

**Financial Integration:
A New Methodology and an Illustration
Robert P. Flood and Andrew K. Rose***

Revised: March 14, 2005

Robert P. Flood
Research Dept, IMF
700 19th St., NW
Washington, DC 20431
Tel: (202) 623-7667

Andrew K. Rose (corresponding author)
Haas School of Business
University of California
Berkeley, CA 94720-1900
Tel: (510) 642-6609

Abstract

This paper develops a simple methodology to test for asset integration, and applies it within and between American stock markets. Our technique relies on estimating and comparing expected risk-free rates across assets. Expected risk-free rates are allowed to vary freely over time, constrained only by the fact that they must be equal across (risk-adjusted) assets in well integrated markets. Assets are allowed to have standard risk characteristics, and are constrained by the Fama-French factor model of covariances over short time periods. We find that internal integration in the S&P 500 market is never rejected and is generally not rejected in the NASDAQ. Integration between the NASDAQ and the S&P, however, is always rejected dramatically.

JEL Classification Numbers: G14

Keywords: risk-free, rate, intertemporal, asset, market, expected, price, stock, conditional.

* Acknowledgments. Flood is Senior Economist, Research Department, International Monetary Fund. Rose is B.T. Rocca Jr. Professor of International Business, Haas School of Business at the University of California, Berkeley, NBER Research Associate, and CEPR Research Fellow. We thank Pedro Rodriguez for assistance with the data, and Patrick Bolton, Ken French, Cam Harvey, Robert Hodrick, and seminar participants at Dartmouth, the Federal Reserve Board, and the IMF and especially an anonymous referee for comments. Rose thanks the Bank of Japan and the IMF for hospitality during the course of this research. This is a dramatically refocused version of NBER WP #9880; the data set, sample E-Views programs and output, and a current version of this paper are available at <http://faculty.haas.berkeley.edu/arose>.

Email addresses: Flood (rflood@imf.org); Rose (arose@haas.berkeley.edu).

1. Defining the Problem

The objective of this paper is to propose and implement an intuitive and simple-to-use measure of asset-market integration. What does *asset-market integration* mean? We adopt the view that financial markets are integrated when assets are priced by the same stochastic discount factor (SDF). More precisely, we define security markets to be integrated if all assets priced on those markets satisfy the pricing condition:

$$p_t^j = E_t(m_{t+1}x_{t+1}^j) \quad (1)$$

where: p_t^j is the price at time t of asset j , $E_t(\cdot)$ is the expectations operator conditional on information available at t , m_{t+1} is the intertemporal marginal rate of substitution (MRS), for income accruing in period $t+1$ (also interchangeably known by many names, including the discount rate, stochastic discount factor, marginal utility growth, and pricing kernel), and x_{t+1}^j is income received at $t+1$ by owners of asset j at time t (the future value of the asset plus any dividends or coupons).

Our object of interest in this study is $E_t m_{t+1}$ the time t *expectation* of the marginal rate of substitution (also interchangeably known as, e.g., the risk-free rate, and/or zero-beta return). Agents behaving according to equation (1) use the entire perceived distribution of m_{t+1} to price assets at t . Nevertheless, we concentrate on its first moment for two reasons. First, $E_t m_{t+1}$ is simple to measure. Second, cross-market differences in estimated values of $E_t m_{t+1}$ turn out in practice to be highly illuminating. In particular, they allow us to use standard risk pricing models to discriminate for differences in market integration.

We emphasize at the outset that our test investigates a *necessary* but not *sufficient* condition for market integration. In other words if two portfolios are well integrated they will pass our test, but passing the test does not imply the portfolios to be well integrated. On the other hand if two portfolios fail the test, the portfolios are not well integrated.

2. Methodology

We use a standard decomposition of equation (1):

$$p_t^j = E_t(m_{t+1}x_{t+1}^j) = COV_t(m_{t+1}, x_{t+1}^j) + E_t(m_{t+1})E_t(x_{t+1}^j). \quad (2)$$

where $COV_t(\cdot)$ denotes the conditional covariance operator. It is useful to rewrite this as

$$\begin{aligned} x_{t+1}^j &= -[1/E_t(m_{t+1})]COV_t(m_{t+1}, x_{t+1}^j) + [1/E_t(m_{t+1})]p_t^j + \varepsilon_{t+1}^j, & \text{or} \\ x_{t+1}^j &= \delta_t(p_t^j - COV_t(m_{t+1}, x_{t+1}^j)) + \varepsilon_{t+1}^j \end{aligned} \quad (3)$$

where $\varepsilon_{t+1}^j \equiv x_{t+1}^j - E_t(x_{t+1}^j)$, a prediction error, and $\delta_t \equiv 1/E_t(m_{t+1})$. The latter is the vector of parameters of interest to us. In an integrated market, it is identical for all assets. Our work below is essentially concerned with exploiting and testing this restriction.

