Characterizing the Tail-Risk of Factor Mimicking Portfolios *

Andreas Johansson† (v2.3)
November 20, 2019

Abstract

By assuming that short-run returns are independent and identically distributed, it is straightforward to extrapolate short-run risks to longer horizons. However, by generalizing the variance-ratio test to include higher co-moments, we establish a significant and sizable intertemporal dependency in all higher moments of equity returns. The intertemporal dependency is strong enough to prevent the convergence to normally distributed returns, at least up to a five-year holding period. We also demonstrate that the intertemporal dependency is both horizon *and* portfolio-specific. Consequently, the common practice of extrapolating the short-run risk by assuming independent and identically distributed returns will severely bias the expected long-run risk.

JEL: G11, G12, C22, C12, C15, C58

Keywords: intertemporal dependency, risk factors, skewness, kurtosis, bootstrap

*I am especially grateful for the support and help from my main adviser Michael Halling and my committee members Jungsuk Han and Per Strömberg. I am also grateful to Peter Schotman, Petri Jylhä, Campbell R. Harvey, and Erik Hjalmarsson for helpful discussions. Also, thanks to participants at the 2019 NFN PhD Workshop, 8th National PhD Workshop in Finance, and the Swedish House of Finance seminars. Finally, I gratefully acknowledge the financial support from the Swedish Bank Research Foundation, Louis Fraenckels Foundation, and Stockholm School of Economics.

†Adress: Stockholm School of Economics, Department of Finance, Drottninggatan 98, 111 60 Stockholm, Sweden, Email: andreas.johansson@phdstudent.hhs.se, Web: Andreas Johansson

A simplifying assumption in finance is that returns are independent and identically distributed (iid). Yes, it is an approximation, but it is usually a surprisingly good approximation, at least when fitting the mean and the volatility. However, when we include higher moments, skewness and kurtosis,[1] the approximation seems to break down. For example, under the assumption of iid monthly returns, we would expect to observe a skewness of -0.09, and excess kurtosis of 0.13 at a five-year horizon, Table 1. Instead, we actually observe a five-year skewness and excess kurtosis of -0.69 and 0.99, respectively. Consequently, when we include the tail risk,[2] it suggests that returns are not iid, not even as an approximation.

Therefore we wonder, how are returns distributed at different horizons? Do returns converge as if they were iid? Moreover, a considerable fraction of the asset pricing literature is based upon the existence of a linear risk factor structure. Since each risk factor should capture a different type of risk, each should also exhibit a different return distribution. Hence, do we observe heterogeneity in the return distributions between different factor mimicking portfolios? And finally, is the seemingly innocent assumption of iid returns reasonable when we extrapolate short-run estimates to a longer horizon?

This paper makes three contributions. First, we establish the empirical facts concerning the return distribution over different holding periods. Under the assumption of iid returns, skewness and kurtosis should converge to zero as short-run log-returns are compounded to a longer horizon. In contrast, we demonstrate that the skewness can sometimes diverge from zero as the holding period gets longer, which is the exact opposite of what is expected from iid returns. The kurtosis is also inflated compared to what would be expected from continuously compounded independent returns, which implies that even long-run portfolio returns are prone to massive shocks. Hence,

[1] The statistical skewness captures the tilt of the distribution; are we more likely to observe positive or negative deviations around the mean? The statistical kurtosis captures the probability of observing massive deviations around the mean, the "fatness" of the tails of the distribution.

[2] There exist slightly different definitions of "tail risk" in the literature. Expected shortfall? Value at Risk? Probability of a three standard deviation shock? In this paper, "tail-risk" will refer to a non-zero skewness and excess kurtosis, which will increase the risk of huge shocks, no matter what tail risk definition we prefer.

we establish that there exists a term structure in the distribution of equity returns.

Second, we generalize the variance ratio test of Lo and Mackinlay (1988) to also test for intertemporal dependencies in higher moments, e.g., skewness and kurtosis. The test demonstrates that there exists a significant intertemporal dependence in returns, no matter the portfolio, and no matter the holding period. The intertemporal dependency also switches sign as we go from short-run returns to long-run returns. Consequently, similar to how there exists a mean-reversal in long-run returns, there is also a reversal pattern in skewness and kurtosis as the holding period expands.

Finally, we demonstrate that the common practice of extrapolating the short-run risk to a longer horizon by assuming independent returns will create a severely biased long-run expectation of the risk exposure. Not just due to the covariance as established by Wang, Yeh, and Cheng (2011), but also due to the coskewness and cokurtosis. Hence, it is not only the volatility that will become biased but also the expected probability of severe shocks, the skewness and kurtosis. Similarly to Neuberger and Payne (2019), we demonstrate that the bias is due to intertemporal dependencies, which will dominate the unconditional higher moments of the high-frequency returns as they are compounded to a more extended holding period.

Since several studies have documented that the expected skewness and excess kurtosis of returns are priced in the market,[3] we anticipate our results to have a direct implication for asset pricing models. A risk-averse investor is prepared to pay a premium to get higher return skewness, while they want to be compensated to carry more exposure to variance and kurtosis. However, under the common assumption of iid returns, the continuously compounded returns should converge to a distribution with zero skewness and zero excess kurtosis (no excess tail-risk). Therefore, we would expect that investors with different investment horizons to perceive a different amount of risk in the

[3] Kraus and Litzenberger (1976), Rietz (1988), Harvey and Siddique (2000), Barro (2006), Mitton and Vorkink (2007), Barberis et al. (2008), Boyer, Mitton, and Vorkink (2010), Rehman and Vilkov (2012), Bali and Murray (2013), Conrad, Dittmar, and Ghysels (2013), Kelly and Jiang (2014), Amaya et al. (2015), Ghysels, Plazzi, and Valkanov (2016), and Schneider, Wagner, and Zechner (2019)

market. Hence, if the tail-risk is priced in the market, it should also create a term structure in the observed risk-premia.

The market also exhibits a seasonality pattern in company reporting, macro events, and the tax code. These types of market events seem to produce abnormal returns (De Bondt and Thaler 1986; Sias 2007; Lucca and Moench 2015; Umar 2017). In other words, there exist reoccurring events that create unusually large price shocks in equities. Using a similar argument as Albuquerque (2012), that the observed negative coskewness between firms is due to reporting seasonality,[4] we argue that the clustering of market events should create an intertemporal dependence in returns. As different holding periods are exposed to different types of market events, it should yield different co-moments on average, which should create a term structure in the return distribution.

In theory, investors should only be compensated for exposure to systematic risk. However, empirically, there seem to exist multiple types of systematic risk, so-called risk factors, which can be approximated by creating factor mimicking portfolios. As these risk factors capture different types of risk, they should also exhibit different return distributions. Moreover, any asset that loads on a specific risk factor will also inherit the tail risk of that particular risk factor. Hence, if the linear factor model is the correct model of risk, then the risk-*premia* distribution of any asset is explained by its exposure to the underlying risk factors. Therefore, by characterizing the return distribution of the risk factors, we can deduce the risk-premia distribution of all assets in the economy.[5]

Fama and French (2018) also tries to characterize the long-run return distribution of equities.[6] However, they use a bootstrap procedure, which

[4]Harvey and Siddique (2000) show that individual firms exhibit positive return skewness, while a portfolio of the same firms has a negative return skewness. They demonstrate that this surprising pattern is due to a negative coskewness between firms, which will push the return skewness downward as we add more assets to our portfolio.

[5]However, if we do not believe in this particular factor structure, MKT, SMB, HML, and MOM, our results will not be generalizable to all assets in the economy. However, we argue that it should still yield a reasonable approximation of the potential risk-premia distributions of equity assets in the economy.

[6]For a in-depth discussion see Appendix D.b.

removes any dependencies between the returns in the sample. In contrast, our approach keeps the dependency structure intact, which allows us to test for intertemporal dependencies explicitly. As our results suggest that there is a strong intertemporal dependence in returns, no matter the holding period, we contend that their conclusions need to be treated with skepticism. Surprisingly though, we both reach a similar conclusion in the end. While their results suggest that returns become normally distributed at holding periods longer than 10-years, our results demonstrate that equity returns are not normally distributed up to a five-year horizon.

In contemporaneous work, Neuberger and Payne (2019) demonstrates that the precision of the skewness and kurtosis estimator can be improved by accounting for the intertemporal dependency in the returns. An essential assumption in the improved estimator is that returns follow a Martingale process, which implies that information today cannot be used to predict the direction of future returns. However, our results suggest that squared and cubed returns today actually have a strong relationship with future returns. Hence, empirically, it seems like returns do not follow a Martingale process, which implies that their estimator will be biased in many real-life applications. Therefore, we suggest that practitioners first estimate the intertemporal dependency in skewness and kurtosis before applying the Neuberger and Payne (2019) estimator. Moreover, since our results require fewer assumptions, while still being able to estimate the intertemporal dependencies consistently, our approach is simpler to employ.[7]

In summary, we develop a new intertemporal dependency test, which makes use of the information contained in the higher moments. The test demonstrates that there exists a significant intertemporal dependency in equity returns, which is strong enough to counteract the implied convergence by iid returns. Moreover, the intertemporal dependency is both horizon *and* portfolio specific. Consequently, we cannot make a general statement about the tail-risk as we move from a short-run return to a longer horizon. Our results demonstrate that the common practice of extrapolating a short-run estimate to a longer horizon by assuming iid returns will severely bias the

[7] For an in-depth discussion see Appendix D.a.

long-run estimate.

