

Skewness Swaps: Evidence from the Cross-Section of Stocks

- INTERNET APPENDIX -

A ETF Ownership Data Construction

Data collection and database merging

The list of ETFs is taken from ETF.com, which is a subsidiary of Cboe Global Markets specialized in ETF analysis, databases, and news. I consider the ETFs based in the US and those with an equity focus. The historical quarterly holdings of the ETFs are taken from the S12 fund file of Thompson Reuters, which contains the mutual fund holdings data filed with the SEC (Security Exchange Commission) every quarter. The matching of the ETFs taken from ETF.com with the funds in Thompson Reuters is primarily done through the ticker of the fund. However, because many funds have a missing ticker in Thompson Reuters, I match the funds by names with the fuzzy logic Levenstein measure of string distances. Finally, I confirm all the matchings by comparing the asset under management (AUM) of the funds. With this procedure, 329 ETFs are matched, which covers 85% of the full ETF scenario, as reported by ETF.com.

Some of the ETFs are share classes inside a portfolio (e.g., Vanguard ETFs), and Thompson Reuters reports the holdings at the portfolio level. For these ETFs, I adjust the holdings with the weight that the ETF share class has in the portfolio.

The data on the share prices and shares outstanding are taken from Compustat, and the matching with the holdings of Thompson Reuters is done via the CUSIP. The data on the shares outstanding of the ETF funds are taken from Bloomberg because of better data coverage.

ETF ownership measure

For every stock, the daily ETF ownership measure is constructed in two steps, following Ben-David et al. (2018).

First, for each stock i and ETF j , I extract the quarterly weight that the stock has in the ETF fund:

$$w_{i,j,q} = \frac{Sh_{i,j,q} * S_{i,q}}{AUM_{j,q}} \quad (1)$$

where $Sh_{i,j,q}$ is the number of shares of stock i that are held by the ETF fund j at the end of quarter q , $S_{i,q}$ is the price of stock i at the end of quarter q , and $AUM_{j,q}$ is the amount of asset under management of the fund j .

Then, I compute a daily measure of ETF ownership according to the following formula:

$$\text{ETF Ownership}_{i,t} = \frac{\sum_{j=1}^J w_{i,j,t} AUM_{j,t}}{MktCap_{i,t}} \quad (2)$$

where J is the set of ETFs that hold stock i ; $w_{i,j,t}$ is the weight of the stock in the ETF j which is extracted with Equation 1 from the most recent quarterly report; $AUM_{j,t}$ is the amount of asset under management of the fund j at time t calculated as the product of the number of shares outstanding of the ETF fund times the price of the ETF share; and $MktCap_{i,t}$ is the market capitalization of stock i at time t .

B Numerical Analysis

This section studies the precision of the skewness swap in measuring the third moment of the asset returns in the Merton jump-diffusion model. The choice of the Merton model is by no means restrictive because, as shown in Hagan et al. (2002), for time horizons less than one year the implied volatility smile generated by the Merton model can satisfactorily reproduce the empirical one. The dynamics of the Merton model is as follows:

$$ds_t = \left(r - \lambda\kappa - \frac{1}{2}\sigma^2 \right) dt + \sigma dW_t + \log(\psi)dq_t \quad (3)$$

where s_t is the logarithm of the stock price, r is the risk-free interest rate, and σ is the instantaneous variance of the return conditional on no jump arrivals. The Poisson process, q_t , is independent of W_t and such that there is a probability λdt that a jump occurs in dt and a $(1 - \lambda)dt$ probability that no jump occurs:

$$q_t = \begin{cases} 1, & \text{with probability } \lambda dt \\ 0, & \text{with probability } (1 - \lambda)dt \end{cases} \quad (4)$$

The parameter λ represents the mean number of jumps per unit of time. The random variable ψ is such that $\psi - 1$ describes the percentage change in the stock price if the Poisson event occurs, and $\kappa = E[\psi - 1]$ is the mean jump size. I further make the standard assumption (e.g., see Bakshi et al. (1997)) that $\log(\psi) \sim N(\mu, \delta^2)$.

