East Asia and Europe during the 1997 Asian collapse: a clinical study of a financial crisis ☆

Rajesh Chakrabarti[a],*, Richard Roll[b]

[a] DuPree College of Management, Georgia Institute of Technology, Room 436, 755 Ferst Drive,
Atlanta GA 30332, USA

[b] The Anderson School, C408, UCLA, Los Angeles, CA 90095-1481, USA

Abstract

Asian stock markets are compared with European markets before and during the 1997 Asian crisis. The clinical issue is whether regional inter-dependence became larger around the crisis, fomenting investor fears of contagion and reducing asset values because of lower diversification potential. Statistical measures are developed to aid in this inquiry. We find that European and East Asian countries were not susceptible to volatility contagion in the pre-crisis era but that susceptibility increased significantly with the onset of the crisis. Covariances, correlations, and volatilities increased from the pre-crisis to the crisis period in both regions, but the percentage increases were much larger in Asia. Diversification potential was better in Asia than in Europe before the crisis; this was reversed during the crisis. The observed decline in diversification potency in Asia is reason enough for large declines in asset values though one cannot prove, of course, that it was the cause rather than the effect of the crisis. Exchange rate volatility played a major role. © 2002 Elsevier Science B.V. All rights reserved.

1. Introduction

On July 2, 1997 the Thai baht broke its peg to the US dollar. In the six months that followed, Thailand lost 65% of its stock market value in dollar terms; Hong Kong 33%, Indonesia 71%, Malaysia 57%, Philippines 58%, Singapore 24%, and South Korea 72%. Both currencies and local stock market indices plunged in most countries. Thailand lost about 33% of its local stock value, Hong Kong 32%, Indonesia 34%, Malaysia 42%, Philippines 32%, Singapore 10%, and South Korea 44%. The nearly simultaneous misfortune of so many countries has raised questions about linkages among stock markets in general. How closely are stock markets in neighboring countries related? Are East Asian stock markets more closely connected than other neighboring countries? Was the rapid spread of the crisis in Asia predictable from strong prior inter-relations or, to the contrary, did the crisis mark a break from the past? How have inter-linkages evolved over time? Finally, how was the international portfolio investor affected by changes, if any, in the regional structure of return co-movement?

To answer these questions, we require a method for characterizing co-movements of national stock returns within a region or a group of countries. Relations between pairs of countries have been studied often, but past literature provides scant guidance about how to measure the overall structure of relations within a group of countries. This paper develops a method for studying co-movements among a group of national stock markets. In doing so, it constructs some new statistical measures. The paper also tries to answer the questions raised in the previous paragraph by analyzing the regional interdependence among stock markets of East Asian countries during the mid-90s and by contrasting it to that among West European countries.

Prior to the crisis, there had been some empirical literature on the long-term linkages among Asian stock markets (e.g., Chung and Liu, 1994) but short-term co-movement and its stability had not been subjected to extensive inquiry. The crisis justifiably engendered theoretical and empirical research attempting to explain the phenomenon; yet little has been done to compare co-movements of Asian markets before and after this landmark event. This subject, which seems of critical importance in understanding the crisis and averting future ones, is where we aim to make a contribution.

The paper is organized as follows. The next section summarizes the important research findings in the areas of stock market inter-linkages and East Asian financial markets. Section 3 describes the data to be examined. In Section 4, we state the basic research questions and design some statistical tools useful in answering them. Section 5 presents the empirical evidence. Section 6 concludes and offers suggestions for future research.

2. Previous literature

Using the then newly developed method of co-integration, Kasa (1992) studied the long-term equilibrium relations among five developed capital markets and extracted a common trend that strikingly captured their movements. Subsequent contributions included Arshanapalli and Doukas (1993) and Chung and Liu (1994). Richards (1995), however, questioned Kasa's conclusions and argued that while national stock markets exhibit predictability, there is little evidence of co-integration.

Another branch of the literature has examined the evolution of correlation among a group of countries over time. Important contributions here include Kaplanis (1988), Koch and Koch (1991) and Longin and Solnik (1995). Kaplanis (1988) traced the stability of the co-movements among monthly stock index returns for ten industrial countries between 1967 and 1982; she found stability in correlations but not in covariances. Koch and Koch (1991) studied co-movements of daily returns in six industrialized and two developing countries between 1972 and 1987; they found strong interactions among markets on the same day and they also documented changes over time in the correlation structure. Among developed countries between 1960 and 1990, Longin and Solnik (1995) concluded that both covariances and correlations were unstable.

In related work, Dumas et al. (2000) study whether observed stock market correlations are consistent with cross-country correlations in national output. They find that equity correlations are indeed better in accord with global integration than with segregation. Since the structural correlations among national outputs could vary over time, their work implies that equity correlations also change over time. This is directly tested by Bansal and Lundblad (2001), who find time variations in cash flow growth rates and risk premiums that are consistent with movements in conditional correlations among country equity returns.

Hamao et al. (1990), King and Wadhwani (1990) and others have investigated "spillovers" in volatility and expected returns from one country to another. The former used a GARCH model to measure spillover while the latter used contemporaneous correlations. Volatility spillover is an important concept because it attempts to identify directional causality between disturbances in two countries. An extension and variation of this approach will be adopted hereafter.

Some studies have looked at East Asian countries in particular. Chung and Liu (1994) found that the US and five East Asian countries have co-integrated stock prices. So et al. (1997) examined the volatility of and correlations among seven Southeast Asian stock markets between 1980 and 1991 and attempted to group countries into interconnected blocks.

The Asian crisis itself has elicited quite a few papers. The role of currency has received much attention. (see Kamin, 1999; Fernald et al., 1999; Masson, 1999). Banks and regulators have received their share of criticism, (see Kane, 2000) but Kho and Stulz (2000) find little evidence that banks as a group or the International Monetary Fund were prime culprits. Few papers have viewed the crisis from the perspective of an international portfolio investor, yet many, notably including the Malaysian prime minister (Mohamad, 1997), have blamed international fund managers for the crisis. Brown et al. (2000) counter this view in the case of hedge funds.

There have also been several recent papers—both theoretical and empirical—focused on financial contagion. Baig and Goldfajn (1998) posit that a rise in correlation levels as compared to "tranquil" periods is an indication of contagion and find such an effect during the Asian Crisis. Boyer et al. (1997) and Forbes and Rigobon (1999) argue that such a finding may be the artifact of the statistical method. Bae et al. (2000), adopting a technique from health sciences, use a multinomial logistic regression to compare the levels of contagion of extreme return shocks across countries in two regions—Latin America and Asia.

In a recent assessment of the literature, Chowdhry and Goyal (2000, p. 135) conclude that most theoretical explanations of the Asian crisis are "disappointing" in their "out-of-sample performance". They call for an alternative diagnosis. Our clinical study responds by comparing the afflicted patient, East Asia, with a healthy Europe. We compare the relative strength of market linkages within each region and attempt to determine whether Asia's linkages underwent a fundamental change during the crisis. While we look at relationship between returns in two regions before and during the Asian crisis like Bae et al. (2000), our focus is on studying the changes in the degree of co-movement in Asian returns during the crisis, rather than the contagion in extreme movements, the central focus of Bae et al. (2000).

3. Data

The available data are daily close-to-close returns for indices of eight East Asian countries that experienced a crisis of some severity in 1997. Data were also collected for a matching sample of eight West European non-crisis countries. The Asian countries are Hong Kong, Indonesia, Malaysia, the Philippines, Singapore, South Korea, Taiwan and Thailand while the European countries are France, Germany, Italy, the Netherlands, Portugal, Spain, Switzerland, and the United Kingdom. The two groups are of equal size for ease of comparison and the European matching sample consists simply of the largest countries augmented with two arbitrarily chosen smaller countries for a total of eight.
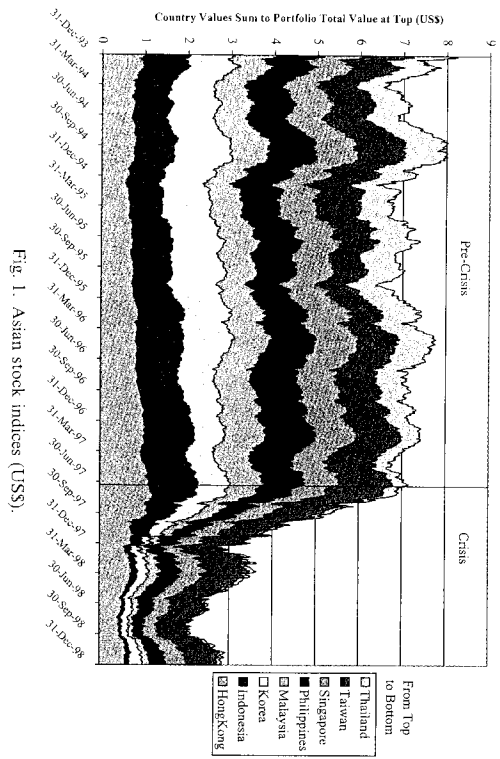
Time zone differences within either region do not exceed two hours; hence, contemporaneous daily data can be employed in computing co-movements. Daily closing stock index levels in local currency are from DATASTREAM. The local currency stock index value is converted to dollars using the daily closing exchange rate obtained from the same source. Dollar returns are computed by using the formula $\ln(P_t/P_{t-1})$ where P_t refers to the dollar value of the index at time t. The data cover five years from 12/31/93 to 12/31/98, 1305 daily returns. We rather arbitrarily selected July 2, 1997 to be the beginning of the Asian crisis. On this date occurred the first major devaluation, that of the Thai baht. There had been previous storm warnings, particularly a "speculative attack" on the baht in mid-May, but July 2 has been generally identified as the first time that a real crisis was apparent; (see, for instance, the chronology of the Asian crisis constructed by Nouriel Roubini[1]). Our data sample thus includes 392 observations during the crisis period and 913 points before. Fig. 1 provides an overall depiction of the pre-crisis and crisis periods by plotting an initially equal-weighted portfolio of the eight Asian countries.

