

The Journal of Portfolio Management

VOLUME 37 NUMBER 4

JPM.pm-research.com

SUMMER 2011

Principal Components as a Measure of Systemic Risk

MARK KRITZMAN, YUANZHEN LI,
SÉBASTIEN PAGE, AND ROBERTO RIGOBON

Principal Components as a Measure of Systemic Risk

MARK KRITZMAN, YUANZHEN LI, SÉBASTIEN PAGE,
AND ROBERTO RIGOBON

MARK KRITZMAN is the president and CEO of Windham Capital Management, LLC, and a senior lecturer at the Sloan School of Management, MIT, in Boston, MA. mkritzman@windhamcapital.com

YUANZHEN LI is a research associate at Windham Capital Management, LLC, in Boston, MA. yzli@windhamcapital.com

SÉBASTIEN PAGE is an executive vice president at PIMCO in Newport Beach, CA. Sebastien.Page@pimco.com

ROBERTO RIGOBON is a professor at the Sloan School of Management, MIT, in Boston, MA, and a research associate of the National Bureau of Economic Research in Cambridge, MA. rigobon@mit.edu

The U.S. government's failure to provide adequate oversight and prudent regulation of the financial markets, together with excessive risk taking by some financial institutions, pushed the world financial system to the brink of systemic failure in 2008. As a consequence of this near catastrophe, both regulators and investors have become keenly interested in developing tools for monitoring systemic risk. But this is easier said than done. Securitization, private transacting, complexity, and "flexible" accounting¹ prevent us from directly observing the many explicit linkages of financial institutions. As an alternative, we introduce a measure of implied systemic risk called the absorption ratio, which equals the fraction of the total variance of a set of asset returns explained or "absorbed" by a fixed number of eigenvectors.²

The absorption ratio captures the extent to which markets are unified or tightly coupled. When markets are tightly coupled, they are more fragile in the sense that negative shocks propagate more quickly and broadly than when markets are loosely linked.

We offer persuasive evidence that the absorption ratio effectively captures market fragility. We show that

1. Most significant U.S. stock market drawdowns were preceded by spikes in the absorption ratio.

2. Stock prices, on average, depreciated significantly following spikes in the absorption ratio and, on average, appreciated significantly in the wake of sharp declines in the absorption ratio.
3. The absorption ratio provided early signs of the consolidation of the U.S. housing market.
4. The absorption ratio systematically rose in advance of market turbulence.
5. Most global financial crises coincided with positive shifts in the absorption ratio.
6. The absorption ratio captures a large fraction of the information of more complex and computationally intensive structural models of financial contagion.

We proceed as follows. We begin with a literature review of systemic risk and related topics. Then we provide a formal description of the absorption ratio. We next present historical estimates of the absorption ratio for a variety of asset markets, and we show how it relates to asset prices, financial turbulence, the global financial crisis, and financial contagion. We conclude with a summary and suggestions about how regulators and investors might use the absorption ratio as an early warning signal of market stress.

LITERATURE REVIEW

De Bandt and Hartmann [2000] provided an extensive review of the literature on systemic risk. Most studies in this review focus on contagion and “financial fragility,” and the literature on contagion itself is quite rich (see, for example, Pavlova and Rigobon [2008]). Recently, the IMF’s *Global Financial Stability Report* [2009] included a chapter on detecting systemic risk in which it stated, “The current crisis demonstrates the need for tools to detect systemic risk” (p. 111), and “Being able to identify systemic events at an early stage enhances policymakers’ ability to take necessary exceptional steps to contain the crisis” (p. 111). As a simple starting point, the IMF report suggests monitoring conditional (stress) correlations.

In a related study, Billio et al. [2010] showed that correlations increase during market crashes. Prior studies have shown that exposure to different country equity markets offers less diversification in down markets than in up markets.³ The same is true for global industry returns (Ferreira and Gama [2004]), individual stock returns (Ang, Chen, and Xing [2002], Ang and Chen [2002], and Hong, Tu, and Zhou [2007]), hedge fund returns (Van Royen [2002a]), and international bond market returns (Cappiello, Engle, and Sheppard [2006]).

Both the IMF’s *Global Financial Stability Report* [2009] and Billio et al. [2010] suggested that an important symptom of systemic risk is the presence of sudden regime shifts. Investors have long recognized that economic conditions frequently undergo abrupt changes. The economy typically oscillates between

- a steady, low-volatility state characterized by economic growth, and
- a panic-driven, high-volatility state characterized by economic contraction.

Evidence of such regimes has been documented in short-term interest rates (Gray [1996], Ang and Bekaert [2002], and Smith [2002]), GDP or GNP (Hamilton [1989], Goodwin [1993], Luginbuhl and de Vos [1999], and Lam [2004]), inflation (Kim [1993] and Kumar and Okimoto [2007]), and market turbulence (Chow et al. [1999], Kritzman, Lowry, and Van Royen [2001], and Kritzman and Li [2010]).

Billio et al. [2010] independently applied principal components analysis to determine the extent to which several financial industries became more unified across two separate regimes. They found that the percentage of the total variance of these industries explained by a single factor increased from 77% during the 1994–2000 period to 83% during the 2001–2008 period. We instead apply principal components analysis to several broad markets and estimate on a rolling basis throughout history the fraction of total market variance explained by a finite number of factors. We call this measure the absorption ratio. We also introduce a standardized measure of shifts in the absorption ratio, and we analyze how these shifts relate to changes in asset prices and financial turbulence. By applying a moving window in our estimation process, we account for potential changes in the risk factors over time. Because Billio et al. divided history into only two periods, they assumed implicitly that these periods are distinct regimes and are stationary within themselves.

