
FELIX GOLTZ
Head of Applied Research, EDHEC-Risk Institute


RENATA GUOBUZAITE
Research Assistant, EDHEC-Risk Institute


LIONEL MARTELLINI
Professor of Finance, EDHEC Business School Scientific Director, EDHEC-Risk Institute

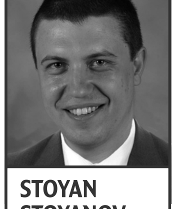
STOYAN STOYANOV
Professor of Finance, EDHEC Business School Head of Research, EDHEC-Risk Institute–Asia

Introducing a New Form of Volatility Index:

THE CROSS-SECTIONAL VOLATILITY INDEX

Recent market turmoil, as well as the presence of ever stricter regulatory constraints, has led investors and asset managers to monitor with increased scrutiny the volatility and downside risk of their equity holdings. In this context, market participants have shown an increasing interest in volatility indices, which are not only heavily used as sentiment indicators, but also serve as underlyings for a number of derivatives contracts that can be used for obtaining long or short exposure to volatility.

While their existence is critical to allow investors to measure and/or trade in volatility, current volatility indices suffer from a number of drawbacks. In particular, they are not based on actual stock index returns, but on auxiliary option markets. As a result, the volatility estimates they provide tend to be polluted by option market factors that have little to do with the underlying stock or stock index returns. Besides, and perhaps more importantly, such measures can only be made available in those rare occasions when liquid option markets exist for the given universe. As a result, very few volatility indices are available, for example, for Asian markets (they have been introduced – recently – only in Japan, Korea, Hong-Kong, India and Australia) or emerging markets beyond Asia (they have been introduced – recently – only in Mexico). Similarly, when it comes to developed equity markets, volatility indices, when they exist, are available only at the broad market level, so that no information is currently available regarding volatility for specific sectors or styles for a given region of the world.

To address these concerns, we introduce a new form of volatility index, the cross-sectional volatility index. Through formal central limit arguments, we show that the cross-sectional variance of stock returns can be regarded as an efficient estimator for the average *idiosyncratic* variance of stocks within the universe under consideration. This measure of idiosyncratic risk is found to be highly correlated to standard measures of systematic risk when they exist, which further justifies its use in the context of equity volatility measurement. Key advantages of this measure over currently available measures such as sample-dependent historical volatility measures or option-based implied volatility measures are: its observability at any frequency, its model-free nature, and its availability for every region, sector and style of the world equity markets, without the need to resort to any auxiliary option market.

The remainder of this note is structured as follows. In section 2, we review the basic definitions of volatility, emphasising the key distinction between historical and implied volatility measures, and also introduce volatility indices. In section 3, we review the motives for trading in volatility, and also the instruments that have been made available for obtaining access to volatility as an investable asset class. In section 4, we introduce a new form of volatility index, which is meant to alleviate existing concerns over currently available volatility indices, and we document the properties of a suitably designed cross-sectional dispersion measure as a model-free unbiased and efficient estimate for specific volatility. In section 5, we provide some interpretation of the cross-sectional volatility index as a proxy for aggregate economic uncertainty, which suggests that the correlation between systematic volatility and average idiosyncratic volatility should be high, which we confirm empirically. Finally, section 6 concludes.

I. Volatility and Volatility Indices

Volatility is a statistical measure of the dispersion of returns for a given security or market index. In other words, volatility refers to the amount of uncertainty or risk about the size of changes in an underlying security or index value. Higher volatility means that the underlying value can potentially be spread out over a larger range of values, signaling a higher riskiness for investors holding the security or index. Indeed, a high volatility means that the price of the security can change dramatically in either direction over a given interval of time. Lower volatility means that a security or index value does not fluctuate dramatically, but changes in value at a steadier pace over a period of time.

I.1. HISTORICAL VERSUS IMPLICIT VOLATILITY

While volatility is unambiguously formally defined as the standard deviation or variance of returns from a given stock or market index, a key distinction exists between two different kinds of volatility measures, namely historical volatility measures and implicit volatility measures.

1.1. Historical Volatility Measures

Historical volatility measures are obtained by estimating the standard deviation of returns from a past sample of equity returns. One advantage of these measures is that they can be estimated directly from time-series of individual stock or stock index returns. One drawback of these measures is that they are not directly observable, and are dependent on a sample of past returns, which makes historical volatility estimates backward-looking measures of current volatility. Besides, historical volatility measures are artificially smooth, in that tomorrow's estimate for historical volatility will only differ from today's estimate for historical volatility by one single data point in a rolling-window analysis. More forward-looking estimates for historical volatility are available through GARCH-type models, which assign a higher weight to more recent observations, but the sample-dependency and artificial smoothing problems remain unsolved.