In Section 1, we establish the necessary theoretical results that are needed for the rest of the analysis, followed by an exposition of the data and empirical findings in Section 2.

1 Theoretical Setup

All analysis is performed by using continuously compounded portfolio returns

$$\begin{aligned} r^{(h)} = \ln(R^{(h)}) &= \ln\left(R_1 \cdot R_2 \cdots R_h\right) \\ &= \ln(R_1) + \ln(R_2) + \ldots + \ln(R_h), \end{aligned} \tag{1}$$

where R is one plus the simple return. The central limit theorem implies that as $h \to \infty$, a sum of well-behaved[8] random variables will converge to a normal distribution. For example, we can view the monthly return as a sum of daily returns, which should be approximately normally distributed. However, we show that up to five-year returns are not normally distributed, both the skewness and excess kurtosis are significantly different from zero.

1.a Theoretical Foundation of Higher Moments

A common assumption/approximation in the literature is that returns are independent and identically distributed (iid), which simplifies most analytical solutions. Moreover, iid returns are also a sufficient condition for the efficient market hypothesis to hold.

If we can assume that returns, r_i, are iid, we also know the aggregation pattern in the first four centralized moments

$$\begin{aligned} \mathbb{E}\left[r_i^{(h)}\right] &= h\mathbb{E}\left[r_i\right], & \mathrm{Var}\left[r_i^{(h)}\right] &= h\mathrm{Var}\left[r_i\right], \\ SK\left[r_i^{(h)}\right] &= \frac{SK\left[r_i\right]}{\sqrt{h}}, & KU^e\left[r_i^{(h)}\right] &= \frac{KU^e\left[r_i\right]}{h}, \end{aligned} \tag{2}$$

[8]Defined higher moments, ergodicity, stationarity, and not too dependent.

where KU^e denotes excess kurtosis. Hence, it implies that we can take the high-frequency estimates and scale them by the horizon to get a consistent estimate of the long-run centralized moments.

Since the skewness and kurtosis are constant under affine transformations

$$\mathrm{SK}\left[r_i\right] = \mathrm{SK}\left[a + b r_i\right], \qquad \mathrm{KU}\left[r_i\right] = \mathrm{KU}\left[a + b r_i\right], \qquad b > 0, \tag{3}$$

it implies that in a one-factor model, e.g., the Capital Asset Pricing Model (CAPM), the asset returns will exhibit the same systematic skewness and kurtosis as the market portfolio. Accordingly, any asset that loads on a specific risk factor will also inherit the tail risk of that specific risk factor. Consequently, if the linear factor model is correct, the risk-premia of any asset is given by the specific loadings on the underlying risk factors. Therefore, the focus of this paper will be on characterizing the return distribution of the factor mimicking portfolios, which is the same thing as characterizing the risk-premia distribution of all assets in the economy (given the assumption that the linear risk-factor model is the correct model of risk).

1.b Linear Combinations and Higher Moments

Most of the analysis will involve continuously compounded returns, which is a linear combination of returns, e.g., $r_t^{(2)} = r_t + r_{t-1}$. From statistical theory[9]

[9] The theory is based on statistical cumulants, which yields much simpler expressions. However, since our focus is on the centralized higher moments, we choose to keep everything consistent throughout the paper, and express everything as centralized higher moments instead.

we know that

$$
\begin{aligned}
\sigma^2_{x+y} = \mathrm{Var}\left[X+Y\right] &= \mathrm{Var}\left[X\right] + \mathrm{Var}\left[Y\right] + 2\mathrm{Cov}\left[X,Y\right], &(4)\\
\mathrm{SK}\left[X+Y\right] &= \frac{1}{\sigma^3_{x+y}}\Big\{\sigma^3_x\mathrm{SK}\left[X\right] + 3\sigma^2_x\sigma_y\mathrm{SK}\left[X,X,Y\right] &(5)\\
&\quad +3\sigma_x\sigma^2_y\mathrm{SK}\left[X,Y,Y\right] + \sigma^3_y\mathrm{SK}\left[Y\right]\Big\},\\
\mathrm{KU}\left[X+Y\right] &= \frac{1}{\sigma^4_{x+y}}\Big\{\sigma^4_x\mathrm{KU}\left[X\right] + 4\sigma^3_x\sigma_y\mathrm{KU}\left[X,X,X,Y\right] &(6)\\
&\quad + 6\sigma^2_x\sigma^2_y\mathrm{KU}\left[X,X,Y,Y\right]\\
&\quad +4\sigma_x\sigma^3_y\mathrm{KU}\left[X,Y,Y,Y\right] + \sigma^4_y\mathrm{KU}\left[Y\right]\Big\},
\end{aligned}
$$

where

$$
\begin{aligned}
\mathrm{SK}\left[X,Y,Z\right] &= \frac{\mathbb{E}\left[(X-\mathbb{E}\left[X\right])(Y-\mathbb{E}\left[Y\right])(Z-\mathbb{E}\left[Z\right])\right]}{\sigma_x\sigma_y\sigma_z},\\
\mathrm{KU}\left[X,Y,Z,V\right] &= \frac{\mathbb{E}\left[(X-\mathbb{E}\left[X\right])(Y-\mathbb{E}\left[Y\right])(Z-\mathbb{E}\left[Z\right])(V-\mathbb{E}\left[V\right])\right]}{\sigma_x\sigma_y\sigma_z\sigma_v},
\end{aligned}
$$

and $X, Y, Z,$ and V are random variables (and by a slight abuse of notation $\mathrm{SK}\left[X\right] = \mathrm{SK}\left[X,X,X\right]$ and $\mathrm{KU}\left[X\right] = \mathrm{KU}\left[X,X,X,X\right]$).

Note that three things drive the skewness of a linear combination of two random variables. First, under an assumption of independence, we know that both the covariance and coskewness will be equal to zero, and the skewness of the linear combination will be

$$
\mathrm{SK}\left[X+Y\right] = \frac{\sigma^3_x\mathrm{SK}\left[X\right] + \sigma^3_y\mathrm{SK}\left[Y\right]}{\sigma^3_{x+y}},
$$

which means that the skewness only depends on the individual skewnesses.

Second, the covariance will adjust the skewness towards or away from zero. Hence, everything else equal we expect the skewness to be closer to zero if the two variables X and Y are positively correlated.

$$\mathrm{SK}\left[X+Y\right]=\frac{\sigma_x^3\mathrm{SK}\left[X\right]+\sigma_y^3\mathrm{SK}\left[Y\right]}{\left(\sigma_x^2+\sigma_y^2+2\rho\sigma_x\sigma_y\right)^{3/2}}.$$

Finally, even if the two random variables are symmetric individually, $\mathrm{SK}\left[X\right]=\mathrm{SK}\left[Y\right]=0$, their combination might still be skewed due to the coskewness

$$\mathrm{SK}\left[X+Y\right]=\frac{3\sigma_x^2\sigma_y\mathrm{SK}\left[X,X,Y\right]+3\sigma_x\sigma_y^2\mathrm{SK}\left[X,Y,Y\right]}{\sigma_{x+y}^3}.$$

When we consider a sum of log-returns over time, one of the coskewness terms, $\mathrm{SK}\left[r_t,r_{t+1},r_{t+1}\right]<0$, is the same thing as the leverage effect, the negative correlation between future volatility and today's return (Engle and Mistry 2014; Neuberger and Payne 2019).

Similarly, we can also decompose the kurtosis of a sum of random variables into its parts. The individual kurtosises will explain the kurtosis if the returns are independent, and a positive covariance will decrease the kurtosis, all else equal. The difference compared to skewness, is which co-moments that will contribute to the kurtosis. The kurtosis contains three cokurtosis terms $\mathrm{KU}\left[X,X,X,Y\right],\mathrm{KU}\left[X,Y,Y,Y\right],$ and $\mathrm{KU}\left[X,X,Y,Y\right]$.

The sum of the first two cokurtosis terms will be positive if big shocks in one variable correlate with a move in the same direction in the other variable, e.g., how yesterday's cubed returns relate to today's returns. Hence, a positive cokurtosis would increase the fatness of the tails compared to what we would expect from the individual kurtosises. The last cokurtosis term depends on the covariance between the squared shocks, which is substantial if there is persistence in variance. As it is already well-established that equity returns exhibit a strong persistence in variance, $\mathrm{KU}\left[r_{t-1},r_{t-1},r_t,r_t\right]>0$,[10] it should also inflate the kurtosis as returns are compounded to a more extended holding period.

