Equity Integration in Times of Crisis

Robert P. Flood and Andrew K. Rose*

Draft: August 21, 2003
Comments Welcome

For *Market Discipline: The Evidence Across Countries and Industries* Chicago, 10/2003

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

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

Abstract

This paper applies a simple new test for asset integration to two episodes of crises in financial 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. The S&P 500 market seems to be generally well integrated, but the level of integration falls temporarily during the LTCM crisis of October 1998. By way of contrast, the Korean stock market remains generally internally integrated through the Asian crisis of 1997. The level of equity integration between Japan and Korea is low and falls further during late 1997.

JEL Classification Numbers: G14

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

* 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. Rose thanks the Bank of Japan for hospitality during the course of this research. The data sets, sample E-Views programs and output, and a current version of this paper are available at http://faculty.haas.berkeley.edu/arose.

The objective of this paper is to implement an intuitive and simple-to-use measure of asset-market integration on stock market data. In particular, we are interested in asking whether equity market performance is affected by financial crises. We examine two important financial crises: the aftermath of Russia/LTCM in the autumn of 1998, and the Asian crisis of late 1997. In both cases, we find that stock market integration is adversely affected by the crisis, but only temporarily.

1: Defining the Problem

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:

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

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). 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. We are especially interested in whether the properties of the MRS are systematically different during periods of financial turbulence.

The substantive point of equation (1) is that all assets share the same marginal rate of substitution. There is no asset-specific MRS in an integrated market (indeed, that's our definition of a "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 use 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. We focus in particular on whether market integration varies systematically between periods of market tranquility and crisis.

2: Methodology

We use the fact that in a well-functioning 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):

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

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

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

$$x_{t+1}^j = \delta_t(p_t^j - COV_t(m_{t+1}, x_{t+1}^j)) + \varepsilon_{t+1}^j \tag{3}$$

where $\delta_t \equiv 1/E_t(m_{t+1})$ and $\varepsilon_{t+1}^j \equiv x_{t+1}^j - E_t(x_{t+1}^j)$, a prediction error.

We then impose two restrictions:

1) *Rational Expectations*: ε_{t+1}^j is assumed to be white noise, uncorrelated with information available at time t, and
2) *Covariance Model*: $COV_t(m_{t+1}, x_{t+1}^j) = \beta_0^j + \Sigma_i \beta_i^j f_{i,t}$, for the relevant sample,

where: β_0^j is an asset-specific intercept, β_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 $\{\delta\}$, the coefficients of interest that represent the expected 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 $\{\beta\}$, 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 $j=1,\ldots,J_0$ and then repeating the analysis for the same period of time with a different set of assets $j=1,\ldots,J_1$ 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, either one by one, or jointly. Under the null hypothesis of market integration, the two sets of $\{\delta\}$ coefficients are equal. We choose our data to focus on periods of time where integration may be expected to break down, that is, periods of financial stress.

Discussion of the Technique

We 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. We simply assume that asset-pricing errors are not *ex ante* predictable at high frequencies. This seems eminently reasonable, even during periods of crisis.

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. Rather than develop our own factor model, we start our analysis by relying on the well-known three-factor model famously deployed by Fama and French (1996). (For our Asian empirics, we 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.) We defend our strategy on three grounds. First, in the applications below, we need maintain the covariance model for only a month or two at a time. It seems intuitively plausible to imagine that the change in an asset's covariance structure does not change much for samples of this length. 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, we 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 $\{\delta\}$, 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 we focus on (3), there are other moments that would help characterize the MRS, $\{\delta\}$; 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. In the examples below, our measure never rejects internal market integration for portfolios of S&P stocks priced in the NYSE during periods of tranquility, or for Korean stocks priced in Seoul. Nevertheless, we *can* reject integration for short periods of financial stress in the NYSE. Further, we reject integration strongly – by an order of magnitude – between the Korean and Japanese markets.

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. 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 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.

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 $\{\beta\}$ 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 mode; 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

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

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, $\delta_t \equiv 1/E_t(m_{t+1}) = (1+i_t)$.
We do not *impose* this assumption; rather we (implicitly) test it (and reject) it.