Pukthuanthong and Roll [2009] independently developed a measure similar to the absorption ratio to measure global market integration. They regressed country returns on principal components estimated from a prior period and averaged the R^2 s to produce a multiple R^2 . They demonstrated that this measure of integration is more reliable than correlation and sometimes contradictory to correlation. We make a similar argument that the absorption ratio is a better measure of systemic than average correlation, but our main focus is to demonstrate that the absorption ratio is a reliable indicator of market fragility.

THE ABSORPTION RATIO

Consider a covariance matrix of asset returns estimated over a particular time period. The first eigenvector is a linear combination of asset weights that explains the greatest fraction of the assets’ total variance. The second eigenvector is a linear combination of asset weights orthogonal to the first eigenvector that explains the greatest fraction of leftover asset variance; that is, variance not yet explained or absorbed by the first eigenvector. The third eigenvector and beyond are identified the same way. They absorb the greatest fraction of leftover variance and are orthogonal to preceding eigenvectors.

It is perhaps more intuitive to visualize eigenvectors. The left panel of Exhibit 1 shows a three-dimensional scatter plot of asset returns with a vector piercing the observations. Each observation is the intersection of returns of three assets for a given period, which might be a day, a month, or a year. This vector represents a linear combination of the assets and is a potential eigenvector. The right panel of Exhibit 1 shows the same scatter plot of asset returns but with a different vector piercing the observations.

Of all the potential vectors piercing the scatter plot, we determine the first eigenvector by perpendicularly projecting the observations onto each potential eigenvector, shown by Exhibit 2.

The first eigenvector is the one with the greatest variance of the projected observations, as shown in Exhibit 3.⁴

In order to identify the second eigenvector, we first consider a plane passing through the scatter plot that is orthogonal to the first eigenvector, shown by Exhibit 4.

The second eigenvector must lie on this orthogonal plane, thereby limiting our search. It is the vector that yields the second highest variance of projected observations. We find the third eigenvector in the same fashion. It is the vector that yields the third greatest variance and is orthogonal to the first two vectors. These three eigenvectors together explain the total variance of the assets.⁵

We may or may not be able to associate these eigenvectors with observable economic or financial variables. In some cases the asset weights of the eigenvector may suggest an obvious factor. For example, if we were to observe short exposures to the airline industry and other industries that consume fuel and assume long exposures to the oil industry and other industries that profit from rising oil prices, we might conclude that this eigenvector is a proxy for the price of oil. Alternatively, an eigenvector may reflect a combination of several influences that came together in a particular way that is unique to the chosen sample of assets, in which case the factor may not be definable other

EXHIBIT 1 Three-Dimensional Scatter Plots of Asset Returns






EXHIBIT 2 Projection of Observations onto Vectors



than as a statistical artifact. Moreover, the composition of eigenvectors may not persist through time. Sources of risk are likely to change from period to period.

In some applications this artificiality or nonstationarity would be problematic. If our intent were to construct portfolios that were sensitive to a particular source of risk, then we would like to be able to identify it and have some confidence of its persistence as an important risk factor. But here our interest is not to interpret sources of risk. Instead, we seek to measure the extent to which sources of risk are becoming more or less compact.


EXHIBIT 3
First Eigenvector



The particular measure we use as an indicator of systemic risk is the absorption ratio, which we define as the fraction of the total variance of a set of assets explained or absorbed by a finite set of eigenvectors, as shown in the following equation:

$$AR = \frac{\sum_{i=1}^n \sigma_{Ei}^2}{\sum_{j=1}^N \sigma_{Aj}^2} \quad (1)$$

EXHIBIT 4
Second Eigenvector



where,

AR = absorption ratio

N = number of assets

n = number of eigenvectors in numerator of absorption ratio

σ_{Ei}^2 = variance of the i^{th} eigenvector

σ_{Aj}^2 = variance of the j^{th} asset

A high value for the absorption ratio corresponds to a high level of systemic risk because it implies that the sources of risk are more unified. A low absorption ratio indicates less systemic risk because it implies that the sources of risk are more disparate. We should not expect high systemic risk to necessarily lead to asset depreciation or financial turbulence. It is simply an indication of market fragility in the sense that a shock is more likely to propagate quickly and broadly when sources of risk are tightly coupled.

The Absorption Ratio versus Average Correlation

One might suspect that the average correlation of the assets used to estimate the absorption ratio provides the same indication of market unity, but it does not. Unlike the absorption ratio, the average correlation fails to account for the relevance of the asset correlations that make up the average.

Exhibit 5 shows an increase in the correlation of two hypothetical assets with relatively high volatility and a decrease in the correlation of two hypothetical assets with relatively low volatility. It turns out that although the average correlation decreases slightly from 0.36 to 0.32, the absorption ratio increases sharply from 0.55 to 0.80. The key distinction is that the absorption ratio accounts for the relative importance of each asset's contribution to systemic risk, whereas the average correlation does not.⁶

EMPIRICAL ANALYSIS: THE ABSORPTION RATIO

We present empirical evidence that the absorption ratio effectively captures market fragility. We show that stock prices, on average, depreciate significantly following spikes in the absorption ratio. We demonstrate

EXHIBIT 5

Shift in Correlations and Volatilities

Absorption Ratio versus Average Correlation					
Period 1	Correlations				Standard
Assets	1	2	3	4	Deviations
1	1.00	0.12	-0.01	0.01	35.16%
2	0.12	1.00	-0.04	-0.03	35.07%
3	-0.01	-0.04	1.00	0.82	4.95%
4	0.01	-0.03	0.82	1.00	5.02%

Period 2	Correlations				Standard
Assets	1	2	3	4	Deviations
1	1.00	0.64	-0.05	-0.01	34.46%
2	0.64	1.00	-0.05	-0.03	34.04%
3	-0.05	-0.05	1.00	0.03	4.92%
4	-0.01	-0.03	0.03	1.00	4.88%

that the absorption ratio of the U.S. housing market provided early signs of the emergence of a national housing bubble. We then show that the absorption ratio systematically rose in advance of market turbulence, and

that most global financial crises coincided with positive shifts in the absorption ratio.