Table 1 provides descriptive statistics of the dollar-valued returns for both groups of countries over the entire sample and for the pre-crisis and crisis

Fig. 1. Asian stock indices (US$).

Table 1
Descriptive statistics, country returns by sub-period

Daily country index returns (%/day) are denominated in US dollars and tabulated for the pre-crisis period, December 31, 1993 through July 1, 1997, (913 observations), and the Asian crisis period, July 2, 1997 through December 31, 1998 (392 observations).

	Hong Kong	Indonesia	Korea	Malaysia	Philippines	Singapore	Taiwan	Thailand
				Pre-crisis				
Mean	0.0240%	0.0189%	−0.0291%	−0.0149%	−0.0006%	−0.0152%	0.0375%	−0.1133%
Standard deviation	1.3003%	1.0190%	1.2513%	1.3010%	1.1475%	0.8777%	1.5531%	1.6187%
Kurtosis	3.882	3.919	1.282	8.243	2.334	4.951	3.471	2.659
Skewness	-0.514	−0.144	−0.022	0.103	−0.242	−0.233	−0.168	−0.085
Minimum	−7.72%	−5.28%	−6.03%	−7.48%	−6.20%	−5.24%	−7.39%	−7.80%
Maximum	5.63%	6.47%	6.04%	9.97%	4.80%	5.52%	6.62%	6.45%
				Crisis				
Mean	-0.1319%	−0.4242%	−0.1118%	−0.2756%	−0.1981%	−0.1095%	−0.1244%	−0.1994%
Standard deviation	2.7300%	6.4720%	4.4184%	4.0963%	2.8091%	2.2490%	1.9586%	3.5978%
Kurtosis	5.908	7.942	3.938	6.655	1.650	3.234	2.078	2.379
Skewness	0.437	−0.947	0.044	0.611	0.117	0.348	−0.131	0.764
Minimum	−13.63%	−42.84%	−19.18%	−22.04%	−10.24%	−9.59%	−8.98%	−9.38%
Maximum	15.37%	21.48%	20.12%	22.20%	11.56%	9.94%	6.72%	17.44%

	France	Germany	Italy	Netherlands	Portugal	Spain	Switzerland	UK
				Pre-crisis				
Mean	0.0252%	0.0432%	0.0357%	0.0730%	0.0586%	0.0647%	0.0643%	0.0421%
Standard deviation	0.8509%	0.8357%	1.3396%	0.7257%	0.7833%	0.9208%	0.8361%	0.6511%
Kurtosis	1.150	1.046	1.431	1.761	4.213	1.716	1.488	0.756
Skewness	−0.135	-0.341	0.028	−0.343	−0.175	-0.399	−0.080	−0.256
Minimum	−3.71%	−3.71%	−5.68%	-4.02%	−5.20%	−5.36%	−3.31%	−2.97%
Maximum	3.09%	2.64%	6.02%	2.43%	3.75%	3.12%	3.55%	2.26%
				Crisis				
Mean	0.0936%	0.0636%	0.1544%	0.0639%	0.0903%	0.0999%	0.0767%	0.0532%
Standard deviation	1.2811%	1.4600%	1.7452%	1.4338%	1.4828%	1.5641%	1.3189%	1.0438%
Kurtosis	1.923	1.239	1.176	1.090	2.912	2.396	1.810	1.028
Skewness	−0.030	−0.371	-0.211	0.061	−0.417	−0.255	0.082	−0.034
Minimum	−4.74%	-5.55%	6.40%	-4.89%	−6.72%	−6.02%	−4.09%	−3.14%
Maximum	5.61%	4.94%	6.69%	5.31%	5.84%	7.16%	6.66%	3.84%

sub-samples. While the Asian countries all experienced negative average returns during the crisis period, no European country shared that experience; indeed, returns were higher than in the pre-crisis period for every European country. Fig. 2a, plots the mean returns to emphasize just how different the crisis period was in the two regions. The volatility of returns went up in the crisis period for every single country in both regions. In Asia, volatility was roughly three times higher (on average) during the crisis period; in Europe, it was about 1.5 times higher (Fig. 2b).

4. Research questions and statistical tools

Our research agenda can be captured in a few questions: How susceptible are countries to volatility spillover and how did this susceptibility change during the Asian crisis? What proportion of intertemporal changes in covariances among countries is attributable to changes in correlation and what proportion to changes in variance? How has the ability to diversify within a region evolved over time? Each of these questions is critically important to the international investor. Volatility spillover, or contagion, implies that countries will become more or less risky around the same time. Evolution of the covariance structure alters the available degree of diversification and thus changes the risk/return tradeoff for portfolio investors. The mere possibility of either phenomenon has the potential to induce investor panic and could conceivably explain why dramatic declines in prices propagate throughout a region.

First, we look at volatility "spillover" from one market to another and seek to determine whether spillover intensity changed during the 1997 crisis. An increase in spillover intensity would be evidence of greater contagion around the crisis. This issue was investigated by King and Wadhwani (1990), among others, during the pandemic October 1987 crash. We change their now standard approach slightly so as to measure the overall susceptibility of countries within a region to volatility shocks imported from neighbors.

To measure volatility spillover, we regress weekly first differences of daily volatility (from dollar-denominated returns of a country's stock index) on lagged values from other countries in the region. Specifically, the regression model is

$$\Delta\sigma_{t,i} = \alpha_i + \Sigma_j[\beta_{i,j}\Delta\sigma_{t,j-1} + \Delta\beta_{i,j}\delta_j\Delta\sigma_{t,j-1}] + \varepsilon_{t,i} \qquad (1)$$

where $\sigma_{t,i}$ is the estimated volatility computed from daily dollar-valued returns during week t in country i, $\Delta\sigma_{t,i} \equiv (\sigma_{t,i} - \sigma_{t,i-1})$ is the weekly first difference, and δ_i is an indicator variable which is zero in the pre-crisis period and unity during the crisis period. The average values of the weekly volatilities of the
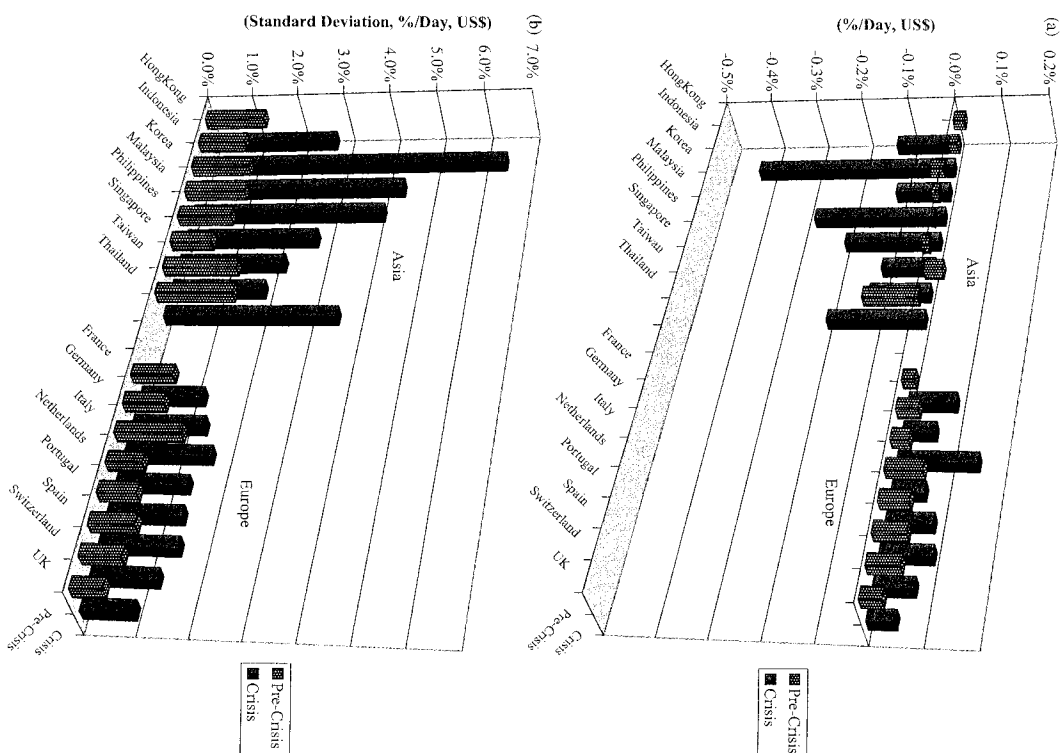
(a)

(%/Day, US$)

(b)

(Standard Deviation, %/Day, US$)

Fig. 2. (a) Mean returns by region and period. (b) Volatility by region and period during the crisis period.

