Financial Analysts Journal
Volume 66 • Number 1
©2010 CFA Institute

Implications for Asset Returns in the Implied Volatility Skew

James S. Doran and Kevin Krieger

This study examined the impact on future asset returns of information contained in the implied volatility skew. Future returns are linked to the discrepancy between call and put volatilities of at-the-money options and to the left side of the volatility skew, calculated as the difference between out-of-the-money and at-the-money puts. The findings discourage the use of skew-based measures for forecasting equity returns without fully parsing the skew into its most basic portions.

Equity and option markets are distinct entities. Securities within each market are traded at different times and locations. Despite the physical constraints, the markets are highly integrated, and information that is revealed in one market should be seamlessly reflected in the other. Black (1975) and others have suggested that informed traders will first go to the option markets to use the leverage of option contracts to achieve higher returns. Consistent with that notion, many recent studies have shown that information contained in option prices has implications for both the returns and the volatility of the equity markets.[1] These results suggest that "information spillover" occurs from the option markets to the equity markets. Our intent is to explain and summarize a subset of the information contained in option volatility and prices, specifically the information contained in the implied volatility skew and its relationship with future returns.

Bali and Hovakimian (2009), Cremers and Weinbaum (forthcoming), and Xing, Zhang, and Zhao (forthcoming) demonstrated that information contained in option volatility skews reveals insight into the future performance of underlying stocks. Bali and Hovakimian (2009) and Cremers and Weinbaum (forthcoming) suggested that stocks with large differences between call and put implied volatilities experience higher future monthly returns. Xing, Zhang, and Zhao (forthcoming) showed that stocks with greater negative skews tend to have lower future returns. On the surface, these results appear to be in direct conflict with each other. How can the option volatility skew imply both higher and lower future equity returns? This article is our attempt to provide a simple answer to that question by separating the volatility skew into parts and detailing how each part can be used to forecast future equity returns.

James S. Doran is the Bank of America Professor of Finance at Florida State University, Tallahassee. Kevin Krieger is assistant professor of finance at the University of Tulsa, Oklahoma.

Data and Methodology

We constructed our sample by using 4,161 companies in the CRSP database with regular shares (codes 10 and 11) that had option data available from OptionMetrics. We calculated five measures on the basis of portions of the implied volatility skew on the last trading day of each month from January 1996 to September 2008. We then investigated the relationship between these five skew measures and the monthly holding period returns of the underlying equities.[2] In addition, we excluded financial companies and utilities and further restricted our sample to only those companies whose shares were trading at $5 or more at the time the measures were calculated.[3] All options used in the analysis had to trade for $0.25 or more. We gathered book values from Standard & Poor's Compustat and calculated book-to-market ratios by following the Fama–French (1993) methodology. Finally, we examined whether our results were consistent among international markets by using different underlying instruments in constructing three portfolios: (1) a portfolio with 229 American Depositary Receipts (ADRs), (2) a portfolio with 17 international index exchange-traded funds (ETFs), and (3) a portfolio with 477 U.S. ETFs.

For all assets, we used the implied volatility of at-the-money (ATM), in-the-money (ITM), and out-of-the-money (OTM) calls and puts, as well as the volume of the contracts traded, to construct five measures.[4] Our first measure, "above minus below" (*AMB*), represents the difference between the mean implied volatility of the option pair whose strike prices are above the current underlying price and the mean implied volatility of the option pair whose strike prices are below the current underlying price. Specifically,

$$AMB = \left[\frac{\left(\sigma_{P,ITM} + \sigma_{C,OTM}\right) - \left(\sigma_{C,ITM} + \sigma_{P,OTM}\right)}{2}\right], \quad (1)$$

where σ is the implied volatility of the given option. We used options with moneyness levels closest to, but not exceeding, a strike-to-spot ratio of 0.95 to represent the OTM (ITM) volatility in the case of puts (calls); we used options with moneyness levels closest to, but not below, a strike-to-spot ratio of 1.05 to represent the ITM (OTM) volatility in the case of puts (calls). If the best matching option had a moneyness level below 0.80 or above 1.20, we deleted the observation. Data for all four options had to exist for *AMB* to be calculated.

The second and third measures capture the difference between the OTM and ATM volatilities (with OMA being shorthand for "out minus at") of calls and puts, respectively:

$$COMA = \sigma_{C,OTM} - \sigma_{C,ATM}, \quad (2)$$

and

$$POMA = \sigma_{P,OTM} - \sigma_{P,ATM}. \quad (3)$$

The ATM options that we used in our calculations had a strike-to-spot ratio closest to 1.00 on the company-date and could not have a ratio less than 0.95 or exceeding 1.05. Our final measures mimic the measures in Cremers and Weinbaum (forthcoming) and Xing, Zhang, and Zhao (forthcoming), respectively:[5]

$$CW = \sigma_{C,ATM} - \sigma_{P,ATM}, \quad (4)$$

and

$$ZZX = \sigma_{P,OTM} - \sigma_{C,ATM}. \quad (5)$$

CW and *ZZX* are simplified versions of the measures presented in the earlier studies because we did not initially weight by volume or open interest. We did, however, require each option to have nonzero trading volume.

Taken together, the five measures should capture all information contained in the cross-section of implied volatilities. **Exhibit 1** describes which part of the cross-section of implied volatilities, or implied volatility skew, each measure captures. Certain pairs of the measures can be combined to form a third measure. For example, *ZZX* plus *CW* equals *POMA*. These constructs allowed us to test for consistency among the measures.

Exhibit 1. Volatility Measures

Measure	Position of the Volatility Skew
AMB	Tails of the volatility skew
COMA	Right and middle side of the volatility skew for calls
POMA	Left and middle side of the volatility skew for puts
CW	Middle of the volatility skew
ZZX	Across the left of the put volatility skew and the middle of the call volatility skew

Notes: Exhibit 1 describes the portions of the implied volatility skew captured by our five measures. We assume that the implied volatility skew starts on the left with low strike prices and moves to higher strikes as the spot price remains constant.