The Absorption Ratio and Stock Returns


In order to estimate the absorption ratio, we use a window of 500 days to estimate the covariance matrix and eigenvectors, and we fix the number of eigenvectors at approximately 1/5th the number of assets in our sample.^{7,8} The variances σ_{Ei}^2 and σ_{Aj}^2 in Equation (1) are calculated with exponential weighting. This approach assumes that the market's memory of prior events fades away gradually as these events recede further into the past. The half-life of the exponential weight decay is set to equal half of the window, that is, 250 days.

Exhibit 6 shows a time series of the absorption ratio estimated from the returns of the 51 U.S. industries (hence, 10 eigenvectors) in the MSCI USA index based on trailing 500-day overlapping windows, along with the level of the MSCI USA price index from January 1, 1998, through January 31, 2010.

Exhibit 6 shows a distinct inverse association between the level of the absorption ratio and the level of U.S. stock prices. It also reveals that the absorption ratio increased sharply to its highest level ever during the global financial crisis of 2008, coincident with a

EXHIBIT 6

Absorption Ratio and U.S. Stock Prices



steep decline in stock prices, and that although stock prices have partially recovered as of the second quarter of 2010, the absorption ratio has fallen only slightly. This continued high level for the absorption ratio, while perhaps worrisome, does not necessarily foretell a renewed sell-off in stocks. It does suggest, however, that the U.S. stock market remains extremely fragile and therefore highly vulnerable to negative shocks.

A casual review of the absorption ratio alongside stock prices suggests a coincident relationship, which perhaps casts doubt on the notion that the absorption ratio might be useful as a signal of impending trouble. Exhibit 7 sheds some light on this question. It shows the fraction of significant drawdowns preceded by a spike in the absorption ratio. We first compute the moving average of the absorption ratio over 15 days and subtract it from the moving average of the absorption ratio over one year. We then divide this difference by the standard deviation of the one-year absorption ratio, as shown. We call this measure the standardized shift in the absorption ratio,

$$\Delta AR = (AR_{15 \text{ Day}} - AR_{1 \text{ Year}}) / \sigma \quad (2)$$

where,

- ΔAR = standardized shift in absorption ratio
- $AR_{15 \text{ Day}}$ = 15-day moving average of absorption ratio
- $AR_{1 \text{ Year}}$ = one-year moving average of absorption ratio
- σ = standard deviation of one-year absorption ratio

As Exhibit 7 shows, all of the 1% worst monthly drawdowns were preceded by a one-standard-deviation spike in the absorption ratio, and a very high percentage of other significant drawdowns occurred after the absorption ratio spiked.

EXHIBIT 7

Absorption Ratio and Drawdowns

Fraction of Drawdowns Preceded by Spike in AR			
	1% Worst	2% Worst	5% Worst
One Day	84.85%	87.69%	70.81%
One Week	84.85%	83.08%	75.78%
One Month	100.00%	98.46%	89.44%

Notes: One standard deviation, 15 days/1 year; 1/1/1998 through 5/10/2010.

We should not conclude from Exhibit 7 that a spike in the absorption ratio reliably leads to a significant drawdown in stock prices. In many instances, stocks performed well following a spike in the absorption ratio. We would be correct to conclude, though, that a spike in the absorption ratio is a near necessary condition for a significant drawdown, just not a sufficient condition. Again, a high absorption ratio is merely an indication of market fragility.

Even though a spike in the absorption ratio does not always lead to a major drawdown in stock prices, on average, stocks perform much worse following spikes in the absorption ratio than they do in the wake of a sharp drop in the absorption ratio. Exhibit 8 shows the average annualized one-day, one-week, and one-month returns following a one-standard-deviation increase or decrease in the 15-day absorption ratio relative to the one-year absorption ratio.

Exhibit 8 offers compelling evidence that significant increases in the absorption ratio are followed by significant stock market losses, on average, while significant decreases in the absorption ratio are followed by significant gains. This differential performance suggests that it might be profitable to reduce stock exposure subsequent to an increase in the absorption ratio and to raise exposure to stocks after the absorption ratio falls, which is what we next test.

Exhibit 9 shows the performance of a dynamic trading strategy in which the stock exposure of an otherwise equal-weighted portfolio of stocks and government bonds is raised to 100% following a one-standard-deviation decrease in the 15-day absorption ratio relative to the one-year absorption ratio and reduced to 0% following a one-standard-deviation increase. These rules are summarized as follows:

Absorption Ratio	Stocks/Bonds
$-1\sigma \geq AR \leq +1\sigma$	50/50
$AR > +1\sigma$	0/100
$AR < -1\sigma$	100/0

These rules are applied daily with a one-day lag following the signal for the period January 1, 1998, through January 31, 2010, using the MSCI USA stock index and Treasury bonds.⁹ Exhibit 9 shows that these rules triggered only 1.72 trades per year on average, which should not be surprising given that one-standard-deviation events occur infrequently. Nonetheless, these infrequent

EXHIBIT 8

Absorption Ratio and Subsequent Returns

Annualized Return after Extreme AR			
	One Sigma Increase	One Sigma Decrease	Difference
One Day	-8.28%	9.27%	-17.56%
One Week	-8.44%	10.06%	-18.50%
One Month	-5.86%	12.16%	-18.02%

Notes: One standard deviation, 15 days/1 year; 1/1/1998 through 5/10/2010.