The characteristic function is given by:

$$\phi(u) = E[e^{iu(s_t - s_0)}] = e^{t(i(r - 0.5\sigma^2 - \lambda\kappa)u - 0.5\sigma^2 u^2 + \lambda(e^{i\mu u - 0.5\delta^2 u^2} - 1))}$$

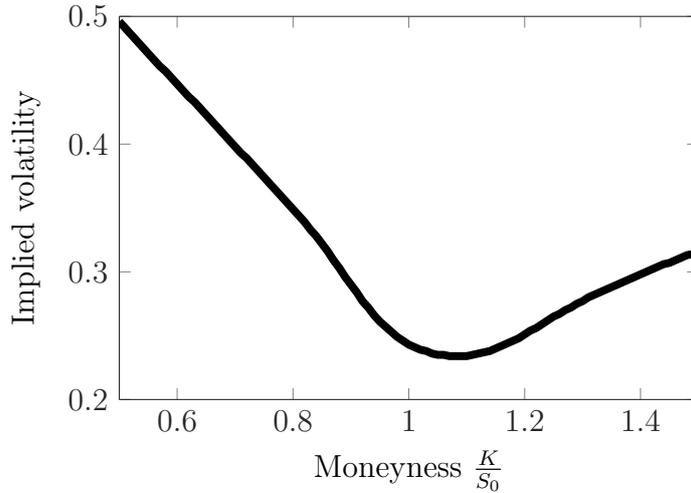


Figure 1. Merton Implied Volatility Smile. The figure displays the one-month implied volatility smile generated by the Merton model with the following set of representative parameters: $r = 0, \mu = -0.05, \delta = 0.08, \sigma = 0.2, \lambda = 3$.

and the moments can be recovered from the property of the characteristic function:

$$E[(s_t - s_0)^k] = (-i)^k \frac{d^k}{du^k} \phi(u)|_{u=0}.$$

As a result, the third moment can be expressed in closed form by the following formula:

$$E[(s_t - s_0)^3] = 3\delta^2 \lambda \mu t + \lambda \mu^3 t - 3(-\delta^2 \lambda - \lambda \mu^2 - \sigma^2)(r + (-\kappa \lambda + \lambda \mu - 0.5\sigma^2))t^2 + (r - \kappa \lambda + \lambda \mu - 0.5\sigma^2)^3 t^3.$$

I chose the following standard parameter values for the simulation study: $r = 0, \mu = -0.05, \delta = 0.08, \sigma = 0.2, \lambda = 3, t = 30/365$. With these parameters, the Merton implied volatility smile is left skewed, as shown in Figure 1. The fixed leg of the swap is defined by Equation 4 in the main text with the discrete approximation given by Equation 10 in the main text; the function Φ is set equal to Φ_S , defined in Equation 23 of Appendix A in the main text. I implement the fixed leg of the swap with the option prices implied by the Merton jump-diffusion model.

Table 1 shows the convergence of the fixed leg of the swap to the true model-based

Table 1. Convergence in the Number of Options.

This table shows the convergence of the fixed leg of the trading strategy to the third moment of the asset returns when the number of options increases. The error is computed as $\frac{|\text{True moment} - \text{Strategy fixed leg}|}{\text{True moment}}$ and is displayed as a percentage. The returns are assumed to follow a Merton jump-diffusion process with standard parameters, i.e., $\mu = -0.05$, $\delta = 0.08$, $\sigma = 0.2$, $\lambda = 3$, $r = 0$, $t = 30/365$. The true moment is computed in closed form and is equal to $-3.12 \cdot 10^{-4}$. In the second column, I consider the skewness swap of Schneider and Trojani (2015) with $\Phi = \Phi_3$, while in the third column I consider the skewness swap introduced in this paper with $\Phi = \Phi_S$.

