Higher-Moment Risk*

Niels Joachim Gormsen and Christian Skov Jensen†

November 13, 2017

Please Click Here for Latest Version

Abstract

We show how the market's higher order moments can be estimated *ex ante* using methods based on Martin (2017). These ex ante higher order moments predict future realized higher order moments, whereas trailing realized moments have little predictive power. Higher-moment risks move together in the sense that skewness becomes more negative when kurtosis becomes more positive. In addition, higher-moment risk is high when volatility is low, suggesting that risk doesn't go away – it hides in the tails. Higher-moment risk has significant implications for investors; for example, the tail loss probability of a volatility-targeting investor varies from 3.6% to 9.7%, entirely driven by changes in higher-moment risk. We empirically analyze the economic drivers of these risks, such as financial intermediary leverage, market and funding illiquidity, and potential bubbles.

*We are grateful for helpful comments from Peter Christoffersen, Robin Greenwood, Sam Hanson, Christian Julliard, Lukas Kremens, David Lando, Ian Martin, Pia Mølgaard, Lasse Heje Pedersen, Davide Tomio, Mads Vestergaard, Paul Whelan as well as from seminar participants at Copenhagen Business School. Both authors gratefully acknowledge support from the FRIC Center for Financial Frictions (grant no. DNRF102) and from the European Research Council (ERC grant no. 312417).

†Gormsen is at Copenhagen Business School, ng.fi@cbs.dk; Jensen is at Copenhagen Business School, csj.fi@cbs.dk.

Times of financial market distress pose threats to the macroeconomy, as we witnessed in the 2008-2009 financial crisis. For policymakers to act in a timely and preemptive manner in the event of financial market distress, it is important to measure the perceived tail risks in real time.

In this paper, we estimate higher moment risk in real time using a new method and arrive at the following five main results: (1) Moments of the market return, measured ex ante using option prices, predict future realized moments. (2) Higher order moments co-move in the sense that skewness (3^{rd} moment) and hyperskewness (5^{th} moment) become more negative when kurtosis (4^{th} moment) and hyperkurtosis (6^{th} moment) become more positive. In other words, there are times when higher-moment risk is high, in the sense that the return distribution is both substantially left-skewed (due to large negative odd-numbered moments) and fat tailed (due to large positive even-numbered moments). (3) Higher-moment risks tend to be high after market run-ups where the variance is low. (4) Higher-moment risk has important implications for investors; for example, the tail loss probability of a volatility-targeting investor and varies from 3.6% to 9.7%, entirely driven by changes in higher-moment risk. (5) The times when higher-moment risks are high are characterized by high market and funding liquidity, high turnover, and low expected future returns.

Our analysis is based on ex ante moments that are estimated from options prices. Using methods based on Martin (2017), we translate risk-neutral moments into physical moments as perceived by an unconstrained power utility investor who wants to hold the market portfolio. Using S&P 500 as a proxy for the market portfolio, we estimate ex ante monthly and quarterly moments. These moments are entirely forward looking and, unlike risk-neutral moments, contain no adjustment for risk, which makes them well suited for studying time-variation in higher-moment risk.

As our first main result, we show that our ex ante moments are positively correlated with ex post realized moments. Consistent with previous research, our ex ante variance predicts ex post realized variance well.[1] More importantly, we show that

[1] Previous literature has shown that ex post realized variance is well predicted by historical variance or option implied variance, e.g. Bollerslev, Tauchen, and Zhou (2009), Andersen, Fusari,

our ex ante higher order moments also predict ex post higher order moments. We show that our ex ante skewness, kurtosis, hyperskewness, and hyperkurtosis all have significant predictive power over ex post realized moments. We further show that our ex ante moments are better at forecasting ex post realized moments than their trailing (lagged) moments.

Next, we show that these predictability results are robust in several ways. First, we show that our results are not driven by the large price moves that occurred during the financial crisis of 2008 to 2009. Second, we show that our moment prediction holds even when controlling for risk-neutral moments. The latter is important because option-implied risk-neutral skewness has been shown to predict ex post realized skewness, e.g. Neuberger (2012).

As our second main result, we find that higher order moments move together in the sense that skewness and hyperskewness are more negative at times when kurtosis and hyperkurtosis are more positive. Indeed, we find that skewness is negatively correlated with kurtosis with a correlation coefficient of -0.80, a negative correlation of -0.66 with hyperkurtosis, and a positive correlation of 0.79 with hyperskewness. These co-movements in higher order moments are so strong that the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis explains 90% of the joint variation in higher order moments.

The first principal component eigenvector has the same signs for skewness and hyperskewness, while the sign is opposite for kurtosis and hyperkurtosis. As shown in Ebert (2013), an investor with power utility has preferences for odd number moments of any order and is averse to even number moments of any order. A high value of the first principal component can therefore be interpreted as times when higher order moment risks are, on average, large (negative for odd moments and positive for even moments). We therefore define the first principal component as a *higher-moment risk index* (HRI).

As our third main result, we find that higher-moment risk varies systematically

and Todorov (2015), and Bollerslev, Hood, Huss, and Pedersen (2016).

with variance. Specifically, the correlation between variance and the HRI is -0.53 with 95% bootstrapped confidence bounds of $[-0.60, -0.48]$, which emphasizes that higher-moment risks tend to be high at times when variance is low. In addition, we find that higher-moment risks tend to be high subsequent to market run-ups, which are usually "calm" times as measured by variance. We find that the HRI is positively related to the past two year return. The relation is statistically significant at a 99% level, showing that the return distribution is more left skewed and fat tailed subsequent to a "good" period where prices have increased significantly.

Fourth, we show that higher-moment risk has large economic implications for investors. To understand the importance of higher-moment risk, we study the portfolio risk of a volatility-targeting investor who holds a portfolio of cash and the market. The investor adjusts the portfolio weights to achieve a constant volatility of $\sigma^{\text{vol target}}$. Despite having constant variance, the riskiness of the portfolio varies substantially over time as higher moment risk varies. Because higher moment risk is high when variance is low, the portfolio is the riskiest when market variance is low.

To understand the economic magnitude of the systematic variation in higher-moment risks, we estimate the probability that the return on the volatility-targeting investor's portfolio is less than $-2\sigma^{\text{vol target}}$. The monthly probability peaked on June 30$^{\text{th}}$ 2014 with a probability of 9.7%, almost three times the size of its low, on February 27$^{\text{th}}$ 2008, where the probability was 3.6%. Furthermore, the average probability of a $-2\sigma^{\text{vol target}}$ event is 6.6%, which is large compared to the 2.5% that is implied by a normal distribution. Similarly, the probability of a portfolio return that is less than $-3\sigma^{\text{vol target}}$ peaked on November 30$^{\text{th}}$ 2006 with a probability of 3.6%, which is four times the size of its low on February 27$^{\text{th}}$ 2008, when the probability was 0.76%. These probabilities are also far above what is implied by a normal distribution, which is 0.13%.

Furthermore, we find that the probability of a portfolio return that is less than $-2\sigma^{\text{vol target}}$ for the volatility-targeting investor is negatively correlated with variance with a correlation coefficient of -0.70 and 95% bootstrapped confidence bounds of

$[-0.78, -0.65]$. This strong negative correlation further emphasizes the importance of considering higher-moment risks in portfolio choice problems. For example, this finding can help explain why Moreira and Muir (2017b) find that investors can earn high Sharpe-ratios by moving wealth into the market at times when variance is low and moving wealth out of the market when variance increases. The relative (to variance) high expected return in calm times may be compensation for elevated higher-moment risks.

Our fifth main result shows how higher-moment risk is associated with several economic drivers. First, our results are closely related to the volatility paradox (Bunnermeier and Sannikov, 2014), which is the notion that systematic risk is high when variance is low. In their model, risk increases when variance is low because specialized investors are more levered. We therefore investigate how the level of financial intermediary leverage is associated with higher-moment risk. In particular, we test if financial intermediaries are more levered when variance is low, and if such variation in financial intermediary leverage can explain our observed variation in higher moment risk. Using the measure of financial intermediary leverage from He, Kelly, and Manela (2016), we find no relation between higher-moment risks and aggregate financial intermediary leverage.

We next investigate how higher-moment risk is related to market illiquidity and funding illiquidity. We find that higher-moment risks are positively associated with both market and funding liquidity. Specifically, using the average value-weighted bid-ask spread of S&P 500 constituents as a proxy for market illiquidity, we find that times when the average bid-ask spread is low are times when higher-moment risks are high. Similarly, using the TED spread as a proxy for funding illiquidity, we find that a low TED spread is associated with high higher-moment risks.

Lastly, we investigate how higher-moment risks are related to previously suggested measures of "bubble" characteristics and market valuation. We consider the "bubble" characteristics: acceleration (Greenwood, Shleifer, and You (2017)), turnover (Chen, Hong, and Stein (2001)), issuance percentage (Pontiff and Woodgate (2008)), and the market valuation measures: CAPE, the dividend-price ratio, and cay (Lettau and Ludvigson (2001)). We find that higher-moment risk is positively related to price acceleration: there is more higher-moment risk when the recent price path is more convex. Also, higher turnover after market run-ups is associated with more higher-moment risk. Furthermore, there is more higher-moment risk when cay (Lettau and Ludvigson, 2001) is high. We find no conclusive relation between higher-moment risks and CAPE, the dividend-price ratio, or equity issuance.

Our paper relates to and extends the existing literature on estimating time-varying market tail risk by integrating two different approaches. Previous research on tail risk is based on either (1) physical moments based on backward looking information or (2) risk-neutral moments based on forward looking option prices. We show that physical higher-moment risks can be estimated in a forward looking manner, and in real time, which complements the existing literature that uses historical (backward looking) returns to estimate tail risks; e.g., using realized returns, Bollerslev and Todorov (2011) suggest using high frequency intraday returns and fit an extreme value distribution to the tails of returns. Also, Kelly and Jiang (2014) estimate market wide tail risks from the cross-section of firm-level returns. Our paper also relates to the literature that studies tail risk using option prices. However, while the existing literature studies tail risk using risk-neutral moments (e.g. Siriwardane (2015), Gao, Gao, and Song (2017), Gao, Lu, and Song (2017), Bates (2000), and Schneider and Trojani (2017)), we study tail risk using physical moments. Thereby, we can investigate physical tail probabilities and study which economic drivers can explain the time-varying patterns in higher-moment risks.

In summary, higher-moment risks can be measured in real time, and a single factor explains 90% of the joint variation in higher order moments. Furthermore, times when higher-moment risks are high are characterized by: (1) low variance, (2) large (and accelerating) recent price run-ups, (3) low market and funding frictions, (4) high turnover, and (5) low future expected returns.

The paper proceeds as follows: Section 1 covers the theory behind how we estimate higher order moments and tail probabilities. Section 2 covers the data and the empirical implementation. Section 3 investigates the relation between our ex ante moments and ex post realized moments. Section 4 studies the commonalities in higher order moments. Section 5 investigates the systematic patterns in higher-moment risks. Section 6 studies the implications of time-varying higher-moment risks for investors. Section 7 studies the economic drivers of higher-moment risks. Section 8 concludes the paper.

1 Inferring Ex Ante Moments from Asset Prices

We consider an economy where agents can trade two assets, a risk-free asset and a risky asset. The risk-free asset earns a gross risk-free rate of return $R^f_{t,T}$ between time t and time T. The risky asset has a price of S and earns a random gross return $R_{t,T}$. The risky asset pays dividends, $D_{t,T}$, between time t and time T such that its gross return is $R_{t,T} = (S_T + D_{t,T})/S_t$.

Starting from the standard asset pricing formula, we can relate risk-neutral and physical expected values of the time T random payoff, X_T, as

$$E_t[X_T m_{t,T}] = E^*_t[X_T]/R^f_{t,T} \tag{1}$$

where the asterisk denotes risk-neutral expectation and $m_{t,T}$ is a stochastic discount factor. If we define the time T random payoff, $X_{t,T}(n)$, in the following way

$$X_{t,T}(n) = R^n_{t,T} m^{-1}_{t,T} \tag{2}$$

then equation (1) implies that the n'th moment of the risky asset's physical return distribution can be expressed in terms of the risk-neutral expectation of $X_{t,T}(n)$:

$$E_t[R^n_{t,T}] = E_t[\underbrace{R^n_{t,T} m^{-1}_{t,T}}_{X_{t,T}(n)} m_{t,T}] = E^*_t[\underbrace{R^n_{t,T} m^{-1}_{t,T}}_{X_{t,T}(n)}]/R^f_{t,T} \tag{3}$$

So if we know the pricing kernel m, then we can derive all moments of $R_{t,T}$ directly from risk-neutral pricing of the claim to $X_{t,T}(n)$. Following Martin (2017), we compute the physical expected value of $R^n_{t,T}$ from the point of view of an unconstrained rational power-utility investor who chooses to be fully invested in the market. This investor has initial wealth W_0 and terminal wealth $W_T = W_0 R_{t,T}$. Given the investor's utility function, $U(x) = x^{1-\gamma}/(1-\gamma)$, with relative risk-aversion, γ, we can determine the investor's stochastic discount factor. Specifically, combining the first order condition from the investor's portfolio choice problem with the fact that the investor holds the market, the stochastic discount factor becomes proportional to $R^{-\gamma}_{t,T}$:

$$m_{t,T} = k R^{-\gamma}_{t,T} \tag{4}$$

for some constant k which is unobservable to us. However, we do not need to learn k to estimate physical moments; we can correct for k by rewriting (3) in the following way. First, setting $n = 0$ in (2) we get $X_{t,T}(0) = m^{-1}_{t,T}$ and the standard asset pricing formula (1) then implies the relation:

$$E^*_t[m^{-1}_{t,T}] = R^f_{t,T} \tag{5}$$

Then, inserting (5) and (4) into (3), we obtain an expression of the n'th physical moment perceived by an unconstrained rational power utility investor who chooses to be fully invested in the market:

$$E_t[R^n_{t,T}] = \frac{E^*_t[R^n_{t,T}\overbrace{R^\gamma_{t,T}/k}^{m^{-1}_{t,T}}]}{E^*_t[\underbrace{R^\gamma_{t,T}/k}_{m^{-1}_{t,T}}]} = \frac{E^*_t[R^{n+\gamma}_{t,T}]}{E^*_t[R^\gamma_{t,T}]} \tag{6}$$

since k is a constant.

The relation between physical and risk-neutral moments shown in (6) is central to our empirical analysis. The key insight is that we can estimate the n'th physical moment directly from risk-neutral pricing of $R_{t,T}^{\gamma}$ and $R_{t,T}^{n+\gamma}$. Furthermore, by pricing claims to the payoffs $R_{t,T}^{m+\gamma}$ for $m \in \{1, ..., n\}$, we can then estimate *standardized* moments.