4: The Russia/LTCM Crisis of 1998

The New York Stock Exchange can reasonably be considered to be the largest and most liquid stock exchange in the world, certainly for the last decade. It is thus a demanding forum in which to check whether our technique can reveal signs of financial turbulence, and where we begin.

The financial crisis of 1998 began with the Russian devaluation and unilateral debt restructuring of August 18 1998. Interest rate spreads for emerging markets and much of the corporate sector rose thereafter, especially following the near failure of Long-Term Capital Management (LTCM). The rescue of LTCM was announced on September 23. The crisis subsided after a loosening of monetary policy throughout much of the OECD that began in late September with an American monetary policy loosening on September 29. This was widely criticized as being inadequate, and was followed on October 15 and November 17 by others.[1]

The American stock market peaked in mid-July and bottomed out in late August, about twenty percent below its peak (though volatility was particular high from mid-September through mid-October). Still, much of the action had been in bond markets. Between mid-August and early October, government bond yields fell by about 110 basis points (implying price gains of between 6-11% for benchmark 7-10 year bonds). A number of spreads also rose during this period, including the spread between "on-the-run" and "off-the-run" treasuries (a standard measure of liquidity; this widened from less than 10 basis points to over 35 basis points in mid-October).

We now examine at some length the integration of the US stock market during this interesting episode of turbulence.

Implementing our Technique

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):

$$x_{t+1}^{j} / p_{t-1}^{j} = \delta_t ((p_t^j / p_{t-1}^j) + \beta_0^j + \beta_1^j f_{1,t} + \beta_2^j f_{2,t} + \beta_3^j f_{3,t}) + \varepsilon_{t+1}^j \tag{4}$$

for assets j=1,…,J, periods t=1,…,T. That is, 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. 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 $COV_t(m_{t+1}, x_{t+1}^j)$, and to ensure stationarity of all variables. 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.[2] For sensitivity analysis, we also examine other simpler covariance models.

Equation (4) 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 one-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. Still, we see no reason why higher- (and/or lower-) frequency data could not be used.

The Data Set

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 and twenty firms from the S&P 500 that did not go ex-dividend during the months in question (September through November). 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 have checked the data set for outliers and excluded holidays.

We are most interested in the period of financial turbulence of October 1998. Still, so as to generate a reasonable comparison set, we separately examine September, October, and November. Also, we perform our analysis for the same three months for four consecutive years: 1996 through 1999. We go back two years rather than one since the stock market turbulence of October 1997 may be of independent interest (the Dow Jones industrial average posted its biggest loss ever on Oct 27, 1997). Thus we use comparable months from 1996 and 1998 as our "control sample" with which to compare our results for 1997 and especially 1998, and also compare October with September and November.

We group our 120 firms into twenty portfolios of six equally weighted 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!).

Results

We start by combining data from all 20 portfolios to estimate the time-varying marginal rate of substitution (i.e., estimates of $\delta_t \equiv [1/E_t(m_{t+1})]$). We 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 ten months. (We do not provide graphics for October 1998 and 1999 for reasons discussed below.) Each month is estimated separately, so as to ensure that the portfolio-specific covariance models are assumed to be constant for only a month at a time.

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

There is one striking feature of the graph. In particular, the time-series variation in delta is 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, the stock market turbulence of October 1997 shows up clearly. Since short-term interest rates in the United States during this period of time were quite low and stable, it is easy to reject the hypothesis that the MRS derived from American equity markets equals the sluggish and low (but positive) short-term American interest return.

Still, 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 American stock market is integrated, and hence delivers a single estimate of $\{\delta_t\}$. Is the latter in fact true?

