

DISCUSSION PAPER SERIES

No. 4027

**FINANCIAL INTEGRATION:
A NEW METHODOLOGY
AND AN ILLUSTRATION**

Robert P Flood and Andrew K Rose

*FINANCIAL ECONOMICS and
INTERNATIONAL MACROECONOMICS*



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP4027.asp

FINANCIAL INTEGRATION: A NEW METHODOLOGY AND AN ILLUSTRATION

Robert P Flood, International Monetary Fund (IMF)
Andrew K Rose, University of California, Berkeley and CEPR

Discussion Paper No. 4027
August 2003

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS and INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Robert P Flood and Andrew K Rose

CEPR Discussion Paper No. 4027

August 2003

ABSTRACT

Financial Integration: A New Methodology and an Illustration*

This Paper develops a simple new methodology to test for asset integration and applies it within and between American stock markets. Our technique is tightly based on a general intertemporal asset-pricing model, and 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 are equal across (risk-adjusted) assets. Assets are allowed to have general risk characteristics, and are constrained only by a factor model of covariances over short time periods. The technique is undemanding in terms of both data and estimation. We find that expected risk-free rates vary dramatically over time, unlike short interest rates. Further, the S&P 500 market seems to be well integrated, and the NASDAQ is generally (but not always) integrated. The NASDAQ, however, is poorly integrated with the S&P 500.

JEL Classification: F32 and G15

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

Robert P Flood
Research Department
International Monetary Fund
700 19th Street NW
Washington DC 20041
USA
Tel: (1 202) 623 7667
Fax: (1 202) 623 6339
Email: rflood@imf.org

Andrew K Rose
University of California at Berkeley
Haas School of Business
Berkeley, CA 94720-1900
USA
Tel: (1 510) 642 6609
Fax: 642 4700
Email: arose@haas.berkeley.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=100612

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=116112

*We thank Pedro Rodriguez for assistance with the data, and Ken French, Cam Harvey, Robert Hodrick and seminar participants at Dartmouth, the Federal Reserve Board, and the IMF for comments. Rose thanks the IMF for hospitality during the course of this research. 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>.

Submitted 16 July 2003

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 rate. 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 as the discount rate, stochastic discount factor, marginal utility growth, pricing kernel, and zero-beta return), 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). We rely only on this standard and general intertemporal model of asset valuation; to our knowledge this Euler equation is present in all equilibrium asset-pricing models.

Our object of interest in this study is m_{t+1} , the marginal rate of substitution, or, more precisely, estimates of the *expected* marginal rate of substitution, $E_t m_{t+1}$. The MRS is the unobservable DNA of intertemporal decisions; characterizing its distribution is a central task of economics and finance. The discount rate ties pricing in a huge variety of asset markets to peoples' saving and investment decisions. The thrust of this paper is to use asset prices and payoffs to characterize important aspects of its distribution.

The substantive point of equation (1) is that all assets in a market share the same marginal rate of substitution. There is no asset-specific MRS in an integrated market, and no market-specific MRS when markets are integrated with each other. Learning more about the MRS is of intrinsic interest, and has driven much research (e.g., Hansen and Jagannathan, 1991, who focus on second moments). Measures of the expected MRS also lead naturally to an intuitive test for integration. In this paper, we propose and implement such a simple test for the equality of $E_t m_{t+1}$ across sets of assets. This is a necessary (but not sufficient) condition for market integration.

2: Methodology

We use the fact that in an integrated market, the MRS prices all assets held by the marginal asset holder. Indeed what we *mean* by asset market integration is that the same MRS prices all the assets. In other words, if we could extract m_{t+1} (or rather, its expectation) independently from a number of different asset markets, *they should all be the same if those markets are integrated*. As Hansen and Jagannathan (1991) show, there may be many stochastic discount factors consistent with any set of market prices and payoffs; hence our focus on the *expectation* of MRS, which is unique.

