Equity Integration in Japan: An Application of A New Method*

Andrew K. Rose*

Draft: July 31, 2003
Comments Welcome

Contact:
Haas School of Business
University of California
Berkeley, CA 94720-1900
Tel: (510) 642-6609
E-mail: arose@haas.berkeley.edu

Abstract

This paper develops a simple new methodology to test for asset integration and applies it to the Japanese stock market. The 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. I find that expected risk-free rates vary dramatically over time, unlike short interest rates. Further, the Tokyo stock market does not always seem to be well integrated in the sense that different portfolios of stocks are priced with different implicit risk-free rates.

JEL Classification Numbers: G12, G15

Keywords: risk-free, intertemporal, asset, financial, market, expected, price, stock, Tokyo.

* 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. This work draws extensively from my ongoing research with Robert Flood. I thank him for our ongoing collaboration, Pedro Rodriguez for assistance with the data, and the Institute for Monetary and Economic Studies at the Bank of Japan 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.

1: Defining the Problem

The objective of this paper is to implement an intuitive and simple-to-use measure of asset-market integration on Japanese stock market data.

What does *asset-market integration* mean? I follow my (2003) work with Flood and 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:

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

where: p_t^j is the price at time t of asset j, $E_t()$ 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). I rely only on this standard and general intertemporal model of asset valuation; to my knowledge this Euler equation is present in all equilibrium asset-pricing models.

The 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 Japanese asset prices

and payoffs to characterize important aspects of its distribution, in particular its time-varying expectation.

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, I 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

I exploit the fact that in an integrated market, the MRS prices all assets held by the marginal asset holder. Indeed what I *mean* by asset market integration is that the same MRS prices all the assets. In other words, if one 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 my focus on the *expectation* of MRS, which is unique.

Consider a generic identity related to (1):

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

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

$$x^j_{t+1} = -[1/E_t(m_{t+1})]COV_t(m_{t+1}, x^j_{t+1}) + [1/E_t(m_{t+1})]\,p^j_t + e^j_{t+1}, \qquad \text{or}$$

$$x^j_{t+1} = d_t(p^j_t - COV_t(m_{t+1}, x^j_{t+1})) + e^j_{t+1} \tag{3}$$

where $d_t \equiv 1/E_t(m_{t+1})$ and $e^j_{t+1} \equiv x^j_{t+1} - E_t(x^j_{t+1})$, a prediction error.

I then impose two restrictions:

1) *Rational Expectations*: e^j_{t+1} is assumed to be white noise, uncorrelated with information available at time t, and
2) *Covariance Model*: $COV_t(m_{t+1}, x^j_{t+1}) = b^j_0 + \Sigma_i b^j_i f_{i,t}$, for the relevant sample,

where: b^j_0 is an asset-specific intercept, b^j_i is a set of I asset-specific factor coefficients and $f_{i,t}$ a vector of time-varying factors.

With just two assumptions, equation (3) becomes a panel estimating equation. I use *cross-sectional* variation to estimate $\{d\}$, 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 the study. I use *time-series* variation to estimate the asset-specific "fixed effects" and factor loadings $\{b\}$, 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,…,J_0 and then repeating the analysis for the same period of time with a different set of assets j=1,…,J_1 delivers two sets of estimates of $\{d\}$, 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 $\{d\}$ coefficients are equal.

Discussion

I make only two assumptions. Both are conventional in the literature, though most of the entire field uses stronger versions of them. While both assumptions can reasonably be characterized as "mild" in the area, it is worthwhile to elaborate on them further.

It seems unremarkable to assume that expectations are rational for financial markets, at least in the very limited sense above. I simply assume that asset-pricing errors are not *ex ante* predictable at high frequencies. This seems eminently reasonable.

The more controversial assumption is that the asset-specific covariances (of payoffs with the MRS) are either constant or depend on a small number of factors. Nevertheless, this is certainly standard practice. I use a model with two asset-specific effects; an intercept and a time-varying factor suggested by the famous Capital Asset Pricing Model (CAPM), namely the market return. I defend it on three grounds. First, in the application below, I need maintain the covariance model for only a month at a time. It seems intuitively plausible to imagine that the change in an asset's covariance structure does not change much from month to month, especially after its response to the market has been taken into account. Second, the literature also makes this assumption, but for much longer spans of time. For instance, Fama and French (1996) assumed that their model worked well for thirty years. Finally, I show below that the key results are insensitive to the exact factor model. This is important; if the technique were sensitive to the factors used to model $\{d\}$, then the integration 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.

