City Research Online

City, University of London Institutional Repository

Citation: Kozhan, R., Neuberger, A. and Schneider, P. (2013). The Skew Risk Premium in the Equity Index Market. *The Review of Financial Studies*, 26(9), pp. 2174-2203. doi: 10.1093/rfs/hht039

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/15209/>

Link to published version: <http://dx.doi.org/10.1093/rfs/hht039>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

The Skew Risk Premium in the Equity Index Market ^{*}

Roman Kozhan[†] Anthony Neuberger[†] Paul Schneider[†]

November 27, 2012

Abstract

We measure the skew risk premium in the equity index market through the skew swap. We argue that just as variance swaps can be used to explore the relationship between the implied variance in option prices and realized variance, so too can skew swaps be used to explore the relationship between the skew in implied volatility and realized skew. Like the variance swap, the skew swap corresponds to a trading strategy, necessary to assess risk premia in a model-free way. We find that almost half of the implied volatility skew can be explained by the skew risk premium. We provide evidence that skew and variance premia are manifestations of the same underlying risk factor in the sense that strategies designed to exploit one of the risk premia but to hedge out the other make zero excess returns.

^{*}We are thankful to Mark Britten-Jones, Peter Carr, Fousseni Chabi-Yo, Bjørn Eraker, Gordon Gemmill, Stewart Hodges, Ciara Rogerson, Thomas Ruf, Mark Salmon, Lucio Sarno, Christian Schlag, Anders Trolle, Andrea Vedolin, Pierre de Vos, Christian Wagner, Liuren Wu and to participants at EFA 2010, AFA 2011, seminars at Cass Business School, Essex University, Lancaster University, Queen's University Belfast, Warwick Business School, Wisconsin Business School, Zurich University, and the WBS conference on Derivatives, Volatility, and Correlation for valuable comments.

[†]Warwick Business School, University of Warwick, Scarman Road, Coventry, CV4 7AL, UK; e-mail: {roman.kozhan,anthony.neuberger,paul.schneider}@wbs.ac.uk

1 Introduction

The implied volatility skew has been a well-recognized feature of the equity index market ever since the crash of 1987. The skew is due to the asymmetric distribution of index returns and to a skew risk premium. This paper seeks to quantify the relative importance of these two factors. It does so by looking at the profitability of a skew swap, a contract that pays the difference between the implied skew and the realized skew, which can be synthesized from standard options and forward contracts. We show that the risk premium accounts for about half the implied skew. We also find that the skew risk premium and the variance risk premium compensate for the same risk.

Risk premia cannot be measured directly because the conditional physical density is not observable; all that can be observed is a particular realization of the distribution. Researchers commonly estimate risk premia by repeating a trading strategy over a number of periods and testing the trading profits for predictability. [Coval and Shumway \(2001\)](#) and [Bakshi and Kapadia \(2003\)](#) look at holding returns on option positions – typically at-the-money (ATM) – that are delta neutral, and show that the average return is significantly negative. Since the principal risk to which delta-hedged options are exposed is volatility risk, this can be seen as evidence for the existence of a volatility risk premium. [Bollen and Whaley \(2004\)](#), and [Driessen and Maenhout \(2007\)](#) have documented the high negative returns earned by the buyers of out-of-the-money (OTM) index puts. This can be interpreted as evidence of a specific premium for crash risk. [Bakshi et al. \(2010\)](#) show that there are also negative returns from holding high strike calls and digitals, and interpret this as evidence for a U-shaped pricing kernel.

These buy-and-hold strategies face a generic problem. As options age, their risk characteristics change. Their exposure to particular risk factors depends not only on the passage of time but also on the evolution of the price level and of the volatility surface. An ATM index straddle may start with zero market exposure and high volatility exposure, but if the market moves far from its initial value, the position will have high market exposure and little volatility exposure. It is hard to draw inferences from the profitability of an option strategy to the pricing of risk factors in the economy when the strategy's exposure to those factors is itself stochastic. Nor is it easy to tell whether two profitable strategies are manifestations of the same or different priced risks. In the case of the negative OTM put returns for example, [Bondarenko \(2003\)](#) and [Broadie et al. \(2009\)](#) question the statistical significance of the trading profits (since they are so highly skewed) and also argue that the negative buy-and-hold OTM put returns could be explained by other risk factors including market and volatility risk.

The obvious solution to the ageing problem is to rebalance the portfolio so that its risk exposure remains constant. The challenge is to do so in a way that is not model-dependent. As [Anderson et al. \(2000\)](#) have noted, use of a misspecified model can generate spurious risk premia. Within a given model, implementation of a dynamic hedging strategy requires the choice of hedging securities, and the profits from the strategy may well be sensitive to the particular securities chosen, making the interpretations of profitability still more difficult. In the case of variance risk, [Carr and Wu \(2009\)](#) solve the problem by using a variance swap which can be synthesized from vanilla options. The key feature of the swap is that innovations in its mark-to-market value are perfectly correlated with innovations in the implied variance of the market - in other words, it is a pure play on variance.

In this paper we use an analogous trading strategy, which we call a skew swap, that is a pure play on skew. It involves holding a portfolio of long OTM calls and short OTM puts. Risk characteristics of the portfolio are maintained by subsequently trading the options. This is done in a way that is effectively model-free. The fixed leg of the skew swap is dependent on the implied skew at inception; if the implied volatility smile is symmetric, the value is identically zero. The floating leg turns out to equal the covariation between the returns on the index and changes in its implied variance. Following [Neuberger \(2012\)](#), we argue that it is natural to call this covariation the realized skew. The skew swap has a simple interpretation: the fixed leg is the implied skew and the floating leg is the realized skew. The expected difference between the implied skew and the realized skew, the expected payoff from the skew swap, is the skew risk premium.

The implied skew measure we develop differs from that of [Bakshi et al. \(2003\)](#). The reason for departing from the standard definition is that we want to use the difference between the implied skew and the physical skew as an estimate of the skew risk premium. With the standard definition, the two can differ in expectation even when there is no risk premium.¹ With our definition, the difference between the physical and implied quantities has a mean of zero in the absence of risk premia.

We apply the skew swap technology to S&P 500 option data, and provide evidence that there does indeed exist a skew risk premium in the US equity index market. We show that almost half of the implied skew in index option prices with one month to maturity can be explained by the risk premium, with the other half reflecting the negative correlation between returns and volatility. Skew risk turns out to be closely related to variance risk. The excess return on skew swaps is highly correlated to the excess return on variance swaps (in our data

¹This claim is supported by an example in [Appendix A](#).

the correlation is nearly 0.9). The profitability of betting on the slope of implied volatility is thus highly correlated with the strategy of betting on its level. But this high correlation depends on continual rebalancing to ensure that the two strategies remain pure bets on skew and variance, and do not pick up other exposures. Under the corresponding buy-and-hold strategies the correlation between the excess returns on skew and variance drops to 0.3.

We also show that the risk premia associated with variance and skew have a common source. When one type of swap is hedged with the other, the expected profit is indistinguishable from zero. Furthermore, there is consistent evidence that the risk premium on both skew swaps and variance swaps is time-varying, and that a common factor drives both premia. This finding is important for asset pricing. It suggests that models of equity market returns, while they need to take account of the clear evidence for the existence of both variance and skew premia, do not need to allow for separate channels to generate them.

Our results are related to a growing literature uncovering spurious risk factors. In particular, [Ferson and Harvey \(1999\)](#) and [Ferson et al. \(1999\)](#) show that the failure to account properly for time variation in market exposure may lead to spurious evidence of additional priced risk factors in cross-section of stock returns. In examining hedge fund performance, [Bollen and Whaley \(2009, Abstract\)](#) note that “The standard measure of performance is the abnormal return defined by a hedge fund’s exposure to risk factors. If exposures are assumed constant when, in fact, they vary through time, estimated abnormal returns may be incorrect.”

The remainder of the paper is organized as follows: Section 2 sets out the model-free approach to constructing the skew swap and shows how to quantify the skew risk premium. Section 3 describes the data. Section 4 presents the empirical results, and Section 5 concludes. Section A of the Appendix discusses the use of central moments in risk premium measurement, Section B contains proofs of propositions not included in the main text, and Section C contains robustness checks.

2 Trading Variance and Skew

This section sets out our approach to defining and estimating the risk premia associated with moments of returns. We show how to construct a trading strategy which is exposed to the desired return moment alone. The risk premium is the expected profit from such a strategy. The estimation of the variance risk premium as the expected profit from a variance swap ([Carr and Wu 2009](#)) is a special case of our general approach.

Previous work on risk premia, whether in equity, bond, currency or other asset markets,

has largely been concerned with the first moment of returns. Researchers examine the predictability of excess returns in order to estimate the asset risk premium. An excess return can be interpreted as the profit from a trading strategy. In this case the strategy is to enter into a forward contract on the asset at time t , and hold it until expiry at time T . In a market free from frictions and other imperfections, the expected excess return is interpreted as a premium for risk bearing.

We extend this approach to estimate risk premia associated with other moments of returns. To do this, one must be able to trade options on the underlying asset as well as forward contracts as we will show below. We take a world with a single risky asset with a forward contract traded on it with fixed maturity T . The forward price of the asset at time t is $F_{t,T}$, the log forward price is $f_{t,T}$ and the log return from t to T is $r_{t,T} \equiv f_{T,T} - f_{t,T}$. With calls and puts available for all strikes, one can create a portfolio of options that has any desired payoff using the results from [Bakshi and Madan \(2000\)](#) and [Carr and Madan \(2001\)](#). In particular, one can buy a portfolio at time t that has payoff at time T of $g(r_{t,T})$ for an arbitrary (twice-differentiable) function g . There is also a bond in the market with price $B_{t,T}$, where $B_{T,T} = 1$.