It is traditional to make equation (3) stationary by dividing the equation by $p_{j,t}$, resulting in:

$$x_{t+1}^j / p_t^j = \delta_t(1 - COV_t(m_{t+1}, x_{t+1}^j / p_t^j)) + \varepsilon_{t+1}^j, \quad (4)$$

where ε_{t+1}^j is redefined appropriately. Dividing through by $p_{j,t}$ also converts equation (3) into an asset-pricing equation – an equation relating one-period asset returns, x_{t+1}^j / p_t^j , to the market $\delta_t \equiv 1 / E_t(m_{t+1})$, and to the asset-specific period risk premium.

Equation (4) is then given economic content by adding two assumptions:

1) *Rational Expectations*: ε_{t+1}^j is assumed to be uncorrelated with information available at time t, and

2) *Covariance Model*: $COV_t(m_{t+1}, x_{t+1}^j / p_t^j) = \sum_i \beta_i^j f_{i,t+1}$, for the relevant sample,

where: β_i^j is a set of I asset-specific factor coefficients and $f_{i,t+1}$ is a vector of time-varying factors.

Both assumptions are common in the literature; Campbell, Lo and MacKinlay (1997) and Cochrane (2001) provide excellent discussions. Our second assumption is clearly the more demanding of the two. It makes sense given two underlying presumptions: a) “SDF Spanning”; an admissible SDF m_{t+1} can be chosen as an affine function of some factors, and b) “Time Invariant Coefficients”; in a conditional affine regression of returns on factors, the conditional coefficients are time-invariant.¹

Combining our two assumptions into equation (4) delivers:

$$x_{t+1}^j / p_t^j = \delta_t (1 + \sum_i \beta_i^j f_{i,t+1}) + \varepsilon_{t+1}^j \quad (5)$$

Equation (5) is now a panel estimating equation. We use *time-series* variation to estimate the asset-specific factor loadings $\{\beta\}$, coefficients that are constant across time. We exploit *cross-sectional* variation to estimate $\{\delta\}$, the coefficients of interest that represent the risk-free return and are time varying but common to all assets.²

Our test for integration is simple. Estimating (5) for a set of assets $j=1,\dots,J_0$ and then repeating the analysis for the same period of time with a different set of assets $j=1,\dots,J_1$ gives us two sets of estimates of $\{\delta\}$, a time-series sequence of estimated discount rates. These can be compared directly, using conventional statistical techniques. Under the null hypothesis of market integration, the two sets of $\{\delta\}$ coefficients are equal. If the two diverge, the hypothesis of market integration between the assets is rejected (jointly with the other two assumptions, of course).³

We emphasize that the assumption of a well-functioning factor model is important because getting it wrong might lead to inconsistent $\delta(t)$ estimates. Accordingly, we take precautions. First, in implementing our new approach, we use the well-known aggregate factors used by Fama and French (1996) to model returns in the traditional approach, and we do robustness checks. Second, we require our factor model to hold with constant coefficients only over relatively short periods – generally two months of daily data. Third, we check our results against a simpler version of the Fama-French returns model, with only a single market factor. (Further sensitivity analysis is available in the working paper version.)

The measurements we produce are discriminating for market integration, yet they seem robust, and confirm both our prior beliefs and previous research (e.g., Chen and Knez, 1995). In the examples below, our measure never rejects internal market integration for portfolios of S&P

stocks priced in the NYSE and seldom rejects for portfolios priced on the NASDAQ, but rejects integration strongly – by an order of magnitude – between NYSE and NASDAQ portfolios.