EXHIBIT 9

Absorption Ratio as a Market-Timing Signal

Performance: 100/0 versus 0/100		
	Dynamic	50/50
Return	9.58%	5.08%
Risk	11.50%	10.89%
Return-Risk	0.83	0.47
Turnover	86.01%	
Number of Trades	1.72	

Notes: 1/1/1998 through 5/10/2010.

shifts improved return by more than 4.5% annually while increasing risk by only 0.61%, thereby raising the return-risk ratio from 0.47 to 0.83.

Although most investors might be reluctant to shift entirely in or out of stocks given a single signal, this experiment does offer persuasive evidence of the potential value of the absorption ratio as a market-timing signal. Exhibit 10 reveals that the absorption ratio would have kept investors out of stocks during much of the dot-com meltdown as well as the global financial crisis. It also reveals that the absorption ratio produced some false positives, but not enough to offset its net beneficial effect.

Although daily MSCI industry data extend back only to 1995, Datastream offers less granular industry data which allows us to include the 1987 stock market crash in our analysis. Exhibit 11 shows how well the absorption ratio served as an early warning signal of the

great stock market crashes of the past three decades. If investors withdrew from stocks within one day of a one-standard-deviation spike in the standardized shift of the absorption ratio, they would have avoided all of the losses associated with the 1987 crash and the global financial crisis, and nearly 70% of the losses associated with the dot-com meltdown.

This trading rule appears to improve performance in other stock markets as well. Exhibits 12 and 13 show the absorption ratio estimated from stock returns in Canada, Germany, Japan, and the U.K. and that applied as a timing signal in those markets yielded similar improvement in total return as well as risk-adjusted return.¹⁰

The Absorption Ratio and the Housing Bubble

Former Federal Reserve Chairman Alan Greenspan stated that although regional housing markets often showed signs of unsustainable speculation resulting in local housing bubbles, he did not expect a national U.S. housing bubble.¹¹ Had the Fed examined the absorption ratio of the U.S. housing market, it might have learned that regional housing markets were becoming more and more unified as early as 1998, setting the stage for a national housing bubble.

Exhibit 14 shows the absorption ratio estimated from 14 metropolitan housing markets in the United States, along with an index of the Case-Shiller 10-City National Composite Index.¹²

Exhibit 14 reveals that the housing market absorption ratio experienced a significant step up from 47.60% in October 1996 to 67.04% in March 1998, just as the national housing bubble got underway. It reached a historic peak of 72.76% in September 1998 and then another peak in July 2004 at 85.63% as the housing bubble continued to inflate. It again reached a historic peak of 89.07% in December 2006 within a few months of the housing bubble peak. Then as the housing bubble burst, the absorption ratio climbed sharply, reaching an all-time high of 94.22% in March 2008. As housing prices stabilized and recovered slightly in 2009, the absorption ratio began to retreat modestly.

EXHIBIT 10

Absorption Ratio Stock Exposure

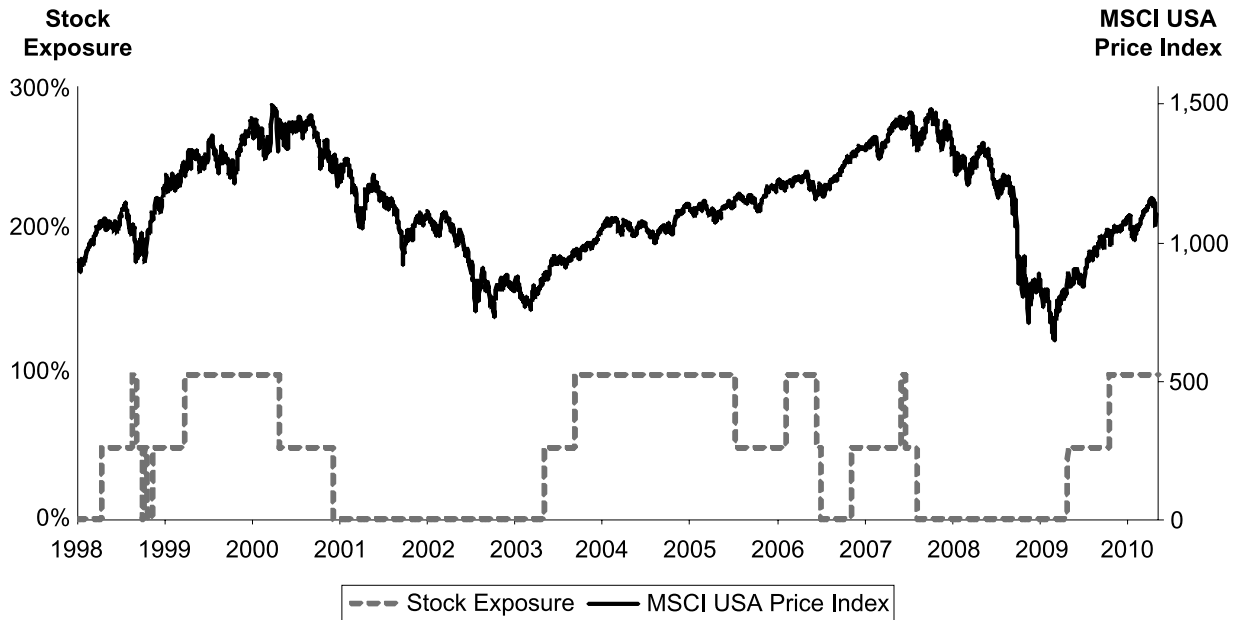



EXHIBIT 11

The Absorption Ratio as an Early Warning Signal of Crashes



It is quite clear from Exhibit 14 that systemic risk in the national housing market increased significantly leading up to the beginning stages of the housing bubble and as the bubble inflated, and that it increased even further after the bubble burst and housing prices tumbled.