To understand how we estimate standardized moments from (6), recall the notion of the n'th standardized moment formula:

$$n\text{'th standardized moment of } R_{t,T} = E_t\left[\left(\frac{R_{t,T} - E_t[R_{t,T}]}{\text{Var}[R_{t,T}]^{1/2}}\right)^n\right] \tag{7}$$

Expanding (7) and replacing physical moments with risk-neutral counterparts as presented in equation (6), we can arrive at expressions for all physical standardized moments as functions of risk-neutral moments. For example, the third standardized physical moment (skewness) can be expressed in terms of risk-neutral moments by first expanding (7) with $n = 3$:

$$\text{Skewness}_{t,T} = \frac{E_t[R_{t,T}^3] - 3E_t[R_{t,T}]E_t[R_{t,T}^2] + 2E_t[R_{t,T}]^3}{(E_t[R_{t,T}^2] - E_t[R_{t,T}]^2)^{3/2}} \tag{8}$$

and then replacing the physical moments in (8) with the risk-neutral counterparts using equation (6). Similar expressions can be written up for other higher order moments of interest, as seen in Appendix A. Importantly, the right-hand-side of (6) consists of asset prices which can be estimated directly from current and observable call and put options written on the risky asset. Hence, higher order moments can be estimated in real time, without using historical realized returns or accounting data.

1.1 Inferring Ex Ante Market Tail Probabilities

Next, we show how we estimate ex ante tail probabilities from option prices written on the market. To understand our approach, note first that the probability at time t of a market return that is lower than α at time T can be written as the physical expectation of an indicator function in the following way

$$P_t(R_{t,T} < \alpha) = E_t[1_{\{R_{t,T}<\alpha\}}] \tag{9}$$

Using the standard asset pricing formula in (1), we can rewrite the probability in terms of the risk-neutral measure by adjusting the right hand side of equation (9) for the inverse of the stochastic discount factor in (4)

$$P_t(R_{t,T} < \alpha) = \frac{E_t^*[R_{t,T}^{\gamma} 1_{\{R_{t,T}<\alpha\}}]}{E_t^*[R_{t,T}^{\gamma}]} \tag{10}$$

The right hand side of (10) is an asset price that has the simple representation presented in Proposition 1, which generalizes Result 2 in Martin (2017) from log-utility to general power utility for any level of relative risk-aversion.

Proposition 1 *For the unconstrained rational power utility investor who wants to hold the market, the conditional physical probability that market return from time t to T is lower than α is:*

$$P_t(R_{t,T} < \alpha) = \frac{R_{t,T}^f}{E_t^*[R_{t,T}^{\gamma}]} \Bigg[\alpha^{\gamma} put'_{t,T}(\alpha S_t - D_{t,T}) - \frac{\gamma}{S_t} \alpha^{\gamma-1} put_{t,T}(\alpha S_t - D_{t,T}) + \int_0^{\alpha S_t - D_{t,T}} \frac{\gamma(\gamma-1)}{S_t^2} \left(\frac{K + D_{t,T}}{S_t} \right)^{\gamma-2} put_{t,T}(K) dK \Bigg]$$

where $put'_{t,T}(\alpha S_t - D_{t,T})$ is the first derivative of the put option price with strike $\alpha S_t - D_{t,T}$.

Proof. The results of Breeden and Litzenberger (1978) imply the equality

$$E_t^*[R_{t,T}^{\gamma} 1_{\{R_{t,T}<\alpha\}}] = R_{t,T}^f \int_0^{\infty} \left(\frac{K + D_{t,T}}{S_t} \right)^{\gamma} 1_{\{K<\alpha S_t - D_{t,T}\}} \mathrm{put}''_{t,T}(K) dK$$

where $\text{put}''_{t,T}(K)$ is the second derivative of the put option price written on the underlying process S. Splitting the integral at $\alpha S_t - D_{t,T}$ we have

$$E_t^*[R_{t,T}^\gamma 1_{\{R_{t,T}<\alpha\}}] = R_{t,T}^f \int_0^{\alpha S_t - D_{t,T}} \left(\frac{K + D_{t,T}}{S_t}\right)^\gamma \text{put}''_{t,T}(K)dK$$

Proposition 1 then follows from using integration by parts twice. ■

2 Data and Empirical Implementation

We use the Ivy DB database from OptionMetrics to extract information on vanilla call and put options written on the S&P 500 index for the last trading day of every month. The data is from January 1996 to December 2015. We obtain implied volatilities, strikes, closing bid-prices, closing ask-prices, and maturities. As a proxy for the risk-free rate, we use the zero-coupon yield curve from the Ivy DB database, which is derived from the LIBOR rates and settlement prices of CME Eurodollar futures. We also obtain expected dividend payments. We consider options with times to maturity between 10 and 720 days, and apply standard filters, excluding contracts with zero open interest, zero trading volume, quotes with best bid below $0.50, and options with implied volatility higher than 100%.

We use daily realized returns to estimate realized daily moments. We also estimate monthly moments from monthly returns. In Appendix A, we discuss the estimation of realized moments in detail.

2.1 Estimating Market Moments

There is a large body of literature devoted to pricing asset derivatives such as those in (6), using observable option prices written on the asset. Indeed, Breeden and Litzenberger (1978), Bakshi and Madan (2000), and Bakshi, Kapadia, and Madan (2003) show that the arbitrage free price of a claim on some future (twice differentiable) payoff can be expressed in terms of a continuum of put and call option prices.

Specifically for our purposes, using the results of Breeden and Litzenberger (1978), Martin (2017) shows that we can write the n'th physical moment of $R_{t,T}$ as

$$E_t[R_{t,T}^n] = \frac{E_t^*[R_{t,T}^{n+\gamma}]}{E_t^*[R_{t,T}^\gamma]} = \frac{(R_{t,T}^f)^{n+\gamma} + R_{t,T}^f\,[p(n+\gamma)+c(n+\gamma)]}{(R_{t,T}^f)^\gamma + R_{t,T}^f\,[p(\gamma)+c(\gamma)]} \tag{11}$$

with

$$\begin{aligned} p(\theta) &= \int_0^{F_{t,T}} \frac{\theta(\theta-1)}{S_t^\theta}\left(S_t R_{t,T}^f - F_{t,T} + K\right)^{\theta-2} \text{put}_{t,T}(K)dK \\ c(\theta) &= \int_{F_{t,T}}^{\infty} \frac{\theta(\theta-1)}{S_t^\theta}\left(S_t R_{t,T}^f - F_{t,T} + K\right)^{\theta-2} \text{call}_{t,T}(K)dK \end{aligned} \tag{12}$$

where $F_{t,T}$ is the forward price and $\text{call}_{t,T}(K)$ and $\text{put}_{t,T}(K)$ are call and put option prices written on the risky asset at time t with horizon $T-t$ and strike K.

In practice, we do not observe a continuum of call and put options and therefore (11) must be numerically approximated. Let $F_{t,T}$ be the forward price and, using the notation from Martin (2017), we can write the price, $\Omega_{t,T}(K)$, at time t of an out-of-the money option with strike K and maturity T as

$$\Omega_{t,T}(K) = \begin{cases} \text{call}_{t,T}(K) & \text{if } K \geq F_{t,T} \\ \text{put}_{t,T}(K) & \text{if } K < F_{t,T} \end{cases}$$

We let $K_1, ..., K_N$ be the (increasing) sequence of observable strikes for the N out-of-the money put and call options and define $\Delta K_i = \frac{K_{i+1}-K_{i-1}}{2}$ with

$$\Delta K_i = \begin{cases} K_{i+1} - K_i & \text{if } i = 1 \\ K_i - K_{i-1} & \text{if } i = N. \end{cases}$$

We approximate the integrals in (12) by observable sums such that the n'th physical moment becomes:

$$E_t[R^n_{t,T}] = \frac{(R^f_{t,T})^{n+\gamma} + R^f_{t,T}\left[\sum_{i=1}^{N} \frac{(n+\gamma)(n+\gamma-1)}{S^{n+\gamma}}(S_t R^f_{t,T} - F_{t,T} + K_i)^{n+\gamma-2}\Omega_{t,T}(K_i)\Delta K_i\right]}{(R^f_{t,T})^{\gamma} + R^f_{t,T}\left[\sum_{i=1}^{N} \frac{\gamma(\gamma-1)}{S^{\gamma}}(S_t R^f_{t,T} - F_{t,T} + K_i)^{\gamma-2}\Omega_{t,T}(K_i)\Delta K_i\right]} \tag{13}$$

In summary, combining equation (13) with the standardized moment formula in equation (7), we can express standardized physical moments in terms of the derivatives prices written on the risky asset.

When we estimate physical moments for a given horizon, say T, for which we do not observe put and call prices, we linearly interpolate the (standardized) moments between the two closest horizons available in the data. In a few cases, we need to extrapolate to obtain moments for the desired horizon.

Our benchmark investor has power utility and a coefficient of relative risk-aversion of 3, that is, $\gamma = 3$. This level of risk-aversion as the benchmark is motivated by the results of Bliss and Panigirtzoglou (2004), i.e., using our sample we replicate their results and find that 3 is the optimal option-implied level of risk aversion when matching realized returns at the monthly horizon. We also estimate moments for the risk-neutral investor, the log-utility investor, and the power utility investor with a risk-aversion coefficient of 5.

Figure 1 shows monthly higher order moments and Table 1 shows the moment summary statistics. The average ex ante estimated skewness is negative for both horizons and all levels of risk aversion, suggesting that the physical distributions are left skewed. Consistent with the results of Neuberger (2012), we find that average skewness is not diminishing in the horizon, in the sense that skewness is close to the same on a monthly and quarterly horizon. Similarly, average kurtosis is larger than 3 for both horizons and all levels of risk aversion, which means that the physical distributions are leptokurtic; that is, the tails of the physical return distributions are fatter than what is implied by a normal distribution.

2.2 Estimating Market Tail Probabilities

The main challenge when implementing Proposition 1 is that we are required to estimate the first derivative of the put option price written on the risky asset at strike $\alpha S_t - D_{t,T}$. To handle a sparse and discrete set of observed option prices, we smoothen observed option prices using a Gaussian kernel smoothening procedure. Specifically, we smoothen implied volatilities around the strike $\alpha S_t - D_{t,T}$ and choose the kernel bandwidth to minimize the squared errors between the observed and estimated implied volatilities under the constraint that the estimated option prices do not allow for arbitrage.

Given a smooth set of option prices around the strike $\alpha S_t - D_{t,T}$, we compute the first derivative as the slope between the two adjacent prices:

$$\text{put}'_{t,T}(\alpha S_t - D_{t,T}) = \frac{\text{put}_{t,T}(\alpha S_t - D_{t,T} + h) - \text{put}_{t,T}(\alpha S_t - D_{t,T} - h)}{2h} \tag{14}$$

where h is the chosen grid step size in the discretization.

Let $K_1, ..., K_M$ be the (increasing) sequence of observable strikes for the M out-of-the money put options where K_M is the observed strike that is closest to $\alpha S_t - D_{t,T}$. We approximate the integral in Proposition 1 by the observable sum:

$$\sum_{i=1}^{M} \frac{\gamma(\gamma-1)}{S_t^2} \left(\frac{K_i + D_{t,T}}{S_t}\right)^{\gamma-2} \text{put}_{t,T}(K_i)\Delta K_i \tag{15}$$

Inserting (14) and (15) into Proposition 1, we can estimate physical probabilities.

3 Estimated Moments Predict Realized Moments

In this section, we show that the ex ante higher order moments estimated using the methods described in Sections 1 and 2 predict ex post realized higher order moments.

We start with a simple sorting exercise. For each moment, we first sort ex post realized monthly returns into a “low” or “high” bucket depending on whether the ex ante moment is lower or higher than its median time series value. Next, we estimate the ex post moments for each bucket; for example, we estimate moments using the monthly ex post returns sorted into the "high" bucket. Figure 2 shows the monthly ex post realized moments of the two buckets for all moments. The ex post realized returns sorted into the "high" buckets exhibit in-sample higher moment values, suggesting that our ex ante moments predict ex post moments, for example, the "high" bucket for kurtosis has an in-sample kurtosis of 5.93, while the "low" bucket has a kurtosis of 2.85.

Next, we test more formally the relation between ex ante and ex post moments. Specifically, we conduct two tests which differ in the way we estimate ex post realized higher order moments. First, we test if the bucket values following our sorting exercise are extreme compared to what a random sample would produce. For each moment, we bootstrap a distribution using permutations and then evaluate where in this distribution our observed "low" and "high" bucket values lie. Panel A of Table 2 reports the values from the ex ante sorting and significance, which is computed from the bootstrapped distribution in the following way: for the "low" buckets, we estimate the frequency at which a random permutation lies below what we observe. For the "high" buckets, we estimate the frequency at which a random permutation lies above what we observe. For example, the -0.83 value for skewness in the "low ex ante" bucket is not in the lower 10% of the bootstrapped distribution and is therefore insignificant at a 90% level. However, the 5.93 value for kurtosis in the "high ex ante" bucket is in the upper 5% of the bootstrapped distribution for kurtosis and is therefore significant at a 95% level. Importantly, our ex ante moments show statistical significance at a 95% level at least once, for every moment except skewness. Comparing these results to the results we get when sorting the ex post realized returns into two buckets based on the trailing monthly moments (estimated using daily returns), we find that our ex ante moments clearly outperform.

Second, we estimate time-varying ex post monthly (and quarterly) realized moments using daily returns; that is, for a given month, we estimate the in-sample moments for that month using the daily returns during that month. The first two columns of Panel B of Table 2 report correlations between our ex ante moments and the ex post realized moments. The latter two columns report correlations between ex post realized moments and their trailing (lagged) moment. We report bootstrapped standard errors in the appendix.

Correlations between our ex ante variances and ex post variances are 49% to 67%, and these correlations are both statistically significant at a 99% level. We also find strong correlations between ex ante and ex post skewness, ranging from 21% to 25%, which are both significantly different from zero at a 99% level. Correlations of our ex ante and ex post hyperskewness are positive and significant at the 99% level. Comparing the correlations of our ex ante moments to those of the trailing moments we find that, on a monthly horizon, trailing moments do not predict either skewness or hyperskewness whereas our ex ante moments do. On a quarterly horizon, trailing moments do predict ex post realized moments, however the correlations are lower than for our ex ante moments.

Neither our ex ante moments nor the trailing moments seem to be able to predict ex post kurtosis or hyperkurtosis. For our ex ante moments, this might be because of the fact that there are fewer available option prices in the right tail of the distribution, that is, deep out-of-the-money call options are traded less frequently than deep out-of-the-money put options. We therefore test if our ex ante kurtosis (and hyperkurtosis) can predict left kurtosis, which is for our purposes the important tail of the distribution to be able to predict. Therefore, we follow Denbee, Julliard, Li, and Yuan (2016), and estimate ex post realized left kurtosis in the following way:

$$\text{Realized left kurtosis}_{t,T} = \frac{\sum_s \left(\frac{\text{Daily return}_s - \text{Realized daily mean}_{t,T}}{\text{Realized daily variance}_{t,T}^{1/2}} \right)^4}{\text{Realized kurtosis}_{t,T}}$$

where s is the days in the month where Daily return < Realized daily mean$_{t,T}$. The realized right kurtosis is defined in the obvious way, where daily returns are larger than the realized mean.

