This product is part of the RAND Corporation occasional paper series. RAND occasional papers may include an informed perspective on a timely policy issue, a discussion of new research methodologies, essays, a paper presented at a conference, a conference summary, or a summary of work in progress. All RAND occasional papers undergo rigorous peer review to ensure that they meet high standards for research quality and objectivity.

OCCASIONAL PAPER

Linkages in World Financial Markets

C. Richard Neu • Jack Clift

CENTER FOR GLOBAL RISK AND SECURITY

International Programs at RAND

This research was sponsored by Harold Brown, chair of the advisory board of the RAND Center for Global Risk and Security (CGRS), and was conducted within CGRS, part of International Programs at the RAND Corporation.

Published 2011 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665
RAND URL: http://www.rand.org
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Preface

This paper had its origins in a conversation among senior analysts at RAND and Harold Brown, chair of the advisory board of the RAND Center for Global Risk and Security (CGRS). At the time, the U.S. and world economies were showing the first indications of recovery from the global financial and economic crisis that had begun in the summer of 2007. Hoping that the worst would soon be behind us, the senior RAND analysts were discussing the possible consequences of the crisis and the policies adopted by various countries to cope with it. In particular, we wondered whether the crisis would have any lasting consequences for U.S. leadership in international economic, financial, or security affairs.

There was a consensus that the full consequences of the crisis would probably take years to become manifest and just as long for scholars to tease out all the intertwined strands of causality. Two questions about the crisis, however, seemed amenable to immediate analysis: Had the crisis made any difference in the way that short-term movements in various financial markets are linked to each other? And had anything happened to the traditional position of U.S. markets in setting the pattern for movements in global financial markets? This paper is an early examination of these questions.

There are, of course, many financial markets in the world and many possible linkages among them. We chose to begin this analysis with an examination of linkages among major equity markets—in the United States, the UK, and Japan. Similar analyses of other financial linkages are possible, and perhaps the findings of this research will raise additional questions and suggest additional analyses.

This work was supported by a generous contribution from Harold Brown, who chairs the CGRS advisory board. This research was conducted within CGRS, part of International Programs at the RAND Corporation. The center aims to improve public policy by providing decisionmakers and the public with rigorous, objective research on critical, crosscutting policy issues affecting global risk and security.

For more information on CGRS, see http://www.rand.org/international_programs/cgrs/ or contact the director (contact information is provided on the web page).

Contents

Figures

Tables

Summary

Assessing the full consequences of the global financial crisis that began in the summer of 2007 will require years. It is possible now, however, to examine how short-term (i.e., day-to-day) linkages in different financial markets have changed since the onset of the crisis. In particular, it might be interesting to ask whether there has been any diminution in the traditional role of U.S. financial markets in leading movements in other financial markets.

We cannot observe directly what events triggered short-term movements in U.S. or other financial markets. We can, however, observe *when* market movements appear to have originated—that is, which markets tend to lead subsequent movements in other markets.

In this short paper, we examine daily movements in three major equity indexes—the U.S. S&P 500, the Japanese Nikkei 225, and the British FTSE 100. We examine how daily changes in one market are correlated with the immediately subsequent changes in the other two. We also examine the size of movements in each market—the "betas"—subsequent to a given change in the other markets before and after the onset of the crisis.

We use two alternative dates for the onset of the crisis: August 1, 2007, when the first signs of trouble with subprime mortgages began to emerge, and September 1, 2008, before the failure, forced sale, or government takeovers of high-profile financial institutions in the United States and elsewhere.

Because UK markets are still open when U.S markets open, we restrict our data to morning trading in the UK, "unpolluted" by news emerging in the couple of hours immediately before the markets open in New York.

Generally, we find that the three national equity indices examined became more highly correlated after the onset of the crisis. This result is robust with respect to the choice of starting date for the crisis and is consistent with the findings of other studies of other financial crises. The returns on individual financial instruments and broader market indices tend to become more correlated during times of crisis.

More novel is our finding that the size of one market's movements subsequent to movements in other markets—the beta of one market with respect to another—also increased.

The links between movements in the U.S. market and subsequent movements in other markets do not strengthen as much as do links originating in the Japanese and the UK markets, but the links from U.S. markets to other markets start at a much higher precrisis level (an indication of the leading role that U.S. markets seem to play in global financial affairs). In our analysis, U.S. influence on the other markets is the strongest influence both before and after the onset of the crisis and does not appear to have diminished in absolute terms in the period immediately following the onset of the crisis.

In an excursion from our basic analysis, we eliminate all Mondays from our sample of daily market movements. The rationale for this excursion is that some important developments might occur over the weekend, allowing Asian or European markets to react before U.S. markets open, even if the news originated in the United States. Generally, "excluding-Mondays" analysis shows the same patterns as the analysis using the full data set. The one interesting exception is that the increased influence of UK market movements on subsequent U.S. market movements observed in the full sample becomes insignificant in the excluding-Mondays analysis. We interpret this to mean that the apparent increase in the influence of UK markets was due mostly to common reactions in the United States and the UK to events that took place over the weekend.

In another excursion, we examined linkages among movements in financial equity subindices in the United States, Japan, and the UK. Our rationale was that, since the recent crisis was particularly threatening to financial stocks, we might expect to see more-dramatic changes in the linkages among financial subindices. Somewhat surprisingly, we found no consistent pattern of changes in linkages among financial subindices. One interpretation of this finding is that developments that were detrimental to, say, U.S. financial institutions might have conferred some advantage on, say, Japanese institutions.

In a final excursion, we allow the date for the onset of the crisis to vary across our entire range of data, searching for the dates that yield the best fit in our two-period (precrisis and after onset of the crisis) models. No single date maximizes the fit for models that seek to explain linkages among all three markets. Generally, though, we find that dates for the onset of the crisis in September and October 2008 yield the best fit. Curiously, a date of January 1, 2007—well before a plausible beginning date for the crisis—yields good fits for models of the influence of UK markets on other markets.

Acknowledgments

We are grateful for the generous financial support of Harold Brown, chair of the advisory board to the RAND Center for Global Risk and Security. All members of this advisory board posed challenging questions and offered helpful suggestions as this work proceeded.

Roger Kubarych of the National Intelligence Council and our RAND colleagues Andrew Weiss and Lionel Galway provided insightful and detailed technical reviews of this document.

The efforts of all of these people made this paper stronger. Any errors that remain despite these efforts are those of the authors.

SECTION 1

Introduction

The international financial crisis that began in 2007 prompted many questions about the structure and functioning of financial markets and the linkages among different markets. The proximate origins of the crisis lay in problems with mortgages in the United States, and large U.S. financial institutions have been the most prominent—although certainly not the only—casualties of the crisis. Nonetheless, no major financial market escaped the turbulence that began with the U.S. subprime-mortgage crisis.

A natural question, perhaps, is whether the crisis has changed the ways in which different markets interact with other. Do developments in one country or one market now exert more or less influence on subsequent developments in other markets? Do policies, institutions, or events in one country now prompt different reactions in markets elsewhere from those they prompted before the onset of the crisis?

Years will pass, of course, before full answers to these questions can be offered, not least because regulatory, policy, and behavioral responses to the crisis are still in progress.

In the immediate aftermath of the crisis, however, we can examine one potentially interesting question: Have the short-term linkages among various financial markets changed in significant ways? More specifically, has there been some observable change in the way in which short-term (say, day-to-day) movements in one market relate to or respond to previous movements in another market?

In this paper, we explore the relationships in day-to-day movements in three major equity markets—in London, New York, and Tokyo—asking how the correlations among price movements in these markets have changed from before the crisis to after the crisis. We also explore how the betas—the percentage changes in one market in response to a previous percentage change in another market—have changed.[1]

Identifying the event or news that causes any particular market movement is difficult. Indeed, it is typically impossible. By taking advantage of the fact that the London, New York, and Tokyo markets are open at different times, however, we have some opportunity to observe *when* a particular market movement began.

It would be going too far to try to infer from the timing of market movements where the "news" that precipitated those movements originated. In today's financial markets, "news"

[1] This is very much a preliminary analysis of how market dynamics changed after the onset of the recent crisis. Much work remains to be done on this subject. We chose equity markets as a place to begin this assessment principally because data on equity-price movements are readily available. Of course, the origins of the recent crisis lay in funding markets. But tracking changes in interest rates in different markets is highly problematic when credit becomes unavailable—at least to some borrowers—at any price. Careful analysis of the ebb and flow of credit to particular borrowers or classes of borrowers will be an important part of the history of this crisis. But that analysis still lies in the future.

is not restricted to the release of economic data or statements by central bankers, government officials, or company spokespersons. Particularly in an unsettled environment, relevant "news" might take the form of judgments by analysts, observations of who seems to be buying or selling what, or just plain rumors. With analysts and traders active around the world and around the clock, markets in futures and other derivatives allow trading in whole asset classes even when the traditional "cash" markets for the underlying assets are closed. "News" can arise anywhere and at any time and be acted upon.

Neither should we take for granted that all relevant news is immediately reflected in market prices. Particularly in confused circumstances in which events are happening rapidly, analysts and traders might require some time to understand the meaning of new information. Consequently, news that moves markets could be hours or possibly days old.[2]

Nonetheless, equity markets are not yet homogeneous. Different shares are traded in Tokyo than in London or in New York. Some price movements might be more pronounced in one market than in another. Movements in some markets might be more potent in influencing subsequent market movements than movements in other markets. And—the proposition we aim to examine in this paper—the patterns of this market-to-market influence might change with the onset of the crisis.[3]

---

[2] In Section 2, we report some analyses of market linkages when Mondays are excluded for the sample. Some of the differences we observe with Mondays in and out of the sample might reflect the fact that analysts and traders have additional time over the weekend to absorb new information.

[3] An interesting and important extension of this work would be to establish the extent to movements in futures markets presage movements in the underlying "cash" markets and influence prices in other markets that are open contemporaneously. We do not have access to detailed time series on futures prices, however. Efforts to understand the recent crisis might result in more available data on futures and other derivative markets.

SECTION 2

Data

We acquired data from a variety of sources to test our hypotheses. Preliminary analyses were conducted using daily closing price data (freely available from Yahoo! Finance [undated]), from August 1, 2005, to June 19, 2009, for the S&P 500, Nikkei 225, and FTSE 100.

To control for fluctuations in currency, we used conversion rates published daily by the U.S. Federal Reserve (Board of Governors of the Federal Reserve System, 2011).

For more-granular analysis of the overlapping U.S. and European markets, we produced hourly prices from tick-by-tick data purchased from TickData (Tick Data, undated).