After constructing each measure, we sorted companies into quintiles by the measures and calculated returns for the following month. We then rebalanced the portfolios on the last day of the following month, which allowed for a tractable and easily replicable portfolio implementation. For consistency, we used only options that expired in the second month following the trading day in question to construct the measures (e.g., those expiring in mid-February to measure end-of-December skew characteristics).[6]

We calculated monthly portfolio returns by quintile, as well as returns and Fama–French (1993) alphas for the zero-cost, long-high, and short-low quintile portfolios for each of the measures. Controlling for company characteristics, we formed double-sorted portfolios and performed Fama–MacBeth (1973) regressions on company-level data to establish a significant link between the measures and month-ahead returns. The Fama–MacBeth (1973) procedure estimates regression coefficients and tests their significance levels via a two-stage approach. First, coefficients are estimated cross-sectionally for each of the 153 monthly data points (the last trading day of every month from January 1996 through September 2008). Next, the 153 cross-sectional coefficients are averaged over time. These coefficients are then tested for significance on the basis of the standard error over time.

For robustness, we examined the impact of the implied skew measures in conjunction with eight company characteristic controls: *SIZE*, *BM*, *LIQ*, *PRICE*, *PSKEW*, *P1MRet*, *P12to2MRet*, and *P36to13MRet*. *SIZE* is the company's market capitalization in millions of dollars on the day of the observation. Constructed in the manner of Fama and French (1993), *BM* is the ratio of book value to market value. Following Amihud (2002) and Bali, Cakici, and Whitelaw (2009), *LIQ* is the measure of liquidity. *PRICE* is the company's share price at the time of measure construction. *PSKEW* is the total skewness measure of Bali et al. (2009). *P1MRet*, *P12to2MRet*, and *P36to13MRet* are the company's returns in the previous month, months –12 to –2, and months –36 to –13, respectively. The size, book-to-market, and momentum variables are standard controls. We incorporated liquidity, price, and total skewness as a proxy for physical skewness in light of evidence of a relationship between these variables and future returns. We were especially interested to see whether the power of physical skewness would diminish the impact of the implied volatility skew measures.

Empirical Results

Our empirical results include descriptive results, single-sorted portfolio results, double-sorted portfolio results, and Fama–MacBeth estimations that control for company characteristics, international and ETF portfolio results, performance over time, and transaction costs.

Descriptive Results. **Table 1** reports the summary statistics for the sample in two panels. Panel A reports the values for each measure on the basis of all available observations on the last trading day of each month over the sample period. Panel B reports the intersection of the sample such that all options were available to calculate all five measures of the companies. The values of our measures are reported as raw volatility differences.[7] The intersection sample consists of 62,076 company-months, with an average of 408 companies per month. Although this total is lower than those in prior studies, our sample conditions are notably more restrictive, particularly for the intersection sample. Irrespective of the sample size, the mean and median values of the *CW* and *ZZX* measures are consistent with values reported in Cremers and Weinbaum (forthcoming) and Xing, Zhang, and Zhao (forthcoming).[8]

The statistics in Panel A (Panel B) show that individual company put options have higher average implied volatility than do call options—*CW* is –0.74 percent (–0.78 percent)—and tend to have negative implied volatility smirks because *AMB* is –3.00 percent (–3.39 percent), the *COMA* mean is –1.15 percent (–1.27 percent), and the *POMA* mean is 2.37 percent (2.30 percent). The mean value of *ZZX* is 3.27 percent (3.08 percent); 74 percent of the value of the negative smirk comes from the difference between OTM and ATM puts, and the other 26 percent is a result of the difference in ATM volatilities of puts and calls. Given the inverse relationship between *ZZX* and future returns and the direct relationship between *CW* and future returns, we can infer whether the left-hand side of the put volatility skew, *POMA*, has any significant power to forecast future returns. If *POMA* does not have such power, then the power of the *ZZX* measure is likely derived only from the difference between the ATM put and call implied volatilities. If *POMA* has a positive relationship with future returns, then the interpretation of *ZZX* could be misleading.

Single-Sorted Results. **Table 2** reports the mean monthly portfolio returns of each quintile for each measure, as well as the raw and abnormal returns of going long the high-quintile portfolio and shorting the low-quintile portfolio. The results for each skew measure are reported in five panels to demonstrate robustness to the construction and sample selection. Panel A reports returns for the

Table 1. Summary Statistics, January 1996–September 2008

	Mean	5th Pct.	25th Pct.	Median	75th Pct.	95th Pct.	*N*
A. Full sample							
AMB	–3.00	–9.52	–5.01	–2.72	–0.71	2.51	109,967
COMA	–1.15	–5.45	–2.53	–1.18	0.05	3.17	99,375
POMA	2.37	–1.40	0.66	2.10	3.71	7.04	97,839
CW	–0.74	–7.31	–2.02	–0.52	0.78	5.24	126,750
ZZX	3.27	–2.36	1.18	2.92	4.92	9.84	98,114
B. N = 62,076 (intersection sample)							
AMB	–3.39	–9.99	–5.47	–3.10	–1.01	2.14	
COMA	–1.27	–5.08	–2.47	–1.22	–0.12	2.35	
POMA	2.30	–1.34	0.67	2.06	3.63	6.76	
CW	–0.78	–6.24	–1.83	–0.53	0.60	3.84	
ZZX	3.08	–2.44	1.10	2.83	4.77	9.33	

Notes: Table 1 shows the descriptive statistics for the full sample in Panel A; Panel B considers the intersection sample, in which all observations have available data to construct each measure. Implied volatilities are shown in percentages.

