Dispersion trading: Empirical evidence from U.S. options markets

Cara M. Marshall *

Queens College of the City University of New York, United States

ARTICLE INFO

ABSTRACT

This paper develops empirical evidence on the viability of a form of volatility trading known as "dispersion trading." The results shed light on the efficiency with which U.S. options markets price volatility. Using end-of-day implied volatilities extracted from equity option prices for the stocks that comprise the S&P 500, the implied volatility of the S&P 500 is computed using a modification of the Markowitz variance equation. This Markowitz-implied volatility is then compared to the implied volatility of the S&P 500 extracted directly from index options on the S&P 500. These contemporaneous measures of implied volatility are then examined for exploitable discrepancies both with and without transaction costs. The study covers the period October 31, 2005 through November 1, 2007.

It is shown that, from a trader's perspective, index option implied volatility tended to be more often "rich" and component volatilities tended to be more often "cheap." Nevertheless, there were times when the opposite was true; suggesting that potential dispersion trades can run in either direction.

© 2009 Published by Elsevier Inc.

1. Introduction

Ever since the publication of the seminal Black–Scholes–Merton option pricing models, volatility[1] has been of growing interest to academics and practitioners alike. Key aspects of volatility have been examined and the findings have been used to help explain various anomalies associated with observed option pricing and to suggest possible trading strategies. Some of the better known dimensions of volatility that have

* Department of Economics, Queens College, 300K Powdermaker Hall, 65-30 Kissena Boulevard, Flushing, NY 11367, United States.
Tel.: +1 718 997 5387; fax: +1 718 997 5466.
 E-mail address: Cara.Marshall@qc.cuny.edu.

[1] Throughout this paper, the term "volatility" will be used to mean the annualized standard deviation of the log price return. This is the customary usage in the options literature.

been studied include (1) the tendency of forward-looking implied volatility to exceed backward-looking realized volatility,[2] (2) the tendency of volatility to be mean reverting,[3] (3) the distribution of volatility,[4] (4) the tendency of volatility to be negatively correlated with the price level of the underlying,[5] (5) the volatility smile or skew[6], and (6) the term structure of volatility. Most recently, a measure of the implied volatility of the S&P 500, known as the VIX, has gained stature as the market's "fear gauge" and is quoted daily in the business press.[7]

This paper contributes to the literature on volatility by developing empirical evidence on a relatively new form of volatility trading, known as "dispersion trading," that is practiced by some quantitatively sophisticated hedge funds and by proprietary trading desks of some banks. In the process, it sheds light on the relative efficiency with which U.S. options markets price index volatility and index component volatilities.

Dispersion trading, as most often defined, involves trading index volatility against index component volatilities. This strategy can make sense if the index volatility implied by index options differs from the index volatility implied by the Markowitz variance equation, as derived from index component volatilities. Those who practice dispersion trading claim that index volatility tends to be more often rich while component volatilities tend to be more often cheap thus suggesting that index options should be written and equity options on the index components should be purchased.

In this study, the end-of-day implied volatility of the S&P 500 derived from index options (referred to here as "index-option-implied volatility" or IOIV) is compared to the contemporaneous end-of-day implied volatility of the S&P 500 derived by applying a modified version of the Markowitz variance equation to the implied volatilities of the index components as extracted from equity options on those components. The latter implied volatility is termed the "Markowitz-implied volatility" or MIV. The study covers the period October 31, 2005 through November 1, 2007. A number of aspects of IOIV and MIV are examined, including the degree to which the two measures of index volatility are correlated, the nature of the distribution of their differences, the frequency and size of their discrepancies, and how the discrepancies are related to the level and to the volatility of index volatility. Violations of the strict "law of one price" (LOP) do not imply trading opportunities unless transaction costs are reasonably estimated and factored into the analysis. For this reason, the transaction costs associated with exploiting mispriced volatility were estimated. In this study, transaction costs are measured in terms of volatility points.

The remainder of this paper is structured as follows. The next section offers a brief description of the dispersion trading strategy together with the claims made by practitioners and arguments that have been offered in the literature. The following section lays out the structure of the empirical tests performed. The final sections report results and conclusions.

2. Dispersion trading

It is difficult to determine the extent to which volatility strategies are used by hedge funds due to the proprietary nature of hedge fund activities and the general hedge fund obsession with secrecy. But it is clear that such strategies are popular and their use is growing. One particular strategy, called "volatility dispersion" or "dispersion trading," has been described in the literature by Nelken (2006). It generally involves selling volatility on a stock index and buying the volatilities of the index components, but usually not the reverse. Such strategies can be effected using equity options and index options or they can be effected using volatility swaps or variance swaps.

The same strategy has been discussed in both the business press and in quantitative research reports circulated by the derivatives industry. For example, Gangahar (2006), writing for the *Financial Times*, describes buying individual equity options (i.e., buying the volatilities of the index components) while simultaneously writing index options (i.e., selling the volatility of the index). He asserts that "if maximum

---

[2] See, for example, Christensen and Prabhala (1998) and Fleming (1998).
[3] See, for example, Goyal and Saretto (2007) and Lehman Brothers (2002).
[4] Many studies have shown that volatility is well-described by a lognormal distribution. See, for example, Dacorogna (2006).
[5] This property of volatility seems to have first been discovered by Black (1976, pages 177–181).
[6] See Ederington and Guan (2002).
[7] See, for example, Whaley (2000). See also Jiang and Tian (2007) and Carr and Wu (2007).

dispersion is realized, the strategy will make more money on the long options on the individual stocks than it loses on the short option position on the index…." In a technical report written by Mougeot (2007), and distributed to clients of Deutsche Bank in May of 2007, dispersion trading is defined as "a short position on index volatility and a long position in index constituent volatilities, which by construction results in being short average correlation."

