129
10	0.024	0.044	0.075	0.118

[a] The table calculates the probability of J_T being less than or equal to its observed sample value according to the $N(0, 1)$ distribution. It uses various preference parameters and stock returns to construct the realized pricing errors. Entries less than 0.025 or greater than 0.975 represent statistically significant evidence against the model.

Singleton (1982)]. In particular, suppose there exists a representative investor who has a power utility function,

$$E_0 \sum_{s=0}^{\infty} \beta^s (c_{t+s})^{1-\gamma} \Big/ (1-\gamma).$$

Then, the first-order conditions of this representative investor imply that the stochastic processes

$$u^s_{t+1} = \beta (c_{t+1}/c_t)^{-\gamma} R^s_{t+1} - 1,$$
$$u^b_{t+1} = \beta (c_{t+1}/c_t)^{-\gamma} R^b_{t+1} - 1$$

are martingale differences and should satisfy (3). In these formulas, c_t is aggregate real consumption in period t, R^s_t is the real gross return to the S&P 500, and R^b_t is the real gross return to the three-month Treasury bill.

Tables 1 and 2 present the results of using the method of moments to test the consumption CAPM using U.S. annual data from 1889 to 1987 on these three variables [the data are an extended version of Mehra and Prescott's (1985) data; for more details, see the Appendix]. To construct the tables, I allow β to be an element of the set[2] {0.97, 0.98, 0.99, 1} and γ to assume integer values between 0 and 10. All of these specifications of the model are rejected at the 5% level by the data on at least one of the pricing errors, especially by the data on the return to the nominally risk-free asset. These findings are consistent with those of Mehra and Prescott (1985) and Hansen and Singleton (1982, 1983), among others. They are

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

TABLE 2. Method-of-moments tests of expected pricing error equaling zero using Treasury-bill returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	0.000	0.077	0.671	0.989
1	0.000	0.000	0.019	0.275
2	0.000	0.000	0.001	0.020
3	0.000	0.000	0.000	0.004
4	0.000	0.000	0.000	0.002
5	0.000	0.000	0.000	0.002
6	0.000	0.000	0.000	0.002
7	0.000	0.000	0.000	0.003
8	0.000	0.000	0.001	0.004
9	0.000	0.000	0.002	0.006
10	0.001	0.002	0.004	0.009

[a] The table calculates the probability of J_T being less than or equal to its observed sample value according to the $N(0, 1)$ distribution. It uses various preference parameters and Treasury-bill returns to construct the realized pricing errors. Entries less than 0.025 or greater than 0.975 represent statistically significant evidence against the model.

commonly interpreted as implying that these parameterizations of the consumption CAPM are not statistically valid models of the comovements of asset returns and consumption.

3. HEAVY-TAILED PRICING ERRORS

In what follows, the following restriction on the shape of the distribution of u_t will be useful.

DEFINITION 1. *Suppose X is a random variable with distribution function F. If X is bounded from below, the distribution function F is Pareto-Levy in the tail (PLIT) with maximal-moment exponent α if there exists a positive constant C and a function δ such that*

$$[1 - F(x)] = C^{\alpha} x^{-\alpha}[1 + \delta(x)],$$

where $\delta(x) \to 0$ as $x \to \infty$.

Stochastic processes with PLIT distributions are convenient because there exists a method of normalizing partial sums of the elements of the stochastic process so that the normalized partial sums have a limiting distribution. (More technically, PLIT stochastic processes lie in the domain of attraction of stable distributions.) Note that a random variable with a PLIT distribution function will have finite mean as long as its maximal moment exponent α is larger than 1; however, the distribution of $T^{-1/2} \sum_{t=1}^{T} u_t$ will not converge to a normal if α is less than 2.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

3.1. Theoretical Plausibility of Heavy Tails

The implication derived by Billingsley (1961) that J_T converges in distribution to $N(0, 1)$ is based on the assumption that the tails of the density of u_t are not too heavy to prohibit application of the Central Limit Theorem. *This assumption has no theoretical justification.* To see this more clearly, it is useful to consider the following (static) example.

Example 1

Consider a two-period economy in which there are J agents with identical preferences represented by the utility function $(c_1^{1-\gamma} + c_2^{1-\gamma})/(1-\gamma)$, $\gamma > 0$. There is a single good. The aggregate endowment in period 1 is equal to 1; the aggregate endowment in period 2, X_2, is distributed over the interval $[0, 1]$ according to the probability density function $x^\eta(1+\eta)$, $\eta > 0$.