Table 2
Weekly volatility of daily returns

The table reports time series mean values of $\sigma_{t,i}$, the estimated standard deviation of daily dollar-denominated returns during week t in country i.

	Pre-crisis (181 weeks)	Crisis (77 weeks)
Asia		
Hong Kong	0.97%	2.04%
Indonesia	0.73%	4.45%
Korea	1.00%	3.08%
Malaysia	0.93%	2.83%
Philippines	0.86%	2.10%
Singapore	0.67%	1.64%
Taiwan	1.19%	1.51%
Thailand	1.22%	2.65%
Europe		
France	0.69%	1.03%
Germany	0.68%	1.21%
Italy	1.06%	1.40%
Netherlands	0.59%	1.16%
Portugal	0.60%	1.11%
Spain	0.74%	1.29%
Switzerland	0.66%	1.03%
UK	0.53%	0.84%

daily returns, the $\sigma_{t,i}$'s, for the different countries in the two regions in the pre-crisis and the crisis periods, are presented in Table 2.

First, we test whether there is evidence of volatility spillover of the same magnitude spanning the pre-crisis and crisis periods. The test's null hypothesis is that the $\beta_{t,i}$ (as opposed to the slope changes, $\Delta\beta_{t,i}$), are all zero. For each region, we estimate all eight regressions (1) as a system and construct a joint test; i.e.,

$$H_0 : \beta_{i,j} = 0 \quad (j \neq i, \quad \forall i, j).\qquad(2)$$

This involves a simultaneous test that 56 coefficients in eight regional equations, are significantly different from zero.[2] Note that this test involves only the *inter*-country coefficients and therefore assesses spillover. Volatility dependence across time *within* a country does not involve spillover and should be assessed separately (see below).

---

[2] This test explicitly allows spillover to be either positive or negative. Although one might intuitively think of volatility increases spilling over as increases elsewhere, or vice versa, the opposite is conceivable.

The "slope dummy" (δ) in (1) allows us to test whether the 1997 Asian crisis was marked by a change in volatility spillover. The null hypothesis is

$$H_0 : \Delta\beta_{i,j} = 0 \quad (j \neq i, \quad \forall i, j).\qquad(3)$$

The test does not restrict the change in spillover susceptibility, if any, to be of a particular sign, though one might intuitively anticipate an increase in susceptibility during a crisis.

To be complete, we also record the evidence on volatility serial dependence within countries during both periods ($H_0 : \beta_{t,i} = 0 \ \forall i$) and test whether there has been a change in such dependence from the pre-crisis to the crisis period ($H_0 : \Delta\beta_{t,i} = 0 \ \forall i$).

Our second research question asks about return covariances among countries within a region and how those covariances changed near the crisis. The importance of the covariance to the international portfolio investor can hardly be over-emphasized; mean-variance efficient portfolios are determined in large part by the covariance matrix. Consequently, significant change in the regional covariance structure can explain why investors might reduce their holdings; the benefits of diversification across those markets have become less compelling.

We also hope to uncover the sources of changes in covariance. Independent but contemporaneous disturbances in different countries will affect covariance but not correlation. Sympathetic movements among stock markets will raise the correlation as well. Understanding the covariance structure requires decomposition into correlation and variance; thus decomposition is proposed and implemented in this paper. The relative importance of one component over the other provides a better understanding of the nature of co-movements among the markets. We estimate whether correlation played the dominant role during the crisis and also whether the correlation/variance decomposition is significantly different in Asia and Europe.

There are no standard methods for testing the structural evolution of correlation or covariance matrices. Apart from long-horizon co-integration, the typical approaches have compared correlation matrices at two points in time (e.g. Longin and Solnik, 1995) or the evolution of pair-wise correlations. The first approach does measure a group effect but only at two points in time; it is not easily adapted for tracing the evolution in between. The latter approach can capture the evolution through time but does not measure any group effect. This situation would be improved if the information in a covariance or correlation matrix could be summarized in a single scalar. The evolution of correlation matrix could be studied with univariate techniques. Moreover, such a scalar could be more easily decomposed into its constituents, correlations and variances, which might provide insights into the fundamental determinants of structural change. We propose a battery of such scalar measures and apply them in studying the Asian crisis.

To summarize information in the covariance matrix in a single number, we propose the geometric mean of absolute values of the covariance between every pair of countries in a region:

$$\gamma_{cov} = \left[\prod_{i<j} |cov_{ij}| \right]^{1/[m(m-1)/2]}$$

where m is the number of countries. A corresponding measure for correlation is

$$\gamma_{corr} = \left[\prod_{i<j} |corr_{ij}| \right]^{1/[m(m-1)/2]}.$$

Finally, we compute the geometric mean of the standard deviations of the individual returns

$$\left(\gamma_a = \left[\prod_{i=1}^{m} \sigma_i \right]^{1/m} \right).$$

These measures not only capture the information in the covariance and correlation matrices but are also connected by the simple identity:[3] $\gamma_{cov} = \gamma_{corr}\,\gamma_a^2$. Thus, the log first difference between two periods is

$$\Delta \log \gamma_{cov} = \Delta \log \gamma_{corr} + \Delta \log \gamma_a^2 \tag{4}$$

which provides a decomposition of covariance change into its volatility and correlation components. Note that log first differences are equivalent to percentage changes (continuously compounded) and are thus dimensionless. Assuming that correlation and volatility change in the same direction, the relative importance of correlation can be measured by

$$\eta = |\Delta \log \gamma_{corr}| / (|\Delta \log \gamma_{corr}| + |\Delta \log \gamma_a^2|) \tag{5}$$

which is bound between 0 and 1. The higher (lower) the value of η, the greater is the contribution of correlation (volatility) to the change in covariance. Evolution of the η statistic should reveal points in time when correlation and variance were relatively more prominent causes of movements in covariance.

Geometric means γ_{cov} and γ_{corr} are intended to compactly summarize covariances and correlations of stock returns in the two regions. This does involve some loss of information as any single number necessarily would if it attempts to capture an array of numbers. The justification is our focus on the general trend of inter-linkages within each regional group rather than on individual countries per se.

The evolution and stability of the correlation matrix of returns within a region is of fundamental importance given the crucial role of the correlation

---

[3] This can easily be verified by expanding the expression for γ_{corr} and reorganizing the terms.

structure in portfolio formation. To measure the evolution of correlation matrices, we adopt a test statistic for the equality of two or more correlation matrices developed in Jennrich (1970). For any two p-variate sample correlation matrices R_1 and R_2, of sample size n_1 and n_2 respectively, the Jennrich χ^2 is given by

$$\chi^2 = \tfrac{1}{2} tr(Z^2) - dg'(Z)S^{-1}dg(Z). \tag{6}$$

where $Z = c^{1/2}\bar{R}^{-1}(R_1 - R_2)$, $\bar{R} = (n_1 R_1 + n_2 R_2)/(n_1 + n_2)$, $c = n_1 n_2/(n_1 + n_2)$ and $S = (\delta_{ij} + \bar{r}_{ij}\bar{r}^{ij})$. Here δ_{ij} is the Kronecker delta, $\bar{r}_{ij}$ is an element of $\bar{R}$ and $\bar{r}^{ij}$ is an element of $\bar{R}^{-1}$. The expressions "tr" and "dg" refer to the trace and diagonal of a matrix, respectively. If the true correlation matrices are equal, χ^2 is distributed asymptotically as χ^2 with $p(p-1)/2$ degrees of freedom. Consequently, by computing correlation matrices from two different time periods, we can test whether they have changed; e.g., from the pre-crisis to the crisis period.

There is an alternative and perhaps more direct approach to measuring changes in investment opportunities around the crisis: calculate diversification potential directly. How well (in mean-variance terms) would an international investor have fared by creating an efficient allocation within a region? Would diversification be more potent in Asia or in Europe, before the crisis or thereafter? These questions are answered by the shape of the regional efficient frontier and in particular by its curvature. The more sharply curved the frontier, the less effective diversification. Hence estimated curvature provides a direct empirical measure of diversification potential, a measure which can be traced over time and compared across periods and regions.

The mathematics of the efficient frontier, (e.g., Roll, 1977) show that the frontier can be described by the equation

$$\sigma_p^2 = (a - 2br_p + cr_p^2)/(ac - b^2), \tag{7}$$

where σ_p^2 is the variance of efficient portfolio p and r_p is its mean return. The efficient set constants are given by $a = E'V^{-1}E$, $b = E'V^{-1}\iota$, and $c = \iota'V^{-1}\iota$, where E is the mean return vector of n assets, V is their variance-covariance matrix, and ι is the unit vector. With a little rearrangement, (7) can be written as

$$(\sigma_p^2 - 1/c) = [c/(ac - b^2)](r_p - b/c)^2 \tag{8}$$

which is the equation for a parabola with curvature $\hat{\partial}^2(\sigma_p^2)/\hat{\partial}(r_p) \equiv C = 2c/(ac - b^2)$. The average level of C and its evolution in the two geographic regions provide the direct measures we seek of diversification potential.