When options are used to effect a dispersion trade, the trader would determine if the trade is warranted by comparing index volatility extracted from index options to the index volatility implied by the volatilities of the index components. The latter requires estimates of the return correlations among the various index components along with the appropriate weighting scheme. The individual component volatilities, correlations, and weights are then combined using the Markowitz variance equation to produce the Markowitz-implied volatility. The correlations employed in the analysis are typically historical correlations, which introduce some inexactness for any forward-looking trade execution.

Option-based dispersion trading strategies most often rely on at-the-money (ATM) options to compute the implied volatilities and to execute the strategy. The strategy is complicated by the fact that an option on a portfolio (e.g., an index option) is not the same thing as a portfolio of options. Further, the moneyness of options is continuously in flux as the prices of the various underlyings change. This necessitates frequent rebalancing of the option portfolio. For this reason, volatility swaps and variance swaps are often substituted to effect the strategy even though the volatilities themselves may be derived from option prices.[8]

Option-based dispersion trades typically take an "algorithmic approach" to trading options. Algorithmic trading is automated, usually electronic, trading that executes buy and sell orders in a rapid-fire fashion according to a set of well-defined trading rules. It typically involves many trades in small lot sizes so as to minimize the impact of the trading on market prices.

The dispersion strategy is most often employed when the premiums (and therefore the implied volatilities) on index options are high relative to the premiums and volatilities of the index components. The trader intends to profit when either (1) implied volatilities come back into equilibrium, through an increase in the volatilities of the components or a decrease in the volatility of the index or some combination thereof; or (2) the options expire and more is earned on the options written than is lost on the options purchased.

Because index volatility is a function of the individual component volatilities but also of the levels of the return correlations among the index components, some have argued that dispersion trading can just as easily be viewed as a form of correlation trading. That is, the claimed tendency of index volatility implied by index options to exceed Markowitz-implied volatility derived from component volatilities may reflect a "correlation risk premium." See, for example, Mougeot (2007). For this reason, it is often difficult to distinguish volatility-based trading from correlation-based trading. Others who have looked at the correlation risk premium include Driessen, Maenhout, and Vilkov (2006) and Coval and Shumway (2001).

3. Empirical analysis

In this section, the methodology employed in computing the various measures of volatility covered by this study is described. The study period was October 31, 2005 through November 1, 2007. It was necessary to employ a more a computationally efficient version of the Markowitz variance equation. The one employed, referred to here as the "modified Markowitz equation," is applicable when the portfolio of interest is structured to replicate the benchmark index. Importantly, the modified Markowitz equation produces the standard deviation of the portfolio's return (i.e., its volatility) directly rather than the variance of the portfolio's return. In this study, the inputs to the equation include (1) the implied volatilities from 30-day ATM equity options for all the stocks in the S&P 500, appropriately weighted by the respective stocks' representations in the index, and (2) the historical return correlations between the individual stocks and the index. The output from the modified Markowitz equation is referred to here as the "Markowitz-implied volatility" (or MIV) for the S&P 500. Separately, the implied volatility of the S&P 500

---

[8] For a discussion of variance and volatility swaps, see Demeterfi, Derman, Kamal, and Zou (1999) and Brockhause and Long (2000).

index is extracted directly from options on the index. This is referred to here as the "index-option-implied   142
volatility" (or IOIV). The two measures of volatility, MIV and IOIV, are then compared, day-by-day over the   143
two-year study period. The goal is, in part, to test the claim made by dispersion traders that index volatility   144
tends to be rich and component volatilities tend to be cheap. That is, the null hypothesis is that, on average,   145
MIV = IOIV versus the alternative hypothesis that, on average, IOIV > MIV. If the null hypothesis can be   146
rejected, deviations from the law of one price (LOP) exist. However, such deviations, if they exist, do not   147
imply profitable dispersion trading opportunities unless they are sufficiently large to cover transactions   148
costs. It is therefore necessary to develop a reasonable, preferably conservative, estimate of transaction cost   149
for strategy execution. Because a dispersion trade can be in either direction, a violation of the LOP, and a   150
trading opportunity, exists whenever MIV > (IOIV + transaction costs) or MIV < (IOIV − transaction costs).   151

The key assumptions in the analysis are (1) that the S&P 500 is sufficiently diversified that unsystematic   152
risk is nonexistent, (2) that historical return correlations are reasonable proxies for current return   153
correlations; (3) that the method of dealing with missing data does not bias the results; and (4) that the   154
estimate of transaction costs is reasonable.   155

3.1. Modified Markowitz equation   156

As shown by Markowitz (1952), the variance of a portfolio's (in this case the S&P 500 index's) return   157
(σ_m^2) is given by: 158

$$\sigma_m^2 = \sum_{i=1}^{500} \sum_{j=1}^{500} w_i w_j \sigma_{i,j} \tag{1}$$

where w_i denotes the weight on stock i within the index, w_j denotes the weight on stock j within the index,   160
and $\sigma_{i,j}$ denotes the covariance of the return on stock i with the return on stock j.[9] If expressed in terms of   161
correlations rather than covariances, Markowitz's variance equation is given by: 162

$$\sigma_m^2 = \sum_{i=1}^{500} \sum_{j=1}^{500} w_i w_j \sigma_i \sigma_j \rho_{i,j} \tag{2}$$

where σ_i and σ_j denote the standard deviations of the returns for stocks i and j, respectively (i.e.,   164
component volatilities) and $\rho_{i,j}$ denotes the return correlation between stock i and stock j.[10]   165
By definition, standard deviation (i.e., index volatility) is given by Eq. (3)   166

$$\sigma_m = \sqrt{\sigma_m^2} \tag{3}$$
168

Using Eqs. (2) and (3) to calculate the index variance and the index volatility would be cumbersome.   169
The correlation matrix would contain 250,000 correlations. Eliminating the 500 unit diagonal points and   170
exploiting the symmetrical property of the correlation matrix would reduce the computations to a still   171
cumbersome 124,750 correlations.   172

However, a portfolio structured to replicate an index will, by construction, have zero unsystematic risk   173
relative to the index.[11] That is, the correlation between the replicating portfolio and the index will be 1.0.   174
This allows us to calculate the volatility of the portfolio by focusing only on its systematic components   175

---

[9] Markowitz's variance equation, as originally formulated, is based on total returns. However, this study uses price returns. That is, dividends are not included. This is necessitated by the fact that exchange-traded options in the United States are not dividend protected. All references to "return" in this study are references to "price return."