The Absorption Ratio and Financial Turbulence

Next, we turn to the relationship between systemic risk and financial turbulence. We define financial turbulence as a condition in which asset prices behave in an uncharacteristic fashion given their historical pattern of behavior, including extreme price moves, decoupling of correlated assets, and convergence of uncorrelated assets. We measure financial turbulence as

$$dt = (y_t - \mu) \Sigma^{-1} (y_t - \mu)' \quad (3)$$

where,

d_t = turbulence for a particular time period t

y_t = vector of asset returns for period t

μ = sample average vector of historical returns

Σ = sample covariance matrix of historical returns

We interpret this formula as follows. By subtracting the historical average from each asset's return, we capture the extent to which one or more of the returns was unusually high or low. In multiplying these differences by the inverse of the covariance matrix of returns, we

EXHIBIT 12

Global Performance of Absorption Ratio


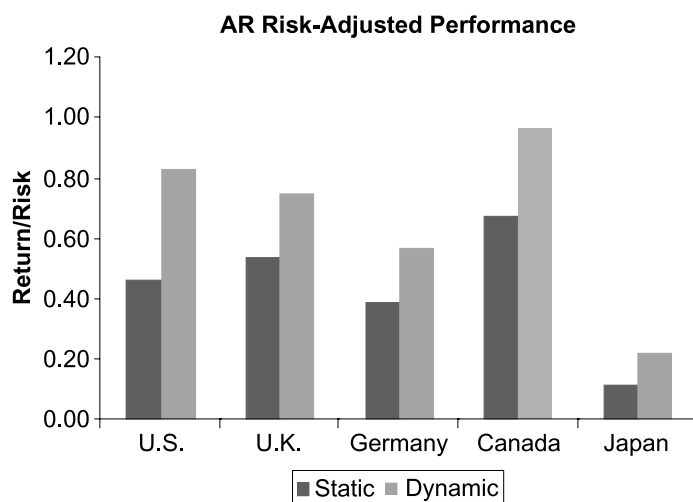


EXHIBIT 13

Global Performance of Absorption Ratio



both divide by variance, which makes the measure scale independent, and we capture the interaction of the assets. By post-multiplying by the transpose of the differences between the asset returns and their averages, we convert this measure from a vector to a single number. Previous research has shown that this statistical characterization of financial turbulence is highly coincident with events in financial history widely regarded as turbulent, as shown in Exhibit 15.¹³

In order to measure the connection between systemic risk and financial turbulence, we first identify the 10% most turbulent 30-day periods, based on average daily turbulence, of the MSCI USA stock index covering the period from January 1, 1997, through January 10, 2010. We then synchronize all of these turbulent events and observe changes in the 15-day absorption ratio relative to the one-year absorption ratio estimated from industry returns as described earlier, leading up to and following the turbulent events. Exhibit 16 shows the results of this event study.

Prior to turbulent events in the stock market, the median of the standardized shift in the absorption ratio increased beginning about 40 days in advance of the event, and continued to rise throughout the turbulent periods. It then fell following the conclusion of the turbulent episodes. This evidence suggests that the absorption ratio is an effective precursor of both the inception and conclusion of turbulent episodes, which could prove to be quite valuable. In addition to persistence, another feature of turbulence is that returns to risk are much lower during turbulent periods than nonturbulent periods.¹⁴

The Absorption Ratio and Global Financial Crises

The previous subsections examined the performance of the absorption ratio in the domestic economy. We now analyze its implications in the global economy. To calculate the global absorption ratio, we collected daily stock market returns for 42 countries (and some regional indices) from February 1995 to December 2009.

Exhibit 17 shows that the global absorption ratio shifts within a range of 65% to 85%. Also, it shows that the global absorption ratio increased in October 1997 (Hong Kong's speculative attack after the Asian financial crises) and August 1998 (the Russian and LTCM collapses), which were two of the most significant emerging market crises in the last 20 years (the other being the Tequila Crisis in 1994, which is outside our sample).

During the recovery in emerging markets (1999–2001), the global absorption ratio decreased. It increased following the boom that came after the severe declines in interest rates that took place after September 11, 2001; the accounting standard scandals; and the dot-com collapse in 2001. Finally, the last significant increase took place