Consider a generic identity related to (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

$$x_{t+1}^j = -[1/E_t(m_{t+1})] COV_t(m_{t+1}, x_{t+1}^j) + [1/E_t(m_{t+1})] p_t^j + e_{t+1}^j, \quad \text{or}$$

$$x_{t+1}^j = \mathbf{d}_t(p_t^j - COV_t(m_{t+1}, x_{t+1}^j)) + \mathbf{e}_{t+1}^j \quad (3)$$

where $\mathbf{d}_t \equiv 1/E_t(m_{t+1})$ and $\mathbf{e}_{t+1}^j \equiv x_{t+1}^j - E_t(x_{t+1}^j)$, a prediction error.

We then impose two restrictions:

- 1) *Rational Expectations*: \mathbf{e}_{t+1}^j is assumed to be white noise, uncorrelated with information available at time t, and
- 2) *Covariance Model*: $COV_t(m_{t+1}, x_{t+1}^j) = \mathbf{b}_0^j + \sum_i \mathbf{b}_i^j f_{i,t}$, for the relevant sample,

where: \mathbf{b}_0^j is an asset-specific intercept, \mathbf{b}_i^j is a set of I asset-specific factor coefficients and $f_{i,t}$ a vector of time-varying factors.

With our two assumptions, equation (3) becomes a panel estimating equation. We exploit *cross-sectional* variation to estimate $\{\mathbf{d}\}$, the coefficients of interest that represent the risk-free return and are time varying but common to all assets. These estimates of the MRS are the focus of our study. We use *time-series* variation to estimate the asset-specific “fixed effects” and factor loadings $\{\mathbf{b}\}$, coefficients that are constant across time. Intuitively, these coefficients are used to account for asset-specific systematic risk (the covariances).

Estimating (3) 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 $\{\mathbf{d}\}$, a time-series sequence of estimated discount rates. These can be compared directly, using conventional statistical techniques, either one by one, or jointly. Under the null hypothesis of market integration, the two sets of $\{\mathbf{d}\}$ coefficients are equal.

Discussion

We make only two assumptions; both are conventional in the literature, though we rely on them less than most of the literature. It seems reasonable to assume that expectations are rational for financial markets, at least in our limited sense that asset-pricing errors are not *ex ante* predictable at high frequencies. Our assumption that the asset-specific covariances (of payoffs with the MRS) are either constant or depend on a small number of factors is more controversial, but standard practice. Rather than develop our own factor model, we rely on the well-known three-factor model famously deployed by Fama and French (1996). We defend it on two grounds. First, in applications we maintain the covariance model for only two months at a time (though we have tried shorter periods with similar results); Fama and French assumed that the same model worked well for thirty years. Second, our results are insensitive to the exact factor model. If our technique were sensitive to the factors used to model $\{\mathbf{d}\}$, then the measure would be no more useful than any of the individual factor models. Indeed, if the measure were factor-model sensitive, it would be preferable to use the factor model itself as the object of measurement. Nevertheless we stress that we must rely on *some* model of covariances.

While we focus on (3), there are other moments that would help characterize the MRS, $\{\mathbf{d}\}$; see e.g., Hansen and Jagannathan (1991). We concentrate on this one for four reasons. First, as the first moment it is the natural place to check first. Second, it is simple to estimate. Third, our estimates and results are robust to the factor model that conditions the measurements. Finally, the measurements are discriminating for market integration, yet they confirm 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 strongly – by an order of magnitude – integration between NYSE and NASDAQ portfolios.

Our methodology has a number of strengths. First, it is based on a general intertemporal theoretical framework, unlike other measures of asset integration such as stock market correlations (see the excellent discussion in e.g., Adam et. al. 2002). Second, standard asset-pricing models are completely consistent with our methodology, and the exact model does not seem to be important in practice. We use the Fama-French workhorse (which subsumes the CAPM), but find that our results are insensitive to the exact choice of model. Third, we do not need to model the MRS directly. The MRS need not be determined uniquely, so long as its expectation is unique. Fourth, our strategy requires only two assumptions; we need not assume e.g., complete markets, homogeneous investors, or that we can model “mimicking portfolios” well. Fifth, the technique requires only accessible and reliable data on asset prices, payoffs, and time-varying factors. Sixth, the methodology can be used at very high frequencies and at low frequencies as well. Seventh, the technique can be used to compare expected discount rates across many different classes of assets including domestic and foreign stocks, bonds, and commodities. Next, the technique is easy to implement and can be applied with standard econometric packages; no specialized software is required. Finally, the technique is focused on an intrinsically interesting object, the expected marginal rate of substitution.