Table 2. Monthly Portfolio Returns, January 1996–September 2008

	Q1	Q2	Q3	Q4	Q5	Q5–Q1	FF α Q5–Q1
A. Full sample							
AMB	0.65	0.50	0.35	0.56	–0.03	–0.68***	–0.77***
COMA	0.64	0.61	0.58	0.57	0.42	–0.22	–0.27
POMA	–0.04	0.36	0.39	0.54	0.80	0.84***	0.48**
CW	–0.04	0.74	0.56	0.74	1.17	1.21***	1.40***
ZZX	0.58	0.43	0.45	0.43	0.20	–0.38**	–0.65***
B. N = 62,076 (intersection sample)							
AMB	0.61	0.49	0.27	0.46	0.02	–0.59***	–0.79***
COMA	0.31	0.50	0.39	0.25	0.39	0.08	–0.28
POMA	0.08	0.33	0.39	0.41	0.63	0.55***	0.45**
CW	–0.23	0.60	0.45	0.51	0.52	0.75***	0.99***
ZZX	0.54	0.31	0.34	0.40	0.06	–0.48**	–0.63***
C. NYSE							
AMB	1.07	0.66	0.53	0.72	0.29	–0.78***	–0.53***
COMA	1.02	0.84	0.57	0.61	0.66	–0.36**	–0.50***
POMA	0.43	0.48	0.64	0.78	0.91	0.48***	0.31*
CW	0.34	0.94	0.85	0.84	1.21	0.87***	1.08***
ZZX	0.78	0.68	0.69	0.66	0.46	–0.32*	–0.75***
D. Non-NYSE							
AMB	0.32	0.31	0.15	0.42	–0.37	–0.69***	–0.72***
COMA	0.27	0.47	0.52	0.48	0.18	0.09	–0.12
POMA	–0.41	0.07	0.09	0.25	0.64	1.05***	0.63***
CW	–0.46	0.56	0.14	0.43	1.32	1.78***	1.55***
ZZX	0.55	0.08	0.08	0.17	–0.19	–0.75**	–0.72***
E. Volume weighted							
AMB	0.78	0.60	0.47	0.55	0.08	–0.70***	0.59***
COMA	0.73	0.66	0.57	0.57	0.47	–0.26	–0.19
POMA	0.22	0.37	0.48	0.68	0.82	0.60***	0.45**
CW	0.10	0.69	0.59	0.69	1.30	1.20***	1.11***
ZZX	0.71	0.57	0.48	0.55	0.27	–0.44*	–0.49**

Note: Table 2 presents the mean monthly raw returns, in percentages, of quintiles following the end-of-month construction of skew measures.

*Significant at the 10 percent level.
**Significant at the 5 percent level.
***Significant at the 1 percent level.

full sample, Panel B reports returns for the intersection sample, Panel C reports returns for NYSE companies only, Panel D reports returns for non-NYSE companies, and Panel E reports returns on the basis of volume weighting each skew measure following Cremers and Weinbaum (forthcoming).[9]

We segmented the sample, and Quintile 1 (Quintile 5) contains the companies with the lowest (highest) values of each measure.[10] In other words, for *AMB* and *COMA*, Quintile 1 contains the most negatively skewed companies, and for *POMA* and *ZZX*, it contains the least negatively skewed companies. For *CW*, the companies with the most negative difference between call and put implied volatilities are found in Quintile 1.[11]

For *AMB* and *ZZX*, the portfolio in Quintile 1 typically generates the highest returns. The *AMB* Quintile 1 portfolio generates 65 bps a month for the full sample, and the returns decrease to –3 bps a month in Quintile 5. This outcome results in a –68 bp return for the *AMB* Q5–Q1 zero-cost portfolio and an abnormal return from the Fama–French (1993) four-factor model of –77 bps. For *ZZX*, the abnormal monthly returns are also negative. Thus, the returns to the long–short portfolios are not related to the market, size, value, or momentum premiums. The returns to the *CW* portfolios are in the opposite direction: The Q5–Q1 portfolio has positive and significant raw and abnormal returns of 1.21 percent and 1.40 percent per month for the full sample, respectively. The *CW* zero-cost portfolio also has the highest absolute raw and abnormal returns in all the subsample panels, with a maximum return of 1.78 percent and 1.55 percent, respectively (Panel E).

These results are initially perplexing because Quintile 1 of *AMB* has the most negatively skewed companies whereas *ZZX* Quintile 1 contains the least negatively skewed companies. Thus, the portfolio returns for the quintiles seem to contradict one another. Perhaps the results are driven by the manner in which *ZZX* or *AMB* is measured. The *AMB* results, however, do not appear to be a function of the ITM options because the *POMA* portfolio returns tend to increase among the quintiles and have positive Q5–Q1 raw and abnormal returns. These results imply that the more negatively skewed portfolios have higher future returns. Moreover, the *AMB* and *POMA* measures and their interpretation are consistent with each other and inconsistent with *ZZX*.

This inconsistency appears to be a result of *ZZX*'s incorporating both *CW* and *POMA*. For example, if *POMA* equals 2 percent and *CW* equals –2 percent, then *ZZX* equals 4 percent. If *CW* increases to –1 percent and *POMA* increases to 3 percent, *ZZX* undergoes no net change. The decrease in *CW* and the more negative skew both imply increases in future returns. The *ZZX* measure, however, would fail to detect either effect because it maintains the same value. Because the average returns of the *ZZX* Q5–Q1 portfolios are negative, the *CW* measure (or, more precisely, –*CW*) appears to impart more weight than does the *POMA* measure. This finding suggests that more influence on future returns lies in the difference between call and put volatilities than in the difference between OTM and ATM volatilities.

Compared with the other four measures, the quintile returns that use the *COMA* measure do not exhibit a monotonic relationship. Consequently, that only two of the five Q5–Q1 raw returns are significantly different from zero, with no significant alphas, is not surprising. This finding suggests that the right-hand side of the implied volatility skew contains little information.

Characteristic Controls. As is fairly standard, we controlled for company characteristics that might be correlated with the implied volatility skew measures to assess whether the information in the skew was truly distinct. We used two approaches. The first approach is to form a double-sorted portfolio that controls for the given characteristic and then for the implied skew measure; the second approach is a company analysis that uses a Fama–MacBeth (1973) methodology that controls for all characteristics and the skew measure jointly.

Under the double-sorted portfolio approach, we first sorted companies into quintiles on the basis of a company control like *SIZE*; then within each *SIZE* quintile, we sorted companies into quintiles on the basis of a skew measure.[12] Following Ang, Hodrick, Xing, and Zhang (2006) and controlling for company characteristics, we then averaged the skew quintile portfolios for each of the company control portfolios. **Table 3** presents the equally weighted, double-sorted portfolio results (controlling for company characteristics) and the high-minus-low skew quintile (Q5–Q1) returns, as well as the Fama–French (1993) four-factor alphas, on the basis of the intersection sample. In addition, χ^2 values are reported for tests of joint significance to determine whether the quintile portfolio alphas equal zero. This step is taken to confirm that the results are not driven by a single portfolio within each sorting.

The results show that irrespective of the company control, the *CW*, *POMA*, and *AMB* skew measures retain all the predictive power as in the single sorts. The Q5–Q1 portfolios have statistically significant raw and abnormal returns, and the tests of the joint alphas for the quintile portfolios are also significantly different from zero. In particular, sorting first on physical skew has a minimal impact on the Q5–Q1 implied skew portfolio returns, which suggests that the implied skew measure has a much stronger link to future returns than does the physical skew measure. In general, the economic significance of the *CW* measure is still the highest—just under 1 percent a month—followed by the *POMA* and *AMB* measures. By comparison, the predictive power of the *ZZX* measure diminishes after controlling for company characteristics, and only four Q5–Q1 alphas are significant. *COMA* retains little, if any, relationship to future returns in most double sorts, although some weak significance is found in the calculations first sorted by momentum.