While I focus on (3), there are other moments that would help characterize the MRS, $\{d\}$; see e.g., Hansen and Jagannathan (1991). I 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, the estimates and results are robust to the factor model that conditions the measurements. Finally, the measurements are discriminating for market integration.

The 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 discussion in e.g., Adam et. al. 2002). Second, standard asset-pricing models are completely consistent with the methodology, and the exact model does not seem to be important in practice. Third, I do not need to model the MRS directly. The MRS need not be determined uniquely, so long as its expectation is unique. Fourth, my strategy requires only two assumptions; I do not assume e.g., complete markets, homogeneous investors, or that I 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. 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 definition 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; the test presented below 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 recently by Adams et. al. (2002), Cochrane (2001), and Campbell, Lo, and MacKinlay (1997). 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 arbitrage pricing theory model, and a large literature has followed. Since I use parametric models to condition the estimation, the work presented below rests on this branch. It follows my (2003) paper with Flood that examines American data.

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.

Differences

My work differs from previous work in four broad ways.

First, I diverge from the finance profession in treating $\{b\}$ as a set of nuisance coefficients. Rather than being of intrinsic interest, they are required only to clear the way to produce estimates of the MRS. Indeed, they are not even really necessary at all, as I shall show below.

Next, I 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 I measure integration by the implied first moment of the stochastic discount rate (that is, the expected MRS). The condition I study, therefore, is a *necessary* condition for integration (though it is not sufficient). 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). I 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 I use is a free parameter, constant across assets but unconstrained across time. My measure is therefore basically cross-sectional, one that I can estimate precisely using a short time-series dimension.

Finally, I 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 government-backed treasury bill), equation (1) implies

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

where: i_t is a risk-less nominal interest rate, and m_{t+1} is a nominal MRS. The tradition inside domestic 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 I were to make this assumption, then it is trivial to estimate the risk-less MRS since $d_t \equiv 1/E_t(m_{t+1}) = (1+i_t)$. The strategy of this technique is not to *impose* this assumption; but rather to *test* it (and reject) it.

4: Empirics

The Factor Model

I begin by estimating a model with asset-specific intercepts and the time-varying market factor. In practice, I divide through by lagged prices (and redefine residuals appropriately):

$$x^j_{t+1} / p^j_{t-1} = d_t((p^j_t / p^j_{t-1}) + b^j_0 + b^j_1 f_{1,t}) + e^j_{t+1} \qquad (4)$$

for assets j=1,…,J, periods t=1,…,T. That is, I allow $\{d_t\}$ to vary period by period, while I use a "one-factor" model and allow $\{b^j\}$ to vary asset by asset. I normalize the data by lagged prices on the argument and belief that $COV_t(m_{t+1}, x^j_{t+1} / p^j_{t-1})$ can be modeled by a simple factor model

with time-invariant coefficients more plausibly than $COV_t(m_{t+1}, x^j_{t+1})$, and to ensure stationarity of all variables. The factors I use is the overall stock market return on the Nikkei 500 (in particular, the daily first-difference of the natural logarithm of the index). For sensitivity analysis, I also examine two other covariance models: one without the time-varying factor but including the asset-specific intercept; and the other asset-specific factors at all (i.e., without any covariance model).

Equation (4) can be estimated directly with non-linear least squares. The degree of non-linearity is not particularly high; conditional on $\{d_t\}$ the problem is linear in $\{b^j\}$ and vice versa. I use robust (heteroskedasticity and autocorrelation consistent "Newey West") covariance estimators.

I use a moderately high frequency approach. In particular, I use one-month spans of daily data. Using daily data allows me to estimate the coefficients of interest $\{d_t\}$ without assuming that firm-specific coefficients $\{b^j\}$ are constant for implausibly long periods of time.

The Data Set

My empirical work examines the integration of the Japanese equity market. Japanese stocks are traded on the Tokyo stock exchange, which seems to be a liquid market. Thus one can reasonably consider *a priori* the TSE to be integrated.