In an attempt to alleviate the concern over artificial smoothing, a large body of recent research has focussed on improving volatility forecasts through the use of high-frequency data. The main challenge here is that neglecting microstructure noise in calculating realized volatility can lead to biased and inconsistent estimates of the integrated volatility as a true measure of daily volatility (see for example Zhang, Mykland and Aït-Sahalia (2012) for a recent reference). Another related promising area of research in the area of historical volatility modelling is the use of data sampled at various frequencies through the Mixed Data Sampling (MIDAS) model introduced by Ghysels, Santa-Clara, and Valkanov (2005) (see also Ghysels, Sinko and Valkanov (2006) for a review of the field and discussions of possible future directions).

1.2. Implied Volatility Measures

More recently, implicit volatility estimates have been obtained from option prices. Indeed, volatility is a key variable in option pricing formulas showing the extent to which the return of the underlying asset will fluctuate between now and the option expiration date. One advantage of these measures is that they are more forward-looking compared to historical volatility measures since they reflect market's expectations about future volatility. One drawback is that they are not based on actual stock index returns, but on an auxiliary option markets. As a result, the volatility estimate is polluted by any factor impacting supply and demand in option markets that have little to do with the underlying stock or stock index returns. Besides, such measures can only be made available in those rare cases when there are liquid option markets for the given universe under consideration.

I. 2. SYSTEMATIC VERSUS SPECIFIC VOLATILITY

Regardless of the method used in estimating volatility, another key distinction exists between systematic and specific volatility. For any given stock, total volatility can be decomposed into systematic volatility, driven by the stock exposure with respect to systematic risk factors, and specific volatility, which is driven by the uncertainty impacting a particular company.

The recent financial literature has paid considerable attention to idiosyncratic volatility. Campbell et al. (2001) and Malkiel and Xu (2002) document that idiosyncratic volatility increased over time, while Brandt *et al.* (2009) show that this trend completely reversed by 2007, falling below pre-1990s levels, and suggest that the increase in idiosyncratic volatility through the 1990s was not a time trend but rather an "episodic phenomenon". Bekaert *et al.* (2008) confirm that there is no trend both for the United States and other developed countries. A second fact about idiosyncratic volatility is also a source of contention. Goyal and Santa-Clara (2003) put forward that idiosyncratic volatility has forecasting power for future excess returns, while Bali *et al.* (2005) and Wei and Zhang (2005) find that the positive relationship is not robust to the sample chosen.

While representing two different underlying risk measures, one expects systematic and average specific volatility risk indicators to be highly correlated, since they both reflect the aggregate uncertainty faced by investors at a given point in time regarding economic fundamentals (see section 5.1).

In what follows, we confirm this intuition and find a high correlation level between the VIX index, a measure of systematic risk based on option prices, and a model-free measure for specific risk which we introduce in this note (see section 5.2). This high correlation is found to be robust with respect to different regions and time periods, and remains stable across market conditions.

I.3. VOLATILITY INDICES

Because information about changes in volatility is of a critical importance to market participants, a number of initiatives have been launched so as to make volatility measures available to investors in form of volatility indices, which are designed to track the aggregate volatility of an asset market. Such indices are typically calculated based on option prices, and as such are based on implicit as opposed to actual volatility measures.

Volatility indices can now be found on some of the major stock indices such as the S&P500, EURO STOXX 50, DAX, Dow Jones Industrial Average and lately (November 2010) the Nikkei 225. The most popular volatility index is the VIX, which is built from prices of equity index options on the S&P 500. The index was introduced by the options exchange CBOE in 1993 and was originally designed to measure the market's expectation of thirty-day volatility implied by prices of at-the-money S&P 100 index options. The implied volatility was based on the Black-Scholes model. Ten years later, CBOE, together with Goldman

Sachs, updated the VIX methodology to reflect a new way to measure expected volatility. The implied volatility is now derived from the implied risk-neutral distribution that can be extracted from call and put options with different strike prices. This volatility is model free: rather than assuming that Black-Scholes holds, the only necessary assumption is that of an absence of arbitrage opportunities.[1]

II. Volatility as an Asset Class

Recent market turbulence, as well as the presence of ever stricter regulatory constraints, has led investors and asset managers to monitor with increased scrutiny the volatility and downside risk of their equity holdings. In this context, investors have shown an increased interest in regarding volatility as an asset class that can be traded, as opposed to a mere statistical indicator measuring stock return uncertainty.

II. 1. MOTIVES FOR TRADING VOLATILITY

One of the main motivations for trading in volatility is to diversify equity risk through a long implied volatility exposure (see Hill (2004) or Szado (2009) for recent references). A key point to note is precisely that volatility of equity returns and equity returns tend to move in opposite directions, i.e. they are strongly negatively correlated. In addition, negative correlation and high volatility are particularly pronounced in stock market downturns, offering protection against stock market losses when it is needed the most and when other forms of diversification are not very effective.