Consequently, to confirm that the results were not driven by cross-correlations in the portfolios, we conducted Fama–MacBeth (1973) regressions on the returns of each company in the sample. We first estimated the coefficient for each measure in a univariate setting (models 1–5) and then incorporated the eight company controls (models 8–12). Finally, we included all the measures in the regressions with and without the company controls (models 6, 7, 13, and 14). Unfortunately, either *ZZX* or *CW* had to be excluded because *ZZX* = *POMA* – *CW*. The results for the intersection sample are presented in **Table 4**.

Table 3. Monthly Portfolio Returns after Controlling for Company Characteristics, January 1996–September 2008

	SIZE			*BM*			*LIQ*			*PRICE*		
	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$
AMB	–0.43**	–0.49**	10.99**	–0.64***	–0.51***	12.90**	–0.59***	–0.49**	11.77**	–0.55***	–0.50***	17.02***
COMA	0.10	0.21	1.28	0.38*	0.21	4.14	0.27	0.19	4.14	0.08	0.21	1.89
POMA	0.59***	0.49**	13.91***	0.79***	0.50***	13.27**	0.38**	0.37*	10.11**	0.51**	0.43**	15.91***
CW	0.52***	0.61***	18.52***	0.66***	0.49***	15.22***	0.83***	0.71***	19.40***	0.92***	0.62***	23.15***
ZZX	–0.28*	–0.37**	9.43*	–0.19	–0.25	4.89	–0.40**	–0.29*	9.19*	–0.28	–0.22	3.95

	PSKEW			*P1MRet*			*P12to2MRet*			*P36to13MRet*		
	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$	Q5–Q1	FF α Q5–Q1	$\chi^2(5)$
AMB	–0.40**	–0.33*	9.89**	–0.52***	–0.60***	16.84***	–0.43**	–0.44**	7.87*	–0.58***	–0.43**	11.23**
COMA	0.31*	0.22	4.67	0.27	0.21	7.99*	0.35*	0.22	11.14**	0.31*	0.19	7.89*
POMA	0.55***	0.57***	18.17***	0.87***	0.75***	21.78***	0.74***	0.69***	18.13***	0.64***	0.70***	10.08**
CW	0.62***	0.49***	16.09***	0.57***	0.55***	16.93***	0.76***	0.53***	18.99***	0.77***	0.59***	21.89***
ZZX	–0.49***	–0.32*	11.01**	–0.38*	–0.43**	10.98**	–0.43**	–0.45**	15.87***	–0.42**	–0.37**	13.89***

Notes: Table 3 presents the equally weighted monthly returns for the Q5–Q1 portfolios and the Fama–French alphas for the five skew measures after first controlling for company characteristics (for the intersection sample: N = 62,076). Returns are shown in percentages.

*Significant at the 10 percent level.
**Significant at the 5 percent level.
***Significant at the 1 percent level.

Table 4. Fama–MacBeth Regression Results, January 1996–September 2008

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
α	0.49	0.75	0.59	0.78	0.87	0.76	0.66	0.23	0.46	0.09	0.55	0.49	0.09	0.44
	(0.92)	(1.49)	(1.21)	(1.63)	(1.82)	(1.48)	(1.08)	(0.33)	(0.66)	(0.23)	(0.81)	(0.69)	(0.05)	(0.77)
AMB	–3.87					–3.00	–2.98	–2.89					0.02	–1.55
	(1.99)					(1.61)	(1.77)	(2.16)					(0.00)	(0.98)
COMA		1.78				2.08	3.11		0.24				0.44	1.55
		(0.71)				(0.71)	(1.31)		(0.09)				(0.14)	(1.01)
POMA			3.88			5.93	6.12			3.16			3.89	4.12
			(1.76)			(2.01)	(2.98)			(1.71)			(1.71)	(1.99)
CW				9.43		9.07					7.78		6.02	
				(7.52)		(4.54)					(5.77)		(3.90)	
ZZX					–7.53		–5.01					–2.10		–2.89
					(3.93)		(3.12)					(1.33)		(1.89)
SIZE								–0.51	–1.02	–0.65	–1.11	–1.33	–1.59	–1.44
								(0.03)	(0.22)	(0.20)	(0.87)	(0.76)	(0.68)	(0.40)
BM								0.67	0.71	0.58	0.66	0.77	0.66	0.68
								(2.27)	(2.39)	(1.89)	(2.00)	(2.31)	(2.41)	(2.23)
P1MRet								–1.56	–1.44	–0.91	–1.55	–1.33	–0.89	–1.32
								(1.49)	(1.49)	(1.02)	(1.54)	(1.41)	(0.98)	(1.22)
P12to2MRet								0.98	0.55	0.41	0.57	0.71	0.61	0.40
								(1.31)	(1.01)	(0.79)	(1.00)	(1.61)	(1.44)	(0.78)
P36to13MRet								–0.21	–0.15	–0.09	–0.13	–0.28	–0.09	–0.07
								(0.19)	(0.18)	(0.04)	(0.16)	(0.28)	(0.06)	(0.15)
PRICE								0.01	0.02	0.01	0.00	0.01	0.03	0.02
								(0.45)	(0.67)	(0.34)	(0.22)	(0.49)	(0.91)	(0.88)
LIQ								5.99	4.89	5.89	6.01	5.70	4.99	5.05
								(0.69)	(0.44)	(0.39)	(0.71)	(0.55)	(0.34)	(0.54)
PSKEW								–0.65	–0.03	–0.02	–0.03	–0.10	–0.07	–0.07
								(0.66)	(0.10)	(0.08)	(0.04)	(0.22)	(0.19)	(0.09)

Note: Table 4 presents the Fama–MacBeth regression results for the intersection sample (N = 62,076); t-statistics are in parentheses.