Number of options	Error (%)	
	Φ_3	Φ_S
8	-18.81%	-5.16%
10	-21.10%	-1.70%
20	-22.70%	-0.72%
50	-22.80%	-0.86%
100	-22.80%	-0.87%

third moment when the number of options used increases from 8 to 100. I compare the performance of the swap defined by Φ_S , that is, the swap used in the empirical application of this paper, with the swap of Schneider and Trojani (2015) defined by Φ_3 (see Equation 22 in Appendix A). The precision of the swap defined by Φ_S is superior because it can reach an error of less than 1%, while the swap defined by Φ_3 always has an error around 20%. This analysis shows that the isolation of the third moment from the fourth is important in order to have a precise measure. It is also interesting to notice that in the swap defined by Φ_S , even with only 8 strikes available (which is the average number of options that I have empirically), the error is only on the order of 5%. The results show that the measurement error is in total less than 1% when the number of options grows to infinity.

Second, I check how the precision of the methodology depends on the range of moneyness available. In this context, the moneyness of an option is defined in standard deviations (SD) as $\log(K/F_{0,T})/(\sigma\sqrt{T})$, where K is the strike price, $F_{0,T}$ is the forward price, σ is the at-the-money implied volatility, and T is the time to maturity. I fix a number of options equal to 10, and I calculate the error of the swap if the ten options span the moneyness range $\pm 1SD$, $\pm 2SD$, $\pm 3SD$, and $\pm 4SD$. Figure 2 shows that the precision changes considerably depending on the moneyness range available. At least three standard deviations are needed in order to have an error around 10%, and four standard deviations are needed to have an error around 1%.

