Higher-Moment Risk*

Niels Joachim Gormsen[†] and Christian Skov Jensen[‡]

September 21, 2021

Abstract

We use methods from Martin (2017) to transform risk-neutral moments of stock returns, calculated from options prices, into physical moments. For the S&P 500 return distribution, even and odd number higher order moments are strongly negatively correlated, creating periods where the return distribution is riskier because it is more left-skewed and fat tailed. The riskiness of the higher order moments varies systematically with the state of the financial markets and peaks when variance is low and prices are high. This time variation has important implications for investors and it is inconsistent with disaster-based models where higher-moment risk peaks when prices are low.

*We are grateful for helpful comments from Adrian Buss (discussant), Peter Christoffersen, Mathieu Fournier (discussant), Robin Greenwood, Sam Hanson, John Heaton (discussant), Jianfeng Hu (discussant), Christian Julliard, Lukas Kremens, David Lando, Ian Martin, Pia Mølgaard, Lasse Heje Pedersen, Davide Tomio, Mads Vestergaard, Grigory Vilkov, Christian Wagner, Paul Whelan as well as from seminar participants at BI Oslo, Bocconi University, Copenhagen Business School, Frankfurt School of Management, Georgetown University, HEC Paris, University of Liechtenstein, London Business School, London School of Economics, Tilburg University, University of Toronto, and Warwick Business School, as well as participants at the Chicago Booth Asset Pricing conference 2018, the European Finance Association 2019, the American Finance Association 2020, and the Canadian Derivatives Conference 2019. We thank the FRIC Center for Financial Frictions (grant no. DNRF102), the European Research Council (ERC grant no. 312417), and the Asness Junior Faculty Fellowship at the University of Chicago Booth School of Business for financial support.

†University of Chicago Booth School of Business. Email: niels.gormsen@chicagobooth.edu.

‡Corresponding author: Christian Skov Jensen. Bocconi University, Italy. Email: christian.jensen@unibocconi.it.

Uncertainty about stock market returns is at the center of asset pricing. Everything else equal, higher uncertainty makes investment less attractive and depresses asset prices.[1] In this regard, investors care not only about the variance of outcomes but also about higher order moments of the distribution of outcomes, and in particular whether very bad outcomes are relatively more likely. To better understand such higher-moment risk, we estimate the higher order moments of the return distribution for the U.S. market portfolio and establish new stylized facts that have broad implications for investors, asset prices, and economic fluctuations more generally.

We base our analysis on ex ante moments that are estimated from options prices. We first estimate the moments of the risk-neutral distribution of stock returns, which reflect a combination of the physical return distribution and the stochastic discount factor. Using methods based on Martin (2017), we translate these risk-neutral moments into physical moments as perceived by an unconstrained power utility investor who invests in the market portfolio. 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.

We first study correlation in higher order moments. We find that the third and the fifth moments (skewness and hyperskewness) are highly positively correlated, and similarly that the fourth and the sixth moments (kurtosis and hyperkurtosis) are highly positively correlated. More importantly, the skewness related moments are negatively correlated with the kurtosis related moments, meaning that there are periods where the return distribution becomes both more left-skewed (due to large negative odd number moments) and more fat tailed (due to large positive even number moments). This comovement in higher order moments is strong as the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis explains 91% of the joint variation in these moments.

[1] Sharpe (1964); Lintner (1965); Black (1972).

The higher order moments vary substantially with the state of the financial markets. In bad times, such as during the financial crisis, there is hardly any higher-moment risk. Excess kurtosis is low and skewness is close to zero, meaning that the return distribution is well approximated by a normal distribution. But going into good times, the shape of the return distribution changes. The distribution becomes left-skewed and fat tailed, meaning additional risk builds in the higher order moments.

We document this result in multiple ways. First, we show that the higher-moment risk is negatively correlated with the variance. When variance decreases, the skewness-related moments become more negative and the kurtosis-related moments increase. Second, higher-moment risk is positively correlated with measures of prices. When prices are high, as measured by a low value of cay (Lettau and Ludvigson, 2001) or high past returns, skewness-related moments are more negative and kurtosis-related moments are more positive.

In the above results, and our analysis in general, we focus on standardized as opposed to raw moments. The advantage of working with standardized moments is that they are informative about the shape of the return distribution. An increase in the fourth standardized moment, for instance, means that the conditional distribution becomes more fat tailed, something that is not necessarily the case for an increase in the raw fourth moment.[2] We note that one disadvantage of the standardized moments is that they do not fully summarize risk without knowledge about the variance, as discussed later.

We explore a series of implications of these new stylized facts about variation in higher-moment risk. First, we emphasize that the variation is inconsistent with consumption-based asset pricing models. Most importantly, the variation in the higher order moments is inconsistent with disaster models that appeal to the shape of the return distribution to explain asset pricing puzzles. Indeed, a higher disaster risk generally makes the return

[2]The raw moments reflects both the shape of the return distribution and the level of variance, meaning an increase in the fourth raw moment could be comletely driven by an increase in variance.

distribution more left-skewed and fat tailed, which is to say that it increases the higher-moment risk. However, in models with time-varying disaster risk, a high disaster-risk is generally associated with *bad times*, as the high disaster risk is the exogenous shock that depresses asset prices and creates bad times. As such, to the extend that price fluctuations are driven by disaster risk, we should expect higher-moment risk to increase in bad times, but we observe the opposite empirically. We make this criqitue precise in simulation studies of the disaster model by Gabaix (2012). In addition, we note that the log-normal models, such as Campbell and Cochrane (1999) and Bansal and Yaron (2004) fail to produce the right higher order moments even unconditionally, as the log-normal distribution imply a right-skewed distribution of returns, whereas we find that the distribution is left-skewed.

The variation in higher-moment risk also has important implications for investors as it influences the tail risk of their portfolios. To analyze the tail risk of different investment strategies, we derive an expression for the probability of a given loss on the market portfolio as perceived by a power utility investor as a function of option prices. This expression allows us to study the tail risk on the market portfolio in real time. The term tail risk may refer to two types of risk: the risk of incurring a large loss, such as 20%, or the risk of incurring a loss that is large relative to the conditional standard deviation of returns. We refer to these types of tail risk as the unconditional and conditional tail risk, and we study both.

We first study unconditional tail risk. We consider the tail risk of losing three or five unconditional standard deviations on the market portfolio. On average, the unconditional tail risk is three times bigger than that implied by the normal distribution, illustrating the importance of higher order moments for tail risk. However, time variation in the unconditional tail risk is largely driven by variance, and not positively correlated with the standardized higher-order moments.

We next study the risk of losing three *conditional* standard deviations. This conditional tail risk is completely driven by higher order moments and would be a constant 0.13% if the returns on the market portfolio followed a normal distribution. The probability of such a tail event varies from from 0.7% to 1.9% over the sample, implying a 271% variation the conditional in tail risk. The tail risk varies systematically over time with the higher-order moments.

The conditional tail risk is particularly relevant for sophisticated investors such as hedge funds, who manage large notionals in constant-volatility portfolios. These portfolios are scaled to have a constant variance over time with the aim of keeping portfolio risk constant. However, the risk in the constant-volatility portfolios varies over time because the higher-moment risk varies. In fact, the tail risk of these portfolios is exactly the same as the conditional tail risk we study. This tail risk peaks when higher-moment risk peaks, which is during calm periods, such as just before the onset of the financial crisis. To the extent that the sophisticated investors who follows volatility targeting strategies have a systematic influence on asset markets,[3] the tendency of these investors to systematically load on tail risk following market run ups during seemingly calm periods may be worrisome.

Given that constant-volatility portfolios are insufficient for inoculating tail risk, we show how to create constant-risk portfolios that has a constant value-at-risk over time. Investors can obtain these portfolios by varying their position in the market portfolio. As such, the constant-risk portfolios are akin to volatility targeting portfolios, only the constant-risk portfolios targets all moments of the return distribution, and not just the variance. We find that the constant-risk portfolio increases its position in the market when variance decreases to offset the decrease in risk. However, it generally increases its position by less than the constant-volatility portfolio does because higher-moment risk on average increase when variance drops, making the market riskier than the variance suggests.

Going beyond the financial markets, higher-moment risk also has potential implications

[3]See He and Krishnamurthy (2013); He, Kelly, and Manela (2017) for evidence on the effect of levered investors on asset prices and the economy.

for the real economy. Indeed, a large literature starting with Bloom (2009) argues that uncertainty about the economy depresses investment because it increases the option value of postponing it: when uncertainty is high, firms prefer to see which state of the economy materializes before they make an investment. Similarly, when higher-moment risk is high, firms may prefer to delay investment until after the risk of a crash diminishes.

We find empirical evidence consistent with the hypothesis that higher-moment risk deters investment. We find empirically that a higher probability of a three-sigma loss predicts lower future industrial production. The effect is stronger the longer the horizon. For the three-year horizon, the probability of a three-sigma event predicts future industrial production with an R^2 of as much as 11%. In addition, the effect of higher-moment risk on industrial production is substantially stronger than the effect of political uncertainty as measured by Baker, Bloom, and Davis (2016), which is only weakly correlated with future industrial production and only at the short horizon. In contrast, higher-moment risk predicts industrial output even more strongly if we skip one year and predict growth in industrial production in year two and three, lending confidence to the idea that higher-moment risk causes low growth, and not the other way around.

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, such as Bollerslev and Todorov (2011) who estimate tail risk using high frequency intraday returns and Kelly and Jiang (2014) who 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. This existing literature studies risk-neutral moments and probabilities

(Siriwardane (2015), Gao, Gao, and Song (2018), Bates (2000), and Schneider and Trojani (2017)), whereas we study physical moments and probabilities. Finally, our paper relates to the literature on the recovery theorem by Ross (2015), which decomposes investor preferences and beliefs directly from observable asset prices.[4] Unlike Ross, we we fix preferences and estimate beliefs directly. In addition, we only recover higher-order moments, not the full return distribution, and thereby avoid a series of challenges with numerical approximations that arise in this regard.

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

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

[4] Ross's results have been sparked a large literature that discusses the empirical validity of the theorem and extended it along multiple dimensions. See e.g. Walden (2017), Jensen, Lando, and Pedersen (2019), Borovicka, Hansen, and Scheinkman (2015), Schneider and Trojani (2019), Jackwerth and Menner (2019), Bakshi, Chabi-Yo, and Gao (2018), Pazarbasi, Schneider, and Vilkov (2019), Martin and Ross (2019), Carr and Yu (2012), and Ghosh and Roussellet (2019).

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 $X_{t,T}(n)$:

$$E_t[R_{t,T}^n] = E_t[\underbrace{R_{t,T}^n m_{t,T}^{-1}}_{X_{t,T}(n)} m_{t,T}] = E_t^*[\underbrace{R_{t,T}^n m_{t,T}^{-1}}_{X_{t,T}(n)}]/R_{t,T}^f \tag{3}$$

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 $R_{t,T}^n$ 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 $R_{t,T}^{-\gamma}$:

$$m_{t,T} = kR_{t,T}^{-\gamma} \tag{4}$$

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 $X_{t,T}(0) = m_{t,T}^{-1}$ and the standard asset pricing formula (1) then implies the relation:

$$E_t^*[m_{t,T}^{-1}] = R_{t,T}^f \tag{5}$$

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:

$$E_t[R_{t,T}^n] = \frac{E_t^*[R_{t,T}^n \overbrace{R_{t,T}^\gamma/k}^{m_{t,T}^{-1}}]}{E_t^*[\underbrace{R_{t,T}^\gamma/k}_{m_{t,T}^{-1}}]} = \frac{E_t^*[R_{t,T}^{n+\gamma}]}{E_t^*[R_{t,T}^\gamma]} \tag{6}$$

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:

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

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 $n = 3$:

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

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:

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

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

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

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

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

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 2017. 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 360 calender days, and apply standard filters, excluding options with implied volatility higher than 100%.

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

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

with

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

where $F_{t,T}$ is the forward price and $\mathrm{call}_{t,T}(K)$ and $\mathrm{put}_{t,T}(K)$ 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. To see how we approach this numerical approximation, let $F_{t,T}$ be the forward price and, then 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.

During our sample period, the number of options traded with different strikes has increased tremendously. In the beginning of our sample (1996), we observe option prices with strikes ranging from the spot price plus/minus three risk-neutral standard deviations. At the end of our sample (2017), we observe option prices with strikes as far as sixteen risk-neutral standard deviations into the tails of the distribution. It is important that the increase in the number of strikes with observable prices does not drive our results. Therefore, when estimating higher order moments, we exclude all prices that are outside the range of the spot price plus/minus three risk-neutral standard deviations. The data we use when estimating higher order moments is therefore comparable over time. In untabulated results, we redo the analysis using all available option prices. The empirical results are largely similar to those reported in this paper but there is a clear trend in the estimated higher order moments which arise due to this increase in traded options. Also, Andersen and Bondarenko (2007) note that model-free implied volatilities can be hard to measure accurately because of sparse coverage of option prices with strikes in the tails of the distribution. This concern is also valid when we estimate higher order moments. Therefore, as robustness, we redo our analysis where we linearly extrapolate the slope for the implied volatilities. Doing so changes the levels of the higher order moments slightly, but the results are similar to those reported in the paper. In conclusion, the choice of cutting observable option prices at the spot plus/minus three risk-neutral standard deviations when estimating option moments is not important for our results.

2.2 Estimating Market Tail Probabilities

The main challenge when implementing Proposition 1 is that we have 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 the best linear prediction of observed implied volatilities (IV) onto their option strikes. Specifically, in the spirit of Dumas, Fleming, and Whaley (1998), we regress observed implied volatilities onto their respective strikes:

$$IV_i = \beta_0 + \beta_1 \text{Strike}_i + \eta_i \tag{14}$$

When fitting the entire implied volatility smile at a single horizon, Dumas, Fleming, and Whaley (1998) suggest using a second degree polynomial in strikes. We choose to use a first order polynomial since we do not fit the entire volatility smile, rather we restrict the regression to the range of observable prices from the price with the lowest strike which is no more than five risk-neutral standard deviations lower than the spot price to the price with the strike that is closest to $\alpha S_t - D_{t,T}$. [5] The first order polynomial approximation is reasonable as long as α is less than one.

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:

$$\text{put}'_{t,T}(\alpha S_t - D_{t,T}) = \frac{\text{put}_{t,T}(\alpha S_t - D_{t,T} + h) - \text{put}_{t,T}(\alpha S_t - D_{t,T} - h)}{2h} \tag{15}$$

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

We evaluate the integral in Proposition 1 using estimated prices implied from equation (14). Specifically, we use the fitted implied volatilities to obtain option prices ranging from 0.1 times the spot price to the strike $\alpha S_t - D_{t,T}$. If 0.1 times the spot price is lower than the lowest observable option strike then we linearly extrapolate the fitted implied volatilities to obtain these option prices.