[10] The Markowitz variance equation routinely employs the portfolio weights, the standard deviations of individual stock returns, and the correlations of paired returns to obtain the variance and standard deviation of the portfolio's return. It is rarely made explicit, however, what compounding assumption is employed in calculating the stock returns from which the standard deviations and correlations are computed. In all option valuation modeling, on the other hand, the standard deviations of stock returns (commonly called the vols) are measured from continuous returns. For this reason, all of the modeling incorporated in this study employs continuous returns (i.e., log returns).

[11] This is equivalent to treating the index as the entire investment universe. It is an appropriate treatment when the portfolio is structured to replicate the index for purposes of trading the portfolio against the index.

(since all the risk is systematic). As shown by Marshall (2008a), the standard deviation (i.e., volatility) of the index replicating portfolio can be obtained directly from Eq. (4)

$$\sigma_m = \sum_{i=1}^{500} w_i \sigma_i \rho_{i,m} \tag{4}$$

where $\rho_{i,m}$ is the return correlation for stock i with the market index. Eq. (4) is the "modified Markowitz equation" and the result of this calculation is referred to as the MIV noted earlier. If the LOP holds, then the MIV would be expected to match the volatility extracted directly from S&P 500 index options referred to here as the IOIV.[12] To the degree that the two volatilities differ, there are, potentially, alpha opportunities in the form of exploitable inefficiencies in the pricing of volatilities.

3.2. Data selection and acquisition

The application of the modified Markowitz equation required the following data: (1) the weightings of the various S&P 500 components on each day of the study period (these change, however slightly, each day as a natural consequence of the index's construction); (2) the historical correlations of the component returns with the S&P 500 return (these were derived from 3-years of rolling monthly returns); and (3) the individual stocks' implied volatilities derived from equity options on each day of the study period. The entire analysis that follows was done twice (once using calls and once using puts).

The end-of-day daily weightings of the 500 stocks that made up the S&P 500 were downloaded from Bloomberg for each day of the study period. It was necessary to download 505 separate days' weightings.

In order to generate the three years of rolling correlations, it was necessary to secure three years of price data prior to the beginning of the study period in addition to the two years of study-period data. The three years prior to the two-year study period is referred to as the pre-study period. Thus, five years of price data were required to generate the correlations. This covered the period October 30, 2002 through November 1, 2007. Over the course of the two-year study period, a number of stocks went into and a number of stocks came out of the S&P 500. As a result, there were a total of 586 stocks employed in the study, but only the relevant 500 were used on any given day. The end-of-day prices for the 586 stocks and for the S&P 500 were downloaded from Bloomberg. These prices were split adjusted but not dividend adjusted.[13]

From the daily price data, monthly returns were generated as the natural log of the ratio of the stock price on day t to the stock price 21 trading days prior (day t–21).[14] The same was done for the index values. The procedural details are worth elaborating.

As per the timeline in Fig. 1, the first day of the study period is day 1 and the last day is day 505. The first day of the pre-study period is day −755 and the last day would be day 0. Because a month is defined as 21 trading days, the first day for which a monthly return could be computed is day −734 (November 29, 2002). The general calculation is as follows:

$$MR_t = ln\left(\frac{S_t}{S_{t-21}}\right) \tag{5}$$

where MR_t denotes the monthly return as measured on day t. For day −734, this would be $MR_{-734} = ln\left(\frac{S_{-734}}{S_{-755}}\right)$. To compute the monthly return for the next day, day −733, each subscript is incremented by 1. Thus, $MR_{-733} = ln\left(\frac{S_{-733}}{S_{-754}}\right)$. This is repeated every day in a rolling fashion through both the pre-study period and the study period, ending with: $MR_{505} = ln\left(\frac{S_{505}}{S_{484}}\right)$. Thirty six of these paired monthly returns (one stock at a time with the S&P 500) were then used to generate a return correlation. On

---

[12] This statement assumes that the historical return correlations are good proxies for current return correlations.

[13] Exchange-traded options in the United States are split adjusted but not dividend adjusted. It was important that the price data employed in this study be consistent with standard practice in the options markets.

[14] In the United States, a trading year averages 252 days which is conveniently broken up into 21 trading day intervals representing trading months. Because of weekends, holidays, and monthly calendar day variations, trading months do not correspond precisely to calendar months.
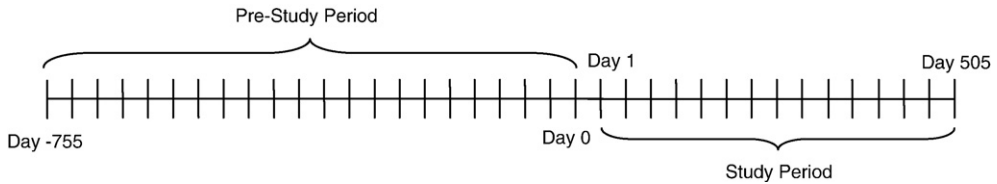

Fig. 1. Timeline.

the following day, the first return observations were dropped and the next return observations were   214
added. This was done for each stock separately.   215

The implied volatilities of each stock for each of the 586 stocks were pulled from Bloomberg for each of   216
the 505 days of the study period. Bloomberg obtains these implied volatilities from synthesized 30-day   217
ATM equity options based on the midpoint of the bid–ask spread.[15] The implied volatilities were gathered   218
separately for calls and puts. The end-of-day implied volatilities for the index, derived from synthesized 30-   219
day ATM index options (underlying ticker symbol SPX), were also downloaded from Bloomberg for each   220
day of the study period.[16] This too was done separately for calls and puts.   221

