Erasmus University Rotterdam
Erasmus School of Economics
Master Thesis Econometrics and Management Science: Quantitative Finance

Variation in Option-implied skewness and stock returns

Mark Rusch (533478)

Supervisor:	A. Quaini
Second assessor:	dr. S.O.R. Lonn
Date final version:	1st June 2024

Abstract

This research investigates the predictive capacity of option-implied skewness variability on the cross section of stock returns. By reviewing existing literature and employing various empirical methodologies—such as portfolio sorts, regression analysis, risk premia assessment, examination of Stochastic Discount Factor (SDF) coefficients, and model misspecification tests—the study explores the relationship between option-implied skewness and stock returns. Portfolio sorts and regression analysis indicate significant turnover within decile portfolios due to fluctuations in skewness variability, suggesting its potential predictability in the cross section of stock returns. Risk premia assessment highlights that incorporating skewness variability alongside traditional skewness measures significantly enhances return prediction models. While acknowledging model misspecification and non-significant SDF loadings, the study reveals an intriguing interaction between the market risk factor and skewness variability, suggesting the need for further investigation. Overall, the findings indicate that including skewness variability in standard pricing models improves the accuracy of return predictions and supports more informed investment decisions. This research contributes to the literature by underscoring the importance of considering variation in option-implied skewness when predicting the cross section of stock returns.

Contents

1 Introduction

Understanding the factors that influence the returns of stocks is paramount in both academic research and practical investment decision-making. The identification of reliable predictors not only enhances our ability to forecast stock performance but also provides valuable insights into market dynamics and investor behavior. Moreover, uncovering these factors contributes to the ongoing debate surrounding market efficiency and the efficacy of traditional asset pricing models.

The ongoing debate surrounding market efficiency and the efficacy of traditional asset pricing models has been a central theme in financial economics literature for decades. Market efficiency, as proposed by Eugene Fama in his seminal work (Fama & MacBeth, 1973), suggests that asset prices fully reflect all available information, leaving no room for investors to consistently outperform the market. However, empirical evidence has challenged the notion of perfect market efficiency, leading to various schools of thought and research streams aimed at understanding market anomalies and identifying factors that can explain stock returns beyond what is predicted by traditional asset pricing models.

One of the most influential contributions to this debate is the Capital Asset Pricing Model (CAPM), which posits that the expected return on an asset is determined solely by its beta, a measure of its systematic risk relative to the market. However, empirical studies have uncovered anomalies, such as the size effect and value effect, which suggest that small-cap stocks and value stocks tend to outperform what would be expected based on their beta alone. These findings have led to the development of alternative models, such as the Fama-French three-factor model, which incorporates additional factors—namely, size and value—to better explain stock returns.

Furthermore, the exploration of other factors, such as momentum, liquidity, profitability, and investment, has expanded our understanding of the drivers of stock returns. Research by Jegadeesh and Titman (1993) on momentum, for example, demonstrated the persistence of stock price trends over intermediate-term horizons, contradicting the efficient market hypothesis. Similarly, studies by Amihud and Mendelson (1986) on liquidity risk highlighted the impact of transaction costs on asset prices and returns, challenging the assumption of frictionless markets.

Moreover, the identification of these factors and their integration into investment strategies has practical implications for investors and portfolio managers. By incorporating factors such as value, momentum, and quality into their investment decisions, practitioners can enhance portfolio performance and manage risk more effectively. The proliferation of factor-based investing and smart beta strategies in recent years underscores the growing recognition of the importance of factor-based models in portfolio construction and asset allocation.

The exploration of factors influencing stock returns not only advances academic research in finance but also has practical implications for investors and market participants. By continually refining our understanding of these factors and their impact on asset prices, I contribute to the ongoing dialogue on market efficiency and the evolution of asset pricing models in finance.

2 Problem Description

Understanding the factors influencing stock returns is critical for both academic research and practical investment decisions. Factors such as firm-specific characteristics, broader market conditions, and investor sentiment all contribute to the complexity of stock returns (Fama & French, 1993; Jegadeesh & Titman, 1993).

One factor that plays a significant role in financial analysis is option implied skewness (Bali & Murray, 2013). Skewneess is a statistical measure that quantifies the asymmetry of the probability distribution of a random variable. Thus, option implied skewness, i.e. skewness provides insights into the distribution of asset returns, with positive skewness indicating a higher likelihood of extreme positive returns and negative skewness indicating a higher likelihood of extreme negative returns. Option implied skewness, derived from option pricing in financial markets, reflects investors' expectations regarding extreme events or tail risks in the underlying asset's returns. It serves as a gauge of market sentiment and risk perception, with lower implied skewness suggesting a heightened perceived risk of negative returns.

The risk-neutral measure, denoted as Q, adjusts real-world probabilities to account for investors' risk preferences in financial markets. Skewness under the risk-neutral measure captures market expectations of future skewness, which may differ from historical skewness due to investors' risk aversion and hedging against tail risks.

Analyzing the risk-neutral skewness of securities based on option implied skewness provides valuable insights into market participants' risk perceptions. Understanding how skewness under the risk-neutral measure differs from historical skewness aids investors and analysts in better assessing the pricing and risk characteristics of financial assets, especially derivative securities such as options. Existing literature underscores the predictive power of low option-implied skewness in predicting stock returns (Bali & Murray, 2013; Jurczenko, Maillet & Negrea, 2002; Schneider, Wagner & Zechner, 2020).

However, while the existing literature has established the predictive power of option-implied skewness, the focus has primarily been on its low values and their association with stock returns. An exploration into the dynamics and variations of skewness, encompassing both positive and negative values, could provide a more nuanced understanding of its impact on the cross-section of stock returns. Park (2013) demonstrates that the volatility of volatility, measured by the VVIX index derived from a cross-section of VIX options, predicts tail risk hedge returns. They also show that the VIX has a direct relationship with the option implied skewness. Additionally, DeLisle, Diavatopoulos, Fodor and Kassa (2021) contribute to this body of knowledge by demonstrating the predictive capacity of implied volatility spread variations in relation to stock returns.

One possible theoretical reason for a shift in option implied skewness is the overvaluation or undervaluation of stocks due to behavioral biases or short-sale constraints. According to Bressan and Weissensteiner (2023), stocks with high (positive) option implied skewness are overvalued because investors are attracted to their lottery-like features, such as high volatility and positive skewness of historical returns. These stocks tend to have low downside risk and high returns. On the other hand, stocks with low (negative) option implied skewness are undervalued because investors are averse to their negative skewness of historical returns. These stocks tend to have high downside risk and low returns. Therefore, a shift in option implied skewness can reflect a correction of mispricing or a change in investor preferences and market expectations.

Another possible theoretical reason for a shift in option implied skewness is the anticipation of future events that can affect the stock price (Kenton, 2023). These events, such as earnings announcements, mergers and acquisitions, product launches, or regulatory decisions, can create uncertainty and volatility in the market, influencing the demand and supply of options. For example, if investors expect a positive event that can boost the stock price, they may buy more call options than put options, increasing the option implied skewness. Conversely, if investors expect a negative event that can lower the stock price, they may buy more put options than call options, decreasing the option implied skewness.

Lastly, negative implied skewness signals an undervalued stock. When this stock has low variability in the skewness and the value of the skewness is negative, the stock is undervalued and will remain undervalued for an extended period, with its price never converging to the fundamental value. Conversely, high variability in the option implied skewness suggests that an undervalued stock would return, within a short time-window, to its fundamental value (Rehman & Vilkov, 2012). The variability in skewness is also linked to arbitrage risk, where high arbitrage risk implies that a misspricing in the stock persists for a longer period (Wurgler & Zhuravskaya, 2002).

In summary, this extended problem description explores the potential role of variability in option-implied skewness as a comprehensive and dynamic predictor of stock returns. By examining how variability in option implied skewness. Explains variations in the cross-section of stock returns. This research seeks to contribute to a more holistic understanding of the predictive power embedded in option-derived information and the underlying dynamics.

3 Methodology

This section first introduces the widely used Bakshi, Kapadia and Madan (2003) method (BKM) for estimating risk-neutral skewness, followed by a definition of a method for measuring variability in risk-neutral skewness. Thereafter, I introduce a framework for testing the influence of the variability in risk-neutral skewness, which is based upon the Fama and MacBeth (1973) factor framework, with extensions.

3.1 Bakshi, Kapadia and Madan Estimators

The BKM is a model-free methodology for estimating risk-neutral moments of an asset return. I derive a risk neutral skewness estimator based on the BKM approach. In practice a version of this approach is adopted by the CBOE, there a shifted version of this estimate is implemented as $SKEW = 100 - 10 \times SKEW_{BKM}$. In this paper the standard BKM will be used as described below.

To enhance the discussion on stock return characteristics and option price structures, let the τ-period return be given by the log return: $R \equiv \ln[S(t+\tau)] - \ln[S(t)]$. I define the volatility contract, the cubic contract, and the quartic contracts to have the payoffs as follows:

$$\begin{aligned} &R^2 : \text{volatility contract (V)} \\ &R^3 : \text{cubic contract (W)} \\ &R^4 : \text{quartic contract (X).} \end{aligned} \tag{1}$$

Let $V \equiv E_Q[e^{-r\tau}R^2]$, $W \equiv E_Q[e^{-r\tau}R^3]$, and $X \equiv E_Q[e^{-r\tau}R^4]$ represent the fair value of the respective payoff. Here, $E_Q[\cdot]$ represents expectation under the risk-neutral probability. The risk-neutral measure, denoted as Q, adjusts real-world probabilities to account for investors' risk preferences in financial markets. Skewness under the risk-neutral measure captures market expectations of future skewness. Then based on these contracts V, W and X, BKM find the following estimator of $SKEW$:

$$Skew_{BKM} \equiv \frac{E_Q(R^3) - 3E_Q(R)E_Q(R^2) + 2E_Q^3(R)}{(E_Q(R^2) - E_Q^2(R))^{3/2}} = \frac{e^{r\tau}W - 3e^{r\tau}\mu V + 2\mu^3}{(e^{r\tau}V - \mu^2)^{3/2}}, \tag{2}$$

where r represent the continuously compounded risk-free rate for the τ-period. The risk-neutral expactation of the squared contract (V), the cubed contract (W), the quartic contract (X), and μ can be calculated as:

$$V = \int_{S^*}^{\infty} \frac{2(1 - \ln(K/S^*))}{K^2} C(K)dK + \int_0^{S^*} \frac{2(1 + \ln(S^*/K))}{K^2} P(K)dK, \tag{3}$$

$$W = \int_{S^*}^{\infty} \frac{3\ln\left(\frac{K}{S*}\right)\left(1 - 2\ln\left(\frac{K}{S^*}\right)\right)}{K^2} C(K)dK - \int_0^{S^*} \frac{3\ln\left(\frac{S^*}{K}\right)\left(1 + 2\ln\left(\frac{S*}{K}\right)\right)}{K^2} P(K)dK, \tag{4}$$

$$X = \int_{S^*}^{\infty} \frac{4\ln^2\left(\frac{K}{S^*}\right)\left(3 - \ln\left(\frac{K}{S^*}\right)\right)}{K^2} C(K)dK - \int_0^{S^*} \frac{4\ln^2\left(\frac{S^*}{K}\right)\left(3 + \ln\left(\frac{S^*}{K}\right)\right)}{K^2} P(K)dK, \tag{5}$$

$$\mu = E_0 \ln\left(\frac{S(\tau)}{S_0}\right) \approx e^{r\tau}\left(1 - e^{-rT} - \frac{V}{2} - \frac{W}{6} - \frac{X}{24}\right), \tag{6}$$

where S^* is an arbitrary stike price that sets the OTM boundary, this is where the delta of a call (put) option closest to 0.5 (-0.5). Further, C(K) and P(K) represents the price of the OTM call and put option with strike price K. In the original model derivation in BKM, each contract (V , W or X) requires the existence of a continuum of options with strike spanning from 0 to infinity. To approximate the integrals in eqs. (3) to (5), it is common to implement a trapezoidal approach to discretize and truncate with available strikes (e.g. see Bali and Murray (2013)):

$$V \approx \sum_i \frac{2\Delta K_i}{K_i^2}\left(1 - \ln\left(\frac{K_i}{F_0}\right)\right) Q(K_i), \tag{7}$$

$$W \approx \sum \frac{3\Delta K_i}{K_i^2}\left(2\ln\left(\frac{K_i}{F_0}\right) - \ln^2\left(\frac{K_i}{F_0}\right)\right) Q(K_i), \tag{8}$$

$$X \approx \sum_i \frac{4\Delta K_i}{K_i^2}\left(3\ln^2\left(\frac{K_i}{F_0}\right) - \ln^3\left(\frac{K_i}{F_0}\right)\right) Q(K_i), \tag{9}$$

where $\Delta K_1 = K_2 - K_1$, $\Delta K_n = K_n - Kn - 1$ and $\Delta K_i = (K_{i+1} - K_{i-1})/2$ for i in $\{2, ..., N-1\}$, and the strike price is ordered from low to high. $Q(K_i)$ is the price of an OTM put (call) option if K_i is smaller (larger) than the forward level F_0. That is, S^* is chosen to be the forward level $F_0 = S_0 e^{(r-q)\tau}$ with an estimated dividend yield q. Based on the results of the study done by Liu and van der Heijden (2016), for the BKM estimate of $SKEW$ it would be sufficient to use the raw data and not use any interpolation method for estimating the $SKEW$ moment.