[5]IWe use a wider grid of observable option prices when estimating tail probabilities compared to when we estimate the moments. The reason to extend observations beyond the spot minus three risk-neutral standard deviations as the follows. When we estimate tail probabilities, we rely on the information solely in the tail of the distribution. Specifically, the integral in Proposition 1 requires observations in the deep left tail. When we run the regression in Equation (14), it is therefore natural to extend the grid of observable options further into the tails. On the other hand, when we estimate moments, we rely on information from the entire distribution. The resulting moments are therefore much less affected by the truncation of option data as we discussed in the previous section.

Inserting (15) and the evaluated integral into Proposition 1, we can estimate physical probabilities.

2.3 Determining Market Risk-Aversion

To estimate physical moments and tail probabilities, we have to choose a level of risk-aversion. Following Bliss and Panigirtzoglou (2004), we use the Berkowitz (2001) test to estimate the constant level of risk aversion that best reconciles the cross-section of options prices with the time-series realizations of the S&P 500 index.[6] Doing so, we find that a $\gamma = 3$ is the best fit to monthly returns in our sample, and we therefore rely of this level of risk aversion in our main analysis. Our estimate is close to the $\gamma = 4$ that Bliss and Panigirtzoglou (2004) find best explains the 1983 to 2001 sample.

Tabel 1 reports moment summary statistics for moments estimated with risk aversion of 0 (risk-neutral investor), 1 (log-utility), 3 (our baseline case), and 5. The average ex ante estimated skewness is negative at both monthly and quarterly 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. In addition, average kurtosis is larger than 3 for both horizons and all levels of risk aversion, suggesting that the physical distributions are leptokurtic (i.e. tails are on average fatter

[6] The idea behind the Berkowitz test is that, for the true value γ, the distribution of $u_{t,T} = F_{t,T}(R_{t,T})$ is uniform and the distribution $y_{t,T} = \Phi^{-1}(u_{t,T})$ is standard normal. Here $F_{t,T}(R_{t,T})$ denotes the distribution function

$$F_{t,T}(r) = \int_{-\infty}^{r} \frac{\pi_{t,T}(x)x^{\gamma}}{\int_{-\infty}^{\infty} \pi_{t,T}(y)y^{\gamma}dy}dx$$

where $\pi_{t,T}$ is the state price density (Arrow-Debreu prices) and $\int_{-\infty}^{\infty} \pi_{t,T}(y)y^{\gamma}dy$ is simply a normalization constant which ensures that the physical return distribution integrates to one. In the Berkowitz test, we estimate the coefficients in the regression model:

$$y_{t,T} = \hat{a} + \hat{\beta} y_{t-1,T-1} + \epsilon_{t,T}, \qquad \epsilon_{t,T} \sim N(0, \hat{\sigma})$$

and perform a likelihood ratio test of the joint hypothesis that $a = \beta = 0$ and $\mathrm{Var}(\epsilon_{t,T}) = 1$.

than the normal distribution). We report in later analysis that our results are not sensitive to the choice of $\gamma = 3$ as our baseline case.

3 New Facts About Higher-Moment Risk

We start the empirical analysis by documenting new facts about the higher order moments of the return distribution for the S&P 500. Section 3.1 studies how the moments relate to on another over time and Section 3.2 studies how higher order moments relate to key state variables of financial markets.

3.1 Time Variation and Comovement in Higher Order Moments

Figure 1 and Figure 2 plot the time series of the skewness, kurtosis, hyperskewness, and hyperkurtosis of the return distribution at the monthly and quarterly horizon. The skewness moments trend down during the sample and the kurtosis moments trend up, suggesting that higher-moment risk increases during the sample; at the quarterly horizon in particular, hyperkurtosis and hyperskewness exhibit a substantial drop/rise during the end of the sample. The results presented throughout the paper are not driven by this latest rise in higher-moment risk – in fact, our results are generally stronger if we exclude the last two years.

On average, the moments deviate from the normal distribution. The skewness is on average negative, and kurtosis is on average well above three, suggesting that the distrubtion is on average left-skewed and fat tailed. These results add to the evidence from the well-known histograms of historical returns. These histograms show that the unconditional distribution of returns is left-skewed and fat tailed, but leave open the possibility that this shape is driven by time-variation in the mean and variance. Our results rule that out by showing that, on most days, the *conditional* distribution is left-skewed and fat tailed.

Table 2 reports the monthly (Panel A) and quarterly (Panel B) pairwise correlations between the moments. The green (lower right) box shows pairwise correlations between higher order moments. Since investors are averse to lower skewness and hyperskewness (i.e., averse to a more left-skewed distribution), we have flipped the signs of these moments such that we can interpret a higher value of the third and fifth moments as higher risk. The correlations between the higher order moments are all positive and large, suggesting that the risk in the individual higher order moments tend to move together.

The strong comovement between 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 3. At both the monthly and quarterly horizon, the first principal component explains more than 90% of the joint variation in higher order moments, underlining the strong comovement in higher-moment risks.

The first principal component eigenvector loads negatively on skewness and hyperskewness and positively on 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 -0.47 (-0.49 quarterly) for skewness, 0.51 (0.50 quarterly) for kurtosis, -0.52 (-0.51 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. We interpret a high value of the first principal component as times when higher-moment risk is high (the higher order moments are on average riskier). Accordingly, we 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.*

3.2 Cyclicality in Higher-Moment Risk

Figure 1 and Figure 2 indicate that higher order moments vary with the state of the financial markets. During the crisis in 2007-2009, the skewness approaches zero and the kurtosis approaches three, meaning that the deviation from the normal distribution is limited. However, in the period leading up to the crisis, 2003 until 2007, skewness drops substantially and kurtosis increases, meaning the distribution deviates substantially from the normal distribution. In particular, the higher order moments become riskier in that the probability mass in the (conditional) left tail increases.[7]

To study such cyclical variation in higher-moment risk, we relate the variation in higher-moment risk to key state variables of financial markets. We first consider the relation between higher-moment risk and the variance. Figure 3 displays the time-series plot of the monthly and quarterly HRIs along with the variances. Both for the monthly and the quarterly horizon, the two times series are negatively correlated: higher-moment risk is high when variance is low. For instance, HRI is low during the burst of the tech bubble and during the financial crisis, during which variance is somewhat high. Conversely, HRI is high during the low variance period from 2004 to 2007 and the low variance period from 2012 an onwards. The higher-moment risk index spikes during 2017 where volatility is historically low.

This negative relation between variance and higher-moment risk is highly statistically significant as shown in Panel A of Table 4. On a monthly horizon, the magnitude of the correlation between variance and HRI is -0.47 with 95% block bootstrapped confidence bounds of $[-0.54, -0.41]$. On the quarterly horizon, the correlation is -0.60 with 95% confidence bounds of $[-0.65, -0.55]$. All the higher order moments contribute to the negative

[7]By left tail, we refer to the conditional left tail, i.e. the left tail measured relative to the conditional standard deviation. This is different than the unconditional tail.

correlation between HRI and variance. Indeed, the blue (upper right) box of Panel A of Table 2 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 -0.48 to -0.40. The results are in line with the results of Bollerslev, Todorov, and Lai (2015), who investigate market tail risk premia in a parametric jump-diffusion model. Within their parametric framework, they find that the left-tail jump intensity increases during low variance periods.

We next study how higher-moment risk relates to past returns. Figure 4 shows time-series plots of the past two year return and the HRI. The past returns and the HRI are positively correlated with correlations of 0.35 on the monthly horizon and 0.26 on the quarterly horizon. Panel A of Table 4 reports the block bootstrapped 95% confidence bounds for the correlations between the HRI and past returns. The confidence bounds are [0.24, 0.45] on the monthly horizon and [0.15, 0.36] on the quarterly horizon. Panel B of Table 4 shows the pairwise correlations between past returns and the individual physical moments. Each of the higher order moments become riskier when past returns are high.

Finally, we study how higher-moment risk relates to the measure of consumption to aggregate wealth, cay, from Lettau and Ludvigson (2001). A high value of cay suggests that prices, or wealth, is relatively low, and that the expected returns are high. As shown in Panel A of Table 4, the higher-moment risk index is highly negatively correlated with the cay variable, with tight confidence intervals. Consistent with the findings on past returns, higher-moment risk appears elevated when prices are higher. Panel C shows that this finding holds for each of the individual moments on their own. We note, however, that higher-moment risk is not particularly correlated with the dividend-price ratio in this sample, mainly due to timing issues during the financial crisis.

Taken together, the results in this section show that higher-moment risk are elevated when variance is low and prices are high. The next sections explore the implications of this finding for asset pricing models, investors, and the economy. Before doing so, we address potential robustenss concerns.

3.3 Robustness

We conduct several robustness analyses to solidify our results on the systematic variation in higher-moment risk. First, a potential concern is that option liquidity is low during periods of financial market turbulence which could be reflected in large option bid-ask spreads during such periods. Such fluctuations in bid-ask spreads could pose a problem when we use mid prices to estimate our higher-moment risk index. Figure 5 shows the higher-moment risk index when using exclusively bid, mid, or ask prices. As seen from the figure, the higher-moment risk index is almost identical when using either bid, mid, or ask prices. The pairwise correlations between the higher-moment risk indexes implied by bid or ask prices with the mid price higher-moment risk index are 0.978 and 0.981 respectively, suggesting that our results are not driven by time-varying bid-ask spreads on option prices.

A second concern is that our results are sensitive to our choice of a constant risk-aversion coefficient equal to three ($\gamma = 3$). Table 8 shows the pairwise correlations between risk-neutral higher order moments. Interestingly, the comovements between risk-neutral higher order moments are similar to those between physical higher order moments. That is, the green (lower right) box of Table 8 shows the pairwise correlations between the negative of (risk neutral) skewness, kurtosis, the negative of hyperskewness, and hyperkurtosis. These correlations are all positive and close to one, suggesting that also under the risk-neutral measure there is a common component in the higher order moments. Table 9 shows the principal components of the space spanned by *risk-neutral* skewness, kurtosis, hyperskewness, and hyperkurtosis. The first principal component explains 95% of the joint variation in risk-neutral higher order moments. The correlation between this first principal component and the first principal component of the space spanned by *physical* skewness, kurtosis, hyperskewness, and hyperkurtosis is 0.97.

3.4 Expected and Realized Moments

We next confirm that the moments we have extracted from option prices indeed are meaningful measures of the actual moments of the return distribution. We do so by showing that the option implied moments predict future realized moments, and that they do so better than trailing realized moments.

To validate a given moment, such as skewness, we create two investment strategies. The first strategy invests in the market during periods where the ex ante skewness is above the median. The second strategy invests in the market during periods where the skewness is below the median. Both strategies keep the ex ante volatility constant by scaling their position during the period by the ex ante volatility. We then estimate the skewness of the excess return on each of the strategies and test if the skewness is higher for the high-skewness strategy.

The motivation for our approach is as follows. First, we cannot use daily returns to estimate the monthly or quarterly skewness, because there is no natural way to aggregate daily skewness to the monthly or quarterly horizon (Neuberger, 2012). Instead, to estimate skewness, we need a set of monthly or quarterly return observations, which we obtain by dividing the sample into two strategies based on the median moment. Second, even within each strategy, the variance might vary substantially between periods, making it difficult to compare returns across periods and therefore difficult to estimate higher order moments. We circumvent this problem by considering constant-volatility strategies.[8]

Panel A of Table 5 shows results where we use the ex ante volatility from option prices to standardize the portfolios. The first column shows the results for sorting on skewness. The skewness of the standardized returns in the "Low" bucket is -0.56 and the skewness in the

[8] We note that our results are similar, but less significant, if we do not scale by volatility.

"High" bucket is -.44. Accordingly, the realized skewness is 0.12 higher in periods where the ex ante skewness is higher. However, the block bootstrapped p-value of the null-hypothesis that High−Low is less than zero is 0.34.

We find more significant, but otherwise similar, results for the other moments. Across the board, our moments predict realized moments with the right sign. The effect is statistically significant for the kurtosis moments on the monthly horizon.

For robustness, we redo our analysis using a parametric estimate of expected variance to standardize the portfolios. We estimate monthly and quarterly variance based on daily returns and fit an AR(1) model on these estimates, and use the predicted values from this regression as the expected variance. The results reported in Panel B are similar to before but the statistical significance is substantially stronger.

Finally, we also ask if trailing realized moments are useful for predicting future moments. Here, we estimate the trailing moments from realized daily returns in a given period. We then undertake the same exercise as before, only sorting on realized moments in the past period instead of the expected moments. We again standardize by the expected variance from the AR(1) regression. The results in Panel C are weak. The trailing realized moments predict future moments with the right sign only five out of eight times, and the effect is generally statistically insignificant. The difficulty of predicting future moments using past realized moments underline the strength of our expected moments.

4 Implications for Consumption-Based Models

In this section, we explore whether the moments of the return distribution produced by leading asset pricing models are consistent with the facts we documented in the previous section.

4.1 Disaster Models

Disaster models appeal to the shape of the return distribution to generate a high and time-varying risk premium. In particular, the models generate the counter-cyclical equity risk premium through counter-cyclical disaster risk (i.e., disaster risk increases in bad times). We argue, however, that this feature is inconsistent with our new facts about higher-moment risk. In particular, higher-moment risk is lower during bad times, which is hard to reconcile with a hightened disaster risk.

To make the critique concrete, we focus on the Gabaix (2012) model. In this model, bad times are periods where the resilience towards disasters is low, which we will refer to as a high disaster risk. The high disaster risk generates a higher equity risk premium, which depresses prices. The disaster risk also changes the return distribution by increasing the variance and by making it more left-skewed and fat tailed. These dynamics are inconsistent with the results in the previous section. Indeed, the dynamics counter-factually imply that bad times are periods with elevated higher-moment risk. More generally, the model implies that higher-moment risk is positively correlated with variance and negatively correlated with price levels, both of which appears inconsistent with the data.

We formalize the critique in simulation studies of the model. Figure 6 shows the results from the simulation studies along with empirical results. Panel A shows relation between skewness and volatility. In red, we show the option-implied skewness in the model for different values of option-implied volatility. The figure shows a strong negative correlation between the two, meaning the distribution is more left-skewed when volatility is higher. This finding is inconsistent with the data, shown in blue, in which there is a positive correlation between volatility and skewness. Similarly, Panel B shows that kurtosis and volatility are positively correlated in the model, which again is inconsistent with the empirical data shown in blue. In the figure, we compare risk-neutral moments from the model with risk-neutral moments in the data to avoid any issues of risk corrections. Simulation details are in the Appendix.