This then completed the data acquisition component of the study.   222

3.3. Missing data   223

In this study, two distinct types of missing data problems were encountered: (1) missing stock price   224
data; and (2) missing option-implied volatility data.   225

With respect to historical stock price data, on average, on any given day, data was missing for 10 stocks.   226
This ranged from a minimum of six stocks to a maximum of fifteen.   227

With respect to missing implied volatility data for individual stocks, there was, on average, on any given   228
day, data missing for nineteen stocks. The minimum was four and the maximum was twenty-eight.   229

While there was missing price and volatility data for individual stocks, there was no missing data for the   230
end-of-day value of the S&P 500 or for the implied volatilities of the S&P 500.   231

For any day that either price or volatility data was missing for a stock, that stock was excluded from the   232
MIV computation. The remaining components of the index were re-weighted by scaling the weights in   233
such a fashion as to maintain proportionality and achieve a total weighting of 100%. There was no reason to   234
believe that the stocks associated with missing price or volatility data had either greater or lesser monthly   235
returns or greater or lesser volatilities than the stocks for which data was not missing. Further, an   236
examination of the specific stocks for which data was missing revealed no clustering by industry.   237

4. Results   238

4.1. Preliminary results in the absence of transaction costs   239

Having fully fleshed out for each of the 505 trading days of the study period (1) the adjusted weightings,   240
(2) the individual stocks' call and put implied volatilities, and (3) the historical return correlations, the   241
modified Markowitz equation was applied and the MIV was calculated. This was straightforward and   242
resulted in two end-of-day MIVs for each day, one generated from calls and one generated from puts. These   243
could then be compared to the same-day end-of-day index-option-implied volatilities (IOIVs) for each of   244
the 505 days of the study period.   245

---

[15] These options are synthesized by Bloomberg by normalizing the expiry and the moneyness through a partially non-linear interpolation. Bloomberg extracts implied volatilities using a Black–Scholes Model for European options, a Roll–Geske Model for American-type options if there is one dividend, and a Trinomial Model if there is more than one dividend.

[16] The same methodology is used to extract the implied volatilities from index options as for equity options.
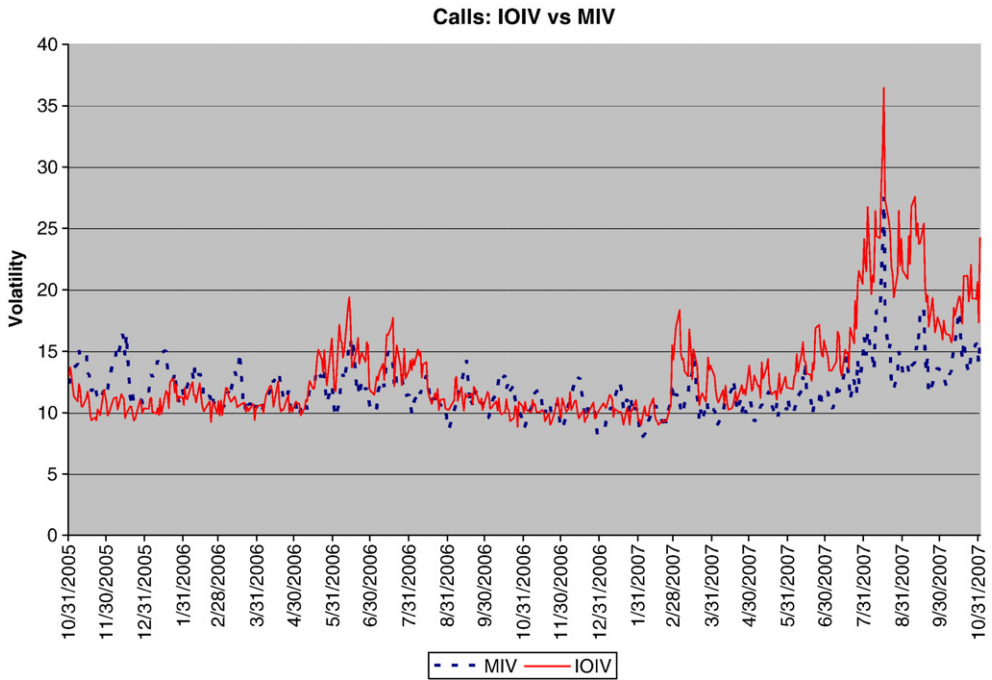
Fig. 2. Day-by-day IOIV versus MIV for calls.

Fig. 2 plots the day-by-day behavior of the MIV and IOIV measures for calls and Fig. 3 plots the day-by-day behavior of these same measures for puts. The Pearson product moment correlation between IOIV and MIV for calls was 0.68 and for puts it was 0.61. Summary statistics for MIV and for IOIV, as well as the difference between the two, for both calls and put are provided in Table 1.

Focusing first on the results for the calls, the IOIV exceeded the MIV on 312 days and the MIV exceeded the IOIV on 193 days. The mean IOIV exceeded the mean MIV by 1.21 vol points (13.16 versus 11.95). This is 5.03% of the call IOIV. That is, the IOIV exceeded the MIV by 5.03% of the IOIV on average. Similar results were found for puts, but the differences between the IOIV and MIV were less extreme. IOIV exceeded MIV on 270 days and MIV exceeded IOIV on 235 days. The mean IOIV exceeded the mean MIV by 0.76 vol points (12.80 versus 12.04). This is 1.40% of the put IOIV.

The results above are consistent with the often made claim by dispersion traders that index volatility (measured here by IOIV) tends to be rich and component volatilities (measured here by MIV) tend to be cheap. That is $IOIV_{average} > MIV_{average}$. This is equivalent to the proposition that $DIF_{average} > 0$ where $DIF_{average} = IOIV_{average} - MIV_{average}$.