3.2 Variability in $SKEW$

Based on the the option implied skewness measure, I need to define an appropriate measure for the variability in the skewness. In DeLisle et al. (2021), the variability in implied volatility spread is measured by the standard deviation of the last 20 days. I expand research into the behavior and performance of variability in option-implied factors in predicting stock returns. I believe it is thus appropriate to adopt a similar methodology to that employed by DeLisle et al. (2021) for standardizing the measurement of variability in option-implied metrics. Thus, as variability measure I use $SKEW\_STD$ which is the standard deviation of the $SKEW$ over the last 20 days.

3.3 Dynamic investigation of $SKEW\_STD$

In order to assess the usefulness of considering the variability in option implied skewness for asset pricing tasks, I employ multiple methods: (i) Fama-French portfolio sorts,(ii) Turnover analyses and (iii) Fama-French Regression.

Fama-French portfolio sorts

First, I use the Fama-French approach of sorting portfolios. After having obtained the variable for $SKEW\_STD$, a Fama French testing framework will be adapted, where the portfolios will be sorted into deciles that are double sorted on $SKEW$ and $SKEW\_STD$ so that I can investigate the variability. Here the $SKEW$ and $SKEW\_STD$ will be ranked into 10 deciles or 5 quantile portfolios, which are sorted from low to high in $SKEW$ or $SKEW\_STD$. Then an equally-weigted portfolio in each (double-)sorted decile (quantile) portfolio will be constructed. Analyses of these portfolio's will provide a first view into the relation of $SKEW$ and $SKEW\_STD$ on the cross section of returns, it will need to confirm a negative relation between $SKEW$ and returns, and potentially some directional relation with $SKEW\_STD$.

Turnover

The variability of $SKEW$, measured by the standard deviation of skew ($SKEW\_STD$), serves as a crucial indicator of the inherent risk structure in assets. The option implied $SKEW$ of an asset reflects its risk profile, with higher variability suggesting a less stable risk environment. In theory, as I use option implied metrics, assets with unstable risk profiles are inherently less predictable, leading to increased risk. In efficient markets, this elevated risk should ideally be compensated by higher returns to entice investors to hold such assets.

When studying the influence of SKEW variability on asset returns, it is essential to consider the turnover of stocks within SKEW deciles. Turnover, in this context, refers to the frequency at which stocks migrate or change skew decile based on their $SKEW\_STD$. As I want to examine if different levels of variability in $SKEW$ leads to different levels of returns, I need to first examine the turnover based on $SKEW\_STD$, since it might be that stocks with high $SKEW\_STD$ exhibit more instability in their skew values and are thus prone to more frequent changes within $SKEW$ deciles. Conversely, an alternative posits that while skew values may vary significantly around a stable mean, they may exhibit a consistent mean skew over time. In such a scenario, assets with high $SKEW\_STD$ may demonstrate stability in their average skew values but experience significant fluctuations around this mean. Consequently, these assets may exhibit high $SKEW\_STD$ despite having relatively stable and predictable mean $SKEW$ values. Examining turnover within $SKEW$ deciles provides insights into the dynamics of skew variability and its implications for asset risk and returns.

Fama-French regressions

After having obtained the $SKEW$ and $SKEW\_STD$ variables for each of the stocks, I perform a factor regression analysis using the methodology from Fama and MacBeth (1973) and Fama and French (1992), where I test the significance of the $SKEW$ and $SKEW\_STD$ variables combined with other benchmark factors. In order to test these variables, factor-mimicking portfolio's are constructed for the $SKEW$ and $SKEW\_STD$ variable. The $SKEW$ factor-mimicking portfolio will be constructed by taking a long position in the lowest $SKEW$ stocks and a short in the highest $SKEW$ stocks (Bali & Murray, 2013). As the hypothesis that is being tested by this research states that a stock exhibiting high turnover within $SKEW$ deciles based on the high value of $SKEW\_STD$ leads to higher returns and quicker realisation of the undervaluation based on low $SKEW$ values. I construct the $SKEW\_STD$ factor-mimicking portfolio by taking a long position in the upper decile $SKEW\_STD$ and a short position in the lower $SKEW\_STD$ decile, this factor-mimicking portfolio is referred to as Model 1. Based on the results of the turnover section, an extra factor mimicking portfolio is constructed, this $SKEW\_STD$ factor-mimicking portfolio is constructed by taking a long position in the high and low $SKEW\_STD$ decile portfolios and a short position in the middle decile portfolio, this factor is referred to as Model 2. The regression then looks like:

$$\begin{aligned} R_i - R_f = \alpha_i &+ \beta_{i,\text{Mkt}}(R_{\text{Mkt}} - R_f) + \beta_{i,\text{SMB}}\text{SMB} + \beta_{i,\text{HML}}\text{HML} \\ &+ \beta_{i,\text{SKEW}}\text{SKEW} + \beta_{i,\text{SKEW\_STD}}\text{SKEW\_STD} + \varepsilon_i. \end{aligned} \tag{10}$$

here standard regressional tests are used to see if the $B_{i,\theta}$ carry significance for $\theta \in \{MKT, SMB, HML, SKEW, SKEW\_STD\}$.

To validate the earlier findings of a negative slope in the alpha on $SKEW$ sorted decile portfolios (Bali & Murray, 2013) and to get a preliminary understanding of the dynamics of $SKEW\_STD$, I will also regress the factors on equally weighted decile and quantile portfolios, sorted on $SKEW$ and $SKEW\_STD$ as:

$$R_{quantileportfolio} - R_f = a_i + \beta_{i,\text{Mkt}}(R_{\text{Mkt}} - R_f) + \beta_{i,\text{SMB}}\text{SMB} + \beta_{i,\text{HML}}\text{HML}. \tag{11}$$

Here $R_{quantileportfolio}$ refers to the sorted portfolio returns.

Testing Model Specification

In the assessment of specification of asset pricing models, the Hansen-Jagannathan Distance (HJD) plays a pivotal role. It quantifies the dissimilarity between theoretical and market return distributions, offering insights into model alignment and potential discrepancies. The HJD is computed as the square root of the Jensen-Shannon divergence between risk-neutral and actual return distributions (Hansen & Jagannathan, 1997; Kan & Robotti, 2008; Barillas, Kan, Robotti & Shanken, 2020). This divergence metric serves to quantify the discrepancy between the market's observed returns and the theoretical expectations under the absence of arbitrage (Hansen & Jagannathan, 1997; Kan & Robotti, 2008; Barillas et al., 2020). Mathematically, it is represented as:

$$\delta = \left[\min_{\gamma}(E[R_t^e] - Cov[R_t, F_t]\gamma)' V[R_t]^{-1}(E[R_t^e] - Cov[R_t, F_t]\gamma)\right]^{1/2}, \tag{12}$$

where $E[R_t^e]$ denotes the excess returns, $E[R_t^e] - Cov[R_t, F_t]\gamma$ denotes pricing errors of a linear SDF as described in eq. (13) , $V[R_t] = E[R^e R^{e\prime}]$ is the variance matrix of returns, and γ are the model parameters. The optimal γ minimize the pricing errors, yielding the minimum HJD.

Additionally, it is crucial to conduct a specification test to assess whether the HJD is significantly different from zero, indicating potential model misspecification. The specification test evaluates the null hypothesis $H_0 : \delta = 0$, suggesting that the model is correctly specified. The test statistic is typically scaled by the square root of sample size T, denoted as $T\delta$, and its asymptotic distribution under the null hypothesis is well understood. For linear factor models, under the null hypothesis, the asymptotic distribution of $T\delta$ follows a chi-squared distribution. This distribution provides easy methodology for calculating p-values and assessment of the significance of the HJD. I use robust and consistent estimations of variances to ensure accurate inference, especially in the presence of model misspecification. These procedures help mitigate potential biases and provide reliable insights into the adequacy of the asset pricing model.

In assessing equality between HJDs of two models, nested or non-nested, statistical tests are conducted. I will test the difference in distance between the CAPM, the Fama-French 3 factor model and the Fama-French with inclusion of the $SKEW$ and $SKEW\_STD$. Thus the last model nests the Fama-French 3 factor model. For nested models, testing $H_0 : \delta_1^2 = \delta_2^2$ involves verifying if parameters corresponding to the additional factors in the second model are zero. Robust estimators and consistent estimations of variances are imperative for accurate inference, especially in the presence of model misspecification (Kan & Robotti, 2008).

3.4 Stochastic Discount Factor (SDF) loadings

If the law of one price holds, then there exists a Stochastic Discount Factor (SDF) M pricing all excess returns, i.e., $E\left[M_t R_t^e\right] = 0$. In order to test if a factor model including the new factors prices the excess returns. I test the loadings of the test factors to this SDF. A factor model for the SDF is represented as:

$$M_t = 1 - \gamma' \left(F_t - E\left[F_t\right]\right). \tag{13}$$

Here, the mean is normalized to one since I work with excess returns. The estimators implemented in Fama and MacBeth (1973) propose to find the candidate factor SDF such that:

$$\gamma = \arg\min_{g \in \mathbb{R}^K} E\left[R_t M_t\right]' E\left[R_t M_t\right]. \tag{14}$$

The expression for γ is given by:

$$\gamma = \left(\text{Cov}\left[R_t, F_t\right]' \text{Cov}\left[R_t, F_t\right]\right)^{-1} \text{Cov}\left[R_t, F_t\right]' E\left[R_t\right]. \tag{15}$$

Additionally, the estimators proposed by Gospodinov, Kan, Robotti and Shanken (2014) aim to find the candidate factor SDF that minimizes the pricing errors, under a weighted L_2-distance:

$$\gamma = \arg\min_{g \in \mathbb{R}^K} E\left[R_t M_t\right]' \text{Var}[R]^{-1} E\left[R_t M_t\right]. \tag{16}$$

Where γ is given by:

$$\gamma = \Big(\text{Cov}\left[R_t, F_t\right]' \text{Var}\left[R_t\right]^{-1} \text{Cov}\left[R_t, F_t\right]\Big)^{-1} \text{Cov}\left[R_t, F_t\right]' \text{Var}\left[R_t\right]^{-1} E\left[R_t\right]. \tag{17}$$

These methodologies are employed to estimate the SDF loadings, providing insights into the risk structures of assets and their implications for investment strategies and portfolio management. Moreover, the analyses provides insights in the importance of the newly tested factors.

Double LASSO

Feng, Giglio and Xiu (2020) propose an alternative approach of testing the significance of new asset pricing factors. They propose a Two-Pass Regression with Double-Selection LASSO. The first selection searches for factors in a large set of factors (h_t) whose covariances with returns are useful for explaining the cross section of expected returns. Then, a selection step is added to search for factors in h_t potentially missed from the first step but that, if omitted, would induce a large omitted variable bias. Factors omitted from both stages of the double-selection procedure must have a small SDF loading and have covariances that correlate only mildly in the cross section with the covariance between factors of interest (g_t) and the returns, these factors can thus be excluded with minimal omitted factor bias ex ante when estimating and testing λ_g. The second step entails an OLS regression, wherein average returns are regressed against the covariances between asset returns and the newly introduced factors, alongside the control factors identified in the initial stage. The model starts with a linear specification for the SDF:

$$m_t = \gamma_0^{-1} - \gamma_0^{-1}\lambda_v^T v_t = \gamma_0^{-1}(1 - \lambda_g^T g_t - \lambda_h^T h_t), \tag{18}$$

where γ_0 is the zero-beta rate, g_t is a $d \times 1$ vector of factors to be tested, and h_t is a $p \times 1$ vector of potentially confounding factors. Without loss of generality, both g_t and h_t are de-meaned. λ_g and λ_h are $d \times 1$ and $p \times 1$ vectors of parameters, respectively. Here λ_g and λ_h are refered to as the SDF loadings of the factors g_t and h_t. In addition, I observe a $n \times 1$ vector r_t of test asset returns. I assume expected returns satisfy:

$$\mathrm{E}\left(R_t\right) = \iota_n \gamma_0 + C_v \lambda_v = \iota_n \gamma_0 + C_g \lambda_g + C_h \lambda_h, \tag{19}$$

where ι_n is a $n \times 1$ vector of 1s, $C_a = \mathrm{Cov}\left(r_t, a_t\right)$, for $a = g, h\ or\ v$. Furthermore, I assume the dynamics of R_t follow a standard linear factor model:

$$R_t = \mathrm{E}\left(R_t\right) + \beta_g g_t + \beta_h h_t + u_t, \tag{20}$$

where β_g and β_h are $n \times d$ and $n \times p$ factor-loading matrices, μ_t is a $n \times 1$ vector of idiosyncratic components with $E(\mu_t) = 0$ and $Cov(\mu_t, v_t) = 0$.