Denoting the physical measure by $\mathbb{P}$ and the forward pricing measure for a maturity T by $\mathbb{Q}$, the forward price of the payoff g is

$$G_{t,T} \equiv \mathbb{E}_t^{\mathbb{Q}} [g(r_{t,T})]. \quad (1)$$

The net profit at time T from the buy-and-hold strategy is

$$g(r_{t,T}) - \mathbb{E}_t^{\mathbb{Q}} [g(r_{t,T})], \quad (2)$$

and the risk premium associated with this strategy is

$$\mathbb{E}_t^{\mathbb{P}} [g(r_{t,T})] - \mathbb{E}_t^{\mathbb{Q}} [g(r_{t,T})]. \quad (3)$$

As a way of measuring the risk premium associated with the function g over the period $[t, T]$, this buy-and-hold strategy has two defects: it is in general exposed to directional risk, so the estimated risk premium is contaminated by any risk premium that is associated with a position in the underlying asset. The second defect is that the risk characteristics of the strategy change over time. At a future time τ the investor is holding a contract that pays $g(r_{t,T})$ and not $g(r_{\tau,T})$, meaning that the risk exposure at time τ depends on the essentially

arbitrary choice of the point of inception of the strategy.

To formalize the notion of a strategy that is a pure bet on g -risk we normalize $g(0) = 0$, which implies that $G_{T,T} = 0$. Given a self-financing trading strategy, we introduce the *wealth process* W_t as the profit from the strategy to date. More specifically, we define W_t as the value of the portfolio at time T if the portfolio were liquidated at time t and the proceeds were invested in the risk-free bond until time T . We normalize $W_0 = 0$. All transactions take place in the forward market, and no cash flows occur prior to time T . Define the *residual process* as

$$Y_{t,T} \equiv W_{t,T} - G_{t,T}. \quad (4)$$

If a trading strategy has a residual process that is predictable, so the stochastic components of the increments in G and W are the same, we call a trading strategy a *perfect hedge* for g . If such a trading strategy exists (and we show below that it does, and how to construct it) then the expected change in W can be interpreted as the risk premium associated with g .

To remove the exposure to the market we introduce delta hedging. We define the *delta* of g as

$$\Delta_{t,T} \equiv \mathbb{E}_t^{\mathbb{Q}}[g'(r_{t,T})] / F_{t,T}. \quad (5)$$

It can be computed at time t from the prices of options, and does not depend on a specific model.

A *delta-hedged g -swap* is the payoff to a dynamic trading strategy where at each time t when the portfolio is rebalanced, the portfolio is invested in the claim with payoff of $g(r_{t,T})$, a holding of $-\Delta_{t,T}$ forward contracts, and the rest of the portfolio is invested in risk-free bonds. Note that with $g(r_{t,T})$ changing over time, this strategy involves trading in the options market as well as in the forward market. Our main result is

Proposition 1. *If the price of the underlying asset and options on it follow a joint diffusion process, the continuously rebalanced delta-hedged g -swap is a perfect hedge for g .*

Proof. First consider a discretely rebalanced delta-hedged g -swap. If it is rebalanced successively at times t and $t + \delta t$, the change of the value of the position is

$$\delta W_{t,T} \equiv W_{t+\delta t,T} - W_{t,T} = \mathbb{E}_{t+\delta t}^{\mathbb{Q}}[g(r_{t,T})] - \mathbb{E}_t^{\mathbb{Q}}[g(r_{t,T})] - \Delta_{t,T}(F_{t+\delta t,T} - F_{t,T}) = \delta G_{t,T} + \delta Y_{t,T}, \quad (6)$$

where

$$\begin{aligned}\delta Y_{t,T} &= \mathbb{E}_{t+\delta t}^{\mathbb{Q}} [g(r_{t+\delta t,T} + \delta f_{t,T}) - g(r_{t+\delta t,T}) - (e^{\delta f_{t,T}} - 1) \mathbb{E}_t^{\mathbb{Q}} [g'(r_{t,T})]] , \text{ and} \\ \delta f_{t,T} &= f_{t+\delta t,T} - f_{t,T} = r_{t,T} - r_{t+\delta t,T}.\end{aligned}\tag{7}$$

Summing over the wealth gains (6), and recalling that $W_0 = G_{T,T} = 0$, we get that² $W_T = \sum \delta Y_{t,T} - G_{0,T}$. This justifies the description of the strategy as a swap, since it synthesizes a contract where the investor receives the floating leg $\sum \delta Y_{t,T}$ and pays the fixed leg $G_{0,T}$. Importantly, the fixed leg is independent of the rebalancing frequency.

As shown in Appendix B.1, if prices follow a diffusion and if the hedge is rebalanced continuously, the residual process $Y_{t,T}$ satisfies the stochastic differential equation

$$dY_{t,T} = dG'_{t,T} df_{t,T} + \frac{1}{2} (G''_{t,T} - G'_{t,T}) (df_{t,T})^2 ,\tag{8}$$

where $G'_{t,T} \equiv \mathbb{E}_t^{\mathbb{Q}} [g'(r_{t,T})]$ and $G''_{t,T} \equiv \mathbb{E}_t^{\mathbb{Q}} [g''(r_{t,T})]$. From (8) it can be seen that $Y_{t,T}$ is predictable, as there are no Brownian increments on the right-hand side, and the delta-hedged g -swap is therefore a perfect hedge for g . $\square$

With the continuously rebalanced g -swap being a perfect hedge for g , it is reasonable to regard the expected profit under the $\mathbb{P}$ measure from a g -swap as the risk premium associated with g . In practice it is only possible to rebalance discretely, and the discretely rebalanced swap only approximates a perfect hedge. In our empirical work we use the average profit from a g -swap that is rebalanced daily as an estimate of the risk premium associated with g .

2.1 A Variance Swap

To create a variance swap, consider the g -function $g^V(r) = -2r$. The fixed leg of the g^V -swap is, from the Bakshi and Madan (2000) formula

$$v_{t,T}^L \equiv G_{t,T}^V = -2\mathbb{E}_t^{\mathbb{Q}} \left[\ln \frac{F_{T,T}}{F_{t,T}} \right] = \frac{2}{B_{t,T}} \left\{ \int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2} dK \right\} ,\tag{9}$$

where $P_{t,T}(K)$ and $C_{t,T}(K)$ are prices of European put and call options at time t on the spot underlying with strike K and maturity T . G^V is increasing in the price of options of all strikes.

It can readily be shown that in a Black-Scholes world with constant volatility σ , the

²The portfolio adjustments take place at discrete times, not necessarily equidistant, between time t and T . To avoid excessive subscripting we write $\sum$ to denote the sum taken over all adjustment times.

fixed leg $G_{t,T}^V = \sigma^2(T - t)$. So $v_{t,T}^L$ is the implied Black-Scholes variance at time t of a *Log Contract* that matures at time T and pays $\ln(F_{T,T})$. It is the same as the model-free implied variance (MFIV) from [Britten-Jones and Neuberger \(2002\)](#) and corresponds to the square of the popular VIX implied variance contract scaled with time. It is therefore reasonable to follow established usage and regard g^V as a measure of variance.

Proposition 1 shows that the continuously rebalanced delta-hedged g^V -swap is a perfect hedge for variance (at least in a diffusion world). From (8) the floating leg of the variance swap can be seen to be

$$dY_{t,T}^V = (df_{t,T})^2. \quad (10)$$

This is identical to the definition of quadratic variation used in [Andersen et al. \(2003\)](#). Where the g^V swap differs slightly from the literature is in the discretely rebalanced case (or when there are jumps), where (7) shows that

$$\delta Y_{t,T}^V = 2 \left(e^{\delta f_{t,T}} - 1 - \delta f_{t,T} \right). \quad (11)$$

[Andersen et al. \(2003\)](#) however measure the realized volatility in discrete time as $\delta f_{t,T}^2$, and it is this definition that is used in practice to compute the floating leg of a variance swap. The formulation in (11) has previously been proposed by [Bondarenko \(2010\)](#). The difference between the two floating legs is locally cubic in the log return, which is why it vanishes in the diffusion world. The g^V -swap where the fixed leg is v^L and the floating leg is given by (11) can be replicated perfectly at zero cost for any data-generating process. This is not true for the standard variance swap.

2.2 A Skew Swap

Now consider $g^S(r) = 6r(1 + e^r)$. We show in Proposition 2 below that in this case the g^S -swap captures skewness. To see this note first that

$$s_{t,T} \equiv G_{t,T}^S = \mathbb{E}_t^{\mathbb{Q}} \left[6 \left(1 + \frac{F_{T,T}}{F_{t,T}} \right) \ln \frac{F_{T,T}}{F_{t,T}} \right]. \quad (12)$$

Using the [Bakshi and Madan \(2000\)](#) formula again, the fixed leg can be replicated from puts and calls

$$s_{t,T} = \frac{6}{B_{t,T}} \left\{ \int_{F_{t,T}}^{\infty} \frac{K - F_{t,T}}{K^2 F_{t,T}} C_{t,T}(K) dK - \int_0^{F_{t,T}} \frac{F_{t,T} - K}{K^2 F_{t,T}} P_{t,T}(K) dK \right\}. \quad (13)$$

The fixed leg s is the value of a position that is long OTM calls and short OTM puts. In a Black-Scholes world, its value is zero. The level of s therefore depends on the skew of implied volatility, and the g^S -swap in this case can be regarded as a hedge for skewness. From (12), s can be written as the difference between the implied Black-Scholes variances of two contracts: the Log Contract from Section 2.1, and the *Entropy Contract* that has payoff $F_{T,T} \ln(F_{T,T})$.

$$s_{t,T} = 3(v_{t,T}^E - v_{t,T}^L), \text{ where} \quad (14)$$

$$v_{t,T}^E \equiv 2\mathbb{E}_t^{\mathbb{Q}} \left[\frac{F_{T,T}}{F_{t,T}} \ln \frac{F_{T,T}}{F_{t,T}} \right].$$