The size of the aggregate endowment in period 2 is the only source of risk in the economy. The agents trade two assets. The first asset is available in zero net supply and pays off one unit of consumption in every state of the world in period 2. The second asset pays off X_2 when the realization of the aggregate endowment is X_2. The J agents are each initially endowed with $(1/J)$ units of consumption in period 1 and $(1/J)$ shares of the second asset. I assume that η and γ satisfy the following condition:

$$(1 + \eta - \gamma) > 0. \tag{4a}$$

This condition is sufficient to guarantee that the typical individual receives a finite amount of utility from his endowment.

In equilibrium, there is no trade in this economy (preferences and endowments are the same). However, there are equilibrium prices. The pricing kernel m is given by the marginal rate of substitution of the representative agent. Hence, $m = X_2^{-\gamma}$. I assume that

$$(3 - 2\gamma + \eta) < 0. \tag{4b}$$

Condition (4b) is sufficient to guarantee that the pricing kernel m has infinite variance.

In the autarkic equilibrium, the prices of the two assets (in terms of period-1 consumption) are given by

$$p_1 = \int_0^1 x^{-\gamma} x^\eta (1+\eta)\, dx = (1+\eta)/(1+\eta-\gamma),$$

$$p_2 = \int_0^1 x^{-\gamma} x^\eta (1+\eta) x\, dx = (1+\eta)/(2+\eta-\gamma).$$

The condition that $(1 + \eta - \gamma) > 0$ is sufficient to guarantee that both asset prices are finite.

The (gross) return to the first asset is given by $1/p_1$ whereas the return to the second asset is given by X_2/p_2. It is easy to see that the asset returns have finite variance (because they are both bounded from above).

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

The pricing errors associated with the two assets are given by

$$u_1 = x^{-\gamma}/p_1 - 1,$$
$$u_2 = x^{1-\gamma}/p_2 - 1.$$

These representations imply that the pricing errors u_1 and u_2 have PLIT distributions with mean zero. The maximal exponent of the bond pricing error is $(1+\eta)/\gamma$ whereas the maximal moment exponent of the stock pricing error is $(1+\eta)/(\gamma-1)$. If conditions (4a) and (4b) are satisfied, then both of these maximal exponents lie in the interval (1, 2). In general, the stock pricing error will have thinner tails than the bond pricing error. This is because stock returns tend to be low when the IMRS of the representative investor is high. ■

In the example, asset returns have finite variance but pricing errors do not. The heavy tail in the density of the pricing kernel is responsible for the heavy tail of the distribution of the pricing errors. In particular, the bond pricing error does not satisfy the Central Limit Theorem as long as $(1 - 2\gamma + \eta) < 0$. This condition is satisfied if η is small or if γ is large. Smaller values of η imply that low consumption growth rates are more likely. Larger values of γ imply that consumers are more risk averse and hence put more weight on the low-consumption-growth states. The example demonstrates that the Central Limit Theorem will not apply to pricing errors in economies in which asset prices are largely determined by the left tail of the distribution of consumption growth.

3.2. Empirical Plausibility

Example 1 shows that as a theoretical matter, it is entirely possible that a pricing error u_t has zero mean (as the consumption CAPM implies) but does not obey the standard Central Limit Theorem. If this is so, use of the method of moments may lead to false rejections. We must verify *empirically* whether the assumption that u_t has finite second moment is correct. To do so, we need to estimate the thickness of the right-hand tail[3] of the distribution of the pricing error u_t (note that limited liability implies that u_t is bounded from below by -1).

Loretan and Phillips (1994) describe how to obtain consistent estimates of the maximal exponent α for martingale differences with PLIT distributions. Suppose we have T observations of u_t. Order the observations in an increasing fashion and let $u_{T,n}$ denote the ordered observations. Then, define

$$\hat{\alpha}_{T,s} = \left\{ s^{-1} \sum_{j=1}^{s} \left(\ln \frac{u_{T,T-j+1}}{u_{T,T-s}} \right) \right\}.$$

This sample statistic is a consistent estimator of α if s goes to infinity with T sufficiently slowly. Essentially, this estimator assumes that the s largest observations are drawn from the Pareto-Levy tails of the distribution. Because we know the shapes of the tails, we can use the largest s observations to estimate the thickness of the tails.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

Table 3. Estimates of maximal exponent of pricing-error using stock returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	3.688	3.800	3.909	4.016
1	2.420	2.503	2.583	2.662
2	1.811	1.878	1.944	2.008
3	2.295	2.370	2.443	2.514
4	1.363	1.412	1.461	1.508
5	1.842	1.896	1.948	2.000
6	1.686	1.731	1.777	1.821
7	1.724	1.766	1.807	1.846
8	1.451	1.485	1.519	1.552
9	1.465	1.496	1.526	1.555
10	1.235	1.261	1.286	1.311

[a] The table contains estimates of the maximal-moment exponent of the pricing error associated with stock returns for various values of the preference parameters. Estimates of the asymptotic standard errors can be calculated by dividing the estimate by sqrt(s) where, in this case, $s = 10$; see Loretan and Phillips (1994).