EXHIBIT 14

Absorption Ratio and the National Housing Bubble

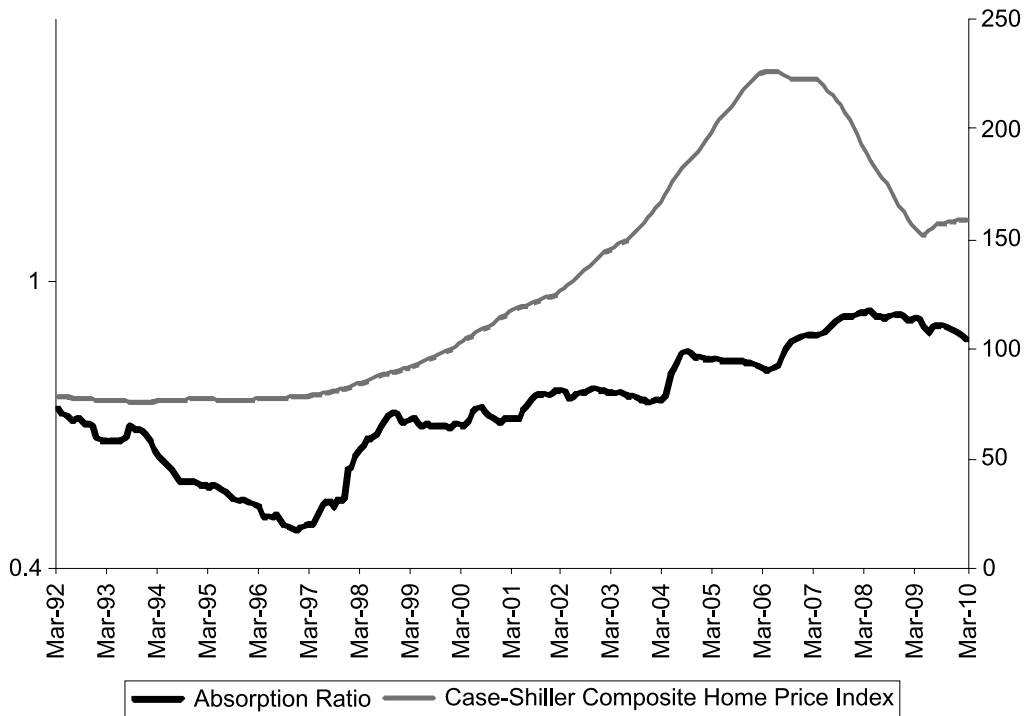
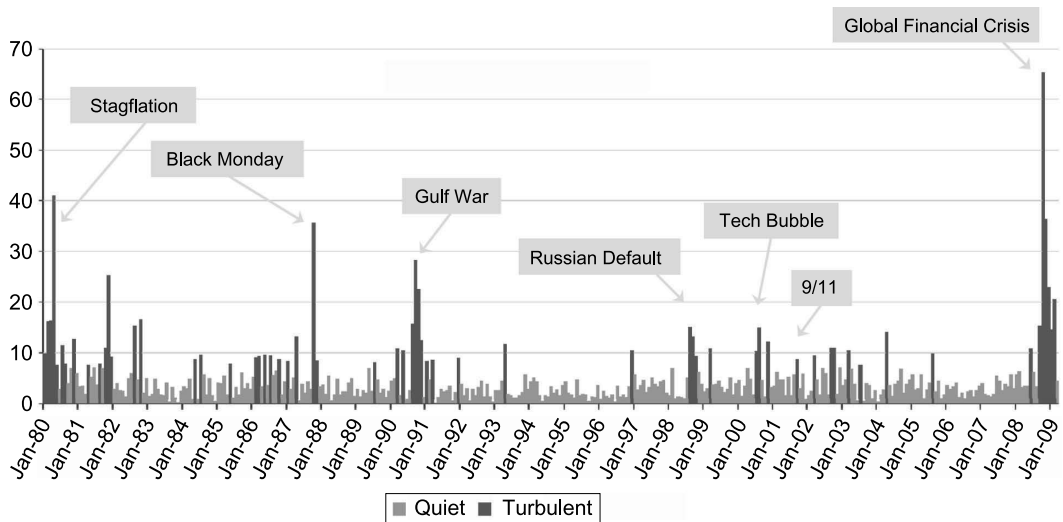


EXHIBIT 15

Financial Turbulence Index



Note: The index is calculated using returns of global stocks, bonds, real estate, and commodities.

EXHIBIT 16

Median Absorption Ratio Around Turbulent Periods





EXHIBIT 17

Global Absorption Ratio



starting in mid-2006 coinciding with the housing bubble, and then the crisis after the Lehman default. This analysis clearly shows that the global systemic risk of the recent crisis was by far the most severe in recent history.

The Absorption Ratio as a Proxy for Financial Contagion

As we mentioned at the beginning of this article, the estimation of systemic risk is extremely challenging

because it is inherently unobservable, and its impact on asset prices is uncertain. One possible approach is to specify a full structural model of asset pricing, such as that of Pavlova and Rigobon [2008], and to estimate the role of a common shock. The advantage of structural models is that, in general, they are over-identified, which means that unobservable shocks can be recovered from the behavior of the observable variables. The main disadvantage is that structural models are extraordinarily difficult to estimate. We present preliminary results of a structural model estimated by Kumar, Pavlova, and Rigobon [2010]. The complexity and computational intensity of their structural model is evidenced by the fact that a single estimation usually takes several months for convergence. Our much simpler measure of systemic risk does a very good job of capturing the information produced by their more complex structural model.

Exhibit 18 shows that the 500-day moving average of the average change in covariance (a preliminary result from the work of Kumar, Pavlova and Rigobon [2010]) closely coincides with the global absorption ratio. Specifically, the two series are 81% correlated.¹⁵

SUMMARY


We have introduced a method for inferring systemic risk from asset prices, which we call the absorption ratio. It is equal to the fraction of a set of assets' total variance explained or absorbed by a finite number of eigenvectors. A high absorption ratio implies that financial markets are relatively compact. When markets are compact they are more fragile because shocks propagate more quickly and broadly. A low absorption ratio suggests that markets are less tightly coupled and therefore more resilient to shocks.

Compact markets do not always lead to asset depreciation, but most significant stock market drawdowns have been preceded by spikes in the absorption ratio. This suggests that spikes in the absorption ratio are a near necessary, but not sufficient, condition for market crashes.

We have shown that stock returns are much lower, on average, following spikes in the absorption ratio than they are in the wake of significant declines in the absorption ratio and that investors could have profited by varying exposure to stocks following significant changes in the absorption ratio. We have demonstrated that the

EXHIBIT 18

Global Absorption Ratio and Smoothed Average Change in Covariance



absorption ratio of the U.S. housing market provided early signs of the emergence of a national housing bubble, long before the Fed recognized this fact. We have presented evidence showing that increases in the absorption ratio anticipate by more than a month subsequent episodes of financial turbulence. Finally, we have shown that variation in the absorption ratio coincided with many global financial crises and tracked other more complex measures of financial contagion. In short, the absorption ratio appears to serve as an extremely effective measure of systemic risk in financial markets.