Panel C of Table 2 shows the correlations between our ex ante kurtosis and the ex post realized left kurtosis and left hyperkurtosis. Both on a monthly and quarterly horizon, our ex ante kurtosis and hyperkurtosis are positively and statistically significantly correlated to ex post realized left kurtosis and left hyperkurtosis. This result should be interpreted in the following way: times when our ex ante kurtosis is high are times when the ex post realized kurtosis can be attributed primarily to the left tail of the return distribution. Comparing the correlations between our ex ante moments and ex post realized left kurtosis and left hyperkurtosis to the correlations between trailing moments and ex post left kurtosis and left hyperkurtosis, we find that while monthly trailing moments do not predict ex post moments, quarterly trailing moments do predict ex post realized left kurtosis and hyperkurtosis, but the correlations are smaller than for our ex ante moments.

Overall, Figure 2 and Panel A, B, and C of Table 2 show that our ex ante moments predict ex post realized moments. It is natural to worry that the results are driven by the large price moves that occurred during the period of financial distress from 2008 to 2009. To address this concern, Panel A of Table 3 shows correlations between our ex ante moments and ex post realized moments when removing observations that overlap with the period August 1st 2008-July 31st 2009. The results are largely unchanged, suggesting that the financial crisis does not drive the strong predictive results.

Panel B of Table 3 shows th correlations for other levels of risk-aversion. The results from the point of view of a log-utility investor or a power-utility investor with a risk-aversion of 5 are not remarkably different from the results presented in Panel B of Table 2 for the power-utility investor with a risk-aversion of 3.

As a second robustness test of moment predictability, we ask if physical higher order moments predict ex post realized moments when controlling for risk-neutral moments. Another way to put it is to ask: do we gain anything in terms of predictability for moving from risk-neutral to physical moments? Table 3 shows the results of the following two-stage procedure. In the first stage, we run the two regressions:

$$\text{Realized Moment}_{t,T} = \alpha_1 + \beta_1 E_t^*[\text{Moment}_{t,T}] + \epsilon_{t,T}$$
$$E_t[\text{Moment}_{t,T}] = \alpha_2 + \beta_2 E_t^*[\text{Moment}_{t,T}] + \eta_{t,T}$$

where $E_t[\text{Moment}_{t,T}]$ is our ex ante physical moment and $E_t^*[\text{Moment}_{t,T}]$ is the corresponding risk-neutral moment. The residuals, ϵ and η, are by construction orthogonal to risk-neutral moments, and their correlation therefore determines whether physical moments can explain the variation in realized moments in excess of what is explained by risk-neutral moments. In the second stage we estimate the correlation between ϵ and η. The first two columns of Panel C of Table 3 report these correlations, and bootstrapped standard errors that correct for the generated regressor problem we face when estimating the residuals in the first stage regressions are in the appendix.

The correlations between ϵ and η on a monthly horizon range from 0.09 to 0.16 and are statistically significant at a 95% level for kurtosis, hyperskewnes, and hyperkurtosis, implying that our monthly ex ante moments still predict ex post realized moments when controlling for risk-neutral moments. The results are weaker for quarterly moments; only hyperskewness is statistically significant and positive.

As a third robustness test of predictability, we test if our ex ante estimated higher order moments predict ex post realized moments when controlling for trailing (lagged) moments. We therefore repeat the two-stage procedure described above. In the first stage we run the following two regressions

$$\text{Realized Moment}_{t,T} = \alpha_3 + \beta_3 \text{Realized Moment}_{t-(T-t),t} + \kappa_{t,T}$$
$$E_t[\text{Moment}_{t,T}] = \alpha_4 + \beta_4 \text{Realized Moment}_{t-(T-t),t} + \psi_{t,T}$$

The residuals, κ and ψ, are by construction orthogonal to the historical moments and their correlation therefore determines whether physical moments can explain the variation in the realized moments in excess of what is explained by historical moments. In the second stage we estimate the correlation between κ and ψ. The last two columns of Panel C of Table 3 report these correlations, and bootstrapped standard errors that correct for the generated regressor problem we face when estimating the residuals in the first stage regressions are in the appendix. Controlling for historical (lagged) moments does not change our results. Our ex ante moments have predictive power for ex post realized moments in excess of what is explained by historical moments. Since trailing quarterly moments do predict ex post realized moments, it is particularly important to notice that our quarterly moments add predictability in excess of what the realized trailing moment counterpart can predict.

4 Commonalities in Higher-Moment Risks

Higher order moments exhibit persistent and interesting time-series co-movements, i.e., higher-moment risks move together, in the sense that skewness and hyperskewness are more negative at times when kurtosis and hyperkurtosis are more positive. To see this, Table 4 shows monthly (Panel A) and quarterly (Panel B) pairwise correlations between the first six moments of the physical return distribution. The green (lower right) box shows pairwise correlations between higher order moments. We have flipped the signs for skewness and hyperskewness such that a higher (more positive) value can be translated into higher risk — recall that lower (more negative) skewness and hyperskewness implies more mass in the left tail of the return distribution and therefore higher probabilities of large down movements. These correlations are all positive and large, suggesting that risk as measured by individual higher order moments tends to be simultaneously high or low.

The strong co-movement of higher order moments suggests that the joint variation in higher order moments can be attributed to a single factor. We therefore estimate the principal components of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis. The four principal components are shown in Table 5. Interestingly, at both the monthly and quarterly horizon, the first principal components explains about 90% of the joint variation in higher order moments, underlining the strong co-movement in higher-moment risks.

As was expected, the first principal component eigenvectors have the same signs for skewness and hyperskewness, while the sign is opposite for kurtosis and hyperkurtosis. We standardize each moment to make the eigenvector loadings comparable. The size of the loadings for the first principal components are very similar across the moments, namely -0.45 (-0.47 quarterly) for skewness, 0.52 (0.51 quarterly) for kurtosis, -0.52 (-0.52 quarterly) for hyperskewness, and 0.50 (0.50 quarterly) for hyperkurtosis, implying that the first principal component is approximately the average of the standardized higher order moments with the signs flipped for skewness and hyperskewness. As shown in Ebert (2013), an investor with power utility has a preference for odd number moments of any order and is averse to even number moments of any order. A high value of the first principal component can therefore be interpreted as times when higher order moments (the moments that add mass to the lower tail of the return distribution) are on average large. It is therefore natural to define the first principal component as a higher-moment risk index.

Higher-Moment Risk Index: *We define a higher-moment risk index (HRI) as the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis.*

5 Systematic Variation in Higher-Moment Risks

Figure 1 displays the time-series plot of the monthly HRI which shows clear systematic variation in higher-moment risk. During the period of financial market distress from 2008 to 2009, HRI was low, whereas during the low variance period from 2004 to 2007, leading up to the financial crisis, monthly HRI was high, suggesting that higher-moment risks are high at times when markets are calm. In this section we investigate these systematic patterns.

The blue (upper right) box of Panel A of Table 4 shows the pairwise correlations between variance and higher order moments. Variance is negatively correlated to the negative of skewness, kurtosis, the negative of hyperskewness, and hyperkurtosis with correlations ranging from 0.41 to 0.54. This finding is interesting because it reveals the systematic variation in higher-moment risks; that is, higher-moment risks are high at times when the market is perceived to be safe and calm as measured by variance. Said differently, risk doesn't go away – it hides in the tails.

Figure 3 shows time-series plots of variance and the HRI. In the years after the high variance period in 2003 (following the dot.com bubble), as the market became more and more safe as measured by variance, higher-moment risks move steadily in the opposite direction, i.e., skewness became more negative, kurtosis became more positive, and overall the HRI increase significantly. Furthermore, as the financial crisis started to reveal itself, following the default of the Bear Sterns hedge funds, then market uncertainty spread through higher variance – as the tail of the distribution diminished, higher-moment risks decreased.

Somewhat surprisingly, the HRI peaked on June 30$^{\text{th}}$ 2014, when monthly ex ante variance was at its lowest point in seven years. This period, which was calm as measured by variance, was associated with high higher-moment risks. The main political and economical uncertainty during this period was associated with the economic sanctions made by the US targeting Russia over Russia's continuing involvement in Crimea.

Panel A of Table 6 shows correlations between variances and the HRI. Generally, higher-moment risk as measured through the HRI is high at times when variance is low. On a monthly horizon, the magnitude of the correlation between variance and HRI is -0.53 with 95% bootstrapped confidence bounds of $[-0.60, -0.48]$. The magnitudes and confidence bounds are quantitatively the same for the quarterly HRI.

Related to the co-movements between variance and higher-moment risks, we also find that higher-moment risks tend to be high after recent market run-ups. To show this, Figure 4 shows time-series plots of the past two year return and the HRI. Past returns and the HRI are positively correlated with correlations of 0.38 and 0.35 on monthly and quarterly horizons respectively. To further investigate the dependencies between market run-ups and subsequent higher-moment risks we run a set of regressions of ex ante moments onto the past two year return,[2] $r_{t-24,t} = R_{t-24,t} - 1$:

$$\mathrm{M}_{t,T} = \beta_0 + \beta_1 r_{t-24,t} + \epsilon_{t,T} \tag{16}$$

where the moments, $\mathrm{M}_{t,T}$, are variance, skewness, kurtosis, hyperkurtosis, hyperskewness, and the higher-moment risk index (HRI). Panel B of Table 6 shows the β_1 coefficients of regression (16) and in Panel C of Table 6 we show β_1 coefficients of regression (16) when controlling for the lagged ex ante moment.

We find a negative and significant relation between past returns and variance. This finding is consistent with the intuition that times after market run-ups are "calm" times where risk, as measured by variance, is low. Looking at skewness, we find a statistically significant and negative relation with past returns, implying that the return distribution tilts to the right and leaves more probability mass in the left tail of the return distribution subsequent to market run-ups. Similarly, kurtosis is statistically significant and positive in past returns, hyperskewness is negative in past returns, and hyperkurtosis is positive in past returns. The results are quantitatively similar for monthly and quarterly moments. Panel C of Table 6 shows that controlling for lagged risk does not change our results. We still find strong significant systematic variation in higher order moments.

6 Implications for Investors

The results presented in Table 4, Table 5, and Table 6 show that times when variance is low are times when the market's return distribution is highly left skewed (due to large negative skewness and hyperskewness) and fat tailed (due to large positive kurtosis and hyperkurtosis). That higher-moment risks are high at times when variance is low runs counter to the way we usually think about risk, i.e., we often equate risk with variance, saying that risk is high at times when variance is high. To better un-

[2] This is similar to the market run-up period of Greenwood, Shleifer, and You (2017).

derstand the importance of higher-moment risks, we next investigate portfolio risks for two investors who both hold a portfolio of cash and the market.

The first investor holds a constant notional in the market. The probability that the investor's portfolio realizes an unexpected return (the shock to the portfolio), $r_{t,T}^{\text{shock}} = R_{t,T} - E_t[R_{t,T}]$, less than α is:

$$P_t(r_{t,T}^{\text{shock}} < \alpha) \tag{17}$$

The constant notional investor is exposed to both time-varying variance risk and time-varying higher-moment risks; that is, the probability that the portfolio realizes an unexpected return less than α depends on both conditional variance and conditional higher order moments.

The second investor targets a constant level of portfolio volatility, i.e., the investor moves wealth in and out of the market such that the portfolio has constant volatility. Such volatility-targeting strategies are common practice and have been shown to generate high risk-adjusted returns (e.g. Moskowitz, Ooi, and Pedersen (2012), Asness, Frazzini, and Pedersen (2012), Moreira and Muir (2017a), and Moreira and Muir (2017b)). If $\sigma_{t,T}$ is the market's conditional volatility, $r_{t,T} = R_{t,T} - 1$ is the return on the market, and $r_{t,T}^{f}$ is the risk-free rate of return, then $r_{t,T}^{\text{vol target}}$ is the return on the volatility-targeting investor's portfolio who targets a constant volatility of $\sigma^{\text{vol target}}$:

$$r_{t,T}^{\text{vol target}} = \underbrace{\frac{\sigma^{\text{vol target}}}{\sigma_{t,T}}}_{\omega_{t,T}} r_{t,T} + \left(1 - \frac{\sigma^{\text{vol target}}}{\sigma_{t,T}}\right) r_{t,T}^{f}$$

where $\omega_{t,T}$ is the fraction of wealth held in the market. If $\omega_{t,T} > 1$, the investor levers up by borrowing cash to invest more than all the initial wealth in the market. We assume for simplicity that the investor is unconstrained. The unexpected return of the volatility-targeting investor's portfolio is $r_{t,T}^{\text{vol target,shock}} = r_{t,T}^{\text{vol target}} - E_t[r_{t,T}^{\text{vol target}}]$

which can be rewritten as:

$$
\begin{aligned}
r_{t,T}^{\text{vol target,shock}} &= \omega_{t,T} r_{t,T} + (1-\omega_{t,T})\, r_{t,T}^{f} - \Big(\omega_{t,T} E_t[r_{t,T}] + (1-\omega_{t,T})\, r_{t,T}^{f}\Big) \\
&= \omega_{t,T} r_{t,T}^{\text{shock}}
\end{aligned}
$$

The probability that the volatility-targeting investor's portfolio realizes an unexpected return less than α is:

$$
\begin{aligned}
P_t(r_{t,T}^{\text{vol target,shock}} < \alpha) &= P_t\left(\frac{\sigma^{\text{vol target}}}{\sigma_{t,T}} r_{t,T}^{\text{shock}} < \alpha\right) \\
&= P_t\left(r_{t,T}^{\text{shock}} < \frac{\alpha}{\sigma^{\text{vol target}}}\sigma_{t,T}\right)
\end{aligned} \tag{18}
$$

For example, if $\sigma^{\text{vol target}} = 5\%$ and $\alpha = -10\%$, then the probability that the volatility-targeting investor's portfolio realizes a return that is 10% lower than expected is $P_t\left(r_{t,T}^{\text{shock}} < -2\sigma_{t,T}\right)$.[3] The volatility-targeting investor's portfolio is only exposed to time-varying higher-moment risks, that is, given a level of $\sigma^{\text{vol target}}$, the probability that the investor's portfolio realizes a return less than α depends only on conditional higher order moments. Time-varying variance risk is eliminated by targeting a constant level of portfolio volatility.

Recall that $\sigma_{t,t+h}$ is the ex ante volatility from time t to $t+h$, and we then define $\bar{\sigma}^h$ as the time series average of $\sigma_{t,t+h}$. For example, the time-series average of monthly volatility for the S&P 500 index is $\bar{\sigma}^h = 5.0\%$. Figure 5 shows time-series plots of monthly probabilities, as shown in (17) and (18), where $\alpha = -2\bar{\sigma}^{\text{month}} = -10.1\%$ and the volatility-target is $\sigma^{\text{vol target}} = 5.0\%$.