3. Relationship to the Literature

Asset-market integration is a classic problem with a large associated literature, one which has grown along two branches. The first branch, based on parametric asset-pricing models, has been surveyed by Adams et. al. (2002), Cochrane (2001), and Campbell, Lo, and MacKinlay (1997). Karolyi and Stulz (2002) provide a survey of open-economy asset-market integration concepts and results. Along this branch, a parametric discount-rate model is used to price asset portfolios. Pricing errors are compared across portfolios. If the portfolios are integrated, the pricing errors should not be systematically identifiable with the portfolios in which they originate. Roll and Ross (1980) tested market integration this way using an (APT) arbitrage pricing theory model, and a large literature has followed, see e.g., Bekaert and Harvey (1995), hereafter “BH”.

The second branch of literature grows from the work of Hansen and Jagannathan (1991) and is represented by Chen and Knez (1995) and Chabot (2000). Along this branch, data from each supposed market is used to characterize the set of stochastic discount factors (SDF) that could have produced the observed data. Testing for cross market integration involves measuring the distance between admissible MRS sets, and asking if, and by how much, they overlap. If a common SDF exists the markets are integrated. If not, measures are available to judge the distance between the market-specific SDF sets.

Our work rests on the first branch, since we use parametric models to condition our estimation. It differs from previous work in three ways.

First, we do not measure integration by the full-blown cross-sectional pricing errors produced by a particular model. BH, working along the first branch used the definition “Markets are completely integrated if assets with the same risk have identical expected returns irrespective of the market.” Our market integration measure is based on a subset of the cross-market conditions demanded by BH. Instead of comparing all aspects of a fully parameterized SDF models, we measure integration by the implied first moment of the SDF. The condition we study, therefore, a necessary condition for integration. It is a subset of the conditions demanded by BH, and also Chen and Knez. Studying it will be valuable, therefore, only if it is simple to produce but still discriminating.

Second, parametric pricing models are often estimated with long data spans and are thus sensitive to parameter instability in time series long enough for precise estimation (e.g., Fama and French (1996); discussion is provided by Cochrane, 2001). We minimize (but do not avoid completely) the instability problem by concentrating attention on a parameter that is conditionally invariant to time-series instability. The measure we use is a free parameter, constant across assets but unconstrained across time. Our measure – borrowed from Roll and Ross – is therefore basically cross-sectional. Thus we can estimate the measure using a short time-series dimension.

Finally, we do not assume that the bond market is integrated with other asset markets. When applied to a bond without nominal risk (e.g., a treasury bill), equation (1) implies

$$1 = E_t(m_{t+1}(1+i_t)) \quad \text{or} \quad \delta_t \equiv 1/E_t(m_{t+1}) = (1+i_t)$$

where: i_t is a risk-free nominal interest rate, and m_{t+1} is a nominal MRS. One tradition, common in Economics and Finance, is to assume that the SDF pricing bonds is the same for all bonds, and identical to that pricing all stocks (and other assets). We do not *impose* the assumption that the treasury bill rate equals the expectation of the MRS; instead we *estimate* the expected MRS for different assets and compare them.

4. Empirical Implementation

We begin by estimating our model (5) with the three time-varying factors used by Fama and French (1996). That is, we estimate:

$$x_{t+1}^j / p_t^j = \delta_t (1 + \beta_1^j f_{1,t+1} + \beta_2^j f_{2,t+1} + \beta_3^j f_{3,t+1}) + \varepsilon_{t+1}^j \quad (6)$$

for assets $j=1, \dots, J$, periods $t=1, \dots, T$. We allow $\{\delta_t\}$ to vary period by period, while we use a “three-factor” model and allow $\{\beta^j\}$ to vary asset by asset. The three Fama-French factors are: 1) the overall stock market return, less the treasury-bill rate, 2) the performance of small stocks relative to big stocks, and 3) the performance of “value” stocks relative to “growth” stocks. Further details and the data set itself are available at French’s website.⁴ We also examine a simpler covariance model below.⁵

Equation (6) can be estimated directly with non-linear least squares. The degree of non-linearity is not particularly high; conditional on $\{\delta_t\}$ the problem is linear in $\{\beta^j\}$ and vice versa. We employ robust (heteroskedasticity and autocorrelation consistent “Newey West”) covariance estimators.

We use a moderately high frequency approach. In particular, we use two-month spans of daily data. Using daily data allows us to estimate the coefficients of interest $\{\delta_t\}$ without assuming that firm-specific coefficients $\{\beta^j\}$ are constant for implausibly long periods of time.