Earlier, we used annual data to reject the null hypothesis that the pricing errors u_t^s and u_t^b have population mean zero. Could this result be due to these pricing errors not satisfying the Central Limit Theorem? To check this possibility, I assume that the pricing errors both have PLIT distributions. I then estimate the maximal-moment exponent α for each specification of the preference parameters (letting $s = 10$; I obtain similar results for $s = 15$). The results are described in Tables 3 and 4. I find that the estimates of $\hat{\alpha}$ are less than 2 for most of the preference specifications, especially when the CRRA is larger than 3. (Note though that net asset returns do not exhibit heavy tails.) If this property of the sample reflects the nature of the population, then the tails of the population are too heavy to warrant application of the Central Limit Theorem.[4]

The estimation results are supported by Figures 1–4. The graphs set $\beta = 0.99$ and depict the tails of the empirical distribution of the pricing errors for two values of γ. The graphs also compare the tails to what they would look like if the pricing errors were normally distributed and to what they would look like if they had the PLIT distributions estimated in Tables 3 and 4. For all of the pricing errors, the tails are heavier than normal; the difference is especially pronounced in the case of the Treasury bill. The estimated PLIT distributions appear to provide a good approximation to the tails of the pricing errors.[5]

The analysis in this section is not definitive: It is certainly difficult to precisely estimate the tail behavior of distributions based on the 10 most extreme observations. It seems reasonable to conclude though that there is not much evidence in the data against the hypothesis that the pricing errors have maximal exponents less

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

TABLE 4. Estimates of maximal exponent of pricing error using Treasury-bill returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	1.448	1.722	1.983	2.233
1	1.431	1.628	1.816	1.998
2	1.144	1.269	1.388	1.503
3	1.444	1.543	1.638	1.731
4	1.481	1.558	1.632	1.704
5	1.498	1.560	1.620	1.679
6	1.503	1.555	1.604	1.653
7	1.500	1.543	1.585	1.626
8	1.491	1.528	1.564	1.599
9	1.462	1.493	1.524	1.554
10	1.426	1.453	1.480	1.506

[a]The table contains estimates of the maximal-moment exponent of the pricing error associated with Treasury-bill returns for various values of the preference parameters. Estimates of the asymptotic standard errors can be calculated by dividing the estimate by sqrt(s) where, in this case, $s = 10$; see Loretan and Phillips (1994).

than 2. [This is especially true given McCulloch's (1997) evidence that, in small samples, estimators of α such as the one used in this paper are likely to be highly biased *upward*.] Even less strongly, these results indicate that it may be desirable to construct tests of the consumption CAPM for pricing errors that do not satisfy the Central Limit Theorem.

4. METHOD OF ATTACK

The preceding section leaves me with the following question. I have T observations of a stochastic process $\{u_t\}_{t=1}^{T}$. My null hypothesis is that u_t has zero mean. How can I test the null hypothesis when the tails of u_t are so thick that it does not satisfy the Central Limit Theorem? The following theorem provides a starting point.

THEOREM 1. *Suppose $\{u_t\}_{t=1}^{\infty}$ is i.i.d. with mean zero and with a PLIT marginal distribution F that has maximal-moment exponent $\alpha \in (1, 2)$. Then, $H_T = (\sum_{t=1}^{T} u_t)/\max(\{u_t\}_{t=1}^{T})$ converges in distribution.*

Proof. Gnedenko and Kolmogorov (1954, p. 175) show that if u_t is PLIT, then $(\sum_{t=1}^{T} u_t)/B(T)$ converges in distribution to a stable distribution with characteristic exponent α, where $B(T)$ is chosen so that $Pr[u_t > B(T)] = T^{-1}$. Athreya (1987) shows in the proof of Theorem B of his Appendix that $\{\max(\{u_t\}_{t=1}^{T})\}^{-1} B(T)$ converges in distribution to a constant. The theorem follows. ■

This theorem suggests the use of H_T as a test statistic.[6] Unfortunately, it is hard to figure out the limiting distribution of H_T. Even if I could, it is then nontrivial