The top figure shows the probabilities of $-2\sigma_{t,t+1}$ drops in the market, which are the probabilities of the volatility-targeting investor's portfolio return. The horizontal line shows the probability of a $-2\sigma_{t,t+1}$ drop in the market implied by a normal distribution, which is 2.5%. The shaded area between the two lines is higher-moment

[3] Notice that this probability is not necessarily the same as the probability of a portfolio return of -10%. In the example, the probability of a portfolio return of -10% is $P_t\left(r_{t,T}^{\text{shock}} < -2\sigma_{t,T} - E_t[r_{t,T}]\right)$.

risk; that is, the excess probability of a tail event due to negative skewness, excess kurtosis, and all other higher order moments. Interestingly, the probabilities, in excess of what is implied by a normal distribution, range from 1.1% to 7.2%, showing that time-varying higher-moment risks have large economic implications for the risk of the volatility-targeting investor's portfolio. The probability of a $-2\sigma_{t,t+1}$ drop peaked on June 30$^{\text{th}}$ 2014 with a probability of 9.7%, almost three times the size of its low on February 27$^{\text{th}}$ 2008, where the probability was 3.6%. The systematic variation in the tail probabilities, from 3.6% at high variance times to 9.7% at low variance times, emphasizes that investors who manage risk by managing variance are implicitly imposing more risk into their portfolio when variance is low.

The bottom figure shows the probabilities of $-2\bar{\sigma} = -10.1\%$ drops in the market along with the probabilities implied by a normal distribution. The shaded area between the two lines is higher-moment risk for the constant notional investor. The probability of a -10.1% drop in the market is, as expected, high when variance is high. Importantly, higher-moment risk also contributes to the portfolio risk for the constant notional investor, and the economic magnitude is large. For example, the probabilities, in excess of what is implied by a normal distribution, range from 0.5% on October 31$^{\text{st}}$ 2006 to 4.8% on August 31$^{\text{st}}$ 2015. On August 31$^{\text{st}}$ 2015, the total probability of a -10.1% drop was 9.90%, which means that, on that day, 48% of the probability mass in the left tail of the return distribution beyond -10.1% was due to higher-moment risk.

Figure 6 shows time-series plots of monthly probabilities, as shown in (17) and (18), where $\alpha = -3\bar{\sigma}^{\text{month}} = -15.1\%$. The probability of a portfolio return that is less than $-3\sigma_{t,t+1}$ peaked on November 30$^{\text{th}}$ 2006 with a probability of 3.6%, which is four times the size of its low on February 27$^{\text{th}}$ 2008, where the probability was 0.8%. These probabilities are far from what is implied by a normal distribution, which is 0.13%. Specifically, the average probability of a $-3\sigma_{t,t+1}$ event is 1.8%, which is fourteen times higher than what is implied by the normal distribution. Figure 6 shows that higher-moment risk is even more important when evaluating the probability of events further out in the lower tail of the return distribution; that is, the relative amount of probability mass in the lower tail that is due to higher-moment risk increases the further we go out in the tail.

The probabilities co-move in the sense that, when the probability of a $-2\sigma_{t,T}$ event is high, then the probability of a $-2\bar{\sigma}$ event is low. To further investigate these patterns, Panel A of Table 7 reports correlations between variance and the probability of a portfolio return that is less than α for the constant notional investor and the volatility-targeting investor. The first column of Panel A shows the correlations between variance and the probability that the market realizes an unexpected return less than $-2\bar{\sigma}^h$ (the probability of a constant notional investor)

$$P_t(r^{\text{shock}}_{t,t+h} < -2\bar{\sigma}^h)$$

The correlations range from 0.94 to 0.98 with tight bootstrapped confidence bounds, showing that the conditional probability that the market realizes a return less than -10.1% monthly or -15.1% quarterly is highly correlated with conditional variance, which would be expected just from looking at Figure 5.

Panel A of Table 7 also reports correlations between variance and the probability of a portfolio return that is less than α for the volatility-targeting investor. Specifically, we estimate the correlations between variance and the probabilities of a $-2\sigma_{t,t+h}$

$$P_t(r^{\text{shock}}_{t,t+h} < -2\sigma_{t,t+h})$$

This probability is equivalent to the probability in (18) with $\frac{\alpha}{\sigma^{\text{vol target}}} = -2$. Interestingly, the correlations in the last two column of Panel A are all negative and range from -0.70 to -0.44, with tight bootstrapped confidence bounds. These high negative correlations show that the portfolio of the volatility-targeting investor is most risky at times when variance is low, even though the investor has eliminated all dependencies on variance in the portfolio.

This finding can help explain why Moreira and Muir (2017a) and Moreira and Muir

(2017b) find that investors can earn high Sharpe ratios by moving wealth into the market at times of low variance and moving wealth out of the market when variance increases (in some sense mimicking a volatility targeting strategy). The relatively (to variance) high expected return in calm periods may be compensation for the elevated higher-moment risks.

To better understand the systematic variation in higher-moment risks, we next investigate the relation between tail probabilities and past returns. Specifically, we regress tail probabilities onto past two year returns, e.g. the probability of a $-2\sigma_{t,t+1}$ drop as

$$P_t(r^{\text{shock}}_{t,t+h} < -2\sigma_{t,t+h}) = \beta_0 + \beta_1 r_{t-24,t} + \epsilon_{t,T} \tag{19}$$

Panel B of Table 7 reports β_1 coefficients from regressions such as in (19). We find that the probability of both a $-2\sigma_{t,t+1}$ and a $-3\sigma_{t,t+1}$ drop in the market is statistically significant and positively related to past returns. The economic magnitude is such that a 50% market run-up over the past two years implies a 1% higher probability of a monthly $-2\sigma_{t,t+1}$ drop in the market. Furthermore, the monthly probability of a -10% drop in the market is negatively related to past returns, which is to be expected, because this probability is highly correlated to variance, as shown in Table 6, and periods after market run-ups are usually associated with low variance. Panel C of Table 7 reports β_1 coefficients from regressions such as in (19) when controlling for lagged probabilities. Controlling for lagged probabilities does not change our results: high past two year returns imply higher current tail probabilities for the volatility-targeting investor.

Our finding that market run-ups are related to contemporaneously higher higher-moment risks supplements the existing literature that relates market run-ups to subsequent (realized) market "crashes", e.g. Greenwood, Shleifer, and You (2017). Specifically, we find that the probability of an $x\%$ drop in the market decreases in past returns. High past returns means low current volatility and a low probability of a subsequent $x\%$ drop in the market price. However, *conditional on variance*, the probability of an $x\%$ drop in the market increases in past returns.

7 What Explains Higher-Moment Risk?

In this section we investigate three possible explanations for the systematic variation in higher-moment risk. First, we investigate how higher-moment risk is associated with financial intermediary leverage. Second, we study how market and funding liquidity relates to higher-moment risk. Third, we investigate how higher-moment risk is associated with common "bubble" characteristics. Throughout this section, we will focus on monthly horizon ex ante higher-moment risk.

7.1 The Volatility Paradox and Intermediary Leverage

The volatility paradox is the phenomenon that endogenous risk is high even though exogenous risk is low (Brunnermeier and Sannikov (2014)). Loosely speaking, exogenous risk can be seen as variance and endogenous risk can be seen as higher-moment risk. When variance is low, investors take on more risk in their positions, for instance through leverage, which creates endogenous risk. This negative relation between higher-moment risk and variance is closely related to our empirical findings, we therefore test if our finding can be linked to the economic drivers suggested by Brunnermeier and Sannikov (2014).

One way in which this endogenous risk may arise is through intermediary leverage.[4] We test if financial intermediary leverage can help explain higher-moment risks by running the following regression:

$$M_{t,T} = \beta_0 + \beta_1 \text{Leverage}_t + \epsilon_{t,T} \tag{20}$$

[4]Several papers have shown that financial intermediary leverage is associated with asset returns, e.g. He, Kelly, and Manela (2016), He and Krishnamurthy (2013), Adrian and Boyarchenko (2012), and Adrian, Etula, and Muir (2014).

where the risk, $\mathrm{M}_{t,T}$, is variance, skewness, kurtosis, hyperkurtosis, hyperskewness, and the higher-moment risk index (HRI). Leverage is the financial intermediary leverage ratio of He, Kelly, and Manela (2016). Regression (20) relates aggregate financial intermediary leverage to contemporaneous higher-moment risks. Panel A of Table 8 shows the results of regression (20). We find that leverage is positively associated with contemporaneous ex ante variance, which is consistent with financial intermediary leverage being counter-cyclical, as noted in He, Kelly, and Manela (2016). The first column of Panel A shows that aggregate financial intermediary leverage is not related to the HRI: we find a regression coefficient pf -0.15 which is statistically insignificant. Decomposing higher-moment risks into individual moments, we do not find a significant relation between financial intermediary leverage and individual higher order moments. Overall, aggregate leverage does not help explain higher-moment risks.

Next, we test if conditional (on variance) financial intermediary leverage is associated with higher-moment risks. We run the regression:

$$\mathrm{M}_{t,T} = \beta_0 + \beta_1 \mathrm{Leverage}_t + \beta_2 \mathrm{Variance}_{t,T} + \epsilon_{t,T} \tag{21}$$

Panel B of Table 8 reports the results of regression (21). Interestingly, we find that, conditioning on ex ante variance, financial intermediary leverage can help explain contemporaneous higher-moment risks. We find that skewness, kurtosis, hyperskewness, and hyperkurtosis all load statistically significantly on financial intermediary leverage, with negative signs for skewness and hyperskewness and positive signs for kurtosis and hyperkurtosis. Furthermore, the HRI is positively related to conditional financial intermediary leverage. Given a level of ex ante variance, higher leverage is associated with higher contemporaneous higher-moment risks.

Panel C in Table 8 reports regression (21) when controlling for lagged risk. Controlling for lagged risk does not change our results. Aggregate financial intermediary leverage is in general not associated with higher-moment risks. Given a level of ex ante variance, and controlling for lagged risk, higher leverage is associated with higher contemporaneous higher-moment risks.

7.2 Market Liquidity and Funding Liquidity.

Several previous papers link market liquidity and funding liquidity to aspects of the stock market's return distribution. Christoffersen, Feunou, Jeon, and Ornthanalai (2016) suggest market illiquidity as an economic factor driving risk-neutral market variance and jump risks, or equivalently, higher order moments. They argue that market illiquidity is the common culprit of market price drops in cases when the price drop happened without news about fundamentals, and it is therefore a reasonable economic driver of market moments. Brunnermeier and Pedersen (2009) show that, from a theoretical point of view, stocks with low market (and funding) liquidity have high variance because they are associated with high margin requirements. Furthermore, Danilova and Julliard (2015) develop a model in which volatility and illiquidity are jointly determined by the same equilibrium forces.

First, we test if high market illiquidity is associated with high contemporaneous ex ante variance. Thereafter, we investigate the relation between market illiquidity and higher-moment risks. When testing the relation between higher-moment risks (or variance) and market illiquidity, we run the regression:

$$\mathrm{M}_{t,T} = \beta_0 + \beta_1 \text{Bid-ask spread}_t + \epsilon_{t,T} \tag{22}$$

where the risk, $\mathrm{M}_{t,T}$, is variance, skewness, kurtosis, hyperkurtosis, hyperskewness, and the higher-moment risk index (HRI). As a proxy for market illiquidity, we follow Christoffersen, Feunou, Jeon, and Ornthanalai (2016), and use the average value-weighted bid-ask spread of constituents of the S&P 500 index.

Panel A of Table 9 reports the results of regression (22). We find that higher market illiquidity is associated with higher contemporaneous ex ante variance, which is consistent with the model of Brunnermeier and Pedersen (2009). The effect is statistically significant at a 99% level and controlling for lagged variance does not change the result.

The HRI is negatively related to market illiquidity with a regression coefficient of −1.22, which is statistically significant at a 99% level. When we control for lagged HRI, we still get a negative relation between market illiquidity and higher-moment risks, but the relation is insignificant. The negative relation between the HRI and market illiquidity shows that higher-moment risks tend to be high at times when the market is most liquid.

Next, we test the relation between funding illiquidity and higher-moment risks. We run the regression:

$$\mathrm{M}_{t,T} = \beta_0 + \beta_1 \text{TED spread}_t + \epsilon_{t,T} \tag{23}$$

where the risk, $\mathrm{M}_{t,T}$, is variance, skewness, kurtosis, hyperkurtosis, hyperskewness, and the higher-moment risk index (HRI). The TED spread is a common proxy for funding illiquidity, e.g. Frazzini and Pedersen (2014). The TED spread is the three month LIBOR intrabank interest rate minus the three month T-bill interest rate and it is available from the St. Louis FED.

Panel B of Table 9 reports the results of regression (23). Contemporaneous ex ante variance is positively related to funding illiquidity, higher TED spread is associated with higher ex ante variance. We find that the HRI is negatively related to funding illiquidity which means that, higher-moment risks are high at times when there is low friction in the funding market. Controlling for the lagged HRI does not change our result.

Figure 7 shows time-series plots of market illiquidity and funding illiquidity with the HRI. Consistent with the results presented in Table 9, we see that the HRI is negatively correlated with both the bid-ask spread and the TED spread.

7.3 “Bubble” Characteristics

A range of macroeconomic variables have been proposed as possible indicators of increased market “crash” risks, or equivalently, increased higher-moment risks. A partial list of the variables include the suggestion of Chen, Hong, and Stein (2001), who suggest turnover, Pontiff and Woodgate (2008), who use issuance as a characteristic, and Greenwood, Shleifer, and You (2017), who propose price acceleration as a higher-moment risk characteristic.

In this section we investigate the relation between common market “crash” indicators and contemporaneous ex ante higher-moment risks. We therefore run regressions on the form:

$$\text{HRI}_{t,T} = \beta_0 + \beta_1 \text{Characteristic}_t + \epsilon_{t,T} \tag{24}$$

where the “bubble” characteristics are: 1) The Greenwood, Shleifer, and You (2017) variable acceleration, which is defined as the annualized past two year return minus the return of the first year of the two year return. Acceleration captures the convexity in the recent price path and a high value of acceleration is intended to be associated with high contemporaneous ex ante higher-moment risks. 2) Issuance as the percentage of firms in the S&P 500 index that issued equity in the past year. We follow Greenwood, Shleifer, and You (2017), and define an equity issuance as the event that a firm's split-adjusted share count increased by five percent or more. 3) Market turnover. The market valuation measures are: 4) CAPE, the Shiller cyclically adjusted price-earnings ratio. 5) The dividend price ratio as the past two year dividends divided by the current market price. 6) Cay, the Lettau and Ludvigson (2001) log consumption - aggregate wealth ratio.

Table 10 reports the results of regression (24). Marginally, we find that cay is negatively and significantly related to the HRI. Calm times when expected returns, as proxied by cay, are low are times when higher-moment risks are high.

Interacting turnover with the past two year return, we find that turnover is positively related to the HRI. This finding is consistent with the findings in Chen, Hong, and Stein (2001), that is, subsequent to market run-ups, a higher turnover is associated with higher higher-moment risks.

For issuance, we find that, subsequent to market run-ups, a higher level of equity issuance implies a lower level of contemporaneous higher-moment risks. This finding is counter to the results of Pontiff and Woodgate (2008). Firms have incentives to issue equity when the stock price is higher than its fundamental value, which should be associated with higher contemporaneous higher-moment risks. Other characteristics show no marginal relation with higher-moment risks.