Our empirical illustration examines the integration of American equity markets. Large American stocks are traded on liquid markets, which we consider *a priori* to be integrated. We begin by examining daily data over a quiet two-month period, April-May 1999 (about a year before the end of the Clinton bull market).⁶ Two months gives us a span of over forty business day observations; this does not appear to stretch our reliance on a factor model of asset covariances excessively, while still allowing us to test financial market integration for an interesting span of data. We see no reason why higher- and/or lower-frequency data cannot be used.⁷

Our data set is drawn from the “US Pricing” database provided by Thomson Analytics. We collected closing rates for the first (in terms of ticker symbol) one hundred firms from the S&P 500 that did not go ex-dividend during the months in question. The absence of dividend payments allows us to set $x_{t+1}^j = p_{t+1}^j$ (and does not bias our results in any other obvious way).

We group our hundred firms into twenty portfolios of five firms each, arranged simply by ticker symbol. We use portfolios rather than individual stocks for the standard reasons of the Finance literature. In particular, as Cochrane (2001) points out, portfolios betas are measured with less error than individual betas because of lower residual variance. They also vary less over time (as size, leverage, and business risk change less for a portfolio of equities than any individual component). Portfolio variances are lower than those of individual securities, enabling more precise covariance relationships to be estimated. And of course portfolios are what investors tend to use (especially those informed by Finance theory!).

Our first sample period consists of 41 days. Since we lose the first and last observations because of lags (p_{t-1}^j) and leads (x_{t+1}^j), we are left with a total of 780 observations in our panel data set (20 portfolios x 39 days). Our data has been checked for transcription errors, both visually and with random crosschecking.

There is no reason that one cannot use more data (longer spans at different frequencies, for larger number of firms and/or portfolios grouped non-randomly). We choose this sample (only two months of daily price data for one hundred firms grouped randomly into twenty portfolios) deliberately to illustrate the power of our methodology and its undemanding data requirements. However, we also check for sensitivity with respect to the sample below.

5. Results

We start by splitting our 20 portfolios of S&P stocks into two sets of 10 portfolios each (simply by ticker symbol) to estimate the expected marginal rate of substitution (i.e., estimates of $\delta_t \equiv [1/E_t(m_{t+1})]$). Since we are interested in testing for integration, we examine a joint test of equality between the two sets of estimated deltas (from the two different sets of ten portfolios). The statistical information is contained in the cells at the top left of Table 1. The log-likelihood of (6) estimated from the first set of 10 portfolios is 1176.9; that from the second set of 10 portfolios is 1177.8. When (6) is estimated from all 20 portfolios simultaneously so that only a single set of $\{\delta_t\}$ is extracted, the log-likelihood is 2334.3. Under the hypothesis of integration (i.e., the same $\{\delta_t\}$ for both sets of assets) and normally distributed errors, minus twice the difference in the log-likelihoods is distributed as a chi-square with 39 degrees of freedom; a likelihood ratio (LR) test. The test statistic is 40.9, consistent with the hypothesis of integration and normal residuals at the .61 confidence level.⁸

It is well known that asset prices are not in fact normally distributed; Campbell, Lo, and MacKinlay (1997). Rather, there is strong evidence of fat tails or leptokurtosis, and this also characterizes our data.⁹ Accordingly, we used a bootstrap procedure to check the probability values for our likelihood ratio tests.¹⁰ The bootstrapped critical values for the test of integration are higher than those of the chi-squared distribution, reinforcing our view that there is no evidence against the null hypothesis of integration.

To check for sample sensitivity, we also consider five other sample periods: July-August 1999, October-November 1999, and the same three two-month samples for the bear market of 2002. Results from these other sample periods are also included in Table 1 and are also consistent with the hypothesis of integration inside the S&P 500 at all reasonable confidence levels.

What about the NASDAQ market for smaller stocks? We follow exactly the same procedures, but using data drawn from the NASDAQ market. We group (again on the basis of ticker symbol) data from 100 NASDAQ firms into 20 portfolios of 10 firms each, and test for equality of deltas (between the two different sets of deltas, estimated from the two sets of ten NASDAQ portfolios) using likelihood ratio tests (again, checking with bootstrapped extreme values). The results are presented in Table 2, and are generally consistent with the null hypothesis of integration inside the NASDAQ. However, two of our samples (July-Aug 1999 and Oct-Nov 1999) are inconsistent with integration at the .05 confidence level (these are marked with an asterisk). Both periods occurred shortly before the collapse of the NASDAQ. We think of these as intuitive, reasonable results, possibly consistent with the existence of “irrational exuberance” manifest in the NASDAQ in the final run-up to the height of the internet bubble.