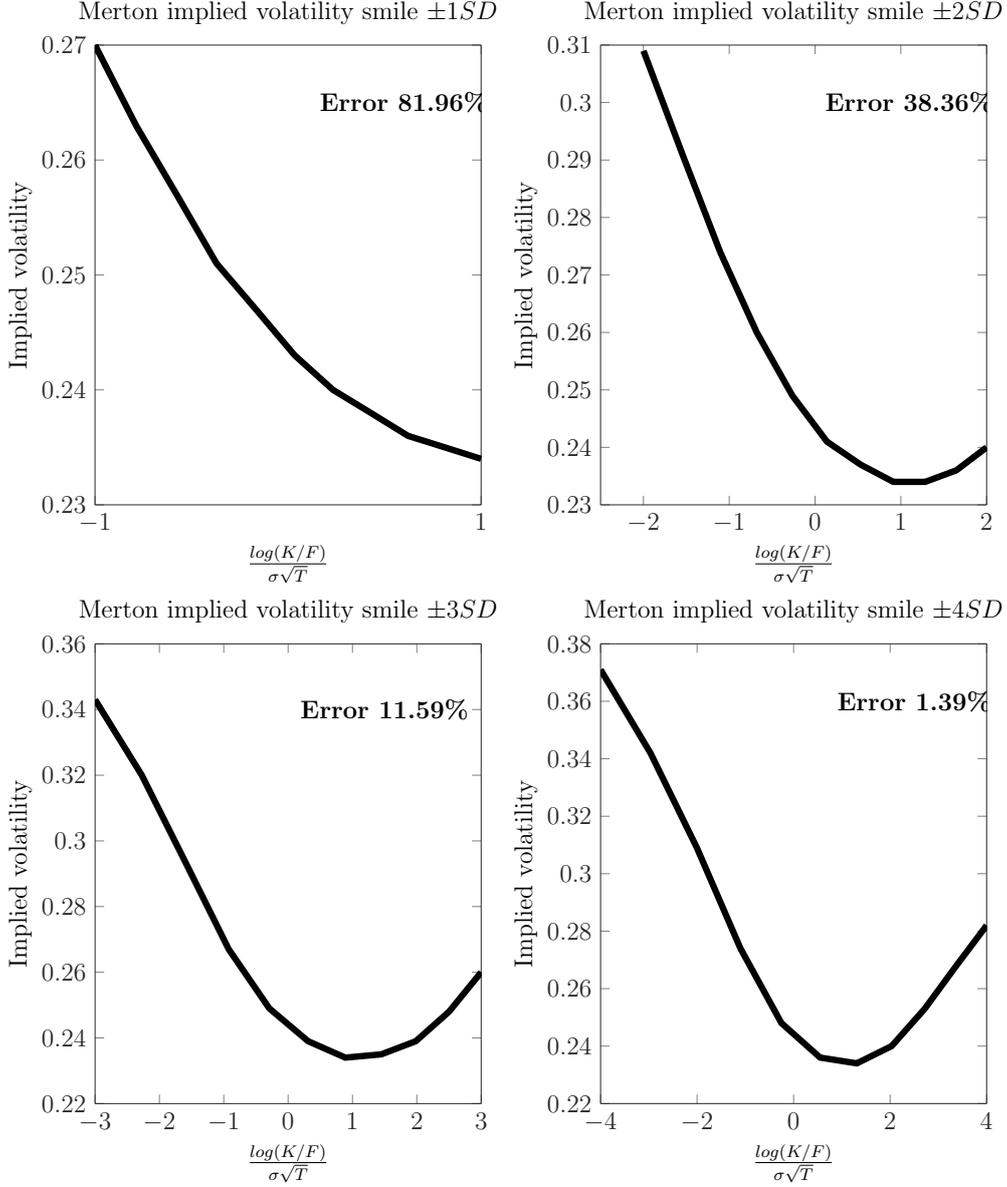


Figure 2. Convergence in the Moneyness Range. This figure shows the volatility smile implied by the Merton jump-diffusion process, with parameters $\mu = -0.05$, $\delta = 0.08$, $\sigma = 0.2$, $\lambda = 3$, $r = 0$, $t = 30/365$, for increasing moneyness ranges. In the first plot, I consider a moneyness range of $[-1SD, 1SD]$, in the second one $[-2SD, 2SD]$, the third one $[-3SD, 3SD]$ and in the fourth one $[-4SD, 4SD]$, where SD is defined as $SD = \log(K/F_{0,T})/(\sigma\sqrt{T})$, where K is the strike price, $F_{0,T}$ is the forward price, σ is the at-the-money implied volatility, and T is the time to maturity. Each plot is a zoom out of the previous one by $1SD$. For each moneyness range, the risk-neutral third moment of the returns is computed with the skewness swap defined by Equation 1 in the main text, and the result is compared with the benchmark model-based risk-neutral third moment, which is recovered in closed form. The error is displayed at the top right of each graph as a percentage.

C CBOE Margin Requirements

The margin requirements for trading options are explained in the CBOE margin manual (<http://www.cboe.com/LearnCenter/pdf/margin2-00.pdf>). There are margin requirements for every individual option and also for popular option strategies like put/call spread or butterfly spread. Because the skewness swap strategy is not a standard trading strategy, I calculate the total margin of the strategy as the sum of the margin of the individual options. There are two types of trades in the option portfolio of the skewness swap: long call, with a maturity less than nine months, and short put, with a maturity less than nine months. The margin requirements are the following:

- **Long call with a maturity less than nine months.** Margin requirement: pay for the option in full.
- **Short put with a maturity less than nine months.** Margin requirement:

$$\max(\text{Option price} + 0.2 \cdot S_0 - |S_0 - K|, \text{Option price} + 0.1 \cdot K)$$

where S_0 is the stock price, K is the strike price, and $|S_0 - K|$ is defined as the out-of-the-money amount.

The above margin requirement is expressed as a function of the level of S_0 and K ; hence, for stocks with a high price S_t , the margin can be much higher than the simple price of the option, particularly for out-of-the-money options.