As a last test, we run a horse race including all "bubble" characteristics. The last column of Table 10 reports the results of the horse race. Jointly, we find that acceleration is statistically significant and positively related to the HRI, and cay is negatively related to the HRI. Conditional on market run-ups, turnover is positively associated with the HRI. Interestingly, issuance changes sign in the horse race compared to the marginal regressions. Indeed, we find that higher issuance is associated with higher contemporaneous higher-moment risks when controlling for other characteristics.

Figure 8 shows time-series plots of the HRI and cay. The two time-series are negatively correlated with an in-sample correlation coefficient of -0.57. Figure 9 shows time-series plots of the HRI and turnover times past two year return. The two time-series are positively correlated with a correlation coefficient of 0.43, implying that, after market run-ups when turnover is high, then so are higher-moment risks.

8 Conclusion: when volatility is low, risk hides in the tails

We show that ex ante physical moments estimated using methods based on Martin (2017) are superior to historical moments and risk-neutral moments at predicting ex ante realized moments.

Ex ante higher order moment risks co-move such that the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis explains 90% of the joint variation. We define this first principal component as a higher-moment risk index (HRI) which captures the time-variation in higher-moment risks.

Interestingly, the HRI is negatively related to variance. We show that times when variance is low are the times when the physical return distribution is most left skewed (due to large negative skewness and hyperskewness) and fat tailed (due to large positive kurtosis and hyperkurtosis). The economic importance of higher-moment risk is most easily understood from the point of view of a volatility-targeting investor. The portfolio risk of this investor is high at times when variance is low, even though the investor has eliminated variance risk in the portfolio. For example, the probability that the investor's portfolio realizes a return less than two standard deviations varies from 3.6% during times of financial distress to 9.7% during periods of low variance.

We show empirically how higher-moment risk is associated with market liquidity, funding liquidity, turnover, and the market valuation variable cay. Times with low liquidity frictions, low cay, and high turnover are times when higher-moment risks are high.

References

Adrian, Tobias, and Nina Boyarchenko, 2012, Intermediary leverage cycles and financial stability, *Working paper*.

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *The Journal of Finance* 69, 2557–2596.

Amaya, Diego, Peter Christoffersen, Kris Jacobs, and Aurelio Vasquez, 2015, Does realized skewness predict the cross-section of equity returns?, *Journal of Financial Economics* 118, 135–167.

Andersen, Torben, Nicola Fusari, and Viktor Todorov, 2015, The risk premia imbedded in index options, *Journal of Financial Economics* 117(3), 558–584.

Asness, Cliff, Andrea Frazzini, and Lasse Heje Pedersen, 2012, Leverage aversion and risk parity, *Financial Analysts Journal* 68, 47–59.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies*.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and derivative-security valuation, *Journal of Financial Economics*.

Bates, D. S., 2000, Post 87 crash fears in the S&P 500 futures options market, *Journal of Econometrics* 94, 181–238.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option implied risk aversion estimates, *Journal of Finance* 1, 407–446.

Bollerslev, Tim, Benjamin Hood, John Huss, and Lasse Heje Pedersen, 2016, Risk everywhere: Modeling and managing volatility, *Duke, working paper*.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *The Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears and risk premia, *Journal of Finance* 66, 2165–2211.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *The Journal of Business* 51, 621–651.

Brunnermeier, Markus, and Yuliy Sannikov, 2014, A macroeconomic model with a financial sector, *American Economic Review* 104, 379–421.

Brunnermeier, Markus .K., and Lasse H. Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Chen, Joseph, Harrison Hong, and Jeremy C Stein, 2001, Forecasting crashes: Trading volume, past returns and conditional skewness in stock prices, *Journal of Financial Economics* 61, 345–391.

Christoffersen, Peter, Bruno Feunou, Yoontae Jeon, and Chayawat Ornthanalai, 2016, Time-varying crash risk: The role of stock market liquidity, *Working paper*.

Danilova, Albina, and Christian Julliard, 2015, Information asymmetries, volatility, liquidity, and the tobin tax, *Working paper*.

Denbee, Edward, Christian Julliard, Ye Li, and Kathy Yuan, 2016, Network risk and key players: A structural analysis of interbank liquidity, *Working paper*.

Ebert, Sebastian, 2013, Moment characterization of higher-order risk preferences, *Theory and Decision* 74, 267–284.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gao, George, Pengjie Gao, and Zhaogang Song, 2017, Do hedge funds exploit rare disaster concerns?, *Review of Financial Studies forthcoming*.

Gao, George, Xiaomeng Lu, and Zhaogang Song, 2017, Tail risk concerns everywhere, *Management Science, forthcoming*.

Greenwood, Robin, Andrei Shleifer, and Yang You, 2017, Bubbles for Fama, *Journal of Financial Economics, forthcoming.*

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2016, Intermediary asset pricing: New evidence from many asset classes, *forthcoming in Journal of Financial Economic.*

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–770.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *The Review of Financial Studies* 27, 2841–2871.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *The Journal of Finance* 56, 815–849.

Martin, Ian, 2017, What is the expected return on the market?, *Quarterly Journal og Economics* 132(1), 367–433.

Moreira, Alan, and Tyler Muir, 2017a, Should long-term investors time volatility?, *Working paper.*

Moreira, Alan, and Tyler Muir, 2017b, Volatility-managed portfolios, *The Journal of Finance* 72, 1611–1644.

Moskowitz, Tobias, Yao Hua Ooi, and Lasse Heje Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.

Neuberger, Anthony, 2012, Realized skewness, *The Review of Financial Studies* 25, 3423–3455.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Pontiff, Jeffrey, and Artemiza Woodgate, 2008, Share issuance and cross-sectional returns, *The Journal of Finance* 63, 921–945.

Schneider, Paul, and Fabio Trojani, 2017, Fear trading, *Working paper*.

Siriwardane, Emil, 2015, The probability of rare disasters: Estimation and implications, *Working paper*.

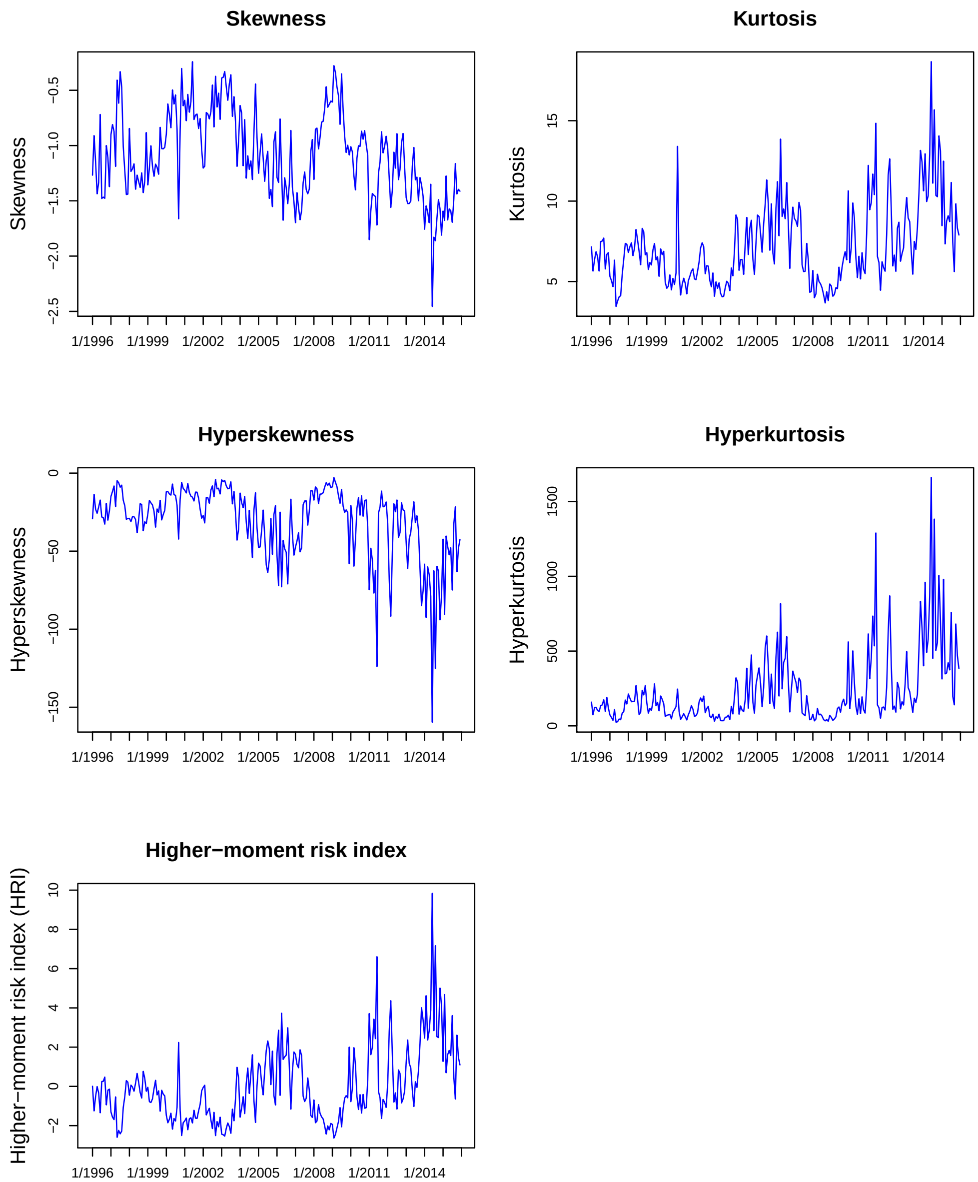

Figure 1: **Higher order moments and the higher-moment risk index.** The figures show a time-series plot of monthly higher order moments and the higher-moment risk index (HRI) for the S&P 500 index. HRI is estimated as the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis. Times when the HRI is high are times when higher-moment risks are high.

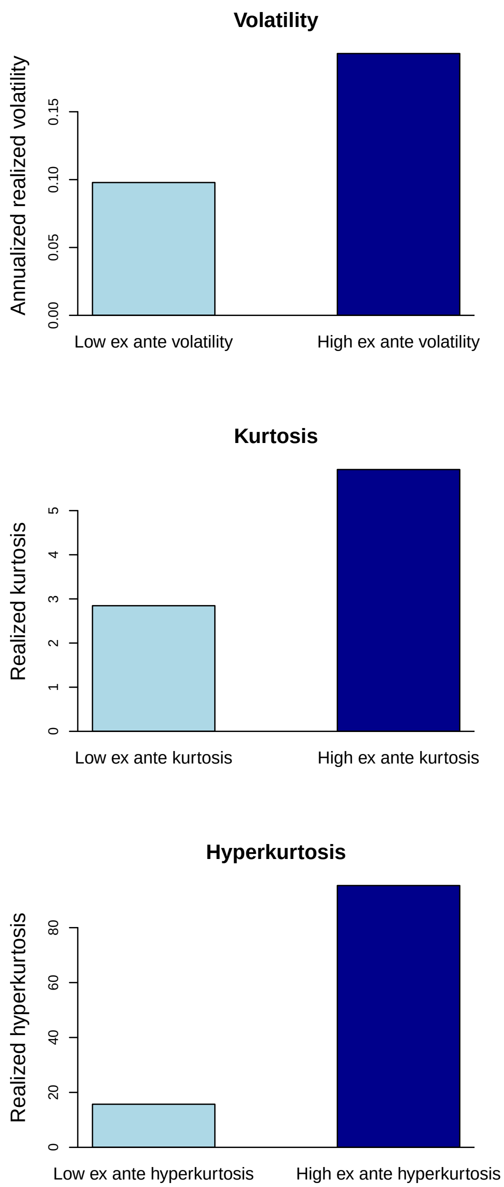

Figure 2: **Moments of ex post realized returns sorted on ex ante expected moments.** For each moment, we sort ex post realized monthly returns into a "low" and "high" bucket based on the median ex ante expected value of that moment. Thereafter, we compute the corresponding realized moment for each bucket.

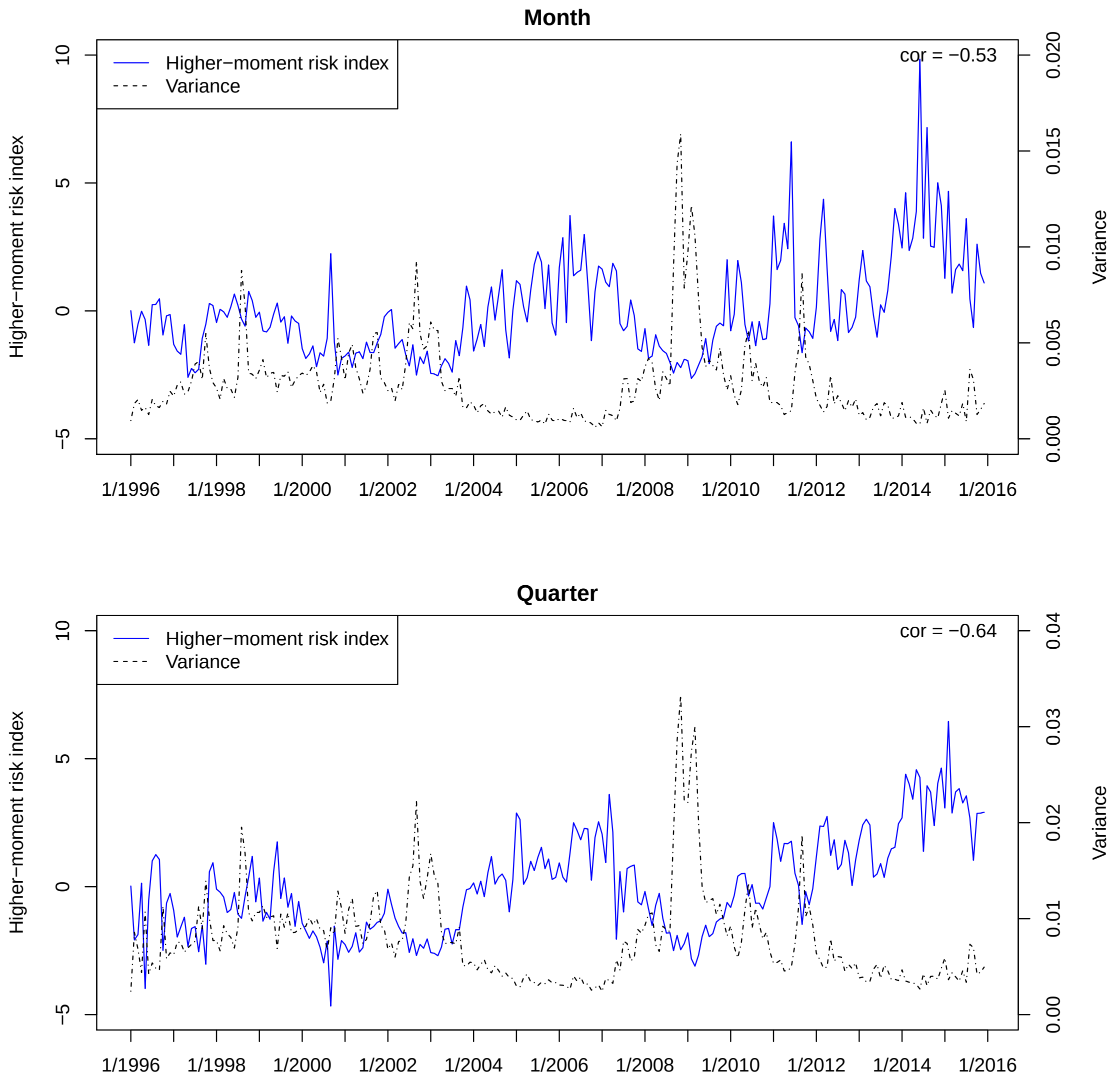

Figure 3: **Higher-moment risk index and variance.** This figure shows time-series plots of monthly and quarterly higher-moment risk index (HRI) and variances. The HRI is high at times when variance is low.