It is simple to test for stock-market integration using the strategy outlined above. One simply estimates $\{\delta_t\}$ from two different samples of assets over the same period of time, and compares them. Consider September 1996. When we estimate (4) from the first ten portfolios, we obtain a log-likelihood of 631.79. When we estimate precisely the same equation using data from the (mutually exclusive) other set of ten portfolios, we obtain a log-likelihood of 602.95. Finally, when we pool observations from all twenty portfolios, we obtain a log-likelihood of 1217.39. This combined estimate of (4) only differs from the two separate estimates of (4) in that a single vector of $\{\delta_t\}$ is estimated instead of two different estimates of the same vector (the portfolio-specific slopes and intercepts $\{\beta^j\}$ are unconstrained). If the NYSE is integrated, the single combined estimate of $\{\delta_t\}$ should be equal to (and more efficiently estimated than) the two different estimates of $\{\delta_t\}$. Statistically, under the hypothesis of normally distributed errors and integration, twice the difference between the combined and separate log-likelihoods is distributed as a chi-square with degrees of freedom equal to the dimensionality of $\{\delta_t\}$; a likelihood ratio (LR) test. Since there were 19 business days in the month, –2*[(631.79+602.95) - 1217.39] = 34.7 should be drawn from the $\chi^2(19)$ distribution under the null of integration and normally distributed errors. As 34.7 is at the .02 tail of the $\chi^2(19)$ distribution, the null hypothesis of integration and normally distributed error is (marginally) rejected.

Such an interpretation is unfair to the hypothesis of market integration, since it is 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. We estimate the deltas from all twenty portfolios, under the null hypothesis of integration. This gives us an estimate of $\{\varepsilon\}$. We then draw with randomly with replacement from this vector to create an artificial vector of $\{\varepsilon\}$ which we then 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 an empirical distribution for the LR test statistic.

Table 1a records the likelihood-ratio tests of integration inside the American stock market. The top panel records 12 test statistics, one for each of the different sample months. The lower the statistic, the more consistent with the null hypothesis of stock market integration. The bootstrapped p-values for the hypothesis are tabulated beneath in parentheses. Only two of the test statistics are high (the p-values are too low), indicating rejection of the null at conventional significance levels; the Octobers of 1998 and 1999.[3] That is, the NYSE seems to be integrated except for the crisis period of October 1998 and also October 1999.

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

Sensitivity Analysis

Two questions of interest emerge immediately. The first is the curious rejection of integration for October 1999; the second is the more important one, and concerns the sensitivity of our results to the precise covariance model employed. We handle them simultaneously.

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 (β_0^j) and factor loadings on 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 1b we provide test statistics (and bootstrapped p-values) to examine tests of integration within the S&P, 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, with one exception; October 1999. For that sample period, the high test statistic falls to a level wholly consistent with integration. In fact, the high value for 10/99 in Table 1a depends critically on an anomalous movement of the SML (returns of small stocks relative to those of large stocks) over two consecutive days in mid-October 1999. Since this is a sensitive finding that stems from a non-traditional factor, we do not take it that seriously. And it is reassuring to us that the vast majority of our findings are robust with respect to the exact covariance model deployed.

Indeed, our findings are more robust than implied by Tables 1a and 1b. Table 1c goes further and drops the market factor from our covariance model, leaving only portfolio-specific intercepts (β_0^j) but no time-varying factors. Again, the results are essentially unchanged (indeed, the actual values of the test statistics are little affected!). Finally, Table 1d is the analogue when no portfolio-specific terms at all are used, and again the results are similar. This robustness is encouraging since it demonstrates the insensitivity of our methodology to even drastic perturbations in the exact factor model employed.

We conclude that the American stock market seems to have been dramatically affected by the crisis of October 1998. Compared with the months and years immediately before and after the period, there is a substantive but transient drop in market integration. This finding seems both sensible and robust.

5: The Asian Crisis

We now apply our methodology to examine the Asian crisis of 1997.

The Asian crisis is often viewed as starting in early July 1997, and subsiding in early 1998. We focus on Korea, the largest of the directly affected economies, and one of the most severely affected. While bad Korean economic news arrived in the spring and summer of 1997, the most dramatic events took place in the late Fall of 1997.[4] After a relatively calm summer and autumn (certainly compared with events further South), the won was attacked beginning on November 6. This quickly lead to a collapse of the Korean stock market, and rippling effects elsewhere in Asia and abroad. The IMF package of early December 1997, international assistance and a decisive Korean policy response started to turned things around by the end of the year, though the real effects continued through 1998.