Based on this model the two-pass estimation proceeds as follows:
(1) Two-Pass Variable Selection
(1.a) Run a cross-sectional LASSO regression of average returns on sample covariances between factors in h_t and returns:

$$\min_{\gamma, \lambda} \left\{ n^{-1} \left\| \bar{r} - \iota_n \gamma - \widehat{C}_h \lambda \right\|^2 + \tau_0 n^{-1} \|\lambda\|_1 \right\}, \tag{21}$$

where $\widehat{C}_h = \widehat{\mathrm{Cov}}\left(r_t, h_t\right) = T^{-1} \bar{R} \bar{H}^{\top}$. This step selects among the factors in h_t, those that best explain the cross section of expected returns. Denote $\left\{\widehat{I}_1\right\}$ as the set of indices corresponding to the selected factors in this step.

(1.b) For each factor j in g_t (with $j = 1, \cdots, d$), run a cross-sectional LASSO regression of $\widehat{C}_{g,,j}$ (the covariance between returns and the j th factor of g_t) on $\widehat{C}_h$ (the covariance between returns and all factors h_t) :

$$\min_{\xi_j, \chi_{j,\cdot}} \left\{ n^{-1} \left\| \left(\widehat{C}_{g,\cdot,j} - \iota_n \xi_j - \widehat{C}_h \chi_{j,\cdot}^{\top} \right) \right\|^2 + \tau_j n^{-1} \left\| \chi_{j,\cdot}^{\top} \right\|_1 \right\}. \tag{22}$$

This step identifies factors whose exposures are highly correlated to the exposures to g_t in the cross-section. This is the crucial second step in the double-selection algorithm, that searches for factors that may be missed by the first step but that may still induce large omitted variable bias in the estimation of λ_g if omitted, due to their covariance properties. Denote $\left\{\widehat{I}_{2,j}\right\}$ as the set of indices corresponding to the selected factors in the j th regression, and $\widehat{I}_2 = \bigcup_{j=1}^{d} \widehat{I}_{2,j}$.

(2) Post-selection Estimation Run an OLS cross-sectional regression using covariances between the selected factors from both steps and returns:

$$\left(\widehat{\gamma}_0, \widehat{\lambda}_g, \widehat{\lambda}_h\right) = \arg\min_{\gamma_0,\lambda_g,\lambda_h} \left\{\left\|\bar{r} - \iota_n\gamma_0 - \widehat{C}_g\lambda_g - \widehat{C}_h\lambda_h\right\|^2 : \quad \lambda_{h,j} = 0, \quad \forall j \notin \widehat{I} = \widehat{I}_1 \bigcup \widehat{I}_2\right\}. \tag{23}$$

The LASSO estimator, like other dimension-reduction methods, relies on a tuning parameter, the penalty parameter τ_0. This parameter is chosen to balance the trade-off between model fit and model sparsity. The robustness of the LASSO selection, in terms of which factors are chosen, is evaluated by exploring how it depends on τ_0.

A key question in this evaluation is determining a reasonable range of values for τ_0 to consider. To address this, a procedure is proposed where the tuning parameter is selected through 5-fold cross-validation (CV). Since these simulations are non-deterministic, the tuning-parameter-selection procedure is run multiple times to explore robustness across different sets of simulations.

The approach gives a more robust estimate of the SDF loadings of the tested factors, by balancing sparsity and accuracy. In the case of a well specified model, these estimates of the SDF loading should then present the pricing dynamics of the factor model.

3.5 Risk Premium Estimation

To understand the compensations for risk factor exposures, I also test for K factors the factor risk premia $\lambda \in \mathbb{R}^K$. They are obtained by the second stage regression of the expected asset returns on regression coefficients as:

$$E[R_t] = \beta\lambda + \epsilon, \tag{24}$$

where $\epsilon \in \mathbb{R}^N$ is the vector of pricing errors for N test assets. Risk premia estimation is performed using the methodology proposed by Kan, Robotti, and Shanken (Kan, Robotti & Shanken, 2013), which accounts for potential model misspecification. The factor risk premia estimates obtained through this approach are given by:

$$\lambda_{FRP} = \left(\beta' W \beta\right)^{-1} \beta' W E\left[R_t\right], \tag{25}$$

where W is a symmetric and positive definite weighting matrix. The tradeable risk premia are calculated as the negative covariance of factors F with the stochastic discount factor (SDF) projection on asset returns, i.e., the minimum variance SDF. The formula used for computing the tradable factor risk premia is:

$$\lambda_{TFRP} = \mathrm{Cov}[F, R] * \mathrm{Var}[R]^{-1} * E[R]. \tag{26}$$

The heteroskedasticity and autocorrelation robust standard errors are computed using the Newey and West (1994) plug-in procedure to select the number of relevant lags. This methodology ensures robust estimation of tradeable risk premia while considering potential model misspecification and accounting for heteroskedasticity and autocorrelation in the data.

Tradeable risk premia of risk factors represent compensation investors receive for bearing specific risks through tradable financial instruments, while risk premia of risk factors encapsulate the underlying sources of risk driving asset returns, informing portfolio construction and risk management decisions. Understanding the distinction is crucial for designing efficient portfolios, managing risks effectively, and developing investment strategies aimed at capturing long-term sources of return.

Oracle Estimator

This section discusses the Oracle tradable risk premium estimators proposed by Quaini, Trojani and Yuan (2023). This estimator is designed to create a testing framework for introducing a new factor into a pricing model.

First, let $R_t := (R_{1t}, ..., R_{Nt})'$ and $F_t := (F_{1t}, ..., F_{Kt})'$ be a vector of excess returns and a vector of candidate asset pricing factors, where $K < N$, observed at times $t = 1, ..., T$. The joint vector $Y_t := (R_t', F_t')$ has moments partitioned as:

$$E[Y_i] = \begin{pmatrix} \mu_R \\ \mu_F \end{pmatrix}, \quad Cov[Y_i, Y_j] = \begin{bmatrix} V_R & V_{RF} \\ V_{FR} & V_F \end{bmatrix}. \tag{27}$$

Given the assumption of a positive definite V_R, the tradable risk premium of vector F_t is defined by:

$$\lambda^i := -Cov[F_t, M_t^i] = V_{FR} V_R^{-1} \mu_R, \tag{28}$$

where M_t^i is the SDF projection on asset returns as the minimum variance SDF. This can be estimated as:

$$\hat{\lambda}^i = \hat{V}_{FR} \hat{V}_R^{-1} \hat{\mu}_R. \tag{29}$$

The sample mean and sample covariance matrix estimators are given by:

$$\hat{\mu} := (\hat{\mu_R}', \hat{\mu_F}')' := \frac{1}{T} \sum_{t=1}^{T} Y_t, \tag{30}$$

and

$$\hat{V} := \begin{bmatrix} \hat{V_R} & \hat{V_{RF}} \\ \hat{V_{FR}} & \hat{V_F} \end{bmatrix} := \frac{1}{T} \sum_{t=1}^{T} (Y_t - \hat{\mu})(Y_t - \hat{\mu})^T. \tag{31}$$

Quaini et al. (2023) show that, unlike other risk premia estimators, the Oracle estimator is consistent and has standard asymptotic normal behavior even in the presence of useless factors.

Borrowing the terminology from high dimensional statistics (see, e.g., Quaini et al. (2023)), I define an Oracle tradable risk premium estimator as an estimator satisfying two key properties. First, it consistently selects in finite samples factors that are not weak or useless. Second, it implies an efficient asymptotic distribution for the estimated risk premia of the selected factors.

Thanks to the convenient asymptotic properties of sample tradable risk premia, Oracle tradable risk premium estimators can be built with a simple approach. Denote by

$$\widehat{\boldsymbol{\rho}} := [\widehat{\boldsymbol{\rho}}_1, \dots, \widehat{\boldsymbol{\rho}}_K] := \widehat{\mathbb{C}or} \left[\boldsymbol{R}_t, \boldsymbol{F}_t' \right], \tag{32}$$

the $N \times K$ matrix of sample correlations between returns and factors. I propose an Oracle tradable risk premium estimator built by means of a convenient minimum distance correction of sample tradable risk premia, in which estimated risk premia of factors having small sample correlations with all asset returns are shrank using a suitable data-driven penalty.

Definition 2. Oracle Estimator Given a penalty parameter $\tau_T > 0$, consider the penalized estimator defined by:

$$\tilde{\boldsymbol{\lambda}}^i := \left(\tilde{\lambda}_1^i, \dots, \bar{\lambda}_K^i \right)' := \underset{\boldsymbol{\lambda} \in \mathbb{R}^K}{\operatorname{argmin}} \left\{ \frac{1}{2} \left\| \hat{\boldsymbol{\lambda}}^i - \boldsymbol{\lambda} \right\|_2^2 + \tau_T \sum_{k=1}^K \frac{|\lambda_k|}{\|\widehat{\boldsymbol{\rho}}_k\|_2^2} \right\}. \tag{33}$$

Oracle Estimator is defined by a penalized minimum Euclidean distance correction of sample tradable risk premia and belongs to the class of proximal estimators studied. The optimization problem in equation (33) is solvable in closed-form and gives rise to the soft-thresholding formula:

$$\tilde{\lambda}_k^i = \operatorname{sign}\left(\widehat{\lambda}_k^i \right) \max \left\{ \left| \hat{\lambda}_k^i \right| - \frac{\tau_T}{\|\hat{\rho}_k\|_2^2}, 0 \right\}; k = 1, \dots K. \tag{34}$$

Therefore, estimator $\tilde{\boldsymbol{\lambda}}^i$ implies a zero estimated risk premium for all factors associated with a sample tradable risk premium $\hat{\lambda}_k^i$ that is smaller in absolute value than scaled penalty parameter $\tau_T / \|\widehat{\boldsymbol{\rho}}_k\|_2^2$.

I next show that an appropriate choice of tuning parameter τ_T in equation (33) implies the Oracle property. To see this, let

$$\mathcal{S} := \{ k \in \{1, \dots, K\} : \boldsymbol{V}_{RF_k} \neq \boldsymbol{0} \}, \tag{35}$$

be the active set indexing components of factor vector $\boldsymbol{F}_t = (F_{1t}, \dots, F_{Kt})'$ that are neither useless nor weak. Accordingly, for any vector $\boldsymbol{x} \in \mathbb{R}^K$ I denote by $\boldsymbol{x}_{\mathcal{S}}$ the subvector consisting only of components of $\boldsymbol{x}$ with index in $\mathcal{S}$. Finally, the estimated active set implied by estimator $\overline{\boldsymbol{\lambda}}^i$ is denoted by

$$\check{\mathcal{S}} := \left\{ k \in \{1, \dots, K\} : \check{\lambda}_k^i \neq 0 \right\}. \tag{36}$$

Using this notation, I characterize in the asymptotic distribution of estimator (33) and its Oracle property.

Under general assumptions on the DGP of returns and factors, the Oracle tradable risk premia estimator is consistent, achieves oracle factor selection (i.e., the probability that the set of nonzero Oracle estimates coincides with the true set of strong factors tends to 1 as the sample size grows) and on the set of strong factors it has a standard Gaussian behavior even in the presence of useless and weak factors. Weak factors are modeled as having a vanishing correlation with asset returns.

Thus I observe that the proximal tradable risk premium estimator possesses the characteristics of an Oracle estimator. The set of factors identified as neither useless nor weak is accurately selected with a probability tending to one as the sample size increases. Moreover, the asymptotic distribution of the estimated risk premia for these selected factors coincides with that of an Oracle sample tradable risk premium estimator, which possesses a priori knowledge of the useless or weak factors. This property enables the Oracle estimator to provide valid inferences for the tradable risk premia of factors demonstrating minimal correlation with returns. Alternatively, it facilitates a consistent initial screening of factors weakly correlated with returns, paving the way for the subsequent application of standard cross-sectional inference methodologies to determine the risk premia of all other factors within an asset pricing model.

Remarkably, while the shrinkage of sample tradable risk premia dictated by the soft-thresholding formula (34) becomes asymptotically negligible for factors identified as neither useless nor weak, there exists the potential for finite-sample bias. This bias can be mitigated by computing sample risk premia exclusively for factors selected within finite samples by our proximal estimator. Notably, under the assumptions outlined, the resulting "relaxed" proximal estimator aligns asymptotically with the original proximal estimator.