I examine daily data over twenty-five periods of one month each. In particular, I separately examine the months of April through August inclusive, for each of the five years 1998 through 2002 inclusive. Each month gives a span of between eighteen and twenty-three business day observations; this does not appear to stretch my reliance on a factor model of asset covariances excessively. However, it still allows me to test financial market integration for an

interesting set of data. I choose April through August to avoid the months when many Japanese companies pay dividends, and the five most recent years (for which the data are available). Still, there is no reason why either higher- and/or lower-frequency data could not be used.

My data set is drawn from Datastream, and is denominated in Japanese yen. Japanese holidays (Greenery Day, April 29; Constitution Memorial Day, May 3; National Holiday, May 4; Children's Day, May 5; and Marine Day, July 20) have been removed from the sample. I collected closing rates for the first (in terms of ticker symbol) 360 firms from the Tokyo Stock Exchange (TSE) that did not go ex-dividend during the months in question (371 was the minimum available for all samples). The absence of dividend payments allows me to set $x_{t+1}^{j} = p_{t+1}^{j}$ (and does not bias the results in any other obvious way). The data set has been checked for transcription errors via visual plotting.

Sorting

I group my 360 firms into twenty portfolios of eighteen firms each. I use three different techniques to create my portfolios to see if sorting has a strong effect on results. First, I use a random approach and group my firms simply on the alphabetical ordering of the firm's name. Second, I assign each of my firms to one of 31 industries, which are then arranged by SIC code; I then create portfolios (twenty for each year) grouped on the basis of industry affiliation. Third, I arrange each of my firms by gross assets for the year in question, and create portfolios (twenty for each year) grouped on the basis of size. Throughout, I weight each of the (eighteen) firms going into each of the (twenty) portfolios equally.

I 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!).

Since I lose the first and last observations of each year because of lags (p^j_{t-1}) and leads (x^j_{t+1}), I am left with a total of 10,540 observations in the panel data set (20 portfolios x 527 business days).[1] I repeat that there is no reason that one cannot use more data (either longer spans at lower frequencies, or shorter spans at higher frequencies)

5: Results

Estimates of the Marginal Rate of Substitution

I start by combining all 20 portfolios to estimate the time-varying marginal rate of substitution (i.e., estimates of $d_t \equiv [1/E_t(m_{t+1})]$). I provide time-series plots of the estimated deltas along with a plus/minus two standard error confidence interval in Figure 1, one graphic for each of the twenty-five months. (Each is estimated separately, so as to ensure that the portfolio-specific covariance models are not assumed to be constant for more than a month at a time).

----- INSERT FIGURE 1 AROUND HERE -----

[1] 1998 has 20 business-day observations for April, 19 for May, 22 for June, 22 for July, and 20 for August. For 1999, the totals are: 20, 18, 22, 21, and 21. For 2000, the totals are: 19, 20, 22, 20, and 22. For 2001, the totals are: 19, 21, 21, 21, and 22. For 2002, the totals are: 20, 21, 20, 23, and 21. Thus for 1998, I have 103 business days, for 1999 102, for 2000 103, for 2001 104, and 105 for 2002.

There is one striking feature of the graph. In particular, the time-series variation in delta is often quite high, consistent with the spirit of Hansen and Jagannathan (1991). While the discount rate moves around the value of unity, it fluctuates considerably. That is, the MRS does not seem to be close to constant at a daily basis. Further, this volatility does not seem to be constant over time. For example, May and August 1998 are unusually volatile while April 1998 is a relatively calm month for MRS. Since short-term interest rates in Japan during this period of time were quite low and stable, it seems easy to reject the hypothesis that the MRS derived from Japanese equity markets equals the short-term Japanese interest return.

Tests of Stock-Market Integration

It is inappropriate to dwell on the characteristics of Figure 1 at this point, since the graphics are implicitly based on the assumption that the Japanese stock market is integrated, and hence delivers a single estimate of $\{d_t\}$. Is the latter in fact true?