For extensions to our analysis, we downloaded daily closing prices for the KBW Large Cap Banking Index from Yahoo! Finance and for the FTSE 350-Banks Index from Yahoo! Finance UK (Yahoo! Finance UK and Ireland, undated), and purchased daily closing prices for the TOPIX-17 Banking Index from the Tokyo Stock Exchange (undated).

SECTION 3

Methodology

Our basic unit of analysis is the change in a stock market index from one period to the next, adjusted for currency fluctuations and expressed on a logarithmic scale for interpretation as a percentage change. We are interested in how a change in one country's index is related to the subsequent change in another country's index—for example, how is a change in the S&P 500 on one day related to the change in the Nikkei 225 on the next day? Moreover, we are interested in whether these relationships among movements of indices in different countries changed with the onset of the global financial crisis. Previous literature in this area was spurred by the effects of crises in the late 1980s and early 1990s; the earliest article in the spirit of our paper analyzed cross-market correlations in the United States, the UK, and Japan after the 1987 stock market crash in the United States, finding a significant increase in correlation after the crash (King and Wadhwani, 1990). A wider literature is surveyed in two more-recent survey articles, with less-conclusive results (Karolyi, 2003; Dungey et al., 2005).

We examine the data using two related but distinct tools: correlation analysis and regression analysis.[1] In both cases, we split the data into two segments, representing periods before and after the onset of the global financial crisis, and compare the relationships seen in the before period with the relationships seen in the after period. We then formally test whether the relationships are statistically significantly different between periods.

There is no clear-cut start date for the onset of financial crisis. We select two candidate dates for the beginning of the crisis and run our analyses using each date, and report results for both.

Our first candidate date is August 1, 2007. Although there were some signs of problems in the subprime-mortgage market earlier in 2007,[2] the wider U.S. economy appeared largely unaffected, with the S&P 500 and Dow Jones Industrial Average continuing to rise steadily. August 2007 was the beginning of the global "credit crunch," with fears over mortgage-backed security holdings freezing liquidity and necessitating coordinated international action from central banks to inject liquidity into markets. The credit crunch led to severe difficulties for British bank Northern Rock, which required liquidity support from the Bank of England. August 2007 also saw American Home Mortgage filing for Chapter 11, and Countrywide

[1] Both methodologies are used to evaluate linear relationships and would not be appropriate if the assumption of linearity were not a reasonable approximation: An examination of data plots and analysis of residuals show no systematic deviations from linearity.

[2] Ownit Mortgage Solutions Inc filed for Chapter 11 in January; Mortgage Lenders Network USA Inc filed for Chapter 11 in February; New Century Financial reported several bombshells in March before filing for Chapter 11 in April.

Financial Corporation drawing on $11.5 billion in emergency lines of credit from a group of 40 banks and seeing its shares plummet.

Our second candidate date is September 1, 2008. Although there were some important events leading up to this date,[3] the events following this date were particularly dramatic: Fannie Mae and Freddie Mac were placed into conservatorship by the U.S. government; Lehman Brothers filed for Chapter 11; the U.S. Federal Reserve provided emergency credit support for major insurer AIG; Merrill Lynch, Washington Mutual, Wachovia in the United States, and HBOS in the UK were all subject to crisis takeovers. The U.S. government put forward legislation for a $700 billion bailout package, and European countries acted to nationalize, part-nationalize or otherwise prop up a number of financial institutions.

These two dates make conceptual sense a priori as potential moments of change for global financial linkages. In Appendix C, we perform a purely statistical model-fitting exercise treating every date from the beginning of the data set (August 1, 2005) to the end (June 19, 2009) as a potential start date for the crisis.

Correlation Analysis

Correlation is a measure of the strength of association between two variables. If market A is always up when market B is up, and always down when market B is down, they are highly (positively) correlated. If market A is always up when market B is down, and vice versa, they are highly negatively correlated. If market A is just as likely to be up or down regardless of what is going on in market B, the markets are uncorrelated.

We expect global markets to be positively correlated—that is, we expect them to tend to move in the same directions as each other—both before and after the onset of the global financial crisis. We are interested in seeing whether markets move more in lockstep after the onset of crisis, or whether the chaos would lead to markets in different countries displaying more-independent behavior.

For intuition, two different correlations are depicted graphically in Figures 1 and 2. In both cases, a one-unit change along the x-axis is associated (on average) with a 0.5-unit change along the y-axis. However, the consistency of the relationship, and the strength of the correlation, is significantly higher in Figure 2.

Our first analysis tests for changes in correlation before and after the onset of the crisis. In order to compare correlations in two separate time periods, we use the procedure developed by Fisher to produce normalized correlation coefficients from the typical Pearson correlation coefficients, and formally test whether any differences are statistically significant (Fisher, 1921).

Normalized correlation coefficients are produced using the following conversion:

$$\rho' = \frac{1}{2}\ln\left(\frac{1+\rho}{1-\rho}\right),$$

where ρ is the standard Pearson correlation coefficient. Testing the difference between two normalized correlation coefficients involves the following test statistic:

[3] Notably the collapse in Bear Stearns' share price and subsequent fire sale acquisition by JP Morgan Chase.

Figure 1
Weaker Correlation (ρ = 0.66)

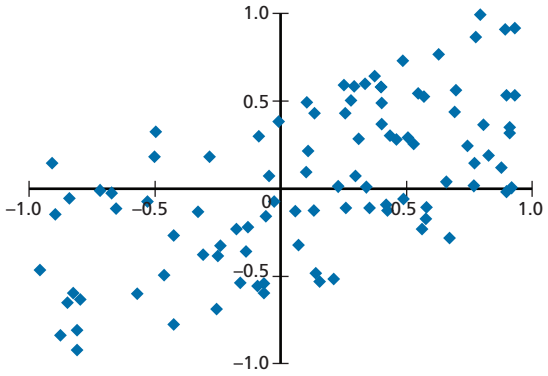

RAND *OP317-1*

Figure 2
Stronger Correlation (ρ = 0.90)

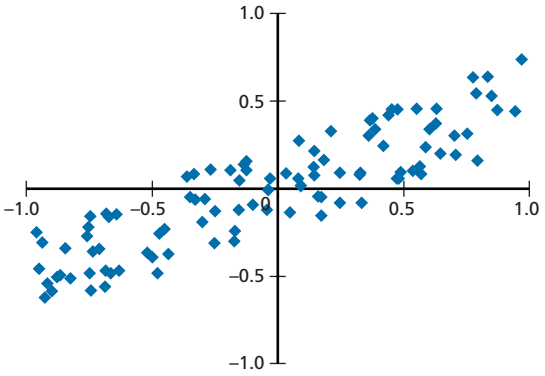

RAND *OP317-2*

$$z = \frac{\rho'_1 - \rho'_2}{\sqrt{\frac{1}{n_1 - 3} + \frac{1}{n_2 - 3}}}$$

.

The *z*-score can then be compared with the standard normal distribution in order to obtain a *p*-value, which expresses the likelihood that a difference in correlations of this size could be produced by random chance even if the "true" underlying correlation is the same in both cases.

Regression Analysis (i): Two-Country Models

Regression analysis allows us to go beyond just the consistency of the relationship between two variables, to examine the "size" of the effect of one variable on another. In our study, this means that we can test whether a specific movement in one market is associated with a larger (or smaller) movement in the following market after the onset of the crisis. We are testing whether a regression on the data before the onset of the crisis produces a different beta-coefficient (or *slope*) from that produced by a regression on the data after the onset of the crisis.

For intuition, scatterplots with different beta-coefficients are presented in Figures 3 and 4. In Figure 3, a one-unit change along the x-axis is associated with a 0.5-unit change along the y-axis, for a beta-coefficient of 0.5. In Figure 4, a one-unit change along the x-axis is associated with a 1.4-unit change along the y-axis, for a beta-coefficient of 1.4.

The regression specification we use to examine these relationships is as follows:

$$Y = \beta_1 X + \beta_2 X^{After} + \varepsilon,$$

where Y is the change in the market index and X is the change in the preceding market—for example, when Y represents the change in the U.S. market, X represents the change in the UK

Figure 3
Depiction of Beta 0.5

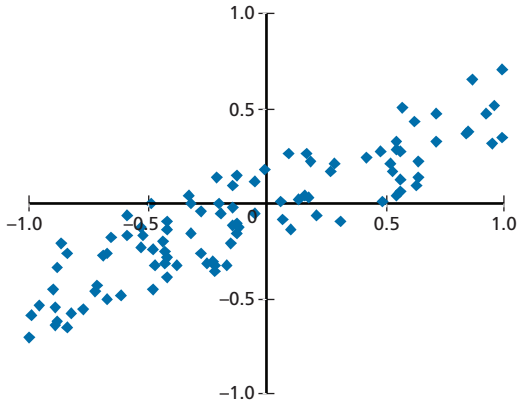

Figure 4
Depiction of Beta 1.4

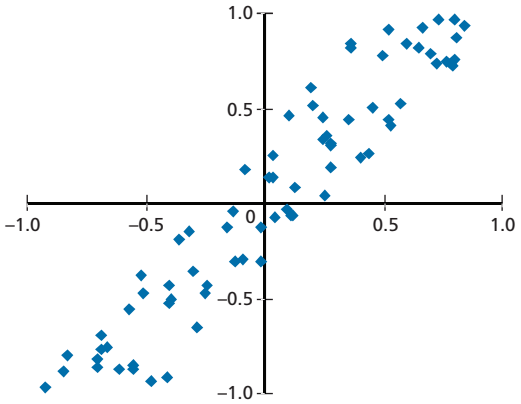

market. X^{After} is an interaction term that equals 0 for dates prior to the onset of the crisis, and otherwise is equal to the change in the preceding market; ε is a random error term.[4]

If the relationship between Y and X is the same after the onset of the crisis as it was before, the relationship will be well explained by the $\beta_1 X$ term, and β_2 will then be close to 0 and statistically insignificant. If the relationship is significantly different after the onset of the crisis, β_2 will be statistically different from 0 and will represent the size of the difference between the relationship in the before and after periods.

Regression Analysis (ii): Three-Country Models

Regression analysis also allows us to consider simultaneously the relationship between market movements in one country and market movements in *both* preceding markets—e.g., we can relate movements in the S&P 500 to movements earlier in the day in both the FTSE 100 and the Nikkei 225.

The regression model is

$$Y = \left(\beta_1 X_1 + \beta_2 X_1^{After} + \beta_3 X_2 + \beta_4 X_2^{After}\right) + \varepsilon,$$

where X_1 represents the market movement in the immediately preceding market and X_2 represents the market movement in the prior market.