3.6 Arbitrage Risk

Finding in Rehman and Vilkov (2012) of more variability into SKEW if the risk of arbitrage is higher, this suggest research into the relation between option implied skewness variability, measured as $SKEW\_STD$ and arbitrage risk. Given a relation, arbitrage risk of a security might subsume the results of $SKEW\_STD$, which would explain the nature of a relationship between $SKEW\_STD$ and the cross-section of returns.

Here I need to use the methodology for *Arbitrage Risk* as Wurgler and Zhuravskaya (2002) describe it. For each stock in our sample, I select, for each month, the three closest *substitute* stocks matched on industry, size and market-to-book ratio Wurgler and Zhuravskaya (2002), I choose the industry classification suggested by Fama and French (1993). To select the closest stocks in terms of size and market-to-book ratios, I compute the sum of the absolute percentage difference of size and market-to-book ratio of each firm with respect to each of the other firms that lie in the same industry. The three firms for which the percentage difference is the smallest are then selected as substitute firms. In order to measure *Arbitrage Risk* every month, I take monthly returns for the five preceding years and estimate the following regression for stock i.

$$R_{it} - R_{ft} = \beta_{1i}(R_{SUB1it} - R_{ft}) + \beta_{2i}(R_{SUB2it} - R_{ft}) + \beta_{3i}(R_{SUB3it} - R_{ft}) \tag{37}$$

where R_{SUB1it}, R_{SUB2it} and R_{SUB3it} denote the returns on three industry, size and market-to-book matched substitute stocks, while R_{ft} denotes the risk-free rate. *Arbitrage Risk* for stock i is then the variance of the residuals from this regression. The higher the variance, the poorer are the substitutes in explaining the returns of stock i and the higher the risk of *arbitrage*.

Then the relation between $SKEW\_STD$ and *Arbitrage Risk* is measured by regressing $SKEW\_STD$ and arbitrage onto each other. This is important in understanding $SKEW$ variability, and it could imply that the higher the *Arbitrage Risk* the higher the speed with which prices in securities are corrected based on under or overpricing based on the $SKEW$ level.

3.7 Mean-reverting SKEW

In the preceding findings I demonstrate that SKEW STD enhances the predictability of returns based on SKEW. It is conceivable that options displaying a significant autoregressive component in SKEW may also exhibit a sizable SKEW STD. For instance, an option with a implied SKEW possessing a highly negative AR(1) coefficient is likely to display pronounced reversal characteristics, potentially resulting in a large standard deviation as SKEW fluctuates frequently. To investigate the hypothesis that the autoregressive nature of SKEW influences our earlier findings, I initially estimate the AR(1) coefficient for each securities SKEW on a monthly basis using daily data. Subsequently, akin to the portfolio analyses conducted in the portfolio sorts section, I categorize securities into portfolios based on their SKEW AR(1) coefficient and SKEW level. The sorting process is performed both sequentially and independently. For each portfolio, I conduct regressions of daily portfolio returns on the Fama-French 3-factor model and estimate the alphas. A lack of any discernible pattern in the long-short portfolio alphas can than suggests that the SKEW AR(1) Coefficient is not a driving factor behind the robust outcomes observed in the SKEW STD analyses.

3.8 Model Identification

To ensure the robustness and validity of our model estimation, I conduct a thorough examination of the identification of model parameters. This step is crucial to confirm that there is no ambiguity in the estimation process, thereby enhancing the reliability of our results. For this purpose, I adopt the Beta Rank Test proposed by Chen and Fang (2019).

The Beta Rank Test is a statistical procedure designed to assess the identification of model parameters in econometric models. It evaluates whether the model parameters can be uniquely determined based on the available data, thereby ensuring the absence of multicollinearity or other issues that may lead to ambiguous parameter estimates.

4 Data

4.1 Screening

In order to accurately estimate BKM estimates I need data of options on securities with a wide range of strike prices. Therefore, I seek a group of stocks that are widely traded. I use the options with underlying securities traded in the S&P 500 and the index itself, the list of securities is based on the securities within the S&P 500 on the 11th of March 2024.

Option Data on S&P 500 stocks will be coming from OptionMetrics. Not all securities have traded Options in the OptionMetrics database. I use 496 out of the 504 traded securities that are traded within the S&P 500, for which data is available on OptionMetrics. The data will span from 01-01-2001 to 28-02-2023, and will have a daily frequency.

There are three primary rationales behind our utilization of daily data to derive weekly estimates for our variables (Bakshi et al., 2003). Firstly, employing daily data helps mitigate the influence of outliers, as it permits the calculation of moments on a daily basis, subsequently allowing for the averaging of these moments over the course of the calendar week. Secondly, the accurate estimation of the slope of the weekly smile for individual equity options necessitates daily data throughout the week to ensure an adequate number of observations for smile estimation. Thirdly, it has been established that daily risk-neutral index skews demonstrate a seasonal pattern, particularly on Mondays, as documented by Harvey and Siddique (1999), thus averaging over a week is a necessity.

In line with the recommendations from Bakshi et al. (2003) to maintain consistency with existing literature, a rigorous data screening process was implemented, incorporating both qualitative and quantitative criteria. This involved the elimination of bid-ask option pairs with missing quotes or zero bids, alongside the removal of option prices that violated arbitrage restrictions. These restrictions are defined by conditions

$$
\begin{aligned}
C_A &\geq \max[0, S-K, S-PV(K)-PV(D)], \\
C_E &\geq \max[0, S-PV(K)-PV(D)], \\
P_A &\geq \max[0, K-S, PV(K)+PV(D)-S], \\
P_E &\geq \max[0, PV(K)+PV(D)-S],
\end{aligned}
\tag{38}
$$

where K is the strike price, S denotes the current stock price, and PVD$[D]$ and PVD$[K]$ represent the present value functions for dividends and strike price, respectively and the lower case A or E denotes American or European options respectively.

Furthermore, to account for potential fluctuations in market activity, options with less than 9 days and more than 60 days to expiration were excluded from consideration. Finally, in accordance with Theorem 1 of Bakshi et al. (2003), only out-of-the-money (OTM) calls and puts were retained. As some stocks might pay dividends, the out-of-the-money (OTM) condition for

Call options is having a $\Delta > 0.5$, and for Put options, it is having a $\Delta < -0.5$.

Based on the results of Aschakulporn and Zhang (2022), I will only use stocks with at least 4 options, and a maximum $\Delta K = 50$. In their simulation study this leads to maximum absolute errors of the BKM skewness estimator of approximately 0.065. This difference should not interfere with the strategy proposed by this research, and it makes sure that there is still enough data on a wide variety of stocks available.

Although each series for skewness and kurtosis pertain to a constant τ, in practice, it is not possible to strictly observe these, as options are seldom issued daily with a constant maturity. Therefore, in our empirical exercises, if an OTM option has remaining days to expiration of 9 to 60 days, it is grouped and deemed to have the same time to maturity. Thus only one classifications of smiles and option portfolios are investigated.

In order to estimate the F_0 contract in the BKM estimators, the risk-free rate and dividend yield are to be obtained. The dividend yield will be estimated based on Dividend rates provided by OptionMetrics. The Fama-French 5 factors, the Moment Factor and the risk-free rate will obtained from the Keneth-French database.

4.2 De-listed securities

To address the inclusion of de-listed or newly listed securities within the database, a decision is made to retain their existing status. This reflects the inherent characteristics of financial markets. When the frequency of the rebalancing strategy exceeds a daily cadence, securities lacking available data between the current and subsequent rebalancing dates are excluded from the portfolio. This proactive measure ensures that securities without pertinent data are not retained within the portfolio, aligning with the objective of maintaining data integrity and reliability in the analysis.

4.3 Missing Values

For option data, missing key values such as strike price, option price, or time to maturity are deemed critical for research purposes. As such, any option lacking these vital attributes is removed from the dataset. This decision was made to mitigate potential noise introduced by imputation methods, ensuring the quality and accuracy of the data analyzed.

However, dividend yield rates within the option data were treated differently. Given that a missing dividend yield rate typically indicates a lack of dividend data for a particular stock. It is assumed that stocks with missing dividend yield rates do not pay dividends. Therefore, missing dividend yield rates were imputed with 0 values. This assumption aligns with the common practice in financial analysis, where a missing dividend yield rate is interpreted as the absence of dividend payments for the respective stock.

Prior to 2008, the dataset exhibited significant missing data in the BKM_skew column, coupled with a relatively small number of securities (Appendix 2). To stabilize the dataset and improve the robustness of subsequent analyses, all data preceding 2008 was removed. This action resulted in a more consistent number of securities and a proportion of missing values, leading to greater stability in the construction of factor mimicking portfolios and subsequent analysis.

Following the removal of pre-2008 data, the dataset still contained missing values, albeit at reduced levels. In particular, the factors contained approximately 0.4% missing values, while the returns had around 0.049% missing values. Given that regression analysis in matrix notation does not permit missing values, imputation methods were employed.

For the factors, a 5-Nearest-Neighbors approach was utilized to impute missing values, leveraging the similarity of neighboring data points. Conversely, for the returns, a time-average imputation method was adopted. This approach involved imputing missing return values on specific dates with the average of observed returns for that date. These imputation strategies were chosen to maintain the integrity of the dataset and facilitate subsequent analyses while minimizing potential biases introduced by missing data.

5 Results

5.1 Turnover of $SKEW\_STD$ within the $SKEW$ decile portfolios

In order to demonstrate that high $SKEW\_STD$ securities migrate more in $SKEW$ decile portfolios than low $SKEW\_STD$ securities, as hypothysized. In Table 1 I show the average turnover for decile portfolios sorted on $SKEW\_STD$, where the turnover is based on migration of $SKEW$ decile portfolios. It shows that the turnover is higher for High $SKEW\_STD$ portfolios than Low $SKEW\_STD$ portfolios at a significance level of 1%. It further shows that this difference is even larger when comparing the Low and High $SKEW\_STD$ decile portfolios with the middle $SKEW\_STD$ decile portfolios. These results show high $SKEW\_STD$ securities migrate across $SKEW$ deciles more often than low $SKEW\_STD$ firms and supports the motivation of the analyses.

Table 1: Turnover SKEW decile portfolios, sorted on SKEW_STD

Decile	Average Turnover (%)
Low	53.89%
Decile 2	38.46%
Decile 3	32.31%
Decile 4	29.67%
Decile 5	29.00%
Decile 6	29.46%
Decile 7	31.52%
Decile 8	35.41%
Decile 9	43.18%
High	64.68%
High - Low	10.79%
P-value	0.00

The average percentage of firms that change SKEW decile portfolios from one period to the next for "All" firms or after sorting firms into deciles based on SKEW STD. Significance level of High - Low is calculated using a standard t-test on comparing the mean of two series.

5.2 Portfolio Regression Analysis

In order to confirm the premium found on low $SKEW$ securities by Bali and Murray (2013), Jurczenko et al. (2002), Schneider et al. (2020). I show in table 2 the alphas that result from a three-factor regression on the equally-weighted decile portfolios based on the $SKEW$ level. The portfolios are rebalanced monthly. The table shows a clear decrease in intercept as the Skew increases and the Low minus High $SKEW$ portfolio also shows a significant positive alpha, confirming that alpha is higher on low $SKEW$ security portfolios.

Table 2: Coefficients of the Three-Factor Model

	Alpha	**Mkt-RF**	**SMB**	**HML**
Low	0.000234 ***	0.977 ***	0.136 ***	0.163 ***
Decile 2	0.000101 ***	1.001 ***	0.129 ***	0.203 ***
Decile 3	0.000227 ***	1.017 ***	0.144 ***	0.208 ***
Decile 4	0.000019 ***	1.014 ***	0.152 ***	0.209 ***
Decile 5	0.000142 ***	1.021 ***	0.111 ***	0.191 ***
Decile 6	0.000278 ***	0.998 ***	0.087 ***	0.165 ***
Decile 7	0.000090 ***	1.013 ***	0.140 ***	0.171 ***
Decile 8	0.000162 ***	1.011 ***	0.113 ***	0.162 ***
Decile 9	0.000252 ***	0.997 ***	0.111 ***	0.132 ***
High	0.000036 ***	1.019 ***	0.103 ***	0.204 ***
Low - High	0.000198 ***	-0.042 ***	0.033 ***	-0.041 ***

Regression coefficients and their statistical significance levels of the 3-factor model on equally weighted SKEW decile portfolios. Results are presented for daily stock returns. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 3 shows the five-factor regression coefficients, where the difference between model 1 and model 2 is the construction of the $SKEW\_STD$ factor mimicking portfolio. Based on the results in Table 1, I construct the model 1 $SKEW\_STD$ factor mimicking portfolio by taking a long position in the high $SKEW\_STD$ portfolio and a short position in the low $SKEW\_STD$ portfolio. In Model 2, I construct the factor mimicking portfolio by a long position in the high $SKEW\_STD$ decile and in the low $SKEW\_STD$ decile and a short position in the middle 2 $SKEW\_STD$ portfolios, as these appear to have the lowest $SKEW$ migration (Table 1). Here one can observe that, opposed to earlier research done, the $SKEW$ factor mimicking portfolio has a significant negative coefficient in both models so adding $SKEW\_STD$ leads to different results for the $SKEW$ factor. The regression analysis also shows for both types of factor construction of $SKEW\_STD$ a significant positive coefficient. Further showcasing that the factor seems to covary with returns of the assets.