3: Relationship to the Literature

The literature is clear that asset markets are integrated when identical cash flows are priced equally across markets (e.g., Adam et. al., 2002 and Cochrane, 2001). This is the asset-market version of economists’ trusty “Law of One Price.” But since no two different assets have

identical cash flows, the integration definition must be extended to be useful. The standard holds two asset markets to be integrated when risks in those markets are shared completely and priced identically. One way to make this definition operational requires identifying the relevant risks. Roll and Ross (1980) recognized the dependence of integration measures on risk identification. They tested asset integration using the argument that two portfolios are integrated only if their implied risk-less returns are the same; our test is similar to theirs in spirit. This simple observation is powerful because it invokes the cross-sectional dimension where every asset in an integrated market implies the same risk-free return.

The literature on asset-market integration 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). Along this branch, a parametric discount-rate models to used to price two 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 arbitrage pricing theory model, and a large literature has followed.

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 market are used to characterize the set of stochastic discount factors that could have produced the observed data. Testing for integration across markets involves measuring the distance between admissible MRS sets, and asking if, and by how much, they overlap.

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

First, we diverge from the finance profession in treating $\{b\}$ as a set of nuisance coefficients. Rather than being of intrinsic interest to us, they are required only to clear the way to produce estimates of the MRS.

Next, we do not measure integration by the cross-sectional pricing errors produced by a particular model; this approach seems relatively non-specific and model-dependent. Instead we measure integration by the implied first moment of the stochastic discount rate (MRS). The condition we study, therefore, is a necessary condition for integration. Studying it will be valuable only if it is a discriminating condition; it turns out to be so.

Third, 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); an excellent 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 is therefore basically cross-sectional, that we estimate precisely using a short time-series dimension.

Finally, we do not assume that (3) holds for the bond market, or 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

$$1 = E_t(m_{t+1}(1+i_t)) \tag{1'}$$

where: i_t is a risk-less nominal interest rate, and m_{t+1} is a nominal MRS. The tradition inside finance is to assume that the MRS pricing bonds is the same for all bonds, and identical to that

pricing all stocks (and other assets). If we make this assumption, $\mathbf{d}_t \equiv 1 / E_t(m_{t+1}) = (1 + i_t)$.

We do not *impose* this assumption; rather we *test* it (and reject) it.

4: Empirical Implementation

We begin by estimating a model with asset-specific intercepts and the three time-varying factors used by Fama and French (1996). In practice, we divide through by lagged prices (and redefine residuals appropriately):

$$x_{t+1}^j / p_{t-1}^j = \mathbf{d}_t ((p_t^j / p_{t-1}^j) + \mathbf{b}_0^j + \mathbf{b}_1^j f_{1,t} + \mathbf{b}_2^j f_{2,t} + \mathbf{b}_3^j f_{3,t}) + \mathbf{e}_{t+1}^j \quad (4)$$

for assets $j=1, \dots, J$, periods $t=1, \dots, T$. That is, we allow $\{\mathbf{d}_t\}$ to vary period by period, while we use a “three-factor” model and allow $\{\mathbf{b}^j\}$ varying asset by asset. We normalize the data by lagged prices since we believe that $COV_t(m_{t+1}, x_{t+1}^j / p_{t-1}^j)$ can be modeled by a simple factor model with time-invariant coefficients more plausibly than $COV_t(m_{t+1}, x_{t+1}^j)$. 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.¹ For sensitivity analysis, we also examine two other covariance models: one with only a single time varying factor, namely the overall market return; and the other without any time-varying factors at all (but, as always, with an asset-specific intercept).