When estimated individually, the *POMA* and *CW* coefficients are positive and significant at the 5 percent level with or without the company controls. The *AMB* and *ZZX* coefficients are negative, but the *ZZX* coefficient is insignificant when all company controls are included. The coefficients on *COMA* are always insignificant. When estimated jointly, the *CW* and *POMA* measures remain significant in the Fama–MacBeth (1973) regressions, whereas the other measures are insignificant. *ZZX* is significant, but this result is misleading because the coefficient on *POMA* increases significantly and is contained within *ZZX*. When *POMA* is excluded, however, the significance on *ZZX* disappears. Our findings suggest that the differences between the ATM calls and puts and between the OTM and ATM puts both capture information for future equity returns.

In interpreting the economic implications of our results, note that a coefficient of 9.43 on the *CW* measure translates to a future monthly return of –7 bps for the average value of *CW* (–0.74 percent). In other words, if the stocks in a portfolio have an average difference between the call and put volatilities of –0.74 percent, then the month-ahead return will, on average, be 0.07 percent lower. When *CW* is equal to 5.24 percent (the 95th percentile value), the future monthly return will result in an additional 49 bps. By comparison, when *CW* is equal to –7.31 percent (the 5th percentile value), the future monthly return will be affected by an additional –68 bps a month. The difference between these two values, 117 bps, is similar to the difference for the raw returns of the long–short portfolio shown in Tables 2 and 3. The other four measures exhibit behavior similar to that of the prior results but are less economically significant. This outcome is evident in the size of the coefficient and the variation in the measures shown in Table 1.

The signs of the control variables are consistent with prior findings and tend to be statistically insignificant, except for *BM* and, to a lesser extent, *P1MRet*. The insignificance of *PSKEW* further highlights that the predictive power of skew is in the risk-neutral, not the physical, measure.

The negative coefficient on *ZZX* again emphasizes that proceeding with care is important when examining implied volatility skew measures. Because *ZZX* combines both the *POMA* and –*CW* measures and the effects are in contrasting directions, interpreting the effect of *ZZX* is difficult. As demonstrated by both *AMB* and *POMA*, the more negatively skewed the implied volatilities, the higher the future returns, which is the exact opposite of what *ZZX* implies. This conclusion also holds for *CW*, where the higher difference between the call and put implied volatilities suggests higher future returns. Thus, combining implied volatility skew measures ultimately muddles the impact and message of two distinct sources of information.

International Results and ETFs. Having demonstrated this effect for U.S. stocks, we then examined whether the skew measure relationship to future returns extends to other underlying instruments and to international markets. If the effect is present only for U.S. stocks, perhaps some idiosyncratic component unique to U.S. equities drives the results. If the results are apparent for ETFs and international stocks, this finding would support the conjecture that the results are a function of the relationship between the underlying instruments and the ability to trade options on those instruments.

To test whether the skew relationship holds for international markets, we constructed two equally weighted portfolios. The first portfolio comprised 229 ADRs, which allowed for a direct comparison between international stocks and the original U.S. equity sample. The second portfolio comprised 17 international index ETFs that were based on the iShares MSCI foreign index funds.[13] These funds are an ideal test of international performance because the objective of the funds is to track the performance of publicly traded securities in the European, Australasian, and Far Eastern markets. We also constructed a portfolio of 477 U.S. ETFs with traded options, which allowed us to compare the results among international and U.S. markets and among various underlying instruments.

Presented in **Table 5**, the results show the mean value of each measure, as well as the value and significance of the coefficient of each measure when it is included in a three-factor Fama–French (1993) regression for the one-month-ahead, equally weighted portfolio returns. Panel A shows the results for the ADR portfolio, Panel B shows the results for the international index ETF portfolio, and Panel C shows the results for the U.S. ETF portfolio. The values for the skew measures are comparable to those reported in Table 1 for the U.S. sample, especially for the *CW*, *AMB*, *POMA*, and *ZZX* measures. The size and direction of the coefficients on the Fama–French regression are similar to those reported in Table 4. For example, a 1 percent absolute increase in the value of the *CW* measure translates to an increase in the one-month-ahead return of 83 bps for the ADR portfolio, 51 bps for the international index ETF portfolio, and 36 bps for the U.S. ETF portfolio. These results suggest that the information contained in the implied volatility skew is present for both U.S. and international markets and for various underlying instruments.

Table 5. ETF and International Evidence, January 1996–September 2008

	AMB	*COMA*	*POMA*	*CW*	*ZZX*
A. ADR sample					
Mean	–1.37%	2.44%	3.87%	–1.69%	4.56%
β	–6.85	3.05	6.02	8.35	–8.48
t-Stat.	1.68	0.27	1.78	2.23	2.38
B. ETF sample					
Mean	–2.15%	–0.34%	1.76%	–1.47%	3.23%
β	–2.39	1.16	2.78	5.18	–3.60
t-Stat.	2.14	0.81	2.29	6.08	3.92
C. U.S. ETF sample					
Mean	–3.09%	–1.11%	2.21%	–1.44%	2.29%
β	–1.89	0.46	2.15	3.60	–2.76
t-Stat.	1.77	0.28	1.91	2.21	2.44

Note: Table 5 presents the mean values for the five volatility skew measures and the beta loadings and *t*-statistics from a three-factor Fama–French regression of one-month-ahead returns, with each of the five measures alternately added to the three-factor model.

Note that the predictive power of the skew measures for the ETFs has a direction similar to that of individual stocks but with a slightly lower magnitude. Why the results should hold for ETFs, given that, by construction, ETFs are more similar to mutual funds or an index, is initially unclear. Mutual funds, however, do not have options traded on them, and an index cannot be held, only replicated. Therefore, because ETFs are traded like individual stocks and have options traded on them, traders with informational advantages will go to the option markets first to take advantage of the leverage offered in the option contracts. Because the options are priced on the market value of the ETF, the implied volatility skew on the ETF represents the underlying risk of its individual stocks. Although ETFs do diversify away the idiosyncratic risk in the individual stocks that compose a given ETF, the individual skew measures are aggregated such that the skew effect is still present.[14] This outcome is reflected in a lower economic, albeit still significant, effect of the volatility skew on the future return on the ETF. This finding supports our conjecture that the results hold for individual stocks and ETFs because of the ability to trade and hold options and the underlying instruments, which allows investors to take advantage of information in either the ETF or individual stocks both in the United States and internationally.