To summarize, most of the observed patterns in skewness and kurtosis

[10]For example, Bollerslev (1986), Nelson (1991), Ding, Granger, and Engle (1993), and Zakoian (1994).

are explained by two parts; the individual return skewnesses/kurtosises and their co-moments (coskewness/cokurtosis).

1.c The Moment-"Ratio" Test

We know that Equation 2 holds under an assumption of iid returns, which means that we can get a consistent estimate of the long-run risk by aggregating short-run estimates. Lo and Mackinlay (1988) use this relation to create a variance-ratio test by comparing the long-run variance against the scaled high-frequency variance as

$$J_r(h) = \frac{\mathrm{Var}\left[r^{(h)}\right]}{k\mathrm{Var}\left[r^{(1)}\right]} - 1. \tag{7}$$

We now want to generalize this test to also include intertemporal dependencies of higher moments, e.g., skewness and kurtosis. However, there are two issues. First, we cannot derive an asymptotic result of higher moments without making some distributional assumption of the returns; different distributions can yield the same skewness and kurtosis, e.g., generalized skewed t-distribution vs. a normal mixture distribution. Since the goal of this paper is to characterize the return distribution, we do not want to set some arbitrary restrictions on plausible distributions of the returns. Second, the skewness and excess kurtosis can be zero, which might create a division by zero.

However, we can create a non-parametric test in a similar vein, because under the null of iid returns, we know that

$$\mathrm{Var}\left[r^{(h)}\right] - h\mathrm{Var}\left[r^{(1)}\right] = 0, \quad \text{“Variance-ratio”} \tag{8}$$

$$\mathrm{SK}\left[r^{(h)}\right] - \frac{1}{\sqrt{h}}\mathrm{SK}\left[r^{(1)}\right] = 0, \quad \text{“Skewness-ratio”} \tag{9}$$

$$\mathrm{KU}^e\left[r^{(h)}\right] - \frac{1}{k}\mathrm{KU}^e\left[r^{(1)}\right] = 0, \quad \text{“Kurtosis-ratio”.} \tag{10}$$

As the test statistics will be estimated under the null of iid returns, we can estimate them by bootstrap, which does not require any distributional assumptions. Hence, to keep the number of assumptions to a minimum, we estimate the test statistics by accelerated bias-adjusted bootstrap.

A rejection of the test implies that the returns are not iid. If we assume that returns are well-behaved;[11] covariance stationary, defined higher moments, and ergodic, a rejection of a specific moment-ratio test is the same thing as a non-zero co-moment. More specifically,

$$\begin{aligned} 2\mathrm{Cov}\left[X, Y\right] \neq 0, \\ 3\sigma_x^2\sigma_y\mathrm{SK}\left[X, X, Y\right] + 3\sigma_x\sigma_y^2\mathrm{SK}\left[X, Y, Y\right] \neq 0, \\ 4\sigma_x^3\sigma_y\mathrm{KU}\left[X, X, X, Y\right] + 6\sigma_x^2\sigma_y^2\mathrm{KU}\left[X, X, Y, Y\right] + 4\sigma_x\sigma_y^3\mathrm{KU}\left[X, Y, Y, Y\right] \neq 0. \end{aligned}$$

Consequently, if we can assume that returns are well-behaved, the test can be used to disentangle which type of intertemporal dependency we observe in different assets and portfolios.

2 Empirical Results

2.a Data Sources

All estimates are based on continuously compounded daily returns of the factor mimicking portfolios, market (MKT), size (SMB), value (HML), and momentum (MOM) downloaded from AQR Capital Management (AQR 2018) and Kenneth French homepage (French 2018). The portfolio construction follows Fama and French (1992, 1993, 1996), Asness, Moskowitz, and Pedersen (2013), and Frazzini and Pedersen (2014). The SMB and HML portfolios are re-sorted in July using the break-points from the end of the previous year, while the MOM portfolio is re-sorted daily.

[11] In this paper we will assume that returns are well-behaved. However, there is probably a related paper that could go down the other rabbit hole, similar to Carr and Wu (2003), by assuming that the basic assumptions of time-series inference breaks down.

2.b Term-Structure in the Distribution of Returns

Engle and Mistry (2014) demonstrates that there exists a term structure in the higher moments in the factor mimicking portfolios. Following their setup, we estimate the unconditional volatility, skewness, and excess kurtosis as[12]

$$\hat{\sigma} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(r_i - \bar{r})^2}, \tag{11}$$

$$\widehat{SK} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{r_i - \bar{r}}{\hat{\sigma}}\right)^3, \tag{12}$$

$$\widehat{KU} = \frac{1}{N}\frac{\sum_{i=1}^{N}(r_i - \bar{r})^4}{\left(\hat{\sigma}^2\right)^2} - 3. \tag{13}$$

The daily excess returns are log-transformed and continuously compounded as in Equation 1 to get returns at different frequencies, from daily to 5-year returns.[13] Then for each observation frequency, the *overlapping* returns are used to get a consistent point estimate of the mean, volatility, skewness, and excess kurtosis of each factor mimicking portfolio, Table 1.

Since we do not know the distribution under the null, we cannot use the asymptotic 95% confidence interval of the skewness and kurtosis. Hence, all the 95% confidence intervals are estimated by accelerated bias-adjusted bootstrap from the *non-overlapping* returns.[14]

[12]The results has been re-produced using bias adjusted skewness and kurtosis estimators, but the results are indistinguishable from each other. And since most of the literature are using the "biased" estimators, we use the same estimators to keep it consistent with previous literature.

[13]Approximately 264 trading days a year. As a robustness check, we also reproduce the results for pre- and post-1963 samples in Appendix A.

[14]If we use the overlapping returns, the confidence interval would be underestimated, especially for more extended holding periods. When comparing the bootstrap intervals against the null of zero skewness and kurtosis, the bootstrap confidence intervals are a bit too wide. However, it still yields a reasonable confidence interval of the uncertainty of the point estimates without imposing any distributional restrictions.

Table 1: The point estimates of the first four centralized moments of each factor mimicking portfolio, from the overlapping continuously compounded returns at different holding periods. It also includes the five-year centralized moments under the assumption of iid daily or monthly returns.

Observation	MKT				SMB				HML				MOM			
Frequency	μ	σ	SK	KU^e	μ	σ	SK	KU^e	μ	σ	SK	KU^e	μ	σ	SK	KU^e
Daily (d)	0.0002	0.01	-0.42	17.66	0.0000	0.01	-0.99	24.72	0.0001	0.01	0.59	15.24	0.0002	0.01	-1.95	34.85
Weekly (w)	0.001	0.02	-0.70	9.04	0.0002	0.01	-0.33	10.95	0.001	0.01	0.60	11.56	0.001	0.02	-1.75	15.18
Monthly (m)	0.01	0.05	-0.71	7.63	0.001	0.03	0.21	7.61	0.003	0.03	0.68	6.00	0.01	0.04	-1.73	11.81
Quarterly (q)	0.02	0.10	-0.46	5.20	0.002	0.05	0.60	4.07	0.01	0.06	0.61	5.76	0.02	0.08	-1.86	12.78
Bi-Annual (b)	0.03	0.14	-0.72	2.27	0.005	0.08	0.28	1.64	0.02	0.09	0.41	2.91	0.03	0.11	-1.97	11.55
Annual (y)	0.06	0.20	-0.97	2.53	0.01	0.11	-0.15	1.23	0.04	0.13	0.10	1.17	0.06	0.15	-1.68	6.57
Two-Year (2y)	0.12	0.29	-1.25	3.11	0.02	0.17	-0.01	0.22	0.08	0.17	-0.06	0.67	0.12	0.22	-1.12	2.25
Three-Year (3y)	0.18	0.34	-1.45	4.25	0.03	0.22	0.15	-0.31	0.12	0.19	0.10	-0.22	0.18	0.27	-0.78	1.01
Four-Year (4y)	0.23	0.39	-1.08	2.85	0.04	0.26	0.14	-0.44	0.16	0.20	0.30	0.06	0.24	0.32	-0.73	0.55
Five-Year (5y)	0.29	0.42	-0.69	0.99	0.06	0.29	0.12	-0.57	0.20	0.23	0.16	-0.21	0.30	0.35	-0.75	0.40
5y (daily iid)	0.31	0.39	-0.01	0.01	0.04	0.21	-0.03	0.02	0.19	0.21	0.02	0.01	0.32	0.27	-0.05	0.03
5y (monthly iid)	0.31	0.42	-0.09	0.13	0.04	0.23	0.03	0.13	0.19	0.26	0.09	0.10	0.31	0.35	-0.22	0.20

2.b.1 Term Structure in the Tail-Risk

Under an assumption of iid returns, we expect the skewness and excess kurtosis to converge towards zero, towards a normal distribution, as the holding period increases. However, empirically, this is not the case. Up to a five-year horizon, both the skewness and excess kurtosis are significantly different from zero in the MKT portfolio, Figure 1. At even longer horizons, the point estimate continues to be different from zero, but the null can no longer be rejected due to the high uncertainty of the point estimates. Hence, we reject the common assumption that equity returns are normally distributed in the long-run, or at least up to a 5-year holding period.