One possible explanation for the negative correlation of equity volatility to the equity market is the "leverage effect" (Black (1976), Christie (1982), Schwert (1989)): a decrease (respectively, an increase) in equity prices increases (respectively, decreases) the company's leverage, thereby increasing (respectively, decreasing) the risk to equity holders and increasing (respectively, decreasing) equity volatility. Another alternative explanation (French *et al.* (1987), Bekaert and Wu (2000), Wu (2001), Kim *et al.* (2004)) is the "volatility feedback effect": assuming that volatility is incorporated in stock prices, a positive volatility shock increases the future required return on equity and stock prices are expected to fall simultaneously.

The presence of sensible economic reasons that explain the inverse relationship between equity return and volatility is a comforting indication of the robustness of the diversification benefits to be expected, which stands in contrast with the well-known lack of robustness of portfolio diversification within the equity universe, where diversification is known to fail precisely when it is most needed because of the convergence of all correlations to one in periods of high market turbulence.

Of course, it should be expected that the risk diversification benefits of long volatility exposure may come at a cost. Recent academic research has found that there is a positive risk premium over time to being short volatility or conversely, that there is a negative risk premium to being long volatility. Because of the negative correlation between market index returns and market index volatility, buyers of options may be willing to pay a premium because a long position in volatility helps hedge market-wide risk (Bakshi and Kapadia (2003)). In other words, because volatility is negatively correlated with the returns to equities, investors are willing to pay a premium to hold this asset. In a recent paper, Carr and Wu (2010) find that the negative correlation between stock index returns and the return variance generates a strongly negative beta, but this negative beta only explains a small portion of the negative variance risk premiums. Other risk factors identified by the recent literature, such as size, book-to-market, and momentum, cannot explain the strongly negative variance risk premiums either, and they conclude that the majority of the market variance risk premium is generated by an independent variance risk factor.

Other motives for trading in volatility include speculative, arbitrage and hedging demands.

– Directional *speculative* bets on volatility changes, implemented by going long volatility exposure when volatility is expected to rise and going short volatility exposure when volatility is expected to fall (see Dash (2005) or Jacob (2009)).

– Non-directional speculative *arbitrage* bets, so as to benefit from mean-reversion to more normal levels in a number of key spread measures such as implied volatility versus historical volatility, three-month implied volatility vs. one-month-implied volatility, etc.

– Volatility exposure *hedging* by hedge funds and mutual funds that are often implicitly short volatility. In particular, benchmarked equity fund managers are short volatility because portfolio tracking error and rebalancing costs increase with an increase in the volatility of equity markets (Hill (2004)).

Obviously, the ability for market participants to implement trades on volatility critically depends on the availability of instruments that can be used to make volatility indices investable quantities. Since volatility is the key determinant of option prices, trading in options is one possible way to get volatility exposure, but this is not a clean bet on volatility alone. With derivatives instruments on volatility, the investors can have a pure play on volatility and can implement their views more precisely. Investors can invest in volatility products though exchange-traded futures and options on the volatility indices. Also, other products on volatility are available like exchange traded notes and OTC variance swaps.

One of the biggest problems with volatility-related products is liquidity. Most of the futures on these volatility indices have low trading volume and open interest. The most liquid are VIX futures which have an open interest of around 50,000 contracts and a daily volume of several thousand contracts. The next most liquid product is the Mini-VIX futures which has very low open interest (~100s) compared to VIX futures and a low daily volume.

■ III. Cross-Sectional Volatility Index as a New Volatility Index

Because they play a central role as market uncertainty measures and because they are used as a basis for investable volatility products, informative and robust volatility indices are critically important to a large number of market participants.

As we have seen, existing volatility indices suffer from a number of shortcomings. On the one hand, volatility indices are not available for an extensive set of markets, because they require the presence of a liquid option market; for example, no volatility index exists for small-cap stocks, growth/value stocks, or various sectors for developed markets, and volatility indices do not exist even at the broad market level in most emerging markets. On the other hand, where and when they exist, implied volatility estimates are plagued by option-market problems that have little to do with underlying equity markets.

In what follows, we introduce a new set of volatility indices, which are meant to be based on observable and model-free volatility measures, obtained using equity market data alone and available for all markets/sectors at all frequencies. These indices are based on the cross-sectional dispersion of observed stock returns on a given date, a readily measurable quantity for any equity universe. Using formal central-limit arguments, we show that under mild simplifying assumptions this cross-sectional measure provides a very good approximation for average idiosyncratic variance.

The conceptual and technical foundations for these indices have been outlined by Garcia, Mantilla-Garcia and Martellini (2011), and we propose below a short summary of the formal arguments they present to motivate the use of cross-sectional dispersion as a measure of volatility.