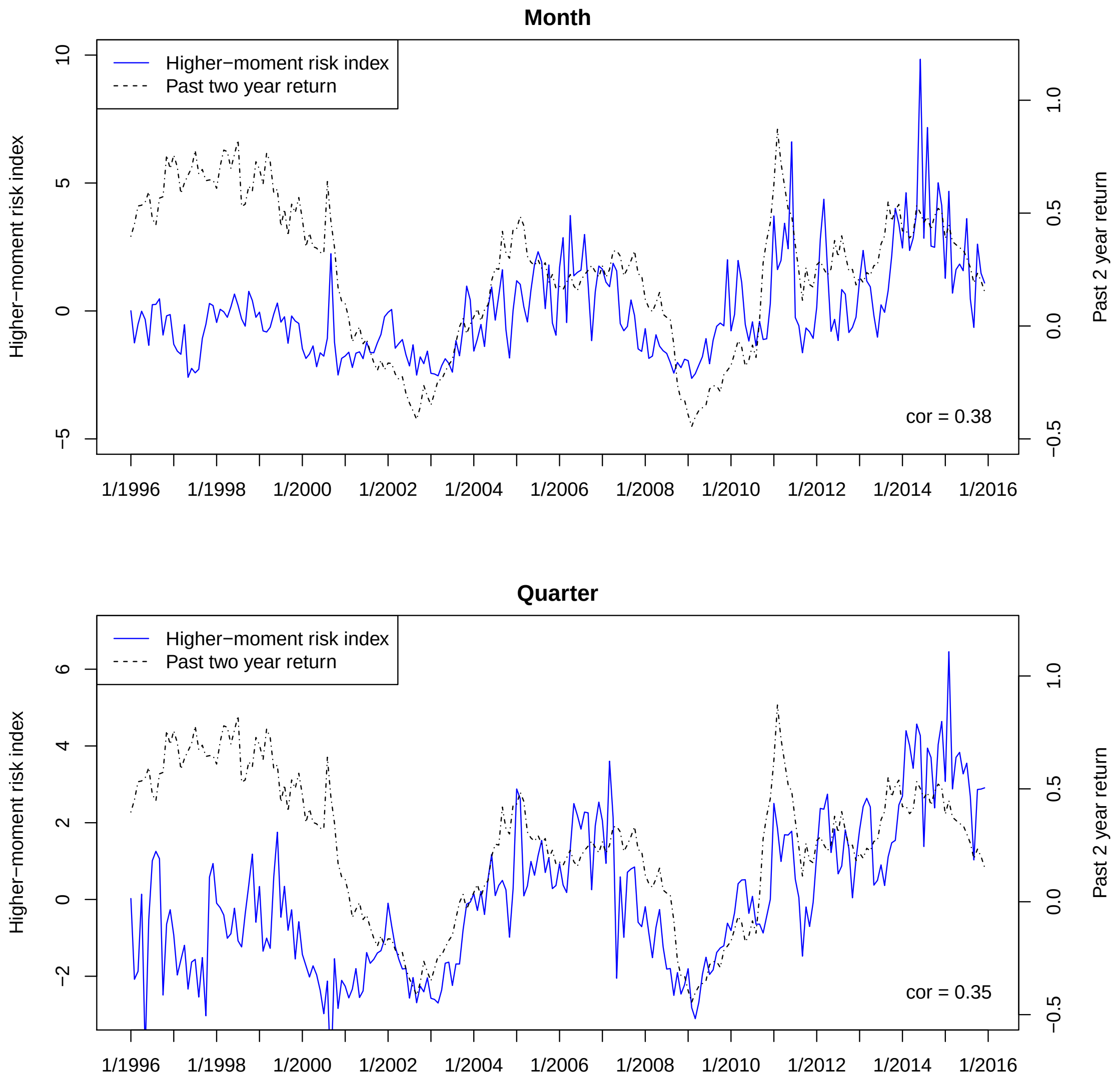

Figure 4: **Higher-moment risk index and the past two year return.** These figures show time-series plots of the past two year return and the S&P 500 higher-moment risk index (HRI). Past return and the HRI are positively correlated, implying that higher-moment risk is high subsequent to market run-ups.

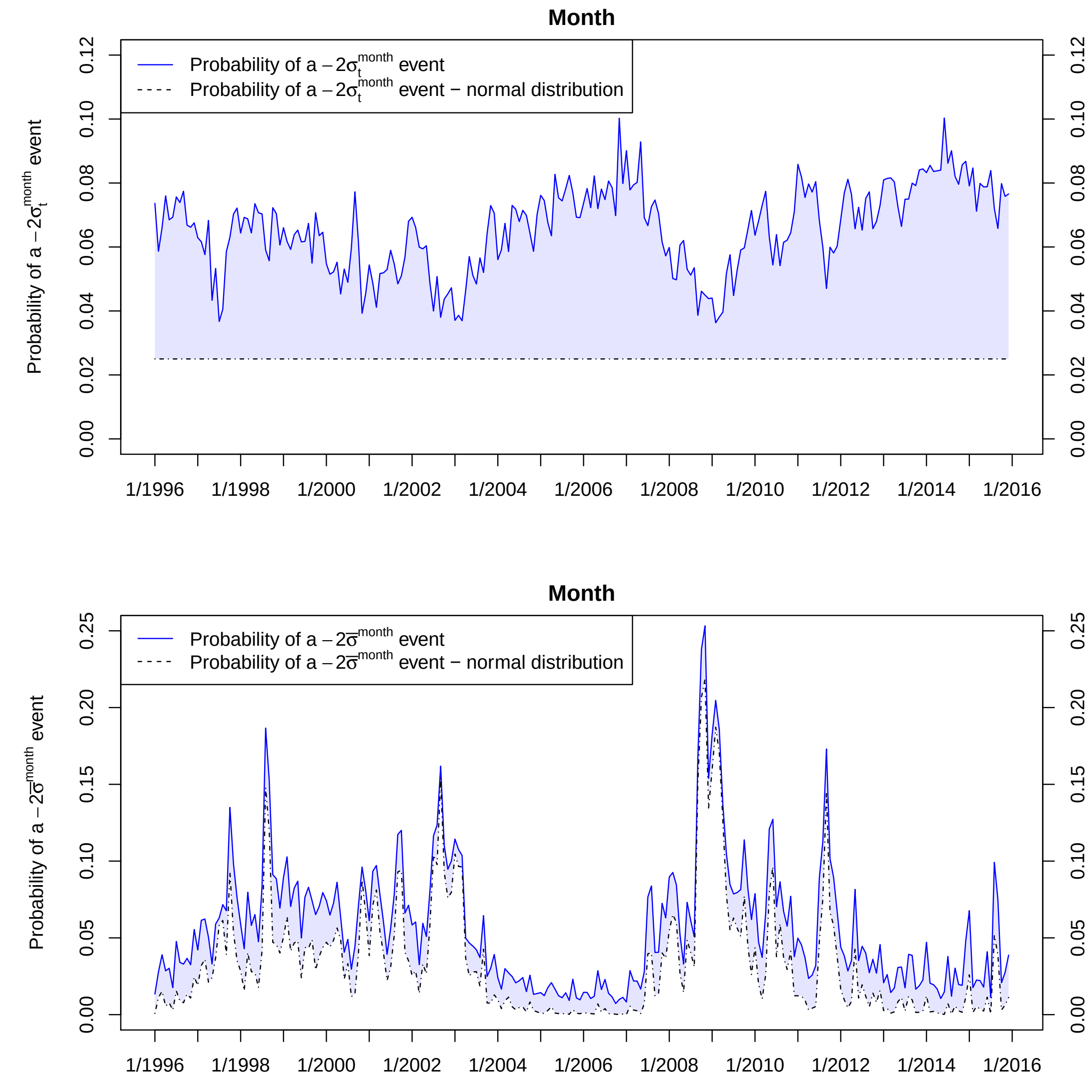

Figure 5: **Market tail loss probabilities – two sigma.** The top figure shows portfolio tail loss probabilities for the volatility-targeting investor; that is, the probability of an unexpected return lower than $-2\sigma_t^{\text{month}}$. The dashed line is the tail loss probabilities implied by a normal distribution. The shaded area between the lines is higher-moment risk, that is, the part of the tail loss probability that is entirely driven by changes in higher order moments. The bottom figure shows portfolio tail loss probabilities for the constant notional investor. Here, σ_t^{month} is the conditional monthly ex ante variance and $\bar{\sigma}^{\text{month}}$ is the time-series average of σ_t^{month}. In our sample, $\bar{\sigma}^{\text{month}} = 5.0\%$.

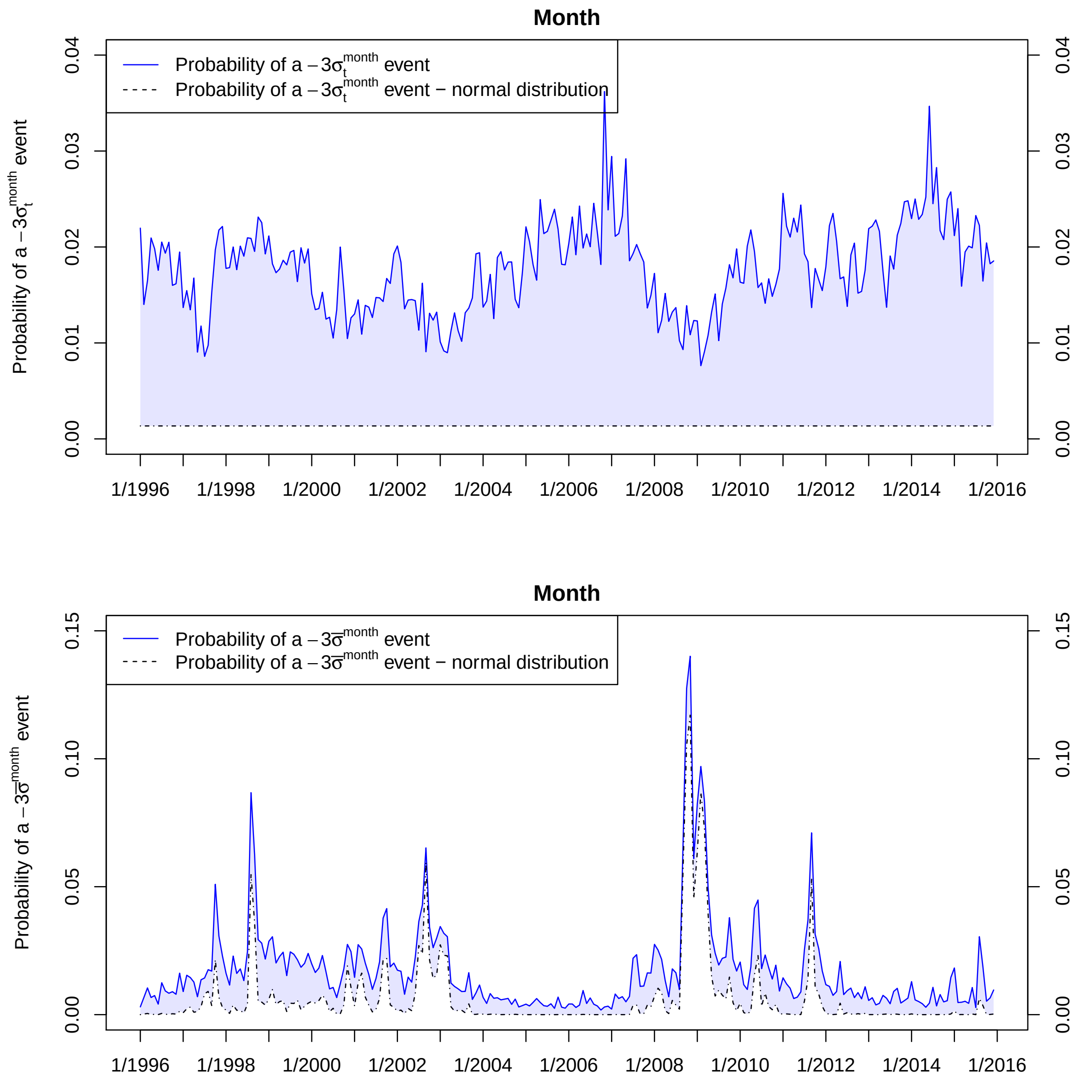

Figure 6: **Market tail loss probabilities – three sigma.** The top figure shows portfolio tail loss probabilities for the volatility-targeting investor; that is, the probability of an unexpected return lower than $-3\sigma_t^{\text{month}}$. The dashed line is the tail loss probabilities implied by a normal distribution. The shaded area between the lines is higher-moment risk, that is, the part of the tail loss probability that is entirely driven by changes in higher order moments. The bottom figure shows portfolio tail loss probabilities for the constant notional investor. Here, σ_t^{month} is the conditional monthly ex ante variance and $\bar{\sigma}^{\text{month}}$ is the time-series average of σ_t^{month}. In our sample, $\bar{\sigma}^{\text{month}} = 5.0\%$.