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

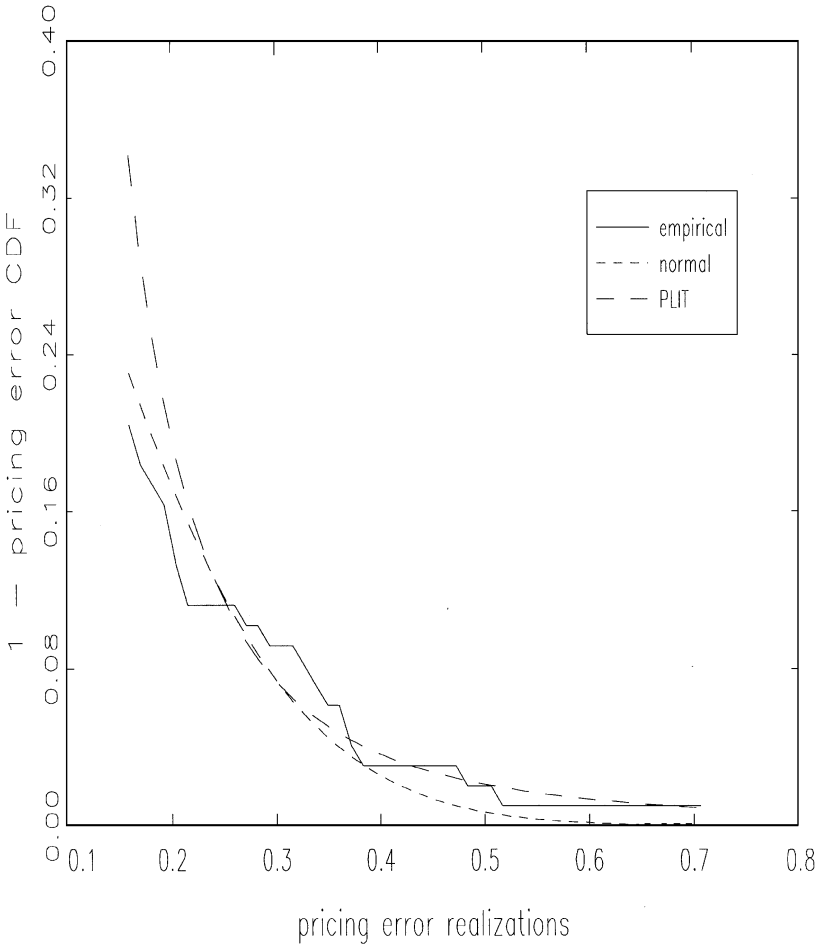


FIGURE 1. Estimates of right tail of stock pricing error CDF (CRRA = 3).

to figure out the critical values associated with that distribution. Thus, how can I find out the (asymptotic) critical values for the test statistic H_T?

The answer to this question lies in some recent work of Politis and Romano (1994). Consider any sample statistic with a limiting distribution. They show that under extremely weak conditions, jackknifing[7] provides an asymptotically valid approximation to the finite sample distribution of the sample statistic.

Here is the explicit method suggested by the work of Politis and Romano. Suppose we have a sample $\{u_t\}_{t=1}^{T}$. We want to test the null hypothesis that u_t has unconditional population mean zero. To do so, let $M(T)$ be some function of the sample size T such that $M(T)/T$ converges to zero and $M(T)$ converges to infinity as T grows large. Typically, I set $M(T)$ equal to $T^{1/2}$.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