In a diffusion world, the floating leg of the skew swap is given by (8), and equals

$$dY_{t,T}^S = 3dv_{t,T}^E df_{t,T}. \quad (15)$$

With discrete rebalancing or with jumps, the floating leg is given by (7) as

$$\begin{aligned} \delta Y_{t,T}^S &= \mathbb{E}_{t+\delta t}^{\mathbb{Q}} \left[g^S(r_{t+\delta t,T} + \delta f_{t,T}) - g^S(r_{t+\delta t,T}) - (e^{\delta f_{t,T}} - 1)G_{t,T}^{S'} \right] \\ &= 3\delta v_{t,T}^E (e^{\delta f_{t,T}} - 1) + h(\delta f_{t,T}), \text{ where} \\ h(x) &= 6(2 - 2e^x + x + xe^x) = x^3 + O(x^4). \end{aligned} \quad (16)$$

As we show below, s locally approximates the third moment of returns. To make it closer to the conventional coefficient of skewness we scale the g^S -swap entered into at time t by dividing both the fixed and the floating legs by $(v_{t,T}^L)^{3/2}$ which gives a fixed leg of

$$skew_{t,T} \equiv 3 \frac{v_{t,T}^E - v_{t,T}^L}{(v_{t,T}^L)^{3/2}} \quad (17)$$

We summarize the properties of the skew swap in the following

Proposition 2. *The scaled difference $skew_{t,T}$ between the entropy implied variance and the log implied variance has the following properties:*

1. *if implied volatility across strikes is symmetric about the forward price, $skew_{t,T} = 0$;*
2. *holding $v_{t,T}^L$ constant, $skew_{t,T}$ is increasing in the implied volatility of high strike (strike*

above the forward price) options and decreasing in the implied volatility of low strike options;

3. to first order, $\text{skew}_{t,T}$ is equal to the skew of the distribution of log returns on the asset (under the pricing measure);

4. with continuous trading in a diffusion world

$$s_{t,T} = \mathbb{E}_t^{\mathbb{Q}} \left[\int_{s \in [t,T]} 3dv_{s,T}^E df_{s,T} \right], \quad (18)$$

and measures the covariation between index returns and changes in index variance (the leverage effect).

Proof. In Appendix B.2. □

The skew swap, in a diffusion world, hedges changes in the implied skew perfectly. In our empirical work we use the average profit from a daily rebalanced scaled skew swap, expressed as a percentage of the initial implied skew, as an estimate of the *skew risk premium*.

2.3 Alternative Formulations

We now explain why we have chosen our measures of variance and skew over perhaps more obvious ones, and why we use raw rather than central moments.

The natural choice for g -functions for second and third moments might appear to be $g(r) = r^2$ and $g(r) = r^3$, respectively. As we will show, while it is indeed possible to design feasible variance and skew swaps using these functions, the mathematics are more complicated, and the corresponding definitions of implied variance and implied skew lack some of the appealing properties of our preferred definitions. In our empirical work we show that changing the definitions would have no significant impact on the measurement of variance and skew risk premia.

The fixed leg of a variance swap under this alternative formulation (a *quadratic swap*) is

$$\begin{aligned} v_{t,T}^Q &\equiv \mathbb{E}_t^{\mathbb{Q}} \left[\left(\ln \frac{F_{T,T}}{F_{t,T}} \right)^2 \right] \\ &= 2 \int_0^{F_{t,T}} \left(\frac{1 + \ln \frac{F_{t,T}}{K}}{B_{t,T} K^2} \right) P_{t,T}(K) dK + 2 \int_{F_{t,T}}^{\infty} \left(\frac{1 - \ln \frac{K}{F_{t,T}}}{B_{t,T} K^2} \right) C_{t,T}(K) dK. \end{aligned} \quad (19)$$

Similarly, the fixed leg of a *cubic swap* is given by

$$\begin{aligned}
s_{t,T}^C &\equiv \mathbb{E}_t^{\mathbb{Q}} \left[\left(\ln \frac{F_{T,T}}{F_{t,T}} \right)^3 \right] \\
&= 3 \int_{F_{t,T}}^{\infty} \frac{\ln \frac{K}{F_{t,T}} \left(2 - \ln \frac{K}{F_{t,T}} \right)}{B_{t,T} K^2} C_{t,T}(K) dK - 3 \int_0^{F_{t,T}} \frac{\ln \frac{F_{t,T}}{K} \left(2 + \ln \frac{F_{t,T}}{K} \right)}{B_{t,T} K^2} P_{t,T}(K) dK.
\end{aligned} \tag{20}$$

The expression for $v_{t,T}^Q$ is not only more complex than its counterpart equation (9), but also v^Q is actually declining in $C(K)$ for very high strikes. Similarly, in a cubic swap the replicating portfolio for the fixed leg is generally short OTM puts and long OTM calls, like the skew swap, but is actually short OTM calls for very high strikes.

In the diffusion case the floating legs are given by

$$dY_{t,T}^Q = -dv_{t,T}^L df_{t,T} + (1 + v_{t,T}^L/2) (df_{t,T})^2, \text{ and} \tag{21}$$

$$dY_{t,T}^C = 3dv_{t,T}^Q df_{t,T} - \frac{3}{2} \left(v_{t,T}^Q + v_{t,T}^L \right) (df_{t,T})^2. \tag{22}$$

These expressions are much more complicated than their counterparts, equations (10) and (15).

As we demonstrate below, the payoffs to the quadratic swap and to the variance swap are highly correlated, and similarly for the cubic and skew swaps. But in view of the substantially greater simplicity of the variance and skew swaps we prefer them to their quadratic and cubic counterparts.

The other aspect of our definitions that is worthy of comment is that we do not demean or center our statistics. The reason is that central moments cannot be perfectly hedged. Proposition 1 does not help us since it only applies to expectations of functions of the return on the asset; central moments cannot be expressed in this way. The implied central moment can of course be computed readily. The problem is with the realized leg. It is not possible to design a trading strategy whose payoff is $(r_{t,T} - \mathbb{E}_t^{\mathbb{P}}[r_{t,T}])^n$ unless $\mathbb{E}_t^{\mathbb{P}}[r_{t,T}]$ is known and this is generally not the case.

3 Data and Methodology

For the empirical analysis of the skew risk premium, we use European options written on the S&P 500 spot index traded on the CBOE. The options mature every month. The data

are from OptionMetrics, and cover the period January 1996 to January 2012. The data set includes closing bid and ask quotes for each option contract along with the corresponding strike prices, Black-Scholes implied volatilities, the zero-yield curve, and closing spot prices of the underlying. From the data we filter out all entries with non-standard settlements.

Suppose that at time t call and put options that mature at the end of month T are traded at $N + 1$ different strike prices K_i which are ordered from K_0 to K_N . Given the set of option prices $C_{t,T}(K_i)$ and $P_{t,T}(K_i)$ for different strike prices taken as the average of the bid and ask quotes we compute the implied variances of the log and entropy contracts, $v_{t,T}^L$ and $v_{t,T}^E$, and the implied third moment $s_{t,T}$ in the following way. Define the function

$$\Delta I(K_i) \equiv \begin{cases} \frac{K_{i+1} - K_{i-1}}{2}, & \text{for } 0 \leq i \leq N \text{ (with } K_{-1} \equiv 2K_0 - K_1, \text{ } K_{N+1} \equiv 2K_N - K_{N-1}) \\ 0, & \text{otherwise.} \end{cases}$$

Then according to formulae (9) and (14) we compute³

$$v_{t,T}^L = 2 \sum_{K_i \leq F_{t,T}} \frac{P_{t,T}(K_i)}{B_{t,T} K_i^2} \Delta I(K_i) + 2 \sum_{K_i > F_{t,T}} \frac{C_{t,T}(K_i)}{B_{t,T} K_i^2} \Delta I(K_i) \quad (23)$$

$$v_{t,T}^E = 2 \sum_{K_i \leq F_{t,T}} \frac{P_{t,T}(K_i)}{B_{t,T} K_i F_{t,T}} \Delta I(K_i) + 2 \sum_{K_i > F_{t,T}} \frac{C_{t,T}(K_i)}{B_{t,T} K_i F_{t,T}} \Delta I(K_i). \quad (24)$$

The implied third moment is $s_{t,T} = 3(v_{t,T}^E - v_{t,T}^L)$ and the implied skew $skew_{t,T}$ is computed using (17). We compute the fixed legs of the quadratic and cubic swaps in a similar fashion.

The floating legs for the variance and skew swaps (realized second and third moments) assume daily rebalancing⁴ and are computed using (11) and (16)

$$rv_{t,T} = \sum_{i=t}^T [2(e^{r_{i,i+1}} - 1 - r_{i,i+1})] \quad (25)$$

$$rs_{t,T} = \sum_{i=t}^T [\delta v_{i,T}^E (e^{r_{i,i+1}} - 1) + 6(2 - 2e^{r_{i,i+1}} + r_{i,i+1} + r_{i,i+1} e^{r_{i,i+1}})] . \quad (26)$$

We define the realized skew as

$$rskew_{t,T} \equiv rs_{t,T} / (v_{t,T}^L)^{3/2}. \quad (27)$$

³The way we compute implied variances is standard (e.g., see [Bollerslev et al. \(2009\)](#)). Some papers however extrapolate volatilities outside the minimum and maximum strike prices (e.g. [Carr and Wu \(2009\)](#)). We perform extrapolation-interpolation procedure for robustness check and present the results in [Appendix C](#).

⁴We consider strategies with less frequent adjustments to decrease transaction costs in [Appendix C](#).

Figure 1: Excess returns from trading skew and variance swaps

The figure shows the time series of the realized excess returns on monthly skew and variance swaps, xs_t and xv_t . The data ranges from January 1996 to January 2012.