Including both preceding markets in the analysis might provide more nuance to our analysis. For example, if the apparent relationship between the Japanese and UK markets is really driven by both Japanese and UK markets reacting to movements in the United States, the three-country model should provide insight into this and help us to identify some amount

[4] We do not include a constant term in this regression or a dummy term for the after period, and so implicitly restrict both values to being 0. Our initial specifications showed that these terms were never significant from a statistical perspective, and, from a conceptual perspective, we prefer a model that passes through the origin.

of correlation as spurious.[5] In each case, it is likely that the immediately preceding market will appear more strongly related than the prior market, as the immediately preceding market has access to all the same news as the prior market and additional news arising after the prior market closed; any deviations from this will be a source of interest.

Calculating the Relevant Index Movements

Our main results focus on movements in broad market indices in the United States (S&P 500), the United Kingdom (FTSE 100), and Japan (Nikkei 225). The opening hours of the Tokyo, London, and New York stock exchanges are as shown in Figure 5, with reference to local London time (LLT).[6]

When comparing movements at the Tokyo Stock Exchange with market movements in London or New York, we can simply calculate the change in closing price from one day to the next in each market.[7] Among our main results, we compare the close–close change in Tokyo with the close–close change in London on the same day, and the close–close change in New York with the close–close change in Tokyo the next day.

However, when exploring the relationship between market movements in London and New York, the situation is complicated by the overlap in opening hours. If the close–close changes appear highly correlated, there is no easy way of distinguishing whether New York is reacting to movements in London, or whether London traders at the end of their day are reacting strongly to the opening trading in New York.

Therefore, we use intraday index price data to evaluate changes other than the standard close–close changes. For our main analysis, we calculate the change from close[t – 1] to

Figure 5
Opening Hours of Tokyo, London, and New York Stock Exchanges

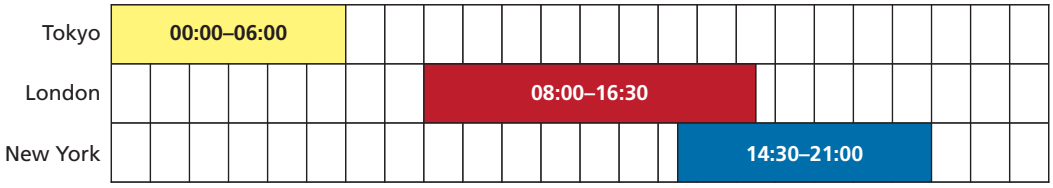

RAND *OP317-5*

[5] The correlation between movements in the two "independent" countries in each regression could lead to concerns about multicollinearity. An examination of variance inflation factor (VIF) statistics for three-country regressions (without the interaction terms included) suggests that multicollinearity is not likely to be a problem for the model.

[6] Due to daylight saving time, it would be incorrect to conflate LLT with Greenwich mean time (GMT). However, the trading hours of the Tokyo market are 0000–0600 GMT all year round, as Japan does not observe daylight saving time; this means that the market is open 2300–0500 LLT during the summer. This does not affect the analysis in any way, as there is always at least a two-hour gap between the close in New York and the open in Tokyo.

[7] For the results that we report, we converted all values into U.S. dollars to control for currency fluctuations. To ensure the robustness of our results, we also repeated our correlation and regression analyses of full index movements before and after the financial crisis without currency conversion. Although estimates changed somewhat, the basic results and patterns of statistical significance remained the same, and the differences in estimates were not systematic either in direction or significance.

12:00 LLT [t] in London and compare this with the change from close[t – 1] to close[t] in New York.[8] In a later analysis testing for effects in the opposite direction, we calculate the change from 16:30 LLT [t – 1] to close[t – 1] in New York and compare this with the change from close[t – 1] to 12:00 LLT [t] in London. In both cases, the idea is to focus purely on market movements that take place when only one of the exchanges is open.[9] We cut the London market at 12:00 LLT—rather than, say, 14:00 LLT, shortly before the market opens in New York—because important U.S. news is often released in the hours before the New York market opens, and we do not want reactions to U.S. economic news to show up in our London market movements before it hits the New York market.

A Note on Timing: A Case of the Mondays

Implicit in the models we use is that relevant "news" comes out in each region during—or slightly prior to—the trading day. Thus, market movements in Japan reflect reaction to previous news that has already affected other markets and additional new information that has arisen in Japan. When looking at how UK market movements are connected to Japanese market movements, the assumption is that all relevant Japanese news has been incorporated into the movements of the Japanese market.

As we noted earlier, the assumption might not always be valid. For example, information relevant to one country might come out too late to affect market prices in that country but early enough to be reflected in the following country's markets. Although these anomalies during the trading week might affect each country similarly,[10] the weekend potentially poses difficulties: Any weekend news will be incorporated into Asian markets first on Monday morning, then European markets, then U.S. markets, regardless of the origin of the news. For example, the difficulties at Lehman Brothers led to the International Swaps and Derivatives Association holding an extraordinary trading session on Sunday, September 14, to allow other firms to offset the counterparty risk associated with their relationships with Lehman Brothers,[11] and Lehman Brothers eventually announced that it would be filing for Chapter 11 bankruptcy protection shortly before 1:00 a.m. in New York on Monday, September 15. The bankruptcy announcement came more than an hour before the end of trading in Tokyo and before the start of the trading day in London, giving both markets some opportunity to react to the news before the New York market opened.

In order to ensure that our results are robust to potential weekend-driven anomalies, we run all our analyses both using the complete data set and using a subsample of the data set that

---

[8] Daylight saving time occurs at slightly different times in the United States and United Kingdom. Our analysis adjusts the calculation of changes to take account of this.

[9] When using three-country models, the situation is slightly more complicated—when investigating the relationship between UK market movements and the two preceding markets, it is necessary to invoke the overlap solutions above, so that the 0800–1200 LLT movement in London is the dependent variable. This means that the Japan coefficients in the three-country models are not directly comparable to the Japan coefficients in the two-country model, in which the dependent variable is the UK close[t – 1] – close[t] change.

[10] This might not be true if there are systematic differences relating to information disclosure, timing of business and government announcements, or other factors across countries. Although an exhaustive analysis of this issue was not possible, we have found no evidence of problematic systematic differences of this nature.

[11] The press release from the International Swaps and Derivatives Association (ISDA) is available online (ISDA, 2008).

excludes Monday market movements. For sake of concision, we report only the complete data set results in full and report the no-Mondays results in the main body of this paper only on the rare occasions where they are substantively different from the complete data set results. A full comparison of the full sample results and no-Mondays results is in Appendix A.

SECTION 4

Main Results

Correlation Results

The formal tests of significance we report for our correlation results are calculated using the simple correlation coefficients normalized using the Fisher transformation, as described in Section 2 of this paper. However, for ease of explication, we will describe correlation changes in terms of the change in R^2 values for the linear relationship between the two variables. The R^2 for a linear relationship between two variables can be calculated by squaring the simple correlation coefficient and has an explicable interpretation as the fraction of variation in one variable explained by the variation seen in the other variable.

In Table 1, we present the results for both the "early cutoff," in which the onset of the crisis is taken as August 1, 2007, and the "late cutoff," in which the onset of the crisis is taken as September 1, 2008.

The results show that market movements tend to be positively correlated and that the positive correlation increased significantly after the onset of the crisis (however defined) with the exception of the correlation of U.S. and Japanese markets using the late cutoff, where there is a statistically insignificant increase. While it might be tempting to conclude that the correlation between UK market movements and the following U.S. market movements is much weaker than the others, it is important to remember that dealing with the overlapping-market problem means that we are using only market movements in the first half of the UK trading day and movements in the second half of the U.S. trading day, which is likely to weaken the size of the

Table 1
Correlation Results

Markets		R^2		Difference	
Following	**Preceding**	**Before**	**After**	**Z-Score**	**P-Value**
Early cutoff					
UK	Japan	0.0470	0.2999	5.6387	0.0000
United States	UK	0.0366	0.1114	2.3274	0.0199
Japan	United States	0.1270	0.2971	3.3821	0.0007
Late cutoff					
UK	Japan	0.1333	0.3350	3.0930	0.0020
United States	UK	0.0319	0.1410	2.5752	0.0100
Japan	United States	0.2144	0.3058	1.3103	0.1901

correlations. Nevertheless, the link between movements in the UK and subsequent movements in the United States appears to be stronger after the onset of the crisis in the model above.

However, a further analysis using the data subset that excludes Monday market movements casts doubt on the U.S.-UK pattern[1] above: Whereas the full sample seems to demonstrate a positive significant change between the before and after periods, the no-Mondays sample shows a small, positive but statistically insignificant change between the before and after periods, regardless of whether it is the early or late cutoff that is used (see Table 2).

As we are using fewer data points when we exclude Monday market movements, we would expect some drop in statistical significance as our estimates become less precise. However, the most-natural reading of this additional analysis is that the original full-sample results were driven by events late on Friday, over the weekend, or early on Monday morning, to which the UK market was able to react before the U.S. market opened. This does not prove that the full-sample results are necessarily spurious—if the relevant weekend events were all UK based, the timing would not be an issue—but does make it difficult to take the full-sample results at face value.[2]

Regression Results (i): Two-Country Models

We divide our regression results into the "main effect" and what we dub the *crisis effect*. Through the structure of the regression, the main effect reflects the relationship between markets before the onset of crisis, and the crisis effect (the interaction term described in the methodology section) reflects the difference between the before and after periods. The relationship between markets after the onset of the crisis is therefore the summation of the main effect and crisis effect. If there is no change in the relationship before and after the onset of the crisis, the crisis effect β and t-score will be close to 0 and the P-value of the crisis effect will be large.[3] If there is a statistically significant difference before and after the crisis, the t-score of the crisis effect will be large and the P-value will be low, and below the critical value.

The regression was based on a logarithmic scale, giving the beta-coefficients the interpretation of the percentage change seen in the following market given a 1-percent change in the preceding market.

The results in Table 3 show that the movement seen in the Japanese market after a movement in the U.S. market is larger than the movement seen in the UK following Japan, or the United States following the UK. This result holds with both the early and late cutoffs and holds both before and after the onset of crisis: The summation of betas for the Japan-U.S. result is higher than the summation of betas for the other two relationships.

However, the Japan-U.S. relationship expressed here does not change significantly after the onset of crisis—the crisis effect is close to 0 and not statistically significant—whereas the

---

[1] In keeping with the layout of the tables, *U.S.-UK pattern* refers to the results in which the United States is the following market and the UK is the preceding market.

[2] A day-by-day analysis of all major newsworthy events would shed some light on this issue but is beyond the scope and resources of this paper.