Performance over Time and Transaction Costs. Finally, we assessed the real performance of using the skew measure to construct a portfolio. We examined the cumulative portfolio returns for the Q5–Q1 (Q1–Q5) zero-cost portfolios of *CW* (*AMB* and *ZZX*) and the performance of the S&P 500 Index over the sample period to assess the overall performance. We also showed how the portfolios performed in certain extreme months. Starting with an initial allocation of $10,000, **Figure 1** highlights the value of each portfolio over time. First, the *CW* portfolio outperformed the S&P 500 with considerably lower risk. That a portfolio can outperform the market on raw and risk-adjusted bases by using information from the option markets indicates that the flow of information seems to be one-sided. Second, the *ZZX* portfolio performed the worst of the zero-cost portfolios. This finding is not surprising given the opposing effects of *CW* and *AMB*. The Q1–Q5 *ZZX* portfolio yielded positive returns because the gains of *CW* dominated the negative effect of *AMB*. Again, this finding highlights the importance of using information from implied volatility skews carefully.

We included the returns to the long–short portfolios in March 2000, April 2000, August 2007, September 2008, and October 2008 to assess the performance of the skew measures when the market is under stress. These months represent the starting points for the bursting of the tech/internet bubble, as well as the collapse of the commodities and real estate markets. In particular, the months of September and October 2008 represent the highest increase in volatility and the largest price declines in a given month since the crash of 1929. Presented in **Table 6**, the results show the average monthly return to all stocks, the level of the Chicago Board Options Exchange Volatility Index (VIX) at the end of the month, the value of each skew measure preceding the month, and the one-month-ahead return.

Clearly, using the skew measure to form portfolios can result in positive returns irrespective of market conditions. For the five months selected, both the *CW* portfolio and the *AMB* portfolio always had positive returns. Moreover, apparently, the more volatile the market, the more pronounced the returns. For example, in October 2008, when the VIX was at more than 50 percent for most of the month, the dispersion in the *CW* measures among all stocks was more than double the average dispersion for a normal month in the sample. So, even though both the Q5 and Q1 portfolios had large negative returns, the Q5 portfolio outperformed the Q1 portfolio by 4.17 percent. In comparison, in August 2007, when volatility was lower, the dispersion in the *CW* measure was less than the average monthly dispersion and was accompanied by lower returns for the long–short portfolio in that month. This finding suggests that the returns to the

Figure 1. Portfolio Values, January 1996–September 2008

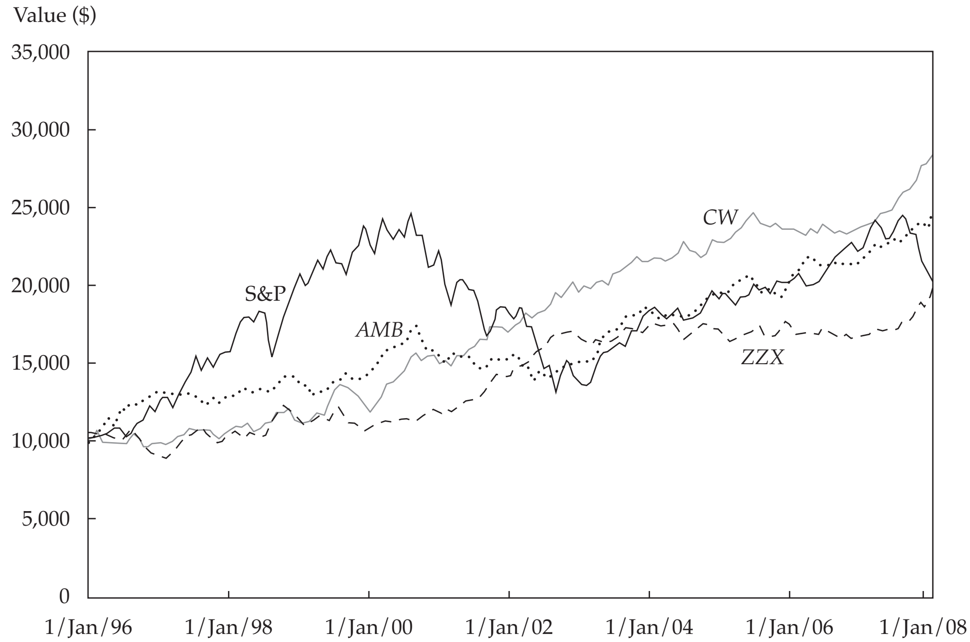

Notes: Figure 1 presents the portfolio values over time for strategies that implement the zero-cost *AMB*, *CW*, and *ZZX* portfolios. The *CW* portfolio is long the Q5 portfolio and short the Q1 portfolio. The *AMB* and *ZZX* portfolios are long the Q1 portfolio and short the Q5 portfolio. The positions are equally weighted and rebalanced monthly. A $10,000 portfolio based on S&P 500 returns is included for comparison purposes.

Table 6. Extreme-Month Portfolio Returns and Portfolio Returns with Transaction Costs

			AMB		*CW*		*ZZX*	
Month	Average Return	VIX	Mean Measure	Return	Mean Measure	Return	Mean Measure	Return
A. Extreme months								
Mar. 2000	2.294%	26.2%	–1.12%	–4.79%	–0.81%	6.00%	1.42%	–3.73%
Apr. 2000	–7.634	23.67	–0.90	–6.16	–0.60	2.23	1.38	1.31
Aug. 2007	–2.964	23.38	–3.02	–2.44	0.04	1.07	2.12	–0.79
Sept. 2008	–13.298	39.89	–4.64	–4.36	–0.95	2.12	3.86	2.03
Oct. 2008	–22.622	59.89	–5.34	–1.14	–4.62	4.17	8.57	–5.83

	AMB	*CW*	*ZZX*	S&P 500
B. Transaction costs				
Geometric return	6.69%	8.43%	4.60%	4.32%
Long position turnover	68.2	66.2	65.9	
Short position turnover	70.1	67.1	66.2	
Return after transaction costs	5.39	7.13	3.29	

Notes: Table 6 presents return results on the basis of long–short positions of the skew measures *AMB*, *CW*, and *ZZX*. The *CW* portfolio is long the Q5 portfolio and short the Q1 portfolio. The *AMB* and *ZZX* portfolios are long the Q1 portfolio and short the Q5 portfolio. Panel A shows results for selected extreme stock movement months. The average return of the market (based on the S&P 500), the VIX level at the beginning of each month, the mean of each skew measure, and the returns of a long–short position for each month are presented. Panel B reports the geometric means for the three long–short *AMB*, *CW*, and *ZZX* portfolios and the S&P 500 from January 1996 to September 2008. The monthly turnover represents the number of companies that leave or enter the portfolio as a percentage of the total number of companies.

skew portfolios are partially driven by high volatility in the market, but it also demonstrates that the long–short skew portfolios are an ideal hedge against market fluctuations.