Let us look at an example that is taken from the CBOE margin manual. An investor is short a put option with strike $K = 80$, option price = 2, $S_t = 95$, and an expiration less than nine months. The option is out-of-the-money. The margin requirement is $\max(2 + 0.2 \cdot 95 - 15, 2 + 8) = 10$, which is five times more than the price of the option.

D Bakshi–Kapadia–Madan (2003) Risk-neutral Skewness

The risk-neutral skewness proposed by Bakshi et al. (2003) has the following expression:

$$SKEW(t, \tau) = \frac{e^{r\tau}W(t, \tau) - 3\mu(t, \tau)e^{r\tau}V(t, \tau) + 2\mu(t, \tau)^3}{[e^{r\tau}V(t, \tau) - \mu(t, \tau)^2]^{3/2}}$$

where $\mu(t, \tau) = e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V(t, \tau) - \frac{e^{r\tau}}{6}W(t, \tau) - \frac{e^{r\tau}}{24}X(t, \tau)$, and $V(t, \tau)$, $W(t, \tau)$, $X(t, \tau)$ are the prices of volatility, cubic, and quartic contracts, respectively, which can be computed with option prices. The price of a volatility contract is given by

$$V(t, \tau) = \int_{S(t)}^{\infty} \frac{2 \left(1 - \log \left[\frac{K}{S(t)}\right]\right)}{K^2} C(t, \tau, K) dK + \int_0^{S(t)} \frac{2 \left(1 + \log \left[\frac{S(t)}{K}\right]\right)}{K^2} P(t, \tau, K) dK$$

where $C(t, \tau, K)$ ($P(t, \tau, K)$) is the price of an European call (put) option quoted on day t with time to maturity τ and strike K . $S(t)$ is the stock price on day t . Similarly, the prices of a cubic and quartic contract are given by the following portfolios of options:

$$W(t, \tau) = \int_{S(t)}^{\infty} \frac{6 \log \left[\frac{K}{S(t)}\right] - 3 \left(\log \left[\frac{K}{S(t)}\right]\right)^2}{K^2} C(t, \tau, K) dK - \int_0^{S(t)} \frac{6 \log \left[\frac{S(t)}{K}\right] + 3 \left(\log \left[\frac{S(t)}{K}\right]\right)^2}{K^2} P(t, \tau, K) dK$$

$$\begin{aligned}
X(t, \tau) = & \int_{S(t)}^{\infty} \frac{12 \left(\log \left[\frac{K}{S(t)} \right] \right)^2 - 4 \left(\log \left[\frac{K}{S(t)} \right] \right)^3}{K^2} C(t, \tau, K) dK \\
& + \int_0^{S(t)} \frac{12 \left(\log \left[\frac{S(t)}{K} \right] \right)^2 + 4 \left(\log \left[\frac{S(t)}{K} \right] \right)^3}{K^2} P(t, \tau, K) dK.
\end{aligned}$$

I calculate the risk-neutral skewness of Bakshi et al. (2003) for each day of the sample period (01/01/2003–31/12/2017) and for each stock. I use the daily options quotes given by the Optionmetrics interpolated implied volatility surface with 30 days to maturity. I convert the implied volatilities given by Optionmetrics in European Black–Scholes prices, and I use these prices in the algorithm ($C(t, \tau, K)$ and $P(t, \tau, K)$). The above expressions are written for a continuum of options, but in practice only a finite number of options is available. I therefore apply the following discrete approximation. Suppose that, on day t , I have N_c out-of-the-money call options available with strikes $S(t) < K_{c,1} < \dots < K_{c,N_c}$ and N_p out-of-the-money put options available with strikes $K_{p,1} < \dots < K_{p,N_p} < S(t)$. The price of a volatility contract is discretized as follows:

$$\begin{aligned}
V(t, \tau) = & \sum_{i=1}^{N_c} \frac{2 \left(1 - \log \left[\frac{K_{c,i}}{S(t)} \right] \right)}{K_{c,i}^2} C(t, \tau, K_{c,i}) \Delta K_i \\
& \sum_{j=1}^{N_p} \frac{2 \left(1 + \log \left[\frac{S(t)}{K_{p,j}} \right] \right)}{K_{p,j}^2} P(t, \tau, K_{p,j}) \Delta K_j
\end{aligned}$$

where

$$\Delta K_i = \begin{cases} K_1 - S(t) & \text{if } i = 1 \\ (K_i - K_{i-1}) & \text{if } i > 1 \end{cases}$$

and

$$\Delta K_j = \begin{cases} (K_{j+1} - K_j) & \text{if } j < N_p \\ S(t) - K_{N_p} & \text{if } j = N_p \end{cases}$$

The cubic and quartic contracts are discretized in the same way.

E Model-based Skewness Swap

As a robustness check, I implement a model-based skewness swap, where, instead of using the actual option prices, I use the option prices calculated from a fitted model. This skewness swap is not tradable because the fitted option prices are not real quotations, but it is nevertheless a useful econometric check for comparing the model-based skewness swap returns with the real tradable skewness swap returns.

There are two main steps in the implementation of the model-based skewness swap: i) calibration of a model and ii) implementation of the swap according to the calibrated model.

Calibration

The model-based skewness swaps are implemented monthly, as the tradable skewness swaps, and they start and end on the third Friday of each month. I exclude the months in which the stocks pay dividends in order to simplify the calculation of the model-based option prices. I recalibrate the model at each start date of the swap and for each stock separately. In detail, at each start date of the swap t and for each stock S_t , I consider all the out-of-the-money options with a maturity 30 days provided by the Optionmetrics implied volatility surface file. This sample constitutes my calibration sample. I choose as benchmark model the Merton jump-diffusion model with Gaussian jump-size distribution, whose dynamics is given by Equation 3. The variable $\log(\psi)$ (the size of the jump) is normally distributed with parameters $N(\mu, \delta^2)$. This model is simple and tractable, and many empirical studies (see e.g., Hagan et al. (2002)) show that it provides a good fit for short-term options data. I then calibrate the parameters of the Merton jump-diffusion model by minimizing the implied volatility mean squared error (IVMSE) as

$$IVMSE(\chi) = \sum_{i=1}^n (\sigma_i - \sigma_i(\chi))^2$$

Table 2. Calibrated Parameters of the Merton Jump-diffusion Model.

This table displays the average calibrated parameters of the Merton jump-diffusion model for the S&P500 and for the cross-section of stocks. The model is calibrated separately for each stock and index, and it is recalibrated monthly at each start date of the swap. The calibration sample includes all the out-of-the-money options with a maturity 30 days quoted by the Optionmetrics interpolated volatility surface file on the calibration day. The numbers displayed are the average calibrated parameters across time and across stocks.

	λ	μ	δ	σ
S&P500	2.59	-0.08	0.05	0.11
All stocks	3.19	-0.07	0.13	0.20

where $\chi = \{\lambda, \mu, \delta, \sigma\}$ is the set of parameters to estimate, $\sigma_i = BS^{-1}(O_i, T_i, K_i, S, r)$ is the market implied volatility provided by Optionmetrics, and $\sigma_i(\chi) = BS^{-1}(O_i(\chi), T_i, K_i, S, r)$ is the model implied volatility, where $O_i(\chi)$ is the Merton model price of the option i . The model implied volatility is obtained by inverting the Black-Scholes formula, where the option price is given by the Merton model price. In the Merton jump-diffusion model, the option prices are available in closed form (see Merton (1976)). The price of a call option is given by:

$$C_{MRT}(t, \tau, K) = \sum_{n=0}^{\infty} e^{-\lambda' \tau + n \log(\lambda' \tau) - \sum_{i=1}^n \log^n} C(S_t, K, r_n, \sigma_n)$$