Finally, we define the excess return (profit per \$1 investment in the fixed leg) as

$$xs_{t,T} \equiv rskew_t/skew_t - 1 \quad (28)$$

$$xv_{t,T} \equiv rv_{t,T}/v_{t,T}^L - 1. \quad (29)$$

Our empirical analysis concentrates on trading strategies that run for a month, from the first trading day after one option expires to the next month's expiration date. To simplify notation we use a single subscript t to relate to the trading month. So $skew_t$ is the implied skew at the beginning of month t and $rskew_t$ is the realized skew over month t .

3.1 Descriptive Statistics

Figure 1 plots the time series for variance and skew swap excess returns. We can see that there is substantial common variation between the two. Panel A of Table 1 contains summary statistics of implied variance, realized variance, implied and realized skew coefficients of the S&P 500 index options and variance and skew premia. The average realized variance is

Table 1: Descriptive statistics

This table presents descriptive statistics for the variables used in the analysis. v^L denotes monthly implied variance of the Log contract on the S&P 500 index options, rv is realized variance and xv denotes the excess return from the variance swap defined as $xv_t \equiv (rv_t/v_t^L - 1)$, $skew$ is the implied skew defined as the fixed leg of the skew contract scaled by $(v^L)^{3/2}$, $rskew$ is the realized skew defined as the floating leg of the skew contracts scaled by $(v^L)^{3/2}$ and xs is the excess return from the skew swap defined as $xs_t \equiv (rskew_t/skew_t - 1)$. In Panel A, columns under Mean, Std. Dev., Min, Median and Max report the sample average, standard deviation, minimum value, median and maximum value, respectively. Panel B presents the correlations among the variables. Bold values are significant at the 1% level. The data are for non-overlapping monthly periods from January 1996 to January 2012.

<i>Panel A: Descriptive Statistics</i>					
Variable	Mean	Std. Dev	Q1	Median	Q3
$v^L \times 100$	0.475	0.508	0.223	0.358	0.523
$rv \times 100$	0.382	0.612	0.113	0.219	0.415
xv	-22.25%	66.06%	-56.29%	-36.83%	-18.53%
$skew$	-1.808	0.722	-2.225	-1.748	-1.243
$rskew$	-1.001	2.241	-0.853	-0.532	-0.323
xs	-42.09%	117.55%	-83.30%	-68.20%	-44.79%
<i>Panel B: Correlations</i>					
	rv	xv	$skew$	$rskew$	xs
v^L	0.731	0.037	0.283	0.025	0.039
rv	1.000	0.574	0.235	-0.345	0.497
xv		1.000	0.048	-0.843	0.897
$skew$			1.000	0.128	0.054
$rskew$				1.000	-0.922

substantially lower than the implied variance. Also, average realized skew is substantially smaller (in absolute terms) than average implied skew, pointing to the existence of a skew risk premium. Implied skew is negative throughout the sample period and realized skew is on average negative. So the writer of a skew swap who receives fixed and pays floating generally receives and pays negative amounts, and on average loses money. Almost half (42.09%) of the implied skew in index option prices can be explained by the risk premium, with the other half reflecting the negative correlation between returns and volatility. The variance risk premium is smaller, accounting for less than a quarter (22.25%) of the implied variance. The sample averages of the excess return of skew and variance swap strategies (estimates of skew and variance risk premia) are negative.

The negative variance risk premium is consistent with the literature (see, for example, [Bakshi and Kapadia 2003](#), [Carr and Wu 2009](#), [Egloff et al. 2010](#)). Writing variance swaps, receiving fixed and paying floating, is on average profitable. The negative return on skew swaps is consistent with [Bakshi et al. \(2003\)](#), who show theoretically that, within a power-

Table 2: Correlations: variance and skew swaps vs. quadratic and cubic swaps

This table presents correlations between the implied moments and premia computed using different methodologies. Panel A contains correlations between the following variables: v^L and v^E (monthly implied variances of the Log and Entropy contracts, respectively), $skew$ (the implied skew defined as the fixed leg of the skew swap scaled by $(v^L)^{3/2}$), v^Q (fixed leg of the quadratic swap), s^C (fixed leg of the cubic swap), v^{cent} and $skew^{cent}$ (implied second and third central moments defined as in Bakshi et al. (2003)). Panel B presents pairwise correlations between excess returns from variance and skew swaps and excess returns from quadratic and cubic swaps defined as in equation (5) of Bakshi et al. (2003) respectively. Bold values are statistically significant at the 1% level. The data are for non-overlapping monthly periods from January 1996 to January 2012.

<i>Panel A: Correlations of Implied Moments</i>						
	v^E	v^Q	v^{cent}	$skew$	$skew^C$	$skew^{cent}$
v^L	0.999	0.999	0.999	0.283	0.240	0.272
v^E	1.000	0.998	0.998	0.299	0.255	0.287
v^Q		1.000	0.999	0.264	0.221	0.253
v^{cent}			1.000	0.265	0.222	0.254
$skew$				1.000	0.984	0.985
$skew^{cent}$					1.000	0.999
<i>Panel B: Correlations of Risk Premia</i>						
	xv	xv^Q	xs	xs^C		
xv	1.000	0.999	0.897	0.874		
xv^Q		1.000	0.909	0.808		
xs			1.000	0.991		

utility economy where returns are leptokurtic, the implied skew is greater in magnitude than the physical skew.

Panel B of Table 1 reports pair-wise correlations between variables. Implied variance is negatively correlated with skew, so when implied variance goes up, the absolute size of the skew also increases. Realized variance is positively correlated with implied variance, as is realized skew with implied skew. Realized variance and realized skew are negatively correlated with each other. Variance and skew swaps excess returns are highly positively correlated.

In the rest of the paper, we report excess returns computed from variance and skew swaps. Table 2 shows that the use of alternative measures of variance and skew based on quadratic and skew swaps, described in Section 2, would make little practical difference. Panel A records that the different measures of implied variance have pairwise correlations in excess of 0.998, while the correlations for implied skew are in excess of 0.98. Panel B shows similar correlations between the corresponding excess returns estimates.

4 Empirical Results

We use the technology developed in this paper to deepen our understanding of the relation between implied and realized skew in the equity index market. In the absence of risk premia, one would expect variations in implied skew to predict variations in realized skew, and we examine whether this is the case. We go on to see whether the excess returns on skew swaps can be explained by a simple one-factor model. [Bakshi et al. \(2003\)](#) have shown theoretically that the size of the skew risk premium should also depend on the higher moments of the conditional distribution of log returns, and in particular on the excess kurtosis. So we use predictive models of the skew risk premium using proxies for the moments as well as for macro-economic factors.

Given the high correlation we have observed between skew and variance swap excess returns, we analyze variance swap returns in parallel with the skew swap returns, and test whether they are driven by the same combination of priced risk factors.

4.1 Predictive Power of Implied Skew

Implied skew, with a mean of -1.81 and a standard deviation of 0.72, varies significantly over time. We run a standard expectations hypothesis regression to test whether the variations in implied skew reflect changes in the expected realized skew,

$$rskew_t = \alpha_0 + \alpha_1 skew_t + u_t. \quad (30)$$

To provide a comparison, we also estimate the corresponding variance regression

$$rv_t = \alpha_0 + \alpha_1 v_t^L + u_t. \quad (31)$$

These are predictive regressions; the implied quantities are known at the beginning of month t , while the realized quantities are only known at the end of the month.

Table 3 shows that implied skew does forecast realized skew and implied variance does forecast realized variance. The slopes are both highly significant, and the two intercepts are insignificantly different from zero.

Given the existence of statistically significant skew and variance risk premia, it is not surprising that we strongly reject the expectations hypothesis that $\alpha_0 = 0$ and $\alpha_1 = 1$ for both regressions. We also reject the simple hypothesis that $\alpha_1 = 1$ in both regressions, suggesting that there is a significant variation in the risk premia.

Table 3: Realized skew vs. implied skew

This table presents the estimation results of regressions $rskew_t = \alpha_0 + \alpha_1 skew_t + u_t$, and $rv_t = \alpha_0 + \alpha_1 v_t^L + w_t$ based on one month S&P 500 index options. Standard t-statistics are given in parentheses. “F ($\alpha_0 = 0, \alpha_1 = 1$)” column provides values of the F-statistic and the corresponding p-values for a Wald test with $H_0: \alpha_0 = 0$ and $\alpha_1 = 1$. The “F ($\alpha_1 = 1$)” column provides values of the F-statistic and the corresponding p-values for a Wald test with $H_0: \alpha_1 = 1$. Standard errors use the [Newey and West \(1987\)](#) correction. The sample period is from January 1996 to January 2012.

	Const	<i>skew</i>	v^L	$\bar{R}^2$	F ($\alpha_0 = 0, \alpha_1 = 1$)	F ($\alpha_1 = 1$)
<i>rskew</i>	-0.282 (-0.75)	0.398 (2.59)		1.11%	12.07 (0.000)	6.46 (0.012)
<i>rv</i>	-0.001 (-1.52)		0.880 (10.1)	53.28%	5.27 (0.006)	3.99 (0.047)

4.2 The Relation between Skew and Variance Risk

To better understand the source of the skew and variance risk premia we next ask whether they can be explained by the Capital Asset Pricing Model. To answer this we estimate

$$xs_t = \alpha_0 + \alpha_1 xm_t + u_t, \quad (32)$$

$$xv_t = \beta_0 + \beta_1 xm_t + w_t, \quad (33)$$

where xm_t denotes the excess market return over month t . The estimation results are provided in Table 4. In both models the market beta is negative and highly significant, suggesting that buyers of variance and skew swaps are exposed to the market. However, the intercepts in both equations remain statistically significant, and they do not differ greatly from the unconditional average excess returns from both swap strategies. Our findings, so far as the variance risk premium is concerned, are consistent with [Carr and Wu \(2009, Table 4\)](#). Thus, we conclude that the simple market model cannot fully explain variance and skew risk premia.