APPENDIX

THE HERFINDAHL INDEX AS A MEASURE OF IMPLIED SYSTEMIC RISK

As alternative to our methodology of estimating the absorption ratio as the fraction of total variance explained by a subset of eigenvectors, we considered using the Herfindahl index to infer systemic risk. Specifically, we squared the fraction of total variance explained by each eigenvector and took the square root of the sum of these squared values,

$$AR_H = \sqrt{\sum_{i=1}^N \left(\frac{\sigma_{Ei}^2}{\sum_{j=1}^N \sigma_{Aj}^2} \right)^2} \quad (\text{A-1})$$

where,

AR_H = absorption ratio estimated from the Herfindahl index

N = number of assets

σ_{Ei}^2 = variance of the i^{th} eigenvector, sometimes called eigenportfolio

σ_{Aj}^2 = variance of the j^{th} asset

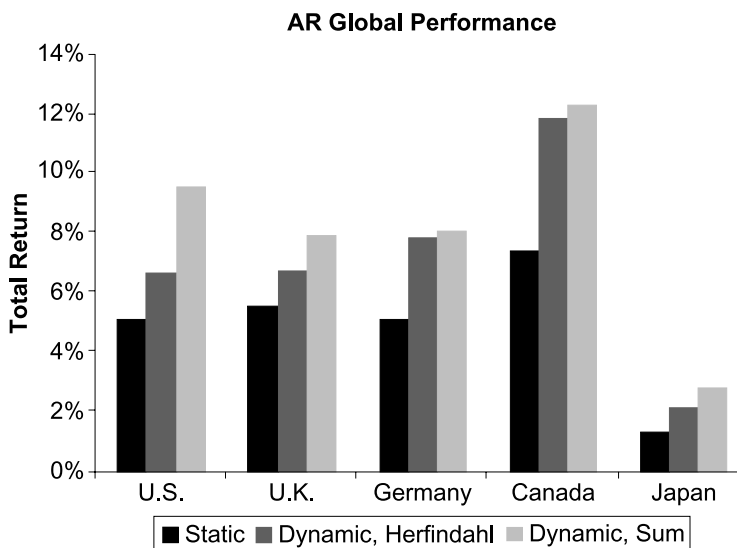
Exhibit A1 clearly shows that our approach leads to better results than using a Herfindahl index to estimate systemic risk. Our conjecture is that the less important eigenvectors, which are included in the Herfindahl index, are relatively unstable and introduce noise to the estimate of systemic risk.

ENDNOTES

We thank Timothy Adler, Robin Greenwood, Richard Roll, and participants of seminars at the European Quantitative Forum, the International Monetary Fund, JOIM, MIT

EXHIBIT A1

Herfindahl Index vs. Sum of Subset



Sloan School, PIMCO, QWAFEFW, and State Street Global Markets in Cambridge, MA, New York, and London for helpful comments.

¹This is a euphemism for accounting practices such as Lehman Brothers' use of swaps to conceal \$50 billion of debt shortly before its demise.

²We could instead calculate the number of eigenvectors required to explain a fixed percentage of variance, but for no particular reason we chose to fix the number of eigenvectors.

³See, for example, Ang and Bekaert [2002], Kritzman, Lowry, and Van Royen [2001], and Baele [2003] on regime shifts; Van Royen [2002b] and Hyde, Bredin, and Nguyen [2007] on financial contagion; and Longin and Solnik [2001], Butler and Joaquin [2002], Campbell, Koedijk, and Kofman [2002], Cappiello, Engle, and Sheppard [2006], and Hyde, Bredin, and Nugyen [2007] on correlation asymmetries.

⁴There are a variety of techniques to identify eigenvectors. We can use matrix algebra to identify eigenvectors given a small set of observations. For larger datasets, it is more efficient to resort to numerical procedures.

⁵The number of eigenvectors never exceeds the number of assets; however, total variation could be explained by fewer eigenvectors than assets to the extent assets are redundant. To be precise, the total number of eigenvectors equals the rank of the covariance matrix.

⁶Pukthuanthong and Roll [2009] provided a formal analysis of the distinction between average correlation and measures of integration based on principal components.

⁷In principle, we should condition the number of eigenvectors on the rank of the covariance. Because the covariance matrices in our analysis are nearly full rank, we are effectively doing this.

⁸As an alternative to measuring market unity by the fraction of total variance explained by a subset of eigenvectors, we could construct a Herfindahl index, by which we square the fraction of variance explained by each of the eigenvectors and sum the squared values. Our experiments show that this latter approach is significantly less informative than our method. We present this evidence in the appendix.

⁹We require two years to estimate the covariance matrix and eigenvectors and another year to estimate the standard deviation of the one-year moving average; hence, our data begin January 1, 1995.

¹⁰As with the U.S. stock market, we set the number of eigenvectors at about 1/5th the number of industries.

¹¹He has since rejected this view.

¹²The returns are computed from the Case–Shiller Seasonally Adjusted Indices that go back to 1987. The metropolitan areas include Los Angeles, San Diego, San Francisco, Denver, Washington, DC, Miami, Tampa, Chicago, Boston, Charlotte, Las Vegas, New York, Cleveland, and Portland. In this case, the covariance matrix is based on five years of monthly returns beginning January 1987 and ending March 2010, and the absorption ratio is based on the first three eigenvectors, roughly 1/5th of the number of assets. The half-life for exponentially weighted variances is set to two-and-a-half years.

¹³For more about this measure of financial turbulence, including its derivation, empirical properties, and usefulness, see Kritzman and Li [2010].

¹⁴See Kritzman and Li [2010].

¹⁵Because both series are quite noisy, we smooth them by taking a 60-day moving average and then compute the correlation of the percentage changes of the smoothed series.