Equation (4) can be estimated directly with non-linear least squares. The degree of non-linearity is not particularly high; conditional on $\{\mathbf{d}_t\}$ the problem is linear in $\{\mathbf{b}^j\}$ 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 $\{\mathbf{d}_t\}$ without assuming that firm-specific coefficients $\{\mathbf{b}^j\}$ 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).² 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.³

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 into two sets of 10 portfolios each (simply by ticker symbol) to estimate discount rates (i.e., estimates of $\mathbf{d}_t \equiv [1 / E_t(m_{t+1})]$). We provide time-series plots of the estimated deltas from the first 10 portfolios along with a plus/minus two standard error confidence interval in Figure 1. We also include the point estimates of delta from the second 10 portfolios, estimated in precisely the same way but using data from the last set of 10 portfolios.

There are two striking features of the graph. First, the time-series variation in delta is high, consistent with the spirit of Hansen and Jagannathan (1991). As shown in Table 1, the log likelihood of our equation estimated on the first 10 S&P portfolios is 1160. In April-May 1999, the US 3-month Treasury bill rate averaged 4.4%, a daily return of 1.00017 (with little time-

series variation). The log likelihood for the default equation estimated with 1.00017 substituted in place of $\{\mathbf{d}_t\}$ is only 1059. Under the null hypothesis of deltas that are constant and equal to the T-bill interest rate, $2*(1160-1059)$ is distributed as a chi-square with 39 degrees of freedom, grossly inconsistent with the null at any reasonable confidence level. (When we use all 20 portfolios, the analogue is $2*(2309-2136)$, again grossly inconsistent with the null.) That is, the hypothesis that the MRS is equal to the short t-bill rate is wildly inconsistent with the data. The MRS seems much more volatile than short-term interest rates.

Second, the estimates of delta from the two different sets of portfolios are similar; the deltas from the second set of portfolios almost always lie within the ± 2 standard error confidence interval of the first estimate of delta. That is, the two different sets of delta are usually statistically indistinguishable on any given day, consistent with the null hypothesis of integration within the S&P.

What about the two sets of delta examined jointly? The ocular evidence leads one to believe that the two sets of deltas are broadly equal. The statistical analogue is contained in the cells at the top left of Table 1. The log-likelihood of (4) estimated from the first set of 10 portfolios is 1160; that from the second set of 10 portfolios is 1166. When (4) is estimated from all 20 portfolios simultaneously so that only a single set of $\{\mathbf{d}_t\}$ is extracted, the log-likelihood is 2309. Under the hypothesis of integration (i.e., the same $\{\mathbf{d}_t\}$ 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 36, consistent with the hypothesis of integration and normal residuals at the .61 confidence level.

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 certainly

characterizes our data.⁴ Accordingly, we use a bootstrap procedure to estimate the probability values for our likelihood ratio tests.⁵ The bootstrapped p-value for the test of integration is even more consistent with the null hypothesis of integration at the .90 level.

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 standard 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 with bootstrapped p-values. The results are presented in Table 2, and are generally consistent with the null hypothesis of integration inside the NASDAQ. However, one of our samples (April-May 2002) is inconsistent with integration at the .03 confidence level (this is marked with an asterisk), while integration is overwhelmingly rejected for Oct-Nov 1999 (two asterisks), 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 just around 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 $\{d_t\}$ estimates when (4) 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 (with bootstrapped p-values) for this hypothesis are presented in Table 3 and are grossly

inconsistent with the null hypothesis of market integration. The LR test statistics are 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).

Time-series plots of $\{\mathbf{d}_t\}$ estimated from all (twenty) S&P and NASDAQ portfolios are provided in Figure 2 for all six sample periods, along with confidence intervals. Figure 3 provides scatterplots of S&P deltas against NASDAQ deltas. All these graphs indicate that there is no single obvious characteristic difference between the S&P and NASDAQ deltas.

6: Sensitivity Analysis

Thus far we have relied on the Fama-French model of asset covariances. That is, the covariance of each asset's return with the MRS is characterized by four parameters: an intercept (\mathbf{b}_0^j) and factor loadings on the market return minus the T-bill rate (\mathbf{b}_1^j), the difference between small and large stock returns (\mathbf{b}_2^j), and the difference between returns of stocks with high and low book to market ratios (\mathbf{b}_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 (and bootstrapped p-values) to examine tests of 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 (while retaining the portfolio intercepts as well). The test statistics and conclusions are essentially unchanged.