To test this more formally, the distributional properties of the DIF_t series, where $DIF_t = IOIV_t - MIV_t$, were examined more closely. The Kolmogorov–Smirnov (K–S) goodness-of-fit test revealed that the DIF_t series was best described as lognormal for both calls and puts.[17] This finding is consistent with earlier research.[18] Additionally, a test for iid revealed significant first order serial correlation when the differencing interval is one day. This too was true for both calls and for puts. For this reason, it was not possible to appeal to the CLT to justify normality for $DIF_{average}$.[19] Consequently a bootstrapping approach

[17] Because lognormal distributions cannot go below 0, but the difference between two lognormal distributions can go below zero, a constant was added to all observations on the difference between IOIV and MIV to shift the values up sufficiently that they were all positive before applying the K–S test for lognormality.

[18] Mitchell (1968) and Mehta, Molisch, Wu, and Zhang (2006).

[19] See Law and Kelton (2000) and Godwin and Zaremba (1961).
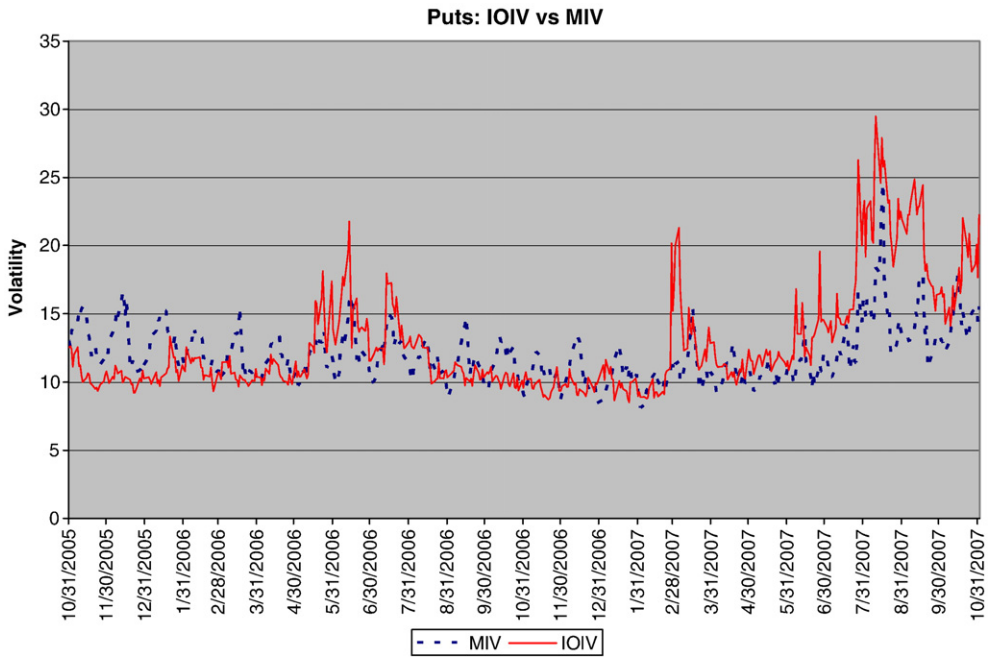
Puts: IOIV vs MIV

Fig. 3. Day-by-day IOIV versus MIV for puts.

was employed to generate 1000 simulated observations on $DIF_{average}$. Each of these simulated averages was   266
based on 505 random draws, with replacement, from the actual DIF_t series. [20] This was done for calls and   267
repeated for puts. The K–S test confirmed the normality of the $DIF_{average}$ for both calls and puts.   268

Based on the mean and the standard deviation of the bootstrapped $DIF_{average}$, we can reject the null   269
hypothesis that $DIF_{average} = 0$ in favor of the alternative hypothesis that the $DIF_{average} > 0$ for both calls and   270
puts at an extraordinarily high level of significance (0.000001).   271

Based on these results, we can conclude that, at least with respect to the study period, the index option   272
implied volatility was, on average, rich and component implied volatilities were, on average, cheap.   273

In industry literature[21] the difference between the means of the two volatility measures (1.21 vol points   274
for calls and 0.76 vol points for puts) is sometimes described as a "correlation risk premium." The logic   275
behind this description is that the MIV measure of index volatility relies on the historical return cor-   276
relations. It is well known that during periods of market stress correlations tend to rise. Thus, one might   277
over-price index volatility (IOIV) as a way of compensating traders for the risk of a correlation spike.   278

4.2. Transaction costs   279

Two prices can differ without giving rise to an arbitrage opportunity if the difference in the prices is   280
insufficient to cover the transaction costs associated with executing the strategy intended to exploit that   281
difference. Transaction costs include three separately identifiable components: (1) the bid–ask spread,   282
(2) the commission (if it has to be paid), and (3) the market impact cost. If the methodology of the   283

---

[20] In statistics, bootstrapping is the practice of estimating the properties of an estimator (such as a variance) by repeated sampling from an approximating distribution. The most common choice for the approximating distribution is the observed empirical distribution (in the present case the 505 observations on DIF). The procedure involves random sampling with replacement from the original dataset the same number of times as the size of the original data set. The principal advantage of bootstrapping over other analytical methods is its simplicity. The principal disadvantage of the method is that it is expensive in terms of computational effort (a problem rendered largely moot by the advances in computer technology).

[21] See the Mougeot (2007) article cited earlier.

Table 1
Summary statistics.