where $\lambda' = \lambda(1 + k)$, $k = e^{\mu + \frac{1}{2}\delta^2} - 1$, $C(S_t, K, r_n, \sigma_n)$ is the Black-Scholes price of an European call with volatility $\sigma_n = \sqrt{\sigma^2 + \frac{n\delta^2}{\tau}}$ and risk-free rate $r_n = r - \lambda k + \frac{n \log(1+k)}{\tau}$. The price of a put option is defined analogously. The choice of the implied volatility mean squared error (IVMSE) loss function follows the argumentation of Christoffersen and Jacobs (2004), who show that the calibration made on implied volatilities is more stable out of sample. Table 2 displays the average calibrated parameters for the S&P500 index and for the cross-section of stocks.

Implementation of the model-based swap

With the calibration of the Merton jump-diffusion model in the previous step, I estimate the parameters $\{\widehat{\lambda}, \widehat{\mu}, \widehat{\delta}, \widehat{\sigma}\}$ on day t for stock S_t . I then compute the Merton option prices of an equispaced grid of strikes covering the moneyness range $[-4SD, +4SD]$, where $SD = \frac{\log(K/S_t)}{\sigma\sqrt{T}}$ is the moneyness of the options measured in standard deviations. In this definition, σ is calculated as the implied volatility of an at-the-money option, that is, $\sigma = BS^{-1}(C_{MRT}, T, S_t, S_t, r)$, where C_{MRT} is the Merton price of a call option with a strike equal to S_t . The equispaced grid is constructed as follows. First, I recover $K_{min} = S_t e^{-4\sigma\sqrt{T}}$ and $K_{max} = S_t e^{4\sigma\sqrt{T}}$, then I divide the range $[K_{min}, S_t]$ and $[S_t, K_{max}]$ in 20 intervals each, and finally I compute the Merton option prices of these out-of-the-money puts and calls. In this way, I keep constant the number of options and the moneyness range because I always have 40 option prices, 20 calls, and 20 puts, which span the range $[-4SD, +4SD]$. The fixed leg of the swap (Equation 4 in the main text) is computed using these model-based option prices, where the discreteness of the options is addressed with the same quadrature-based approximation of the integral explained in Equation 10 in the main text. The floating leg of the swap (Equation 5 in the main text) is given by the sum of the payoff of the same option portfolio and a continuous delta-hedge in the forward market. The return of the swap is computed as in Section 2.3.3 in the main text.

Table 8 in the main text reports the average return of the model-based skewness swap for the S&P500 index and for the cross-section of stocks in the pre-crisis and post-crisis subsamples.

F Idiosyncratic Skewness Risk Premium and Book Leverage. Robustness Checks.

As a robustness check, this section studies whether the relation between idiosyncratic skewness risk premium and book leverage is driven by the CAPM-beta of the stocks or by an overvaluation channel.

The portfolio sort analysis presented in Table 3 and Table 4 documents that the stocks with the lowest book leverage have the highest idiosyncratic skewness risk premium, even after controlling for the CAPM beta of the stocks and overvaluation measures.

Table 3. Book Leverage and CAPM-beta.

This table shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of book leverage and the CAPM-beta of the stocks. At each start date of the swaps t , stocks are independently sorted in ascending order according to their book leverage estimate and CAPM-beta estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t + 1$ (i.e., ex-post risk premium). The book leverage is calculated for each stock at each start date of the swaps as the logarithm of the ratio of the book value of the assets to the book value of equity. The CAPM-beta is calculated for each stock at each start date of the swaps by regressing the daily excess returns of the stock in the previous month on the daily excess return of the S&P500. The t-statistics of the difference portfolios are provided in parentheses. All the t-statistics are adjusted using the method of Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	Book leverage 1 (low)	Book leverage 2	Book leverage 3 (high)	Difference
CAPM-beta 1 (low)	25.42%	23.18%	14.62%	-10.79%*** (-3.90)
CAPM-beta 2	21.81%	22.62%	8.02%	-13.78%*** (-5.23)
CAPM-beta 3 (high)	23.07%	19.58%	-10.48%	-23.3%*** (-5.74)
Difference	-2.35% (-1.24)	-3.59% (-1.55)	-4.14% (-1.31)	

Table 4. Book Leverage and Overvaluation Measures.