[3] There is some redundancy in presenting t-scores and P-values, but both are included in this paper because different readers might be familiar with one or the other. We treat a P-value of 0.05 or lower as the standard measure of statistical significance, with P-values between 0.05 and 0.10 noted as marginally significant.

Table 2
Comparison of Full-Sample and No-Monday Results, U.S.-UK Correlation

U.S.-UK Correlation	R^2		Difference	
Sample	**Before**	**After**	**Z-Score**	**P-Value**
Early cutoff				
Full sample	0.0366	0.1114	2.3274	0.0199
No Mondays	0.0359	0.0514	0.5575	0.5772
Late cutoff				
Full sample	0.0319	0.1410	2.5752	0.0100
No Mondays	0.0335	0.0580	0.6626	0.5076

UK-Japan and U.S.-UK effects become significantly larger after the onset of the crisis. With the early cutoff, a 1-percent movement in Japan before the crisis is followed by a 0.15-percent movement in the UK, rising to a 0.47-percent movement in the UK following a 1-percent movement in Japan after the onset of crisis. The change in effect size between the before and after periods is highly statistically significant for UK-Japan with both the early and late cutoffs and for the U.S.-UK late cutoff result; the change in effect size for the U.S.-UK early-cutoff result is marginally significant ($0.05 < p < 0.10$).

Again, the Japan-U.S. and UK-Japan relationships are not sensitive to restricting the data to exclude Monday movements, but the U.S.-UK relationship does seem less clear when the analysis is rerun on the restricted data set.

The results in Table 4 suggest that the main effect is relatively insensitive to the choice of sample or before/after cutoff; but the crisis effect is much smaller and statistically insignificant when the no-Mondays data subsample is used. While it is possible that the full-sample results are correct, the restricted results raise the possibility that the apparent increase in effect size after the onset of crisis is an artifact of weekend news being incorporated into UK markets before U.S. markets open, rather than a reflection of increased influence of UK market movements on U.S. markets.

Regression Results (ii): Three-Country Models

The results in the three-country model, as with the two-country models, are divided into main effects and crisis effects, with the main effect representing the relationship between market movements before the onset of the crisis, and the crisis effect representing the change in the relationship after the onset of the crisis. In the three-country model, we denote the dependent market as Y, the immediately preceding market movement as X_1, and the prior market as X_2. Thus, the meaning of X_1 and X_2 varies conditionally on which market is treated as the dependent market.

The results in Table 5 add some additional nuance to the previous results. The following features are notable:

Table 3
Regression Results

Markets		Main Effect				Crisis Effect			
Following	**Preceding**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Early cutoff									
UK	Japan	0.1510	0.0593	2.5476	0.0110	0.3215	0.0653	4.9270	0.0000
United States	UK	0.1742	0.0945	1.8427	0.0657	0.1682	0.1004	1.6756	0.0942
Japan	United States	0.6519	0.1310	4.9776	0.0000	0.0415	0.1375	0.3017	0.7629
Late cutoff									
UK	Japan	0.2685	0.0384	6.9855	0.0000	0.2531	0.0503	5.0289	0.0000
United States	UK	0.1753	0.0583	3.0045	0.0027	0.2103	0.0695	3.0241	0.0026
Japan	United States	0.7407	0.0742	9.9776	0.0000	–0.0718	0.0880	–0.8159	0.4148

NOTE: SE = standard error.

Table 4
Comparison of Full-Sample and No-Monday Results, U.S.-UK Regression

U.S.-UK Regression	Main Effect				Crisis Effect			
Sample	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Early cutoff								
Full sample	0.1742	0.0945	1.8427	0.0657	0.1682	0.1004	1.6756	0.0942
No Mondays	0.1732	0.1027	1.6864	0.0921	0.0635	0.1100	0.5772	0.5640
Late cutoff								
Full sample	0.1753	0.0583	3.0045	0.0027	0.2103	0.0695	3.0241	0.0026
No Mondays	0.1844	0.0642	2.8705	0.0042	0.0656	0.0784	0.8376	0.4025

Table 5
Three-Country Model Results

Markets			X_1 Main Effect				X_1 Crisis Effect				X_2 Main Effect				X_2 Crisis Effect			
Y	X_1	X_2	β	SE	t-Score	P-Value	β	SE	t-Score	P-Value	β	SE	t-Score	P-Value	β	SE	t-Score	P-Value
Early cutoff																		
U.S.	UK	Japan	0.1494	0.1016	1.4704	0.1418	0.1419	0.1108	1.2806	0.2007	0.0376	0.0590	0.6371	0.5242	–0.0004	0.0673	–0.0066	0.9948
Japan	U.S.	UK	0.5736	0.1495	3.8379	0.0001	0.0796	0.1578	0.5041	0.6143	0.1754	0.1146	1.5309	0.1262	–0.0858	0.1232	–0.6963	0.4864
UK	Japan	U.S.	0.0770	0.0470	1.6367	0.1021	0.2555	0.0524	4.8727	0.0000	0.7811	0.1308	5.9698	0.0000	–0.3400	0.1390	–2.4461	0.0146
Late cutoff																		
U.S.	UK	Japan	0.1391	0.0656	2.1215	0.0342	0.2121	0.0835	2.5401	0.0113	0.0656	0.0401	1.6372	0.1020	–0.0576	0.0564	–1.0208	0.3077
Japan	U.S.	UK	0.6949	0.0861	8.0673	0.0000	–0.0719	0.1038	–0.6922	0.4890	0.1098	0.0702	1.5630	0.1184	–0.0121	0.0878	–0.1378	0.8905
UK	Japan	U.S.	0.1494	0.0312	4.7848	0.0000	0.2358	0.0417	5.6606	0.0000	0.6616	0.0794	8.3348	0.0000	–0.2668	0.0953	–2.7988	0.0052

1. X_1 main effects: All are statistically significant with the later cutoff, but only the effect of the United States on Japan is significant with the earlier cutoff.
2. X_1 crisis effects: There is no significant change in the U.S. effect after the start of the crisis, with either early or later cutoff; the Japan effect on the UK increases after the start of the crisis with both early and later cutoffs; the UK effect on the United States increases after the onset of crisis if we use the later cutoff, but not with the early cutoff.
3. X_2 results: The relationship between the UK movements and prior U.S. movements is the only statistically significant main effect, regardless of which cutoff is used.
4. X_2 crisis effects: The crisis effect for the UK-U.S. X_2 relationship is statistically significant and *negative*, superficially suggesting that a movement in the U.S. markets is associated with a smaller movement in the following day's UK markets after the onset of the crisis (but see the Discussion section for caveats to this interpretation).
5. All other X_2 crisis effects are negative (although insignificant).

For comparison, Table 6 compares the immediately preceding market effects for both the two-country and three-country models. The two-country models presented in this table differ slightly from those previously presented in Table 3 because, for the Japan-UK relationship, we now restrict analysis to the morning market movements in the UK, in order to maintain consistency between the two-country model (which does not require this truncation) and the three-country model (which does).

Overall, the effects of the immediately preceding market appear to change little with the addition of the prior market into the model. In general, the coefficients remain very similar in size, but they do not always maintain the same level of statistical significance due to a loss of precision.[4] The relationship between movements in Japan and subsequent movements in the UK might decline slightly when the prior movements in the United States are included in the model, which might suggest that some of the original Japan-UK relationship is due to both markets reacting to market movements in the United States. However, the change is not large.

As with the correlation analysis and two-country regression analysis, the restriction of analysis to a data set excluding market movements on Mondays has little effect on the models with Japan or the UK as the following market, but it has some effect on the models that feature the United States following the UK: Whereas the full data set suggests a significant change in effect after the onset of the crisis, the no-Mondays data set shows no significant change in effect.[5]

Discussion

It would seem that the three major equity markets we studied were increasingly responding to the same news after the onset of the crisis. Taking the two-country correlation analyses and two- and three-country regression analyses together, relationships between movements in adjacent markets generally appeared to become stronger and more consistent after the onset

[4] This is consistent with expectation because we are adding independent regressors in the three-country model that are correlated with the regressors already present in the two-country model.

[5] Full comparison of results is provided in Appendix A.

Table 6
Comparison of Immediately Preceding Market Effects, Two-Country Versus Three-Country Models

Markets		2-Country Main Effect				2-Country Crisis Effect				3-Country X_1 Main Effect				3-Country X_1 Crisis Effect			
Foll.	**Prec.**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Early cutoff																	
U.S.	UK	0.1742	0.0945	1.8427	0.0657	0.1682	0.1004	1.6756	0.0942	0.1494	0.1016	1.4704	0.1418	0.1419	0.1108	1.2806	0.2007
Japan	U.S.	0.6492	0.1326	4.8975	0.0000	0.0577	0.1393	0.4143	0.6787	0.5736	0.1495	3.8379	0.0001	0.0796	0.1578	0.5041	0.6143
UK	Japan	0.1411	0.0490	2.8777	0.0041	0.2803	0.0540	5.1906	0.0000	0.0770	0.0470	1.6367	0.1021	0.2555	0.0524	4.8727	0.0000
Late cutoff																	
U.S.	UK	0.1753	0.0583	3.0045	0.0027	0.2103	0.0695	3.0241	0.0026	0.1391	0.0656	2.1215	0.0342	0.2121	0.0835	2.5401	0.0113
Japan	U.S.	0.7490	0.0752	9.9646	0.0000	−0.0672	0.0894	−0.7518	0.4524	0.6949	0.0861	8.0673	0.0000	−0.0719	0.1038	−0.6922	0.4890
UK	Japan	0.2325	0.0317	7.3324	0.0000	0.2396	0.0415	5.7721	0.0000	0.1494	0.0312	4.7848	0.0000	0.2358	0.0417	5.6606	0.0000

of crisis; it appears that national financial markets became more-strongly linked to each other after the crisis.[6]

The correlations seen in the R^2 values increased in all cases, and increased statistically significantly in all but one case, indicating that more variation in the following market was explained by variation in the preceding market after the onset of crisis. This finding is consistent with the results of earlier correlation studies (see papers cited above), which generally show increased correlations among the returns on various financial instruments and financial markets during times of crisis.

More novel is our finding that, in most cases, the size of the reactions—the "betas"—of markets to earlier movements in other markets also increased statistically significantly after the start of the crisis, with the exception being in the relationship between the United States and Japan—the typical movement in Japanese markets seen after a given movement in U.S. markets was quite high before the crisis, and remained high (without increasing or decreasing in a statistically significant way) after the crisis began. In both the two-country main effects and the three-country X_1 effects, the impact of U.S. market movements appeared to be larger than the impact of other market movements, both before and after the onset of crisis.