Our Asian data set is drawn from Datastream. Since Tokyo and Seoul are in the same time-zone, we collected closing rates for the first (in terms of ticker symbol) 400 stocks from the

Tokyo stock exchange and the first 360 stocks from the Korean stock exchange that did not go ex-dividend during the months in question (November and December, 1996 through 1998). We converted Korean prices in won into Japanese yen using a matched nominal exchange rate. Again, we have checked the data set for outliers and excluded holidays, and again we group our firms into twenty portfolios each for both Japan and Korea, arranged simply by ticker symbol. As the single time-varying factor, we use the first-difference in the natural logarithm of the stock market index (the Nikkei 500 for Tokyo, and the Korea SE Composite or "KOSPI" for Seoul). So for each of our (twenty Japanese and twenty Korean) portfolios, the covariance of payoffs with the discount rate is parameterized by a portfolio-specific intercept and also a sensitivity with respect to the appropriate (Japanese/Korean) market return. That is, we estimate the following equation:

$$x_{t+1}^{j} / p_{t-1}^{j} = \delta_t((p_t^j / p_{t-1}^j) + \beta_0^j + \beta_1^j f_{1,t}) + \varepsilon_{t+1}^j \tag{4'}$$

Estimates of the marginal rate of substitution (i.e., $\{\delta_t\}$) for Japan and Korea are provided in Figure 2. For each country, three time-series plots of the estimated deltas along with a plus/minus two standard error confidence interval is provided, one graphic for each year. Each plot thus spans a two-month interval (though the plots look similar when they are estimated on samples of a single month).

----- INSERT FIGURE 2 AROUND HERE -----

As with Figure 1, the time-series volatility in delta is high. This is especially true of 1997 (the Korean crisis of December 1997 is readily apparent), though the other years also show a lot of variation. Still, an even more striking feature of the graphics in Figure 2 is that the Japanese and Korean sets of $\{\delta_t\}$ do not look similar. This is especially true of 1997, though the 1998 series also look quite different. Of course, if the two markets are integrated then only a single expected marginal rate of substitution should characterize both markets.

To confirm this ocular impression, we now compute analogues to the likelihood ratio tests of integration presented in Table 1. Where Table 1 focused on the internal integration of American stocks traded on the NYSE, Tables 3-4 check the cross-market integration of Japanese and Korean stocks.

Before we examine the degree of international integration, it is important to check for the degree of internal (domestic) financial integration. After all, if American stock exchanges can be disrupted by financial crises, less liquid Asian ones can be as well! Thus Table 2 is an analogue to Table 1 in that it tests for internal integration of the Korean stock exchange during the height of the financial crisis of 1997, as well as the years before and after. Surprisingly (at least to us), there is no indication of any breakdown in financial integration from the likelihood ratio tests (which again use bootstrapped p-values). This is true when we examine November and December separately, and also when we pool the data across the two months.

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

Still, the real question of interest to us concerns international integration; do the Korean and Japanese stock exchanges share a common expected MRS? As we expect from Figure 2, the answer is decisively negative in Table 3. In this table, we estimate (4') but imposing equality across the two sets of {δ}. (The portfolios of e.g., Japanese stocks, are each parameterized with a portfolio-specific intercept, and a coefficient on the aggregate Japanese index.) The likelihood ratio tests of market integration are an order of magnitude larger than those in Table 1 or Table 2, and reject the hypothesis of market integration at all reasonable significance levels. They jump considerably during the crisis of 1997, but from already enormous levels in 1996. They then fall in 1998, but only back to the huge levels experienced before. That is, the Korean and Japanese stock markets become less integrated during the financial crisis of late 1997, but they are never integrated during the period we consider.

----- INSERT TABLES 3 AND 4 AROUND HERE -----

Under the null hypothesis of integration, (4') is not strictly correct since both Japanese and Korean stocks should be affected by both the aggregate Japanese and Korean markets (or a composite of the two). We allow for this by estimating:

$$x_{t+1}^{j} / p_{t-1}^{j} = \delta_t((p_t^j / p_{t-1}^j) + \beta_0^j + \beta_1^j f_{1,t} + \beta_2^j f_{2,t}) + \varepsilon_{t+1}^j \qquad (4'')$$

where the second factor represents the foreign aggregate market return (Korea for the Japanese portfolios, Japan for the Koreans). The analogues for this model are tabulated in Table 4 and delivers similar point-estimates and identical conclusions to those of Table 3. This is no surprise, since the message throughout is that the hypothesis of international integration is grossly at odds with the data.