Taken as a battery of statistics, the above empirical constructs should help bring out a comprehensive picture of the short-term relations among stock markets in East Asia and how they differed from those among European stock markets. We hope also that these constructs will refine methods for studying and characterizing regional interdependence of markets in general.

5. Empirical results

5.1. A graphical portrayal

To depict graphically the evolution of relations among stock markets, we make calculations within overlapping windows of pre-specified length. Each window excludes the first daily observation from the previous window and adds one more daily observation in sequence. Hence successive values of the statistics are analogous to moving averages.

For regional (geometric) averages of covariance, correlation, and volatility, (γ_{cov}, γ_{corr} and γ_σ^2), a six-month window gives the values plotted in Figs. 3–5. The spike in γ_{cov} for the East Asian countries in Fig. 3 coincides with the Asian crisis and appears to represent a significant break from the past. Before the crisis, Asia had a slightly lower covariance than Europe on average, but this was reversed dramatically during the crisis. Around the middle of the crisis period, γ_{cov} fell substantially in Asia, though it remained at a level much higher than the pre-crisis period and considerably higher than in Europe.

The European countries too display an increase in covariance during the crisis period. The largest part of the rise in Europe came later and was less pronounced than in Asia. Unlike Asia, no mid-crisis period downturn is visible in Europe; there is a sharp second upward movement instead.

Did correlations within the two regions move as much or can some part of the spikes in covariance be explained by higher volatility? Figs. 4 and 5 present
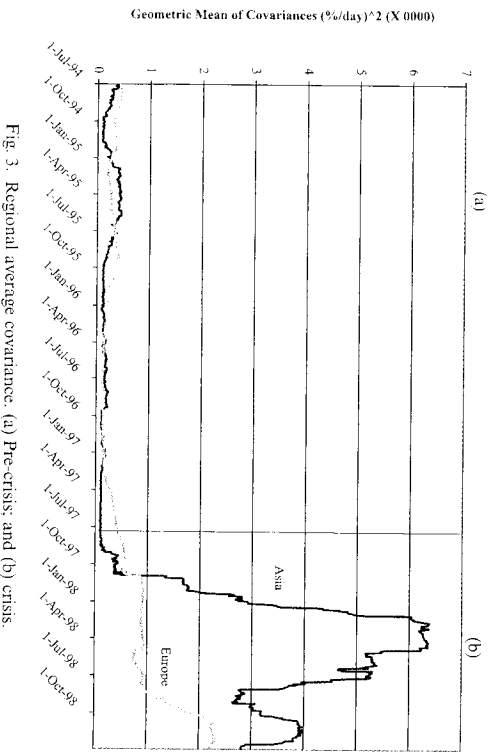

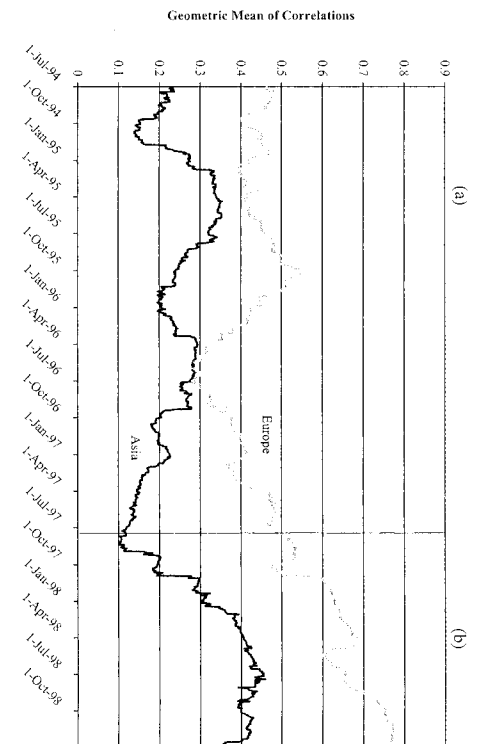Fig. 3. Regional average covariance. (a) Pre-crisis; and (b) crisis.

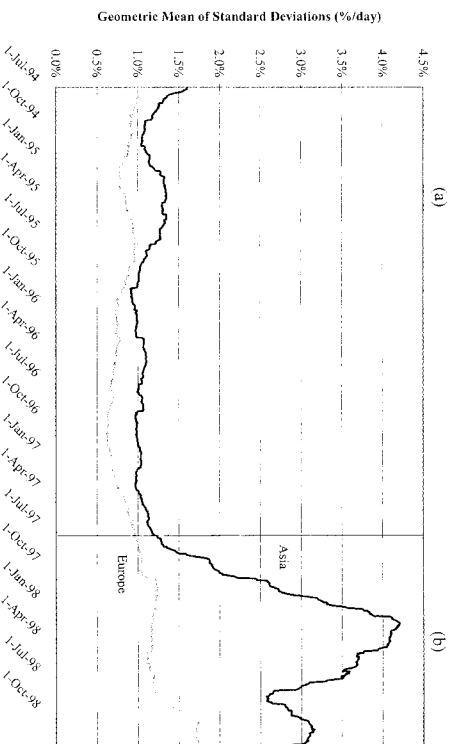Fig. 4. Regional average correlation. (a) Pre-crisis; and (b) crisis.

Fig. 5. Regional average volatility. (a) Pre-crisis; and (b) crisis.

visual evidence about this question. Fig. 4 shows an upward trend in correlation among European countries which appears to begin up to a year before the crisis and continues during the crisis period. This rise in correlations in Europe perhaps had more to do with the movement towards a monetary union than the Asian crisis. During the crisis, European and Asian countries displayed concurrent upward trends in correlation. Asia's trend was sharper near the beginning of the crisis but it leveled off and began to turn down in late 1998 while Europe's continued upward after a temporary dip in late April of 1998.[4]

Fig. 5 plots average volatility. The similarity of the pattern to covariance (Fig. 3) is obvious even visually. Prior to the crisis, volatility was slightly higher in Asia than in Europe. The gap widened substantially as the crisis deepened and then narrowed toward the end of 1998.

The Jennrich χ^2 statistic, measuring the stability of the correlation matrix, is computed in a different fashion. Here one-year overlapping windows are each divided into two six-month sub-periods; the statistic compares correlation matrices between the sub-periods. Hence, each one-year window yields one value of the Jennrich χ^2. Successive windows drop the first day from the previous window and add the day after the previous last day. Fig. 6 plots the p-value from the Jennrich test of the hypothesis that the two adjacent six-month periods have the same correlation matrix, so values near zero imply a strong likelihood that the matrix has changed.

For both Europe and Asia during the pre-crisis period, there are many large p-values: i.e., little evidence of an evolving correlation matrix. There are some near zero but then this statistic like any other has unavoidable sampling variation. The striking feature of Fig. 6 is that Asia's p-value falls essentially to zero during the crisis and remains there throughout.[5] The Jennrich χ^2 p-value is also relatively low for Europe during the crisis period, but not persistently as low as for Asia.

Fig. 7 depicts the evolution of the sample curvatures of efficient frontiers within the Europe and Asia regions before and during the 1997 crisis. Each point gives the curvature of a sample efficient frontier constructed from daily return data for six months ending on the plotted date. Higher curvature implies less effective diversification within the region.

For about 18 months prior to the crisis and for six months into the crisis, Asia's curvature was lower than Europe's. This no doubt reflects similarities

---

[4] Because of the rolling window, there is considerable dependence across time in these plotted numbers. For instance, the dip in European correlation in the last few days of April 1998 was probably induced by dropping a few observations from the window. These same observations induced the sharp rise in the computed average correlation around late October 1997, six months earlier.