The United States also appeared to be the only market that exerted an independent influence on a market other than that which immediately follows it chronologically, as the movements in the UK appeared to be affected by U.S. market movements even after controlling for the intervening influence of Japanese market movements. The independent effect of the United States on the UK appeared to decrease after the onset of crisis, but this should be interpreted with caution: This could indicate that UK market movements were less related to U.S. market movements after the onset of crisis, but it might also indicate that the more-consistent relationship between market movements in the United States and Japan leaves less variation in U.S. markets that differs from Japanese markets, and thus makes it harder for the United States to have an effect on the UK after controlling for the influence of Japan. An examination of the effect of U.S. market movements on UK market movements without controlling for the intervening Japanese market movements gives some support to the latter explanation.[7]

Interpretation of the relationship between market movements in the UK and the United States is challenging for several reasons. Due to the overlapping trading hours, we focus only on the morning trading in the UK, which reduces the likelihood of seeing significant effects and makes direct comparison with the other market relationships difficult. Although we do see some interesting results in the full data-set analyses—and these are consistent with the general patterns described above—the results do not hold in the more-restricted data set that excludes Monday market movements, and so must be interpreted with caution. A detailed analysis of specific market-moving news on a date-by-date basis could explore the reasons for the unstable results, but that is beyond the scope of our analysis.

One further caveat is necessary: The increased market volatility immediately following the onset of the crisis presents some challenges to our methodological approach. Forbes and Rigo-

---

[6] In particular, it appears that movements in preceding world markets became a more-important source of news for each market than each country's own idiosyncratic news.

[7] With both early and late cutoffs, the relationship between U.S. market movements and UK market movements showed very small (and statistically insignificant) positive crisis effects and increases in the R^2 value, if we ignore the intervening Japanese market movements. This does not provide strong evidence for an *increase* in influence of the United States on the UK, but it does undermine the case for a *decrease* in influence.

bon have illustrated that a move from low volatility to high volatility can produce an apparent increase in correlation even if the fundamental linkages between markets are unchanged (Forbes and Rigobon, 2002). Intuitively, the point is that, if some fixed portion of market movements in one country tends to be transmitted to the following country, then when movements in the first country increase, the signal of dependent movements in the following country might stand out more above the noise of idiosyncratic movements in the following country (if idiosyncratic movements in the following country do not increase commensurately). Forbes and Rigobon suggest imposing additional restrictions on the model in order to correct for this potential bias, but these restrictions seem neither plausible nor usefully applicable to our modeling.[8] We draw some comfort from the consistency of results between our correlation analysis and regression analysis because the potential effects of bias differ between the models. Nevertheless, it is necessary to see our two approaches as complementary methods whose results are better considered in combination than relied upon in isolation.

[8] They assume no exogenous shocks to the system and no feedback from one country to the other; they also require one country to be designated as the source of the increased volatility, with bivariate relationships between that country and other countries the only valid comparisons; their correction also implicitly assumes that the beta-coefficient remains constant between the before and after period, which is not easily reconciled with our findings.

SECTION 5

Extensions

Financial Subindices

The global financial crisis had its origins in the U.S. subprime housing market but spread to the U.S. financial sector and the financial sectors of other countries through the proliferation of mortgage-backed securities and other financial instruments that were sensitive to the U.S. housing market. Although the movements in markets over the period we study reflect changes in stock values across a variety of industries in each country, we wished to explore whether focusing on the share prices specifically of financial institutions might provide additional insight into the linkages between markets in different countries.

Unfortunately, it was not possible to acquire any intraday data on indices reflecting the market performance of financial institutions, so it was not possible to conduct a full analysis of relationships between markets in different countries (due to the overlapping opening hours in London and New York). However, we were able to acquire closing-price data for the U.S. KBW Large Cap Banking Index, the UK FTSE 350-Banks Index, and the Japanese TOPIX-17 Banking Index, with which we were able to analyze the relationship between the United States and Japan and between Japan and the UK.[1] Although we could not perform any three-country analyses, we did examine the relationship between Japanese markets and later movements the same day in the United States and the relationship between UK markets and movements the next day in Japan, to see whether this provided any additional insight.

Financial Subindices: Correlation Results

Using the early cutoffs, the results we see in the consecutive-market correlation analysis of financial subindices are broadly consistent with the results we saw for the whole-market indices, but seem a little weaker (see Table 7). The correlation between consecutive subindices increases after the onset of crisis, but the eventual R^2 values are not as high as those seen in the whole-market results in Table 1. The late-cutoff results are a little different: The UK-Japan relationship shows an increase after the onset of crisis, but, unlike in the whole-market results, the change is not statistically significant; the Japan-U.S. late-cutoff result for the financial subindices shows a small and statistically insignificant *decrease* in correlation, compared with a statistically insignificant increase in correlation in the whole-market results.

[1] These indices include only banks headquartered in the country (i.e., they do not include shares of major foreign banks that are cross-listed) and do not include banks whose primary business is investment banking. Some large financial conglomerates with investment banking arms, such as Mitsubishi UJF Financial Group in Japan and Bank of America in the United States, are included in these indices, but major investment banks, such as Nomura Holdings in Japan and Goldman Sachs in the United States, are excluded.

Table 7
Financial Subindices, R-Squared Comparisons

Markets		R^2		Difference	
Preceding	**Following**	**Before**	**After**	**Z-Score**	**P-Value**
Consecutive, early cutoff					
Japan	UK	0.0371	0.2244	4.4653	0.0000
United States	Japan	0.0247	0.1783	4.1345	0.0000
Consecutive, late cutoff					
Japan	UK	0.1506	0.2443	1.4937	0.1353
United States	Japan	0.1659	0.1587	–0.1352	0.8925
Once removed, early cutoff					
UK	Japan	0.0170	0.1252	3.2314	0.0012
Japan	United States	0.0045	0.0089	0.3084	0.7577
Once removed, late cutoff					
UK	Japan	0.1046	0.1221	0.3096	0.7569
Japan	United States	0.0142	0.0059	–0.5791	0.5625

The correlation analysis of once-removed markets does not produce many results of interest. The only significant change appears in the relationship between market movements in the UK and market movements the next day in Japan (with the early cutoff only); however, without being able to deal with the overlapping data between the UK and United States, it is possible (though unprovable with available data) that this result reflects UK markets reacting to U.S. markets during the trading overlap, with Japanese markets then reacting to U.S. markets later and having a spurious correlation with UK markets.

Financial Subindices: Regression Results

The regression analysis of the financial subindices provides some potentially interesting results (Table 8).

In these results, there is a significant increase in the size of a movement in the UK financial subindex for a given Japanese financial subindex movement, with both the early and late cutoffs. More surprisingly, given the other results we have seen, the Japan-U.S. results show basically no change before and after the onset of crisis when using the early cutoff, but they show a large and statistically significant decrease in effect size after the onset of crisis when using the late cutoff: A 1-percent change in the U.S. financial subindex is associated with a 0.4465-percent increase in the Japanese financial subindex in the before period, decreasing by 0.1946 percent to a total after-crisis effect size of 0.2519 percent.

Again, the once-removed results are difficult to interpret—the only significant crisis effect is in the Japan-UK late-cutoff result. Again, the similarity of the change here to the change in the Japan-U.S. late-cutoff result suggests that the Japan-UK result might be driven by the UK reacting at the end of its trading day to early movements in the United States, and the Japa-

Table 8
Financial Subindices, Regression Results

Markets		Main Effect				Crisis Effect			
Foll.	**Prec.**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Consecutive, early cutoff									
UK	Japan	0.1137	0.0854	1.3315	0.1834	0.4273	0.0929	4.5986	0.0000
Japan	U.S.	0.3059	0.1473	2.0767	0.0381	–0.0046	0.1495	–0.0306	0.9756
Consecutive, late cutoff									
UK	Japan	0.2861	0.0483	5.9210	0.0000	0.3633	0.0672	5.4092	0.0000
Japan	U.S.	0.4465	0.0492	9.0703	0.0000	–0.1946	0.0570	–3.4132	0.0007
Once removed, early cutoff									
Japan	UK	0.2203	0.1483	1.4858	0.1378	0.0901	0.1519	0.5934	0.5531
U.S.	Japan	0.0350	0.1017	0.3441	0.7308	0.0952	0.1125	0.8459	0.3978
Once removed, late cutoff									
Japan	UK	0.4338	0.0667	6.5065	0.0000	–0.1657	0.0760	–2.1810	0.0295
U.S.	Japan	0.1045	0.0610	1.7139	0.0869	0.0169	0.0871	0.1947	0.8457

nese market subsequently reacting to the United States. However, it should be noted that there is no way to distinguish this definitively from the alternative hypothesis that the U.S. results are being driven by movements in the UK, and that the Japan-U.S. results are the results that are questionable. Given the three-country analyses we were able to perform with our more-detailed whole-market data set, this latter option appears unlikely, but we cannot fully rule it out using our limited financial subindex data.

Comparison with 1997 Asian Financial Crisis

The 1997 Asian financial crisis is a useful point of reference to give context to our other results: Are the patterns we see in the global financial crisis common to other periods of turmoil, or do patterns change differently based on the nature of the crisis? The 1997 Asian crisis allows us to look at how one crisis originating outside the United States compares with the more-recent crisis originating in U.S. subprime mortgages and associated financial instruments. Importantly, the 1997 Asian crisis led to a period of fluctuation in global markets over a number of weeks, with a large number of dates contributing to our statistical estimates—we conducted a similar analysis of the late 1980s, but the statistical estimates were dominated by a single data point ("Black Monday," October 19, 1987) and are thus difficult to generalize. The late-1980s analysis is included in Appendix B for completeness.

As with the financial subindices, it was not possible to acquire any intraday data for the period of the 1997 crisis with which we could conduct our full analysis of relationships between markets in different countries. Therefore, we are restricted to analyzing the relationships between the United States and Japan, between Japan and the UK, and the once-removed U.S.-Japan and Japan-UK markets. We also select a single date for the start of the crisis, July 1,

1997, the day before the devaluation of the Thai baht, which triggered a run on Southeast Asian currencies and capital flight from troubled Southeast Asian economies.

Previous research on cross-market effects of the Asian crisis has focused on how shocks were transmitted between the emerging economies at the heart of the crisis, and (to a lesser extent) between the East Asian countries and Latin American emerging economies, rather than on the effects on relationships between major developed economies. The most-cited work focused on the former question found significant increases in cross-border links in both the equity and currency markets among the East Asian countries (Baig and Goldfajn, 1998).

Asian Financial Crisis: Correlation Results

The Asian financial-crisis correlation results have some similarities and differences with the global financial-crisis results. The consecutive market correlations increase in a statistically significant way in both cases. However, the starting R^2 values, and the increases in R^2 values, are somewhat lower for the Asian financial crisis (see Table 9).