6: Summary and Conclusions

This paper developed a simple method to test for asset integration, and then applied it within and between equity markets during times of crises. The technique 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.

If our integration findings hold up to further scrutiny, the interesting question is not whether financial markets with few apparent frictions (such as the Seoul and Tokyo stock exchanges) are poorly integrated but why? What are the mechanisms that seem to break down during periods of market stress such as those that affected the New York Stock Exchange during the financial crisis of 1998? We leave such important questions 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.

International Monetary Fund (various) *World Economic Outlook.*

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

A: Fama-French-Factor Model (intercepts, 3 time-varying factors)

	September	October	November
1996	34.7 (.33)	26.6 (.83)	32.3 (.40)
1997	39.7 (.17)	37.5 (.39)	32.2 (.05)
1998	34.4 (.40)	55.5 (.02)	27.6 (.56)
1999	16.6 (.97)	57.1 (.00)	30.3 (.43)

B: One-Factor Model (intercepts, market return factor)

	September	October	November
1996	24.2 (.55)	25.9 (.64)	41.2 (.05)
1997	29.6 (.34)	36.3 (.26)	29.7 (.19)
1998	25.4 (.49)	53.1 (.01)	22.6 (.61)
1999	15.9 (.95)	24.7 (.65)	26.3 (.42)

C: Model without Time-Varying Factors (intercepts only)

	September	October	November
1996	20.0 (.67)	25.8 (.59)	28.1 (.31)
1997	21.1 (.66)	38.3 (.66)	28.5 (.22)
1998	25.1 (.42)	54.2 (.01)	21.4 (.55)
1999	12.9 (.95)	26.4 (.50)	23.6 (.47)

D: Model without Asset-Specific Covariances

	September	October	November
1996	20.0 (.65)	24.8 (.61)	27.8 .25)
1997	20.6 (.59)	36.8 (.17)	31.1 (.10)
1998	20.9 (.60)	52.3 (.00)	22.4 (.43)
1999	10.9 (.98)	28.5 (.34)	22.2 (.47)

Table 1: Integration inside the American S&P 500

Likelihood-ration test statistics (bootstrap P-value)

	November	December	November-December
1996	32.8 (.19)	27.3 (.32)	50.8 (.26)
1997	33.3 (.20)	27.0 (.29)	51.6 (.23)
1998	54.1 (.13)	10.5 (.97)	31.7 (.92)

Table 2: Integration inside the Korean Stock Exchange
One-Factor Model (intercepts, market return factor)
Likelihood-ration test statistics (bootstrap P-value)

	November	December	November-December
1996	389.9	259.2	640.3
1997	639.1	1716.2	2480.5
1998	269.3	591.3	876.3

Table 3: Integration between Korea and Japan
One-Factor Model (intercepts, domestic market return factor)
Likelihood-ration test statistics (all p-values =.00)

	November	December	November-December
1996	439.0	261.8	645.6
1997	626.5	1595.1	2401.1
1998	284.7	519.0	814.3

Table 4: Integration between Korea and Japan
Two-Factors Model (intercepts, domestic and foreign market return factors)
Likelihood-ration test statistics (all p-values =.00)