It is simple to test for stock-market integration using the strategy outlined above. One simply estimates $\{d_t\}$ from two different samples of assets over the same period of time, and compares them. Consider April 1998. When I estimate (4) from the first ten (randomly sorted) portfolios, I obtain a log-likelihood of 636.1. When I estimate precisely the same equation using data from the (mutually exclusive) other set of ten (randomly sorted) portfolios, I obtain a log-likelihood of 558.8. Finally, when I pool observations from all twenty portfolios, I obtain a log-likelihood of 1167. This combined estimate of (4) only differs from the two separate estimates of (4) in that a single vector of $\{d_t\}$ is estimated instead of two different estimates of the same vector (the portfolio-specific slopes and intercepts $\{b^j\}$ are unconstrained). If the TSE is integrated, the single combined estimate of $\{d_t\}$ should be equal to (and more efficiently

estimated than) the two different estimates of $\{d_t\}$. Statistically, under the hypothesis of normally distributed errors and integration, twice the difference between the separate and combined log-likelihoods is distributed as a chi-square with degrees of freedom equal to the dimensionality of $\{d_t\}$; a likelihood ratio (LR) test. But –2*((636.1+558.1)-1167)=55.9, which is inconsistent with the null hypothesis of integration and normally distributed error at the .00 significance level.

----- INSERT TABLE 1 AROUND HERE -----

Table 1 records the likelihood-ratio tests of integration inside the Japanese stock market. The top panel records 25 test statistics, one for each of the 25 different sample months. The lower the statistic, the more consistent with the joint null hypothesis of normally distributed errors and stock market integration. The p-values for the null hypothesis are tabulated beneath in square brackets. Many of the test statistics are too high (the p-values are too low), indicating rejection of the null hypothesis at conventional significance levels.

This is unfair to the hypothesis of market integration, since the null hypothesis is also being maintained jointly with the assumption of normality. 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. Accordingly, it is more appropriate to use a bootstrap procedure to estimate the probability values for the likelihood ratio tests.

The bootstrap procedure I employ is as follows. I estimate the deltas from all twenty portfolios, under the null hypothesis of integration. This gives me an estimate of $\{\varepsilon\}$. I then draw with randomly with replacement from this vector to create an artificial vector of $\{\varepsilon\}$ which

I then use to construct an artificial regressand variable {x}. Using this artificial data I 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. I then repeat this procedure a large number of times to generate an empirical distribution for the LR test statistic.

The bootstrapped p-values for the test of integration are tabulated in parentheses below the likelihood-ratio test statistics (low values are inconsistent with the null hypothesis, indicating rejection of integration). Rejections of the null hypothesis of integration that are significant at the .05 (.01) level are marked with an (two) asterisk(s).

The top panel of Table 1 shows that the hypothesis of integration works better with bootstrapped than normal confidence intervals. Nevertheless, there are still two individual months that reject the hypothesis of integration at the .01 level, and another couple that are inconsistent at the .05 level.

The middle panel of Table 1 is the analogue to the top panel, but is produced using industry-sorted portfolios rather than randomly allocated portfolios. The bottom panel uses size-based portfolios but is otherwise analogous to the top. Sorting makes a difference! When industry-based portfolios are used, twelve of the test statistics are inconsistent with integration at the .01 level, and another four at the .05 level. The rejections are spread across all years and months. Size-based portfolios give a similar story; nine of the twenty-five test statistics reject integration at the .01 level, and another four at the .05 level. That is, the hypothesis that the Japanese stock market is integrated in the sense that it delivers the same estimated risk-less discount rate is often rejected. This is especially true when the data are sorted by industry, but it is also true of portfolios that have been randomly created and portfolios formed on the basis of size.

Sensitivity Analysis: The Covariance Model

Table 1 provides likelihood ratio tests of market integration when tests are estimated on the basis of (4). In this model, each portfolio of stocks is modeled with a portfolio-specific intercept (b_0^j) and a slope for the response of the portfolio to the aggregate stock market (b_1^j). Still, this may be an inappropriate way to model the covariance of the portfolio's payoff with the MRS. It is therefore appropriate to examine the sensitivity of results with respect to the factor model I use for the covariances.

I examine two restrictions of (4):

$$x_{t+1}^j / p_{t-1}^j = d_t((p_t^j / p_{t-1}^j) + b_0^j) + e_{t+1}^j \qquad (4')$$

and

$$x_{t+1}^j / p_{t-1}^j = d_t(p_t^j / p_{t-1}^j) + e_{t+1}^j \qquad (4")$$

Of these, only the first is a "serious" model; the second model contains no asset-specific covariance terms at all. The first model simply throws away the responsiveness of the portfolio to movements in the aggregate stock market. This is admissible if a static Capital Asset Pricing Model (CAPM) works well, so that the covariance between the discount rate and the payoff is simply an asset-specific constant.