	Calls				Puts			
	IOIV	MIV	Difference		IOIV	MIV	Difference	
			Vol points	% of IOIV			Vol points	% of IOIV
Mean	13.16	11.95	1.21	5.03	12.80	12.04	0.76	1.40
Median	11.60	11.53	0.69	6.25	11.24	11.66	0.28	2.57
Standard deviation	4.10	2.11	3.09		3.91	2.03	3.13	
Standard error of the mean			0.1374				0.1391	
# of standard errors (Z score)			8.804				5.487	
Min absolute value			0.00				0.01	
Max absolute value			13.51				11.27	
# cases IOIV > MIV			312				270	
# cases MIV > IOIV			193				235	
Correlation between IOIV and MIV			0.68				0.61	

dispersion traders' approach necessitates periodic rebalancing of the portfolio, a rebalancing cost will also be incurred. This rebalancing cost has elements of all three types of the aforesaid transaction costs.

The principal transaction cost associated with option-based dispersion trading is the bid–ask spread that the trader must absorb when entering and subsequently exiting a position. Professional dispersion traders, which include some hedge funds, generally negotiate preferential commission schedules which can take a variety of forms. For example, the commission can take the form of a flat monthly fee, irrespective of the volume of trading, in which case the commission fee can be viewed as a fixed cost. Alternatively, the commission can be a function of the monthly volume such that the "per contract" commission gets smaller as the monthly volume gets larger. The third component of transaction cost, known in the academic world as "market impact cost," but known among arbitrageurs as "slippage," is the impact on market prices brought about by attempts to trade "in size." That is, if the dispersion trader attempts to trade more contracts than are currently available at the bid or ask, the trader will transact some of the contracts at prices outside the current bid and ask (i.e., below the current bid or above the current ask). Because dispersion trading done with options often takes an algorithmic approach, the market impact cost is likely negligible. Indeed, that is in large part the logic behind algorithmic trading. For these reasons, the dominant component of the transaction cost incurred by dispersion traders, using options to express their views, is the bid–ask spread.

In the world of options, the bid and ask prices can be stated in either monetary units or in volatility points. That is, suppose that a stock is trading at $50, there are 30-day ATM calls on this stock that are bid at $1.75 and asked at $1.85. The bid and ask are quoted in dollars, i.e., monetary units.[22] Call the $0.10 difference between these two prices the "monetary bid–ask spread" or simply a "monetary spread." From the option prices, one can back out the bid–ask spread in terms of "volatility points." For example, assuming that the annual risk-free interest rate is 4% and the stock does not pay dividends, the implied bid volatility is 29.20% and the implied ask volatility is 30.96%.[23] That is, one can buy volatility at 30.96 vol points and one can sell volatility at 29.20 vol points. The bid–ask spread in terms of volatility points is then 1.76 (i.e., 1.76%). Call this the "volatility bid–ask spread" or simply the "volatility spread."

Trying to estimate a reasonable volatility spread is complicated by a number of factors. First, for most of the two years studied, the minimum permissible tick size, and therefore the minimum monetary spread, was $0.05 for options trading below $3.00 and $0.10 for options trading above $3.00.[24] Second, the liquidity of options can vary tremendously. Options on heavily traded stocks tend to be more liquid than options on less heavily traded stocks. More liquid options have tighter spreads. Further, options that are ATM tend to

---

[22] The bid and ask prices, and therefore the bid–ask spread, are quoted per share covered by the option. Each option actually covers 100 shares. It is standard practice in options trading to quote prices on a per share covered basis.
[23] These were extracted using a Black/Scholes (B/S) Model.
[24] On January 29, 2007 the options exchanges began a penny pilot program in which options on a few stocks were allowed to trade in ticks of $0.01. Thus allowing the bid–ask spread to shrink to $0.01. However, this only affected a few stocks during the relevant two-year period, and only during the latter part of the two-year period. It also produced a reduction in the "size" that dealers would quote.

be more liquid than both options that are in-the-money (ITM) and out-of-the-money (OTM). Third, a given   315
monetary spread can translate into very different volatility spreads.   316

The point of this is that it can be particularly difficult to develop reasonable estimates of the bid–ask spread   317
in terms of implied volatilities, and thus difficult to know, absent real-time data on bid and ask prices, how   318
large a volatility discrepancy is necessary in order to create a profitable dispersion trade. Nevertheless, one can   319
reason as follows: Over the relevant two-year period, the International Securities Exchange (the first   320
electronic options exchange and the one whose technology is credited with a tightening of spreads across all U.   321
S. options markets after it opened in 2000)[25] reported that the average bid–ask spread for all options series   322
(which includes ATM, OTM, and ITM options on all names traded for all expiration dates) was just under $0.20.   323
Since the present study uses ATM options for the more liquid names, which tend to have the tightest spreads,   324
one can take this $0.20 as the upper bound for the monetary bid–ask spread with some room to spare to cover   325
the other types of transaction costs noted above.   326

Two different approaches were used to translate $0.20 monetary spread into a volatility spread. In the   327
first approach, $0.10 was added to and $0.10 was deducted from the Bloomberg quoted end-of-day option   328
values to obtain end-of-day ask and bid prices, respectively. From these ask and bid prices, the ask and the   329
bid implied volatilities were backed out. The difference was then the bid–ask volatility spread for the   330
option at a monetary spread of $0.20. This value averaged 3.56 vol points when all the options were   331
included for all 505 days. The second approach employed a worst-case scenario in which the implied   332
volatility spread was derived for a single option based on a composite stock assumed to be priced at the   333
lowest average daily stock price. This produced a highest possible, or "worst-case," volatility spread of   334
4.09 vol points.[26] In the interest of drawing more conservative conclusions, this worst-case volatility   335
spread of 4.09 vol points was employed.   336

4.3. Results in the presence of transaction costs   337

The 4.09 vol point transaction cost can be viewed as a "hurdle volatility discrepancy." That is, the   338
dispersion trader must find a discrepancy between IOIV and MIV of at least this size before he/she can be   339
confident that a trade will produce a profit after transaction costs. Because the trade can be done either way   340
(i.e., buy component volatilities and sell index volatility or sell component volatilities and buy index   341
volatility), the criterion for a trade is MIV < (IOIV − 4.09) or MIV > (IOIV + 4.09); where both MIV and IOIV   342
are measured in volatility points.   343