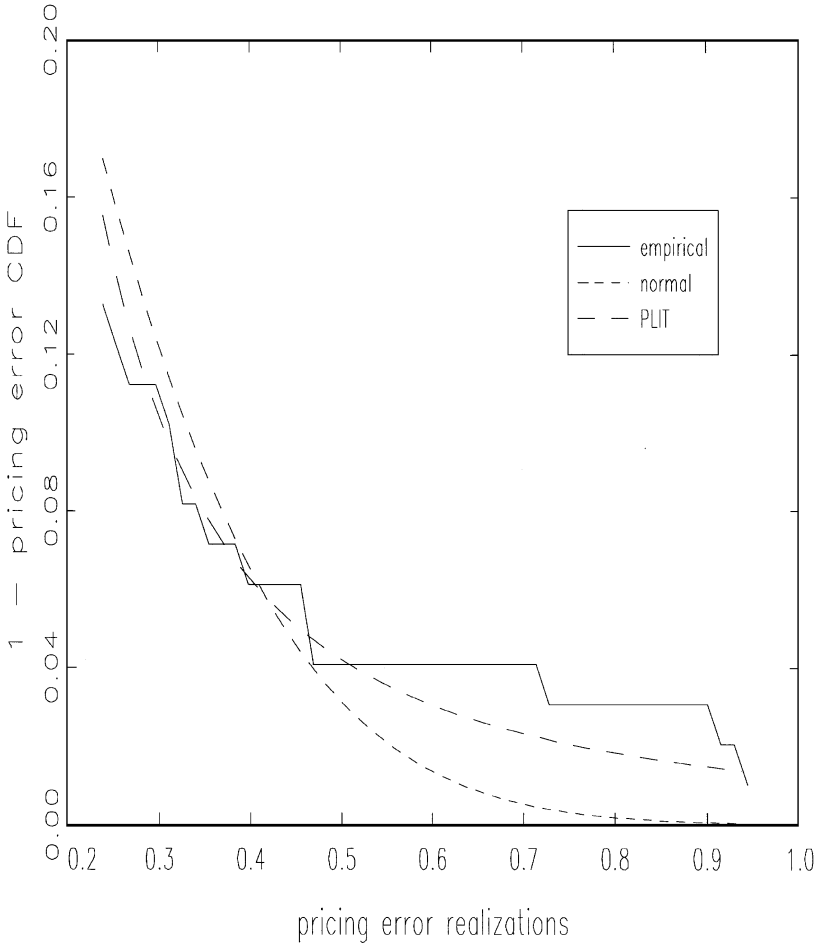


FIGURE 2. Estimates of right tail of stock pricing error CDF (CRRA = 7).

(1) Find the sample mean $\bar{u}$ and $H_T = T\bar{u}/\max(\{u_t\}_{t=1}^{T})$.
(2) Construct the $[T - M(T) + 1]$ subsamples $u^i = (u_i, \ldots, u_{i+M(T)-1})$. Let $I = T + M(T) - 1$.
(3) For each subsample i, define $H_T^i = [M(T)\bar{u}^i - M(T)\bar{u}]/\max(u^i)$.
(4) Treat the I H_T^i's as the true distribution of H_T; in other words, reject the null hypothesis if

$$(1/\mathrm{I}) \sum_{i=1}^{I} 1\{\mathrm{H}_\mathrm{T}^\mathrm{i} \geq \mathrm{H}_\mathrm{T}\}$$

is either less than $p/2$ or greater than $(1 - p/2)$.

Politis and Romano (1994) prove that if $E(u_t) = 0$, then the probability of rejection using this testing procedure converges to p as long as the stochastic process

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

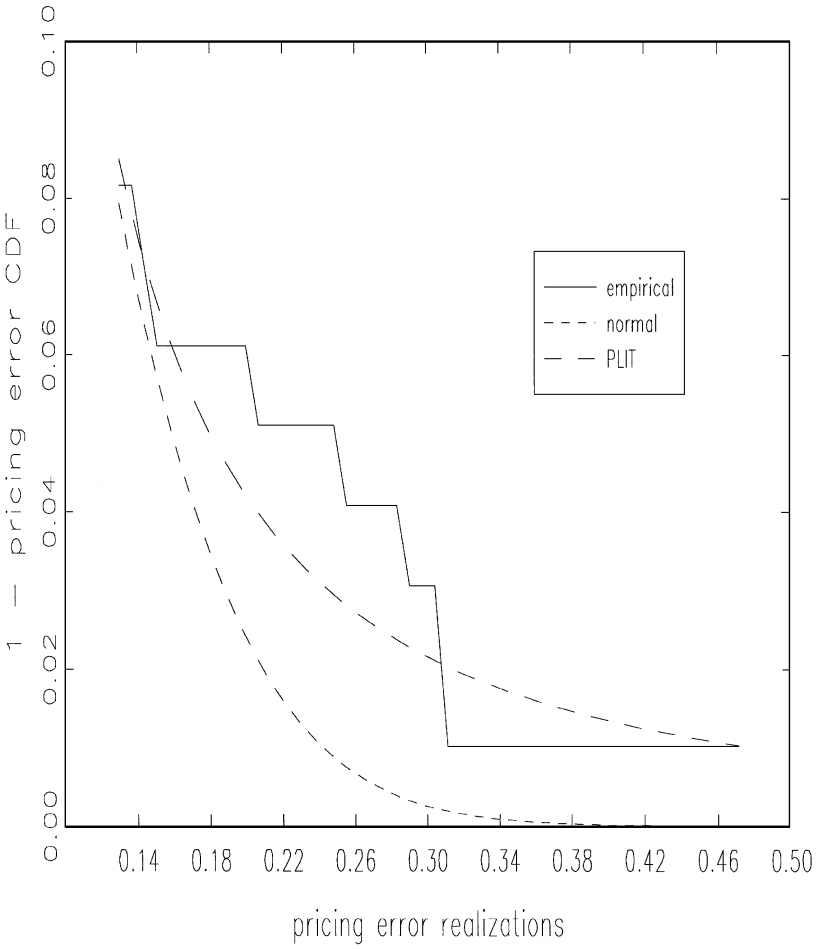


FIGURE 3. Estimates of right tail of T-bill pricing error CDF (CRRA = 3).

generating u_t is stationary and strongly mixing. In this sense, the procedure is asymptotically valid.