Table 3: Fama-Macbeth Regression

Variable	**Model 1**	**Model 2**
Intercept	0.0001 ***	0.0001 ***
Mkt-RF	1.0015 ***	0.8565 ***
SMB	0.1281 ***	0.1047 ***
HML	0.1815 ***	0.1442 ***
Skew	−0.0201 ***	−0.0263 ***
Skew_STD	0.0657 ***	0.1531 ***

The coefficients of the intercept and each factor when regressed on daily returns. Skew is a factor mimicking portfolio constructed by taking a long position in low skew securities and a short position in high skew securities. SKEW_STD is a factor mimicking portfolio that takes in model 1 a long position in high skew std securities and a short in low skew std securities. In Model 2, it takes a long position in the high and low skew std securities, and a short position in the middle skew std securities.

Table 4 shows alpha, which results from regressing the equally-weighted portfolio returns in a double sorted quantile on the three-factor model. The high-low $SKEW$ portfolio yields a three-factor alpha, which is almost in all cases not significantly different from 0. Eventhough not significant at a 5% level, the table below supports the hypothesis tested in this paper, it shows that high $SKEW\_STD$ improves the $SKEW$ low minus high portfolio alpha. Table 4 confirms the hypothesis that $SKEW\_STD$ can enhance the predictability of the SKEW factor for returns. The dependent sort indicates that the alpha is higher for the low-minus-high $SKEW$ portfolio within the highest $SKEW\_STD$ quantile compared to the lowest $SKEW\_STD$ quantile. Sequential sorting ensures there is a similar number of stocks in each portfolio, but it could result in smaller variation of $SKEW$ levels across $SKEW$ quintiles. Therefore, for robustness purposes, I repeat the previous analyses using independent sorting of $SKEW$ levels and $SKEW\_STD$. The results confirm that the low minus high $SKEW$ portfolio yields a higher alpha within the highest $SKEW\_STD$ quantile.

Table 4: Alpha Three-Factor Table with P-Values for Low-High Skew Quantile Differences

	SKEW_STD				
SKEW	Low	2	3	4	High
			Panel A: Sequential Sorts		
Low	0.000134	0.000153*	0.000128	0.000146*	0.000256***
2	0.000313**	0.000034	0.000240***	0.000046	-0.000004
3	0.000141	0.000329**	0.000117	0.000091	0.000452*
4	0.000216**	0.000066	0.000023	0.000210***	0.000068
High	0.000249**	0.000212**	0.000180	0.000037	0.000166*
Low-High	-0.000115	-0.000059	-0.000051	0.000109	0.000090
			Panel B: Independent Sorts		
Low	0.000151***	0.000084	0.000277***	0.000072	0.000401*
2	0.000235*	0.000158**	0.000055	-0.000005	-0.000050
3	0.000217*	0.000198**	0.000075	0.000280***	0.000382
4	0.000196	0.000078	0.000116	0.000145**	0.000056
High	-0.000018	0.000707**	0.000423**	0.000162**	0.000100*
Low-High	0.000169***	-0.000624***	-0.000147	-0.000090	0.000301

Three-factor alphas of double sorted portfolios, on daily returns. Panel A has double sorted data where the first sort is SKEW and the second is SKEW_STD. Panel B shows independent sorts on both SKEW and SKEW_STD. *, **, *** indicate significance at 10%, 5%, and 1% level, respectively.

Risk Premia

In order to argue if the newly obtained factors are significantly priced, I use three approaches to check if the Risk Premia are significanlty priced.(i) The first one is the standard two-pass approach as outlined in Fama and MacBeth (1973)(FRP). (ii) Secondly, I adopt the tradeable factor risk premium approach (TFRP), (iii) and the Oracle Tradeable (Oracle TFRP) approach as outlined in Quaini et al. (2023).

Figure 1 shows the results for the first three approaches. Note that Model 1 uses the construction of the $SKEW\_STD$ factor mimicking portfolio, by taking a long position in the highest decile $SKEW\_STD$ portfolio and a short position in the lowest $SKEW\_STD$ decile portfolio. In Model 2 the factor mimicking portfolio is constructed by taking a long position in the highest and lowest $SKEW\_STD$ portfolio's and taking a short position in the middle $SKEW\_STD$ portfolio.

The Figure shows that the $Mkt - RF$ factor is consistently and significantly priced in every model, where the HML factor is not significanlty priced in either approach and has a Oracle TFRP of 0, suggesting that it is a useless or weak factor. The SMB has a significant positive risk premium in the FRP and TFRP approach, but also yields a 0 Oracle TFRP. With the inclusion of a $SKEW\_STD$ factor, the performance of a $SKEW$ factor seems weak in all results, it yields an FRP and TFRP that is small and not significant at a 5% level and it has a 0 Oracle TFRP in both models. The $SKEW\_STD$ factor as constructed in Model 1 also yields a small and insignificant FRP and TFRP, it also has a 0 Oracle TFRP. However, the Model 2 $SKEW\_STD$ factor yields more promising results, the FRP and teh TFRP are both positive and significant at a 5% level, the Oracle TFRP is also non-zero and significant, suggesting it is not a weak or useless factor.

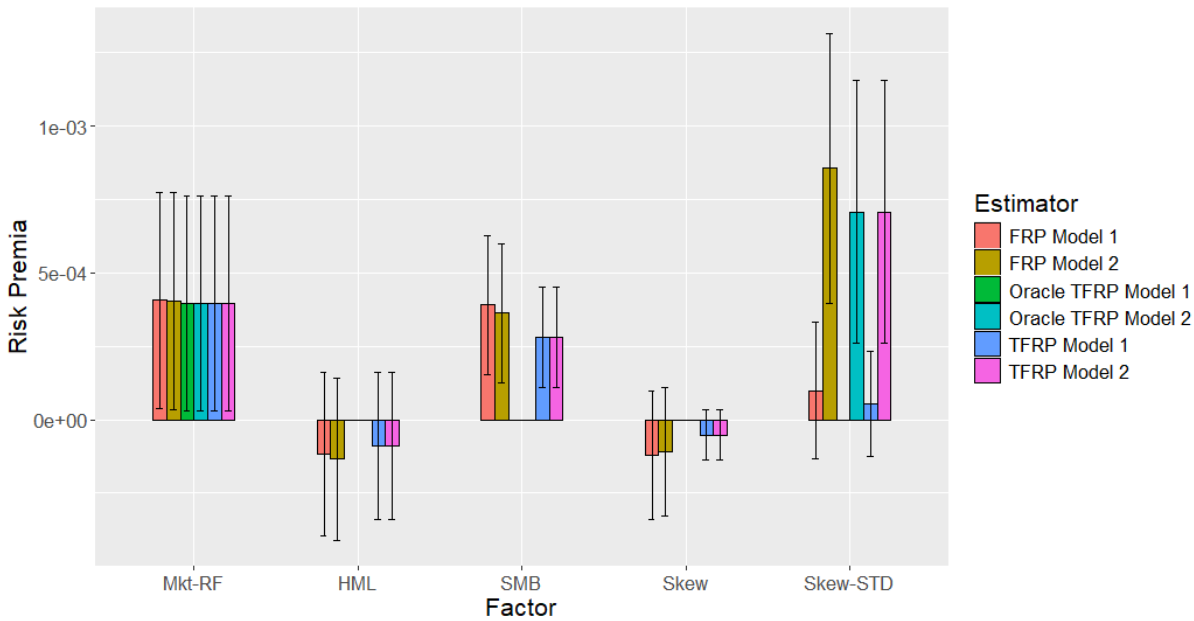

Figure 1: Risk premia using FRP, TFRP and the Oracle Estimator, with their 95% confidence bands.

Hansen-Jagannathan Distance

In order to check if the addition of the $SKEW$ and $SKEW\_STD$ to the Fama-French 3-factor model, causes the model to price the returns. I test the Hansen-Jagannathan (HJ) misspecification distance (Kan & Robotti, 2008). Where a well specified model would give a distance of 0. Figure 2 shows the squared HJ distance for Model 1, Model 2 and the Fama-French 3-Factors and the 95% confidence intervals. It shows that none of the models is well specified as all the distances are significantly larger than 0. A piecewise comparison of these distances also suggest that the models have similar specification, as the HJ distance values and 95% confidence bounds are very similar. The values for Model 2 are somewhat lower, but do not suggest any significant difference.

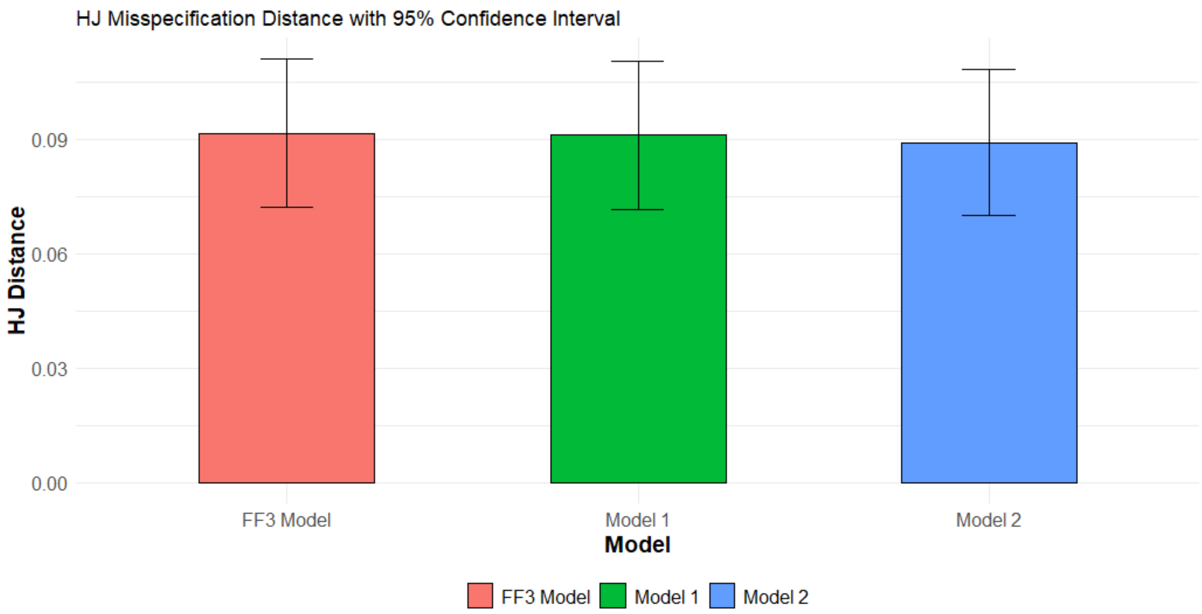

Figure 2: The figure displays the HJ misspecification distance along with 95% confidence intervals for Model 1, Model 2, and Fama-French 3-factors.

Stochastic Discount Factor Coefficients

If the law of one price holds, then there is a Stochastic Discount Factor (SDF) M pricing all excess returns, i.e., $E[M_t R_t] = 0$. A factor model for the SDF is of the form, where the mean is normalized to one since I work with excess returns. In this research that would mean that the SDF coefficient of the new factors would have a non-zero coefficients which are signficant at a 5% level. Figure 3 shows the SDF coefficients and their 95% confidence interval. It shows that in Model 1, the Fama-French 3-factors have significant coefficients that are all non-zero, here the $SKEW$ parameter has a non-zero coefficient which is nearly significant at a 5% level. The $SKEW\_STD$ factor however is not significantly different from 0 in both model 1 and 2.

In Model 2, one can observe that there is a clear interaction between the $Mkt - RF$ and $SKEW\_STD$ factor, in this model the $Mkt - RF$ appears to be negative but not significant and the $SKEW\_STD$ factor appears to be more positive than in Model 1 but it also has a larger standard error. This suggests that the $SKEW\_STD$ has a strong relationship with the average market excess returns. Overall, the figure shows, that the SDF loadings of the new $SKEW\_STD$ factors are non-zero but not significant at a 5% level. In the Appendix 3, I show that the $SKEW$ factor is significant at a 10% confidence interval. Using the Feng et al. (2020) 3-step approach, I obtain similar insignificant results as above (Appendix 3).