Panel C and D show the relation between higher order moments and valuation ratios in the model. The model again gets the correlation wrong. As shown in Panel C, skewness is positively correlated with valuation ratios, meaning the distribution is more left-skewed when prices are lower. This is inconsistent with the data, plotted in blue. Similarly, Panel D shows that kurtosis and prices are possitivly correlated in the model, which again is inconsistent with the data shown in blue.

It is important to emphasize that the disaster model fails to reproduce the stylized facts about standardized moments, not raw moments. Raw higher order moments indeed becomes riskier in bad times, both empirically and in the disaster model. However, empirically, this happens mainly because volatility increases in bad times, not because the shape of the return distribution becomes riskier. To the extend that disaster models appeal to a time-varying shape of the return distribution to explain asset pricing movements, and not simply a time-varying second moment, this ability to only reproduce the raw moments is less of a victory.

In conclusion, our empirical facts are conceptually challenging for distaster models. The models rely on disaster risk being high in bad times, which is hard to reconcile with the fact that higher-moment risk goes down in bad times.

4.2 Log-Normal Models

Log-normal asset pricing models, such as Campbell and Cochrane (1999) and Bansal and Yaron (2004) are generally inconsistent with the higher-moment risk we document in the paper. In both models, the risk-neutral distribution is log-normal, which means that it is right skewed. On every single day in the sample, we estimate a left-skewed risk-neutral distribution, which is a direct rejection of log-normal distributions. These results are anolagous to the results in Martin (2017), who document that VIX is higher than SVIX in every month in the sample, which is inconsistent with log-normal distributions.

5 Implications for Investors

We next analyze the implications of our new stylized facts about higher-moment risk for investors. We first study how higher-moment risk influences the probability of a tail event on the market portfolio. We define the tail risk as the probability that the realized market return is three or five standard deviations below the expected value. We consider tail risk probabilities relative to both the conditional and unconditional standard deviation of the market return. Variation in the tail risk relative to the conditional standard deviation is entirely driven by higher order moments because it is measured relative to the conditional expected return and variance (i.e., the first two moments). One the other hand, the tail risk relative to the unconditional standard deviation is also influenced by the conditional variance, as it is measured relative to the conditional expected return and the unconditional variance. Throughout the section, we estimate probabilities using Proposition 1.

Figure 7 Panel A plots the time series of the probability that the market return is three *conditional* standard deviations below the expected returns. The solid blue line is the tail risk observed from option prices. This tail risk peaked on October 31$^{\text{th}}$ 2017 with a probability of 1.9%, more than twice the size of its low on May 30$^{\text{th}}$ 1997, where the probability was 0.7%. The probability was also high before the financial crisis.

The horizontal dotted line is the probability implied by a normal distribution. The line is flat, reflecting that the risk of a three-sigma event is constant for a normal distribution. The gray area between the two lines is higher moment risk: it is the additional risk of a tail event that arises from the higher order moments of the return distribution. This probability of a tail loss in excess of what is implied by a normal distribution ranges from 0.6% to 1.8%, showing that higher order moments constitute most of the tail risk for investors and drives the time variation.

Panel B of Figure 7 plots the time series of the probability that the market return is three *unconditional* standard deviations below the expected returns. Three unconditional standard deviations is approximately 15%. The solid blue line is the tail risk observed from option prices and the dotted line is the probability implied by a normal distribution. The shaded area between the two lines is again higher moment risk. In many parts of the sample, higher order moments constitute the majority of the probability of the 15% percent drop However, most of the variation in tail risk comes from variation in the conditional volatility.

Figure 9 shows similar results for the probability of a five-sigma loss on the market portfolio. The probability that realized returns are more than five conditional standard deviation below expected return is shown in Panel A. It peaked on October 31st 2017 with a probability of 0.5%, which is thirty times the size of its low on November 28th 2008, where the probability was 0.015%. These probabilities are well above what is implied by a normal distribution, which is $2.9 \times 10^{-5}\%$.

Figure 8 and Figure 10 shows the same probabilities at the quarterly rather than monthly horizon. The results are similar, except that the importance of higher-moment risk for tail loss probabilities is larger than at the monthly horizon.

To make the analysis more tangible and to ease the interpretation, we relate the tail risk probabilities to two different investors. The tail risk relative to the unconditional standard deviation can be thought of as the risk of a certain percentage loss for a constant notional investor who keeps the same weight in the market over time. In contrast, the tail risk relative to the conditional standard deviation can be thought of as the probability of a certain loss for a volatility targeting investor who varies the position in the market over time to keep a constant portfolio volatility (and invests the remainder in the risk-free asset). Such a volatility targeting strategy is common practice among hedge funds and often used in academic research as well (Moskowitz, Ooi, and Pedersen (2012), Asness, Frazzini, and Pedersen (2012), Moreira and Muir (2018), and Moreira and Muir (2017)).

To make the tail risk of the two types of investors more precise, consider the unexpected return from an investment in the return on the market and the risk-free asset,

$$r_{t,T}^{\text{i, shock}} = \omega_{t,T}^{i}(R_{t,T} - E_t[R_{t,T}])$$

where $\omega_{t,T}^{i}$ is the conditional portfolio weight on the market portfolio (and $1 - \omega_{t,T}^{i}$ is the weight in the risk free asset). The conditional portfolio weight for the constant notional investors is constant. For convenience and without loss of generality, we assume that the weight is 1. The portfolio weight of the volatility targeting investor is

$$\omega_{t,T}^{\text{vol target}} = \frac{\sigma^{\text{target}}}{\sigma_{t,T}}$$

where $\sigma_{t,T}$ is the ex-ante conditional volatility on the market. For convenience, we assume that the target volatility (σ^{target}) is the same as the time series average of the market portfolio, $\bar{\sigma}$. The probability of an unexpected loss which is lower than α for either investor is

$$\begin{aligned} P_t(r_{t,T}^{\text{i, shock}} < \alpha) &= P_t\left(\omega_{t,T}^{\text{i}}(R_{t,T} - E_t[R_{t,T}]) < \alpha\right) \\ &= P_t\left(R_{t,T} - E_t[R_{t,T}] < \frac{\alpha}{\omega_{t,T}^{\text{i}}}\right) \end{aligned}$$

If we, for instance, study the probability of realizing an unexpected loss that is more than three times the unconditional standard deviation of the market (i.e. $\alpha = 3\bar{\sigma}$), then the probabilities for the constant notional and volatility targeting investors are

$$P_t(r_{t,T}^{\text{constant, shock}} < -3\bar{\sigma})) = P_t\left(R_{t,T} - E_t[R_{t,T}] < -3\bar{\sigma}\right))$$

for the constant notional investor and

$$P_t(r_{t,T}^{\text{target, shock}} < -3\bar{\sigma}) = P_t\left(R_{t,T} - E_t[R_{t,T}] < -3\sigma_{t,T}\right)$$

for the volatility targeting investor. Accordingly, for the constant notional investor, the probability is the same as the probability that the return to the market is three unconditional standard deviation below expectations. For the volatility targeting investor, the probability is the same as the probability that the return to the market is three conditional standard deviation below expectations. We thus interpret the results on the unconditional probabilities as the tail loss of a constant notional investor and the results on the conditional probabilities as the tail loss of the volatility targeting investor.

5.1 Cyclical Variation in Tail-Loss Probabilities

We next study how different investors tail probabilities relate to the financial state variables studied in Section 3.2. We first consider the conditional variance. Panel A of Table 6 reports correlations between log-variance and the probability of a portfolio return that is less than the threshold α. The first two columns of Panel A shows the correlations between log-variance and the probability that the market realizes an unexpected return less than $-3\bar{\sigma}$ and $-5\bar{\sigma}$. The correlations range from 0.65 to 0.84 with tight block bootstrapped confidence bounds, showing that the conditional variance plays a large role in this probability.

More importantly, Panel A of Table 6 also reports correlations between log-variance and the probability of a portfolio return of $3\sigma_{t,T}$ and $5\sigma_{t,T}$, which reflect the portfolio risk of the volatility targeting investor. These probabilities are negatively correlated with log-variance at -0.81 to -0.47, with tight block bootstrapped confidence bounds. These negative correlations show that the portfolio of the volatility targeting investor is exposed to substantially more risk at times when variance is low. This finding can potentially help explain why Moreira and Muir (2018) and Moreira and Muir (2017) 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 risk.

We next investigate the relation between tail probabilities and past returns. Specifically, we regress tail probabilities onto past two year returns in the following regression

$$P_t(r^{\text{shock}}_{t,T} < -3\sigma_{t,T}) = \beta_0 + \beta_1 r_{t-24,t} + \epsilon_{t,T} \tag{16}$$

Panel B of Table 6 reports β_1 coefficients from regressions such as in (16). We find that the probability of both a $-3\sigma_{t,T}$ and a $-5\sigma_{t,T}$ 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 0.18 percentage point higher probability of a monthly $-3\sigma_{t,T}$ drop in the market.

Finally, Panel C of Table 6 reports slope coefficients from similar regressions of the tail probabilities onto the ex ante value of cay. The conditional tail probabilites are negatively related to cay, suggesting that the conditional tail risk is higher when cay is low (that is, prices are high). Unconditional tail probabilities are on the other hand positively related to cay. The slope coefficients are all statistically significant.

5.2 Value-at-Risk Targeting

Given that volatility targeting does not inoculate tail risk, it is natural to ask what portfolio might. If investors want the tail risk to be constant over time, how should they invest?

We answer this question by solving for the portfolio weights that create a constant value-at-risk. Formally, for a desired value-at-risk, say c, we determine the fraction of wealth invested in the market, $\omega^{\text{constant risk}}_{t,T}$, such that there is a constant probability, p, for a portfolio shock less than c:[9]

$$P_t(\omega_{t,T}^{\text{constant risk}}(R_{t,T} - E_t[R_{t,T}]) < c) = p$$

Panel A of Figure 11 shows time-variation in the tail risk of such a constant-risk portfolio. The portfolio depicted targets a 1.31% probability of a 15% loss.[10] The tail risk is constant by construction at 1.31%. We compare this tail risk to that of a volatility targeting investor who targets a 5% monthly volatility. The tail risk of the volatility targeting investor varies from 0.7% to 1.9% over our sample which is a variation of 271%. Panel A of Figure 12 shows similar results for a 25% loss on the quarterly horizon. Here, we assume the volatility targeting investor targets a 8.33% volatility and the value-at-risk targeting investor targets a constant probability, 1.17%, for the shock. The tail risk probability of the volatility targeting investor varies from 0.5% to 2.2% over our sample period, equivalent to a 440% variation in tail risk.

Panel B of Figure 11 shows the weight the constant-risk investor puts in the market portfolio over time. It generally decreases in bad times, such as during the financial crisis, as the variance increases thereby making a 15% loss more likely for any given weight. In this sense, the constant-risk portfolio trades in the same direction as the constant-volatility investor, who also decrease her weight in the market when variance goes up.

Importantly, however, the constant-risk investor decrease her weight in the market portfolio by less than the volatility targeting investor when volatility goes up. To illutstratate this finding, Panel C of Figure 11 shows the ratios of wealth invested in the market for the two types of investors:

$$\frac{\omega_{t,T}^{\text{constant risk}}}{\omega_{t,T}^{\text{vol target}}} - 1$$

[9]See Appendix C for details on how we estimate portfolio weights.

[10]The 1.31% is the average tail risk of a volatility targeting investors who targets 5% monthly volatility.

When volatility is high, e.g. during the financial crisis, the constant-risk investor has 13% more wealth in the market than the volatility targeting investor. When volatility is low, we find that the constant-risk investor has 13% less invested into the market relative to the volatility targeting investor. Panel C of Figure 12 shows similar results on a quarterly horizon.

The constant-risk investor trades less than the constant-volatility investor. When volatility increases, both investors decrease their positions in the market as second-moment risk goes up, but the constant-risk investor decreases the position by less, as the higher-moment risk goes down simultaneously. The same dynamics occur when variance decrease. Over our sample, the difference in trading is substantial. On the quarterly horizon, for example, the constant-risk investor trades 10% less than the constant-volatility investor.[11] This lower trading activity suggests that the constant-risk portfolio is easier and cheaper to implement than the constant-volatility strategy.

6 Economic Implications: Higher-Moment Risk, Uncertainty, and Business Cycles

Time variation in higher-moment risk may also have important implications for the real economy. Indeed, the growing literature on uncertainty shocks argues that uncertainty about the economy may drive business cycles because higher uncertainty deters firms from investing. This may happen through multiple mechanisms, but often the option value of postponing investments is stressed. Indeed, higher uncertainty about a project increases the option

[11]Estimated as:

$$\frac{\sum_{t=2}^{264} |\omega_{t,T}^{\text{constant risk}} - \omega_{t-1,T-1}^{\text{constant risk}}|}{\sum_{t=2}^{264} |\omega_{t,T}^{\text{vol target}} - \omega_{t-1,T-1}^{\text{vol target}}|} - 1$$

over the 264 monthly observations from Jan. 1996 until Dec. 2017.

value of investment projects, and managers may thus prefer to hold on to that option rather than undertaking the investment (Dixit, Dixit, and Pindyck, 1994). Similarly, more higher-moment risk increases the option value of an investment project and therefore increases the incentive to postpone investment (Dixit, Dixit, and Pindyck, 1994).[12] We therefore test whether more higher-moment risk predicts lower future industrial production.

At each time t, we regress the change in industrial production between period t and $t+i$ onto the ex ante probability of a three-sigma loss:

$$\mathrm{IND}_{t+i} - \mathrm{IND}_t = \beta_0 + \beta_1 P_t(r_{t,T}^{\text{shock}} < -3\sigma_{t,T}) + \epsilon_{t,T}$$

The results are presented in Table 7 for the three year horizon. The estimate is negative and statistically significant when measuring tail risk at the quarterly horizon. Accordingly, a high tail risk predicts lower future industrial production. For shorter horizons the coefficients are negative but insignificant. However, we note that the lack of predictability on the one and two year horizons could arise from the fact that it takes time for firms to adjust their investments, and it takes additional time for their adjustments to ripple through the economy. Finally, we also find that political uncertainty does not predict changes in industrial production.