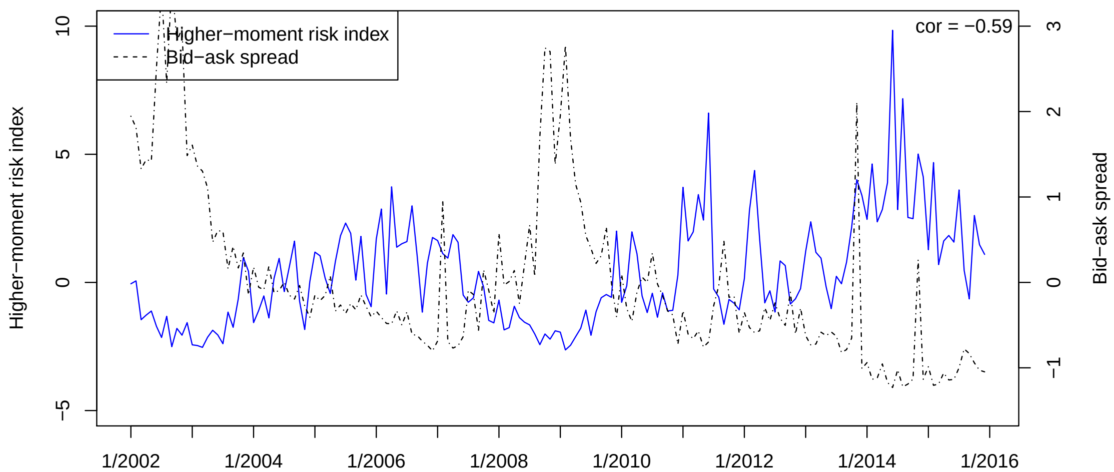

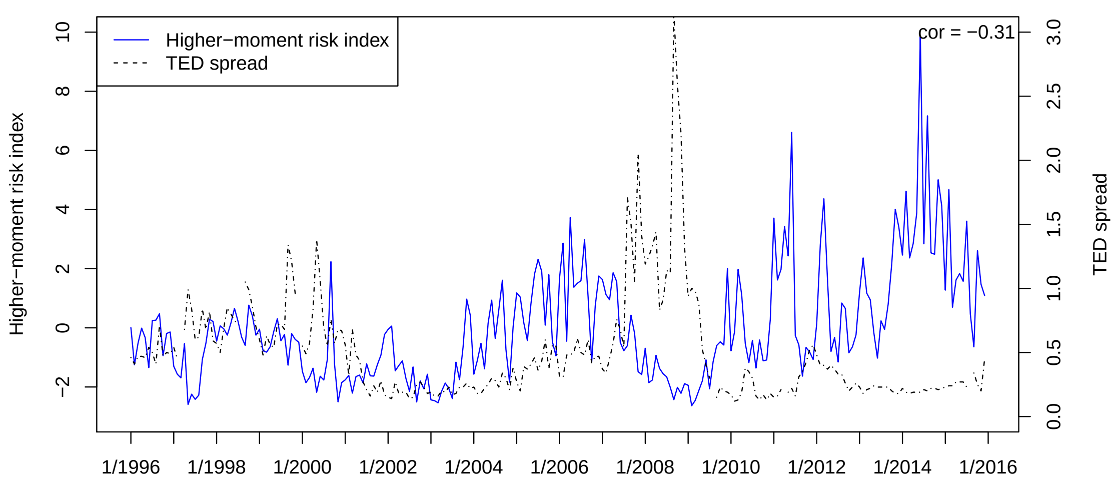

Figure 7: **Higher-moment risk index and market and funding illiquidity.** The top figure shows time series plots of the HRI and market illiquidity (proxied by the average value-weighted bid-ask spread of S&P 500 constituents). The bottom figure shows time series plots of the HRI and funding illiquidity (proxied by the TED spread).

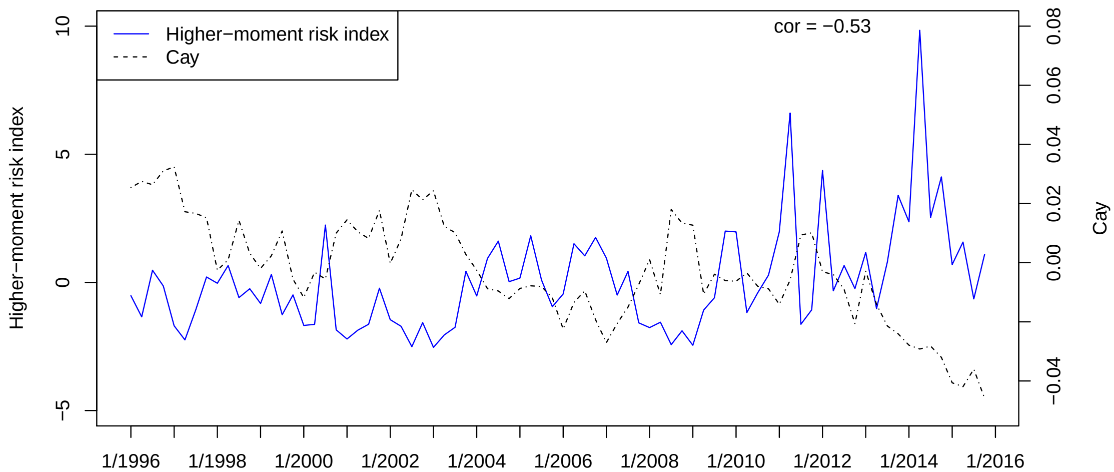

Figure 8: **Higher-moment risk index and cay.** The figure shows time series plots of the HRI and cay.

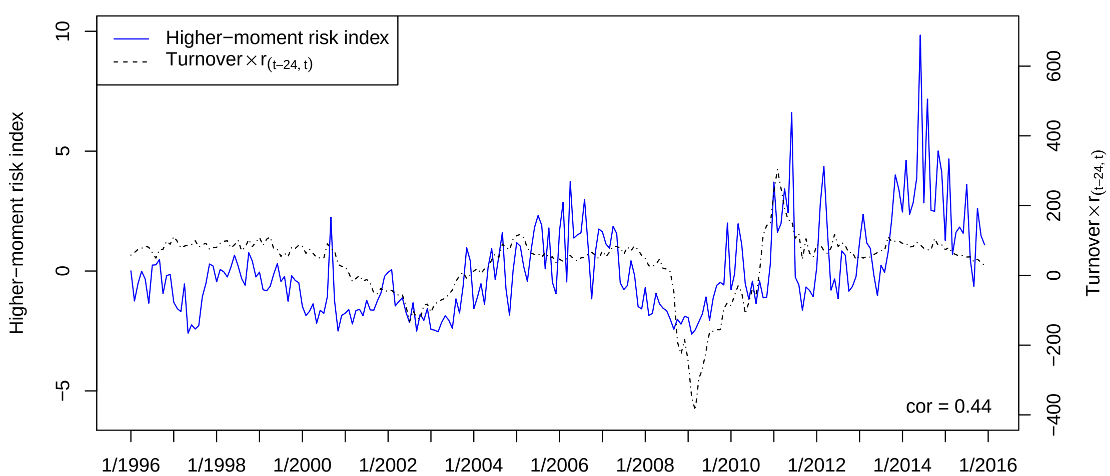

Figure 9: **Higher-moment risk index and turnover.** The figure shows time series plots of the HRI and turnover times the past two year return.

Table 1: **Moment Summary Statistics.** In this table we report the average time-series values for ex ante estimated moments: excess return (ER−Rf), standard deviation (St. dev.), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). We estimate ex ante moments from the point of view of a risk-neutral investor ($\gamma = 0$), a log-utility investor ($\gamma = 1$), and two power-utility investors ($\gamma = 3$, $\gamma = 5$).

		Annualized (%)					
Horizon	Risk-aversion	ER−Rf	St. dev.	Skew	Kurt	Hskew	Hkurt
Month	$\gamma = 0$	0	21.07	-1.45	8.90	-46.58	347.41
Month	$\gamma = 1$	4.44	19.89	-1.31	8.25	-41.01	307.44
Month	$\gamma = 3$	12.00	18.33	-1.08	7.20	-31.16	233.07
Month	$\gamma = 5$	18.36	17.32	-0.89	6.43	-23.47	175.88
Quarter	$\gamma = 0$	0	21.07	-1.17	5.78	-20.58	110.36
Quarter	$\gamma = 1$	4.44	19.59	-1.09	5.57	-18.64	97.45
Quarter	$\gamma = 3$	11.48	17.55	-0.95	5.23	-15.56	80.69
Quarter	$\gamma = 5$	17.12	16.25	-0.81	4.94	-12.93	68.19

Table 2: **Ex Ante Conditional Moments Predict Ex Post Realized Moments.** Panel A reports ex post moments for monthly returns sorted into a low or high bucket based on the ex ante moment. Panel B reports correlations between our ex ante moments and ex post realized moments. Panel C reports correlations between our ex ante kurtosis and hyperkurtosis with ex post left kurtosis and left hyperkurtosis. We also report correlations between historical moments and ex post moments. We report bootstrapped standard errors in the appendix and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$.

Panel A: Sorting on ex ante monthly moments

	Our moments		Historical moments	
	Low ex ante	High ex ante	Low ex ante	High ex ante
Variance (%)	0.08***	0.31***	0.07***	0.28***
Skewness	−0.83	−0.48	−0.49	−0.75
Kurtosis	2.85	5.93**	3.35	4.46
Hyperskewness	−19.03**	−4.24	−5.40	−9.01
Hyperkurtosis	15.70	95.24***	15.34	40.10

Panel B: Correlation between ex ante moments and ex post realized moments

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Variance	0.67***	0.49***	0.72***	0.46***
Skewness	0.21***	0.25***	0.07	0.24***
Kurtosis	−0.01	0.00	0.06	0.04
Hyperskewness	0.17***	0.20***	0.07	0.15***
Hyperkurtosis	−0.03	0.06	0.03	0.01

Panel C: Left kurtosis and left hyperkurtosis

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Left kurtosis	0.19***	0.26***	0.02	0.14**
Left hyperkurtosis	0.17***	0.21***	0.05	0.13**

Table 3: **Ex Ante Conditional Moments Predict Ex Post Realized Moments — Robustness.** Panel A reports correlations between our ex ante moments and ex post realized moments when we remove observations that overlap with the period from August 1st 2008 to July 31st 2009. Panel B reports correlations between our ex ante moments (estimated with different levels of relative risk aversion) and ex post realized moments. Panel C reports correlations when controlling for risk-neutral moments or historical moments. We report bootstrapped standard errors in the appendix and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$..

Panel A: Excluding August 1st 2008 to July 31st 2009

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Variance	0.52***	0.49***	0.51***	0.42***
Skewness	0.23***	0.23***	0.11	0.25***
Kurtosis	−0.01	0.00	0.07	0.06
Hyperskewness	0.18***	0.18***	0.09	0.13***
Hyperkurtosis	−0.03	0.04	0.04	0.01

Panel B: Other levels of risk-aversion

	$\gamma = 1$		$\gamma = 5$	
	Month	Quarter	Month	Quarter
Variance	0.67***	0.48***	0.67***	0.48***
Skewness	0.20***	0.25***	0.21***	0.23***
Kurtosis	−0.03	0.00	0.01	0.02
Hyperskewness	0.13***	0.17***	0.20***	0.22***
Hyperkurtosis	−0.06	0.04	−0.01	0.07

Panel C: Marginal correlations

	Controlling for risk-neutral moments		Controlling for historical moments	
	Month	Quarter	Month	Quarter
Variance	0.09	0.12	0.18**	0.20***
Skewness	0.09	0.01	0.20***	0.17***
Kurtosis	0.11**	−0.02	−0.02	0.01
Hyperskewness	0.16***	0.14***	0.16***	0.17**
Hyperkurtosis	0.09**	0.07	−0.03	0.05

Table 4: **Correlations Between S&P 500 Moments.** Panel A reports pairwise correlations between monthly S&P 500 moments. Expected return (Er), variance (Var), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). Panel B shows the correlation between quarterly horizon moments. We report 95% bootstrapped confidence bounds in brackets.

Panel A: Month

	Er	Var	−Skew	Kurt	−Hskew	Hkurt
Er	1	0.99	−0.46	−0.50	−0.48	−0.41
		[0.99,1]	[−0.56,−0.37]	[−0.57,−0.46]	[−0.55,−0.43]	[−0.48,−0.37]
Var		1	−0.52	−0.54	−0.51	−0.44
			[−0.60,−0.43]	[−0.60,−0.50]	[−0.58,−0.47]	[−0.51,−0.40]
-Skew			1	0.80	0.78	0.66
				[0.76,0.84]	[0.74,0.82]	[0.60,0.72]
Kurt				1	0.97	0.93
					[0.95,0.98]	[0.90,0.95]
-Hskew					1	0.98
						[0.97,0.98]
Hkurt						1

Panel A: Quarter

	Er	Var	−Skew	Kurt	−Hskew	Hkurt
Er	1	0.99	−0.46	−0.54	−0.58	−0.56
		[0.99,0.99]	[−0.55,−0.37]	[−0.60,−0.49]	[−0.64,−0.54]	[−0.62,−0.52]
Var		1	−0.54	−0.62	−0.65	−0.62
			[−0.62,−0.47]	[−0.67,−0.57]	[−0.71,−0.61]	[−0.68,−0.57]
-Skew			1	0.83	0.86	0.74
				[0.79,0.86]	[0.82,0.90]	[0.68,0.79]
Kurt				1	0.96	0.94
					[0.95,0.97]	[0.92,0.96]
-Hskew					1	0.97
						[0.96,0.98]
Hkurt						1

Table 5: **Principal Components of Higher-Moment Risks.** We estimate the four principal components (PC) spanning the space of monthly (Panel A) and quarterly (Panel B) skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). Panel A reports the loadings on each of the monthly moments. Panel B reports the loadings on each of the quarterly moments. The last column of Panel A shows that the first principal component (PC 1) explains 89% of the variation in monthly higher order moments. Similarly, the last column of Panel B shows that 91% of the variation in quarterly higher order moments is captured by the first principal component.

Panel A: Month

	Skew	Kurt	Hskew	Hkurt	Variation explained
PC 1 eigenvector	−0.45	0.52	−0.52	0.50	89%
PC 2 eigenvector	0.85	0.07	−0.20	0.48	10%
PC 3 eigenvector	−0.23	−0.83	−0.16	0.48	1%
PC 4 eigenvector	−0.13	0.19	0.81	0.54	0%
PC 1 correlation	−0.85	0.98	−0.99	0.95	

Panel B: Quarter

	Skew	Kurt	Hskew	Hkurt	Variation explained
PC 1 eigenvector	−0.47	0.51	−0.52	0.50	91%
PC 2 eigenvector	−0.84	−0.16	0.11	−0.51	7%
PC 3 eigenvector	−0.13	−0.84	−0.32	0.40	2%
PC 4 eigenvector	−0.25	0.01	0.78	0.57	0%
PC 1 correlation	−0.89	0.98	−0.99	0.96	

Table 6: **Cyclicality in Higher-Moment Risks.** Panel A reports correlations between ex ante variance and the higher-moment risk index (HRI). We report bootstrapped 95% confidence intervals in brackets. Panel B reports β_1 coefficients when regressing physical moments onto the past two year returns:

$$M_{t,T} = \beta_0 + \beta_1 r_{t-24,t} + \epsilon_{t,T}$$

where the moment $M_{t,T}$ is variance (Var), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), hyperkurtosis (Hkurt), and the higher-moment risk index. Panel C reports the regression when controlling for lagged moments. We report t-statistics in parentheses and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. We correct standard errors for autocorrelation using Newey and West (1987).