Table 2 is an analogue to Table 1 that uses (4'), so that it restricts $(b_1^j = 0)$. As in Table 1, likelihood ratio test statistics are tabulated, along with bootstrapped p-values for the null hypothesis of market integration. It seems that allowing the portfolios covariances to depend on the Nikkei's market return makes little difference. Using random portfolios, one of the (25) samples is inconsistent with integration at the .01 level, and another four at the .05 level (in Table 1, the analogous figures were two and two). Fourteen of the industry based portfolio tests reject integration at .01, and three at .05 (twelve and four in Table 1); eight (at .01) and four (at .05) using size-based portfolios (nine and four in Table 1).

----- INSERT TABLE 2 AROUND HERE -----

Table 3 is an analogue that uses (4"), i.e., it restricts $(b_0^j = b_1^j = 0)$ and thus employs no covariance model at all. Again, the results are largely unchanged. Indeed, many of the actual test statistic change very little! This robustness is encouraging since it demonstrates the insensitivity of the methodology to reasonable perturbations in the exact factor model employed. Succinctly, the exact factor model used to model the covariances seems to make little difference to the results. Indeed, completely dropping the factor model seems to make little difference.

----- INSERT TABLE 3 AROUND HERE -----

6: Summary and Conclusions

This paper developed a simple method to test for asset integration, and then applied it within the Japanese equity market. It relies on estimating and comparing the expected risk-less returns implied by different sets of assets. My 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.

I illustrated this technique with an application to Japanese stocks. I use daily data from 360 equities traded on the Tokyo stock exchange, and examine twenty-five separate months, April through August, 1998 through 2002. I sort the stocks into portfolios in three different ways: randomly, and grouped by both industry and size. I find that the time-series variation in the Marginal Rate of Substitution is often high on a daily basis. Even more strikingly, I find that Japanese stocks are not integrated in the sense that the risk-less discount rates implied by different portfolios differ by statistically significant amounts. That is, the TSE does not always seem to be integrated. This result is found most strongly on the basis of portfolios sorted by industry, but is present in both randomly- and size-based portfolios, and seems insensitive to the exact factor model employed.

If my finding of lack of integration within the Japanese stock market holds up to further scrutiny, the interesting question is not whether this market with few apparent frictions is poorly integrated but why? I 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.

Flood, Robert P. and Andrew K. Rose (2003) "Financial Integration: A New Methodology and an Illustration" unpublished.

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.

	April	May	June	July	August
Random Portfolios					
1998	55.9*	16.0	131.1**	26.6	43.0
	(.03)	(.91)	(.00)	(.65)	(.08)
	[.00]	[.65]	[.00]	[.23]	[.00]
1999	26.8	16.7	61.5	49.4*	24.2
	(.47)	(.87)	(.08)	(.02)	(.81)
	[.14]	[.55]	[.00]	[.00]	[.28]
2000	28.3	23.7	32.4	47.7**	43.5
	(.34)	(.58)	(.32)	(.00)	(.10)
	[.08]	[.26]	[.07]	[.00]	[.00]
2001	35.3	40.4	41.0	33.9	37.1
	(.12)	(.07)	(.08)	(.20)	(.23)
	[.01]	[.01]	[.01]	[.04]	[.02]
2002	21.8	14.5	21.3	18.4	17.2
	(.69)	(1.0)	(.69)	(.95)	(.93)
	[.37]	[.85]	[.38]	[.57]	[.70]
Industry-Based Portfolios					
1998	85.5**	56.4**	143.0**	71.4**	94.5**
	(.00)	(.00)	(.01)	(.01)	(.00)
	[.00]	[.00]	[.00]	[.00]	[.00]
1999	41.6	34.9	88.3*	24.7	24.0
	(.06)	(.10)	(.04)	(.60)	(.81)
	[.00]	[.01]	[.00]	[.27]	[.29]
2000	75.8**	101.4**	79.3**	35.4	59.9**
	(.00)	(.00)	(.00)	(.16)	(.00)
	[.00]	[.000	[.00]	[.02]	[.00]
2001	48.5**	41.0*	71.3**	38.3	50.5*
	(.00)	(.05)	(.00)	(.10)	(.02)
	[.00]	[.01]	[.00]	[.02]	[.00]
2002	48.4**	39.4**	37.4	21.8	30.6
	(.03)	(.00)	(.15)	(.83)	(.33)
	[.00]	[.01]	[.01]	[.37]	[.08]
Size-Based Portfolios					
1998	66.3**	89.6**	61.6*	47.2	112.5**
	(.00)	(.00)	(.02)	(.09)	(.00)
	[.00]	[.00]	[.00]	[.00]	[.00]
1999	94.0**	123.0**	48.1	119.5**	138.5**
	(.00)	(.00)	(.13)	(.00)	(.00)
	[.00]	[.00]	[.00]	[.00]	[.00]
2000	38.7*	16.5	29.9	32.4*	38.8
	(.04)	(.91)	(.43)	(.02)	(.12)
	[.01]	[.69]	[.13]	[.04]	[.02]
2001	32.2	75.7**	30.9	48.2*	28.0
	(.20)	(.00)	(.33)	(.02)	(.58)
	[.03]	[.00]	[.08]	[.00]	[.18]
2002	48.1	58.4**	34.3	30.2	38.1
	(.09)	(.00)	(.20)	(.52)	(.11)
	[.00]	[.00]	[.03]	[.09]	[.02]