Finally, to gain perspective on the economic magnitude of implementing the various skew-based positions, we calculated the annual geometric average returns for the long–short portfolios of *CW*, *AMB*, and *ZZX* and the S&P 500. We also wanted to account for the cost of monthly rebalancing and the monthly turnover, as well as the potential cost of short selling. We used the Ofek, Richardson, and Whitelaw (2004) estimate of the cost of short selling, which they referred to as the rebate rate, or the actual cost of borrowing a stock. They showed that the cost of shorting a stock, on average, is 0.30 percent. Although this estimate is an implicit cost, it is also an upward-biased estimate because it assumes that all stocks in the portfolio that are shorted experience the same short-selling cost.[15] For the cost of trading, we used 0.34 percent for the buy-initiated trade and 0.64 percent for the sell-initiated trade, similar to those costs reported in Keim and Madhavan (1998). Because our companies were larger and had more trading volume than those in Keim and Madhavan, the cost of rebalancing overstates the actual cost of trading.

The portfolio turnover statistics for the skew measures are reported in Table 6, together with the returns accounting for the cost of trading. The portfolio turnover is high, but that is expected because the information contained in the volatility skew represents a short-term opportunity. Before accounting for the transaction costs, we calculated the returns for the four portfolios. The mean annual returns are 6.69 percent, 8.43 percent, 4.60 percent, and 4.32 percent for the *AMB*, *CW*, and *ZZX* portfolios and the S&P 500, respectively. The long–short *AMB* and *CW* portfolios clearly outperformed the market. After accounting for the portfolio turnover, the cost of short selling, and all other potential trading costs, the *AMB* and *CW* portfolios still outperformed the S&P 500 by approximately 1.07 percent and 2.81 percent. After accounting for transaction costs, the *ZZX* portfolio underperformed the S&P 500; this finding is not surprising given that the measure incorporates two skew effects that work in opposite directions.

These results demonstrate two important aspects of trading on information within the implied volatility skew. Earning abnormal returns is possible but requires significant monthly portfolio turnover, which incurs annual transaction costs of more than 1 percent. Second—and consistent with the evidence presented earlier in the article—using the information in the volatility skew carefully is important. Otherwise, the potential returns to a portfolio formed on the wrong information can lead to significant underperformance.

Explaining the Volatility Skew Predictability

Although the empirical evidence reveals strong predictability in the *CW*, *AMB*, and *POMA* measures, explaining how the information from the option markets translates to future equity and ETF returns is worthwhile. We propose two mechanisms to explain this predictability.

First, traders who have superior information or talent will be motivated to go to the option markets first to exploit their superior information or talent by using the leverage in option contracts. Simply put, if a trader believes a stock is going to appreciate (depreciate) in price, then, as Black (1975) suggested, the trader is better off buying calls (puts) than the underlying stock because the call (put) option provides greater levered profits. In response to the increased demand for the option, the option market maker will increase the price of the option as more buy orders are entered. This action increases the disparity between the call and put volatilities, reflected in the skew measures, but the underlying asset's price remains unchanged. The underlying asset will appreciate/depreciate only as information about the underlying stock is revealed to equity traders, who presumably have less information. This mechanism is driven by the conjecture that most equity and option traders do not trade among markets and that the disparity is driven by open-buy purchases.

The second potential explanation concerns option traders who hedge short open positions. If a trader wishes to become delta neutral to limit a short call option position, increases in call volatilities relative to put volatilities (reflected in the *CW* measure) require the option trader to purchase more of the underlying stock to hedge the increase in delta. This action requires purchasing more of the underlying stock, which results in buy pressure on the underlying stock and driving up the price. How dynamic the delta-hedge rebalancing is translates to how quickly the underlying stock price reacts to the differences in the implied volatilities. This explanation also describes the predictability in the *POMA* measure. Increases in the steepness of the put volatility skew correspond to increasing pressure to hedge OTM puts relative to ATM puts. As the option approaches expiration, however, the OTM delta decreases at a quicker rate than the ATM delta. Delta-hedge traders then buy to cover their short underlying stock positions, which causes an increase in the buying pressure and resulting price appreciation of the stock.[16]

The difference in the second explanation is that the mechanism that drives the predictability comes from traders who trade in both equity and option markets and who buy the underlying stock to cover short, open option positions. The two explanations, however, are not mutually exclusive, and both can simultaneously explain this phenomenon. Specifically, implied volatility skew predictability can be a function of (1) open-buy purchases from segmented traders and (2) buy-to-cover positions from short-option traders who trade in both option and equity markets.

Conclusion

Using the option market to reveal information about future equity and ETF returns relies on the implicit assumption that information starts in the option market and flows to the equity market. The speed with which option-based information is incorporated into equity and ETF prices has a considerable impact on whether a portfolio that uses this information realizes significant returns. We demonstrated that several option-based measures of the volatility skew have strong predictive power in forecasting the future direction of the underlying asset price.

Specifically, we showed that the construction of measures that incorporate information from the volatility skew must be undertaken with care. Information is contained in various parts of the volatility skew, most notably in two places. First, future underlying returns are positively linked to the higher differences in volatilities between ATM calls and puts. Second, higher returns are positively related to more-negative skews as measured on the left-hand side of the volatility skew, between the OTM and ATM puts. We found the success of such strategies somewhat puzzling because these measures are easy to calculate and a portfolio that uses these strategies is easy to implement. We postulated two explanations for this phenomenon. First, it is a function of distinct traders of long positions who trade only in either the option or equity market whereby superior information is traded on initially in the option market. Second, it is also a result of delta-hedge traders who use the equity market to hedge short-option positions. Both mechanisms can explain why information appears to flow from the option market to the equity market.

We acknowledge the helpful comments and suggestions of Gurdip Bakshi, Jim Carson, Don Chance, Martijn Cremers, Sanjiv Das, Dave Humphrey, Ehud Ronn, Keith Vorkink, David Weinbaum, Xiaoyan Zhang, and participants at Florida State University and University of Tulsa seminars.

This article qualifies for 1 CE credit.