The once-removed market analyses suggest very little correlation between nonconsecutive markets; this might have increased slightly after the onset of crisis, but the results are not statistically significant in the U.S.-Japan relationship and are marginally significant in the Japan-UK relationship.

Asian Financial Crisis: Regression Results

The Asian financial-crisis consecutive market regression results also have similarities and differences with global financial-crisis regression results, with the same basic patterns but lower starting values (see Table 10). In both cases, the Japan-U.S. main effect is much larger than the UK-Japan effect before the start of the crisis, but it has an increase after the crisis that is not statistically significant; in both cases, the UK-Japan main effect is small, but the larger crisis effect shows a statistically significant increase in the effect size after the onset of the crisis.

The once-removed market results show small effects—the main effects are negligible, and the crisis effects show increases that are either not statistically significant (Japan-UK) or very marginally significant (U.S.-Japan).

Table 9
Asian Crisis Correlation Results

Markets		R^2		Difference	
Following	Preceding	Before	After	Z-Score	P-Value
Consecutive					
UK	Japan	0.0040	0.0667	3.0052	0.0027
Japan	United States	0.0381	0.1100	2.2274	0.0259
Once Removed					
Japan	UK	0.0002	0.0118	1.7672	0.0772
United States	Japan	0.0008	0.0148	1.5263	0.1269

Table 10
Asian Crisis Regression Results

Markets		Main Effect				Crisis Effect			
Foll.	**Prec.**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Consecutive									
UK	Japan	0.0298	0.0285	1.0445	0.2966	0.1326	0.0380	3.4869	0.0005
Japan	U.S.	0.4000	0.1020	3.9205	0.0001	0.0878	0.1209	0.7266	0.4677
Once Removed									
Japan	UK	−0.0281	0.1115	−0.2520	0.8011	0.2023	0.1339	1.5104	0.1313
U.S.	Japan	0.0137	0.0303	0.4513	0.6519	0.0680	0.0406	1.6752	0.0943

SECTION 6

Conclusions

Taking our main analyses and extensions together, the effect of the global financial crisis on the interconnectedness of global markets is complicated. We can draw several key points from our analyses:

1. Movements in the benchmark stock indices for New York, London, and Tokyo are more-closely linked after the onset of the global financial crisis than they were before. The beta effects and correlations between markets and the immediately following market generally increase. The statistically independent effects of markets on non–immediately following markets might have decreased, but this is an unsurprising statistical finding given stronger links between adjacent markets.
2. The links between the benchmark U.S. market movements and movements in the following markets do not increase as much as other market links, but they start from a much stronger position. In our analysis, U.S. influence on other markets is the strongest influence both before and after the onset of crisis and does not appear to have slipped in absolute terms following a crisis that began in the U.S. housing and financial sectors.
3. The pattern seen with the global financial crisis is consistent with historic precedent: When the Asian financial crisis caused significant turbulence in world financial markets, the interconnectedness between markets apparently increased in a similar fashion. Although the two crises originated in different regions, our analysis does not show much difference in the effects of the crises.
4. While our main focus was the change in market relationships between 2005 and 2009, world markets appear to have been more interconnected in the years immediately prior to the global financial crisis than they were in the years immediately prior to the Asian financial crisis. Although we saw a significant change with the onset of the global financial crisis, there might be a more-general upward trend in interconnectedness of the benchmark indices, perhaps due to a broader process of globalization.
5. Although global factors affecting whole markets might have led to broad market indices moving together, the banking subindices tell a slightly different story, with the share value of Japanese banks seemingly less connected to the share value of U.S. banks after the onset of crisis. This might indicate that competitive shocks in the global banking industry might somewhat offset the effect of global shocks: What's good for the U.S. banking industry might be good for the United States, and what's good for the United States might be good for the global economy—but some bad news for the U.S. banking industry might be good news for rival financial companies based in Tokyo, London, and elsewhere.

Statistical analyses, such as those reported in this paper, reveal relationships among markets and how these relationships have changed. By their nature, these analyses are silent on *why* these relationships might have changed. It might be interesting to speculate, however, about the reasons for these changes.

One of the defining characteristics of the recent crisis was the difficulty that market participants and regulators initially encountered in identifying and quantifying the exposure of particular financial institutions and investors to deteriorating conditions in the U.S. housing sector. Financial innovations in the preceding years had allowed risks associated with mortgage lending to be distributed widely and in different patterns from any that markets had previously seen. Moreover, as the crisis unfolded, it became apparent that rating agencies, regulators, and many market participants had misjudged the risks associated with some of the new financial instruments. As the realization spread that risks were greater, more-widely distributed, and more difficult to quantify than had been previously understood, all market participants came under suspicion of possibly being exposed. With many financial institutions across many countries potentially exposed, lending everywhere slowed as financial intermediaries sought to conserve cash. This, of course, damaged the prospects of nonfinancial companies in need of credit. Since ability to distinguish relative levels of exposure was limited, share prices in multiple markets tended to react to the same news—either positive or negative—about the extent of credit problems and government efforts to deal with them.

Under the circumstances prevailing at the time, it is perhaps not surprising that benchmark equity indexes became more-closely linked. A more-interesting question might be whether this new correlation of markets will persist. Has recent experience demonstrated that risks and rewards are more-broadly dispersed than in the past, and will this broad dispersion be a lasting feature of the global financial system? If so, we might expect different national markets to remain more-closely linked. And as economic phenomena—both real and financial—become increasingly global and the assets traded in different markets become more similar or more susceptible to similar influences, we should expect to see global linkages. Our analysis suggests that, even before the onset of the recent crisis, equity markets were more-closely linked than at the start of the Asian crisis.

Alternatively, the regulatory reforms still being formulated might eventually make it easier to localize risks and rewards to particular institutions and companies, allowing prices of particular equities and national market indices to move more independently. Perhaps the simple passage of time and the easing of concerns about systemic weakness will allow market participants to focus more clearly on factors that distinguish assets traded in different markets.

Finally, it is possible that correlation of markets will decrease once this crisis passes, only to increase again at the next crisis. Although the circumstances surrounding each crisis were quite different, both the recent global financial crisis and the earlier Asian financial crisis seem to have been followed by a period of increased correlation accompanying the increased volatility. One consequence of closer linkages between markets is that the benefits of diversifying an investment portfolio by buying broad indices of domestic and foreign stocks might be reduced in times of crisis—just when investors will most need the stability of diversification.

It is simply too early to resolve these issues. We hope to return to them in future analyses.

An Agenda for Future Analyses

The analysis presented here is of a distinctly preliminary nature, limited by the data and resources available. A number of questions—some of possible importance to policy—remain.

How, for example, does the growth of markets for derivatives related to equities trading on the traditional exchanges influence the linkages among the major exchanges? Does the fact that, say, S&P futures are trading at times when the Japanese markets are open serve to connect Japanese and New York markets to a degree that we have not seen before? Performing analyses similar to those presented in this paper with the inclusion of futures markets might shed light on such phenomena.

It would also be interesting to undertake a similar analysis of the linkages among interbank lending rates to major banks in different countries. Because banks in different jurisdictions will operate under different (although increasingly similar) regulatory regimes and have access to different backstop credit arrangements, we might expect national differences in interbank rates. Did these rates become more-closely linked during the crisis?

Extending the analysis to other markets beyond the three major markets investigated here might reveal different patterns of linkages. Equity markets in emerging economies might be the most-interesting places to look next.

These additional analyses might raise challenging data issues. But some other potentially useful analyses can be pursued with data already in hand.

Finer divisions of the period after the onset of the crisis, for example, might allow some assessment of policies put in place at different times. Did limits on short selling, for example, have any effect on linkages among markets?

Further investigation of linkages on days when markets saw particularly large movements might also prove interesting. We do not fully understand the extent to which the results reported here reflect routine interaction of markets, as opposed to the consequences of a handful of very large moves.

APPENDIX A

No-Mondays Analysis

No Mondays: Correlation Analysis

Models are as run in the main body of the paper. Sample NM does not include changes between market close Friday and market close Monday.

No Mondays: Regression Analysis

Models are as run in the main body of the paper. Sample NM does not include changes between close Friday and close Monday.

Table 11
No-Mondays Correlation Analysis

Markets			R^2		Difference	
Following	Preceding	Sample	Before	After	Z-Score	P-Value
Early cutoff						
UK	Japan	Full	0.0470	0.2999	5.6387	0.0000
		NM	0.0432	0.2665	4.7086	0.0000
United States	UK	Full	0.0366	0.1114	2.3274	0.0199
		NM	0.0359	0.0514	0.5575	0.5772
Japan	United States	Full	0.1270	0.2971	3.3821	0.0007
		NM	0.1313	0.2927	2.9259	0.0034
Late cutoff						
UK	Japan	Full	0.1333	0.3350	3.0930	0.0020
		NM	0.1070	0.3083	2.9281	0.0034
United States	UK	Full	0.0319	0.1410	2.5752	0.0100
		NM	0.0335	0.0580	0.6626	0.5076
Japan	United States	Full	0.2144	0.3058	1.3103	0.1901
		NM	0.1870	0.3162	1.7026	0.0886

Table 12
No-Mondays Regression Analysis

Markets			Main Effect				Crisis Effect			
Foll.	**Prec.**	**Sample**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Early cutoff										
UK	Japan	Full	0.1411	0.0490	2.8777	0.0041	0.2803	0.0540	5.1906	0.0000
		NM	0.1317	0.0551	2.3905	0.0171	0.2473	0.0604	4.0937	0.0000
U.S.	UK	Full	0.1742	0.0945	1.8427	0.0657	0.1682	0.1004	1.6756	0.0942
		NM	0.1732	0.1027	1.6864	0.0921	0.0635	0.1100	0.5772	0.5640
Japan	U.S.	Full	0.6492	0.1326	4.8975	0.0000	0.0577	0.1393	0.4143	0.6787
		NM	0.6341	0.1436	4.4161	0.0000	0.0435	0.1506	0.2893	0.7725
Late cutoff										
UK	Japan	Full	0.2325	0.0317	7.3324	0.0000	0.2396	0.0415	5.7721	0.0000
		NM	0.2116	0.0359	5.8906	0.0000	0.2074	0.0461	4.4949	0.0000
U.S.	UK	Full	0.1753	0.0583	3.0045	0.0027	0.2103	0.0695	3.0241	0.0026
		NM	0.1844	0.0642	2.8705	0.0042	0.0656	0.0784	0.8376	0.4025
Japan	U.S.	Full	0.7490	0.0752	9.9646	0.0000	−0.0672	0.0894	−0.7518	0.4524
		NM	0.6817	0.0833	8.1840	0.0000	−0.0109	0.0974	−0.1116	0.9112

APPENDIX B

Black Monday

A separate analysis was run on the period surrounding Black Monday (October 19, 1987). However, the market crash on Black Monday itself dwarfs all other stock-market moves in the previous and following two years, making the results of the analysis highly sensitive to that one data point—for example, the increase in R^2 value for the U.S.-Japan analysis disappears almost entirely when Black Monday itself is excluded. The results with and without that outlying data point are below.