Table 1: Likelihood-Ratio Tests for Integration on Tokyo Stock Exchange
Covariance Model includes portfolio-specific intercepts and market factors

(bootstrapped p-values for null hypothesis in parentheses) * (**) indicates rejection of null hypothesis at .05 (.01)
[normal p-values for null hypothesis in brackets]

	April	May	June	July	August
Random Portfolios					
1998	54.0*	14.0	133.8**	36.7	24.9
	(.03)	(.92)	(.00)	(.23)	(.62)
1999	25.9	12.4	55.3	39.7	24.1
	(.41)	(.92)	(.10)	(.08)	(.67)
2000	23.4	23.9	31.8	43.6	46.6
	(.49)	(.50)	(.24)	(.03)	(.02)
2001	31.6	22.4	40.4*	38.1	37.8
	(.18)	(.68)	(.03)	(.09)	(.15)
2002	20.6	13.7	19.3	13.2	17.3
	(.69)	(.94)	(.77)	(1.0)	(.85)
Industry-Based Portfolios					
1998	84.2**	55.2**	147.0**	98.1**	96.8**
	(.00)	(.00)	(.00)	(.00)	(.00)
1999	31.9	44.8**	92.0**	30.7	23.6
	(.19)	(.01)	(.01)	(.22)	(.75)
2000	69.8**	97.4**	80.7**	33.8	58.4**
	(.00)	(.00)	(.00)	(.13)	(.00)
2001	33.0	40.4	67.7**	42.7	50.6
	(.12)	(.06)	(.00)	(.03)	(.04)
2002	51.8**	40.3**	42.7	20.0	31.0
	(.00)	(.00)	(.05)	(.85)	(.22)
Size-Based Portfolios					
1998	58.9**	65.5**	57.8*	31.3	104.1**
	(.00)	(.00)	(.02)	(.40)	(.00)
1999	93.1**	106.7**	43.7	112.3**	136.7**
	(.00)	(.00)	(.18)	(.00)	(.00)
2000	34.3	15.7	27.4	27.4	42.9*
	(.07)	(.93)	(.45)	(.36)	(.05)
2001	26.0	62.6**	29.3	41.8	27.4
	(.37)	(.01)	(.24)	(.09)	(.44)
2002	44.4*	46.6**	26.0	34.9	37.0
	(.05)	(.01)	(.56)	(.26)	(.13)

Table 2: Likelihood-Ratio Tests for Integration on Tokyo Stock Exchange
Covariance Model includes only portfolio-specific intercepts

(bootstrapped p-values for null hypothesis in parentheses)

* (**) indicates rejection of null hypothesis at .05 (.01)