Previous users of the jackknife may be surprised at the use of subsamples of *consecutive* data to construct the approximate distribution of H_T (step 2, above). Using consecutive data allows the method to capture the effects of fairly arbitrary temporal dependence in the data. If the data were known to be i.i.d., then we could simply draw arbitrary subsamples of size $M(T)$ without replacement to construct the blocks.

5. TESTING THE CONSUMPTION CAPM IN THE HEAVY-TAILED CASE

In this section, I apply the jackknife method described above to actual data. As before, I look at annual data from 1889 to 1987 on real annual consumption, c_t; the

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

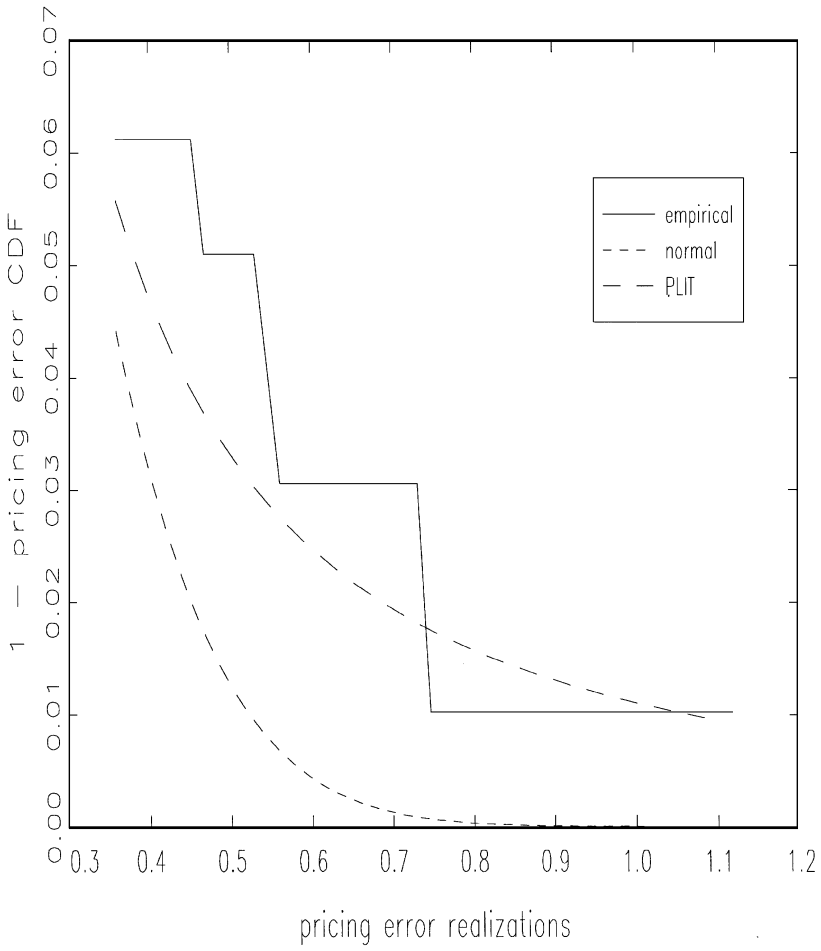


FIGURE 4. Estimates of right tail of T-bill pricing error CDF (CRRA = 7).

real rate of return on the S&P 500, R_t^s; and the real rate of return on the three-month Treasury bill, R_t^b. Define

$$u_t^s(\beta, \gamma) = \beta(c_t/c_{t-1})^{-\gamma} R_t^s - 1,$$

$$u_t^b(\beta, \gamma) = \beta(c_t/c_{t-1})^{-\gamma} R_t^b - 1.$$

If there is a representative investor with power utility, then there exists (β, γ) such that the unconditional population mean of both quantities is zero.

I use the jackknife procedure to test the null hypothesis that the pricing errors have mean zero at the 5% level for $\beta \in \{0.97, 0.98, 0.99, 1\}$ and $\gamma \in \{0, 1, \ldots, 10\}$. Tables 5 and 6 contain the jackknifed p-value of the test statistic H_T. They should be compared to Tables 1 and 2, which implicitly argue that the data are inconsistent

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

Table 5. Tests of consumption CAPM allowing for heavy-tailed errors using stock returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	1.000	1.000	1.000	1.000
1	1.000	1.000	1.000	1.000
2	0.809	1.000	1.000	1.000
3	0.461	1.000	0.955	1.000
4	0.236	0.663	0.618	0.798
5	0.180	0.371	0.337	0.562
6	0.157	0.213	0.213	0.326
7	0.157	0.202	0.213	0.248
8	0.157	0.202	0.213	0.225
9	0.180	0.191	0.191	0.225
10	0.157	0.191	0.225	0.236

[a] The table describes the approximate probability that H_T is less than the value it assumes in U.S. annual data. The distribution of H_T is approximated using the jackknife procedure described in the text. Entries less than 0.025 or greater than 0.975 represent significant evidence against the model.