Notes

1. Manaster and Rendleman (1982) showed that stocks whose prices lie below their implied values from the option markets significantly outperform those stocks that are overvalued relative to implied price. Moreover, Pan and Poteshman (2006) demonstrated that option characteristics beyond implied price contain information relevant to future equity prices. Others have also documented the informational content in option markets, including Diavatopoulos, Doran, and Peterson (2008); Poteshman (2001); and Christensen and Prabhala (1998).
2. We also examined weekly returns and arrived at similar conclusions.
3. Following Battalio and Schultz (2006), we calculated returns by skipping the first day after the measurement of the volatility skew to eliminate market microstructure issues regarding options and equity returns.
4. We adjusted the implied volatilities for the early-exercise feature because all the options were U.S. Refer to OptionMetrics for a discussion of the methodology.
5. Our *CW* measure was motivated by the Cremers and Weinbaum (forthcoming) measure (from their Equation 7) and by the Bali and Hovakimian (2009) measure (*CVOL* – *PVOL*). Note that our *CW* measure captures not only potential information in future returns but also the deviations in put–call parity. Our *ZZX* measure was motivated by the measure found in Equation 1 of Xing, Zhang, and Zhao (forthcoming).
6. Our results did not materially change when we used options that expired in the next month or when we used the average skew measures for the next month and the following month, weighted by the volume of contracts traded.
7. We also replicated the analysis by using percentage differences (versus raw differences), whereby the raw differences are scaled by the ATM volatilities. The results did not change materially.
8. Cremers and Weinbaum (forthcoming) reported a mean value of –0.770, whereas Xing, Zhang, and Zhao (forthcoming) reported a mean value of 5.97.
9. We used options that expired between 10 and 60 days in the future and weighted each measure by the volume traded.
10. The portfolio returns are equally weighted. The value-weighted returns are higher, with magnitudes similar to those of Cremers and Weinbaum (forthcoming) and Xing, Zhang, and Zhao (forthcoming).
11. *COMA* captures the right-hand side of the volatility skew, and thus, the interpretation of "negatively skewed" is slightly misleading.
12. Following Amihud (2002), we calculated *LIQ* such that the most liquid companies were in Quintile 1. We also looked at alternative measures of physical skew, but the results were unchanged.

13. These funds include iShares MSCI Australia Index (EWA), iShares MSCI Canada Index (EWC), iShares MSCI Sweden Index (EWD), iShares MSCI Germany Index (EWG), iShares MSCI Hong Kong Index (EWH), iShares MSCI Italy Index (EWI), iShares MSCI Japan Index (EWJ), iShares MSCI Belgium Investable Market Index (EWK), iShares MSCI Switzerland Index (EWL), iShares MSCI Malaysia Index (EWM), iShares MSCI Netherlands Investable Market Index (EWN), iShares MSCI Austria Investable Market Index (EWO), iShares MSCI Spain Index (EWP), iShares MSCI France Index (EWQ), iShares MSCI Singapore Index (EWS), iShares MSCI United Kingdom Index (EWU), and iShares MSCI Mexico Investable Market Index (EWW).
14. The results are also independent of the tracking error of the ETF: the difference between the price of the ETF and the value of the ETF's underlying stocks.
15. Although some of the stocks may not be available for short selling, this possibility is incorporated in the cost estimate.
16. For example, assume an option on a $100 stock with 50 days to expiration, 30 percent ATM volatility, and 50 percent OTM volatility. With 20 days to expiration, the change in the ATM delta is from –0.45 to –0.47, whereas the change in the OTM delta is from –0.24 to –0.16.

References

Amihud, Y. 2002. "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets*, vol. 5, no. 1 (January):31–56.

Ang, A., R.J. Hodrick, Y. Xing, and X. Zhang. 2006. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, vol. 61, no. 1 (February):259–299.

Bali, T., and A. Hovakimian. 2009. "Volatility Spreads and Expected Stock Returns." *Management Science*, vol. 55, no. 11 (November):1797–1812.

Bali, T.G., N. Cakici, and R.F. Whitelaw. 2009. "Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns." NBER Working Paper 14804 (March).

Battalio, R., and P. Schultz. 2006. "Options and the Bubble." *Journal of Finance*, vol. 61, no. 5 (October):2071–2102.

Black, F. 1975. "Fact and Fantasy in the Use of Options." *Financial Analysts Journal*, vol. 31, no. 4 (July/August):36–41.

Christensen, B.J., and N.R. Prabhala. 1998. "The Relation between Implied and Realized Volatility." *Journal of Financial Economics*, vol. 50, no. 2 (November):125–150.

Cremers, M., and D. Weinbaum. Forthcoming. "Deviations from Put-Call Parity and Stock Return Predictability." *Journal of Financial and Quantitative Analysis*.

Diavatopoulos, C., J. Doran, and D. Peterson. 2008. "The Information Content in Implied Idiosyncratic Volatility and the Cross-Section of Stock Returns: Evidence from the Option Markets." *Journal of Futures Markets*, vol. 28, no. 11 (November):1013–1039.

Fama, E.F., and K.R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, vol. 33, no. 1 (February):3–56.

Fama, E.F., and J.D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy*, vol. 81, no. 3 (May/June):223–237.

Keim, D., and A. Madhavan. 1998. "The Cost of Institutional Equity Trades." *Financial Analysts Journal*, vol. 54, no. 4 (July/August):50–69.

Manaster, S., and R. Rendleman, Jr. 1982. "Option Prices as Predictors of Equilibrium Prices." *Journal of Finance*, vol. 37, no. 4 (September):1043–1057.

Ofek, E., M. Richardson, and R. Whitelaw. 2004. "Limited Arbitrage and Short Sales Restrictions: Evidence from the Options Markets." *Journal of Financial Economics*, vol. 74, no. 2 (November):305–342.

Pan, J., and A. Poteshman. 2006. "The Information in Option Volume for Future Stock Prices." *Review of Financial Studies*, vol. 19, no. 3 (Fall):871–908.

Poteshman, A.M. 2001. "Underreaction, Overreaction, and Increasing Misreaction to Information in the Options Market." *Journal of Finance*, vol. 56, no. 3 (June):851–876.

Xing, Y., X. Zhang, and R. Zhao. Forthcoming. "What Does Individual Option Volatility Smirk Tell Us about Future Equity Returns?" *Journal of Financial and Quantitative Analysis*.