There is also a term structure in the higher moments, especially in the skewness. For example, as the returns are compounded to a longer horizon, the skewness can both increase or decrease depending on the horizon. Over some horizons, the skewness can even diverge from zero, which is in direct contradiction to what is expected from iid returns. In Section 2.c, we show that a significant and huge intertemporal dependency drives these unexpected patterns in returns.

2.b.2 Tail-Risk Heterogeneity Among Factor Mimicking Portfolios

A common assumption in asset pricing is that there exist latent systematic risk factors, which capture different types of risk in the economy. As the centralized moments are constant under affine transformations, it implies that any asset that loads on a specific risk factor also inherits the skewness and kurtosis of that particular risk factor. Moreover, if risk factors capture different types of systematic risk, they should also exhibit different return distributions. Consequently, we want to empirically compare the return distributions, the skewness and kurtosis, of different factor mimicking portfolios.

Figure 2 suggests that different factor mimicking portfolios exhibit different term structures in skewness and kurtosis, at least up to an annual observation frequency. For example, while the HML portfolio has positive skewness no matter the holding period, the MOM portfolio has negative

Figure 1: The black line is the point estimate of skewness and excess kurtosis from the *overlapping* continuously compounded returns at different holding periods, from daily to five-years. The shaded area is the 95% accelerated bias-adjusted bootstrap confidence interval estimated from the *non-overlapping* continuously compounded return.

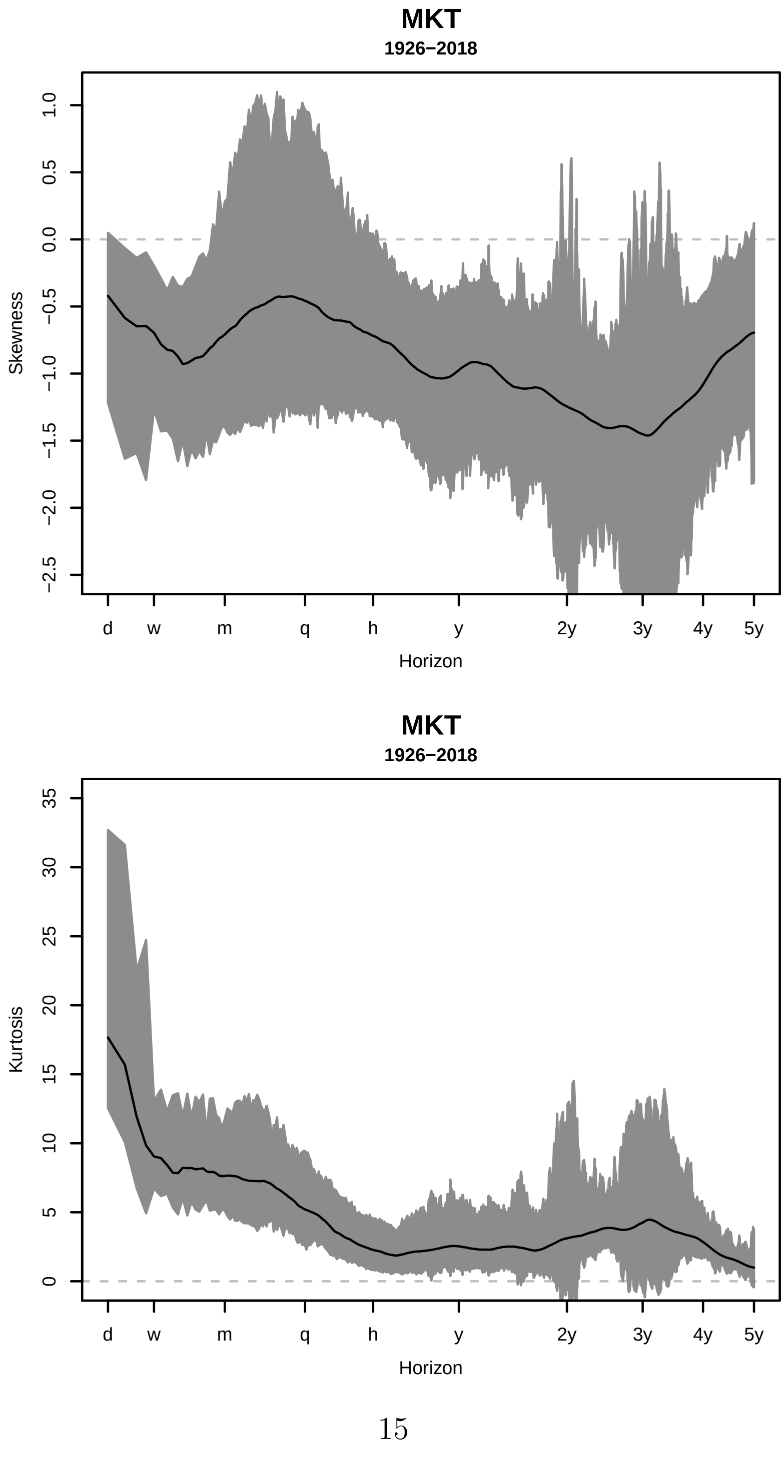

skewness no matter the holding period.[15] In contrast, the MKT and SMB portfolios have increasing skewness over certain holding periods, and decreasing skewness over others.

The observed heterogeneity between different risk factors demonstrates that the tail-risk is portfolio specific — any asset or portfolio that loads on a particular risk factor will also inherit the same tail-risk. Moreover, the tail-risk exhibits a term structure that is different between different factor mimicking portfolios, which indicate that the tail-risk is both portfolio *and* horizon specific. Consequently, we cannot generalize a short-term risk estimate to a longer horizon, without making some strong assumptions about the term structure in the particular portfolio. Hence, we expect this result to have direct implications for risk management, asset pricing, and portfolio optimization.

2.c Moment-Ratio Test - Application

As explained in Section 1.c, we generalize the variance-ratio test, which allows us to test for intertemporal dependencies in returns, Figure 3. First, no matter which moment-ratio test, horizon, or portfolio we chose, we still reject the null of iid returns at a 5% significance level. Equity returns are *not* iid. Second, different portfolios exhibit different types of intertemporal dependencies. For example, while the MKT and MOM portfolios have positive coskewness at short holding periods, the SMB and HML have negative coskewness at short horizons. Finally, with few exceptions, most intertemporal dependencies reverse sign at longer horizons. There seem to exist reversals, not just in expected returns, but in all higher moments.

2.c.1 Significant Intertemporal Dependency - Economic Implications

It is well established that returns are positively correlated over short horizons, and negatively correlated over longer horizons, i.e., momentum in the

[15] The MOM portfolio stands out, it consistently has lower skewness and higher excess kurtosis than the other factor mimicking portfolios, which confirms the momentum crash risk observed in Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016). See Appendix B for an in-depth discussion.

Figure 2: The black line is the point estimate of skewness and excess kurtosis from the *overlapping* continuously compounded returns at different holding periods, from daily up to a five-year horizon. The shaded area is the 95% accelerated bias-adjusted bootstrap confidence interval estimated from the *non-overlapping* continuously compounded return.