III. 1. CROSS-SECTIONAL DISPERSION AND SPECIFIC VOLATILITY

We first assume without any loss of generality the following single conditional factor model for (excess) stock returns: $r_{it} = \beta_{it} F_t + \varepsilon_{it}$, where F_t is the factor excess return at time t, β_{it} is the beta of stock i at time t, and ε_{it} is the residual or specific return on stock i at date t, with $E(\varepsilon_{it}) = 0$ and $cov(F_t, \varepsilon_{it}) = 0$. We assume that the factor model under consideration is a strict factor model, that is $cov(\varepsilon_{it}, \varepsilon_{jt}) = 0$ for $i \neq j$.

We further make the following two simplifying assumptions (we will argue below that they come with very little loss of generality):

– Homogeneous beta assumption: $\beta_{it} = \beta_t$ for all i;

– Homogeneous residual variance assumption: $E(\varepsilon^2_{it}) = \sigma^2_\varepsilon(t)$ for all i.

Under these assumptions, Garcia, Mantilla-Garcia and Martellini (2011) show that cross-sectional variance converges towards specific variance in the limit of an increasing large number of constituents:

$$CSV^2_{t,w_t} = \sum_{i=1}^{N_t} w_{it}\left(r_{it} - \overline{r}_{t,w_t}\right)^2 \xrightarrow[N_t \to \infty]{} \sigma^2_\varepsilon(t)$$

where $\overline{r}_{t,w_t}$ is the weighted-return on the portfolio with weights w_{it} at date t, CSV_t is the cross-sectional volatility, and N_t is the number of constituents in the universe for a given date t. The proof of this result is based upon central limit arguments, and we refer the interested reader to Garcia, Mantilla-Garcia and Martellini (2011) for more details.

This result is important because it draws a formal relationship between the dynamics of the cross-sectional dispersion of realised returns and the dynamics of idiosyncratic variance.[2] Note that this asymptotic result holds for any (non-trivial) weighting scheme. Of course, for a finite number of constituents N_t, different weighting schemes will generate different proxies for idiosyncratic variance. In fact, it can be shown that the estimator with equal-weights is the best estimator for idiosyncratic variance within the class of estimators obtained under a strictly positive weighting scheme. To see this, we report the following result from Garcia, Mantilla-Garcia and Martellini (2011) about the first two moments of the estimator for a finite number of constituents N_t:

$$E\left[CSV^2_{t,w_t}\right] = \sigma^2_\varepsilon(t)\left(1 - \sum_{i=1}^{N_t} w^2_{it}\right)$$

$$Var\left[CSV^2_{t,w_t}\right] = 2\sigma^4_\varepsilon(t)\left(\left(\sum_{i=1}^{N_t} w^2_{it}\right)^2 + \sum_{i=1}^{N_t} w^2_{it} - 2\sum_{i=1}^{N_t} w^3_{it}\right)$$

Hence the squared cross-sectional volatility estimator appears to be a biased estimator for idiosyncratic variance, with a bias given by the multiplicative factor $1 - \sum_{i=1}^{N_t} w^2_{it}$, which can be easily corrected for (see section 4.2). In the end, the bias and variance of the estimator appear to be minimum for the equally weighted (EW) scheme, which corresponds to taking $w_{it} = 1/N_t$ at each date t.

For the equally-weighted scheme, we thus have:

$$E\left[CSV^2_{t,EW}\right] = \sigma^2_\varepsilon(t)\left(1 - \frac{1}{N_t}\right) \xrightarrow[N_t \to \infty]{} \sigma^2_\varepsilon(t)$$

$$Var\left[CSV^2_{t,EW}\right] = 2\sigma^4_{\varepsilon_t}\left(\frac{N_t - 1}{N^2_t}\right) \xrightarrow[N_t \to \infty]{} 0$$

Unlike most previous measures of average idiosyncratic variance, the cross-sectional one offers two main advantages: it can be computed directly from observed returns, with no need to estimate other parameters such as betas, and it is available at any frequency and for any universe of stocks. For any given weighting scheme (in particular EW or cap-weighted), the corresponding cross-sectional measure has the obvious advantage of being readily computable at any frequency from observed returns. This stands in contrast with the previous approaches that have used monthly measures based on time series regressions on daily returns. The second important feature of the cross-sectional estimator is its model-free nature, since we do not need to specify a particular factor model to compute it.

We now discuss the implications of relaxing the two simplifying assumptions of homogenous residual variances and homogenous betas.

If we first relax the homogeneous residual variance assumption, we obtain that:

$$E\left[CSV_{t,EW}^{2}\right]=\left(\frac{1}{N_t}\sum_{i=1}^{N_t}\sigma_{\varepsilon_i}^{2}(t)\right)\left(1-\frac{1}{N_t}\right)$$

Hence, the assumption of homogenous residual variances comes with no loss of generality. In the general case with non-homogenous variances, the cross-sectional estimator simply appears to be a (biased) estimator for the average idiosyncratic variance of the stocks in the universe.