	April	May	June	July	August
Random Portfolios					
1998	59.8* (.02)	17.6 (.87)	121.1** (.00)	33.7 (.35)	25.1 (.58)
1999	24.6 (.54)	12.0 (1.00)	50.6 (.15)	32.8 (.26)	23.0 (.84)
2000	24.8 (.54)	20.8 (.80)	30.8 (.40)	43.6* (.02)	40.2 (.14)
2001	29.3 (.27)	27.2 (.50)	39.5 (.13)	32.9 (.25)	37.2 (.18)
2002	19.0 (.81)	14.0 (.98)	18.1 (.88)	13.7 (.99)	18.2 (.87)
Industry-Based Portfolios					
1998	85.8** (.00)	57.2** (.01)	145.4** (.00)	91.3** (.00)	99.1** (.00)
1999	30.5 (.32)	47.2** (.01)	83.8** (.01)	24.8 (.57)	19.9 (.92)
2000	72.9** (.00)	99.7** (.00)	72.7** (.00)	33.0 (.18)	63.0** (.00)
2001	28.3 (.36)	34.1 (.23)	63.3** (.01)	39.7 (.09)	52.6* (.02)
2002	53.1* (.02)	40.6 (.07)	41.9 (.09)	20.2 (.89)	30.0 (.39)
Size-Based Portfolios					
1998	37.4 (.15)	58.0** (.01)	53.2* (.02)	33.8 (.41)	49.3* (.03)
1999	99.1** (.00)	88.9** (.00)	44.6 (.20)	107.8** (.00)	134.6** (.00)
2000	34.9 (.10)	14.3 (.97)	27.2 (.59)	24.4 (.51)	42.9 (.07)
2001	22.4 (.60)	60.0** (.00)	25.3 (.54)	41.8 (.08)	24.3 (.75)
2002	42.4 (.09)	46.2* (.02)	24.9 (.58)	32.0 (.46)	37.2 (.13)

Table 3: Likelihood-Ratio Tests for Integration on Tokyo Stock Exchange
Covariance Model includes no portfolio-specific features

(bootstrapped p-values for null hypothesis in parentheses)

* (**) indicates rejection of null hypothesis at .05 (.01)