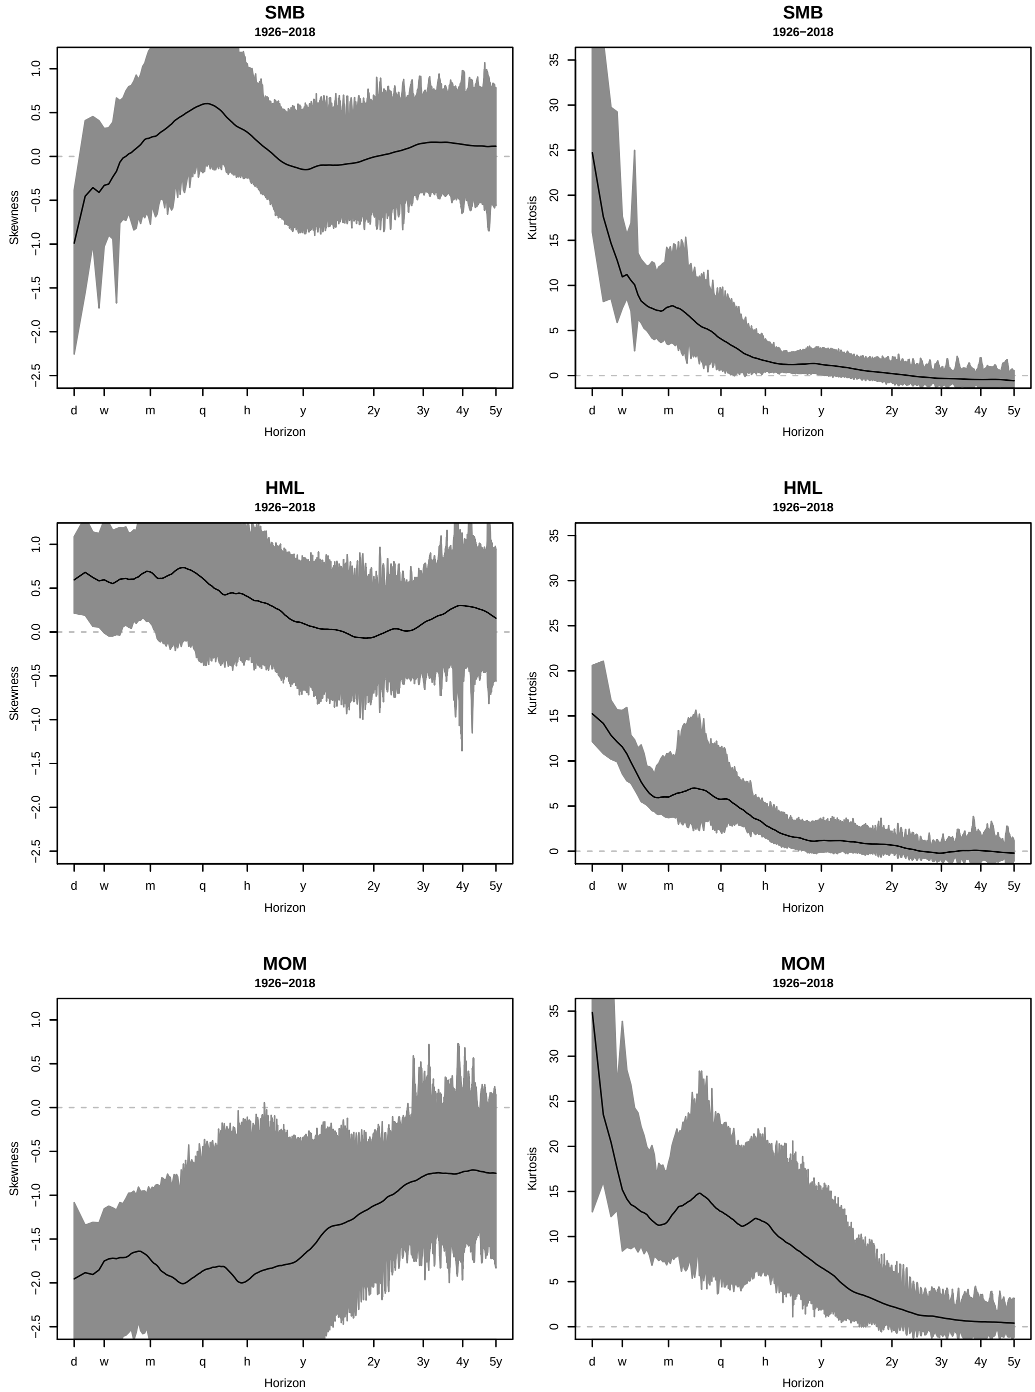

Figure 3: The accelerated bias-adjusted bootstrap estimate of the difference between the scaled short-run estimate and the long-run point estimate. Under the null hypothesis, $H_0 =$ iid returns, the difference should be equal to zero. The x-axis is the observation frequency, from daily (d) up to three years (3), where each marker indicates a different number of compounded returns of that specific observation frequency. Black markers indicate a significant difference at 5% significance level, a rejection of the null of iid returns of that specific observation frequency. The sign of the difference also implies the sign of the co-moment, e.g., a positive difference in the "Variance-Ratio" implies that $\text{Cov}\left[r_t, r_{t-1}\right] > 0$.

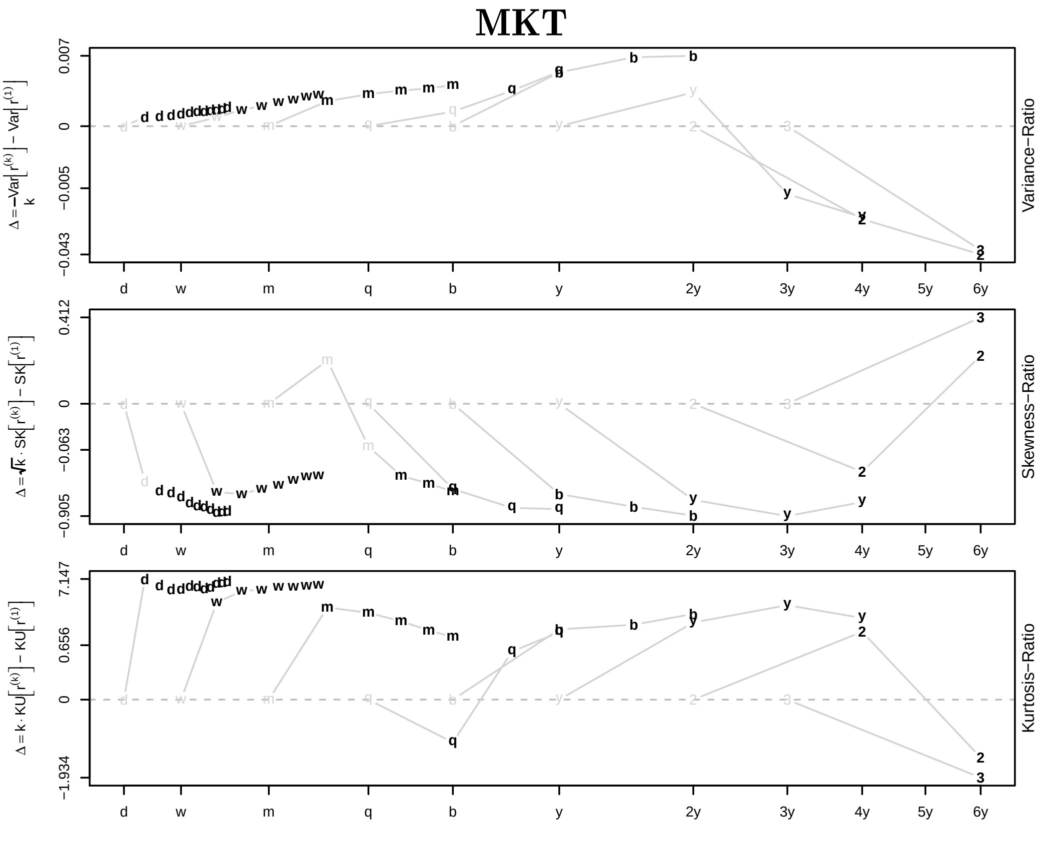

SMB

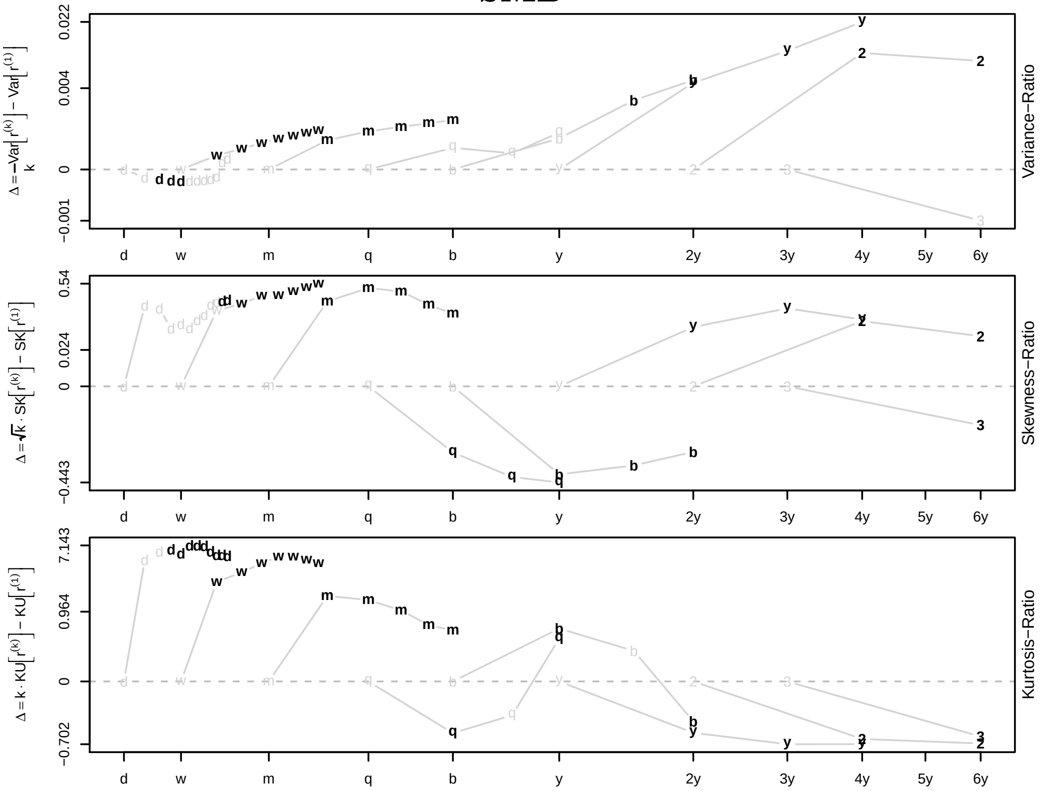
Variance−Ratio
Skewness−Ratio
Kurtosis−Ratio
d w m q b y 2y 3y 4y 5y 6y