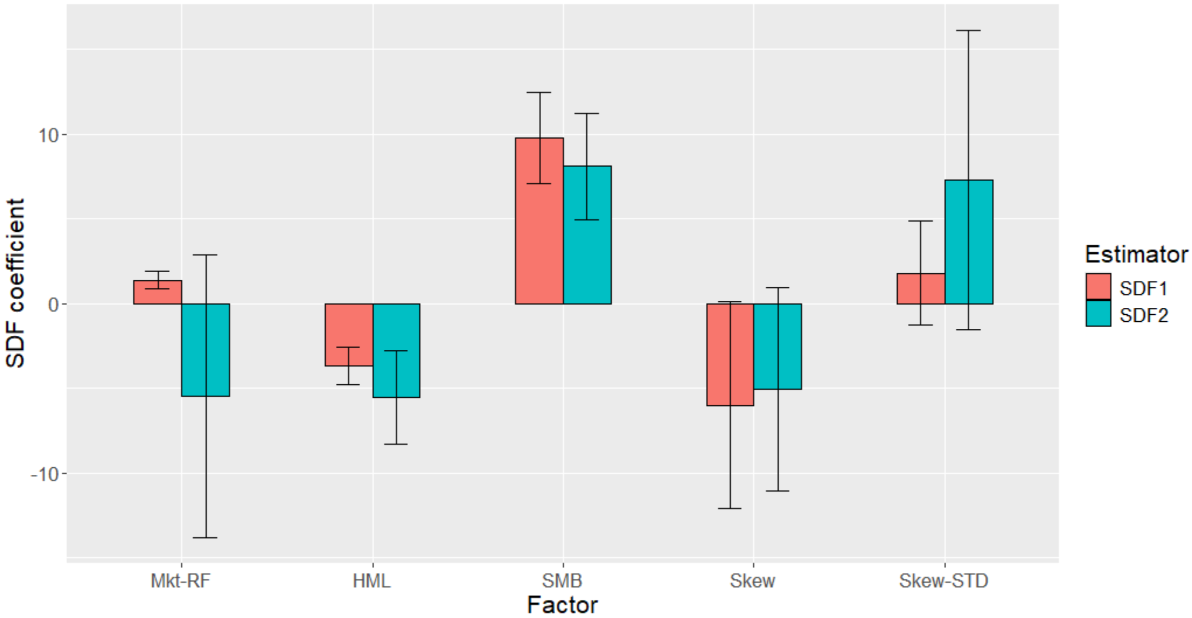

Figure 3: SDF Loadings of the five factors in Model 1 and Model 2, and there 95% confidence bands.

Based on these findings, and considering the model's lack of adequate specification, I conclude that the factors do not effectively price the assets under consideration. Therefore, further refinement and reassessment of the model's structure and variables may be warranted to improve its explanatory power and predictive accuracy.

Model Identification

At last, I perform a check on the identification of the model parameters to confirm that there is no ambiguity in the estimation of the model parameters. For this purpose, I adopt the Chen and Fang (2019) Beta Rank Test. This test is particularly insightful as it assesses the presence of a reduced rank in the matrix of regression loadings for test asset excess returns on risk factors. Intuitively, a reduced rank in this matrix suggests linear dependencies among the risk factors, indicating potential multicollinearity issues in the model estimation.

The Beta Rank Test examines the structure of the covariance matrix of the estimated coefficients. When the beta rank is reduced, it directly implies that the covariance matrix of the estimated coefficients is reduced as well. This reduction in the covariance matrix can render it non-invertible. This would cause problems as the methodology used in this paper assumes an invertible covariance matrix.

Table 5 presents the results for Models 1 and 2 obtained from the Chen-Fang (2019) Beta Rank Test. The statistically significant results with p-values close to zero indicate rejection of the null hypothesis of a reduced rank, suggesting that the models are well-specified and there is no ambiguity in the estimation of the model parameters, it also means that the covariance matrix is of full rank and thus invertible.

Table 5: Chen-Fang (2019) Beta Rank Test Results

Test	**Model 1**	**Model 2**
Statistic	1457.039	1454.691
p-value	0	0

The Chen-Fang (2019) Beta Rank Test is used to assess model identification. It tests the null hypothesis of reduced rank in the matrix of regression loadings for test asset excess returns on risk factors.

5.3 Arbitrage Risk

Based on the methodology outlined in Rehman and Vilkov (2012) and Wurgler and Zhuravskaya (2002), I aimed to investigate the relationship between option implied skewness variability ($SKEW\_STD$) and *Arbitrage Risk* in the financial markets. The rationale behind this inquiry stemmed from the hypothesis that higher arbitrage risk might lead to increased variability in skewness, reflecting the dynamic nature of market corrections and pricing inefficiencies.

To operationalize *Arbitrage Risk*, I adopt the methodology proposed by Wurgler and Zhuravskaya (2002). For each stock in our sample, I identified three substitute stocks every month, matched on industry, size, and market-to-book ratio. The selection process involved computing the sum of absolute percentage differences in size and market-to-book ratio for each stock within the same industry. The three stocks with the smallest percentage differences are chosen as substitutes. Arbitrage risk is then measured by estimating a regression model for each stock, using the returns of the selected substitute stocks and the risk-free rate. The variance of the residuals from this regression represents the level of arbitrage risk, with higher variances indicating poorer substitutes in explaining the returns of the focal stock, hence higher arbitrage risk.

To stabilize the variance and ensure homoscedasticity in the regression model, a log transformation is applied to the $SKEW\_STD$ variable. This transformation helps mitigate the effects of heteroscedasticity, where the variance of the dependent variable changes across different levels of the independent variable. Additionally, the Box-Cox (Box & Cox, 1964) transformation is utilized on the *Arbitrage_Risk* variable to normalize its distribution and address potential skewness or non-normality. By transforming the dependent variable, I aim to meet the assumption of normality in the residuals of the regression model, which is crucial for obtaining reliable parameter estimates and valid statistical inference. In appendix 5 I further show that the regression is not spurious.

Upon conducting the regression analysis, the results, as summarized in Table 6, reveal that the coefficient of $SKEW\_STD$ is not statistically significant. The low R-squared and F-statistic values further suggest that the variability in option implied skewness does not significantly explain the variation in *Arbitrage Risk* and vice versa. These findings indicate that, contrary to my initial hypothesis, there is no discernible relationship between $SKEW\_STD$ and *Arbitrage Risk* based on the selected sample and methodology.

In conclusion, the absence of a significant relationship between $SKEW\_STD$ and *Arbitrage Risk* implies that high variability in option implied skewness does not stem from *Arbitrage Risk*.

Table 6: GLS Regression Results

Variable	Coefficient	Standard Error	P-value
Constant	-4.4249	0.010	< 0.001
log_SKEW_STD	-0.0090	0.005	0.082
R-squared	0.002		
Adj. R-squared	0.001		
F-statistic	3.015		
Prob (F-statistic)	0.827		

GLS regression of *Arbitrage Risk* on log_SKEW_STD with a constant. Standard errors are reported in the second column. The F-statistic and its associated probability are provided in the last two rows.

5.4 Mean Reverting SKEW

I find, as illustrated in Table 7, a weakly significant relationship between the $SKEW\_STD$ and $AR(1)$ coefficients. The regression results show that the coefficient for the $SKEW\_STD$ variable is statistically significant but has a very small coefficient compared to the intercept value. Additionally, the adjusted R-squared value is very low at 0.021, suggesting that the $AR(1)$ coefficient does not really contribute significantly to explaining the variability in SKEW STD. In appendix 6 you find the analysis on the assumptions of the regression.

Table 7: OLS Regression Results

Variable	Coefficient	Standard Error	P-value
Constant	0.9415	0.003	0.000
SKEW_STD	−0.0013	7.5×10^{-4}	0.000
R-squared	0.021		
Adj. R-squared	0.021		
F-statistic	5791		
Prob (F-statistic)	0.000		

Regression results of regressing $SKEW\_STD$ on AR(1) coefficients. Standard Errors are shown in scientific notation. R-squared, F-statistic and Prob (F-statistic) are provided as additional information from the regression results.

Moreover, the results from Table 8 present the alphas of double sorted portfolios based on both $SKEW$ and the $AR(1)$ coefficient. The alphas for the portfolios of Low-High $SKEW$ are similar across all the $SKEW$ $AR(1)$ Coefficient quintiles in both Panels A and B. The absence of any discernable pattern in the Low-High portfolio alphas indicates that the SKEW AR(1) Coefficient is not driving the results found in the $SKEW\_STD$ analyses

Table 8: Alpha Three-Factor Table with P-Values for Low-High Skew Quantile Differences (AR(1) coefficients (value x 10^{-4}))

	AR(1) Coefficient (value x 10^{-4})				
SKEW	Low	2	3	4	High
			Panel A: Sequential Sorts		
Low	-0.07	0.98	2.49***	-0.45	-0.67
2	2.22**	2.49	0.44	0.34	0.91
3	1.78**	0.23	2.60*	2.05**	1.84**
4	2.09***	1.14	1.59	1.96**	0.96
High	2.09***	1.87**	3.22***	1.24	2.52***
Low-High	-2.83***	-0.89	-0.73	-1.69	-3.20**
			Panel B: Independent Sorts		
Low	-0.23	1.87**	1.71**	2.77***	2.03***
2	1.49*	0.17	1.44	1.27	1.78**
3	1.88**	0.68	2.13	0.15	4.64*
4	0.07	-0.18	1.30	2.34**	1.30
High	-0.07	0.82	1.77**	1.89**	2.01*
Low-High	-0.15	1.05	-0.06	0.87	0.02

Three-factor alphas of double sorted portfolios, on daily returns. Panel A has double sorted data where the first sort is SKEW and the second is the AR(1) coefficient. Panel B shows independent sorts on both SKEW and the AR(1) coefficient. *, **, *** indicate significance at 10%, 5%, and 1% level, respectively.

5.5 Transaction costs

In understanding the implications of our findings, it's crucial to consider transaction costs, which can significantly impact the practical implementation of investment strategies. To evaluate this aspect, an examination of the migrations of firms across deciles based on their $SKEW$ and $SKEW\_STD$ values was conducted. Table 9 provides insights into the average percentage of firms transitioning between $SKEW$ or $SKEW\_STD$ decile portfolios from one period to the next, with portfolios rebalanced at monthly frequencies.

It's noteworthy that higher turnover rates are observed for extremity portfolios, particularly those associated with high $SKEW$ and $SKEW\_STD$ values. This implies that if these strategies were to be employed in practice, the transaction costs incurred during portfolio rebalancing could be quite substantial. Moreover, the findings of higher alpha in stocks with high $SKEW\_STD$ and high $SKEW$ values also comes with the caveat of potentially higher transaction costs associated with frequent rebalancing of the strategy.

However, it's essential to acknowledge that the earlier results presented did not incorporate these transaction costs. Hence, future studies should delve deeper into turnover dynamics and explore ways to mitigate transaction costs while still capitalizing on the alpha opportunities presented by $SKEW$ and $SKEW\_STD$ factors. By optimizing trading strategies and considering transaction costs in portfolio construction, practitioners can enhance the efficiency and effectiveness of their investment approaches.

Table 9: Turnover of Decile Portfolios for SKEW and SKEW_STD

Decile	SKEW Turnover (%)	SKEW_STD Turnover (%)
Low	46.75%	54.15%
Decile 2	31.72%	37.06%
Decile 3	26.71%	31.51%
Decile 4	24.43%	29.26%
Decile 5	23.58%	28.72%
Decile 6	23.69%	29.27%
Decile 7	24.57%	31.13%
Decile 8	27.13%	34.67%
Decile 9	32.65%	42.52%
High	50.51%	63.79%

Note: The table presents the average percentage of firms that change SKEW or SKEW_STD decile portfolios from one period to the next for "All" firms. Portfolios are rebalanced at monthly frequencies.

6 Conclusion

This study aimed to explore the predictive capacity of option implied skewness variability concerning stock returns. Employing a multifaceted approach involving portfolio sorts, regression analysis, risk premia assessment, examination of Stochastic Discount Factor (SDF) coefficients, model misspecification and identification tests, I have gained substantial insights into the interplay between option implied skewness and stock returns.

The findings present robust evidence supporting the pivotal role of option-implied skewness variability, as captured by $SKEW\_STD$, in explaining the cross-section of stock returns. While previous studies have predominantly focused on option implied skewness ($SKEW$) in isolation, my investigation underscores the significance of incorporating skewness variability as an additional determinant in models explaining the cross-section of returns.

The portfolio sorts analysis revealed notable turnover within decile portfolios corresponding to fluctuations in $SKEW$ or $SKEW\_STD$ values, suggesting potential predictability associated with skewness variability. Additionally, regression analysis showed a discernible increase in the intercept as skewness variability ($SKEW\_STD$) increases, with the low-minus-high SKEW portfolio displaying a significant positive alpha.

Moreover, my assessment of risk premia through diverse methodologies highlighted the augmenting effect of skewness variability on the predictive power of return models. While the market risk factor ($Mkt - RF$) retained consistent and significant pricing, the incorporation of skewness variability in Model 2 yielded a substantial and positive risk premium, accentuating its influence on returns. Using the Oracle Estimator, which discards useless and weak factors, I conclude that the skewness variability factor is neither useless nor weak, implying that the factor holds significant explanatory power.