Table 13
Black Monday Correlation Analysis, Full-Sample Results

Markets		R^2		Difference	
Following	**Preceding**	**Before**	**After**	**Z-Score**	**P-Value**
Consecutive					
UK	Japan	0.0012	0.0826	4.9040	0.0000
Japan	United States	0.0545	0.2829	5.2136	0.0000
Once removed					
Japan	UK	0.0074	0.0888	3.2345	0.0012
United States	Japan	0.0065	0.0203	0.9942	0.3201

Table 14
Black Monday Correlation Analysis Excluding Black Monday

Markets		R^2		Difference	
Following	**Preceding**	**Before**	**After**	**Z-Score**	**P-Value**
Consecutive					
UK	Japan	0.0012	0.0001	0.6875	0.4918
Japan	United States	0.0545	0.0856	0.9501	0.3421
Once removed					
Japan	UK	0.0074	0.0510	2.1170	0.0343
United States	Japan	0.0065	0.0663	2.7704	0.0056

Table 15
Black Monday Regression Analysis, Full-Sample Results

Markets		Main Effect				Crisis Effect			
Foll.	**Prec.**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Consecutive									
UK	Japan	−0.0335	0.0489	−0.6846	0.4938	0.3114	0.0653	4.7671	0.0000
Japan	U.S.	0.3007	0.0609	4.9342	0.0000	0.1484	0.0700	2.1204	0.0343
Once removed									
Japan	UK	0.0898	0.0537	1.6714	0.0950	0.2655	0.0757	3.5054	0.0005
U.S.	Japan	0.0628	0.0529	1.1873	0.2354	0.1043	0.0696	1.4984	0.1344

Table 16
Black Monday Regression Analysis, Excluding Black Monday Itself

Markets		Main Effect				Crisis Effect			
Foll.	**Prec.**	**β**	**SE**	**t-Score**	**P-Value**	**β**	**SE**	**t-Score**	**P-Value**
Consecutive									
UK	Japan	−0.0335	0.0459	−0.7293	0.4660	0.0418	0.0662	0.6307	0.5284
Japan	U.S.	0.3007	0.0600	5.0088	0.0000	−0.0131	0.0757	−0.1731	0.8626
Once removed									
Japan	UK	0.0898	0.0537	1.6714	0.0950	0.1452	0.0714	2.0300	0.0730
U.S.	Japan	0.0628	0.0516	1.2188	0.2233	0.2831	0.0726	3.9018	0.0001

APPENDIX C

Model Fitting

The analyses presented in the main body of this paper suggest that the relationships among major equity markets changed significantly after the onset of the recent financial crisis. Of course, there is no single date that can clearly be identified as the *day* when the crisis began. Rather a succession of events spread over a period of more than a year served cumulatively to define the crisis. The analyses in the main body of the paper propose two plausible—but nonetheless somewhat arbitrary—dates for the beginning of the crisis: August 1, 2007, and September 1, 2008.

In the appendix, we a test a range of dates for the division between the precrisis and post-crisis periods, looking for the dates that yield the cleanest break in the relationships among markets.

Operationally, we show how the statistical fit of the model varies with different start dates for the onset of the crisis. For each country relationship, we divide the data into two pieces using a particular cut date, run the regression model described in the main analysis sections of this paper, and record the fraction of variance explained by the model (the R^2 value). We then iterate this process using every date from the start of our data set to the end as the cut date, and plot the resulting R^2 values. In the presentation below, we first show the results for the two-country models, and then the results for the three-country models. To maintain comparability between the two-country and three-country models, we use only the morning market movements in the UK when modeling the two-country relationship between the UK and Japan.

Two-Country Models

Interestingly, the patterns displayed across the two-country models differ from country to country: No single date maximizes the model fit across all three countries. All three models have a peak in the fall of 2008, but the fall peak date is quite different when we look at the response of Japanese markets to U.S. markets (Figure 8) compared with the UK response to Japan (Figure 6) and U.S. response to UK (Figure 7). The peak dates for the latter models are from the third week of September 2008, the week in which Lehman Brothers filed for bankruptcy protection, the Federal Reserve lent $85 billion to AIG to avert bankruptcy, and the $700 billion bailout package was proposed by Treasury Secretary Henry Paulson. The peak date in the Japan-U.S. model is from the fourth week in October 2008, a week in which Wachovia announced a quarterly loss of $24 billion (the largest of any bank) and the UK government confirmed that the UK economy had its biggest drop in GDP since 1990; the end of the week saw stock markets plummeting in Asia, Europe, and the United States.

Figure 6
Variance Explained by Model with Varying Cut Date, UK-Japan Two-Country Model

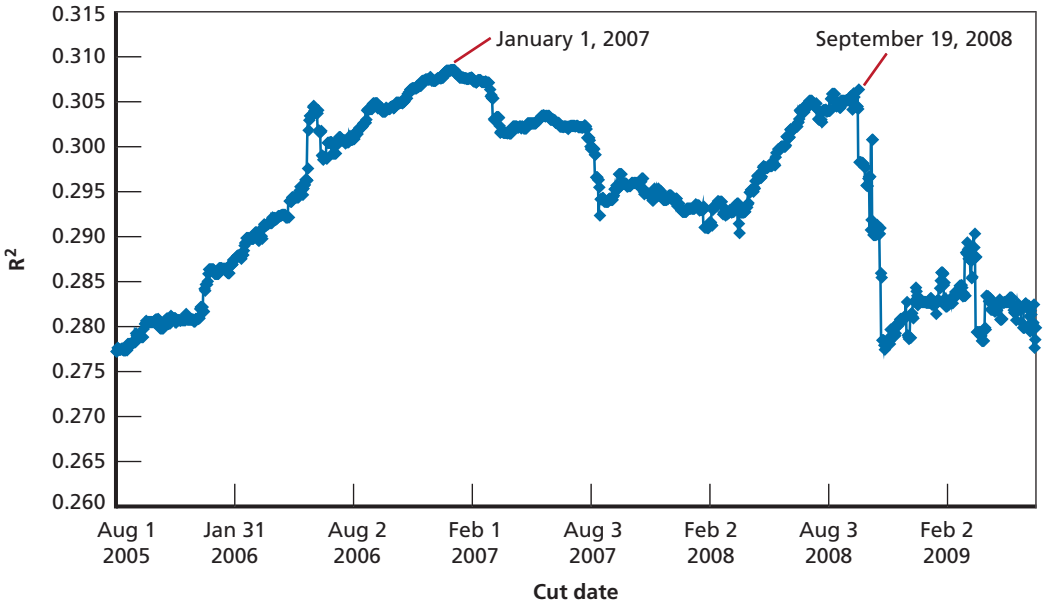

RAND *OP317-6*

Figure 7
Variance Explained by Model with Varying Cut Date, U.S.-UK Two-Country Model

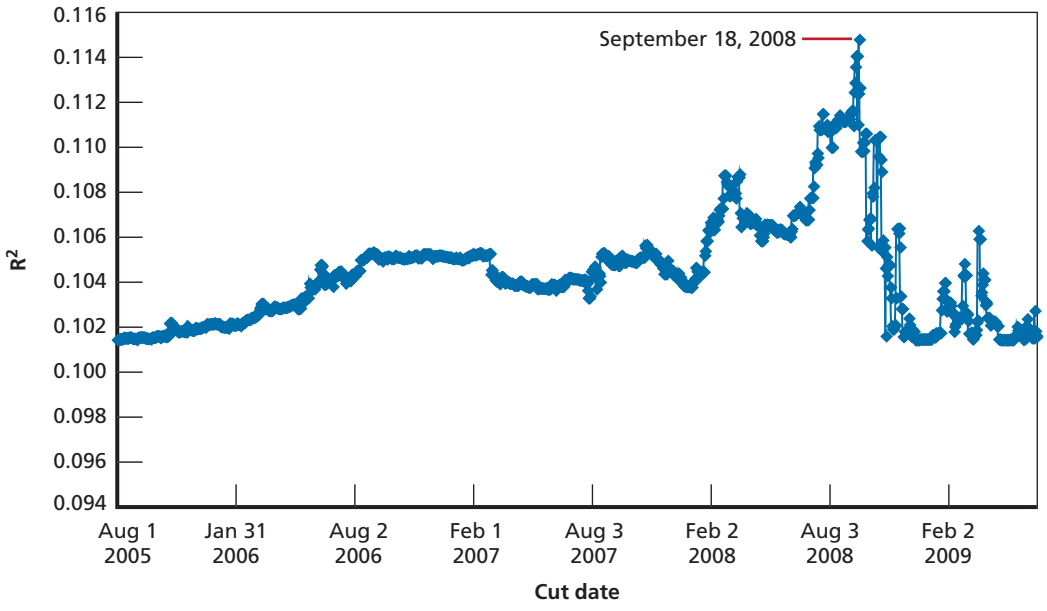

RAND *OP317-7*

In addition to the fall 2008 peak, the model of the UK regression on Japan has a peak at the turn of the new year in 2007. This peak does not appear to be related to any of the large and notable events in the financial crisis, and is generally not so pronounced as the fall 2008 peaks in all three models.