HML

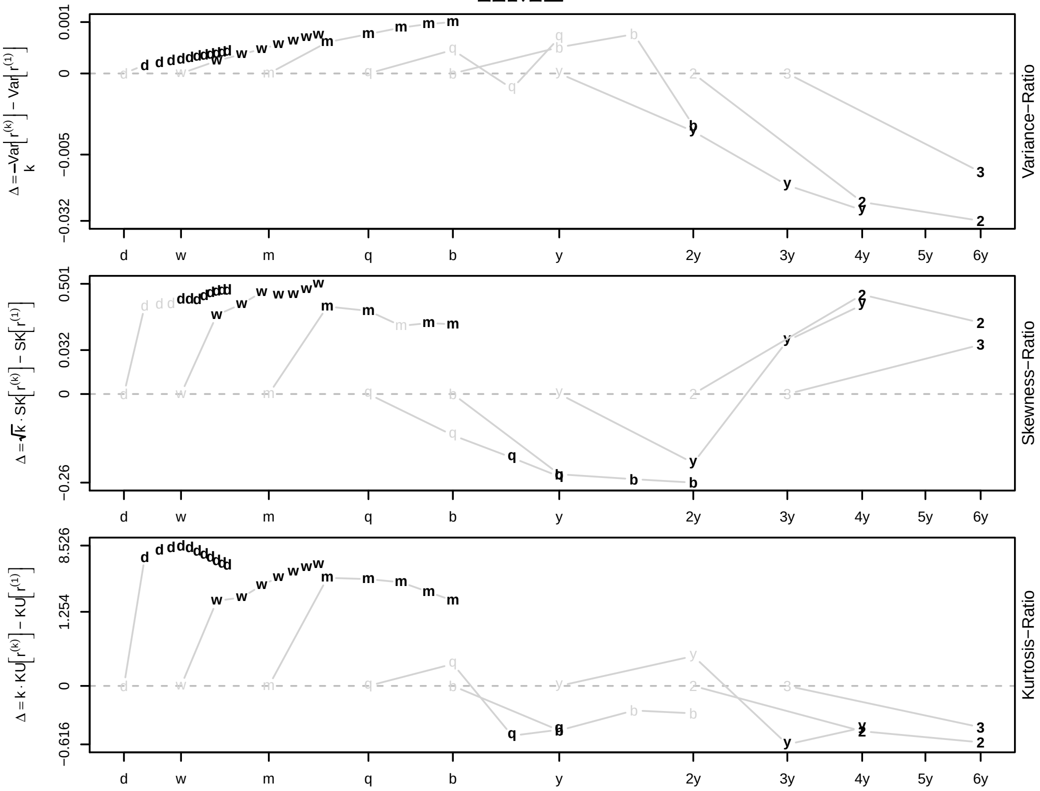
Variance−Ratio
Skewness−Ratio
Kurtosis−Ratio
d w m q b y 2y 3y 4y 5y 6y

(Figure Continued)

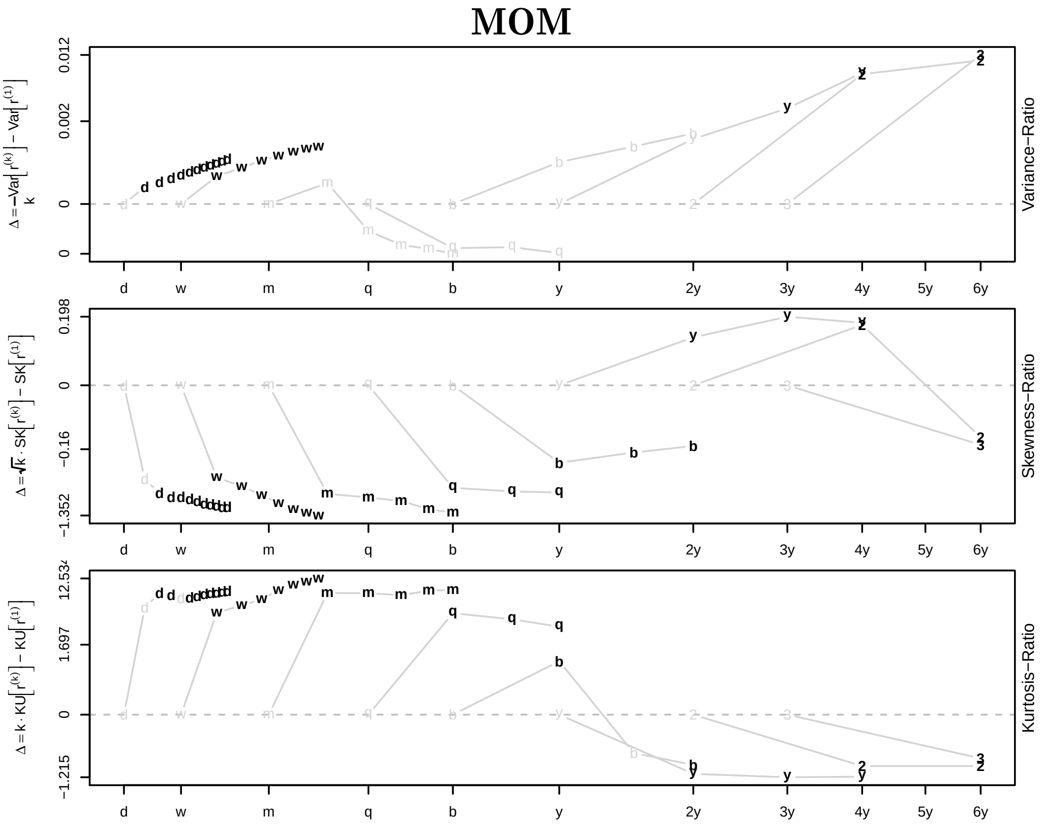

short-run and mean-reversal in the long-run.[16] Our test confirms this result; the MKT and HML portfolios have significant and positive covariance at short horizons and significantly negative covariance at long horizons. However, both the SMB and MOM portfolios exhibit positive covariance even at long horizons. Since the MOM portfolio is re-balanced daily and is supposed to capture the trend, we would expect it to have positive covariance. However, it is harder to explain why the SMB portfolio also continues to exhibit momentum at more extended holding periods.

The literature has also established that the equity returns should exhibit a leverage effect; higher expected risk tomorrow decreases the returns today. Hence, the leverage effect implies that $\text{Cov}\left[r_t^2, r_{t-1}\right] < 0$, which should induce a negative coskewness. Consequently, if investors are risk-averse and forward-looking, we would expect a negative coskewness, which is exactly what we observe in the MKT and MOM portfolios over short horizons. However, in the SMB and HML portfolio, we observe the exact opposite; higher volatility tomorrow is associated with higher returns today, which is in direct contrast to the leverage effect.

There are three possible explanations for this counterintuitive result. First, the co-skewness has two terms, where one is the leverage effect, and the other is how volatility today affects prices tomorrow. Hence, a positive coskewness might also imply that an unusually large shock to returns today would increase the expected returns tomorrow.[17] Second, the leverage effect is derived under the classic mean-variance trade-off, which ignores higher moments. Therefore, it is plausible that we can get "inconsistent" patterns in the mean-variance trade-off if investors also care about higher moments. Finally, in this paper, we assume that returns are well-behaved, which implies that a rejection of the null hypothesis must be due to a non-zero co-moment. However, if returns are non-stationary, non-ergodic, or have undefined higher moments, we would still reject the null hypothesis, but it would imply that

[16] Stoll and Whaley (1990), Ding, Granger, and Engle (1993), Jegadeesh and Titman (1993), Lewellen (2002), and Jiang and Tian (2005)

[17] When estimating each co-moment separately, it is not clear that we can ignore one of the terms. Hence, it is not obvious that we can ignore the second coskewness term, $\text{Cov}\left[r_t, r_{t-1}^2\right]$, as in Engle and Mistry (2014) and Neuberger and Payne (2019).

our current interpretation is faulty.