Still, the most interesting question to us is: Is the market for large (S&P 500) stocks integrated with the NASDAQ? It is easy to ask the question by comparing $\{\delta_t\}$ estimates when (6) is estimated with: a) the twenty S&P portfolios; b) the twenty NASDAQ portfolios; and c) all forty portfolios pooled together (which is most efficient if the two markets are integrated). Our LR tests for this hypothesis are presented in Table 3 and are grossly inconsistent with the null hypothesis of market integration. The LR test statistics are often an order of magnitude bigger than those of Tables 1 and 2. That is, while the S&P always seems integrated and the NASDAQ is generally integrated, the S&P is never integrated with the NASDAQ. This result is similar to that of Chen and Knez (1995).

Thus far we have relied on the Fama-French model of asset covariances. That is, the covariance of each asset's return with the expectation of the MRS is characterized by three parameters or factor loadings: the market return minus the T-bill rate (β_1^j), the difference between small and large stock returns (β_2^j), and the difference between returns of stocks with high and low book to market ratios (β_3^j). Are our results sensitive to the number of factors used? It turns out that the answer is negative.

In Table 4 we provide test statistics to examine integration within the S&P and NASDAQ and between the two markets, but using only the return on the market instead of the three Fama-French factors. The test statistics and conclusions are essentially unchanged.¹¹

6. Summary and Conclusions

This paper developed a simple method to test for asset integration, and then applied it within and between American equity markets. It relies on estimating and comparing the expected risk-less returns implied by different sets of assets. Our technique has a number of

advantages over those in the literature and relies on just two assumptions: 1) rational expectations in financial markets; and 2) covariances between discount rates and returns that can be modeled with a small number of factors for a short period of time.

We illustrated this technique with an application to stocks drawn from the NYSE and the NASDAQ, and found that: a) the NYSE always seems to be integrated; b) the NASDAQ is usually (but not always) integrated; and c) the NYSE and NASDAQ do not seem close to being integrated. Our results seem reasonably insensitive to the exact sample and conditioning model used.

If our finding of integration within but not across stock markets holds up to further scrutiny, the interesting question is not whether financial markets with few apparent frictions are poorly integrated but why? We leave that important question for future research.

References

Adam, Klaus, Tullio Jappelli, Annamaria Menichini, Mario Padula, and Marco Pagano (2002) “Analyse, Compare, and Apply Alternative Indicators and Monitoring Methodologies to Measure the Evolution of Capital Market Integration in the European Union” University of Salerno manuscript.

Bekaert, Geert and Campbell R. Harvey (1995) “Time-Varying World Market Integration” *The Journal of Finance* 50-1, pp. 403-444.

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

Chabot, Benjamin (2000) “A Single Market? The Stock Exchanges of the United States and London: 1866-1885” University of Michigan working paper.

Chen, Zhiwu and Peter J. Knez (1995) “Measurement of Market Integration and Arbitrage” *Review of Financial Studies* 8-2, 287-325.

Cochrane, John H. (2001) *Asset Pricing*, Princeton University Press.

Fama, Eugene and Kenneth R. French (1996) “Multifactor Explanations of Asset Pricing Anomalies” *Journal of Finance* 51-1, 55-84.

Hansen, Lars Peter and Ravi Jagannathan (1991) “Implications of Security Market Data for Models of Dynamic Economies” *Journal of Political Economy* 99-2, 225-262.

Karolyi, G. Andrew and Rene’ M. Stulz (2002) “Are Financial Assets priced locally or globally” in G. Constantinides, M. Harris and R. Stulz (eds.) *Handbook of the Economics of Finance*, North Holland.