In the case of call options, of the 505 days in the study, the total number of days (based on end-of-day   344
values only) in which a profitable opportunity occurred was 84. Of these, 73 were situations in which the index   345
options were rich and the equity options were cheap. That is MIV < (IOIV − 4.09). And 11 were situations in   346
which the index options were cheap and the equity options were rich. That is, MIV > (IOIV + 4.09).   347

In the case of put options, the total number of days in which a profitable opportunity presented itself   348
(based again on end-of-day volatilities) was 91. Of these, 71 are situations in which the index options were   349
rich and the equity options were cheap. That is MIV < (IOIV − 4.09). And 20 were situations in which the   350
index options were cheap and the equity options were rich. That is, MIV > (IOIV + 4.09).   351

Looking back at Figs. 2 and 3, it is clear that the levels of the volatility measures IOIV and MIV were   352
higher overall during the second year of the study period than during the first year of the study period. In   353
addition, the volatilities of IOIV and MIV were slightly higher during the second year than during first year.   354
That is to say, volatility levels changed more dramatically over short periods of time during the second year   355
than during the first year.   356

One would expect more arbitrage-like[27] dispersion trading opportunities to arise when implied   357
volatility is less stable than when it is more stable. A simple check on this proposition is to look at the   358
number of times the hurdle volatility discrepancy was exceeded during each of the two study-period years.   359

---

[25] See, for example, Simaan and Wu (2007).
[26] For those interested in the conversion procedure to obtain the worst-case volatility spread from the monetary spread, see
Marshall (2008b).
[27] "Arbitrage-like" refers to situations in which one buys the "cheap" and sells the "rich." However, as noted earlier, unless the risk
can be completely removed, which is rarely the case, the trade is not a "true arbitrage" in the academic sense of the term "arbitrage."
The term "arbitrage-like" is often used to describe trades that bear a surface similarity to arbitrage, but which are not completely
free of risk.

Focusing on the call options, at a hurdle volatility discrepancy of 4.09 vol points, a total of 13 trading　360
opportunities occurred during the first year while 71 trading opportunities occurred during the second　361
year. The results were similar for puts: there were 26 trading opportunities in the first year and 65 in the　362
second year.　363

4.4. Cluster analysis　364

These results suggested that a more thorough month-by-month cluster analysis would likely be　365
beneficial in an effort to judge the effect of volatility levels and the effect of volatility on the appearance of　366
alpha opportunities. Because Pearson's product moment correlation assumes normality and is a measure of　367
linear correlation only, Spearman's rank order correlation was used instead. This measure of correlation is　368
distribution free and does not assume that the underlying relationship between the variables is linear.　369

Two tests were conducted. In the first, the null hypothesis was that there is no correlation between the　370
frequency of alpha opportunities and the level of IOIV against the alternative hypothesis that the two　371
variables are correlated. In the second, the null hypothesis was that there is no correlation between the　372
frequency of alpha opportunities and the volatility of IOIV against the alternative hypothesis that the two　373
variables are correlated. This was done separately for the calls and the puts. Thus, there were four cluster　374
tests performed.　375

The test procedure was straight forward. For each of the 24 months of the study period, the number of　376
alpha opportunities was counted for the month after allowing for transactions costs. For each of the　377
24 months the average IOIV for the month was also calculated. Finally, for each of the 24 months the　378
volatility of the IOIV for the month was calculated. The months were then ranked from highest to lowest for　379
each of the three variables. Spearman's rho for the four correlation tests are reported in Table 2.[28] Pearson's　380
product moment correlation appears in parentheses for comparison, but is not used in judging statistical　381
significance.　382

For a sample size of 24 (i.e., 24 months in this case), the critical Spearman value for a test of the null　383
hypothesis at a 10% level of significance is 0.344, at a 5% level of significance it is 0.407, and at a 1% level of　384
significance it is 0.521.[29] These levels of significance are indicated in Table 2 by *, **, and ***, respectively.　385

It is clear from this cluster analysis that both the level of IOIV and the volatility of IOIV positively　386
influence the number of alpha opportunities. The relationship between alpha opportunities and IOIV level　387
is clearly much stronger than the relationship between alpha opportunities and IOIV volatility.　388

5. Conclusions　389

Not surprisingly, the number of profitable dispersion trading opportunities declines as the transaction　390
cost rises. However, even when employing a worst-case-scenario transaction cost of 4.09 vol points,　391
dispersion trading opportunities repeatedly appeared. Further, because this study only looked at　392
volatilities at one moment in the day (i.e., the end-of-day), it is virtually certain that many other　393
opportunities occurred at other times of the day. Importantly, the results for both the calls and the puts are　394
consistent with the claim by Nelken and others, that the implied volatilities of index options are generally　395
higher than the Markowitz-implied volatilities derivable from the volatilities of the components.　396

Because these results are not consistent with the Efficient Markets Hypothesis, one is inclined to ask　397
what the reasons for the apparent mispricing of volatilities might be. Several possible explanations　398
immediately come to mind. First, it is possible that transaction costs have been underestimated. For　399
example, no attempt was made to estimate or to explicitly include rebalancing costs.[30] This argument,　400

---

[28] Spearman's rho (ρ_s) is calculated by first defining a variable $d_t = R_{1,t} - R_{2,t}$ where $R_{1,t}$ is the rank of the tth observation on variable 1 and $R_{2,t}$ is the rank of the tth observation on variable 2. Then ρ_s is given by: $\rho_s = 1 - \frac{(6 \sum d_t^2)}{n \times (n^2 - 1)}$.

[29] Note that in the Spearman rank order correlation test, the degrees of freedom are two less than the sample size (that is $df = n - 2$). These critical values are taken from Ramsey (1989).