[5] The plot stops in mid-1998 because the Jennrich statistic requires a six-month sample before and after the plotted date
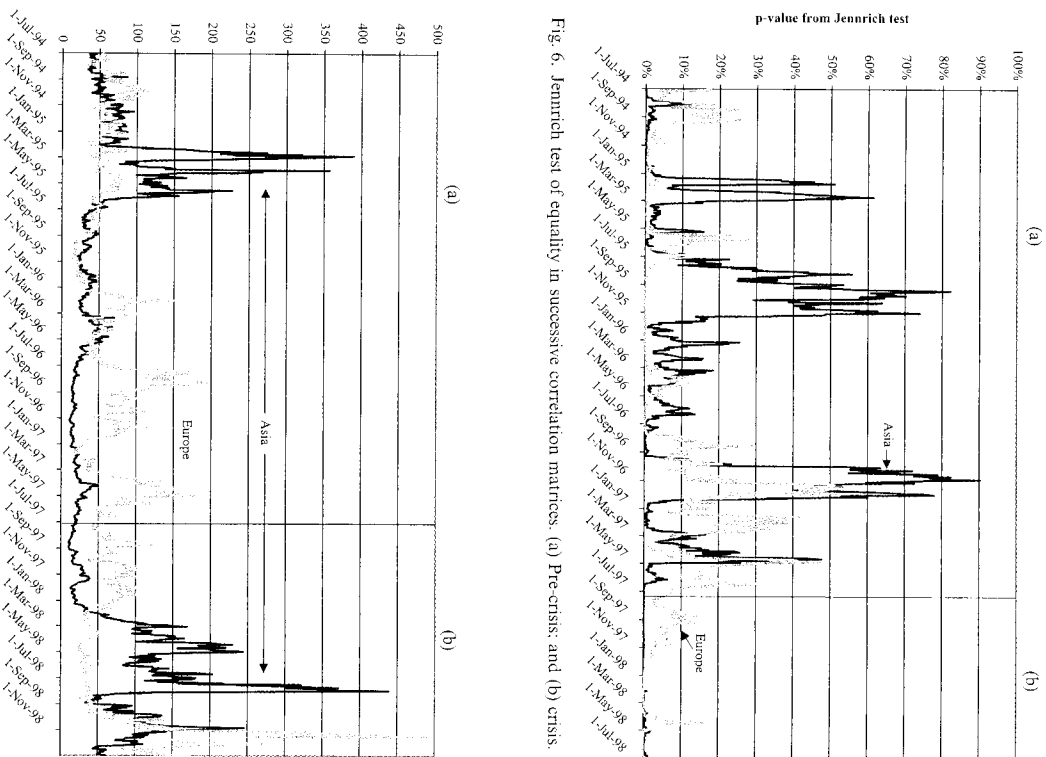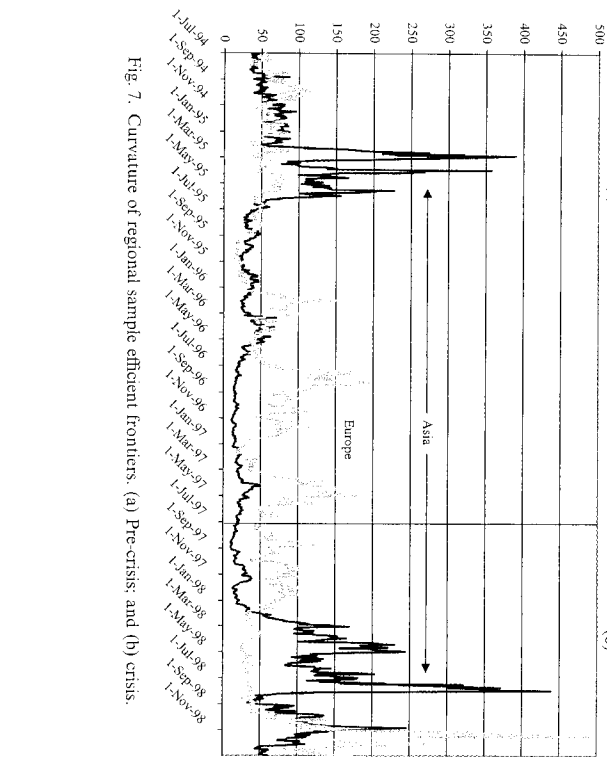
Fig. 6. Jennrich test of equality in successive correlation matrices. (a) Pre-crisis; and (b) crisis.

Fig. 7. Curvature of regional sample efficient frontiers. (a) Pre-crisis; and (b) crisis.

among European economies and the concomitant smaller potential for diversification. By early 1998, however, Asia's curvature had increased dramatically and exceeded Europe's by a sizeable margin. By this time at least six months of the crisis period fell within the window's calculation. Perhaps surprising, late 1998 saw a sharp reduction in Asia's curvature followed by an even sharper increase in curvature for Europe. This is consistent, of course, with the relative movements of the average covariances in the two regions (Fig. 3) around the same time.

5.2. Statistical tests

The pictures discussed above seem to portray remarkable changes within Asia during the 1997 crisis and less dramatic but still noticeable changes in Europe. But in dealing with statistics, a picture's value may not rise to its purported worth of a thousand words. Indeed, intra-ocular tests based on pictures are often misleading. To be assured of reliable inferences, one is obliged to adopt more formal methods.

5.2.1. Volatility spillover

Table 3 reports tests for the presence of volatility spillover effects before and during the crisis in both regions. The first hypothesis, (1), that volatility spillover was present in pre-crisis and crisis periods and had the same intensity in both, is not supported by the evidence in either region. Neither p-value in the first row of Table 3 is even close to significant.

In contrast, there is strong evidence of spillover in Asia during the crisis, a p-value of 0.000 rejecting hypothesis (3). There is also evidence, (p-value 0.00221) that volatility spillover increased in Europe.

Given previous literature, it is not surprising that volatility changes are highly serially dependent within countries after accounting for spillover; this is documented by the p-values in the third row of Table 3; they are zero to three decimal places in both regions. Interestingly, the "own" coefficients $\beta_{i,i}$ from model (1) (not reported) are mostly negative. There is no evidence of a change within Asia of this intra-country dependence from pre-crisis to crisis, (the p-value is 0.222 for a test that $\Delta\beta_{i,i} = 0$ jointly for all i). There is some evidence, however, that Europe experienced a change (p-value 0.0310). It would seem that Europe became more generally serially dependent after the beginning of the Asian crisis both within each country and across countries. Asia differed from Europe because there was no significant increase in volatility serial dependence intra-country.

Whatever the underlying causes of Asia's crisis might have been, they apparently induced a structural change in the propensity of Asian countries to import recent volatility disturbances from their neighbors. Remember that (1) measures very short-term contagion, volatility that appeared in neighboring

Table 3
Tests for volatility spillover and for serial dependence

For each of two regions, Asia and Europe, a system of eight simultaneous equations is estimated using US$-denominated country index daily returns. There is one equation per country in the form $\Delta\sigma_{i,t} = \alpha_i + \Sigma_j\beta_{i,j}\Delta\sigma_{j,t-1} + \Delta\beta_{i,j}\delta_t\Delta\sigma_{j,t-1} + \varepsilon_{i,t}$, where $\sigma_{i,t}$ is the estimated volatility computed from daily returns during week t in country i, $\Delta\sigma_{i,t} = (\sigma_{i,t} - \sigma_{i,t-1})$ is the weekly first difference, δ_t is an indicator variable which is zero in the pre-crisis period and unity during the crisis period, and $\varepsilon_{i,t}$ is a disturbance not necessarily independent across equations (countries). Joint tests that the $\Delta\beta_{i,j}$'s for $j \neq i$ are all zero would be rejected if there is evidence of volatility spillover in both the pre-crisis and crisis periods of the same intensity. Joint tests that the $\beta_{i,j}$'s for $j \neq i$ are all zero would be rejected if volatility spillover changed (and was statistically more significant) during the crisis period. Joint tests that the $\beta_{i,j}$'s are all zero would be rejected if there was a change in within-country volatility dependence from the pre-crisis to the crisis period. The sixteen countries are listed in Table 1. There are 258 weekly observations. To save space, the estimated coefficients and other statistics are not reported but will be provided to interested readers upon request.

Null hypothesis	Asia		Europe	
		p-value		
Spillover across countries				
$\beta_{i,j} = 0, (j \neq i, \forall_{i,j})$	0.999		0.548	
$\Delta\beta_{i,j} = 0, (j \neq i, \forall_{i,j})$	0.000		0.00221	
Serial dependence within countries				
$\beta_{i,i} = 0, \forall_i$	0.000		0.000	
$\Delta\beta_{i,i} = 0, \forall_i$	0.222		0.0310	

countries just a few days earlier (i.e., in the previous week). In the pre-crisis period, there was no evidence at all of such contagion. It would be interesting to know whether such a marked increase in contagion is a general feature of crises.

5.2.2. Variability, co-movement, and efficient frontier curvature

Regional geometric averages of covariances, correlations, volatilities, and efficient frontier curvature are non-linear and would have complicated sampling distributions even if returns were Gaussian. This is all the more problematic because kurtosis statistics (Table 1) show the departure from Gaussian typical of asset returns. Hence, tests of statistical significance would be very difficult to develop analytically. It seems doubtful that GMM or similar methods can be exploited to obtain standard errors for our non-linear constructs, particularly in light of the large cross-section and relatively limited time period of the sample. Hence, the simplest approach appears to be re-sampling from the original data; i.e., employing the "bootstrap".[6]

---

[6] A basic reference on the bootstrap is Efron (1982).

Our resampling plan adheres faithfully in every instance to the procedures actually used in computing the sample point estimate from the data. To give a concrete example, suppose we wish to test whether covariances within a region, say Asia, increased from the pre-crisis to the crisis period. The point estimate would simply be the difference between the geometric averages for the Asian region in the two periods; i.e., $\gamma_{cov,Asia,Crisis} - \gamma_{cov,Asia,Pre-Crisis}$. The first of these averages is calculated from 913 observations while the second is based on 392 observations, those being the numbers of daily returns available in the two sub-periods.

Resampling for this statistic thus involves first selecting 913 calendar dates at random (with replacement) along with the actual daily returns for the eight Asian countries on those dates. Next, a second random selection of 392 calendar dates and their corresponding returns are selected. Covariances among the eight countries are computed for the two samples separately. Their geometric averages are then differenced to produce $\gamma_{cov,Asia,Sample\#2} - \gamma_{cov,Asia,Sample\#1}$. The random selections are repeated 10,000 times, thereby generating a sampling distribution for 10,000 geometric mean covariance differences. Statistical significance is gauged by the position of the actual data statistic within the fractiles of the resampled distribution. Fig. 8 plots the bootstrapped sampling distribution[7] for this illustrated example, (pre-crisis to crisis geometric mean covariance difference for Asia). The point estimate (3.63) lies completely outside the sampling distribution and is thus statistically significant.