To complicate the interpretation further, the coskewness switches sign at longer horizons, creating a reversal pattern in skewness similar to the mean-reversal in the variance-ratio test. However, in contrast to the variance-ratio, the skewness-ratio shows no general pattern. At the same horizon, some portfolios will exhibit positive coskewness, while others will have negative coskewness. Consequently, we cannot generalize a short-run skewness estimate to a longer horizon without making a strong assumption about the intertemporal dependency in the returns of our specific portfolio.

In contrast, it seems like all equity portfolios have a positive cokurtosis at shorter horizons, which implies that the probability of outliers is higher than we would expect from an assumption of iid returns. There is also an extensive literature that has demonstrated that variance is persistent,[18] which implies that $\text{Cov}\left[r_t^2, r_{t-1}^2\right] > 0$. Hence, if there is persistence in variance, it should inflate the kurtosis, which is precisely what we observe in our tests, at least over short horizons. Similarly to how there is a reversal in variance and skewness, the kurtosis also “reverses” in the long-run.

In summary, equity returns are strongly dependent over time;[19] we reject the null of iid returns up to a three-year holding period for all factor mimicking portfolios.[20] The MKT returns exhibit positive covariance, negative coskewness, and positive cokurtosis, up to a quarterly observation frequency. At longer horizons, all the co-moments switch sign, similar to a "mean-reversal" in all higher moments. Finally, there exists heterogeneity in the co-moments between the different factor mimicking portfolios, some demonstrate positive coskewness in the short run, SMB and HML, while others have negative coskewness, MKT and HML.

[18]Bollerslev (1986), Nelson (1991), Ding, Granger, and Engle (1993), Zakoian (1994), Harvey and Siddique (2000), and Hansen and Lunde (2005)

[19]For a discussion of the relative size between the co-moments and the individual centralized moments, see Appendix C.

[20]There seem to exist intertemporal dependencies at even longer horizons. However, some of the estimators became unstable at horizons of six years or longer.

3 Conclusion

Given the importance of the distribution of short- vs. long-horizon returns, we wonder, how are the equity returns distributed over different horizons? We make three contributions. First, we establish that there exists a term structure in the higher moments, which is specific to each factor mimicking portfolio. Since the skewness and kurtosis are constant under an affine transformation, it implies that any asset that loads on a particular risk factor will also inherit the tail-risk of that risk factor. Consequently, we cannot make a general statement about the long-run tail-risk of a specific portfolio without knowing which specific risk factors it is exposed to.

Second, to explain the observed term structure in the tail-risk, we assume that there exists an intertemporal dependence in returns. To test this, we generalize the variance-ratio test of Lo and Mackinlay (1988) to include higher moments. By applying this new test, we can reject the assumption of independent and identically distributed (iid) returns for all portfolios and all observation frequencies. Returns are *not* iid.

Moreover, similar to how there exists a mean-reversal in returns, there is also a significant reversal in skewness and kurtosis at longer holding periods. It even seems like the coskewness is oscillating between positive and negative in the SMB, HML, and MOM portfolio as the horizon expands. Even more surprising is that the coskewness can be positive, which implies that higher volatility tomorrow increases returns today, which is the direct opposite of the leverage effect. A plausible reason is that this "risk-return trade-off" is derived in a mean-variance setting, and if investors also care about higher moments, it might show up as an inconsistent risk-return trade-off.

Finally, we show that the conventional approach of extrapolating a short-run estimate to a longer horizon by assuming iid returns creates a biased estimate of the long-run risk. Even if the bias is lessened by considering a simple dependency structure, e.g., an autoregressive process, it will still yield a bias in many situations since the dependency can switch signs. Consequently, ignoring the intertemporal dependence in returns will create a severely biased long-run estimate of the risk exposure. Hence, we expect our results to have a direct impact on asset prices, risk management, and optimal portfolio allocation.

References

Albuquerque, Rui. 2012. "Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns." *Review of Financial Studies* 25 (5): 1630–1673.

Amaya, Diego, et al. 2015. "Does realized skewness predict the cross-section of equity returns?" *Journal of Financial Economics* 118 (1): 135–167.

AQR, Capital Management. 2018. "Data Sets - AQR Capital Management." Visited on 04/05/2018. https://www.aqr.com/Insights/Datasets.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen. 2013. "Value and Momentum Everywhere." *Journal of Finance* 68 (3): 929–985.

Bali, Turan G., and Scott Murray. 2013. "Does Risk-Neutral Skewness Predict the Cross-Section of Equity Option Portfolio Returns?" *Journal of Financial and Quantitative Analysis* 48 (04): 1145–1171.

Barberis et al. 2008. "Stocks as Lotteries: The Implication of Probability Weighting for Security Prices." *American Economic Review*: 2066–2100.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century*." *Quarterly Journal of Economics* 121 (3): 823–866.

Barroso, Pedro, and Pedro Santa-Clara. 2015. "Momentum has its moments." *Journal of Financial Economics* 116 (1): 111–120.

Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroscedasticity." *Journal of Econometrics* 31:307–327.

Boyer, Brian, Todd Mitton, and Keith Vorkink. 2010. "Expected idiosyncratic skewness." *Review of Financial Studies* 23 (1): 170–202.

Campbell, John Y, and Luis M Viceira. 2005. "The Term Structure of the Risk-Return Trade-Off." *Financial Analysts Journal* 61 (1): 34–44.

Carr, Peter, and Liuren Wu. 2003. "The Finite Moment Log Stable Process and Option Pricing." *Journal of Finance* 58 (2): 753–777.

Conrad, Jennifer, Robert F Dittmar, and Eric Ghysels. 2013. "Ex Ante Skewness and Expected Stock Returns." *The Journal of Finance* LXVIII (1): 85–124.

Dahlquist, Magnus, Adam Farago, and Romeo Tedongap. 2017. "Asymmetries and portfolio choice." *Review of Financial Studies* 30 (2): 667–702.

Daniel, Kent, and Tobias J. Moskowitz. 2016. "Momentum crashes." *Journal of Financial Economics* 122 (2): 221–247.

De Bondt, Werner F. M., and Richard H. Thaler. 1986. "Further Evidence on Investor Overreaction and Stock Market Seasonality." *The Journal of Finance* 42 (3): 557–581.

Ding, Zhuanxin, Clive W.J. Granger, and Robert F. Engle. 1993. "A long memory property of stock market returns and a new model." *Journal of Empirical Finance* 1 (1): 83–106.

Engle, Robert, and Abhishek Mistry. 2014. "Priced risk and asymmetric volatility in the cross section of skewness." *Journal of Econometrics* 182 (1): 135–144.

Fama, Eugene F., and Kenneth R. French. 1993. "Common risk factors in the returns on stocks and bonds." *Journal of Financial Economics* 33 (1): 3–56.

— . 2018. "Long-Horizon Returns." *The Review of Asset Pricing Studies* 8 (2): 232–252.

— . 1996. "Multifactor Explanations of Asset Pricing Anomalies." *Journal of Finance* 51 (1): 55–84.

— . 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance* 47 (2): 427–465.

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

French, Kenneth R. 2018. "Kenneth R. French - Data Library." Visited on 03/20/2018. http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Ghysels, Eric, Alberto Plazzi, and Rossen Valkanov. 2016. "Why Invest in Emerging Markets? The Role of Conditional Return Asymmetry." *Journal of Finance* 71 (5): 2145–2192.

Hansen, Peter R., and Asger Lunde. 2005. "A forecast comparison of volatility models: Does anything beat a GARCH(1,1)?" *Journal of Applied Econometrics* 20 (7): 873–889.

Harvey, Campbell R., and Akhtar Siddique. 2000. "Conditional skewness in asset pricing tests." *Journal of Finance* 55 (3): 1263–1295.

Jegadeesh, Narasimhan, and Sheridan Titman. 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance* 48 (1): 65–91.

Jiang, George J., and Yisong S. Tian. 2005. "The model-free implied volatility and its information content." *Review of Financial Studies* 18 (4): 1305–1342.

Jondeau, Eric, Ser-Huang Poon, and Michael Rockinger. 2008. *Financial Modeling Under Non-Gaussian Distributions.*

Kelly, Bryan, and Hao Jiang. 2014. "Tail risk and asset prices." *Review of Financial Studies* 27 (10): 2841–2871.

Kraus, Alan, and Robert H. Litzenberger. 1976. "Skewness Preference and the Valuation of Risk Assets." *The Journal of Finance* 31 (4): 1085–1100.

Lewellen, Jonathan. 2002. "Momentum and Autocorrelation in Stock Returns." *Review of Financial Studies* 15 (2): 533–564.

Lo, Andrew W, and A Craig Mackinlay. 1988. "Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification test." *The Review of Financial Studies* 1 (1): 41–66.