We now seek to establish whether the skew risk premium reflects an independent risk factor from the variance risk premium, or whether the profitability of buying skew swaps can be explained by exposure to both equity and variance risk. To do this, we regress the skew swap excess returns on the variance swap and market excess returns

$$xs_t = \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + u_t. \quad (34)$$

We also estimate a similar regression for the variance risk premium to verify if it can be

Table 4: Skew and variance excess returns vs. the market return

This table presents OLS estimation results of the skew and variance risk premia regressions: $xs_t = \alpha_0 + \alpha_1 xm_t + u_t$ and $xv_t = \beta_0 + \beta_1 xm_t + w_t$. Variables xs and xv denote excess returns from skew and variance swaps, xm denotes the market excess return. [Newey and West \(1987\)](#) t-statistics are given in parentheses. The data are monthly, from January 1996 to January 2012.

	Const	xm	$\bar{R}^2$
xv	-0.193 (-4.66)	-8.531 (-4.69)	34.68%
xs	-0.382 (-4.46)	-11.217 (-3.43)	18.69%

explained by the skew risk premium⁵

$$xv_t = \beta_0 + \beta_1 xs_t + \beta_2 xm_t + w_t. \quad (35)$$

The high correlation between skew and variance swap returns makes it hard to find evidence of separate risk premia through contemporaneous regressions. To improve efficiency we estimate the two equations as a system using the Seemingly Unrelated Regression (SUR) approach.⁶ The estimation results are in Panel A of Table 5. They highlight the close relationship between variance and skew swaps. Table 1 shows that the buyer of a skew swap on average earns a risk premium equal to 42.09% of the implied skew. But in doing so, the buyer takes on significant variance risk. Once that is hedged by buying variance swaps, and market risk is also hedged, most of the premium is lost, and the buyer of the skew swap earns a return of 0.216% that is both economically small and statistically insignificant. Similarly the writer of a variance swap is exposed to skew risk (the connection of the variance risk premium and the well-known leverage effect, has been highlighted in the literature ([Carr and Wu 2009](#), [Bollerslev et al. 2009](#), [Driessen and Maenhout 2007](#), [Jones 2003](#))), and once that is hedged away the profit from writing variance swaps also vanishes.

The other striking feature about Table 5 is the very high R^2 for all the regressions. Around 80% (84%) of the risk of a skew (variance) swap can be hedged using variance (skew) swaps, and index forwards.

We do not observe the expected profit (risk premia) from the skew and variance swaps but only the actual realization. This creates an errors-in-variables problem. Moreover, the

⁵[Chabi-Yo \(2009\)](#) finds that Fama-French factors also explain variation in index options returns. We have estimated the above specification but including Fama-French and momentum factors and the results remain unchanged. We report them in Appendix C.

⁶The results of OLS estimations of each equation separately are qualitatively similar.

Table 5: Skew vs. variance risk premia

This table presents results for the skew and variance risk premia regressions: $xs_t = \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + u_t$ and $xv_t = \beta_0 + \beta_1 xs_t + \beta_2 xm_t + w_t$. Variables xs and xv denote excess returns from skew and variance swaps, xm denotes the market excess return. Panel A presents Seemingly Unrelated Regression estimation results. Panel B shows results using three-stage least squares, where we use the following predetermined instrumental variables: ted_t , the TED spread defined as the difference between 3 month LIBOR and the 3 month U.S. T-bill rate at the beginning of month t ; ks_t , the implied normalized excess kurtosis as defined in equation (36); and s_t , the implied third moment. Newey and West (1987) t-statistics are given in parentheses. The data are monthly, from January 1996 to January 2012.

Panel A: Seemingly Unrelated Regressions

Dep.Var.	Const	xv	xs	xm	$\bar{R}^2$
xv	-0.002 (-0.11)		0.500 (38.4)	-2.927 (-6.88)	84.41%
xs	-0.002 (-0.06)	1.969 (38.4)		5.584 (6.16)	80.59%

Panel B: Three-Stage LS Regressions

<i>First Stage Regressions</i>						
	Const	ks_t	ted_t	$s_t \times 100$	$\bar{R}^2$	F
xv	-0.366 (-2.47)	0.144 (0.70)	0.500 (1.99)	1.197 (1.70)	12.62%	10.10
xs	-0.474 (-1.89)	0.712 (1.80)	0.770 (2.13)	1.659 (1.58)	10.71%	8.56
<i>Third Stage Regressions</i>						
	Const	xv	xs	xm	$\bar{R}^2$	
xv	0.003 (0.17)		0.514 (25.0)	-2.768 (-5.84)	83.94%	
xs	-0.007 (-0.16)	1.946 (25.1)		5.386 (5.16)	80.80%	

estimators based on the SUR approach might suffer from simultaneity bias. To deal with this, we use instrumental variables and estimate the system of equations by three-stage least squares.

As candidates for instrumental variables we use predetermined variables which are theoretically motivated and can predict excess returns of the two swap strategies. Bakshi et al. (2003) and Bakshi and Madan (2006) show that skew and variance risk premia can be explained by higher moments of the distribution of log returns. So we include as candidates implied variance, implied skew and implied kurtosis, realized variance and skew, and variance of variance (defined as monthly variance of daily levels of VIX index). Theorem 2 of Bakshi et al. (2003) has an explicit expression for the skew risk premium in terms of the kurtosis of returns, so we include as a candidate instrumental variable

$$ks_t = \frac{(kurt_t^{cent} - 3) \sqrt{v_t^{cent}}}{skew_t^{cent}}, \quad (36)$$

where $kurt_t^{cent}$ is the implied excess kurtosis defined in Equation (6) of [Bakshi et al. \(2003\)](#). We use implied moments rather than physical moments, because the conditional physical moments are not directly observable, and we divide by skew since the dependent variable is the swap return rather than the difference between implied and realized skew.

Other candidates we consider as instrumental variables include lagged macroeconomic variables (the levels and changes in price-earnings, price-dividend and consumption-wealth ratios, TED, term and default spreads, and the excess market return), and market variables (put-call trading volume ratio, and option open interest).

There are only three variables which have power to predict excess returns (are significant at the 10% level in the first stage regressions): ted_t , the spread between 3 month LIBOR and the 3 month U.S. T-bill rate at the beginning of month t ; ks_t , defined in equation (36); and s_t , the beginning of the month implied third moment. We drop the other variables to avoid multicollinearity issues and the loss of efficiency of our estimators.

Estimation results are given in Panel B of Table 5. In the first stage, we regress the risk premia on the chosen instruments. The F-statistics from the first-stage regressions are significant at the 1% level, firmly rejecting the hypothesis that our chosen instrumental variables are weak. The third-stage estimation results are qualitatively the same as in the SUR estimation confirming our earlier conclusions.

In Appendix C we demonstrate that the empirical results are robust to the choice of time period, the method of computing the implied moments, the omission of illiquid options and the presence of transaction costs.

4.3 Common Time Variation in Two Premia Series

In the previous section we have shown that unconditionally one cannot make significant profits by buying a skew swap and simultaneously hedging against variance and market risk. We can go further and test whether the same result holds conditionally. Given that there is significant time variation in premia, we can ask whether the time variation is the same for the two types of risk.

Following [Cochrane and Piazzesi \(2005\)](#) we test the specification

$$xs_t = \lambda_0 (\gamma_0 + \gamma_1 z_{t-1}^1 + \dots + \gamma_n z_{t-1}^n) + u_t, \quad (37)$$

$$xv_t = \lambda_1 (\gamma_0 + \gamma_1 z_{t-1}^1 + \dots + \gamma_n z_{t-1}^n) + w_t, \quad (38)$$

for predetermined predictors z_t^i , $i = 1, \dots, n$. Here, we use the same predictive variables as in

Table 6: Skew and Variance Risk Premia

This table presents the estimation results of the following two predictive regressions (columns xs -unrestr and xv -unrestr)

$$\begin{aligned} xs_t &= \alpha_0 + \alpha_1 ks_t + \alpha_2 s_t + \alpha_3 ted_t + u_t, \\ xv_t &= \beta_0 + \beta_1 ks_t + \beta_2 s_t + \beta_3 ted_t + w_t. \end{aligned}$$

xs and xv denote excess returns from skew and variance swap strategies, ks is the variable defined in equation (36), s denotes the implied third moment and ted is the TED spread (the difference between 3 month LIBOR and the 3 month U.S. T-bill rate) at the beginning of month t . Columns xs -restr and xv -restr present the corresponding coefficients from the following system of restricted regressions

$$\begin{aligned} xs_t &= \lambda_0 (\gamma_0 + \gamma_1 ks_t + \gamma_2 s_t + \gamma_3 ted_t) + u_t, \\ xv_t &= \lambda_1 (\gamma_0 + \gamma_1 ks_t + \gamma_2 s_t + \gamma_3 ted_t) + w_t, \end{aligned}$$

where the coefficients $\gamma_0, \dots, \gamma_3$ are estimated from the model $(xs_t + xv_t)/2 = \gamma_0 + \gamma_1 ks_t + \gamma_2 s_t + \gamma_3 ted_t + \varepsilon_t$. t -stat column presents the values of t -statistics for testing the individual coefficients in the unrestricted model being equal to the value of the corresponding coefficients in the restricted model. F -stat row reports the values of the Wald test statistics that the above restrictions are not binding.

	xs -unrestr	xs -restr	t -stat	xv -unrestr	xv -restr	t -stat
Const	-0.474 (-1.89)	-0.544	0.28	-0.366 (-2.47)	-0.296	-0.47
ks_t	0.712 (1.79)	0.534	0.45	0.144 (0.69)	0.321	-0.86
$s_t \times 100$	1.659 (1.59)	1.783	-0.12	1.197 (1.70)	1.073	0.17
ted_t	0.770 (2.14)	0.793	-0.06	0.500 (1.99)	0.477	0.18
R^2	12.13%	12.07%		14.01%	13.83%	
F -stat		0.03			0.10	

Table 5. As noted by [Cochrane and Piazzesi \(2005\)](#), the coefficients λ_i and γ_i are not separately identified by (37) and (38), and we impose the additional constraint that $(\lambda_0 + \lambda_1)/2 = 1$.