In the above regressions, one may be concerned that causality runs the opposite way, and that it is recessions that cause uncertainty and not the other way around (Berger, Dew-Becker, and Giglio, 2020). For instance, in the financial crisis volatility spiked in October 2008, which was probably after the great recession had set in motion, meaning that the uncertainty was unlikely to have started the recession. To mitigate such concerns, we next run similar regressions where we skip the first year of industrial production, meaning we

[12] Merton (1976) shows that, keeping the first moment constant, higher tail risk increases the value of a call option.

forecast the change in production from one year ahead and onwards:

$$\text{IND}_{t+i} - \text{IND}_{t+12} = \beta_0 + \beta_1 P_t(r_{t,T}^{\text{shock}} < -3\sigma_{t,T}) + \epsilon_{t,T} \tag{17}$$

These results from regression (17) are presented in columns (4) and (5) of Table 7. Tail risk still predicts changes in industrial production negatively at the three year horizon. The results are substantially stronger than in Panel A, with R^2 as high as 14%, suggesting that tail risk may play an important role in business cycle fluctuations. In constrast, political uncertainty predicts future indstrial production positively, not negatively, questioning its importance in understanding business cycle fluctuations. The strong relation between tail risk and industrial production is depcited in Figure 13.

Finally, as a word of caution, the results of the predictive regressions in Table 7 are based on a small sample and the inference might thus be subject to small sample issues. For instance, the t-statistics based on Newey-West errors are going to be biased upwards. In untabulated results, we redo the analysis using standard errors of Lazarus, Lewis, and Stock (2017), that are robust to small samples, and find that our results remain significant.

7 Conclusion

We employ a new method to estimate forward looking higher order moments of the return distribution for the S&P 500 index from option prices.

Using this method, we document new stylized facts about higher-moment risk: (1) The riskiness of higher order moments comove strongly in the sense that the third to sixth moment tend to become riskier at the same time; (2) the comovement is strong as 91% of the comovements in higher order moments is explained by their first principal component; (3) higher-moment risk is high when the second moment, variance, is low; and (4) higher-moment risk is high when past returns are high and cay is low.

The movement in higher-moment risk is difficult to align with consumption-based asset pricing models, and disaster models in particular. Indeed, in disaster models such as Gabaix (2012), disaster risk peaks in bad times, which runs counter to the fact that higher-moment risk is low during bad times. If anything, disaster risk appears particularly elevated in good times.

The movement in higher-moment risk also has important implications for investors. The variation in higher-moment risk causes the risk of a three-sigma event to vary from 0.7% to 1.9% over the sample. This probability is higher in periods of low variance, creating additional risk for investors who lever their portfolios when variance is low, such as volatility targeting investors. We show how volatility targeting investors can take this additional higher-moment risk into account when creating portfolios by targeting a particular value-at-risk instead of a particular level of volatility. Going beyond investors, financial regulators should worry that volatility targeting investors, such as hedge funds, load up on tail risk in calm periods, thereby exposing themselves and the financial system to more risk. These results echo the findings of (Brunnermeier and Sannikov, 2014), who shows theoretically that the economy and the financial system becomes more vulnerable when volatility is low because sophisticated investors apply more leverage and creates endogenous risk.

Finally, higher-moment risk is also related to future industrial output. When higher-moment risk is high, firms produce less during the subsequent three years. The correlation between future industrial output and higher-moment risk is as high as -0.37, suggesting a potentially important role for higher-moment risk in understanding firm investments. In addition, higher-moment risk predicts industrial production better than political uncertainty, suggesting that higher order moments of returns are relevant measures for understanding the impact of uncertainty and uncertainty shocks on the business cycle.

References

Andersen, Torben G., and Oleg Bondarenko, 2007, Construction and interpretation of model-free implied volatility, *Working paper*.

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

Baker, Scott R., Nicholas A. Bloom, and Steven J. Davis, 2016, Measuring economic policy uncertainty, *The Quarterly Journal of Economics* 131, 1593–1636.

Bakshi, Gurdip, Fousseni Chabi-Yo, and Xiaohui Gao, 2018, A recovery that we can trust? deducing and testing the restrictions of the recovery theorem, *The Review of Financial Studies* 31, 532–555.

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*.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

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

Berger, David, Ian Dew-Becker, and Stefano Giglio, 2020, Uncertainty shocks as second-moment news shocks, Working paper, .

Berkowitz, Jeremy, 2001, Testing density forecats, with applications to risk management, *Journal of Business and Economic Statistics* 19, 465–474.

Black, Fischer, 1972, Capital market equilibrium with restricted borrowing, *Journal of Business* pp. 444–455.

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

Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

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

Bollerslev, Tim, Viktor Todorov, and Xu. Lai, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–134.

Borovicka, Jaroslav, Lars Peter Hansen, and Jose A. Scheinkman, 2015, Misspecified recovery, *Journal of Finance* forthcomming.

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

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Carr, Peter, and Jiming Yu, 2012, Risk, return, and Ross recovery, *Working paper.*

Dixit, Avinash K, Robert K Dixit, and Robert S Pindyck, 1994, *Investment under uncertainty.* (Princeton university press,).

Dumas, Bernard, Jeff Fleming, and Robert E. Whaley, 1998, Implied volatility functions: Empirical tests, *Journal of Finance* 53, 2059–2106.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gao, George, Pengjie Gao, and Zhaogang Song, 2018, Do hedge funds exploit rare disaster concerns?, *Review of Financial Studies* 31, 2650—-2692.

Ghosh, Anisha, and Guillaume Roussellet, 2019, Identifying beliefs from asset prices, *Working paper*.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economic* 126, 1—-35.

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

Jackwerth, Jens C., and Marco Menner, 2019, Does the Ross recovery theorem work empirically?, *Journal of Financial Economics, forthcoming*.

Jensen, Christian Skov, David Lando, and Lasse Heje Pedersen, 2019, Generalized recovery, *Journal of Financial Economics* 133, 154–174.

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

Lazarus, Eben, Daniel J Lewis, and James H Stock, 2017, The size-power tradeoff in har inference, Working paper, Harvard University Mimeo.

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

Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *The review of economics and statistics* pp. 13–37.

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

Martin, Ian, and Steve Ross, 2019, Notes on the yield curve, *Journal of Financial Economics* 134, 689–702.

Merton, Robert C, 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

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

Moreira, Alan, and Tyler Muir, 2018, Should long-term investors time volatility?, *Journal of Financial Economics forthcoming.*

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.

Newey, Whitney K., and Kenneth D. West, 1994, Automatic lag selection in covariance matrix estimation, *Review of Economic Studies* 61, 631–653.

Pazarbasi, Altan, Paul Schneider, and Gregory Vilkov, 2019, Sentimental recovery, *Working paper*.

Ross, Steve, 2015, The recovery theorem, *The Journal of Finance* 70, 615—648.

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

Schneider, Paul, and Fabio Trojani, 2019, (almost) model-free recovery, *Journal of Finance* 74, 323–370.

Sharpe, William F, 1964, Capital asset prices: A theory of market equilibrium under conditions of risk*, *The journal of finance* 19, 425–442.

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

Walden, Johan, 2017, Recovery with unbounded diffusion processes, *Review of Finance* 21, 1403–1444.

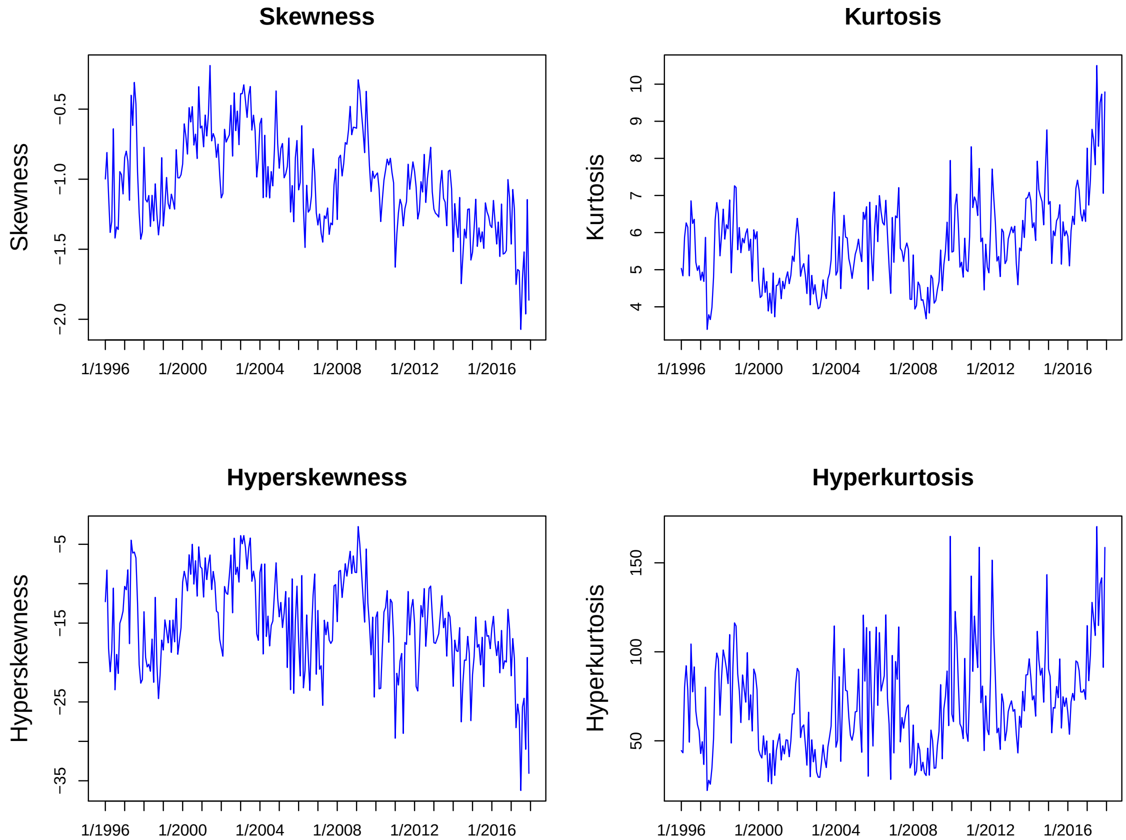

Figure 1: **Higher Order Moments of the Return Distribution of the S&P 500 Index (Monthly Horizon).** The figures show time series plots of monthly higher order moments for the S&P 500 index.

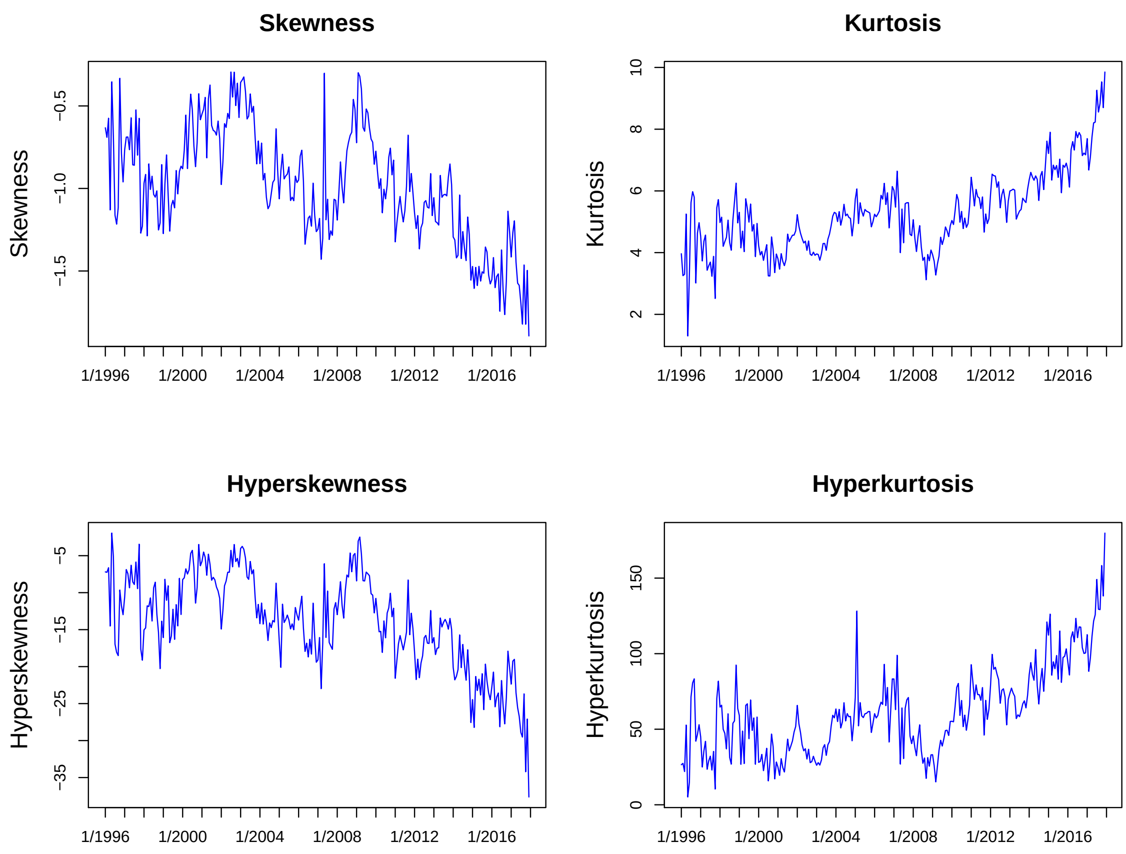

Figure 2: **Higher Order Moments of the Return Distribution of the S&P 500 Index (Quarterly Horizon).** The figures show time series plots of quarterly higher order moments for the S&P 500 index.