A similar but distinct method is used to develop a resampling distribution for testing across regions. For example, we might like to know whether Asia and Europe had statistically different covariances within the crisis period. The point estimate is $\gamma_{cov,Asia,Crisis} - \gamma_{cov,Europe,Crisis}$. In this case, a pair of countries is selected at random from the sixteen available and their covariance is computed from the 391 crisis period observations. The pairwise selection is repeated (with replacement) to obtain 56 resampled covariances, a number ordained by the fact that each region of eight countries has $8(7)/2 = 28$ distinct **true** covariances. The 56 resampled covariances are then divided randomly into two groups, the geometric mean is computed within each group and their difference, $\gamma_{cov,Group\#1,Crisis} - \gamma_{cov,Group\#2,Crisis}$ is taken. Repeating this operation 10,000 times provides the sampling distribution required.

For inter-regional comparison of efficient frontier curvature for a given period, a slight modification is required. If the same country were selected twice in a bootstrap pseudo-region, the covariance matrix would be singular by construction, thus implying infinite curvature. Consequently, eight *different* countries must be selected at random. Efficient frontier curvature is computed
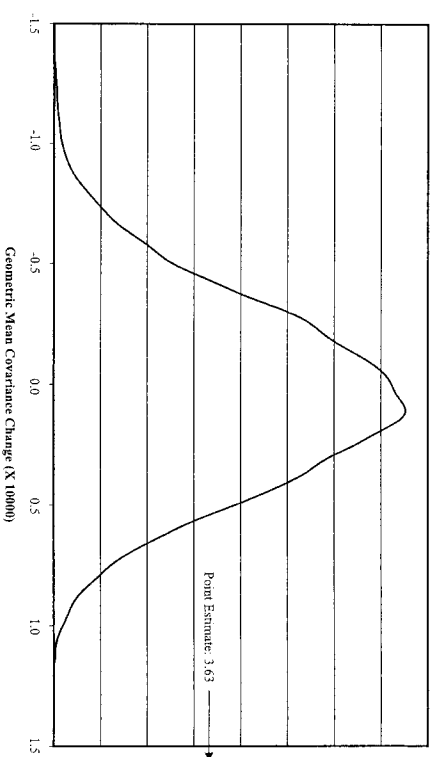
Fig. 8. Non-parametric estimated density for change in covariance from pre-crisis to crisis period, Asia.

using the covariance matrix of those eight and compared with the corresponding curvature for the eight countries not selected. The difference in the two curvatures, tabulated 10,000 times, provides the bootstrap resampled distribution.

Table 4 reports all the results. The two left columns provide evidence on the differences between the crisis and pre-crisis periods for each region considered separately. The two right columns give comparisons between the Asian and European regions by sub-period.

There are strong indications that covariances, correlations, and volatilities increased during the crisis in both regions.[8] All the p-values are 0.0000, thus firmly rejecting the hypotheses that these parameters remained constant over time. Prior to the crisis, covariances were not significantly different in Europe and Asia but they became significantly higher in Asia during the crisis. Correlations were significantly higher in Europe in both periods. Volatility was significantly higher in Asia in both periods though the magnitude of the difference and its level of significance increased during the crisis.

The η statistic, which measures the relative contribution of correlation to the change in covariance between periods, (see (5)), is 0.309 and 0.359 for Asia and

Table 4
Comparisons across periods and regions of diversification potential

Geometric averages of covariances, correlations, and volatilities are calculated from US$-denominated daily index returns for eight Asian and eight European countries before and during the 1997 Asian crisis. Efficient frontier curvatures are calculated for the countries within each region. The η statistic measures the relative importance of correlation to an observed change in covariance. The pre-crisis period was December 31, 1993 through July 1, 1997. The crisis period was July 2, 1997 through December 31, 1998. Changes in the statistics across the two periods and differences in the statistics across regions within each period are tabulated and compared against resampled (bootstrapped) probability distributions. The p-values (in parentheses) indicate the fractile of the bootstrapped distribution corresponding to the observed point estimate. The first column presents results for tests comparing the pre-crisis data and the crisis data while the second column presents those comparing Asia with Europe for the entire sample period.

		Crisis–pre-crisis	Asia–Europe
Covariance[a]			
Asia	Pre-crisis	3.63	−0.0648
		(0.0000)	(0.7869)
Europe	Crisis	1.08	2.48
		(0.0000)	(0.0055)
Correlation			
Asia	Pre-crisis	0.206	−0.251
		(0.0000)	(0.9928)
Europe	Crisis	0.298	−0.343
		(0.0000)	(0.9952)
Standard deviation (%/day)			
Asia	Pre-crisis	2.069%	0.387%
		(0.0000)	(0.0030)
Europe	Crisis	0.553%	1.903%
		(0.0000)	(0.0016)
Efficient frontier curvature			
Asia	Pre-crisis	160.00	−65.1
		(0.0309)	(0.6710)
Europe	Crisis	5.27	89.5
		(0.1744)	(0.0093)
η statistic			
Asia		0.309	
		(0.8422)	
Europe		0.359	
		(0.6053)	

[a] Covariances have been scaled upward by a multiple of 10^4.

Europe, respectively. η would be 1.0 if correlation were entirely responsible for the observed change in covariance from the pre-crisis to the crisis period while η would be zero if volatility were entirely responsible; hence, it appears that volatility was somewhat more responsible than correlation. However, the

p-values are insignificant for both regions. This suggests that both correlation and volatility played a role, neither being dominant in a statistically reliable sense.

Efficient frontier curvature increased significantly in Asia (p-value 0.0309) during the crisis, which confirms that diversification opportunities declined substantially. In Europe, the efficient frontier's curvature also increased during the crisis period but only moderately and with an insignificant p-value of 0.1744. As measured by curvature, diversification was actually better in Asia than in Europe before the crisis, though the difference was not statistically significant. Diversification's drastic fall in Asia during the crisis reversed the inter-regional difference, which then became significant (p-value 0.0093).[9]

The impressive decline in diversification potential in Asia at the time of the crisis is consistent with rational investors driving down asset values. Of course there is no way to ascertain whether a reduction in diversification actually caused the crisis or even whether it was an effect rather than a cause.[10] Nonetheless, it appears to be an interesting smoking gun.

5.3. Decomposing the crisis: currencies vs. local returns [11]

To this point, our perspective has been that of an international portfolio manager who focuses on returns in a home currency. This explains why we translated individual country returns into a common currency, the U.S. dollar. Another perspective might be that of an international economist or of a central banker who would naturally wonder about the role played by exchange rates in a crisis, as distinguished from the role played by local currency-denominated real equity returns.

Real returns cannot be easily computed on a daily basis because most countries report inflation only for monthly or even longer intervals. We can, however, examine nominal local returns. This should provide a fairly reliable picture because inflation has low volatility compared to equity returns in most countries. Rather than carry out additional formal tests, we thought it would

[9] The reader may wonder why the Asia–Europe curvature difference of −65.1 was not significant pre-crisis while the +89.5 crisis difference was highly significant even though it is not all that much larger in absolute value. This seeming puzzle is explained by differences between the bootstrapped resampling distributions. The pre-crisis period had more than twice as many observations (913 vs. 392) and the efficient frontier curvature's sample volatility increases with sample size. Intuitively this phenomenon must be related to higher apparent (as opposed to true) diversification potential in smaller samples; i.e., mean curvature also increases with sample size, ceteris paribus.

[10] It is, of course, conceivable that diversification potential within a region taken in isolation masks an offsetting potential of countries in the region when considered in the context of a globally diversified portfolio. However, this would only occur if the regional countries experience decreases in covariances on average with non-regional countries.

[11] We are grateful to the referee for bringing up the issue discussed in this section.

be sufficient to simply plot our statistical constructs computed from local returns and compare them with the previous plots. Figs. 2 (bis) through 6 (bis) provide the results. Each figure is plotted on the same scale as its corresponding dollar-denominated Figs. 2–6.

Figs. 2 and 2 (bis) show, not surprisingly, that currency depreciation and currency volatility played a major role in the Asian crisis. Mean returns are much more negative when denominated in dollars (Fig. 2a) than when denominated locally (Fig. 2 (bis) a). In contrast, there is only a minor difference between the two figures for European countries. Similarly, while European volatility is virtually indistinguishable between dollar and local returns, Asian volatility is considerably higher for most countries in dollars, thus revealing that currency fluctuations represented a substantial portion of the generally higher crisis period volatility. However, the degree of this currency volatility effect differs markedly among the Asian countries. Hong Kong, for instance, has none because it maintained a fixed exchange rate with the dollar throughout. Taiwan also has roughly the same volatility in local currency as in dollars. All other countries have higher volatility in dollars. Korea, Malaysia, and Thailand have about 20–25% more volatility in dollars while Indonesia's is twice as large.
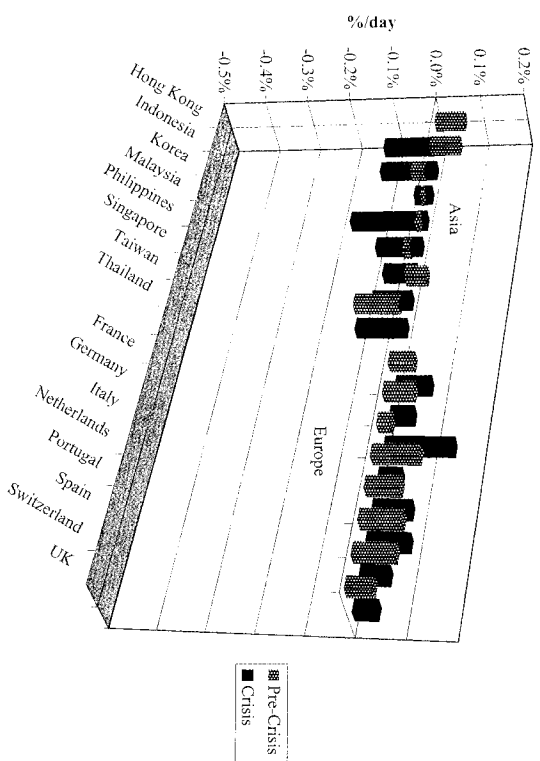