Table 6. Tests of consumption CAPM allowing for heavy-tailed errors using Treasury-bill returns[a]

	β			
γ	0.97	0.98	0.99	1.00
0	0.011	0.213	0.607	1.000
1	0.079	0.056	0.191	0.494
2	0.079	0.079	0.191	0.191
3	0.079	0.191	0.157	0.169
4	0.067	0.169	0.157	0.169
5	0.067	0.146	0.157	0.169
6	0.169	0.146	0.169	0.180
7	0.236	0.191	0.202	0.180
8	0.236	0.202	0.225	0.213
9	0.247	0.225	0.225	0.225
10	0.247	0.236	0.225	0.225

[a] See footnote of Table 5.

with the consumption CAPM. In contrast, Tables 5 and 6 document that there is little evidence against the consumption CAPM once we assume that the pricing errors are sufficiently heavy-tailed that they do not satisfy the Central Limit Theorem.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

6. CONCLUSION

This paper argues that the right tail of the distributions of the pricing errors associated with the consumption CAPM are so heavy as to preclude valid application of the Central Limit Theorem. The paper provides a test of whether the errors have mean zero that is valid as long as they lie in a particular large class of stochastic processes that do not satisfy the Central Limit Theorem. It shows that there is little evidence against the consumption CAPM when this alternative procedure is applied to annual data.

I should emphasize the importance of the sample period to my results. Annual consumption data include the Great Depression whereas available quarterly or monthly series do not. Because consumption growth is low during the Depression, the marginal rate of substitution of the representative investor is high, especially when the investor is highly risk averse. Including the Depression thus serves to thicken the right tail of the distribution of the pricing errors. Judging from the popularity of doomsday investment books, investors today certainly seem to believe that Depression-like falls in consumption are possible. If this is true, data on consumption growth that is exclusively from the "good times" since World War II delivers an inadequate picture of the true IMRS of the representative investor.

Mehra and Prescott (1985) argue that consumption growth is insufficiently variable to explain the wide differential between average stock returns and average bond returns in the United States. This paper attributes their finding to an insufficiently rich model of the consumption growth process. Mehra and Prescott assume that annual consumption growth is bounded from below by -1.8% (ignoring the realizations of consumption during the Great Depression). This paper instead assumes that there is no lower bound on the possible realizations of consumption growth; although they may be unlikely or hitherto unobserved, extremely low consumption growth rates are possible. This model of consumption growth implies that bonds are more valuable than in Mehra and Prescott's model because risk-averse agents want to buy bonds as a hedge against the states of the world in which consumption growth is exceedingly low.[8]

Some readers might criticize the model of consumption growth described in the above paragraph as being unrealistic because the realizations of annual consumption growth in the United States are bounded below by -10%. Example 1 demonstrates, though, that the heaviness of the tails of the pricing errors is determined not by the *magnitude* of the probability of consumption growth being below some ε, but by the *derivative* of this probability with respect to ε. Thus, the consumption CAPM pricing errors may have heavy tails even if the probability of consumption growth falling below -10% is very low.

There is a definite connection between this paper and the work of Hansen and Jagannathan (1991). They show that the population standard deviation and mean of the pricing kernel must lie above a bound that can be characterized in terms of the joint distribution of asset returns. This requirement is trivially satisfied by any

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

asset pricing model in which the second moment of the pricing kernel is infinite, including models that are inconsistent with the data. Hence, in contrast to the method described in this paper, the Hansen–Jagannathan approach cannot be used to falsify any asset pricing model that has a pricing kernel with an infinite variance.[9]

The paper provides some empirical evidence in support of the consumption CAPM. More important, though, it constructs a simple procedure to test arbitrary asset pricing models in situations in which the pricing errors are heavy-tailed. The current procedure only exploits information in the marginal distributions of pricing errors. Note that this limited amount of information is sufficient to provide significant evidence against reasonable parameterizations of the consumption CAPM if we maintain the assumption that the maximal exponent is larger than 2 (see Tables 1 and 2). In principle, it should be possible to construct more powerful tests by using sample information about the joint distribution of the pricing errors. Practical statistical techniques to handle this additional information are not available at this time.

NOTES

1. It is easy to show that u_t is not autocorrelated under the null hypothesis that the consumption CAPM is true, and so, it is straightforward to obtain a consistent estimate of the standard deviation of u_t.

2. Elsewhere [Kocherlakota (1990)], I have shown that using higher values of β and γ can resolve the equity premium puzzle implicit in Tables 1 and 2. However, I recognize that not all members of the financial economics community view these higher values as economically plausible. For the purposes of this paper, I adopt Mehra and Prescott's (1985) definition of economically plausible.