This table shows the idiosyncratic skewness risk premium of bivariate stock portfolios constructed on the basis of book leverage and one of the three overvaluation measures: *BM*, *MAX*, and *MISP*. At each start date of the swaps t , stocks are independently sorted in ascending order according to their book leverage estimate and overvaluation estimate and are assigned to one of the three tercile portfolios for each sorting variable. The intersection of these classifications yields nine portfolios. I then calculate the average idiosyncratic skewness risk premium of these nine portfolios at the end of the following month $t + 1$ (i.e., ex-post risk premium). The book leverage is calculated for each stock at each start date of the swaps as the logarithm of the ratio of the book value of the assets to the book value of equity. *BM* is the book-to-market ratio, *MAX* is the maximum daily return over the previous month, and *MISP* is the mispricing score of Stambaugh and Yuan (2017). All measures are contemporaneous. The t-statistics of the difference portfolios are provided in parentheses. All the t-statistics are adjusted using the method of Newey and West (1987) with the optimal bandwidth suggested by Andrews and Monahan (1992).

	Book leverage 1 (low)	Book leverage 2	Book leverage 3 (high)	Difference
BM 1 (low)	26.7%	25.54%	22.87%	-3.8%** (-2.79)
BM 2	22.8%	18.93%	18.45%	-4.3% (-1.83)
BM 3 (high)	16.09%	19.89%	-7.25%	-23.3%*** (-5.79)
Difference	-10.61%*** (-5.56)	-5.66%** (-3.13)	-30.12%*** (-6.79)	
	Book leverage 1 (low)	Book leverage 2	Book leverage 3 (high)	Difference
MAX 1 (low)	23.92%	21.16%	5.85%	-18.07%*** (-4.91)
MAX 2	21.72%	20.6%	9.2%	-12.53%*** (-4.54)
MAX 3 (high)	24.75%	22.41%	16.57%	-8.18%*** (-3.44)
Difference	0.08% (0.49)	1.26% (0.53)	10.72%* (2.38)	
	Book leverage 1 (low)	Book leverage 2	Book leverage 3 (high)	Difference
MISP 1 (low)	24.62%	24.92%	22.96%	-1.66% (-0.99)
MISP 2	23.77%	22.81%	7.11%	-16.66%*** (-4.50)
MISP 3 (high)	20.97%	19.34%	6.28%	-14.69%*** (-4.34)
Difference	-3.65%* (-2.07)	-5.58%* (-2.45)	-16.68%*** (-4.70)	

References

- D. W. K. Andrews and J. C. Monahan. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica*, 60(4):953–966, 1992.
- G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049, 1997.
- G. Bakshi, N. Kapadia, and D. Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143, 2003.
- I. Ben-David, F. Franzoni, and R. Moussawi. Do etfs increase volatility? *The Journal of Finance*, 73(6):2471–2535, 2018.
- P. Christoffersen and K. Jacobs. The importance of the loss function in option valuation. *Journal of Financial Economics*, 72(2):291–318, 2004.
- P. S. Hagan, D. Kumar, A. S. Lesniewski, and Woodward D. E. Managing smile risk. *The Best of Wilmott*, page 249, 2002.
- R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.
- W. Newey and K. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.
- P. Schneider and F. Trojani. Fear trading. *Swiss Finance Institute Research Paper*, 15(3), 2015.
- R. F. Stambaugh and Y. Yuan. Mispricing factors. *The Review of Financial Studies*, 30(4):1270–1315, 2017.