Table 5 goes even further and drops the market factor from our covariance model, leaving only portfolios-specific intercepts (\mathbf{b}_0^j) but no time-varying factors. Again, the results

are essentially unchanged. This robustness is encouraging since it demonstrates the insensitivity of our methodology to reasonable perturbations in the exact factor model employed.

7: 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 relatively weak 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 S&P 500 and the NASDAQ, and found that a) the time-series variation in the Marginal Rate of Substitution is high; b) the S&P always seems to be integrated; c) the NASDAQ is usually (but not always) integrated; and d) the S&P 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.

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.

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	1160.	1302.	1157.
Second 10 portfolios	1166.	1299.	1172.
All 20 portfolios	2309.	2574.	2303.
Test (bootstrap P-value)	36 (.90)	54 (.37)	51 (.43)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 10 portfolios	1438.	1255.	1247.
Second 10 portfolios	1405.	1302.	1227.
All 20 portfolios	2805.	2525.	2456.
Test (bootstrap P-value)	75 (.06)	62 (.24)	37 (.90)

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	881.	1066.	757.
Second 10 portfolios	816.	990.	945.
All 20 portfolios	1677.	2023.	1625.
Test (bootstrap P-value)	42 (.83)	65 (.20)	153** (.00)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
First 10 portfolios	1052.	1061.	991.
Second 10 portfolios	1174.	1003.	962.
All 20 portfolios	2185.	2035.	1919.
Test (bootstrap P-value)	82* (.03)	58 (.45)	69 (.08)

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

Log Likelihoods	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
20 S&P Portfolios	2309.	2574.	2303.
20 NASDAQ Portfolios	1677.	2023.	1625.
Combined	3706.	4396.	3633.
Test (bootstrap P-value)	559** (.00)	403** (.00)	590** (.00)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
20 S&P Portfolios	2805.	2525.	2456.
20 NASDAQ Portfolios	2185.	2035.	1919.
Combined	4735.	4352.	4170.
Test (bootstrap P-value)	511** (.00)	416** (.00)	410** (.00)

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

Test Statistics (bootstrap P-value)	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
Within S&P	36 (.93)	48 (.75)	30 (.99)
Within NASDAQ	47 (.79)	65 (.27)	127** (.00)
S&P vs. NASDAQ	548** (.00)	388** (.00)	594** (.00)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
Within S&P	44 (.88)	55 (.61)	35 (.98)
Within NASDAQ	80 (.09)	58 (.61)	72 (.13)
S&P vs. NASDAQ	497** (.00)	432** (.00)	422** (.00)

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

Test Statistics (bootstrap P-value)	April-May 1999	July-Aug. 1999	Oct.-Nov. 1999
Within S&P	33 (.97)	46 (.71)	34 (.94)
Within NASDAQ	42 (.80)	62 (.28)	114** (.00)
S&P vs. NASDAQ	534** (.00)	378** (.00)	591** (.00)
	April-May 2002	July-Aug. 2002	Oct.-Nov. 2002
Within S&P	46 (.76)	47 (.77)	36 (.95)
Within NASDAQ	86* (.03)	52 (.63)	68 (.12)
S&P vs. NASDAQ	506** (.00)	416** (.00)	419** (.00)

Table 5: Integration within and between S&P 500 and NASDAQ, Only Asset Intercepts




Figure 1: Estimates of Marginal Rate of Substitution from two sets of (10) S&P portfolios

Deltas from Different Markets and Samples




Figure 2: Estimates of Marginal Rate of Substitution from sets of (20) portfolios




Figure 3: Estimates of Marginal Rate of Substitution from sets of (20) portfolios

Endnotes

¹ http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

² 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.

³ 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.

⁴ Jarque-Bera tests are inconsistent with the null hypothesis for $\{\epsilon\}$ at all reasonable confidence levels.

⁵ 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 $\{\epsilon\}$. We then draw with randomly with replacement from this vector to create an artificial vector of $\{\epsilon\}$ 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.