[30] Without tick-by-tick market data (i.e., extremely high frequency data) observed in the context of a live trading operation, it is not possible to ascertain how large such costs would be. These costs and complexities are the key reason why many volatility traders would, in practice, employ volatility or variance swaps to express their volatility views, even though their analysis was based on option data.

Table 2
Spearman rank order correlation for alpha opportunities.

	Calls		Puts	
	Spearman	Pearson	Spearman	Pearson
Level of IOIV	0.659***	0.916	0.557***	0.893
Vol of IOIV	0.353*	0.554	0.365*	0.201

Transaction cost: 4.09 vol points.

however, is weakened by the fact that a worst-case transaction cost scenario was employed. Second, the calculation of the MIV relies upon historical return correlations and projects those correlations forward. It is quite possible that the index option market is, on average, assuming higher return correlations going forward than have been experienced in the recent past. Alternatively, consistent with the discussion earlier in this paper, the index option market might be pricing in a "correlation risk premium." This latter argument would be consistent with the findings of many earlier studies that implied index volatility (a forward-looking measure) tends to exceed realized index volatility (a backward-looking measure).

Another possible explanation for ongoing dispersion trading opportunities is the sheer complexity of the execution. Unlike simple forms of arbitrage in which one trades a single asset in one market against a single asset in another market, dispersion trading of the type described in this study is extraordinarily complex. It requires access to real-time data for thousands of options, an ability to crunch the numbers very quickly, and an ability to engage in rapid-fire execution of the buy and sell orders for the various components. Any delays in processing the numbers or in executing the orders will introduce risk because the implied volatilities are continuously changing. To be persuaded to bear these risks, traders would be expected to demand a greater volatility discrepancy than simply one large enough to cover transaction costs. Additionally, because algorithmic trading is still a relatively new field, and even less so when applied to dispersion trading, it is likely that an insufficient number of traders have sought to exploit the mispricing opportunities to completely eliminate them.

Finally, the evidence offered by the cluster analysis indicates, as one would expect, that the higher the level of volatility and the more unstable volatility is (i.e., the more volatile the volatility) the more frequently inefficiencies in the pricing of volatility arise.

Acknowledgments

References

Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 meetings of the American Statistical Association, Business and Economics Statistics Section* (pp. 177—181).
Brockhause, O., & Long, D. (2000). Volatility swaps made simple. *Risk*, *19*(1), 92—95.
Carr, P., & Wu, L. (2007). A tale of two indices. *Journal of Derivatives*, *13*(3), 13—29.
Christensen, B. J., & Prabhala, N. R. (1998). The relationship between implied and realized volatility. *Journal of Financial Economics*, *22*, 125—150.
Coval, J., & Shumway, T. (2001). Expected option returns. *Journal of Finance*, *56*, 983—1009.
Dacorogna, M. "Lessons learned from studying high frequency data in finance," presented at the Centre for Computational Finance and Economic Agents, University of Essex, October 26, 2006.
Demeterfi, K., Derman, E., Kamal, M., & Zou, J. (1999). A guide to volatility and variance swaps. *Journal of Derivatives*, *4*, 9—32.
Driessen, J., Maenhout, P., & Vilkov, G. (2006, December). Option-implied correlations and the price of correlation risk. *EFA 2005 Moscow Meetings*.
Ederington, L. H., & Guan, W. (2002, Winter). Why are those options smiling? *The Journal of Derivatives*, *10*(2), 9—34.
Fleming, J. (1998). The quality of market volatility forecast implied by S&P 500 index option prices. *Journal of Empirical Finance*, *5*, 317—345.
Gangahar, A. (2006, August 21). *Smart money on dispersion* : Financial Times.
Godwin, H. J., & Zaremba, S. K. (1961). *A central limit theorem for partly dependent variables* : Institute of Mathematical Statistics (get more complete citation).

Goyal, A. and Saretto, A. "Option returns and volatility mispricing." (February 2007). Working paper. Available at SSRN: http://ssrn.com/abstract=889947

Jiang, G. J., & Tian, Y. S. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. *Journal of Derivatives*, *14*(3), 35—60.

Law, A. M., & Kelton, W. D. (2000). *Simulation modeling and analysis*, Third edition. McGraw–Hill.

Lehman Brothers (2002, July 11). On mean reversion in implied volatility time series. *Global Foreign Exchange and Local Market Strategies*.

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, *7*(1), 77—91.

Marshall, C. M. (2008). Isolating the systematic and unsystematic components of a single stock's (or portfolio's) standard deviation. *Working paper, Queens College of the City University of New York*.

Marshall, C.M. (2008b). "Volatility trading: Hedge funds and the search for alpha (new challenges to the efficient markets hypothesis)," Doctoral Dissertation, Fordham University, September 2008.

Mehta, N. B., Molisch, A. F., Wu, J., & Zhang, J. (2006, June). *Approximating the sum of correlated lognormal or lognormal-rice random variablesIEEE International Conference on Communications (ICC), vol. 4.* (pp. 1605—1610) 8164–9547.

Mitchell, R. L. (1968, February). Permanence of the log-normal distribution. *Defense Technical Information Center*.

Mougeot, N. (2007, May 14). Hidden assets' investing series: Dispersion trading. *Technical report: Equity Derivatives Strategy Group. Deutsche Bank*.

Nelken, I. (2006). Variance swap volatility dispersion. *Derivatives Use, Trading & Regulation*, *11*(4), 334.

Ramsey, P. H. (1989). Critical values for Spearman's rank order correlation. *Journal of Educational Statistics*, *14*(3), 245—253.

Simaan, Y. E., & Wu, L. (2007, Winter). Price discovery in the U.S. stock options market. *Journal of Derivatives*, *15*(2), 20—38.

Whaley, R. (2000). The investor fear gauge. *Journal of Portfolio Management*, *26*, 12—17.