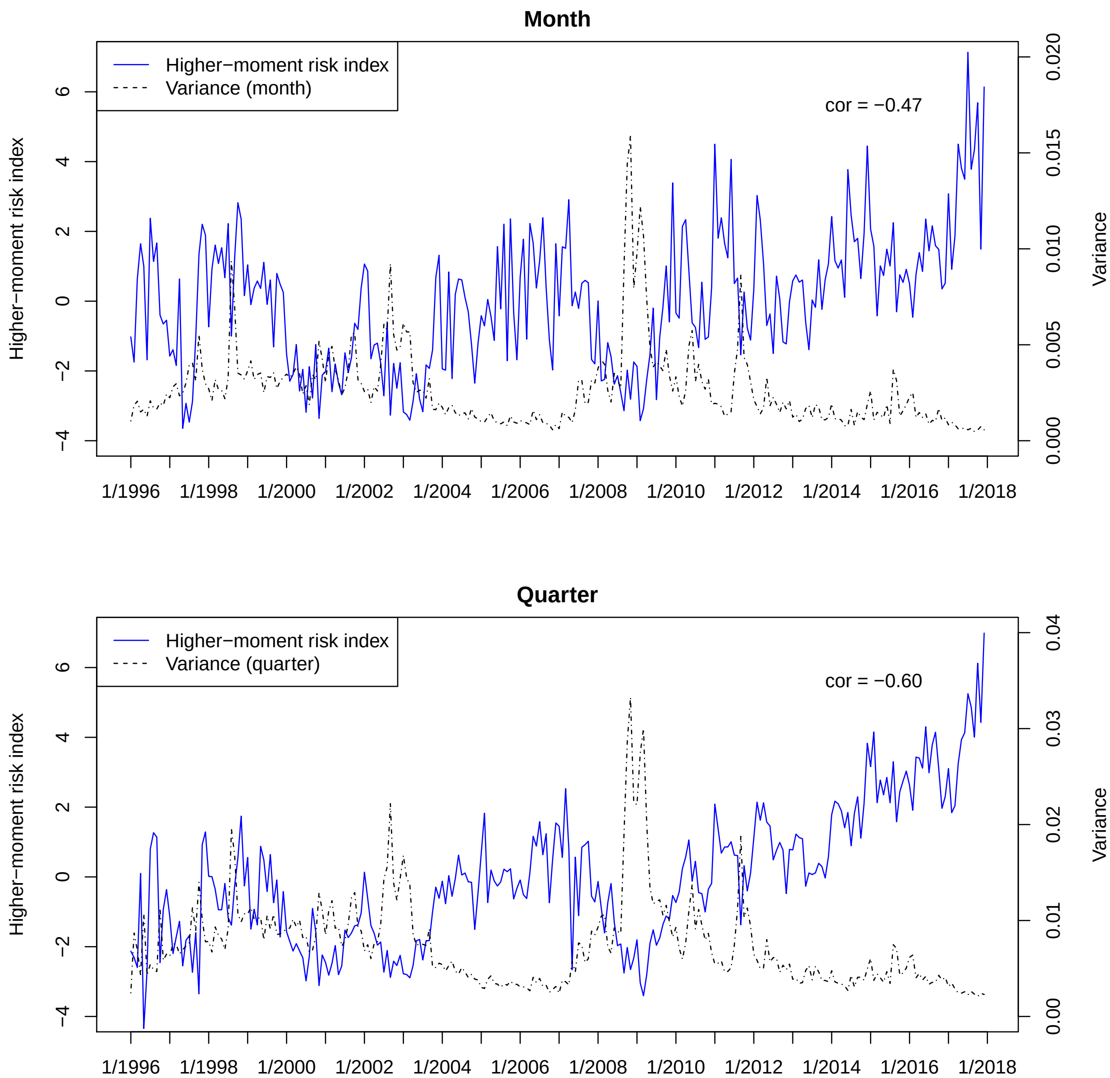

Figure 3: **The Higher-Moment Risk Index and Variance.** This figure shows time series plots of the monthly and quarterly higher-moment risk index and the second moment (variance) of the return distribution. The higher-moment risk index is the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis. The higher the index the more higher-moment risk.

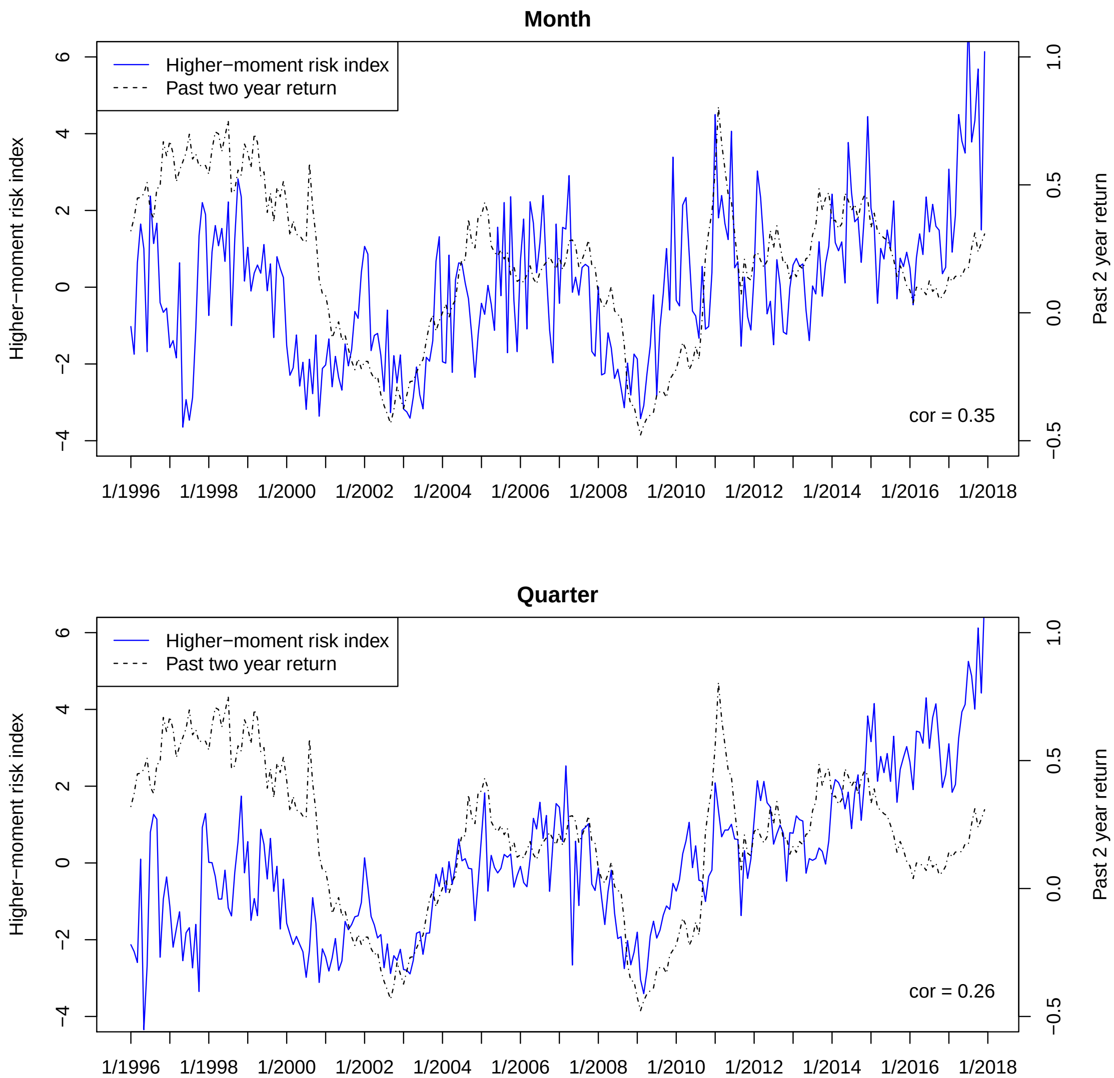

Figure 4: **The Higher-Moment Risk Index and the Past Two Year Returns.** This figure shows time series plots of the monthly and quarterly higher-moment risk index and the past two-year return on the S&P 500 Index. The higher-moment risk index is the first principal component of the space spanned by skewness, kurtosis, hyperskewness, and hyperkurtosis. The higher the index the more higher-moment risk.

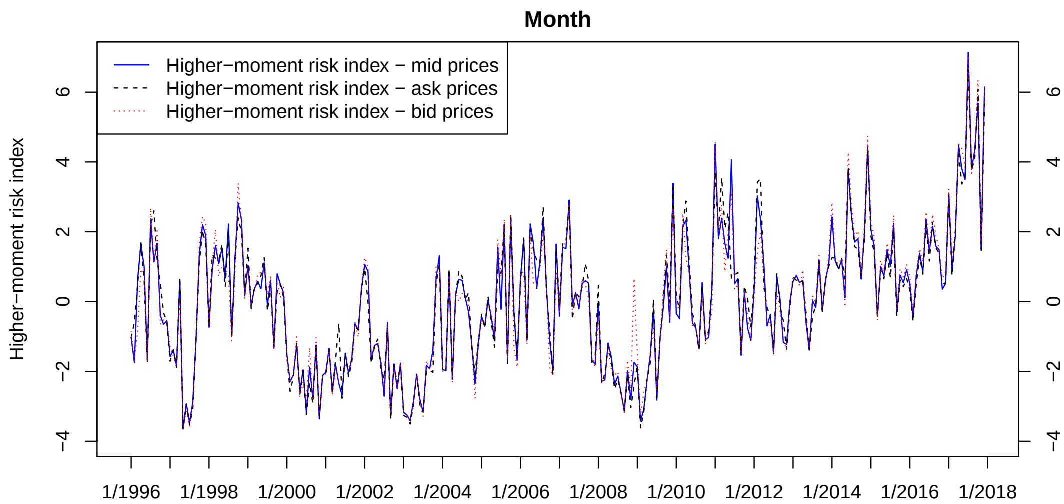

Figure 5: **The Higher-Moment Risk Index Using Bid or Ask Prices.** This figure shows the monthly horizon higher-moment risk indexes when using either bid, ask, or mid prices. The pairwise correlation between the time series range from 0.978 to 0.981, suggesting that our choice of mid prices does not affect our results.

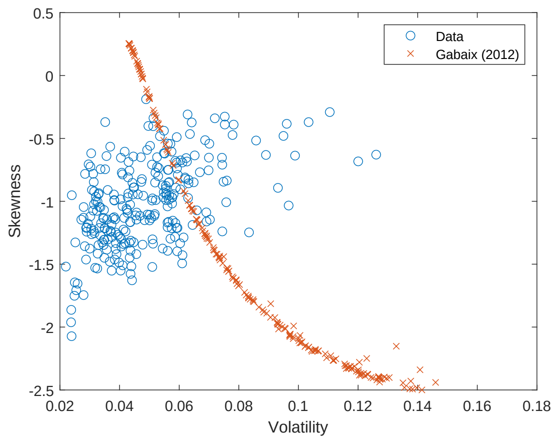

(a)

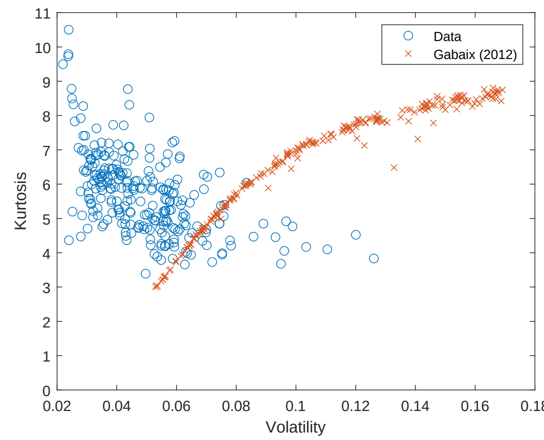

(b)

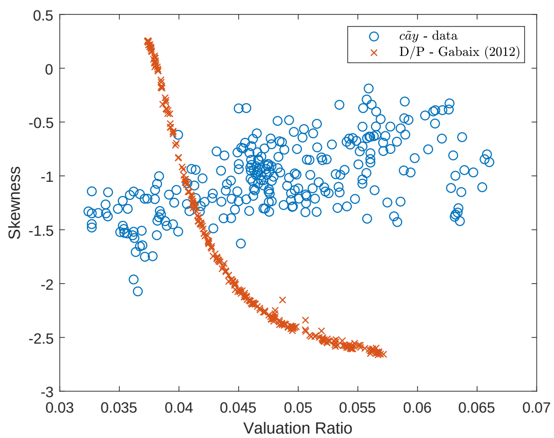

(c)

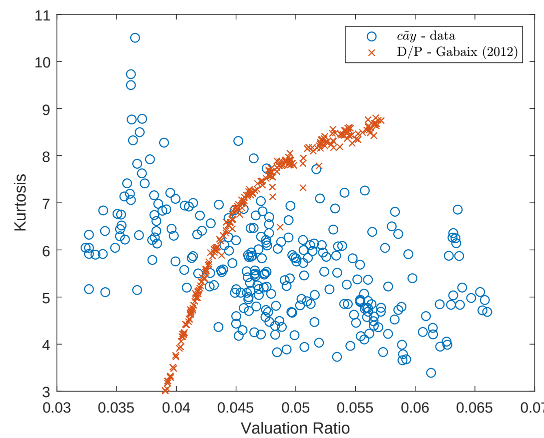

(d)

Figure 6: **Consumption-Based Asset Pricing Models and Higher-Moment Risk.** Panel (a) shows a scatterplot of simulated stock market skewness and volatility in Gabaix (2012) along with a scatterplot of skewness and volatility in the data. Panel (b) shows a scatterplot of simulated stock market kurtosis and volatility in Gabaix (2012) along with a scatterplot of kurtosis and volatility in the data. Panel (c) shows a scatterplot of simulated stock market skewness and the dividend-price ratio in Gabaix (2012) along with a scatterplot of skewness and cay from the data. Panel (d) shows a scatterplot of simulated stock market kurtosis and the dividend-price ratio in Gabaix (2012) along with a scatterplot of kurtosis and cay from the data. To set cay and the dividend price ratio on the same scale, we transform the observed cay as: $\tilde{\text{cay}} = (\text{cay} - 0.1)/2$. Simulation details of the Gabaix (2012) model are in Appendix B.