Roll, Richard, and Stephen A. Ross (1980) “An Empirical Investigation of the Arbitrage Pricing Theory” *Journal of Finance* 35-5, 1073-1103

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
First 10 portfolios	1176.9	1292.9	1158.6
Second 10 portfolios	1177.8	1298.2	1161.1
All 20 portfolios	2334.3	2569.7	2306.2
LR Integration Test (df)	40.9 (39)	42.7 (41)	27.2 (40)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 10 portfolios	1408.9	1236.9	1230.8
Second 10 portfolios	1392.7	1300.6	1209.8
All 20 portfolios	2786.8	2519.6	2424.8
LR Integration Test (df)	29.4 (42)	35.7 (42)	31.4 (41)

Table 1: Integration inside the S&P 500, Fama-French-Factor Model

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
First 10 portfolios	909.7	1076.3	749.7
Second 10 portfolios	861.4	999.9	931.0
All 20 portfolios	1742.4	2030.0	1622.0
LR Integration Test (df)	57.5 (39)	88.2* (41)	116.9* (40)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 10 portfolios	1034.8	1048.6	994.0
Second 10 portfolios	1182.3	990.6	944.9
All 20 portfolios	2176.3	2016.8	1907.8
LR Integration Test (df)	81.8 (42)	44.9 (42)	62.2 (41)

Table 2: Integration inside the NASDAQ, Fama-French -Factor Model

Note:) indicates significantly different from zero at the .05 level, bootstrapped confidence interval.

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
20 S&P Portfolios	2334.3	2569.7	2306.2
20 NASDAQ Portfolios	1742.4	2030.0	1622.0
Combined	3829.5	4406.5	3632.4
LR Integration Test (df)	494.3** (39)	390.6** (41)	592.0** (40)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
20 S&P Portfolios	2786.8	2519.6	2424.8
20 NASDAQ Portfolios	2176.3	2016.8	1907.8
Combined	4713.0	4329.8	4129.8
LR Integration Test (df)	500.2** (42)	413.4** (42)	405.6** (41)

Table 3: Integration between S&P 500 and NASDAQ, Fama-French -Factor Model

Note: ** indicates significantly different from zero at the .01 level, bootstrapped confidence interval.

LR Integration Test (df)	April-May 1999 (39)	July-Aug. 1999 (41)	Oct.-Nov. 1999 (40)
Within S&P	33.2	47.7	39.5
Within NASDAQ	65.8	57.7	118.9**
S&P vs. NASDAQ	513.5**	374.8**	595.9**
	April-May 2002 (42)	July-Aug. 2002 (42)	Oct.-Nov. 2002 (41)
Within S&P	50.9	42.6	38.0
Within NASDAQ	114.7**	53.9	54.9
S&P vs. NASDAQ	548.1**	338.5**	379.2**

Table 4: Integration within and between S&P 500 and NASDAQ, One-Factor Model

Note: ** indicates significantly different from zero at the .01 level, bootstrapped confidence interval.

Endnotes

¹ Thus, $g_{i,t} = \text{Cov}_t(m_{t+1}, f_{i,t+1})$, as pointed out by an anonymous referee.

² Thus our estimator is outside the scope of traditional empirical asset pricing models, since they typically set the expected MRS to an appropriate short-term riskless interest rate; more on this below.

³ Other methods of estimating $\{\delta\}$ are discussed in the earlier version of this paper.

⁴ http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

⁵ Other covariance models are used in the earlier version of the paper, and deliver similar conclusions.

⁶ We choose these months to avoid January (and its effect), February (a short month), and March (a quarter-ending month), but test for sample sensitivity extensively below.

⁷ For instance, we could use data at five-minute intervals for a day, making our assumption of constant asset-specific effects even more plausible; but the question of whether financial markets are integrated over hours (not weeks) is less interesting to us.

⁸ More evidence, including time-series plots derived from a different model, is provided in the working paper version.

⁹ Jarque-Bera tests are inconsistent with the null hypothesis for $\{\varepsilon\}$ at all reasonable confidence levels.

¹⁰ Our bootstrap procedure is as follows. We estimate the deltas from (say) all 20 portfolios under the null hypothesis of integration. This gives us an estimate of $\{\varepsilon\}$. We then draw randomly with replacement from this vector to create an artificial vector of $\{\varepsilon\}$ which we use to construct an artificial regressand variable $\{x\}$. Using this artificial data we then generate a likelihood ratio test by estimating the model from the first set of 10 portfolios, the second set of 10 portfolios, and the combined set of 20. We then repeat this procedure a large number of times to generate a distribution for the LR test statistic.

¹¹ Integration is now rejected for the NASDAQ in April-May 2002 instead of July-Aug 1999.