If we now relax the homogeneous beta assumption, we obtain that:

$$E\left[CSV_{t,EW}^{2}\right]=\left(\frac{1}{N_t}\sum_{i=1}^{N_t}\sigma_{\varepsilon_i}^{2}(t)\right)\left(1-\frac{1}{N_t}\right)+F_t^2\sigma_{t,\beta}^2$$

where

$$\sigma_{t,\beta}^2=\frac{1}{N_t}\sum_{i=1}^{N_t}\left(\beta_{it}-\frac{1}{N_t}\sum_{i=1}^{N_t}\beta_{it}\right)^2$$

Hence, the cross-sectional estimator appears to be a biased estimator for the average specific variance, even for an increasingly large number of stocks, when the homogenous beta assumption does not hold. Empirically, however, this bias is found to be small. In particular, Garcia, Mantilla-Garcia and Martellini (2011) estimate that the median value for this term is a mere .348% of specific variance based on daily data for US stocks in the CRSP data base on the sample period ranging from July 1963 to December 2006.

III. 2. ESTIMATION METHODOLOGY

To construct the volatility index, we first gather information on the entire constituent universe. In particular, we require information on past returns data for purposes of filtering stocks, as well as on current returns for the construction of the current volatility index value.

For the case of the cross-sectional volatility index, we compute the following transformation of cross-sectional variance, which we have shown to be an unbiased estimator for specific volatility within the universe:

$$CSV_t=\sqrt{\frac{\sum_{i=1}^{N_t}\left(r_{it}-\bar{r}_t^{EW}\right)^2}{N_t-1}}$$

where r_t^{EW} is the return on the equally-weighted portfolio at date t.

This natural estimator, however, is very sensitive to outliers and is not appropriate. We use instead a quantile-based variance estimator which is meant to enhance the robustness of the dispersion estimator by reducing its sensitivity to the presence of outliers. The idea behind the construction of quantile-based variance indicators is to take advantage of the elementary relationship between quantiles and the spread of a distribution. For example, for i.i.d. Gaussian returns, a robust quantile-based estimator for the standard-deviation is given by *(Q95%-Q5%)/(2×1.645)*, where *Qk%* denotes the k[th] quantile of returns (see for example David (1970)). It is worth noting at this stage that more refined approaches to make the cross-sectional volatility index estimate less sensitive to the presence of outliers could be implemented, e.g., an approach based upon a robust regression model (see Appendix 1 for more details).

IV. Interpretation of the Cross-Sectional Volatility Index

In this section, we argue that the cross-sectional volatility index can be regarded as a proxy for economic uncertainty. We also show that it exhibits a high correlation with the VIX index.

IV. 1. CROSS-SECTIONAL VOLATILITY AS A PROXY FOR ECONOMIC UNCERTAINTY

As argued below, average specific volatility can be regarded as a proxy that reflects the aggregate uncertainty faced by investors at a given point in time regarding economic fundamentals. To put this analysis in the proper context, we should go back to the very nature of idiosyncratic risk. In an asset pricing model, it represents the risk that belongs to an individual firm, after accounting for the sources of risk that are common to all firms. In the previous sections, we have shown that the cross-sectional variance of returns provides a very good measure of this idiosyncratic risk, even if it ignores the risk exposures to the usual common risk factors such as the market return or the Fama-French factors.

To get a first sense of the relationship between the cross-sectional volatility index and economic conditions, we first plot the time-series of the cross-sectional volatility index against NBER recessions for the period 1990 to 2010 (see Garcia, Mantilla-Garcia and Martellini (2011) for a longer sample). The shaded areas in figure 1, which time stamp the NBER recession periods, indicate that the peaks in the probability of remaining in the high-mean high-variance regime coincide most of the times with the contraction periods. Therefore, the cross-sectional volatility measure appears to be counter-cyclical, with the dispersion of returns being high and quite variable when economic growth subsides.

To confirm the interpretation of cross-sectional volatility index as a measure of aggregate economic uncertainty, Garcia, Mantilla-Garcia and Martellini (2011) further compare the cross-sectional volatility index series to a measure of consumption volatility used in the academic literature (see Bansal and Yaron (2004) or Tedongap (2010)). More specifically, they use the Federal Reserve Economic Data (FRED) personal consumption expenditures of non-durables and services monthly series, divided by the consumer price index and the population values to obtain a per-capita real consumption series, the volatility of which is filtered in a GARCH model. While the cross-sectional volatility series is much noisier than consumption-growth volatility, the coincident movements between the two series are found to be quite remarkable, and the correlation between the two series reaches .401. After a highly volatile period just before 2000, both series show a marked downward trend after the turn of the century.

A reasonable explanation for this strong correlation is again to think about a common factor (aggregate economic uncertainty) affecting the idiosyncratic variance of each security. Aggregating over all securities will make the cross-sectional volatility a function of economic uncertainty. The interpretation of the cross-sectional volatility index as a proxy for economic uncertainty may suggest that an investable form of the cross-sectional volatility index may serve as a hedge against economic slowdown uncertainty.