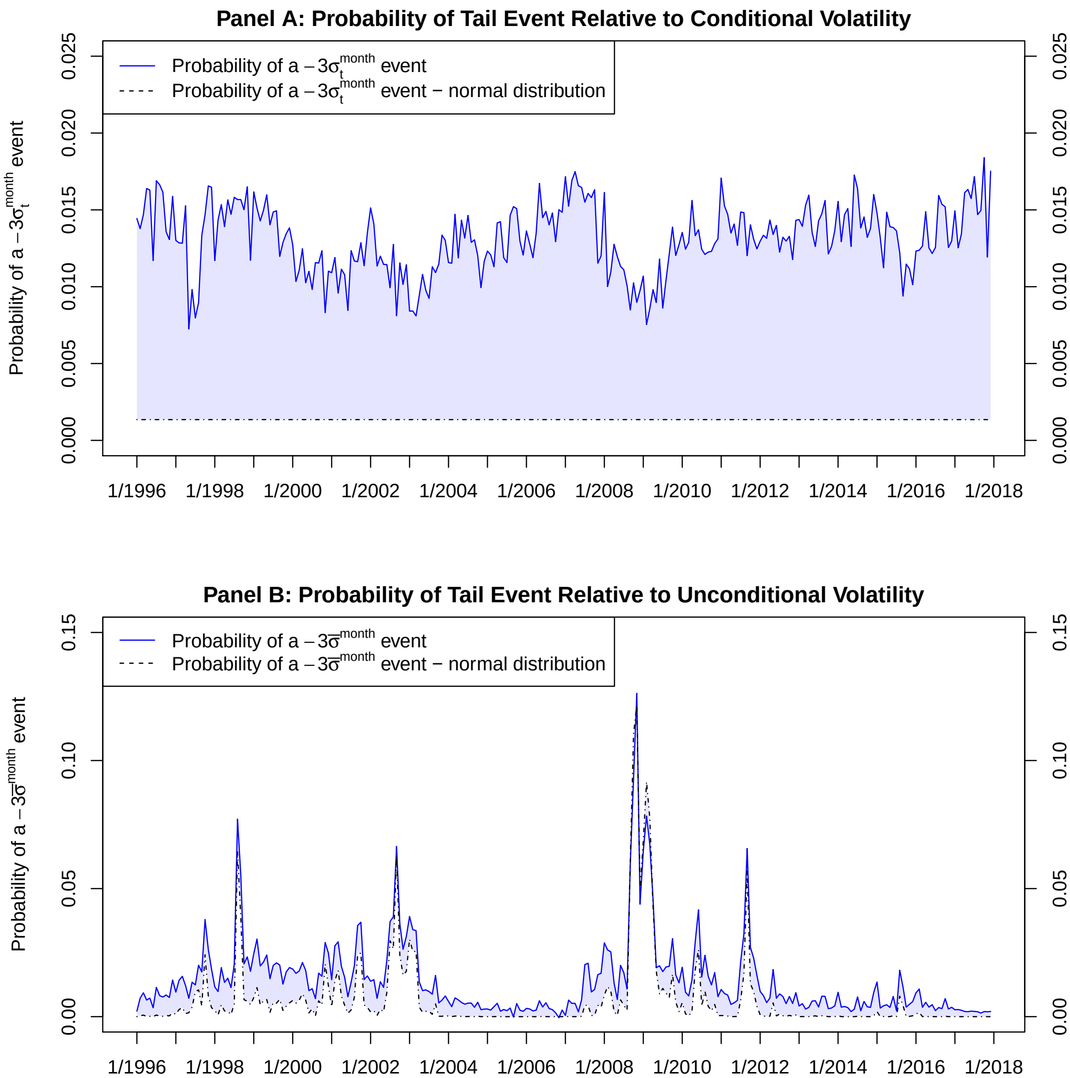

Figure 7: **Probabilities of a Three-Sigma Loss on the S&P 500 Index (Monthly Horizon).** Panel A shows the probability that the monthly S&P 500 return is three conditional standard deviations ($3\sigma_t^{\text{month}}$) below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution, which is a constant 0.13%. The shaded area between the lines represents higher-moment risk. Panel B shows the probability that the monthly S&P 500 return is three unconditional standard deviations below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution. The unconditional standard deviation is the time series average of 5.0% per month. The shaded area between the lines represents higher-moment risk.

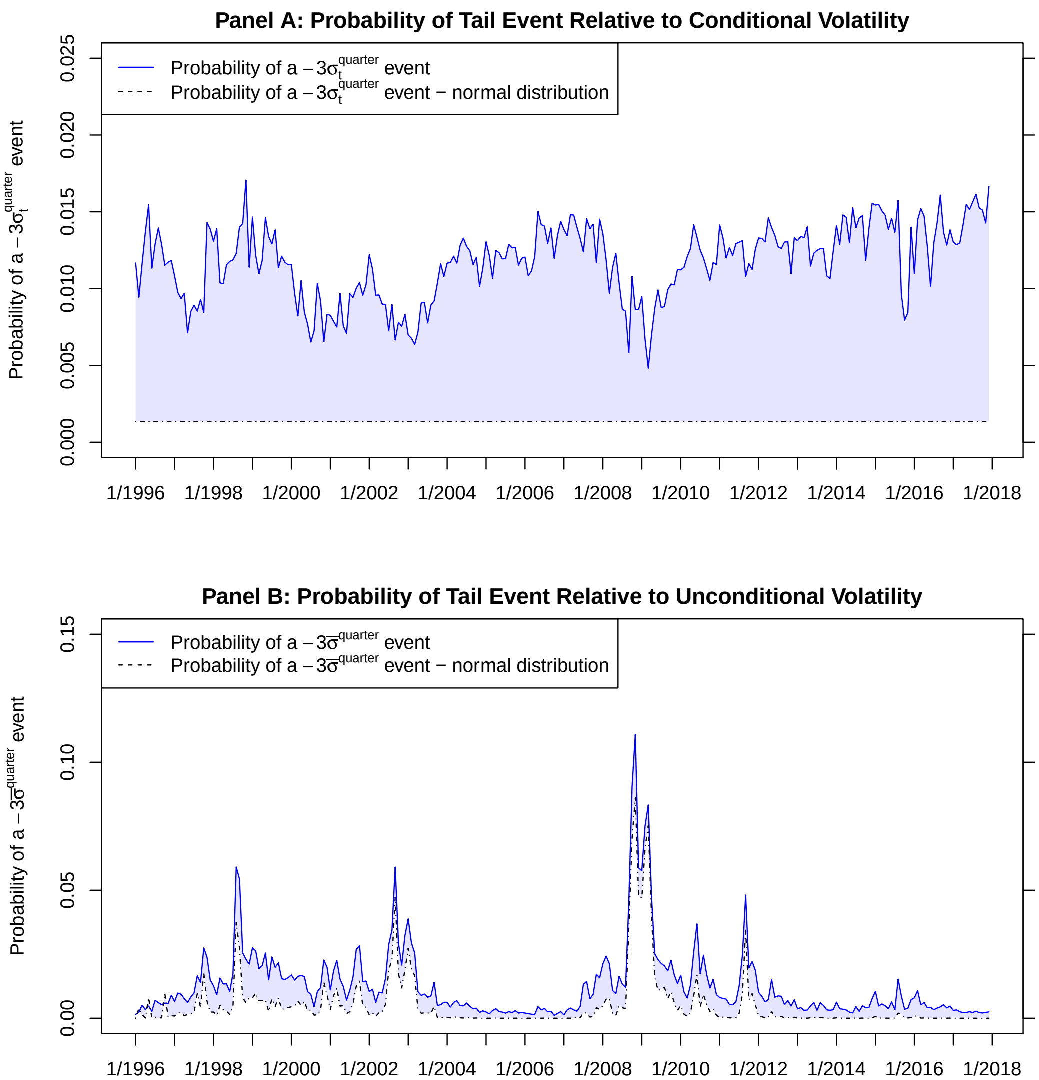

Figure 8: **Probabilities of a Three-Sigma Loss on the S&P 500 Index (Quarterly Horizon).** Panel A shows the probability that the quarterly S&P 500 return is three conditional standard deviations ($3\sigma_t^{\text{month}}$) below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution, which is a constant 0.13%. The shaded area between the lines represents higher-moment risk. Panel B shows the probability that the quarterly S&P 500 return is three unconditional standard deviations below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution. The unconditional standard deviation is the time series average. The shaded area between the lines represents higher-moment risk.

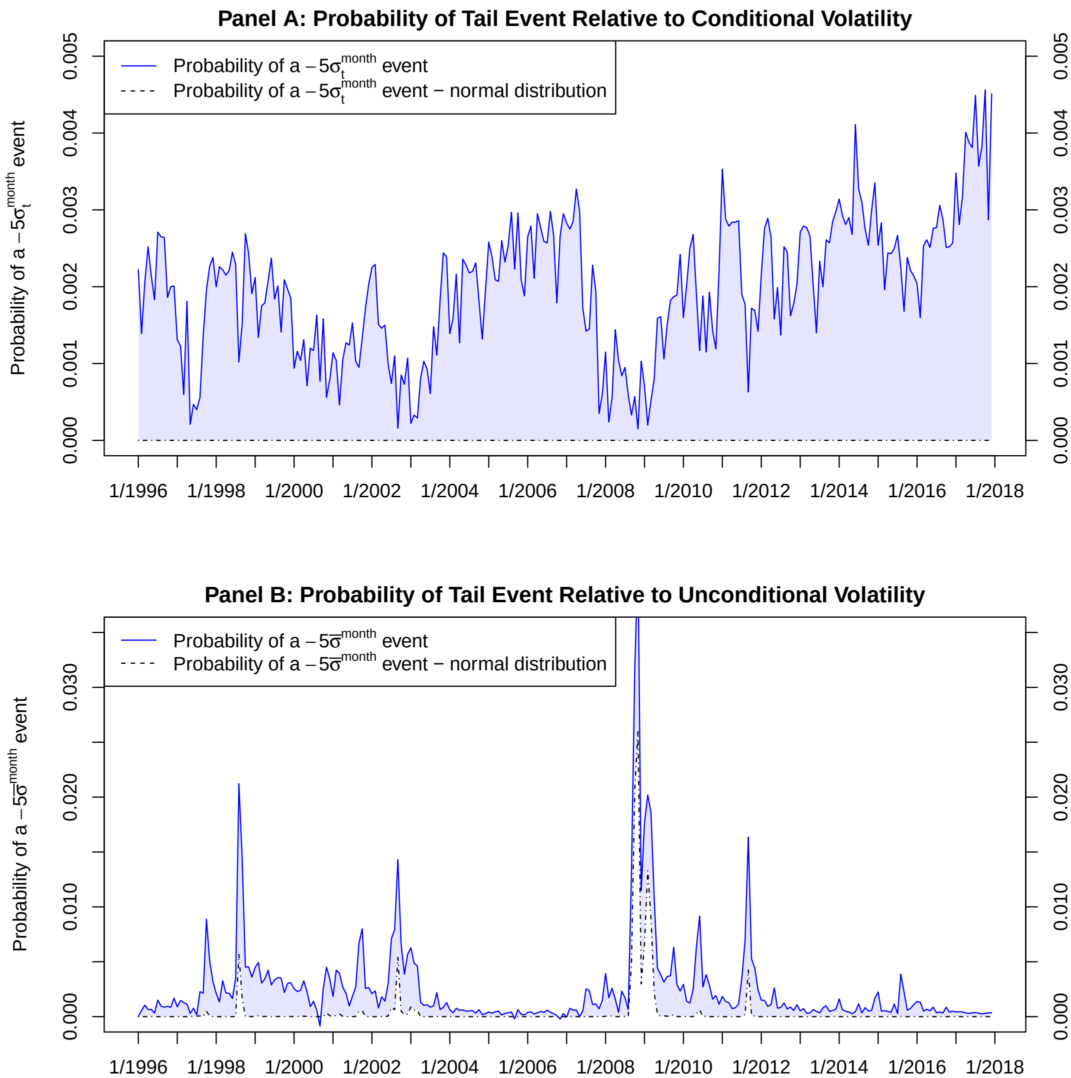

Figure 9: **Probabilities of a Five-Sigma Loss on the S&P 500 Index (Monthly Horizon).** Panel A shows the probability that the monthly S&P 500 return is five conditional standard deviations ($5\sigma_t^{\text{month}}$) below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution, which is a constant $2.9 \times 10^{-5}\%$. The shaded area between the lines represents higher-moment risk. Panel B shows the probability that the monthly S&P 500 return is five unconditional standard deviations below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution. The unconditional standard deviation is the time series average of 5.0% per month. The shaded area between the lines represents higher-moment risk.

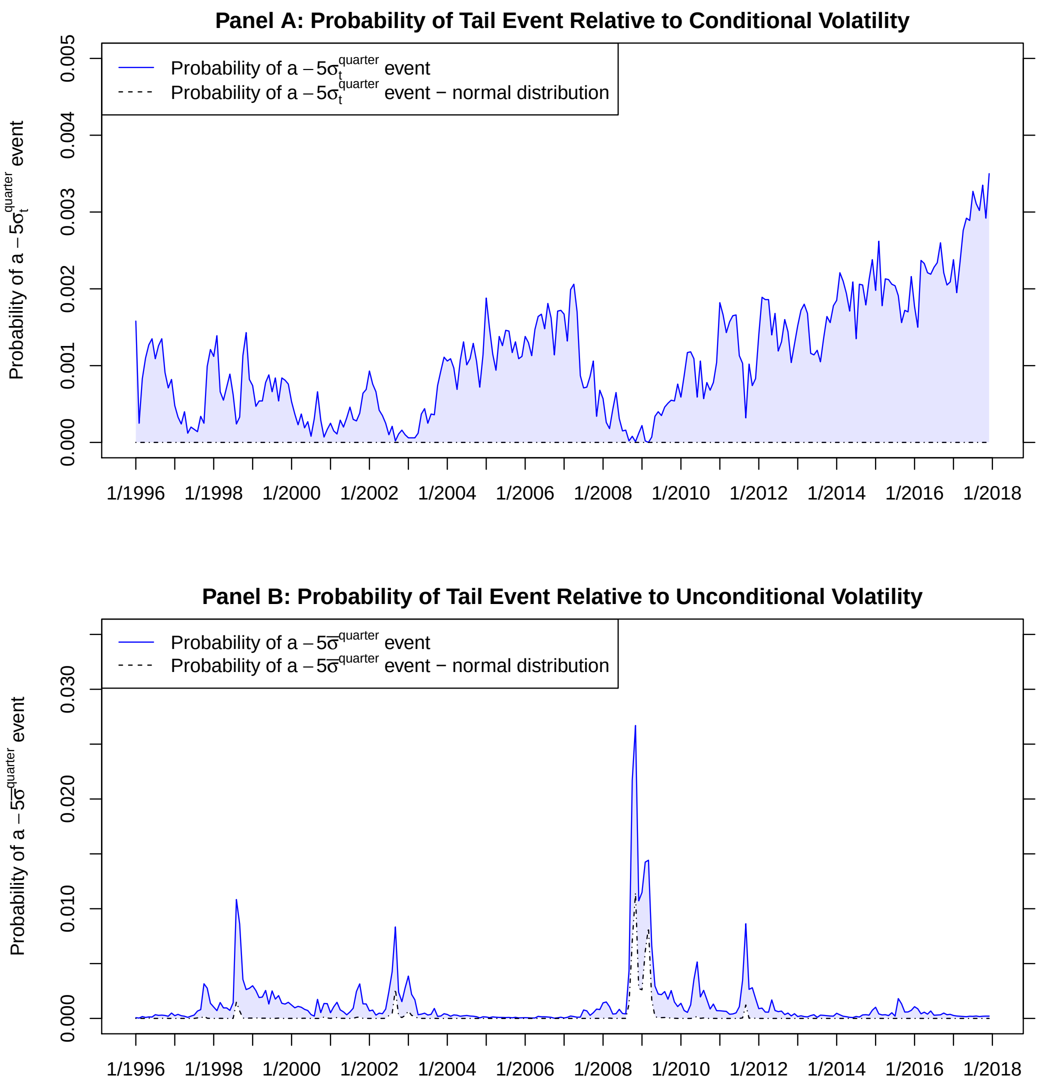

Figure 10: **Probabilities of a Five-Sigma Loss on the S&P 500 Index (Quarterly Horizon).** Panel A shows the probability that the quarterly S&P 500 return is five conditional standard deviations ($5\sigma_t^{\text{month}}$) below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution, which is a constant $2.9 \times 10^{-5}\%$. The shaded area between the lines represents higher-moment risk. Panel B shows the probability that the quarterly S&P 500 return is five unconditional standard deviations below its expected value. The solid blue line shows the probability as derived from option prices and the dotted line shows the probability implied by the normal distribution. The unconditional standard deviation is the time series average. The shaded area between the lines represents higher-moment risk.