3. When a stochastic process has infinite variance, its sample variance also has infinite variance. Hence, it is not possible, using sample estimates, to try to gauge whether the population variance is close to infinite.

4. I have constructed versions of Tables 3 and 4 for the demeaned pricing errors $(u_t - 1/T \sum_{t=1}^{T} u_t)$ as well; the estimates of α are generally higher, but mostly less than 2, especially for high degrees of risk aversion.

5. To calculate the estimated PLIT distribution, I assume that $\delta(x) = 0$ for the s largest sample points and use the following estimate of the scaling coefficient C:

$$\hat{C} = (s/T)^{1/\hat{\alpha}}(u_{T,T-s}),$$

where δ and C are defined as in Definition 1.

6. Although the theorem is proved only for PLIT i.i.d. sequences $\{u_t\}_{t=1}^{\infty}$ (and not to all martingale differences), I assume that it applies to the pricing errors.

7. "Jackknifing" is a generic name given to the use of subsamples that are drawn *without* replacement to construct an approximation to the finite sample distribution of a sample statistic. In bootstrapping, the subsamples are drawn *with* replacement. Athreya (1987) suggests a method of obtaining a bootstrapped approximation to the distribution of H_T when u_t is i.i.d. with a stable distribution.

8. This explanation of the equity premium puzzle is essentially the same as that of Rietz (1988), but note that Rietz's model does not explain the rejections described in Tables 1 and 2 because the Central Limit Theorem applies to the pricing errors in his model.

9. The preceding paragraph in text is slightly misleading because it is possible to derive Hansen–Jagannathan bounds on the pth moment of the pricing kernel for any $p > 1$, not just $p = 2$. However, for any fixed p, the Hansen–Jagannathan approach cannot falsify asset pricing models that have pricing kernels with infinite pth moment.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press

REFERENCES

Athreya, K. (1987) Bootstrap of the mean in the infinite variance case. *Annals of Statistics* 15, 724–731.
Billingsley, P. (1961) Lindeberg–Levy theorem for martingales. *Proceedings of the American Mathematical Society* 12, 788–792.
Breeden, D. (1979) An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7, 265–296.
Fama, E. (1965) The behavior of stock market prices. *Journal of Business* 38, 34–105.
Gnedenko, B.V. & A.N. Kolmogorov (1954) *Limit Distributions for Sums of Independent Random Variables*. Reading, MA: Addison–Wesley.
Hansen, L. & R. Jagannathan (1991) Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.
Hansen, L. & K. Singleton (1982) Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269–1286.
Hansen, L. & K. Singleton (1983) Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy* 91, 249–265.
Ibbotson, R. & R. Sinquefield (1987) *Stocks, Bonds, Bills and Inflation—1987 Yearbook*. Chicago: Ibbotson Associates.
Kocherlakota, N. (1990) On tests of representative consumer asset pricing models. *Journal of Monetary Economics* 26, 285–304.
Kocherlakota, N. (1996) The equity premium: It's still a puzzle. *Journal of Economic Literature* 34, 42–71.
Loretan, M. & P. Phillips (1994) Testing the covariance stationarity of heavy-tailed time series. *Journal of Empirical Finance* 1, 211–248.
Lucas, R.E., Jr. (1978) Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
Mandelbrot, B. (1963) The variation of certain speculative prices. *Journal of Business* 36, 394–419.
McCulloch, J.H. (1978) Continuous time processes with stable increments. *Journal of Business* 51, 601–619.
McCulloch, J.H. (1997) Measuring tail thickness to estimate the stable index α: A critique. *Journal of Business Economic Statistics* 15, 74–81.
Mehra, R. & E.C. Prescott (1985) The equity premium: A puzzle. *Journal of Monetary Economics* 15, 145–161.
Politis, D. & J. Romano (1994) A general theory for large sample confidence regions based on subsamples under minimal assumptions. *Annals of Statistics* 22, 2031–2050.
Rietz, T. (1988) The equity premium: A solution. *Journal of Monetary Economics* 22, 117–132.
Weil, P. (1989) The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics* 24, 401–422.

APPENDIX

The data on real per-capita consumption growth and real returns from 1889 through 1978 are the same as those of Mehra and Prescott (1985). The data on real returns from 1979 through 1987 are from Ibbotson and Sinquefield. The data on per-capita consumption growth from 1979 through 1987 are calculated using 1988 National Income and Product Accounts data on real consumption nondurables plus services and 1991 population figures from CITIBASE. The complete data set is available on request.

https://doi.org/10.1017/S136510059700401X Published online by Cambridge University Press