We estimate the model in two steps. First, we estimate the γ 's by running a regression of the average of the two risk premia on all the z 's,

$$(xs_t + xv_t)/2 = \gamma_0 + \gamma_1 z_{t-1}^1 + \dots + \gamma_n z_{t-1}^n + \varepsilon_t. \quad (39)$$

Then we estimate the λ 's by running the two regressions

$$xs_t = \lambda_0 z_{t-1} + u_t \quad (40)$$

$$xv_t = \lambda_1 z_{t-1} + w_t, \quad (41)$$

where $z_t = \hat{\gamma}_0 + \hat{\gamma}_1 z_t^1 + \dots + \hat{\gamma}_n z_t^n$. The single-factor model is a restricted model and we can test for these restrictions using a Wald test.

Consistent with the results in Table 3, Table 6 confirms that there is significant predictability in both risk premia. Both risk premia decline in absolute value after the stock market

Table 7: Correlations: dynamic vs. static strategies

This table presents pair-wise correlations among excess returns from dynamically hedged variance (xv) and skew swaps (xs) and excess returns from delta-hedged buy-and-hold variance (xv^{stat}) and skew (xs^{stat}) swaps. Bold values are statistically significant at the 1% level. The data are for non-overlapping monthly periods from January 1996 to January 2012.

	xv	xv^{stat}	xs	xs^{stat}
xv	1.000	0.643	0.897	0.003
xv^{stat}		1.000	0.487	0.311
xs			1.000	0.032

declines, when the implied distribution is less negatively skewed, and when the TED spread is high. The Wald test results show that one cannot reject the null that both risk premia are driven by a single factor. Moreover, none of the individual coefficients in the unrestricted model is statistically different from its counterpart in the restricted model. So, both the conditional and unconditional evidence is consistent with the hypothesis that skew and variance risk premia are driven by a common factor.

4.4 Static Strategies

We have argued for the importance of rebalancing to ensure that the risk characteristics of the strategy remain constant over time. In this section we compare the performance of the daily rebalanced skew swap with its static counterpart. At the beginning of the month, the writer of the skew swap buys OTM calls and writes OTM puts, and goes short forward contracts so as to have a delta-neutral position. The position is a pure bet on the slope of the implied volatility surface. Suppose now that the index rises substantially. With no rebalancing, the OTM puts would become deep out-of-the-money, and the portfolio would be dominated by the long call position, so the position would be long volatility. The delta would also become positive. To maintain a pure exposure to skew, the swap writer sells options at all strikes and also sells forward contracts.

Under the static strategy the portfolio is not rebalanced, and the skew swap picks up exposure to implied volatility (positive or negative exposure depending on whether the index is above or below its initial level) and also exposure to the underlying (with the exposure always being positive and increasing as the price moves further away from its initial level). Similarly, the static variance swap picks up delta exposure, with the direction of the exposure depending on whether the price is above or below the initial level. So both swaps pick up risks

Table 8: Skew vs. variance risk premia: static strategies

This table presents estimation results of the following skew and variance risk premia regressions: $xs_t^{stat} = \alpha_0 + \alpha_1 xv_t^{stat} + \alpha_2 xm_t + u_t$ and $xv_t^{stat} = \beta_0 + \beta_1 xs_t^{stat} + \beta_2 xm_t + w_t$. Variables xs_t^{stat} and xv_t^{stat} denote excess returns from static buy-and-hold skew and variance swap strategies, xm denotes the market excess return. Panel A presents Seemingly Unrelated Regression estimation results, Panel B contains estimation results using three-stage least squares. For the latter method we use the following predetermined instrumental variables: ted_t , the TED spread defined as the difference between 3 month LIBOR and the 3 month U.S. T-bill rate at the beginning of month t ; ks_t , the implied normalized excess kurtosis as defined in equation (36), and s_t , the beginning of the month third moment. T-statistics are given in parentheses. Standard errors use the Newey and West (1987) correction. The data are monthly, from January 1996 to January 2012.

Panel A: Seemingly Unrelated Regressions

Dep.Var.	Const	xv^{stat}	xs^{stat}	xm	$\bar{R}^2$
xv^{stat}	0.139 (2.47)		0.746 (17.3)	-14.19 (-13.6)	38.86%
xs^{stat}	-0.331 (-5.65)	1.023 (17.3)		16.59 (13.6)	38.36%

Panel B: Three-Stage LS Regressions

<i>First Stage Regressions</i>						
	Const	ks_t	ted_t	$s_t \times 100$	$\bar{R}^2$	F
xv^{stat}	-0.591 (-3.35)	0.105 (0.36)	0.479 (1.65)	1.806 (2.72)	7.29%	5.96
xs^{stat}	-0.749 (-3.86)	0.469 (0.96)	0.384 (1.37)	0.874 (0.90)	3.26%	3.12
<i>Third Stage Regressions</i>						
	Const	xv^{stat}	xs^{stat}	xm	$\bar{R}^2$	
xv^{stat}	-0.010 (0.17)		0.560 (13.4)	-12.56 (-12.2)	45.84%	
xs^{stat}	0.017 (0.16)	1.784 (12.6)		22.41 (11.1)	-26.76%	

other than skew and variance, and this masks the strong relationship we have seen between variance and skew risk. The correlation between excess returns on the variance and skew swaps drops from 0.897 (Table 1) when rebalanced daily to 0.311 (Table 7) when not rebalanced at all.⁷

Table 8 is similar to Table 5 except that it uses excess returns from static rather than dynamic strategies. It shows that if the excess returns of the skew and variance swaps are regressed on each other and the market, the R^2 's drop from over 80% with daily rebalancing to under 40%. More importantly, the intercepts now become both highly significant, and large in absolute terms. The estimated skew risk premium, after hedging variance and market risk, is -33%, while the hedged variance risk premium is +14%. The static results, taken in

⁷We define the realized leg of the static buy-and-hold variance and skew swaps according to formulae (11) and (16), but using monthly forward return and monthly change in the entropy variance rather than daily rebalancing.

isolation, could be interpreted as evidence that the skew and variance risk premia are rewards for taking distinct risks. But the fact that the intercepts from the daily rebalanced strategies are both small and insignificant shows that these intercepts are driven by the random risk exposures generated by not rebalancing rather than by skew and variance.

Panel B of Table 8 supports this conclusion. When we use instrumental variables with the same set of instruments as before, the intercepts in both equations become insignificant. This confirms that the significant intercepts in the SUR model were attributable to an errors-in-variables problem. Even though the instrumental variable method deals well with the endogeneity problem, there is a substantial decline in the first-stage F-statistics, showing that the instruments are weaker in predicting static rather than dynamic excess returns, and this leads to higher standard errors.

5 Conclusions

We have developed a trading strategy for investigating the skew risk premium in the equity index market. The strategy involves buying low-strike puts and writing high-strike calls, and subsequently trading options and forwards on the underlying. The strategy has a number of advantages over a simple buy-and-hold strategy. It has a risk profile that is stable over time. It has a simple interpretation that is independent of any assumptions about asset price dynamics; it replicates a skew swap which pays the difference between the implied skew in option prices and the realized skew as determined by the covariance between returns and variance shocks.

We apply the technology to options on the US S&P 500 index, which exhibit a pronounced and consistent skew. We show that the skew is only partly accounted for by the negative correlation between returns and volatility (the leverage effect). We find that there is a substantial skew premium. The finding is robust to the presence of transactions costs.

The absolute size of the skew is highly correlated with the level of volatility and the profit from buying skew swaps is highly correlated with the profit from writing variance swaps. When implied variance rises unexpectedly, the writer of the variance swap loses money (on a mark-to-market basis); but the rise in implied variance tends to be associated with an increase in the covariation between variance and returns, thereby causing the buyer of the skew swap to lose money. This suggests that variance swaps and skew swaps create similar kinds of risk exposure. A strategy of buying skew swaps and hedging out variance risk by buying variance swaps generates no significant excess return.

Further evidence for the close connection between variance and skew risk comes in the time series. Both, skew and variance risk premia exhibit time variation. We present evidence that the same combination of priced risk factors drives both risk premia.

These results depend heavily on the use of dynamic hedging to maintain a constant risk exposure. Static strategies to exploit the skew and variance risk premia are exposed to stochastic risk exposures that give the spurious appearance that the two premia are rewarding different risks. The problem of identifying priced risk factors when the trading strategy's loading on the factors is time-varying is not unique to option trading. It has been much discussed in hedge fund performance analysis ([Bollen and Whaley \(2009\)](#)). In option trading however it is somewhat easier to allow for since we can design dynamic strategies to have constant risk exposures.

We conclude that there is powerful evidence for the existence both of variance and skew risk premia in the equity index market. The magnitudes are sufficiently large that one would expect them to be reflected in an asset pricing model. They do not appear to be explained by the equity risk premium, but the evidence presented here suggests that a single channel is sufficient to explain both premia. There is no need to provide separate channels for volatility and skew risk premia.

References

- Andersen, T., Bollerslev, T., Diebold, F. and Labys, P.: 2003, Modeling and forecasting realized volatility, *Econometrica* **71**(2), 579–625.
- Anderson, E. W., Hansen, L. P. and Sargent, T. J.: 2000, Robustness, detection and the price of risk. Working paper, UNC, Stanford, and Chicago.
- Bakshi, G. and Kapadia, N.: 2003, Delta-hedged gains and the negative market volatility premium, *Review of Financial Studies* **16**, 527–566.
- Bakshi, G., Kapadia, N. and Madan, D.: 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* **16**, 101–143.
- Bakshi, G. and Madan, D.: 2000, Spanning and derivative-security valuation, *Journal of Financial Economics* **55**, 205–238.
- Bakshi, G. and Madan, D.: 2006, A theory of volatility spreads, *Management Science* **52**, 1945–1956.