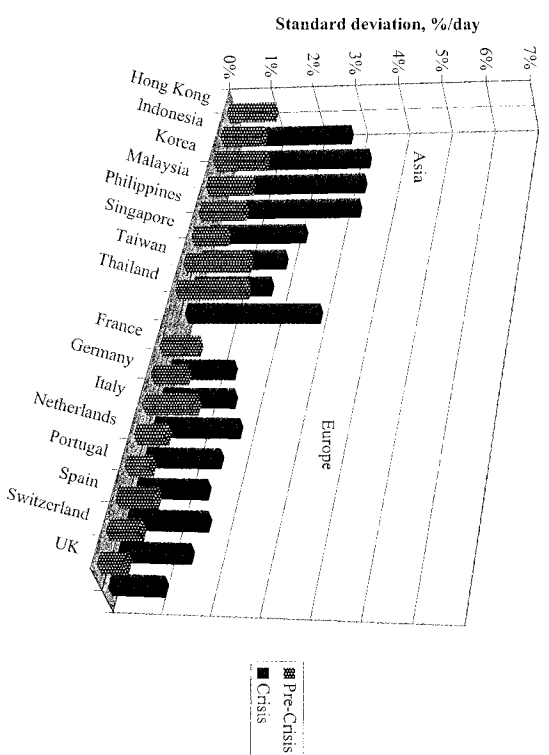Fig. 2. Panel A (bis). Mean returns by region and period in local currencies.

Fig. 2. Panel B (bis). Volatility by region and period in local currencies.

Turning now to the contrast between dollar-denominated and local currency denominated statistics, Figs. 4 and 4 (bis) show that average correlations were very similar. They are slightly lower in dollars, but any differences in the time pattern are barely distinguishable. Thus, our previous conclusion, that (1) European correlations were higher than Asian correlations both before and during the crisis and (2) Both regions experience an increase in correlation in the crisis, is unaffected by the currency denomination.

This is definitely not true for covariances. Fig. 3 shows the geometric average covariance in Asia rising to well over 6 ($\times 10^{-4}$) during the crisis when returns are denominated in dollars, but Fig. 3 (bis) shows an increase of only half that magnitude. In contrast, the European rise in covariance is actually slightly higher near the end of the crisis if returns are locally-denominated. Clearly, currency had a dramatic effect on the covariance matrix of Asian returns. As discussed above in describing Figs. 2a and b, volatility was also influenced substantially by currency for many Asian countries, though not for European countries. This is apparent in a comparison of the geometric average volatilities in Figs. 5 and 5 (bis).

Finally, the Jennrich statistic, a measure of changes in the covariance matrix, is similar whether calculated in dollar or in local returns. Prior to the crisis, the p-value was above 10% for most periods, indicating little statistically reliable
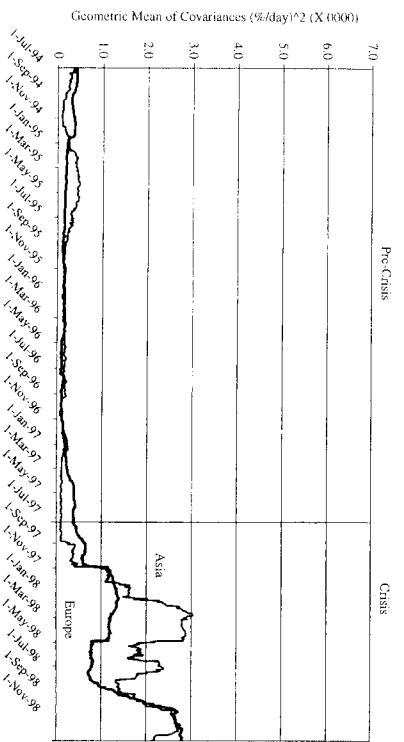
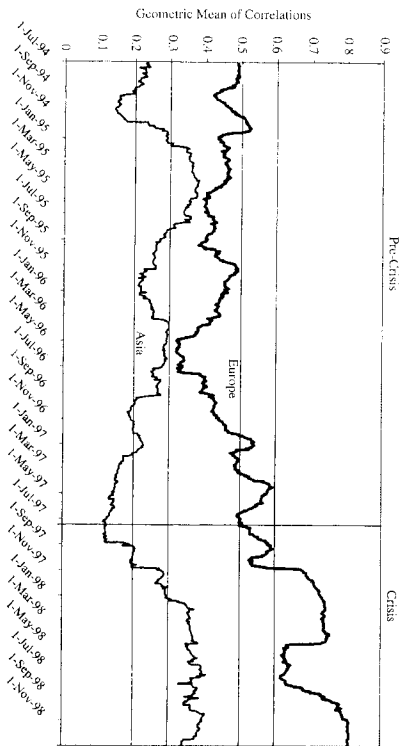Fig. 3. (bis). Regional average covariances in local currencies.

Fig. 4. (bis). Regional average correlations in local currencies.

evidence of variation in the covariance matrix. During the crisis, the p-values were quite small for both Europe and Asia, thereby suggesting formally that the entire covariance matrix was changing intertemporally within both regions. For Asia, the low p-value persists throughout the crisis period when returns are dollar-denominated. As shown in Fig. 6 (bis), however, it rises above 10% toward the end of the crisis; this suggests that intertemporal variations in the
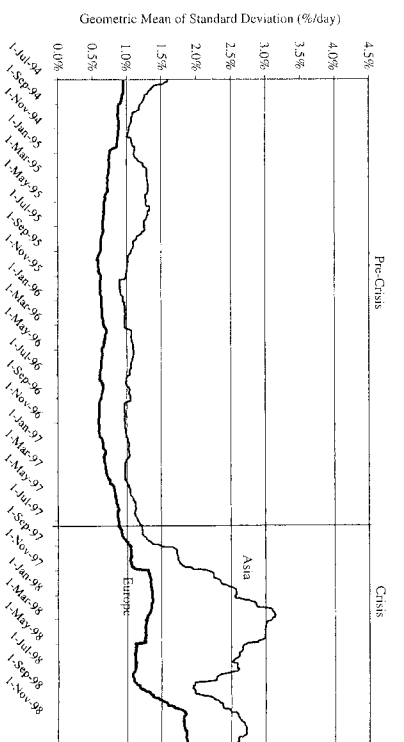
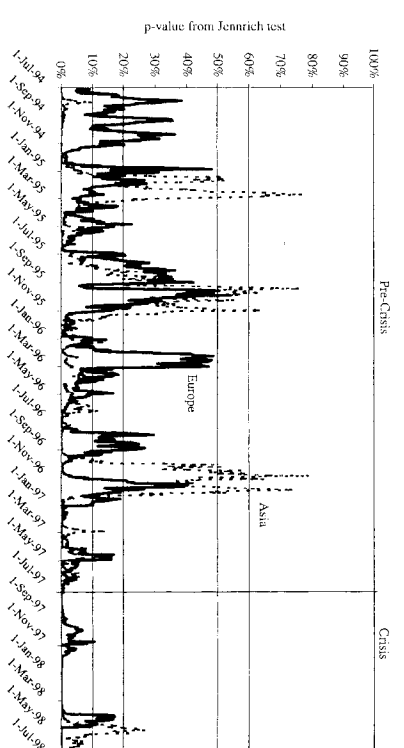Fig. 5. (bis). Regional average volatility in local currencies.

Fig. 6. (bis). Jennrich test of equality in successive covariance matrices; in local currencies.

covariance matrix became mainly attributable to currencies at that point in time.

6. Conclusions and future research

The 1997 Asian "crisis" brought staggering declines in equity values and exchange rates, which other regions escaped. Though painful to Asian

investors, this episode provides an opportunity for analysis and study. Perhaps we can learn something about precipitating factors in financial debacles and uncover ways to mitigate or even avoid them. Here, we offer a clinical comparison of the diseased patient, Asia, and a healthy patient represented by Europe. Our focus is on volatility spillover and diversification potential within each region. The basic idea is to examine and compare the two patients' financial fluids in an effort to understand the affliction.

To aid in this diagnostic endeavor, we design and apply some new statistical equipment to succinctly measure patterns of return co-movement and volatility within a region. These include general measures of covariance, correlation, volatility, and diversification potential.

In the three-and-a-half years prior to the crisis, Asian countries as a group had slightly higher volatility but somewhat lower correlation than their European counterparts. The general level of covariance was roughly the same in the two regions. There was no evidence of volatility spillover; i.e., return shocks were not propagated from a country to its neighbors in either region.