Extending previous work on mean reverting unit daily GARCH process (Engle and Rangel (2007)) and on mixed data sampling (MIDAS) (e.g. Ghysels, Santa-Clara, and Valkanov (2005), they find.

IV. 2. COMPARISON BETWEEN VIX AND CROSS-SECTIONAL VOLATILITY INDEX

As argued above, while representing *a priori* two different underlying risk measures, systematic and average specific volatility indicators should be highly correlated, since they both reflect the aggregate uncertainty faced by investors at a given point in time regarding economic fundamentals. In what follows, we confirm this intuition and find a high correlation level between the VIX index, an imperfect measure of systematic risk based on option prices, and the cross-sectional volatility index, a model-free efficient and unbiased proxy for specific risk.

Figure 1 also shows the time-series of the quantile-based cross-sectional volatility index, in comparison with the VIX index, based on daily stock return data for the S&P500 universe on the sample period ranging from January 1991 to January 2011. We find a high correlation (.70) on the sample period that confirms that option-implied volatility is closely related to its average idiosyncratic volatility counterpart. We have also noted that the high correlation is robust to changes in market conditions, with a conditional correlation that tends to be higher in down markets. For example, the correlation between VIX and cross-sectional volatility is estimated to be .73, slightly higher than the unconditional estimate of .70, when daily market returns are one standard deviation below the mean.

In order to provide further robustness checks for the relationship between VIX and cross-sectional volatility, and check whether our results remained consistent over time, we have generated a ten-year rolling window correlation estimate for the two series on the period ranging from January 1991 to January 2011 (see figure 2). These results suggest that the correlation level is systematically high, whatever the time period under consideration. They also suggest evidence of an upward-sloping trend, with a correlation level that has increased to .77 for the most recent ten year period (2001-2011). Not only is correlation between the cross-sectional volatility index and the option-implied volatility index robust across time, but it is also robust across different regions, as evidenced by the results in table 1.

Figure 1: Time evolution of the cross-sectional volatility index and VIX and recession periods. The shaded areas are the NBER recessions.

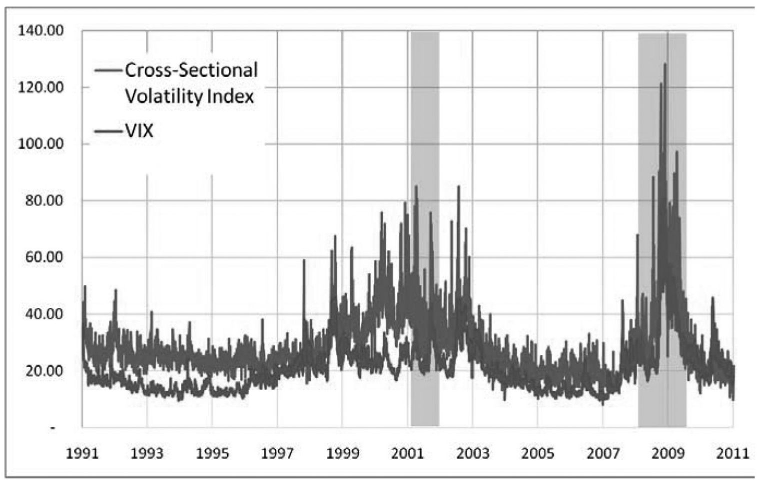

V. Conclusion

We introduce a new form of volatility index, the cross-sectional volatility index. Through formal central limit arguments, we show that the cross-sectional dispersion of stock returns can be regarded as an efficient estimator for the average idiosyncratic volatility of stocks within the universe under consideration. Amongst the key advantages of the cross-sectional volatility measure over currently available measures are its observability at any frequency, its model-free nature, and its availability for every region, sector and style of the world equity markets, without the need to resort to any auxiliary option market. We also provide some interpretation of the cross-sectional volatility index as a proxy for aggregate economic uncertainty, which suggests that the cross-sectional volatility index should be closely related to option-based implied volatility measures. We confirm this intuition by reporting a high correlation level between the VIX index and the corresponding cross-sectional volatility index based on the S&P500 universe. We also find the high correlation between the two volatility measures to be robust with

Table 1: Option-implied volatility indices and their correlation with corresponding cross-sectional volatility index series

Index	Ticker	Correlation	Sample period
CAC 40 Volatility	VCAC	0.6555	01.2000 – 12.2010
FTSE 100 Volatility	VFTSE	0.7235	01.2000 – 12.2010
DAX Volatility	VDAX-NEW	0.6864	01.1992 – 12.2010
SMI Volatility	VSMI	0.6251	06.1999 – 12.2010
EURO STOXX 50 Volatility	VSTOXX	0.6783	01.1999 – 12.2010
NIKKEI Volatility Index	VNKY	0.7032	01.2001 – 12.2010
India NSE VIX	INVIXN	0.6709	11.2007 – 12.2010
KOSPI 200 Volatility Index	VKOSPI	0.5863	01.2003 – 12.2010
Mexico Volatility Index	VIMEX	0.5727	03.2004 – 12.2010

respect to changes in sample period and across different regions. Overall, these results suggest that the cross-sectional volatility index is intimately related to other volatility measures where and when such measures are available, and that it can be used as a reliable proxy for volatility when such measures are not available.