Figure 8
Variance Explained by Model with Varying Cut Date, Japan-U.S. Two-Country Model

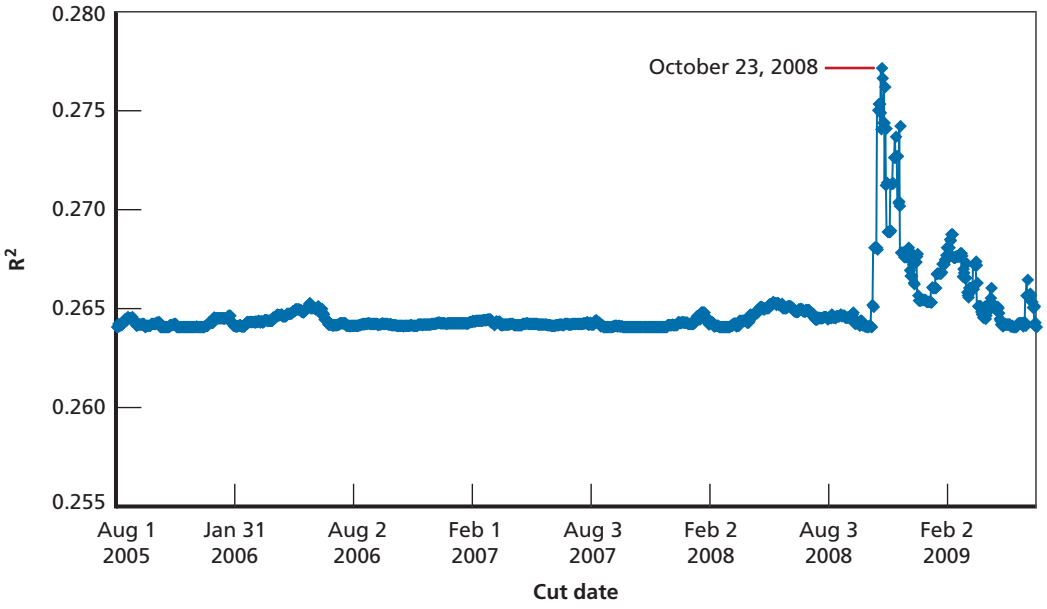

RAND *OP317-8*

Three-Country Models

The three-country model results are generally similar to the two-country results: All three models show a peak in fall 2008, with the UK model (Figure 9) having an additional peak at

Figure 9
Variance Explained by Model with Varying Cut Date, UK-Japan-U.S. Three-Country Model

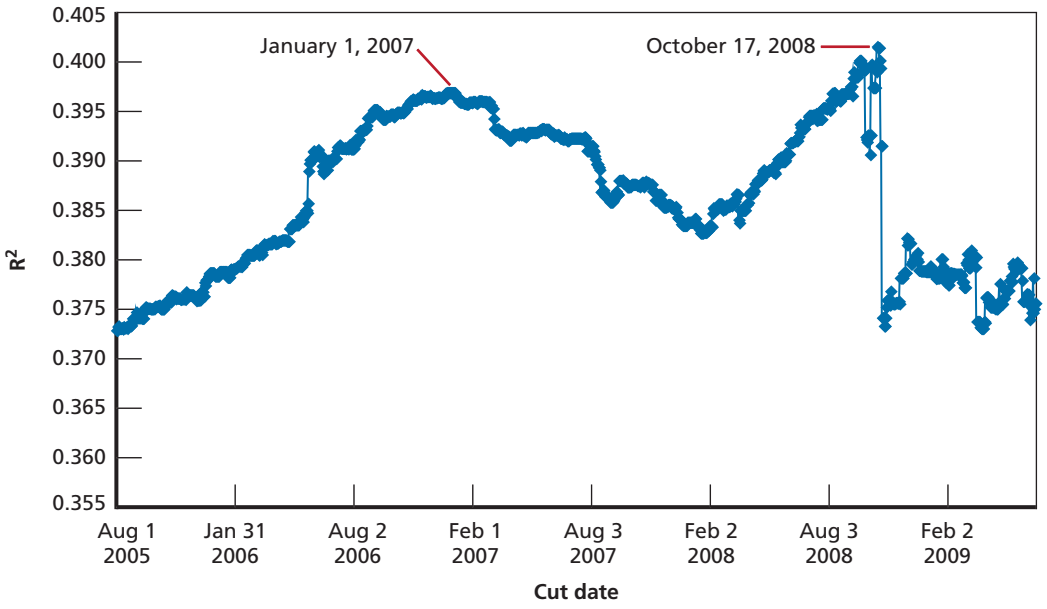

RAND *OP317-9*

the turn of the year 2007. The peaks in the United States (Figure 10) and Japan (Figure 11) models are in basically the same place as in the two-country models, in late September 2008 and late October 2008, respectively.

Figure 10
Variance Explained by Model with Varying Cut Date, U.S.-UK-Japan Three-Country Model

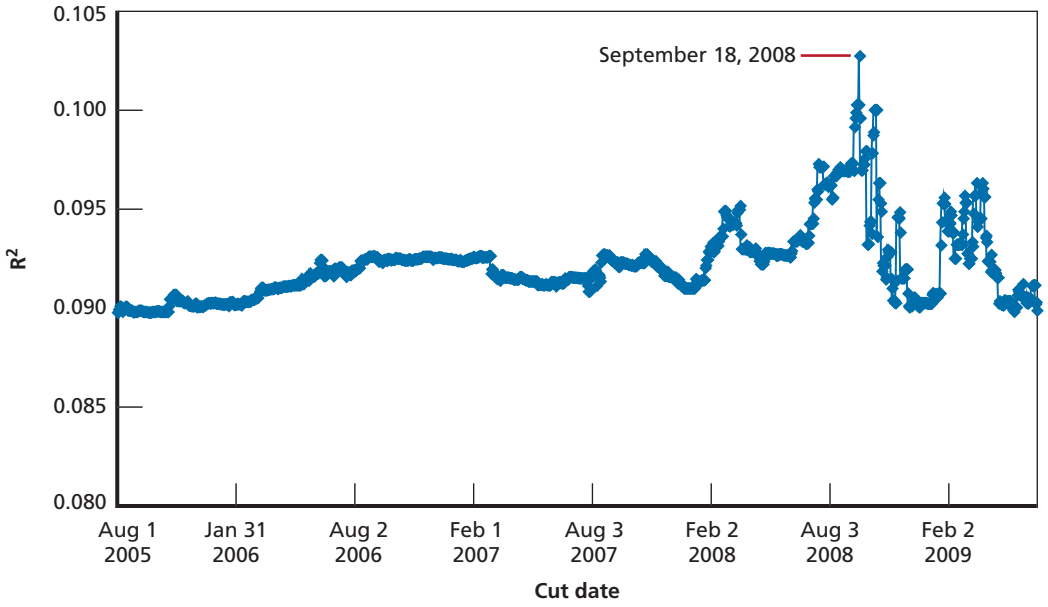

RAND *OP317-10*

Figure 11
Variance Explained by Model with Varying Cut Date, Japan-U.S.-UK Three-Country Model

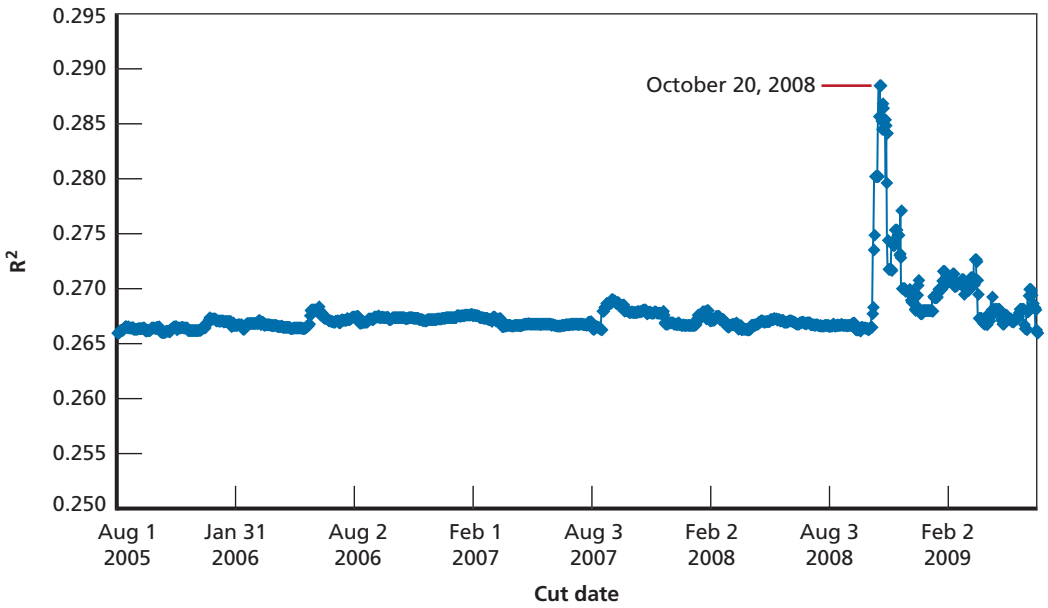

RAND *OP317-11*

The most-notable difference between the two- and three-country models is in the timing of the fall 2008 peak in the UK model. In the two-country model, this peak is in late September; in the three-country model, the peak is in late October. It appears that adding the U.S. market movements as an additional explanatory variable in this model affects the timing of the model fit peak, while adding the UK market movements to the Japan model, or the Japanese market movements to the U.S. model, has no effect. This provides additional evidence that U.S. market movements generally have more global influence (or reflect more-important global news) than market movements in other countries.

References

Baig, Taimur, and Ilan Goldfajn, *Financial Market Contagion in the Asian Crisis*, Washington, D.C.: International Monetary Fund, Asia and Pacific Department, Working Paper 98/155, November 1998. As of January 28, 2011:
http://bibpurl.oclc.org/web/24285/wp98155.pdf

Board of Governors of the Federal Reserve System, "Foreign Exchange Rates," last updated January 18, 2011. As of January 24, 2011:
http://www.federalreserve.gov/RELEASES/H10/Hist/

Dungey, Mardi, Renée Fry, Brenda González-Hermosillo, and Vance L. Martin, "Empirical Modelling of Contagion: A Review of Methodologies," *Quantitative Finance*, Vol. 5, No. 1, February 2005, pp. 9–24.

Fisher, R. A., "On the Probable Error of a Coefficient of Correlation Deduced from a Small Sample," *Metron*, Vol. 1, 1921, pp. 3–32.

Forbes, Kristin J., and Roberto Rigobon, "No Contagion, Only Interdependence: Measuring Stock Market Co-Movements," *Journal of Finance*, Vol. 57, No. 5, October 2002, pp. 2223–2261.

International Swaps and Derivatives Association, "Lehman Risk Reduction Trading Session and Protocol Agreement," news release, September 14, 2008. As of January 28, 2011:
http://www.isda.org/press/press091408lehman.html

Karolyi, G. Andrew, "Does International Financial Contagion Really Exist?" *International Finance*, Vol. 6, No. 2, July 2003, pp. 179–199.

King, M. A., and S. Wadhwani, "Transmission of Volatility Between Stock Markets," *Review of Financial Studies*, Vol. 3, No. 1, 1990, pp. 5–33.

Tick Data, undated home page. As of January 24, 2011:
http://www.tickdata.com/

Tokyo Stock Exchange, undated home page. As of January 24, 2011:
http://www.tse.or.jp/english/

Yahoo! Finance, undated home page. As of January 24, 2011:
http://finance.yahoo.com

Yahoo! Finance UK and Ireland, undated home page. As of January 24, 2011:
http://uk.finance.yahoo.com/