An increased level of return volatility and co-movement accompanied the onset of the crisis. Both Asia and Europe experience statistically significant increases in covariance, correlation, and variance. Correlation actually increased somewhat more in Europe than in Asia,[12] but this was swamped by much larger increases in covariance and in volatility in Asia. The volatility and covariance increases were both almost four times larger in Asia. This was exacerbated by the sudden appearance of volatility spillover, highly statistically significant in the case of Asia.

Diversification potential, as measured (inversely) by the curvature of the sample efficient frontier, tells a compelling story. Before the crisis, diversification was actually more potent within the Asian region, due perhaps to closer ties among countries within the European community than among the more diverse Asian economies. During the crisis this situation was reversed. Asian diversification potential became much worse. Its change from the pre-crisis to the crisis period was statistically significant and it fell behind Europe's diversification potential by a statistically significant margin. A large loss in diversification potency reduces the overall benefit of investing. That alone could precipitate sizable declines in asset values, though obviously we cannot know whether the observed change was a cause or an effect of the crisis.

A better understanding of regional interdependence will require something beyond stock returns. Capital movements, and event studies of major crisis incidents would be informative. A better understanding of the seeming link

---

[12] However, the relative increase in correlation was higher in Asia since it was considerably lower in the pre-crisis period. Within each region, volatility was slightly more responsible for the covariance increase than was correlation. (This is measured by our η statistic, defined in equation 5 of the text.)

between stock market movements and currency devaluations is also crucial. But no one is quite sure yet just how and why currency changes impact stock returns or vice versa. It would also be interesting to ponder theoretical reasons for the empirical findings in our clinical investigation.

The empirical procedures introduced here could be used unmodified in the study of industry linkages—to ascertain how industry or sectoral indices relate to one another and how that relation evolves over time. They could be applied with equal ease to study the relations among different stock exchanges within a country. They could measure the impact of legislation, policy changes, or technological innovations on the relations among industries within a country. Thus, we hope this clinical study will help in the diagnosis and resolution of other interesting cases.

Appendix A. Higher correlation in more volatile times: Bias or Fact?

Forbes and Rigobon (1999) (FB) argue that increased correlation during more volatile periods could be an artifact. They derive a "bias" in the correlation coefficient and apply an adjustment for the bias to three episodes, the 1987 crash in the U.S., the 1994 Mexican peso collapse, and the 1997 Asian crisis. Since the adjustment is strictly toward zero, they derive smaller correlations in all instances and thereafter find little evidence that these crises were characterized by "contagion," i.e., by higher positive inter-market correlations.

An appendix to their paper entitled "Proof of the Bias in the Unadjusted Correlation Coefficient" contains everything necessary to understand their argument. Adopting their notation, the true structural model is assumed to be

$$y_t = \alpha + \beta X_t + \varepsilon_t,$$

where $\beta \neq 0$ and the disturbance ε has standard spherical properties and is uncorrelated with x.

Now suppose we have two groups of observations, one with higher x volatility than the other. FB show, though it's rather obvious, that when the volatility of x increases while the volatility of ε remains the same, the correlation between x and y must increase (in absolute value).

It is crucial to understand a point not mentioned by FB. The two "groups" with differing x volatility could be (a) two different regimes during which x has disparate true volatility or (b) two sets of observations sorted by observed sample volatility. A bias exists only in the latter case.

In the former case (a), an increase in the volatility of x, ceteris paribus, obviously increases the true R^2 of the structural model and raises the absolute value of the correlation between x and y; (the sign of the correlation depends, of course, on β). In such a circumstance, higher volatility of x would be

associated *correctly* with higher volatility of y and higher correlation between x and y. On the other hand, if the true volatility of ε increased while x's volatility remained unchanged, the correlation between x and y would decrease even though the volatility of y would increase.

The above structural model specialized to a group of markets can be written as

$$r_{j,t} = \alpha_j + \beta_j f_t + \xi_{j,t}$$

where the subscript j indexes the country, r is a return, and f is some underlying driving factor. (Since countries are usually positively correlated, the β_j's are positive). Clearly, the general level of volatility (in the r's) could increase during a crisis because f's volatility increases or because the volatilities of the ξ_j's increase, or both. Correlation among the r's would increase if the root cause were f but it would decrease if the cause were ξ. This implies that there is no necessary connection between a general increase in volatility during a crisis and an increase in correlation. The fact that both volatilities and correlations seem to have increased during many crises is merely consistent with a rise in f volatility relative to ξ volatility.

FB seem to have in mind case (b) where both f and ξ have constant true volatilities but data miners have picked out a sub-sample of abnormally volatile f's. Since the true correlation is constant in this scenario, the correlation coefficient estimated from the sub-sample really is biased away from zero.

The basic issue, then, is whether crises have been identified because volatility just happened to be larger than usual. The crises studied by FB were characterized most obviously by dramatic declines in asset values. But a negative average return is a phenomenon related to the mean of the distribution, not necessarily to the variance. The increase in volatility after the onset of these crises was a fact discovered later, after they had already become infamous episodes. We do not, therefore, think any adjustment is necessary and, indeed, that such an adjustment would inappropriately bias estimated correlations downward from their true values.

References

Arshanapalli, B., Doukas, J., 1993. International stock market linkages: evidence from the pre- and post-October 1987 period. Journal of Banking and Finance 17, 193–208.

Bae, K., Karolyi, G.A., Stulz, R.M., 2000. A new approach to measuring financial contagion. Dice Center Working Paper, Fisher College of Business, Ohio State University.

Baig, T., Goldfajn, I., 1998. Financial market contagion in the Asian crisis. IMF working paper, International Monetary Fund, Washington DC.

Bansal, R., Lundblad, C., 2001. Market efficiency, fundamental values, and asset returns in global equity markets. Working paper, Fuqua School of Business, Duke University, February.

Boyer, B.H., Gibson, M.S., Loretan, M., 1997. Pitfalls in tests for changes in correlations. International Finance Discussion Papers No. 597, Board of Governors of the Federal Reserve System, Washington DC.

Brown, S.J., Goetzmann, W.N., Park, J.M., 2000. Hedge funds and the Asian currency crisis. Journal of Portfolio Management 26, 95–101.

Chowdhry, B., Goyal, A., 2000. Understanding the financial crisis in Asia. Pacific-Basin Finance Journal 8, 135–152.

Chung, P.J., Liu, D.J., 1994. Common stochastic trends in pacific rim stock markets. The Quarterly Review of Economics and Finance 34, 241–259.

Dumas, B., Harvey, C.R., Ruiz, P. 2000. Are correlations of stock returns justified by subsequent changes in national output? Working paper, Fuqua School of Business, Duke University, August.

Efron, B., 1982. The Jackknife, the Bootstrap and Other Resampling Plans. Society for Industrial and Applied Mathematics, Philadelphia.

Fernald, J., Edison, H., Loungani, P., 1999. Was China the first domino? Assessing links between China and other Asian economies. Journal Of International Money And Finance 18, 515–535.

Forbes, K., Rigobon, R., 1999. No contagion, only interdependence: measuring stock market co-movements. Working paper 7267, National Bureau of Economic Research, July.

Hamao, Y., Masulis, R.W., Ng, V., 1990. Correlations in price changes and volatility across International stock markets. The Review of Financial Studies 3, 281–307.

Jennrich, J.I., 1970. An asymptotic chi-square test for the equality of two correlation matrices. Journal of the American Statistical Association 65, 904–912.

Kamin, S.B., 1999. The current international financial crisis. Journal Of International Money And Finance 18, 501–514.

Kasa, K., 1992. Common stochastic trends in international stock markets. Journal of Monetary Economics 29, 95–124.

Kane, E.J., 2000. Capital movements, banking insolvency, and silent runs in the Asian financial crisis. Pacific-Basin Finance Journal 8, 153–175.

Kaplanis, E.C., 1988. Stability and forecasting of the co-movement measures of international stock market return. Journal of International Money and Finance 8, 63–75.

Kho, B., Stulz, R.M., 2000. Banks, the IMF, and the Asian crisis. Pacific-Basin Finance Journal 8, 135–152.

King, M.A., Wadhwani, S., 1990. Transmission of volatility between stock markets. The Review of Financial Studies 3, 5–33.

Koch, P.D., Koch, T.W., 1991. Evolution in dynamic linkages across national stock indexes. Journal of International Money and Finance 10, 231–251.

Longin, F., Solnik, B., 1995. Is the correlation in international equity returns constant: 1960–1990? Journal of International Money and Finance 14, 3–26.

Masson, P., 1999. Contagion. Journal Of International Money And Finance 18, 587–602.

Mohamad, M., 1997. Highwaymen of the global economy. The Wall Street Journal, September 23, A22.

Richards, A.J., 1995. Comovements in national stock market returns: evidence of predictability, but not cointegration. Journal of Monetary Economics 36, 631–654.

Roll, R., 1977. A critique of the asset pricing theory's tests. Journal of Financial Economics 4, 129–176.

Scott, D.W., 1992. Multivariate Density Estimation. Wiley, New York.

So, M.K.P., Lam, K., Li, W.K., 1997. An empirical study of volatility in seven Southeast Asian stock markets using ARV models. Journal of Business Finance and Accounting 24, 261–274.