ACKNOWLEDGEMENTS

This research is related to the companion paper “Idiosyncratic Risk and the Cross-Section of Stock Returns” by R. Garcia, D. Mantilla-Garcia and L. Martellini, which provides a detailed analysis of the theoretical grounding for the use of cross-sectional dispersion as a measure of idiosyncratic volatility. It has benefited from the support of the “Produits Structurés et Produits Dérivés” research chair at EDHEC-Risk Institute supported by the Fédération Bancaire Française. We thank Michel Crouhy, Robert Engle, Stephane Gregoir, Abraham Lioui, Harrison Hong, Frederic Smadja, Georges Tauchen, Stephane Tyc, Vijay Vaidyanathan, Volker Ziemann, as well as seminar and conference participants at EDHEC Business School, the Financial Econometrics Conference (Toulouse), the Financial Econometrics Workshop (SMU), the Conference in Asset and Risk Management in the Aftermath of the Financial Crisis (HEC Lausanne), for discussions and feedback. We also thank an anonymous referee for useful comments. ■

Figure 2: 10-year rolling window correlation between the cross-sectional volatility index and the VIX index using daily data on the sample period ranging from January 1991 to January 2011.

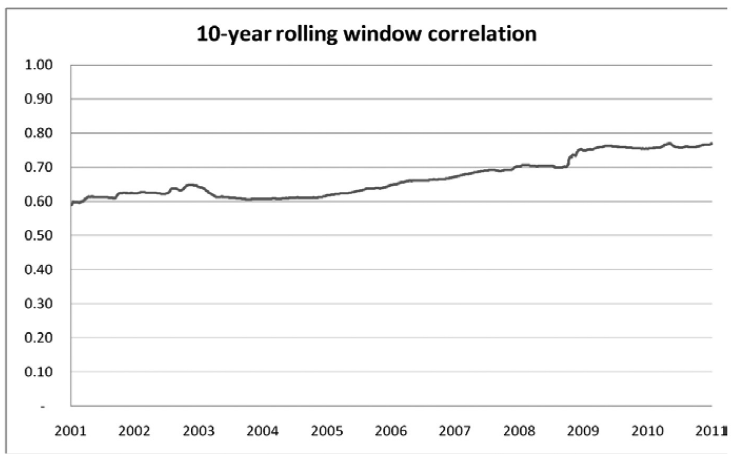

1 Other exchanges have developed volatility indices that are similar in methodology to those provided by the CBOE, see Goltz *et al.* (2011) for additional information.

2 In the practitioners’ literature, cross-sectional dispersion of returns is called “variety” and is used in performance analysis with no formal link to specific volatility (see DiBartolomeo (2006)).