As I still work with low dimensional models, I find that the models are all misspecified, meaning that the model does not price all the assets. These findings are confirmed by the non-significant SDF loadings, which confirms that indeed using these models the law of one price does not hold and the factors do not linearly price the assets. However it is interesting to observe a pronounced interaction between the market risk factor and skewness variability, proposing further research into the dynamics of the skewness variability.

In summary, the research adds new understanding to the current body of knowledge by highlighting the importance of considering variation in option implied skewness when predicting stock returns. By incorporating $SKEW\_STD$ alongside the traditional $SKEW$ measure in models predicting returns, investors may improve the precision of their forecasts and make wiser investment choices. Future studies could focus on refining existing models and exploring additional variability factors to boost predictive accuracy in financial markets. It could also focus on the dynamics between option implied skewness variability and market returns.

References

Almeida, C. & Garcia, R. (2017). Economic implications of nonlinear pricing kernels. *Management Science*, *63*(10), 3361–3380.

Amihud, Y. & Mendelson, H. (1986). Asset prices and the bid-ask spread. *Journal of financial economics*, *17*(2), 223–249.

Ang, A., Hodrick, R. J., Xing, Y. & Zhang, X. (2009). High idiosyncratic volatility and low returns: International and further us evidence. *Journal of Financial Economics*, *91*(1), 1–23.

Aschakulporn, P. & Zhang, J. E. (2022). Bakshi, kapadia, and madan (2003) risk-neutral moment estimators: A gram–charlier density approach. *Review of Derivatives Research*, *25*(3), 233–281.

Bakshi, G., Kapadia, N. & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, *16*(1), 101–143.

Bali, T. G., Cakici, N., Chabi-Yo, F. & Murray, S. (2017). The risk-neutral distribution of option returns. *Georgetown McDonough School of Business Research Paper*(2902209).

Bali, T. G. & Hovakimian, G. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, *99*(2), 427–446.

Bali, T. G. & Murray, S. (2013). Does risk-neutral skewness predict the cross section of equity option portfolio returns? *The Journal of Financial and Quantitative Analysis*, *48*(4), 1145–1171. doi: 10.1017/S0022109013000410

Barillas, F., Kan, R., Robotti, C. & Shanken, J. (2020). The cross-section of risks in the covered interest parity condition. *Journal of Financial Economics*, *138*(3), 730–759.

Bekaert, G. & Hoerova, M. (2014). Understanding aggregate risk: The capm with countercyclical idiosyncratic risk. *Journal of Political Economy*, *122*(5), 1093–1140.

Bintara, R., Wahyudi, S. M. et al. (2019). Analysis of fundamental factors on stock price. *International Journal of Academic Research in Accounting, Finance and Management Sciences*, *9*(4), 169–181.

Box, G. E. & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, *26*(2), 211–252.

Bressan, S. & Weissensteiner, A. (2023). Option-implied skewness and the value of financial intermediaries. *Journal of Financial Services Research*, *64*(2), 207–229.

Britten-Jones, M. & Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, *55*(2), 839–866.

Chabi-Yo, F. (2012). Pricing kernels with stochastic skewness and volatility risk. *Management Science*, *58*(3), 624–640.

Chen, X. & Fang, X. (2019). An improved bootstrap-based test of asset pricing models. *Journal of Financial Econometrics*, *17*(4), 521–549.

Chordia, T., Lin, T.-C. & Xiang, V. (2021). Risk-neutral skewness, informed trading, and the cross section of stock returns. *Journal of Financial and Quantitative Analysis*, *56*(5), 1713–1737.

Cremers, M. & Weinbaum, D. (2010a). Active portfolio management and portfolio construction. *Journal of Portfolio Management*, *36*(3), 121–130.

Cremers, M. & Weinbaum, D. (2010b). Deviations from put-call parity and stock return predictability. *The Journal of Financial and Quantitative Analysis*, *45*(2), 335–367. Retrieved 2024-01-24, from http://www.jstor.org/stable/27801488

DeLisle, R. J., Diavatopoulos, D., Fodor, A. & Kassa, H. (2021). Variation in option implied volatility spread and future stock returns. *The Quarterly Review of Economics and Finance*, *83*, 152–160. doi: 10.1016/j.qref.2021.12.004

Durbin, J. & Watson, G. M. (1950). Testing for serial correlation in least squares regression. i. *Biometrika*, *37*(3/4), 409–428.

Fama, E. F. & French, K. R. (1992). The cross-section of expected stock returns. *the Journal of Finance*, *47*(2), 427–465.

Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, *33*(1), 3–56.

Fama, E. F. & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy*, *81*(3), 607–636.

Fassas, A. P. & Siriopoulos, C. (2021). Implied volatility and stock returns: A review of the literature. *International Review of Financial Analysis*, *74*, 101681.

Feng, L., Giglio, S. & Xiu, D. (2020). Taming the factor zoo: A test of new factors. *The Review of Financial Studies*, *33*(5), 2231–2285.

Gibbons, M. R., Ross, S. A. & Shanken, J. (1989). A test of the efficiency of a given portfolio. *Econometrica: Journal of the Econometric Society*, 1121–1152.

Gonzalez-Perez, M. T. (2015). Stock return predictability and the implied volatility spread. *Journal of Banking & Finance*, *58*, 221–234.

Gospodinov, N., Kan, R., Robotti, C. & Shanken, J. (2014). A new specification test for the validity of the stochastic discount factor using hansen-jagannathan bounds. *The Journal of Finance*, *69*(6), 2737–2784.

Hansen, L. P. & Jagannathan, R. (1997). Assessing specification errors in stochastic discount factor models. *The Journal of Finance*, *52*(2), 557–590.

Harvey, C. R. & Siddique, A. (1999). Autoregressive conditional skewness. *Journal of financial and quantitative analysis*, *34*(4), 465–487.

Harvey, C. R. & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, *55*(3), 1263–1295. Retrieved 2024-01-24, from http://www.jstor.org/stable/222452

Jegadeesh, N. & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, *48*(1), 65–91.

Jurczenko, E., Maillet, B. & Negrea, B. (2002). Skewness and kurtosis implied by option prices: a second comment.

Kan, R. & Robotti, C. (2008). Model comparison using the hansen-jagannathan distance. *Journal of Financial Econometrics*, *6*(1), 123–170.

Kan, R. & Robotti, C. (2010). On the hansen-jagannathan distance with a no-arbitrage constraint. *Journal of Business & Economic Statistics*, *28*(2), 312–328.

Kan, R., Robotti, C. & Shanken, J. (2013). Pricing model performance and the two-pass cross-sectional regression methodology. *The Journal of Finance*, *68*(6), 2617–2649.

Kenton, W. (2023). *Volatility skew: How it can signal market sentiment.* https://www.investopedia.com/terms/v/volatility-skew.asptoc-why-would-volatility-skew. (Updated September 06, 2023. Reviewed by Gordon Scott. Fact checked by Kirsten Rohrs Schmitt.)

Le Courtois, O. & Xu, X. (2023). Efficient portfolios and extreme risks: a pareto–dirichlet approach. *Annals of Operations Research*, 1–32.

Liu, Z. F. & van der Heijden, T. (2016). Model-free risk-neutral moments and proxies. *Thijs, Model-Free Risk-Neutral Moments and Proxies (July 4, 2016).*

Mittal, S. K. & Srivastava, N. (2021). Mean-variance-skewness portfolio optimization under uncertain environment using improved genetic algorithm. *Artificial Intelligence Review*, 1–22.

Newey, W. K. & West, K. D. (1994). Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, *61*(4), 631–653.

Ni, Z. & Wang, L. (2023). The predictability of skewness risk premium on stock returns: Evidence from chinese market. *International Review of Economics and Finance*, *87*, 576–594. doi: 10.1016/j.iref.2023.05.010

Park, Y.-H. (2013). Volatility of volatility and tail risk premiums. *Finance and Economics Discussion Series*, *2013-542*. Retrieved from https://www.federalreserve.gov/pubs/feds/2013/201354/201354pap.pdf

Quaini, A., Trojani, F. & Yuan, M. (2023). Tradable factor risk premia and oracle tests of asset pricing models. *Swiss Finance Institute Research Paper*(23-81).

Rehman, Z. & Vilkov, G. (2012). Risk-neutral skewness: Return predictability and its sources. *Available at SSRN 1301648*.

Schneider, P. & Trojani, F. (2015). Fear trading. (Retrieved from https://ssrn.com/abstract=1994454)

Schneider, P., Wagner, C. & Zechner, J. (2020). Low-risk anomalies? *The Journal of Finance*, *75*(5), 2673–2718.

Ulrich, M. & Walther, S. (2020). Option-implied information: What's the vol surface got to do with it? *Review of Derivatives Research*, *23*(3), 323–355.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: Journal of the Econometric Society*, *48*(4), 817–838.

Wurgler, J. & Zhuravskaya, E. (2002). Does arbitrage flatten demand curves for stocks? *The Journal of Business*, *75*(4), 583–608.

1 Appendix A: missing securities of the S&P 500 in the OptionMetrics database

The securities below have no available data on traded options in the OptionMetrics database: NVR, VLTO, KVUE, BRK.B, FI, RVTY, WHR, EG, BF.B, DAY.

2 Appendix B: Data Summary

Table 10 displays for each year the number of recorded days in that year, in order to check for daycount conventions and irragulaties. It also shows the signficant increase in number of securities recorded in the data. The second last column shows the proportion of NA values in the SKEW column of the created dataset. A NA values appears, if there is not sufficient data available on certain securities in order too accurately calculate the SKEW as descriped in the paper. As the factor mimicking portfolio's are based on decile portfolio's of the securities sorted on the SKEW factor, I also know the size of the average decile portolio. Which is calculated as

$$Size = UniqueSecurities * (1 - NAProportion/100)/10. \tag{39}$$

One can observe that the decile size is quickly increasing and becomes somewhat stable after 2008.

Table 10: Summary of Data Characteristics by Year

Year	Trading Days	Unique Securities	NA Proportion (%)	Average decile Size
2000	250	336	17.98	27.56
2001	248	338	18.73	27.47
2002	252	361	20.04	28.87
2003	252	367	21.28	28.89
2004	252	381	21.13	30.05
2005	252	387	18.66	31.48
2006	251	403	16.05	33.83
2007	251	416	13.05	36.17
2008	253	422	1.52	41.56
2009	252	426	1.93	41.78
2010	252	432	3.04	41.89
2011	252	444	5.12	42.13
2012	250	448	7.06	41.64
2013	252	460	0.29	45.87
2014	252	466	0.65	46.30
2015	252	474	2.25	46.33
2016	252	478	3.22	46.26
2017	251	481	3.95	46.20
2018	251	484	0.50	48.16
2019	252	490	0.69	48.66
2020	253	495	1.37	48.82
2021	252	495	0.35	49.33
2022	251	495	0.66	49.17
2023	39	496	0.89	49.16

This table displays the number of trading days, unique securities, NA proportion in the SKEW column, and average decile size by year in the dataset.

3 Appendix C: SDF Coefficients

In order to test for weak significance in the $SKEW$ factor I also test the factors SDF coefficient with a 10% confidence interval. Figure 4 displays the results and shows that the $SKEW$ factor has become significant in both models.

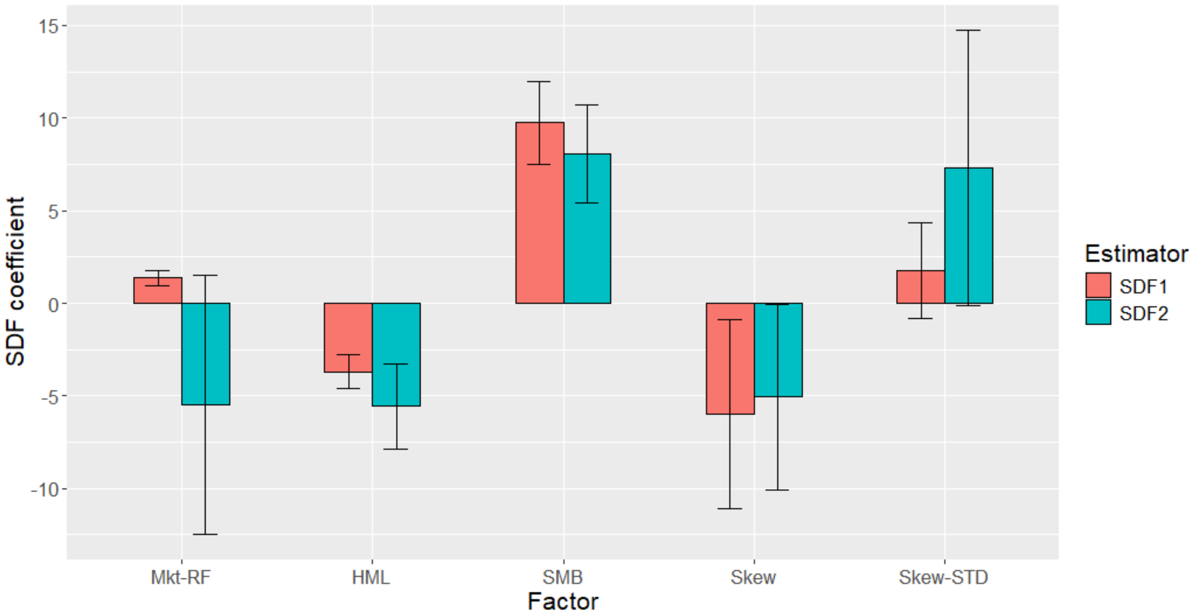

Figure 4: SDF Loadings of the five factors in Model 1 and Model 2, and there 90% confidence bands.