- Bakshi, G., Madan, D. and Panayotov, G.: 2010, Returns of claims on the upside and the viability of u-shaped pricing kernels, *Journal of Financial Economics* **97**, 130–154.
- Bollen, N. P. B. and Whaley, R. E.: 2009, Hedge fund risk dynamics: Implications for performance appraisal, *Journal of Finance* **64**, 985–1035.
- Bollen, N. and Whaley, R.: 2004, Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance* **59**, 711–753.
- Bollerslev, T., Tauchen, G. and Zhou, H.: 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* **22**, 4463–4492.
- Bondarenko, O.: 2003, Why are put options so expensive?, *Technical report*, University of Illinois at Chicago.
- Bondarenko, O.: 2010, Variance trading and market price of variance risk. Working Paper, University of Illinois.
- Britten-Jones, M. and Neuberger, A.: 2002, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* **55**, 839–866.
- Broadie, M., Chernov, M. and Johannes, M.: 2009, Understanding index options returns, *Review of Financial Studies* **22**, 4493–4529.
- Carr, P., Ellis, K. and Gupta, V.: 1998, Static hedging of exotic options, *Journal of Finance* **53**, 1165–1190.
- Carr, P. and Madan, D.: 2001, Towards a theory of volatility trading, in E. Jouini, J. Cvitanic and M. Musiela (eds), *Option Pricing, Interest Rates and Risk Management*, Handbooks in Mathematical Finance, Cambridge University Press, pp. 458–476.
- Carr, P. and Wu, L.: 2009, Variance risk premiums, *Review of Financial Studies* **22**, 1311–1341.
- Chabi-Yo, F.: 2009, Expected returns and volatility of fama-french factors. Working paper, Fisher College of Business.
- Cochrane, J. and Piazzesi, M.: 2005, Bond risk premia, *American Economic Review* **94**(1), 138–160.
- Coval, J. and Shumway, T.: 2001, Expected option returns, *Journal of Finance* **56**, 983–1009.

- Driessen, J. and Maenhout, P.: 2007, An empirical portfolio perspective on option pricing anomalies, *Review of Finance* **11**, 561–603.
- Egloff, D., Leippold, M. and Wu, L.: 2010, The term structure of variance swap rates and optimal variance swap investments, *Journal of Financial and Quantitative Analysis* **45**, 1279–1310.
- Ferson, W. E. and Harvey, C. R.: 1999, Conditioning variables and the cross section of stock returns, *Journal of Finance* **54**, 1325–1360.
- Ferson, W. E., Sarkissian, S. and Simin, T.: 1999, The alpha factor asset pricing model: A parable, *Journal of Financial Markets* **2**, 49–68.
- Jiang, G. J. and Tian, Y. S.: 2005, The model-free implied volatility and its information content, *Review of Financial Studies* **18**, 1305–1342.
- Jones, C. S.: 2003, The dynamics of stochastic volatility: evidence from underlying and options markets, *Journal of Econometrics* **116**, 181–224.
- Neuberger, A.: 2012, Realized skewness, *Review of Financial Studies* . forthcoming.
- Newey, W. and West, K.: 1987, A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**, 703–708.

APPENDIX

A Implied Central vs. Observed Population Moments

The conditional distribution of returns under $\mathbb{P}$ is not directly observable. In practice, it is common to make inferences about the $\mathbb{P}$ distribution by looking at returns over many periods. The question then arises about the relation between the $\mathbb{P}$ moments estimated from a long series of returns, and the average conditional $\mathbb{Q}$ moments computed over the same period from option prices.

To analyze this, assume a world with no risk premia. In that case

$$\mathbb{E}^{\mathbb{P}} [r_{t,T}^n] = \mathbb{E}^{\mathbb{P}} [\mathbb{E}_t^{\mathbb{Q}} [r_{t,T}^n]], \quad (42)$$

so that mean option-implied moments converge to the (unconditional) population moment. The conclusion does not apply to central moments, however, since equation (42) does not hold due to the centering conditional $\mathbb{P}$ expectation. Indeed it is easy to think of examples where there are no risk premia and population moments differ from average implied moments. Take for instance the discrete-time model

$$r_{t,T} = -\frac{1}{2}v_t + \sqrt{v_t}\varepsilon_{t,T}, \quad (43)$$

where v_t is known at time t and ε_T is standard normal. In the absence of risk premia, the implied variance in period $[t, T]$ is v_t , so the mean implied variance is $\bar{v}$ (the unconditional mean variance). The population variance is

$$\mathbb{E}^{\mathbb{P}} [(r_{t,T} - \mathbb{E}[r_{t,T}])^2] = \bar{v} + \text{var}(v_t)/4. \quad (44)$$

This example shows that significant differences between average implied central moments and population central moments do not themselves show the presence of risk premia.

B Proofs

B.1 Trading Strategy with Diffusions

Start from (7)

$$\delta Y_{t,T} = \mathbb{E}_{t+\delta t}^{\mathbb{Q}} [g(r_{t+\delta t,T} + \delta f_{t,T}) - g(r_{t+\delta t,T}) - (e^{\delta f_{t,T}} - 1) \mathbb{E}_t^{\mathbb{Q}} [g'(r_{t,T})]] .$$

Expanding we obtain

$$\begin{aligned} \delta Y_{t,T} = \mathbb{E}_{t+\delta t}^{\mathbb{Q}} & \left[\delta f_{t,T} g'(r_{t+\delta t,T}) + \frac{1}{2} \delta f_{t,T}^2 g''(r_{t+\delta t,T}) + O(\delta f_{t,T}^3) \right] \\ & - \left(\delta f_{t,T} + \frac{1}{2} \delta f_{t,T}^2 + O(\delta f_{t,T}^3) \right) \mathbb{E}_t^{\mathbb{Q}} [g'(r_{t,T})] . \end{aligned} \quad (45)$$

With $\delta f_{t,T}$ known at time $t + \delta t$, and using the definitions of G' and G'' gives

$$\delta Y_{t,T} = \delta f_{t,T} G'_{t+\delta t,T} + \frac{1}{2} \delta f_{t,T}^2 G''_{t+\delta t,T} - \delta f_{t,T} G'_{t,T} - \frac{1}{2} \delta f_{t,T}^2 G'_{t,T} + O(\delta f_{t,T}^3) \quad (46)$$

Rearranging we obtain

$$\delta Y_{t,T} = \delta f_{t,T} \delta G'_{t,T} + \frac{1}{2} (\delta G''_{t,T} - G'_{t,T} + G''_{t,T}) \delta f_{t,T}^2 + O(\delta f_{t,T}^3). \quad (47)$$

Finally, taking the limit as $\delta t \rightarrow 0$ under the diffusion assumption we obtain

$$dY_{t,T} = df_{t,T} dG'_{t,T} + \frac{1}{2} (G''_{t,T} - G'_{t,T}) df_{t,T}^2 \quad (48)$$

B.2 Proof of Proposition 2

Part 1. Through a change of variables (13) can be written as

$$s_{t,T} = \frac{6}{F_{t,T} B_{t,T}} \left(\int_1^\infty \left(\frac{a-1}{a^2} \right) \{C_t(aF_{t,T}) - aP_t(F_{t,T}/a)\} da \right). \quad (49)$$

If there is put-call symmetry⁸ in the sense of Carr et al. (1998), so that the implied volatility plotted against log strike is symmetrical about the forward price, then the risk reversal term in curly brackets in the second line costs nothing. This means that in the presence of put-call

⁸Carr et al. (1998) define put-call symmetry to mean that prices of out-of-the-money puts and calls are related

$$C_t(\alpha F_{t,T}) = \alpha P_t(F_{t,T}/\alpha), \text{ all } \alpha > 0.$$

Put-call symmetry is equivalent to requiring that $\sigma_t^I(k)$, the Black-Scholes implied volatility of a call or put option with strike k and maturity T at time t , is symmetrical about the forward price

$$\sigma_t^I(\alpha F_{t,T}) = \sigma_t^I(F_{t,T}/\alpha).$$

symmetry $s_{t,T} = 0$.

Part 2. The claim follows from the positivity of the integrands in (13).

Part 3. From the definition of the Log and the Entropy Contract

$$v_{t,T}^L = -2\mathbb{E}_t^{\mathbb{Q}} [\ln F_{T,T} - \ln F_{t,T}], \text{ and } v_{t,T}^E = 2\mathbb{E}_t^{\mathbb{Q}} \left[\frac{F_{T,T}}{F_{t,T}} (\ln F_{T,T} - \ln F_{t,T}) \right]. \quad (50)$$

Writing x_t for the return $\ln \frac{F_{T,T}}{F_{t,T}}$ we have

$$v_{t,T}^L = -2\mathbb{E}_t^{\mathbb{Q}} [x_t] = -2\mathbb{E}_t^{\mathbb{Q}} [x_t - e^{x_t} + 1], \text{ and } v_{t,T}^E = 2\mathbb{E}_t^{\mathbb{Q}} [x_t e^{x_t}] = 2\mathbb{E}_t^{\mathbb{Q}} [x_t e^{x_t} - e^{x_t} + 1]. \quad (51)$$

Using a Taylor series expansion

$$v_{t,T}^L = \mathbb{E}_t^{\mathbb{Q}} [x_t^2 + x_t^3/3 + \dots], \quad v_{t,T}^E = 2\mathbb{E}_t^{\mathbb{Q}} [x_t^2 + 2x_t^3/3 + \dots], \quad (52)$$

and therefore

$$s_{t,T} = 3(v_{t,T}^E - v_{t,T}^L) = \mathbb{E}_t^{\mathbb{Q}} [x_t^3 + O(x_t)^4]. \quad (53)$$

Part 4. With a self-financing portfolio strategy, starting from zero wealth, we have absence of arbitrage $\mathbb{E}_0^{\mathbb{Q}} [W_T - W_0] = 0$ and as a consequence

$$\underbrace{G_{0,T}}_{\text{fixed leg}} = \mathbb{E}_0^{\mathbb{Q}} \left[\underbrace{\int_{t \in [0,T]} \left\{ dG'_{t,T} df_{t,T} + \frac{1}{2} (G''_{t,T} - G'_{t,T}) (df_{t,T})^2 \right\}}_{\text{realized leg}} \right]. \quad (54)$$

This is the payoff to a swap contract in continuous time. Using $g = g^S$ gives the desired result.