References

- Ait-Sahalia Y., Mykland P. and Zhang L. (2012), "Ultra high frequency volatility estimation with dependent microstructure noise", *Journal of Econometrics*, forthcoming.
- Bakshi G. and Kapadia N. (2003), "Delta-Hedged Gains and the Negative Market Volatility Risk Premium", *Review of Financial Studies*, 16, 2, 527–566.
- Bansal R. and Yaron A. (2004), "Risks for the long run: A potential resolution of asset pricing puzzles", *Journal of Finance*, 59, 4, 1481-1509.
- Bekaert G., Hodrick R.J. and Zhang X. (2008), "Is There a Trend in Idiosyncratic Volatility?", SSRN eLibrary.
- Bekaert G. and Wu G. (2000), "Asymmetric Volatilities and Risk in Equity Markets", *Review of Financial Studies*, 13, 1, 1-42.
- Black F. (1976), Studies in Stock Price Volatility Changes, Proceedings of the 1976 Business, Meeting of the Business and Economic Statistics Section, American Statistical Association, p. 177-181, 22.
- Brandt M., Brav A., Graham J. and Kumar A. (2009), "The Idiosyncratic Volatility Puzzle: Time Trend or Speculative Episodes", forthcoming, *Review of Financial Studies*.
- Campbell J., Lettau M., Malkiel B., and Xu Y. (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk", *Journal of Finance*, 56, 1, 1-43.
- Carr P. and Wu L. (2009), "Variance Risk Premiums", *Review of Financial Studies*, 22, 3, 1311-1341.
- Christie A. (1982), "The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects", *Journal of Financial Economics*, 10, p. 407-432.
- Dash S. and Moran M.T. (2005), "VIX as a Companion for Hedge Fund Portfolios", *Journal of Alternative Investments*, 8, 3, 75-80.
- David H. (1970), *Order Statistics*, John Wiley & Sons, Inc., New York, NY.
- Engle R., Ghysels E. and Sohn B. (2008), "On the Economic Sources of Stock Market Volatility", *working paper*.
- Ghysels E., Santa-Clara P. and Valkanov R. (2005), "There is a Risk-Return Tradeoff after all", *Journal of Financial Economics*, 76, 509-548.
- Ghysels E., Sinko A. and Valkanov R. (2006), "MIDAS Regressions: Further Results and New Directions", *Econometric Reviews* 26, 53-90.
- French K., Schwert G. and Stambaugh R. (1987), "Expected Stock Return and Volatility, *Journal of Financial Economics*, 19, 3-30.
- Goyal A. and Santa-Clara P. (2003), "Idiosyncratic Risk Matters!", *Journal of Finance*, 58, 3, 975-1008.
- Garcia R., Mantilla-Garcia D. and Martellini L. (2011), "Idiosyncratic Risk and the Cross-Section of Stock Returns", *working paper*, EDHEC-Risk Institute.
- Hill J. and Rattray S. (2004), *Volatility as a Tradable Asset: Using the VIX as a market signal, diversifier and for return enhancement*, Equity Product Strategy, Goldman Sachs & Co.
- Huber P. and Ronchetti E. (2009), *Robust Statistics*, 2nd edn, Wiley.
- Jacob J. and Rasiel E. (2009), *Index Volatility Futures in Asset Allocation: A Hedging Framework*, Lazard Investment Research.
- Kim C., Morley J. and Nelson C. (2004), "Is there a Positive Relationship between Stock Market Volatility and the Equity Premium?", *Journal of Money, Credit, and Banking*, 36, 3, 339-360. Malkiel B.G. and Xu Y. (2002), "Idiosyncratic risk and security returns", *working paper*, Princeton University.
- Schwert G. (1989), "Why Does Stock Market Volatility Change over Time?", *Journal of Finance*, 44, 5, 1115-1153.
- Szado E. (2009), "VIX Futures and Options – A Case Study of Portfolio Diversification During the 2008 Financial Crisis", *working paper*, University of Armhest.
- Tédongap R. (2010), "Consumption Volatility and the Cross-Section of Stock Returns", unpublished paper, Stockholm School of Economics.
- Wu G. (2001), "The Determinants of Asymmetric Volatility", *Review of Financial Studies*, 14, 837-859.

Appendix 1 – Using robust regressions to generate a cross-sectional volatility estimates

Under the assumptions in Section 4.1, we can adopt the following model for the estimation of the cross-sectional volatility index,

$r_{it} = a_t + \varepsilon_{it}$

where $a_t = \beta_t F_t$ and ε_{it} are residuals with a mean of zero. We assume that for *t* fixed, the distribution of ε_{it} is contaminated with outliers. To deal with this problem, we estimate the intercept and the dispersion through a robust method. Following the M-estimation technique (see for example Huber and Ronchetti (2009)), the general form of the estimator is:

$$\left(\hat{a}_t, CSV_{t,w_t}\right) = \arg\min_{\left(a_t,\sigma_{\varepsilon t}\right)} \sum_{i=1}^{n} \rho\left(\frac{r_{it} - a_t}{\sigma_{\varepsilon t}}\right)$$

where the function $\rho(e)$ measures the contribution of each residual. In fact, the first derivative of $\rho(e)$ can be interpreted as a weighting function. The standard notation is

$$\rho'\left(e\right) = \psi\left(e\right) = w\left(e\right)e\ .$$

where $w(e)$ denotes the weighting function. The actual calculations are done through the weighted least squares method and a bias-efficient choice for $\rho(e)$. As an approximate solution to the minimization problem, we first find a robust estimate for a_t and the weights corresponding to the observations, and then we use those weights to estimate CSV_{t,w_t} which stands for the cross-sectional volatility index observed at time *t*. We follow the algorithm below:

a) Estimate the intercept a_t through as the sample average, i.e. start with equal weights as in the natural estimator provided in section 4.2.

b) At each iteration *k* calculate the residuals and the weights from the previous iteration.

c) Solve for the new weighted least squares estimates, $\hat{a}^{(k)} = w^{(k-1)}r$ where $w^{(k-1)}$ is a vector of weights.

d) Repeat until convergence. Denote the vector of final weights by *w*.

Finally, calculate the cross-sectional volatility index through

$$CSV_{t,w} = C\sqrt{\sum_{i=1}^{n} w_i \left(r_{it} - \sum_{i=1}^{n} w_i r_{it}\right)^2}\ .$$

where *C* is a calibration constant that depends on the particular function $\rho(e)$. *C* can be calibrated by removing the bias for a particular target distribution (e.g. the Gaussian distribution) for a given function $\rho(e)$. Because the final weight of a given stock at time *t* depends on whether the observed return is an outlier at *t*, it will vary across time; that is, it will not be static.