Panel A: Variance and the higher-moment risk index

Horizon	HRI
Month	−0.53
95% CI	[−0.60,−0.48]
Quarter	−0.64
95% CI	[−0.69,−0.59]

Panel B: Past return and higher-moment risks

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	2.26***	−0.30*	−0.61***	2.75***	−24.57***	201.65**
(sd)	(0.79)	(0.16)	(0.14)	(1.01)	(9.41)	(98.06)
Quarter	2.13**	−0.68*	−0.45***	1.02*	−9.62**	50.63**
(sd)	(0.88)	(0.39)	(0.12)	(0.56)	(3.83)	(25.23)

Panel C: Past return and higher-moment risks — controlling for lagged risk

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	0.94***	−0.07**	−0.23***	1.16***	−11.11**	109.07*
(sd)	(0.34)	(0.03)	(0.06)	(0.39)	(4.53)	(57.91)
Quarter	0.47**	−0.14**	−0.15***	0.25*	−2.04**	11.98**
(sd)	(0.19)	(0.06)	(0.04)	(0.14)	(0.86)	(5.33)

Table 7: **Constant Notional, Volatility-Targeting, and Higher-Moment Risks.** Panel A reports correlations between ex ante variance and tail loss probabilities. The probabilities are $P(r_t^h < -2\bar{\sigma}^h)$ and $P(r_t^h < -2\sigma_t^h)$ where $r_t^h = R_{t,t+h} - E_t[R_{t,t+h}]$, σ_t^h is the ex ante volatility from time t to $t+h$, and we define $\bar{\sigma}^h$ as the time series average of σ_t^h. We report bootstrapped 95% confidence intervals in brackets. Panel B reports regression slope coefficients when regressing physical tail loss probabilities onto the past two year returns. Panel C reports coefficients when controlling for lagged probabilities. We report t-statistics in parentheses and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. We correct standard errors for autocorrelation using Newey and West (1987)

Panel A: Correlations between variance and tail probabilities

Horizon	$P(r_t^h < -2\bar{\sigma}^h)$	$P(r_t^h < -3\bar{\sigma}^h)$	$P(r_t^h < -2\sigma_t^h)$	$P(r_t^h < -3\sigma_t^h)$
Month	0.97	0.98	−0.70	−0.54
95% CI	[0.96,0.98]	[0.96,0.99]	[−0.78, −0.65]	[−0.62,−0.47]
Quarter	0.97	0.94	−0.58	−0.44
95% CI	[0.96,0.98]	[0.92,0.96]	[−0.66,−0.51]	[−0.53,−0.35]

Panel B: Tail probabilities (%) and past return

Horizon	$P(r_t^h < -2\bar{\sigma}^h)$	$P(r_t^h < -3\bar{\sigma}^h)$	$P(r_t^h < -2\sigma_t^h)$	$P(r_t^h < -3\sigma_t^h)$
Month	−5.28*	−2.23	2.02***	0.65***
(sd)	(3.03)	(1.44)	(0.63)	(0.17)
Quarter	−5.67	−2.09	1.24**	0.38**
(sd)	(3.61)	(1.46)	(0.57)	(0.18)

Panel C: Tail probabilities (%) and past return - controlling for lagged probabilities

Horizon	$P(r_t^h < -2\bar{\sigma}^h)$	$P(r_t^h < -3\bar{\sigma}^h)$	$P(r_t^h < -2\sigma_t^h)$	$P(r_t^h < -3\sigma_t^h)$
Month	−1.17*	−0.56*	0.55***	0.26***
(sd)	(0.60)	(0.30)	(0.20)	(0.09)
Quarter	−0.96*	−0.44*	0.38**	0.11**
(sd)	(0.54)	(0.26)	(0.16)	(0.05)

Table 8: **Financial Intermediary Leverage and Higher-Moment Risk.** Panel A reports regression slope coefficients when regressing higher-moment risks onto the financial intermediary leverage of He, Kelly, and Manela (2016). Panel B reports coefficients when conditioning on ex ante variance. Panel C reports coefficients when conditioning on ex ante variance and controlling for lagged risk. We report standard errors in parentheses and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. We correct standard errors for autocorrelation using Newey and West (1987).

Panel A: Leverage and higher-moment risks

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	−0.14	0.11***	0.05	−0.22	1.27	−4.65
(sd)	(0.29)	(0.03)	(0.07)	(0.30)	(3.41)	(28.54)
Quarter	−0.03	0.24**	0.01	−0.00	0.23	0.37
(sd)	(0.51)	(0.10)	(0.06)	(0.39)	(2.00)	(15.64)

Panel B: Conditional leverage and higher-moment risks

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	0.51**	—	−0.07	0.65***	−6.59***	68.43***
(sd)	(0.20)	—	(0.05)	(0.24)	(2.49)	(24.94)
Quarter	0.76***	—	−0.10***	0.51***	−3.28***	22.29***
(sd)	(0.19)	—	(0.03)	(0.08)	(0.80)	(2.60)

Panel C: Conditional leverage and higher-moment risks — controlling for lagged risk

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	0.29***	—	−0.03*	0.38***	−3.93***	45.12***
(sd)	(0.10)	—	(0.02)	(0.13)	(1.35)	(15.42)
Quarter	0.34***	—	−0.05**	0.26***	−1.44***	10.16***
(sd)	(0.10)	—	(0.02)	(0.08)	(0.37)	(2.37)

Table 9: **Market Liquidity, Funding Liquidity, and Higher-Moment Risks.** This table reports the results when regressing higher-moment risks onto market liquidity (Panel A) and funding liquidity (Panel B). We use the value-weighted bid-ask spread of S&P 500 constituents as a proxy for market illiquidity. We use the TED spread as a proxy for funding illiquidity. Panel C and Panel D report results when controlling for lagged market or funding illiquidity respectively. We report t-statistics in parentheses and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. We correct standard errors for autocorrelation using Newey and West (1987).

Panel A: Market illiquidity

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	-1.22^{***}	0.17^{***}	0.27^{***}	-1.53^{***}	14.18^{***}	-128.26^{***}
(sd)	(0.20)	(0.05)	(0.04)	(0.25)	(2.47)	(28.29)
Quarter	-1.44^{***}	0.40^{***}	0.24^{***}	-0.85^{***}	6.30^{***}	-37.05^{***}
(sd)	(0.19)	(0.11)	(0.03)	(0.12)	(0.81)	(5.85)

Panel B: Funding illiquidity

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	-1.41^{***}	0.25^{***}	0.18	-2.12^{***}	14.19^{***}	-128.26^{***}
(sd)	(0.39)	(0.11)	(0.11)	(0.52)	(2.47)	(28.28)
Quarter	-1.37^{***}	0.54^{**}	0.08	-1.07^{***}	6.30^{***}	-37.05^{***}
(sd)	(0.51)	(0.22)	(0.11)	(0.32)	(0.81)	(5.85)

Panel C: Market illiquidity — controlling for lagged risk

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	-0.68^{***}	0.08^{***}	0.14^{***}	-0.78^{***}	8.37^{***}	-83.77^{***}
(sd)	(0.17)	(0.02)	(0.02)	(0.20)	(2.15)	(24.75)
Quarter	-0.55^{***}	0.15^{***}	0.11^{***}	-0.31^{***}	2.51^{***}	-14.96^{***}
(sd)	(0.11)	(0.05)	(0.02)	(0.07)	(0.51)	(3.93)

Panel D: Funding illiquidity — controlling for lagged risk

Horizon	HRI	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	-0.57^{***}	0.10^{***}	0.08^{***}	-0.88^{***}	8.37^{***}	-83.77^{***}
(sd)	(0.18)	(0.02)	(0.03)	(0.23)	(2.15)	(24.75)
Quarter	-0.38^{***}	0.20^{***}	0.05^{*}	-0.31^{***}	2.51^{***}	-14.96^{***}
(sd)	(0.12)	(0.07)	(0.03)	(0.08)	(0.51)	(3.93)

Table 10: **"Bubble" Characteristics and Higher-Moment Risks.** This table reports the results when regressing the higher-moment risk index onto "bubble" characteristics. These are: acceleration, CAPE, dividend-price ratio, cay, turnover, and issuance. We report t-statistics in parentheses and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. We correct standard errors for autocorrelation using Newey and West (1987).

	Dependent variable: Monthly HRI								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Acceleration	2.90								3.86***
(sd)	(1.62)								(3.52)
CAPE		-0.03							-0.01
(sd)		(-0.76)							(-0.39)
Dividend-price ratio			-23.73						162.08
(sd)			(-0.29)						(1.45)
Cay				-64.10***					-45.80***
(sd)				(-3.63)					(-3.12)
Turnover					-3.52***	1.58			-0.00
(sd)					(-2.81)	(1.33)			(-1.07)
Turnover$\times r_{t-24,t}$						9.90***			0.01**
(sd)						(3.94)			(2.11)
Issuance							-2.12	-0.72	6.20**
(sd)							(-0.54)	(-0.18)	(1.97)
Issuance$\times r_{t-24,t}$								-27.27***	-4.97
(sd)								(-3.19)	(-0.67)
$r_{t-24,t}$						0.28		10.77	0.90
(sd)						(0.40)		(4.00)	(0.30)
No. obs.	240	238	240	80	240	240	240	240	80
Adj. R^2	0.02	0.00	-0.00	0.32	0.06	0.18	0.00	0.27	0.46

Appendix A Ex Ante Physical Moments, Risk-Neutral Pricing, and Realized Moments

Ex ante physical moments and risk-neutral pricing

Using equation (6) we can represent physical ex ante moments in terms of asset prices:

$$E_t[R_{t,T}^i] = \frac{E_t^*[R_{t,T}^{i+\gamma}]}{E_t^*[R_{t,T}^{\gamma}]}$$

for $i \in \{1, ..., 6\}$. These asset prices can be used to estimate ex ante physical moments by expanding the standardized moment formula in equation (7). We estimate kurtosis, hyperskewness, and hyperkurtosis in the following way:

$$\text{Kurtosis}_{t,T} = \frac{E_t[R_{t,T}^4] - 3E_t[R_{t,T}]^4 + 6E_t[R_{t,T}]^2 E_t[R_{t,T}^2] - 4E_t[R_{t,T}]E_t[R_{t,T}^3]}{(E_t[R_{t,T}^2] - E_t[R_{t,T}]^2)^2}$$

$$\text{Hyperskewness}_{t,T} = \frac{E_t[R_{t,T}^5] + 4E_t[R_{t,T}]^5 + 10E_t[R_{t,T}]^2 E_t[R_{t,T}^3] - 10E_t[R_{t,T}]^3 E_t[R_{t,T}^2] - 5E_t[R_{t,T}]E_t[R_{t,T}^4]}{(E_t[R_{t,T}^2] - E_t[R_{t,T}]^2)^{5/2}}$$

$$\text{Hyperkurtosis}_{t,T} = \frac{E_t[R_{t,T}^6] - 5E_t[R_{t,T}]^6 + 15E_t[R_{t,T}]^4 E_t[R_{t,T}^2] - 20E_t[R_{t,T}]^3 E_t[R_{t,T}^3] + 15E_t[R_{t,T}]^2 E_t[R_{t,T}^4]}{(E_t[R_{t,T}^2] - E_t[R_{t,T}]^2)^3} - \frac{6E_t[R_{t,T}]E_t[R_{t,T}^5]}{(E_t[R_{t,T}^2] - E_t[R_{t,T}]^2)^3}$$

Estimating ex post realized moments

Let N be the number of daily realized returns between time t and T and denote the daily return between day s and $s+1$ as $r_{s,s+1}$. The realized moments between time t and T are estimated from daily realizations in the following way:

$$\mu_{t,T} = \frac{1}{N}\sum_{i=1}^{N} r_{i-1,i}$$

$$\sigma_{t,T}^2 = \frac{N\sum_{i=1}^{N}(r_{i-1,i} - \mu_{t,T})^2}{N-1}$$

$$\text{Realized Skewness}_{t,T} = \frac{N^{1/2}\sum_{i=1}^{N}\left(r_{i-1,i} - \mu_{t,T}\right)^3}{\sigma_{t,T}^3}$$

$$\text{Realized Kurtosis}_{t,T} = \frac{N\sum_{i=1}^{N}\left(r_{i-1,i} - \mu_{t,T}\right)^4}{\sigma_{t,T}^4}$$

$$\text{Realized Hyperskewness}_{t,T} = \frac{N^{3/2}\sum_{i=1}^{N}\left(r_{i-1,i} - \mu_{t,T}\right)^5}{\sigma_{t,T}^5}$$

$$\text{Realized Hyperkurtosis}_{t,T} = \frac{N^{2}\sum_{i=1}^{N}\left(r_{i-1,i} - \mu_{t,T}\right)^6}{\sigma_{t,T}^6}$$

This is similar to the methods used by Amaya, Christoffersen, Jacobs, and Vasquez (2015).

Appendix B Tables

Table AI: **Ex Ante Conditional Moments Predict Ex Post Realized Moments (Test statistics).** This Table reports test statistics for the results reported in Table 2. Panel A reports p-values from the bootstrapped distribution. Panel B reports bootstrapped standard errors for the correlation coefficient between ex ante moments and ex post realized moments. Panel C reports bootstrapped standard errors for the correlation coefficient between ex ante kurtosis and ex post realized left kurtosis. Panel C reports also reports correlations for hyperkurtosis.

Panel A: Sorting on ex ante monthly moments

	Our moments		Historical moments	
	Low ex ante	High ex ante	Low ex ante	High ex ante
Variance	0.00	0.00	0.00	0.01
Skewness	0.19	0.46	0.54	0.74
Kurtosis	0.21	0.03	0.46	0.20
Hyperskewness	0.02	0.45	0.47	0.75
Hyperkurtosis	0.39	0.01	0.38	0.20

Panel B: Correlation between ex ante moments and ex post realized moments

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Variance	0.07	0.09	0.09	0.10
Skewness	0.06	0.06	0.07	0.06
Kurtosis	0.06	0.07	0.05	0.04
Hyperskewness	0.05	0.06	0.06	0.04
Hyperkurtosis	0.05	0.06	0.04	0.03

Panel C: Left kurtosis and left hyperkurtosis

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Left kurtosis	0.06	0.07	0.07	0.06
Left hyperkurtosis	0.06	0.06	0.07	0.06

Table AII: **Ex Ante Conditional Moments Predict Ex Post Realized Moments — Robustness (Test statistics).** Panel A reports bootstrapped standard errors for the correlations between our ex ante moments and ex post realized moments when we remove observations that overlap with the period from August 1, 2008 to July 31, 2009. Panel B reports bootstrapped standard errors for the correlations between our ex ante moments (estimated with different levels of relative risk aversion) and ex post realized moments. Panel C reports bootstrapped standard errors for the correlations when controlling for risk-neutral moments or historical moments. We report bootstrapped standard errors in the appendix and significance as; * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$..

Panel A: Excluding August 1, 2008 to July 31, 2009

	Our moments		Historical moments	
	Month	Quarter	Month	Quarter
Variance	0.06	0.05	0.07	0.08
Skewness	0.06	0.05	0.07	0.06
Kurtosis	0.06	0.08	0.06	0.07
Hyperskewness	0.05	0.06	0.07	0.07
Hyperkurtosis	0.05	0.07	0.06	0.05

Panel B: Other levels of risk-aversion

	$\gamma = 1$		$\gamma = 5$	
	Month	Quarter	Month	Quarter
Variance	0.07	0.10	0.06	0.09
Skewness	0.06	0.06	0.06	0.06
Kurtosis	0.06	0.06	0.05	0.06
Hyperskewness	0.05	0.05	0.05	0.07
Hyperkurtosis	0.05	0.05	0.05	0.06

Panel C: Marginal correlations

	Controlling for risk-neutral moments		Controlling for historical moments	
	Month	Quarter	Month	Quarter
Variance	0.17	0.10	0.12	0.06
Skewness	0.06	0.06	0.06	0.06
Kurtosis	0.05	0.06	0.06	0.07
Hyperskewness	0.05	0.06	0.05	0.06
Hyperkurtosis	0.04	0.06	0.05	0.06