By using a factor selection approach as outlined in Feng et al. (2020) where a double lasso is utilized to select strong factors, the third step is a OLS of average returns on the covariances between asset returns and the new factors. I show in table 11 the SDF coefficients and there standard error of the three-step procedure proposed by Feng et al. (2020). It shows that when using the Fama-French 3-factors (Mkt-Rf, HML and SMB) as control factors, the $SKEW$ and $SKEW\_STD$, in both Model 1 and 2, yield non-zero but non-significant SDF loadings. Suggesting that the factors do not have an SDF loading and thus do not covary with the risk-neutral density.

Table 11: Feng, Giglio and Xiu SDF factor selection

Factor	SKEW	SKEW_STD_Model_1	SKEW_STD_Model_2
Coefficient	-18.819	1.571	1.146
Standard Errors	180.934	17.745	23.898

The SDF coefficients for the FGX model, including standard errors. Using Mkt-RF, SMB and HML as control factors.

4 Appendix D: Robustness

To assess the robustness of the risk premia associated with the second derived factor, $SKEW\_STD$, I evaluate its performance within the context of established factor-based pricing models. Specifically, I examine its behavior within the frameworks of the Fama-French 5-factor model and the Carhart 4-factor model, both of which represent extensions of the seminal Fama-French 3-factor model.

The Fama-French 5-factor model expands upon the Fama-French 3-factor model by incorporating two additional factors: profitability (RMW) and investment (CMA). In my analysis, I augment the Fama-French 5-factor model with the inclusion of the $SKEW$ and $SKEW\_STD$ factors. The $SKEW\_STD$ factor is constructed by taking a long position in the highest and lowest deciles of $SKEW\_STD$ and a short position in the middle decile portfolio. Remarkably, the results obtained align closely with those presented in the earlier sections. Figure 5 illustrates the risk premia derived within this extended Fama-French 5-factor framework.

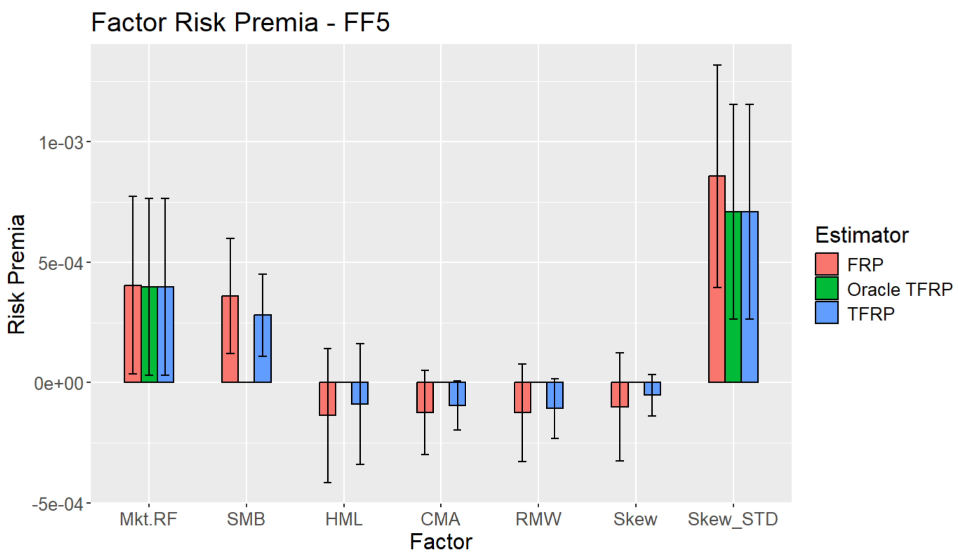

Figure 5: Risk Premia for the Fama-French 5-Factor Model

Similarly, I conduct a parallel examination within the Carhart 4-factor model context, which builds upon the Fama-French 3-factor model by adding a momentum (MOM) factor. Once more, the analysis reveals consistent results for the risk premia associated with both $SKEW$ and $SKEW\_STD$.

This rigorous evaluation underscores the stability and robustness of the risk premia estimates across various factor-based pricing frameworks, affirming the reliability of the findings presented herein.

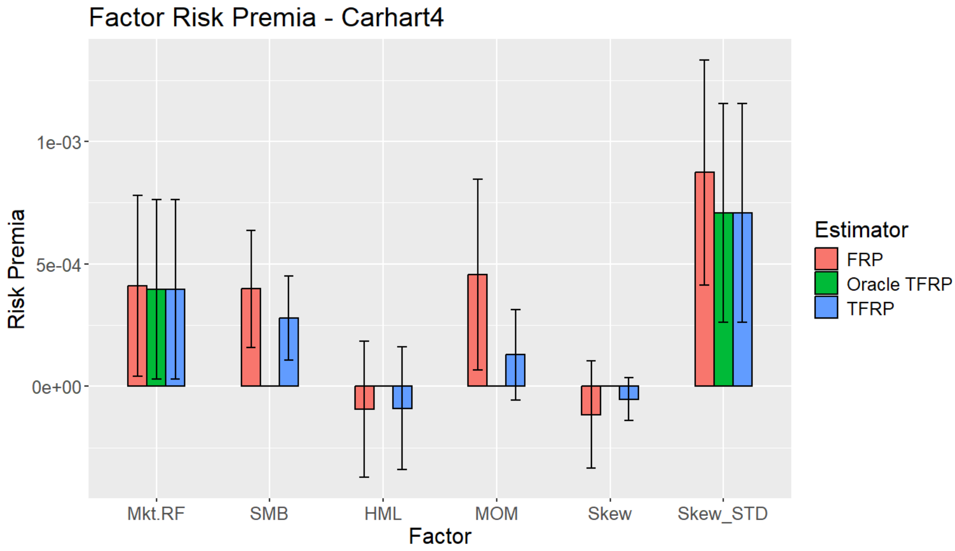

Figure 6: Risk Premia for the Carhart 4-Factor Model

5 Appendix E: Arbitrage Risk residuals

The regression analysis conducted to investigate the relationship between option implied skewness variability ($SKEW\_STD$) and *Arbitrage Risk* produced results consistent with the underlying assumptions of linear regression.

Firstly, the Q-Q plot of the residuals (see Figure 7) exhibits a pattern closely aligned with the diagonal line, indicating that the residuals follow a normal distribution. This suggests that the assumption of normality of residuals is not violated, providing confidence in the robustness of the regression analysis.

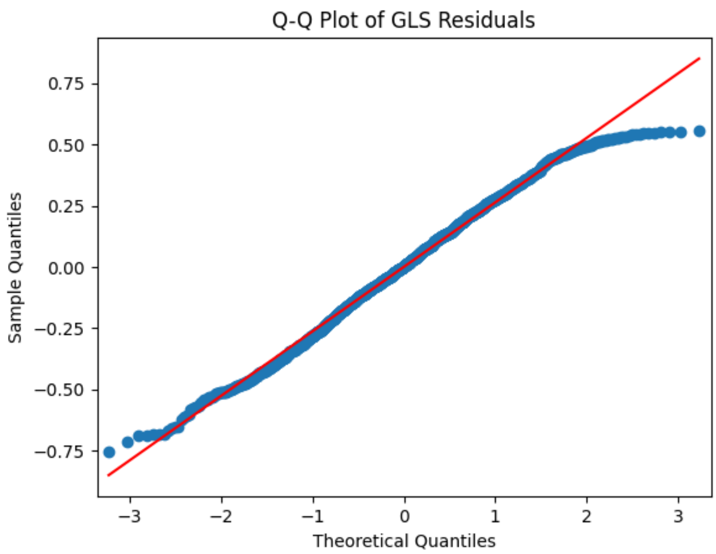

Figure 7: Q-Q Plot of Residuals

Secondly, the scatter plot of residuals against fitted values (see Figure 8) displays a random scatter pattern without any discernible trends or patterns. The spread of the residuals appears constant across different levels of the independent variable, indicating that the assumption of homoscedasticity is met. Therefore, the variance of the residuals remains consistent, validating the reliability of the regression results. These diagnostic plots reassure that the regression model

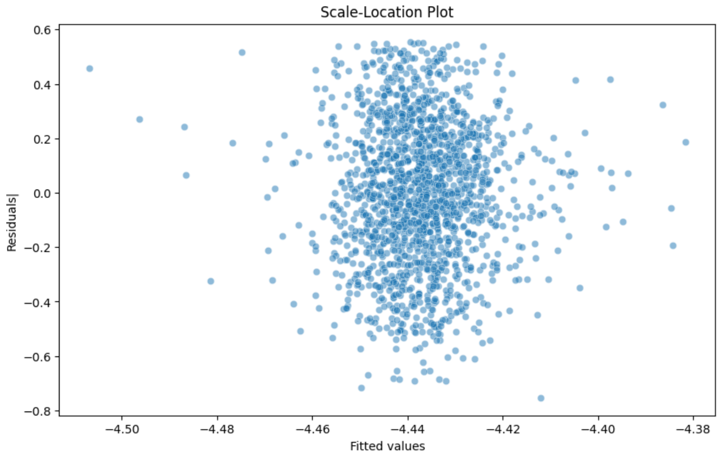

Figure 8: Scatter Plot of Residuals vs Fitted Values

is not spurious and provides a valid representation of the relationship between $SKEW\_STD$ and *Arbitrage Risk* in the financial markets.

The Durbin-Watson (Durbin & Watson, 1950) in 6 with a test statistic of 1.245 suggests the absence of significant autocorrelation in the residuals. With a value close to 2, the residuals exhibit no apparent pattern of autocorrelation, supporting the assumption of independence of observations.

To ensure for homoscedasticity, I use robust residuals employing the Huber-White (White, 1980) sandwich estimator. This approach enhances the robustness of the regression analysis, particularly in the presence of non-constant variance in the residuals.

6 Appendix F: Mean reversion regression

The regression analysis conducted to investigate the relationship between option implied skewness variability ($SKEW\_STD$) and the $AR(1)$ $Coefficient$ produced results consistent with the underlying assumptions of linear regression. To manage outliers, I applied Winsorization, capping extreme values at two standard deviations from the mean.

The Q-Q plot of the residuals (see Figure 9) exhibits a pattern closely aligned with the diagonal line, indicating that the residuals follow a normal distribution. This suggests that the assumption of normality of residuals is not violated, providing confidence in the robustness of the regression analysis.

The Durbin-Watson test (Durbin & Watson, 1950) in 6 with a test statistic of 1.922 suggests

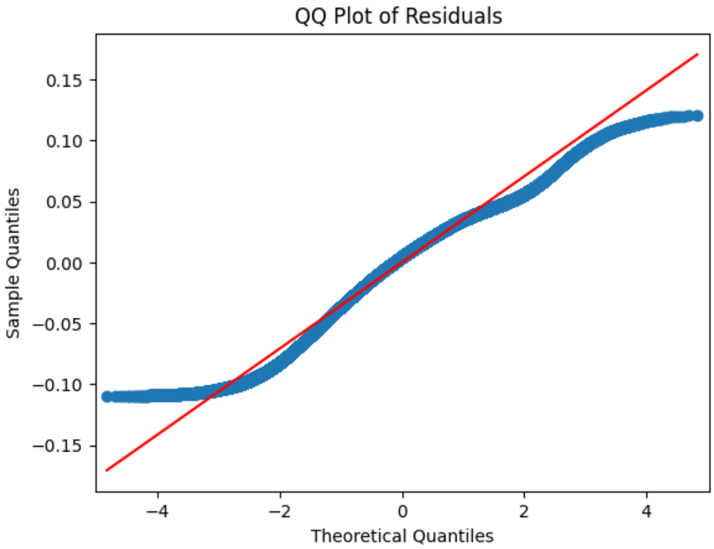

Figure 9: Q-Q Plot of Residuals

the absence of significant autocorrelation in the residuals. With a value close to 2, the residuals exhibit no apparent pattern of autocorrelation, supporting the assumption of independence of observations.

To ensure for homoscedasticity, I use robust residuals employing the Huber-White (White, 1980) sandwich estimator. This approach enhances the robustness of the regression analysis, particularly in the presence of non-constant variance in the residuals.