C Robustness Checks

In this section we check whether the results are driven by the the choice of risk factors, specific sample selection, discreteness of strike prices, and whether the results disappear once account is taken of transaction costs.

Table 9: Skew vs. variance risk premia: Fama-French and momentum factors

This table presents estimation results of the following skew and variance risk premia regressions:

$$\begin{aligned} xs_t &= \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + \alpha_3 smb_t + \alpha_4 hml_t + \alpha_5 umd_t + u_t \\ xv_t &= \beta_0 + \beta_1 xs_t + \beta_2 xm_t + \beta_3 smb_t + \beta_4 hml_t + \beta_5 umd_t + w_t. \end{aligned}$$

xs and xv denote excess returns from skew and variance swaps, xm denotes the market excess return, smb and hml denote Fama-French factors and umd is the momentum risk factor. The system is estimated using three-stage OLS. The instrumental variables used for the estimations are ks_t , s_t and ted_t . Market excess return, Fama-French and momentum factors are treated as exogenous. Newey-West t-statistics are given in parentheses. The data are monthly, from January 1996 to January 2012.

	Const	xv	xs	xm	smb	hml	umd	$\bar{R}^2$
xv	0.011 (0.60)		0.526 (22.3)	-2.947 (-5.28)	0.231 (0.36)	0.187 (0.32)	-0.581 (-1.49)	83.47%
xs	-0.021 (-0.53)	1.901 (22.4)		5.602 (4.63)	-0.439 (-0.36)	-0.355 (-0.31)	1.104 (1.46)	81.14%

C.1 Fama-French and Momentum Factors

We start by checking whether the inclusion of Fama-French and momentum risk factors into equations (34) and (35) changes our conclusions. Specifically, we re-estimate the following system of equations using three-stage OLS

$$xs_t = \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + \alpha_3 smb_t + \alpha_4 hml_t + \alpha_5 umd_t + u_t \quad (55)$$

$$xv_t = \beta_0 + \beta_1 xs_t + \beta_2 xm_t + \beta_3 smb_t + \beta_4 hml_t + \beta_5 umd_t + w_t, \quad (56)$$

where smb and hml denote Fama-French factors and umd is the momentum risk factor. As before, we use the three instrumental variables: ks_t , s_t and ted_t . The market excess return, Fama-French and momentum factors are treated as exogenous. The estimation results are given in Table 9. As before, excess returns from skew and variance swaps and the market returns remain significant in both equations, while Fama-French and momentum factors do not. The intercepts in both equations remain economically and statistically insignificant.

C.2 Subsample Analysis

We look separately at the first and the second half of our dataset: from January 1996 to December 2003 and from January 2004 to January 2012. For each sub-period we estimate equations (34) and (35) using three-stage OLS. The estimation results are given in Table 10. The strong correlation between skew and variance risk premia is manifest in both sub-periods,

Table 10: Skew and variance risk premia: sub-samples

This table presents estimation results of the skew and variance risk premia regressions $xs_t = \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + u_t$, and $xv_t = \beta_0 + \beta_1 xs_t + \beta_2 xm_t + w_t$ across two periods: January 1996 – December 2003 and January 2004 – January 2012. xs and xv denote excess returns from skew and variance swaps, xm denotes the market excess return. The system is estimated using three-stage OLS. The instrumental variables used for the estimations are ks_t , s_t and ted_t . Newey-West t-statistics are given in parentheses.

<i>January 1996 – December 2003</i>					
	Const	xv	xs	xm	$\bar{R}^2$
xv	-0.010 (-0.20)		0.512 (5.86)	-2.107 (-2.80)	62.27%
xs	0.007 (0.07)	1.900 (6.01)		3.871 (1.98)	63.94%
<i>January 2004 – January 2012</i>					
	Const	xv	xs	xm	$\bar{R}^2$
xv	0.016 (0.59)		0.507 (10.0)	-3.507 (-3.53)	91.48%
xs	-0.032 (-0.62)	1.969 (10.0)		6.876 (2.67)	88.45%

with the relationship being stronger in the second half. In both sub-samples the intercepts are economically small and statistically insignificant.

C.3 Robustness in Computation of Implied Moments

Exact computation of implied skew and implied variance requires a continuum of option prices at all strikes. To check whether the results are sensitive to the methodology used to compute moments we use a different approach. We generate 1,500 option prices for strikes from \$10 to \$3,000, with step \$5, from the available option quotes. At each maturity we interpolate implied volatilities across all available strike prices using a cubic spline method. For strike prices below and above the lowest and the highest available strikes we use the implied volatilities at the lowest and the highest strike price respectively; this is the same procedure as in [Jiang and Tian \(2005\)](#) and [Carr and Wu \(2009\)](#).

Having recomputed the skew and variance risk premia, we estimate equations (34) and (35). The regression results are in Panel A of Table 11, which are similar to those in Table 5.

Our results rely on the accuracy of prices in the database. Option prices for contracts having no open interest and zero trading volume may be inaccurate or stale. We replicate our results using only options with non-zero open interest and trading volume. Panel B of Table 11 contains the estimation results. Coefficient estimates remain qualitatively similar. Both intercepts are statistically insignificant and R^2 's remain high suggesting that our results are not driven by the prices of illiquid options.

Table 11: Skew and variance risk premia: robustness in computations

This table presents estimation results of the skew and variance risk premia regressions $xs_t = \alpha_0 + \alpha_1 xv_t + \alpha_2 xm_t + u_t$ and $xv_t = \beta_0 + \beta_1 xs_t + \beta_2 xm_t + w_t$. xs and xv denote excess returns from skew and variance swap strategies, xm denotes the market excess return. Panel A uses the interpolation-extrapolation techniques from Jiang and Tian (2005) and Carr and Wu (2009) and Panel B computes implied moments using options with non-zero open interest and trading volume. The system is estimated using three-stage OLS. The instrumental variables used for the estimations are ks_t , s_t and ted_t . Newey-West t-statistics are given in parentheses. The data are monthly, from January 1996 to January 2012.

	Const	xv	xs	xm	$\bar{R}^2$
<i>Panel A: Interpolation-Extrapolation</i>					
xv	0.012 (0.60)		0.505 (31.0)	-2.871 (-6.00)	82.04%
xs	-0.025 (-0.61)	1.981 (31.0)		5.688 (5.51)	77.92%
<i>Panel B: Non-zero Interest</i>					
xv	0.008 (0.39)		0.527 (23.6)	-2.731 (-5.59)	83.35%
xs	-0.016 (-0.39)	1.899 (23.7)		5.184 (4.92)	80.68%

C.4 Transaction Costs

Replication of the variance swap involves setting up an option position at inception and then dynamically trading the underlying asset. By contrast, replication of the skew swap requires dynamic option trading as well as dynamic trading in the underlying asset. With significant bid-ask spreads in the options market, the question naturally arises as to whether the results obtained previously, which assume that options can be traded at the mid-price between the bid and the ask, go away once transaction costs are taken into account. Since transaction costs in the spot index market are very much lower than in the options market we continue to assume they are zero.

Consider the case of an agent replicating a long position in a skew swap. At inception, the agent buys high-strike options and writes low-strike options. With the negative skew in option prices, this is cash-flow positive. The position is delta-hedged in the forward market. As the underlying rises, the agent sells entropy contracts - by selling conventional options at all strikes; conversely when the underlying falls, the agent buys conventional options. In the presence of transaction costs, the agent buys options at the ask and sells at the bid. The transaction costs reduce the profits from the strategy. The agent could reduce the impact of transaction costs by rebalancing less frequently, but this would make the replication noisier.

Rather than solving for the optimal hedging strategy under transaction costs, we assume that the agent rebalances at the end of a day only when the index has moved by more than some threshold of $\kappa\%$ from the last time the portfolio was rebalanced. The skew swap continues

Figure 2: Returns and standard deviations of skew swap strategy under transaction costs vs. no transaction costs

The figure presents average excess returns and their standard deviations for a long position in one-month skew swap under different rebalancing frequencies. $std(xs)$ and xs correspond to standard deviation and excess returns for the case of no transaction costs (options traded at the mid-quotes) while $std(xs_{tc})$ and xs_{tc} denote standard deviation and returns under transaction costs (options traded at bid and ask prices). The x-axis is the filter which controls the rebalancing frequency. The trader rebalances the portfolio only when the index prices changes more than $\kappa\%$. The sample period is January 1996 to January 2012.

to be a self-financing strategy after transaction costs subject to the interpretation that v_t^E is the bid entropy variance (computed using the bid prices of all options) at time t if the index has risen since the portfolio was last rebalanced, and the ask if the index has fallen. The fixed leg s_t is computed from the bid price of low strike puts and the ask price of high strike calls.

If the threshold is set low enough, the portfolio is rebalanced daily. The effect of transaction costs is then substantial. The average monthly return⁹ drops to 4.78%. If the threshold κ is raised, the portfolio is rebalanced less frequently, leading to smaller transaction costs but more noise. Figure 2 presents the means and standard deviations of monthly returns for different values of κ . The variation of the excess returns increases as the rebalancing frequency decreases. The results suggest that the skew risk premium is a real phenomenon that remains even after accounting for transaction costs.

⁹We compute returns as profit net of transaction costs divided by the absolute value of the mid-quote implied skew.