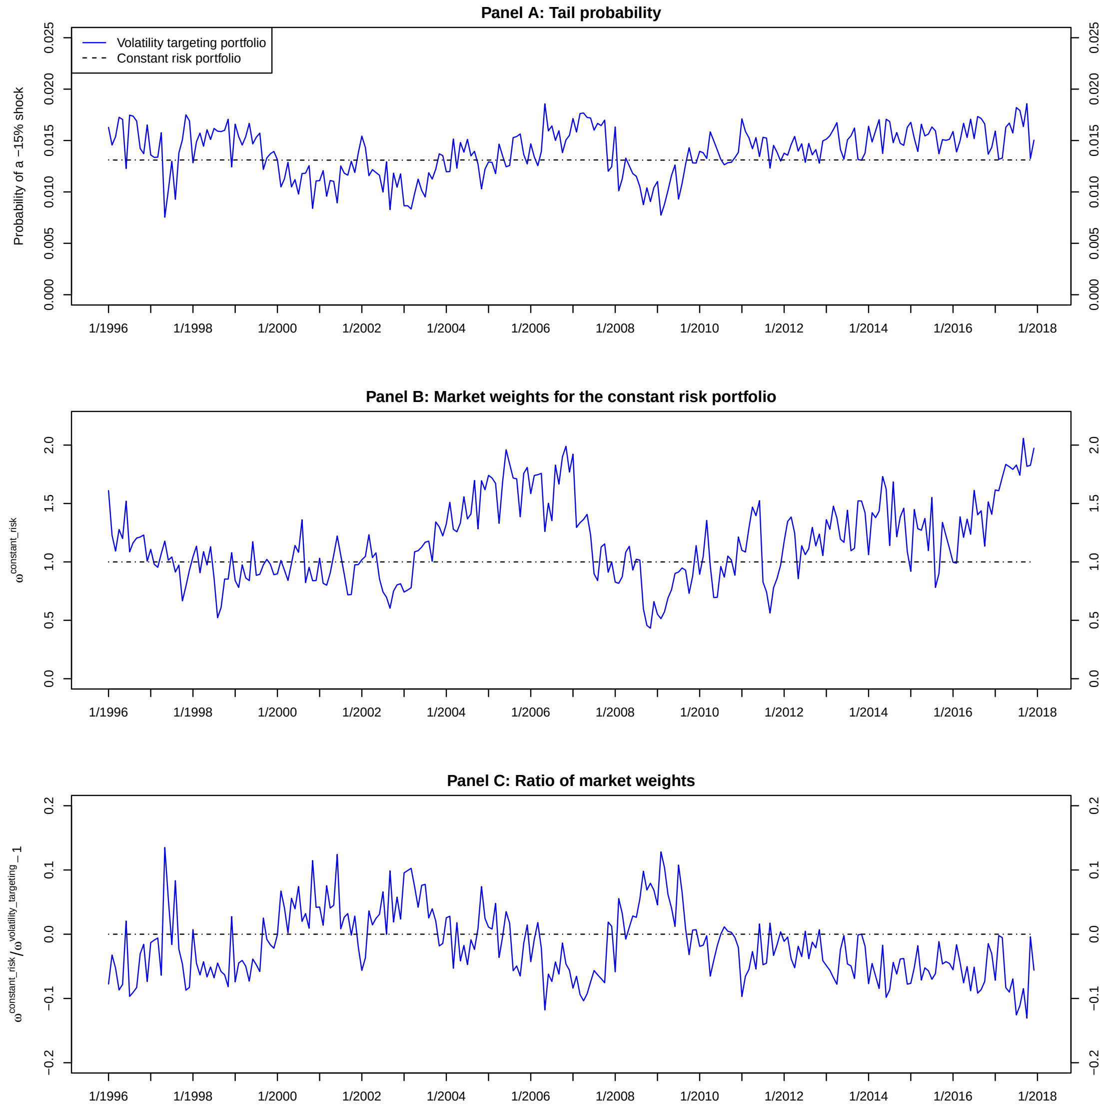

Figure 11: **A constant-risk portfolio (Monthly Horizon).** Panel A shows the conditional probability of a 15% shock to the monthly return of: (i) a volatility targeting investor who targets 5% monthly volatility (solid blue line) and (ii) a constant-risk investor who targets a constant probability of a 15% shock (dashed black line). Panel B shows the weight invested in the market by the constant-risk investors. Panel C shows the ratio of market weights between the constant-risk and the volatility targeting investors (solid blue line).

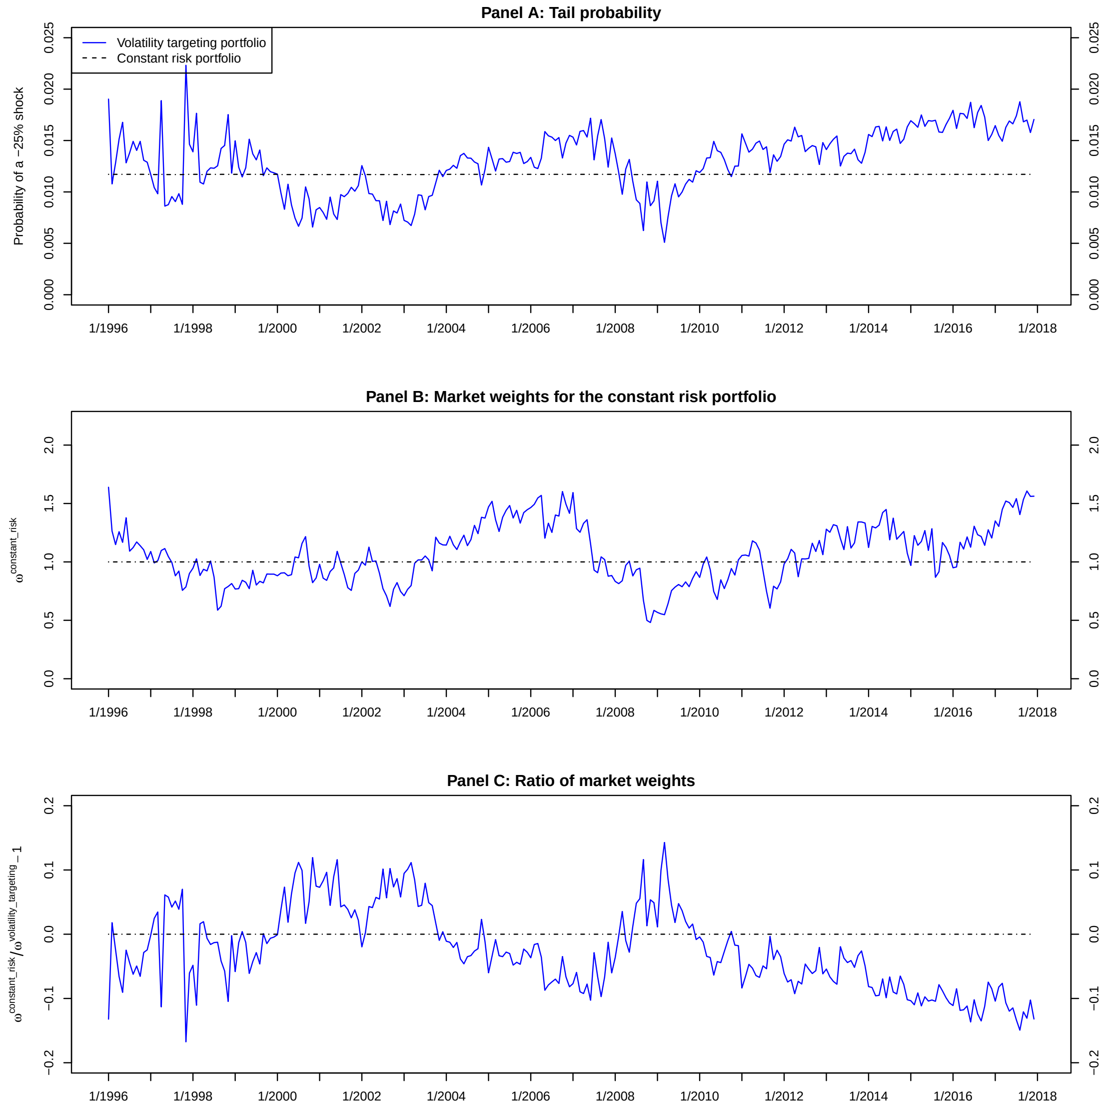

Figure 12: **A constant risk portfolio (Quarterly Horizon).** Panel A shows the conditional probability of a 25% shock to the quarterly return of: (i) a volatility targeting investor who targets 8.33% quarterly volatility (solid blue line) and (ii) a constant-risk investor who targets a constant probability of a 25% shock (dashed black line). Panel B shows the weight invested in the market by the constant-risk investors. Panel C shows the ratio of market weights between the constant-risk and the volatility targeting investors (solid blue line).

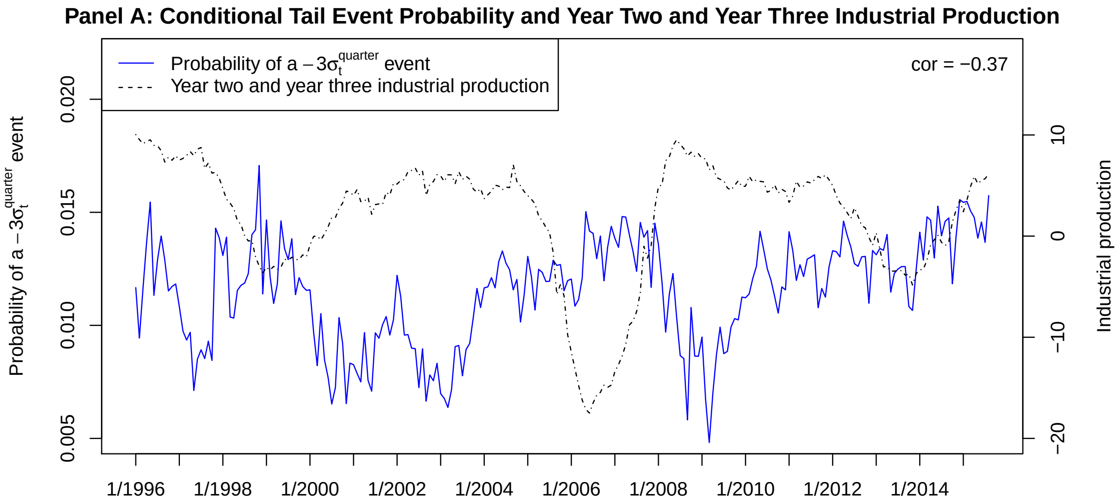

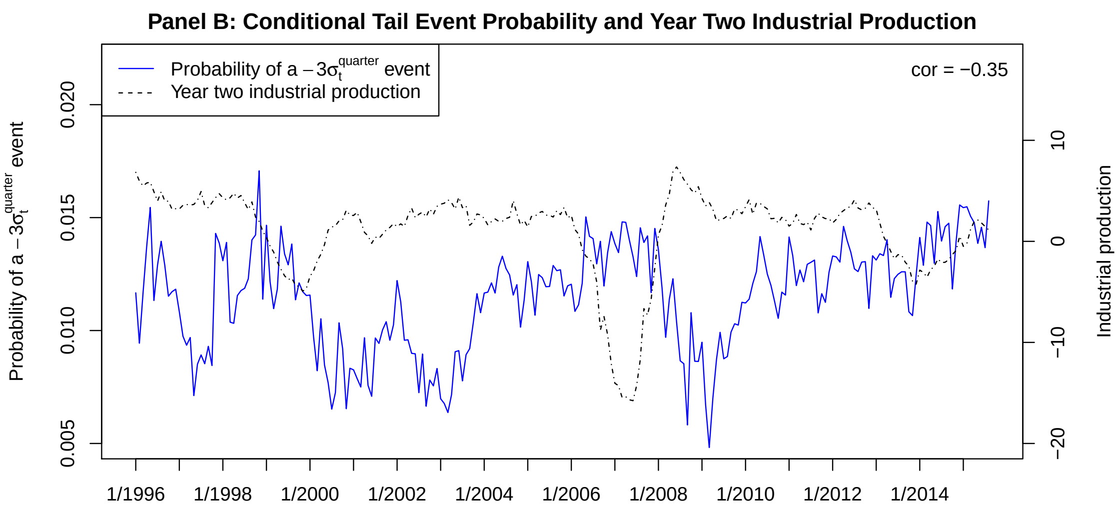

Figure 13: **Conditional Tail Loss Probabilities and Industrial Production.** Panel A shows time series of the probability of a $-3\sigma_t^{\mathrm{month}}$ event (solid) along with the level of industrial production three years ahead minus the level of industrial production one year ahead (dashed). Panel B shows time series of the probability of a $-3\sigma_t^{\mathrm{month}}$ event (solid) along with the level of industrial production two years ahead minus the level of industrial production one year ahead (dashed).

Table 1: **Summary Statistics of S&P 500 Moments.** This table reports the average time-series values of the following ex ante moments of the return distribution of the S&P 500: 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	18.71	-1.11	5.39	-14.03	58.35
Month	$\gamma = 1$	4.01	18.07	-1.09	5.53	-14.58	63.13
Month	$\gamma = 3$	11.09	16.97	-1.02	5.71	-15.07	70.98
Month	$\gamma = 5$	17.25	16.11	-0.91	5.76	-14.74	75.84
Quarter	$\gamma = 0$	0	19.56	-1.07	4.66	-11.14	40.06
Quarter	$\gamma = 1$	4.19	18.40	-1.07	4.97	-12.56	48.41
Quarter	$\gamma = 3$	11.01	16.60	-0.98	5.26	-13.92	61.36
Quarter	$\gamma = 5$	16.47	15.34	-0.83	5.19	-13.13	65.25

Table 2: **Correlations Between S&P 500 Moments.** This table reports the pairwise correlations between the following S&P 500 moments: expected return (Er), variance (Var), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). Panel A reports results for monthly moments and Panel B reports results for the quarterly moments. We report 95% block bootstrapped confidence bounds in brackets.