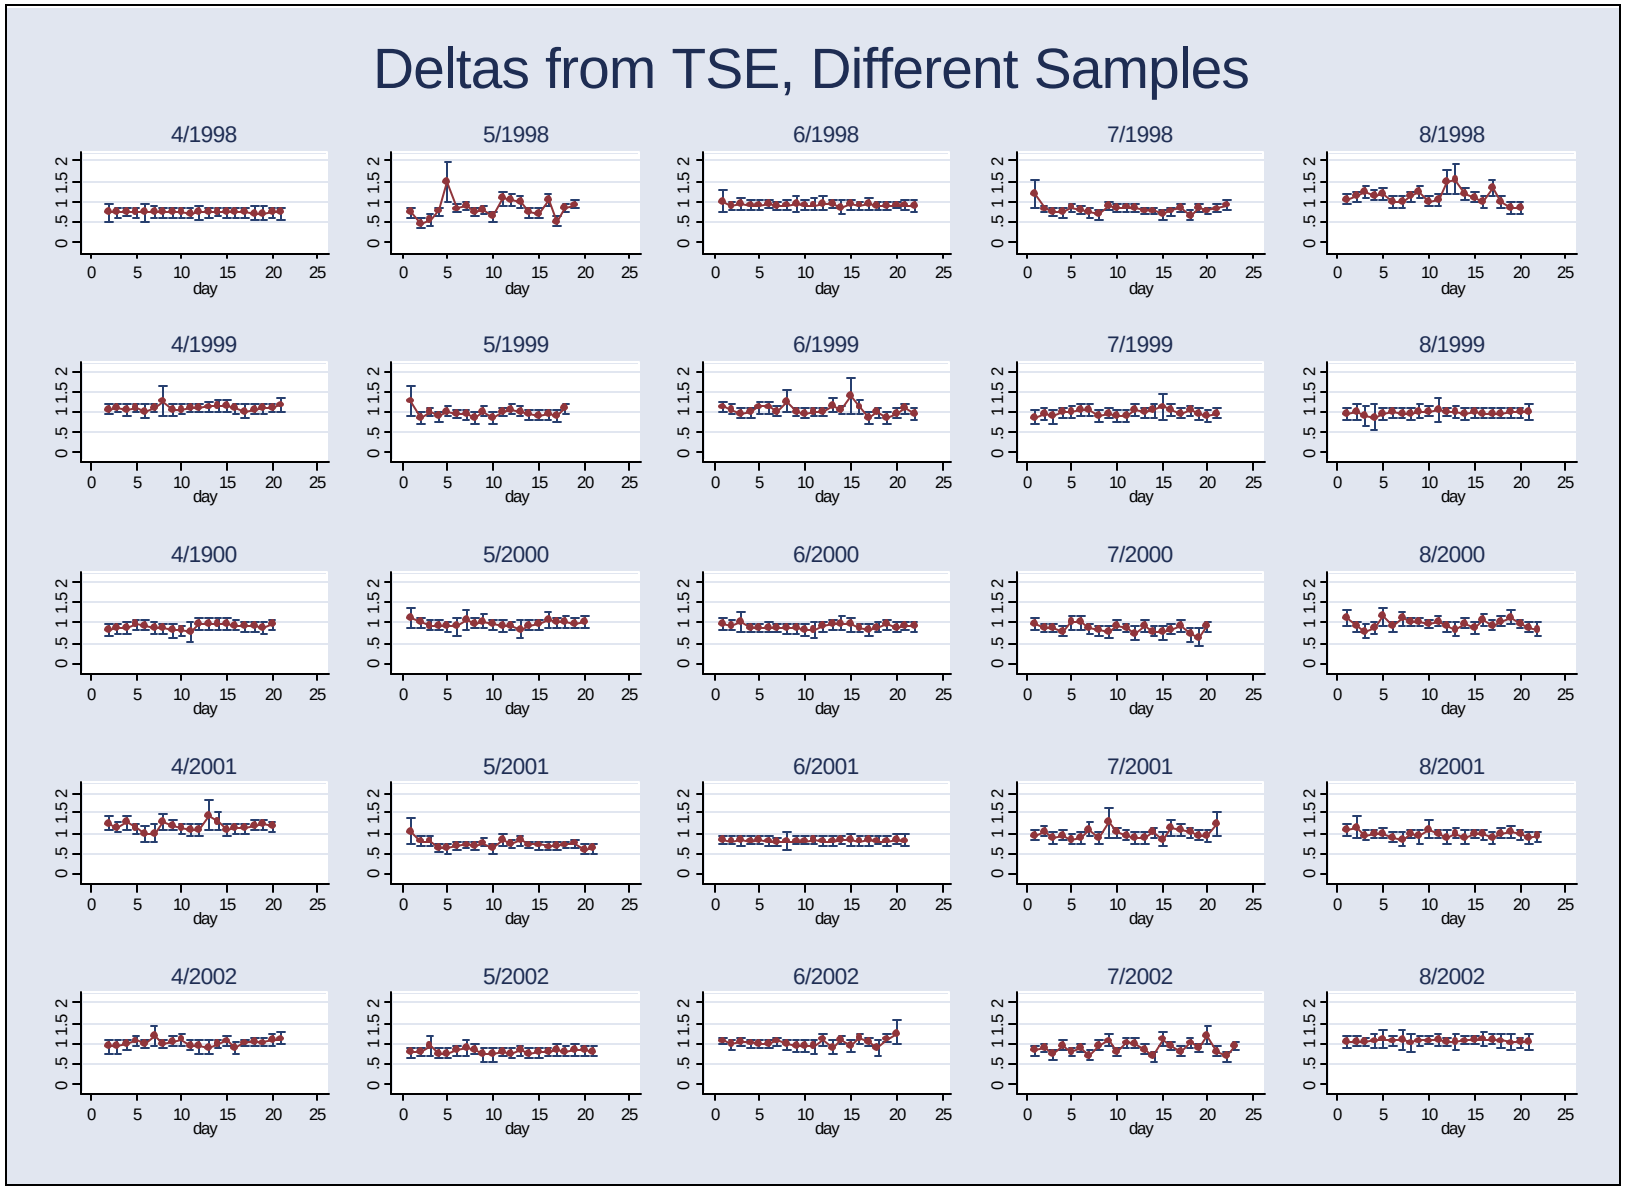

Figure 1: Estimates of Marginal Rate of Substitution from sets of (20) portfolio