REFERENCES

- Ang, A., and G. Bekaert. “International Asset Allocation with Regime Shifts.” *Review of Financial Studies*, Vol. 15, No. 4 (2002), pp. 1137-1187.
- Ang, A., and J. Chen. “Asymmetric Correlations of Equity Portfolios.” *Journal of Financial Economics*, Vol. 63, No. 3 (2002), pp. 443-494.
- Ang, A., J. Chen, and Y. Xing. “Downside Correlation and Expected Stock Returns.” Working paper, Columbia University, 2002.
- Baele, L. “Volatility Spillover Effects in European Equity Markets: Evidence from a Regime Switching Model.” Working Paper No. 33, United Nations University, Institute for New Technologies, 2003.
- Billio, M., M. Getmansky, A. Lo, and L. Pelizzon. “Measuring Systemic Risk in the Finance and Insurance Sectors.” MIT Sloan Research Paper No. 4774-10, 2010. Available at SSRN: <http://ssrn.com/abstract=1571277>.
- Butler, K., and D. Joaquin. “Are the Gains from International Portfolio Diversification Exaggerated? The Influence of Downside Risk in Bear Markets.” *Journal of International Money and Finance*, Vol. 21, No. 7 (2002), pp. 981-1011.
- Campbell, R., K. Koedijk, and P. Kofman. “Increased Correlation in Bear Markets.” *Financial Analysts Journal*, Vol. 58, No. 1 (2002), pp. 87-94.
- Cappiello, L., R. Engle, and K. Sheppard. “Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns.” *Journal of Financial Econometrics*, Vol. 4, No. 4 (2006), pp. 385-412.
- Chow, G., E. Jacquier, M. Kritzman, and K. Lowry. “Optimal Portfolios in Good Times and Bad.” *Financial Analysts Journal*, Vol. 55, No. 3 (1999), pp. 65-73.
- De Bandt, O., and P. Hartmann. “Systemic Risk: A Survey.” European Central Bank Working Paper No. 35, 2000.
- Ferreira, M., and P. Gama. “Correlation Dynamics of Global Industry Portfolios.” Working paper, ISCTE Business School, 2004.
- Goodwin, T. “Business-Cycle Analysis with a Markov-Switching Model.” *Journal of Business and Economic Statistics*, Vol. 11, No. 3 (1993), pp. 331-339.
- Gray, S. “Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process.” *Journal of Financial Economics*, 42 (1996), pp. 27-62.
- Hamilton, J. “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” *Econometrica*, 57 (1989), pp. 357-384.
- Hong, Y., J. Tu, and G. Zhou. “Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation.” *Review of Financial Studies*, Vol. 20, No. 5 (2007), pp. 1547-1581.

Hyde, S., D. Bredin, and N. Nguyen. "Correlation Dynamics between Asia-Pacific, EU and US Stock Returns." In *Asia-Pacific Financial Markets: Integration, Innovation and Challenges*, Vol. 8, pp. 39-61, M. McKenzie and S.-J. Kim, eds. *International Finance Review*, 2007.

International Monetary Fund (IMF). *Global Financial Stability Report: Responding to the Financial Crisis and Measuring Systemic Risks*, April 2009. Available at <http://www.imf.org/external/pubs/ft/gfsr/2009/01/pdf/text.pdf>.

Kim, C.-J. "Unobserved-Component Time Series Models with Markov-Switching Heteroscedasticity." *Journal of Business and Economic Statistics*, Vol. 11, No. 3 (1993), pp. 341-349.

Kritzman, M., and Y. Li. "Skulls, Financial Turbulence, and Risk Management." *Financial Analysts Journal*, Vol. 66, No. 5 (2010), pp. 30-41.

Kritzman, M., K. Lowry, and A.-S. Van Royen. "Risk, Regimes, and Overconfidence." *The Journal of Derivatives*, Vol. 8, No. 3 (2001), pp. 32-43.

Kumar, M., and T. Okimoto. "Dynamics of Persistence in International Inflation Rates." *Journal of Money, Credit and Banking*, Vol. 39, No. 6 (2007), pp. 1458-1479.

Kumar, P., A. Pavlova, and R. Rigobon. "A Structural Estimation of Portfolio Constraints." Working paper, MIT, 2010.

Lam, P. "A Markov-Switching Model of GNP Growth with Duration Dependence." *International Economic Review*, Vol. 45, No. 1 (2004), pp. 175-204.

Longin, F., and B. Solnik. "Extreme Correlation of International Equity Markets." *Journal of Finance*, Vol. 56, No. 2 (2001), pp. 649-676.

Luginbuhl, R., and A. de Vos. "Bayesian Analysis of an Unobserved-Component Time Series Model of GDP with Markov-Switching and Time-Varying Growths." *Journal of Business and Economic Statistics*, Vol. 17, No. 4 (1999), pp. 456-465.

Pavlova, A., and R. Rigobon. "The Role of Portfolio Constraints in the International Propagation of Shocks." *Review of Economic Studies*, Vol. 75, No. 4 (2008), pp. 1215-1256.

Pukthuanthong, K., and R. Roll. "Global Market Integration: An Alternative Measure and Its Application." *Journal of Financial Economics*, 94 (2009), pp. 214-232.

Smith, D. "Markov-Switching and Stochastic Volatility Diffusion Models of Short-Term Interest Rates." *Journal of Business and Economic Statistics*, Vol. 20, No. 2 (2002), pp. 183-197.

Van Royen, A.-S. "Hedge Fund Index Returns." *Hedge Fund Strategies*, 36 (2002a), pp. 111-117.

———. "Financial Contagion and International Portfolio Flows." *Financial Analysts Journal*, Vol. 58, No. 1 (2002b), pp. 35-49.

To order reprints of this article, please contact David Rowe at d.rowe@pageantmedia.com or 646-891-2157.