Panel A: Month

	Er	Var	−Skew	Kurt	−Hskew	Hkurt
Er	1	0.99	−0.42	−0.44	−0.42	−0.37
		[0.99,1]	[−0.51,−0.33]	[−0.51,−0.38]	[−0.49,−0.33]	[−0.44,−0.30]
Var		1	−0.46	−0.48	−0.46	−0.40
			[−0.55,−0.39]	[−0.54,−0.43]	[−0.53,−0.38]	[−0.47,−0.33]
−Skew			1	0.83	0.90	0.72
				[0.78,0.86]	[0.88,0.93]	[0.65,0.79]
Kurt				1	0.93	0.94
					[0.91,0.94]	[0.93,0.96]
−Hskew					1	0.93
						[0.91,0.95]
Hkurt						1

Panel B: Quarter

	Er	Var	−Skew	Kurt	−Hskew	Hkurt
Er	1	0.99	−0.51	−0.52	−0.54	−0.51
		[0.99,0.99]	[−0.58,−0.44]	[−0.57,−0.47]	[−0.60,−0.48]	[−0.57,−0.46]
Var		1	−0.58	−0.58	−0.61	−0.57
			[−0.64,−0.52]	[−0.64,−0.54]	[−0.66,−0.56]	[−0.62,−0.53]
−Skew			1	0.87	0.95	0.86
				[0.85,0.90]	[0.93,0.96]	[0.82,0.89]
Kurt				1	0.95	0.98
					[0.94,0.96]	[0.96,0.99]
−Hskew					1	0.97
						[0.96,0.98]
Hkurt						1

Table 3: **Principal Components of Higher Order Moments.** This table presents the results of a principal component analysis of the ex ante skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt) of the S&P 500. Panel A reports results on for the monthly moments and Panel B reports results for the quarterly moments. The rightmost column shows the amount of variation explained by the various principal components. The last row in each panel shows the correlation between the first principal component and the different moments.

Panel A: Monthly horizon

	Skew	Kurt	Hskew	Hkurt	Variation explained
PC 1 eigenvector	−0.47	0.51	−0.52	0.50	91%
PC 2 eigenvector	−0.78	−0.24	−0.07	−0.58	7%
PC 3 eigenvector	0.07	−0.78	−0.56	0.29	2%
PC 4 eigenvector	0.40	0.28	−0.65	−0.58	0%
PC 1 correlation	−0.96	0.93	−0.95	0.82	

Panel B: Quarterly horizon

	Skew	Kurt	Hskew	Hkurt	Variation explained
PC 1 eigenvector	−0.49	0.50	−0.51	0.50	95%
PC 2 eigenvector	−0.79	−0.39	−0.10	−0.46	4%
PC 3 eigenvector	0.19	−0.74	−0.50	0.41	1%
PC 4 eigenvector	−0.33	−0.22	0.69	0.60	0%
PC 1 correlation	−0.97	0.89	−0.92	0.83	

Table 4: **Cyclicality in Higher-Moment Risk.** Panel A reports correlations between ex ante variance, past two-year return, and cay (Lettau and Ludvigson (2001)) with and the higher-moment risk index (HRI). Panel B reports correlations between the past two year returns of the S&P 500 index and the individual moments: variance (Var), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). Panel C reports correlations between cay and the individual moments. We report block bootstrapped 95% confidence intervals in brackets.

Panel A: Correlations between HRI and variance, past 2-year return, and cay

Horizon	Variance	Past 2-year return	Cay
Month	−0.47	0.35	−0.50
95% CI	[−0.54, −0.41]	[0.24, 0.45]	[−0.59, −0.41]
Quarter	−0.60	0.26	−0.74
95% CI	[−0.65, −0.55]	[0.15, 0.36]	[−0.80, −0.68]

Panel B: Correlations between past 2 year return and individual higher order moments

Horizon	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	−0.42	−0.41	0.28	−0.37	0.27
95% CI	[−0.51, −0.30]	[−0.51, −0.30]	[0.18, 0.38]	[−0.47, −0.26]	[0.15, 0.38]
Quarter	−0.42	−0.37	0.15	−0.28	0.19
95% CI	[−0.52, −0.30]	[−0.47, −0.26]	[0.05, 0.26]	[−0.38, −0.17]	[0.09, 0.29]

Panel C: Correlations between cay and individual higher order moments

Horizon	Var (%)	Skew	Kurt	Hskew	Hkurt
Month	0.47	0.55	−0.52	0.46	−0.38
95% CI	[0.4, 0.55]	[0.46, 0.64]	[−0.59, −0.43]	[0.36, 0.55]	[−0.47, −0.28]
Quarter	0.52	0.71	−0.73	0.72	−0.71
95% CI	[0.45, 0.60]	[0.64, 0.78]	[−0.79, −0.68]	[0.66, 0.78]	[−0.77, −0.64]

Table 5: **Expected and Realized Moments on Constant-Volatility Portfolios.** This table reports the results of time-series portfolios sorted on ex ante moments. The High (Low) moment portfolio is a strategy that invests in the market in the periods where the ex ante moment is above (below) the time-series median. The position in the market is scaled by the ex ante conditional volatility to create constant-volatility portfolios. We then calculate standardized moments of realized excess returns to the two portfolios. We report realized moments of the low- and high-moment portfolios, the difference between the two, and block bootstrapped p-values for the hypothesis that the difference between the high- and low-moment portfolios is non-positive. We report significance as: * when $p < 0.1$, ** when $p < 0.05$, and *** when $p < 0.01$. Panel A and B sort on option implied moments whereas Panel C sort on past realized moments.

Panel A: Portfolios sorted on ex ante moments, weighing by option implied volatility

	Monthly horizon				Quarterly horizon			
	Skew	Kurt	Hskew	Hkurt	Skew	Kurt	Hskew	Hkurt
Low-moment portfolio	−0.56	2.29	−6.38	7.59	−0.75	2.24	−6.26	7.64
High-moment portfolio	−0.44	4.09	−2.15	33.04	−0.48	3.70	−2.17	19.37
High−Low	0.12	1.80***	4.23	25.45***	0.27	1.46**	4.09	11.73
p-value (H_0: High−Low ≤ 0)	0.34	0.007	0.16	0.002	0.24	0.05	0.11	0.11

Panel B: Portfolios sorted on ex ante moments, weighing by AR(1) volatility

	Monthly horizon				Quarterly horizon			
	Skew	Kurt	Hskew	Hkurt	Skew	Kurt	Hskew	Hkurt
Low-moment portfolio	−0.83	2.27	−12.52	7.14	−0.74	2.24	−6.82	7.43
High-moment portfolio	−0.27	5.00	−1.35	60.62	−0.35	4.01	−1.31	21.07
High−Low	0.56*	2.73***	11.17**	53.18***	0.39	1.77**	5.51*	13.64*
p-value (H_0: High−Low ≤ 0)	0.09	0.001	0.02	0.001	0.16	0.02	0.06	0.09

Panel C: Portfolios sorted on past realized returns, weighing by AR(1) volatility

	Monthly horizon				Quarterly horizon			
	Skew	Kurt	Hskew	Hkurt	Skew	Kurt	Hskew	Hkurt
Low-moment portfolio	−0.70	3.38	−8.90	24.35	−0.63	2.76	−3.06	10.61
High-moment portfolio	−0.26	2.63	−1.26	10.24	−0.52	3.27	−3.77	14.19
High−Low	0.44	−0.75	7.64	−14.11	0.11	0.51	−0.71	3.58
p-value (H_0: High−Low ≤ 0)	0.17	0.65	0.25	0.64	0.38	0.33	0.63	0.48

Table 6: **Conditional and Unconditional Tail Risk.** Panel A reports correlations between ex ante log-variance and tail loss probabilities. The probabilities are $P(r_{t,T} < -3\bar{\sigma}_{t,T})$ and $P(r_{t,T} < -3\sigma_{t,T})$ where $r_{t,T} = R_{t,T} - E_t[R_{t,T}]$, $\sigma_{t,T}$ is the ex ante volatility from time t to T, and we define $\bar{\sigma}$ as the time series average of $\sigma_{t,T}$ at the appropriate horizon. We report block bootstrapped 95% confidence intervals in brackets. Panel B reports regression slope coefficients for regressions of physical tail loss probabilities onto the past two year returns. Panel C reports slope coefficients for regressions of tail loss probabilities onto cay (Lettau and Ludvigson (2001)). 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). We use the automatic bandwidth procedure of Newey and West (1994) to select the optimal number of lags.

Panel A: Tail probabilities and log-variance (correlation coefficients)

Horizon	$P_t(r_{t,T} < -3\bar{\sigma})$	$P_t(r_{t,T} < -5\bar{\sigma})$	$P_t(r_{t,T} < -3\sigma_{t,T})$	$P_t(r_{t,T} < -5\sigma_{t,T})$
Month	0.84	0.71	−0.47	−0.79
95% CI	[0.82,0.89]	[0.66,0.78]	[−0.56, −0.36]	[−0.83,−0.75]
Quarter	0.83	0.65	−0.59	−0.81
95% CI	[0.80,0.87]	[0.60,0.72]	[−0.67,−0.50]	[−0.84,−0.78]

Panel B: Tail probabilities (%) and past return (slope coefficients)

Horizon	$P_t(r_{t,T} < -3\bar{\sigma})$	$P_t(r_{t,T} < -5\bar{\sigma})$	$P_t(r_{t,T} < -3\sigma_{t,T})$	$P_t(r_{t,T} < -5\sigma_{t,T})$
Month	−2.05*	−0.54	0.36***	0.11**
(s.e.)	(1.15)	(0.34)	(0.08)	(0.05)
Quarter	−2.00	−0.34	0.34***	0.07
(s.e.)	(1.28)	(0.25)	(0.12)	(0.06)

Panel C: Tail probabilities (%) and cay (slope coefficients)

Horizon	$P_t(r_{t,T} < -3\bar{\sigma})$	$P_t(r_{t,T} < -5\bar{\sigma})$	$P_t(r_{t,T} < -3\sigma_{t,T})$	$P_t(r_{t,T} < -5\sigma_{t,T})$
Month	42.51***	8.82**	−3.69**	−3.10***
(s.e.)	(14.15)	(3.92)	(1.85)	(0.68)
Quarter	35.58**	4.63*	−7.74***	−3.31***
(s.e.)	(14.33)	(2.53)	(1.90)	(0.55)

Table 7: **Tail Risk and Industrial Production.** This table reports the results of regressing industrial production (IND) onto the conditional probability of a tail event. 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). We use the automatic bandwidth procedure of Newey and West (1994) to select the optimal number of lags.

	Dependent variable:					
	$\mathrm{IND}_{t+36} - \mathrm{IND}_t$			$\mathrm{IND}_{t+36} - \mathrm{IND}_{t+12}$		
	(1)	(2)	(3)	(4)	(5)	(6)
$P_t(r_{t,t+1} < -3\sigma_{t,t+1})$	-801.99			-1021.74^{***}		
(s.e.)	(549.14)			(353.49)		
$P_t(r_{t,t+3} < -3\sigma_{t,t+3})$		-1002.72^{*}			-982.21^{*}	
(s.e.)		(511.71)			(544.20)	
Economic policy uncertainty			521.91			(544.20)
(s.e.)			(704.29)			516.53
No. obs.	235	235	235	235	235	235
Adj. R^2	0.06	0.11	0.07	0.14	0.14	0.10

Table 8: **Correlations Between S&P 500 Moments Under Risk-Neutrality.** This table reports pairwise correlations between monthly risk-neutral S&P 500 moments: variance (Var), skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). We report 95% block bootstrapped confidence bounds in brackets.

	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

Table 9: **Principal Components of Higher-Moment Risks Under Risk-Neutrality.** We estimate the four principal components (PC) spanning the space of monthly risk-neutral skewness (Skew), kurtosis (Kurt), hyperskewness (Hskew), and hyperkurtosis (Hkurt). The last column of Panel A shows that the first principal component ($PC1^{RN}$) explains 95% of the variation in monthly risk neutral higher order moments.

	$Skew^{RN}$	$Kurt^{RN}$	$Hskew^{RN}$	$Hkurt^{RN}$	Variation explained
$PC1^{RN}$ eigenvector	−0.48	0.51	−0.51	0.50	95%
$PC2^{RN}$ eigenvector	0.78	−0.08	−0.27	0.57	5%
$PC3^{RN}$ eigenvector	−0.41	−0.76	0.10	0.49	0%
$PC4^{RN}$ eigenvector	−0.00	−0.38	−0.81	−0.44	0%
$PC1^{RN}$ correlation	−0.94	0.99	−0.99	0.97	

Appendix A 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:

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

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

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

Appendix B Details on Simulations of Gabaix (2012)

Within the Gabaix (2012) disaster model, we estimate the markets risk-neutral moments by: (i) simulating option prices, (ii) converting these prices into a risk-neutral density using Breeden and Litzenberger (1978), (iii) compute moments of the risk-neutral distribution. In our simulation, we let variation in option prices arise from a time-varying recovery rate of the market in case of a disaster. Specifically, we let the size of recovery in case of a disaster be uniform on the interval from 0.1 to 0.95, i.e., if a disaster occurs then the value of the stock market drops by 10% to 90%. Table 10 summarizes monthly horizon parameter values used in the simulation.

Table 10: **Variables used in the simulation**

Variables	Values
Time preference parameter	0.55%
Risk aversion	3
Growth rate of dividends	0.21%
Volatility of dividends	3.2%
Probability of disaster	3.63%
Normal time stock volatility	4%
Stocks recovery rate	Varies from 0.10 to 0.95

Appendix C Estimating weights in constant-risk portfolios

We consider an investor who targets a specific probability ('target probability') of a certain size shock ('shock') to her portfolio. We estimate the 'constant-risk' weights using the following algorithm:

1. Make an initial guess on the fraction of wealth invested in the market (initial weight, $\omega_{t,T}^{\text{initial}}$).

2. Using Proposition 1, estimate the probability of a 'shock' to the portfolio given the initial weight, $P_t(R_{t,T}^{\text{portfolio}} < \text{'shock'}|\omega_{t,T}^{\text{initial}})$.

3a. If the absolute difference between the 'target probability' and the estimated probability of a 'shock' is greater than 0.0001 then go to 4.

3b. If the absolute difference between the 'target probability' and the estimated probability of a 'shock' is smaller than 0.0001 then end.

4. Update the fraction of wealth invested in the market as:

$$\begin{aligned}\omega_{t,T}^{\text{updated}} &= \omega_{t,T}^{\text{initial}} + 0.0001 \times 1_{\{P_t(R_{t,T}^{\text{portfolio}} < \text{'shock'}|\omega_{t,T}^{\text{initial}}) < \text{'target probability'}\}} \\ &\quad - 0.0001 \times 1_{\{P_t(R_{t,T}^{\text{portfolio}} < \text{'shock'}|\omega_{t,T}^{\text{initial}}) > \text{'target probability'}\}}\end{aligned}$$

5. Replace the initial weight in 2. with the updated weight and repeat 2.-4. until

$$|P_t(R_{t,T}^{\text{portfolio}} < \text{'shock'}|\omega_{t,T}^{\text{updated}}) - \text{'target probability'}| < 0.0001.$$