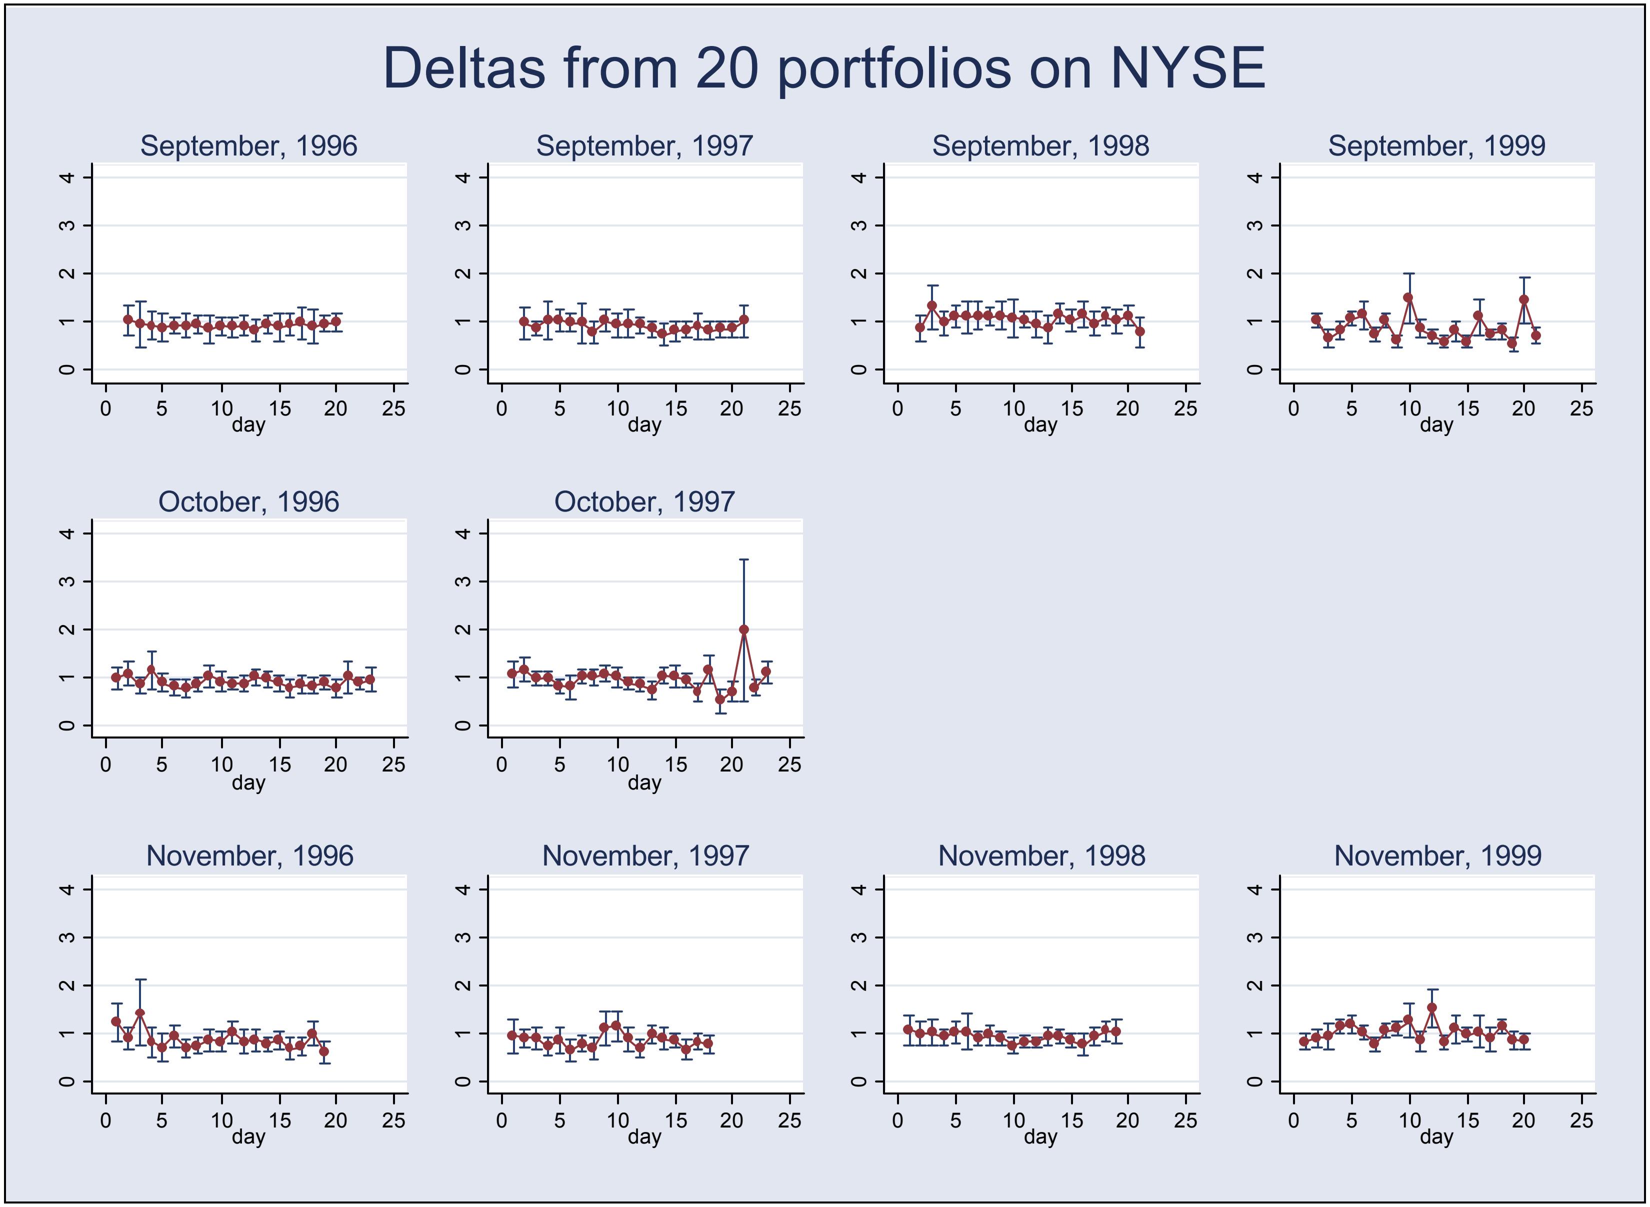

Figure 1: Estimates of Expected Discount Rate from LTCM/Russia Crisis

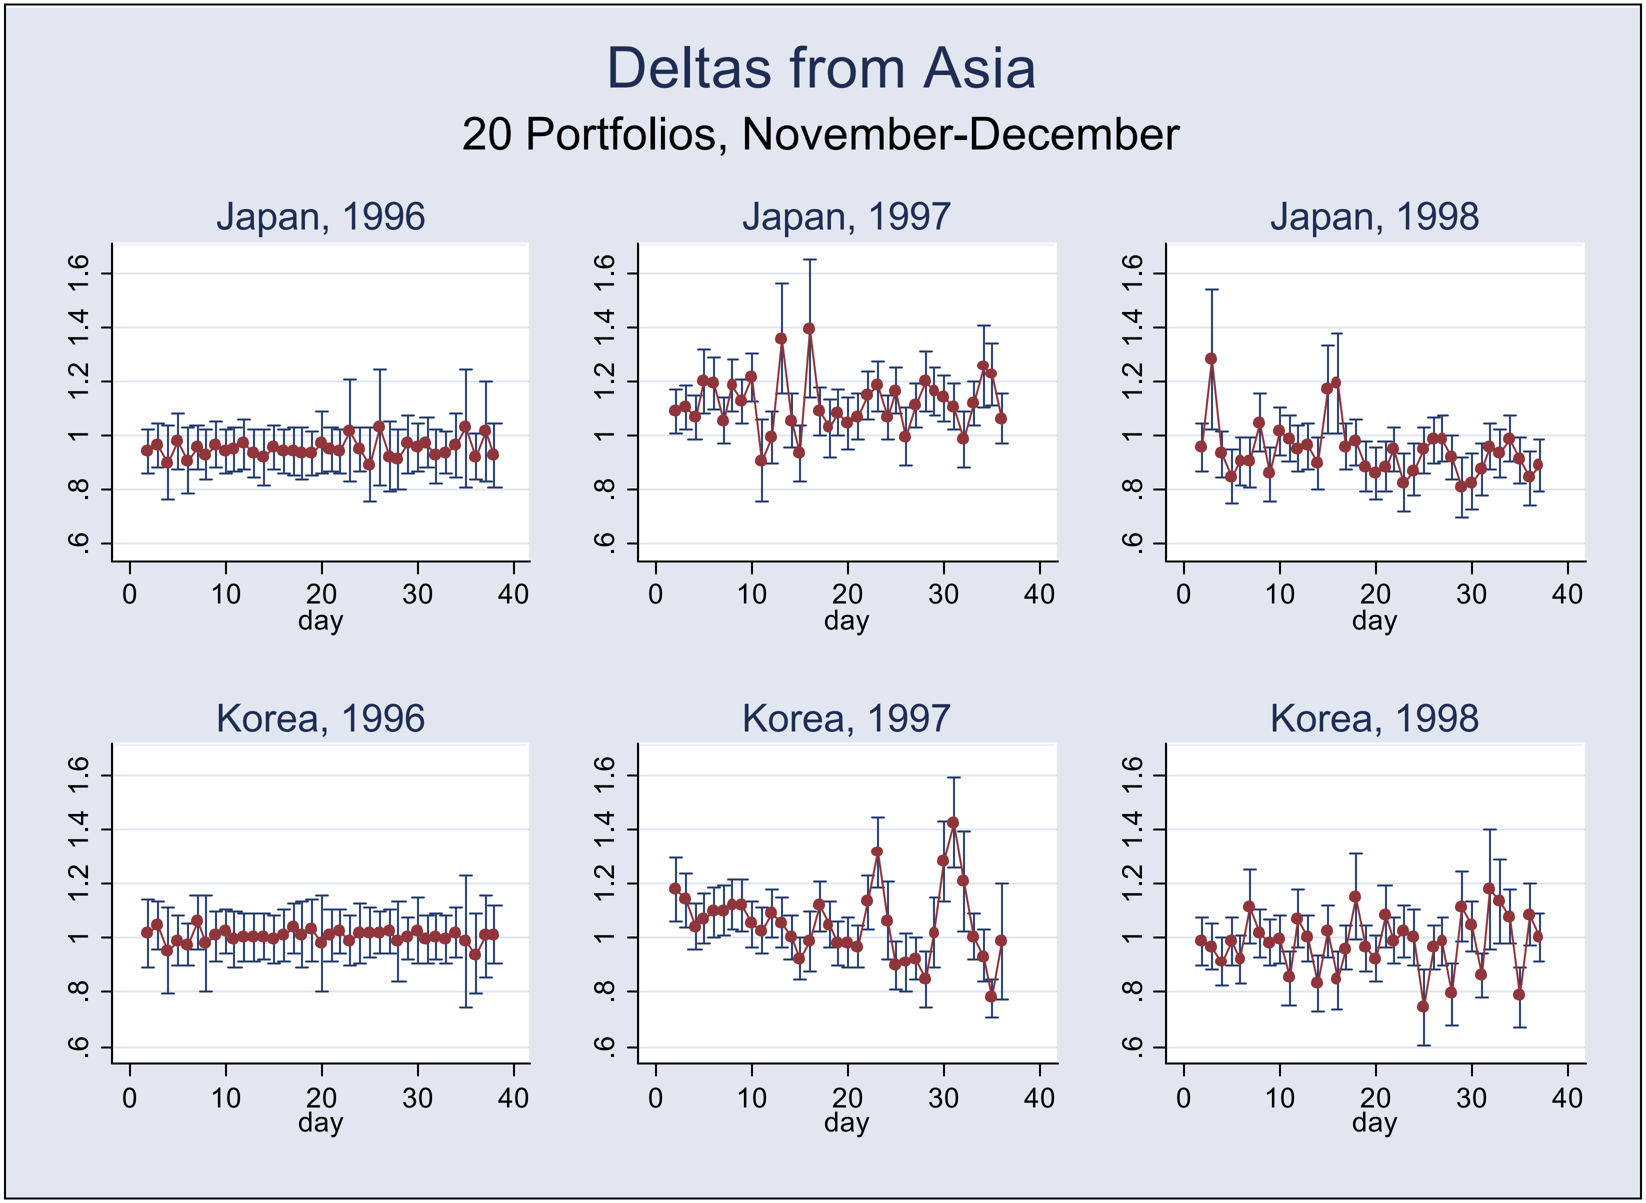

Figure 2: Estimates of Expected Discount Rate from Asian Crisis

Endnotes

[1] See e.g., the IMF's *World Economic Outlook and International Capital Markets Interim Assessment* of December 1998.

[2] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

[3] Hence the missing graphics of Figure 1.

[4] Nouriel Roubini's *Global Macroeconomic and Financial Policy Site* is the standard background reference; http://www.stern.nyu.edu/globalmacro.