Lucca, David O., and Emanuel Moench. 2015. "The Pre-FOMC announcement drift." *Journal of Finance* 70 (1): 329–371.

Mitton, T, and K Vorkink. 2007. "Equilibrium underdiversitication and the preference for skewness." *Review of Financial Studies* 20 (4): 1255–1288.

Nelson, Daniel B. 1991. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica* 59 (2): 347–370.

Neuberger, Anthony. 2012. "Realized skewness." *Review of Financial Studies* 25 (11): 3424–3455.

Neuberger, Anthony, and Richard G. Payne. 2019. "The Skewness of the Stock Market at Long Horizons."

Pástor, Ľuboš, and Robert F. Stambaugh. 2012. "Are Stocks Really Less Volatile in the Long Run?" *Journal of Finance* 67 (2): 431–478.

Rehman, Zahid, and Grigory Vilkov. 2012. "Risk-Neutral Skewness: Return Predictability and Its Sources." *SSRN Electronic Journal*, no. November 2008.

Rietz, Thomas A. 1988. "The equity risk premium a solution." *Journal of Monetary Economics* 22 (1): 117–131.

Schneider, Paul, Christian Wagner, and Josef Zechner. 2019. "Low Risk Anomalies?"

Sias, Richard. 2007. "Causes and seasonality of momentum profits." *Financial Analysts Journal* 63 (2): 48–54.

Stoll, Hans R., and Robert E. Whaley. 1990. "The Dynamics of Stock Index and Stock Index Futures." *Journal of Financial & Quantitative Analysis* 25 (4): 441–468.

Umar, Tarik. 2017. "Complexity Aversion When Seeking Alpha."

Wang, Jying Nan, Jin Huei Yeh, and Nick Ying Pin Cheng. 2011. "How accurate is the square-root-of-time rule in scaling tail risk: A global study." *Journal of Banking and Finance* 35 (5): 1158–1169.

Zakoian, Jean Michel. 1994. "Threshold heteroskedastic models." *Journal of Economic Dynamics and Control* 18 (5): 931–955.

A Split sample

As a robustness check, we split the sample into pre-1963 and post-1963 samples. The patterns in skewness and kurtosis are very similar between the full sample and the split sample, Figure 6 and Figure 7. The MOM portfolio has large negative skewness no matter the holding period, while the HML portfolio is slightly positive no matter the holding period. Moreover, the SMB portfolio skewness increases as we move from daily to quarterly observation frequency, and then flattens out. However, the MKT portfolio in the post-1963 sample has a different pattern in skewness compared to the full sample. There is no noticeable decrease in skewness as we go from daily to weekly observations, and the hump in skewness around quarterly returns are also gone.

Overall, the patterns in skewness and kurtosis seem to be robust between different periods, and using the argument made by Albuquerque (2012), it suggest that the seasonality patterns of market events is similar both before and after 1963. However, there seem to be a change in how the MKT portfolio behaves, which might indicate that the systematic risk might have changed over time. Another possible explanation is a inefficiency story, where the discovery of the CAPM in the early 1960s removed this pricing inefficiency, which changes the unconditional distribution in the post 1963 sample. While the discovery by Fama and French (1992) and Jegadeesh and Titman (1993) is too new to have had a real impact on the unconditional estimates so far.

B Momentum Crash Risk

We observe a very distinct pattern in skewness and kurtosis of the MOM portfolio compared to the other factor mimicking portfolios, Figure 8. The momentum strategy is more exposed to adverse tail events, more negative skewness and higher excess kurtosis, no matter which time frame or market we look at. Consequently, the momentum strategy is more exposed to massive losses, which is the same conclusion as in Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016). However, our results expand this

Figure 6: Split-sample estimate pre-1963. The black line is the point estimate of skewness and excess kurtosis from the *overlapping* continuously compounded returns at different holding periods (number of trading days). The shaded area is the 95% accelerated bias adjusted bootstrap confidence interval estimated from the *non-overlapping* continuously compounded return.

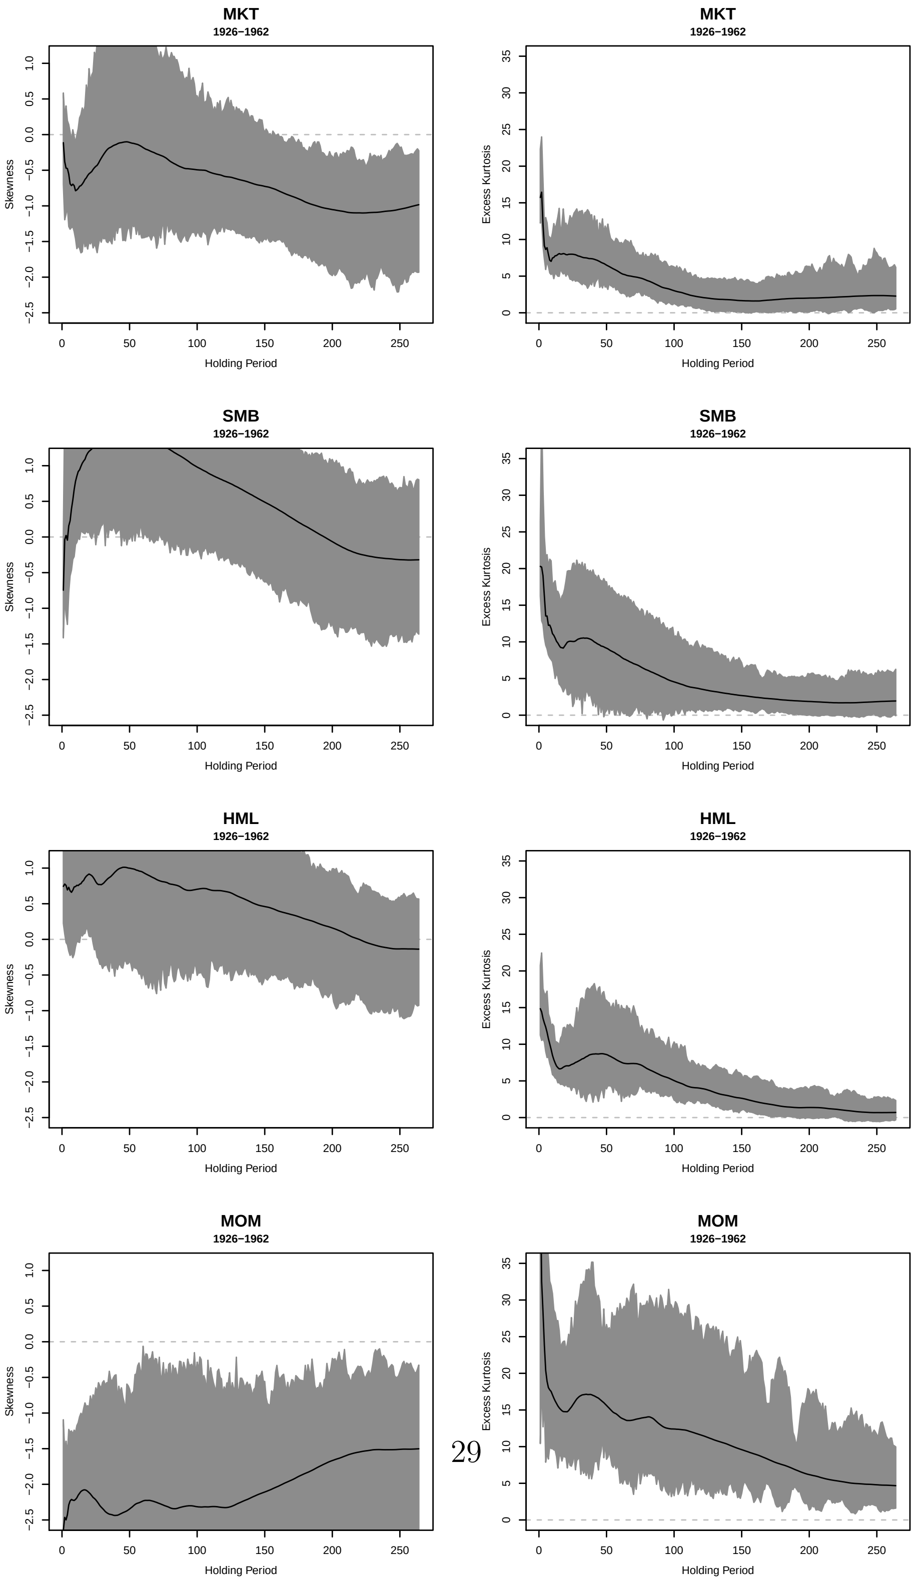

Figure 7: Split-sample estimate post-1963. The black line is the point estimate of skewness and excess kurtosis from the *overlapping* continuously compounded returns at different holding periods (number of trading days). The shaded area is the 95% accelerated bias adjusted bootstrap confidence interval estimated from the *non-overlapping* continuously compounded return.

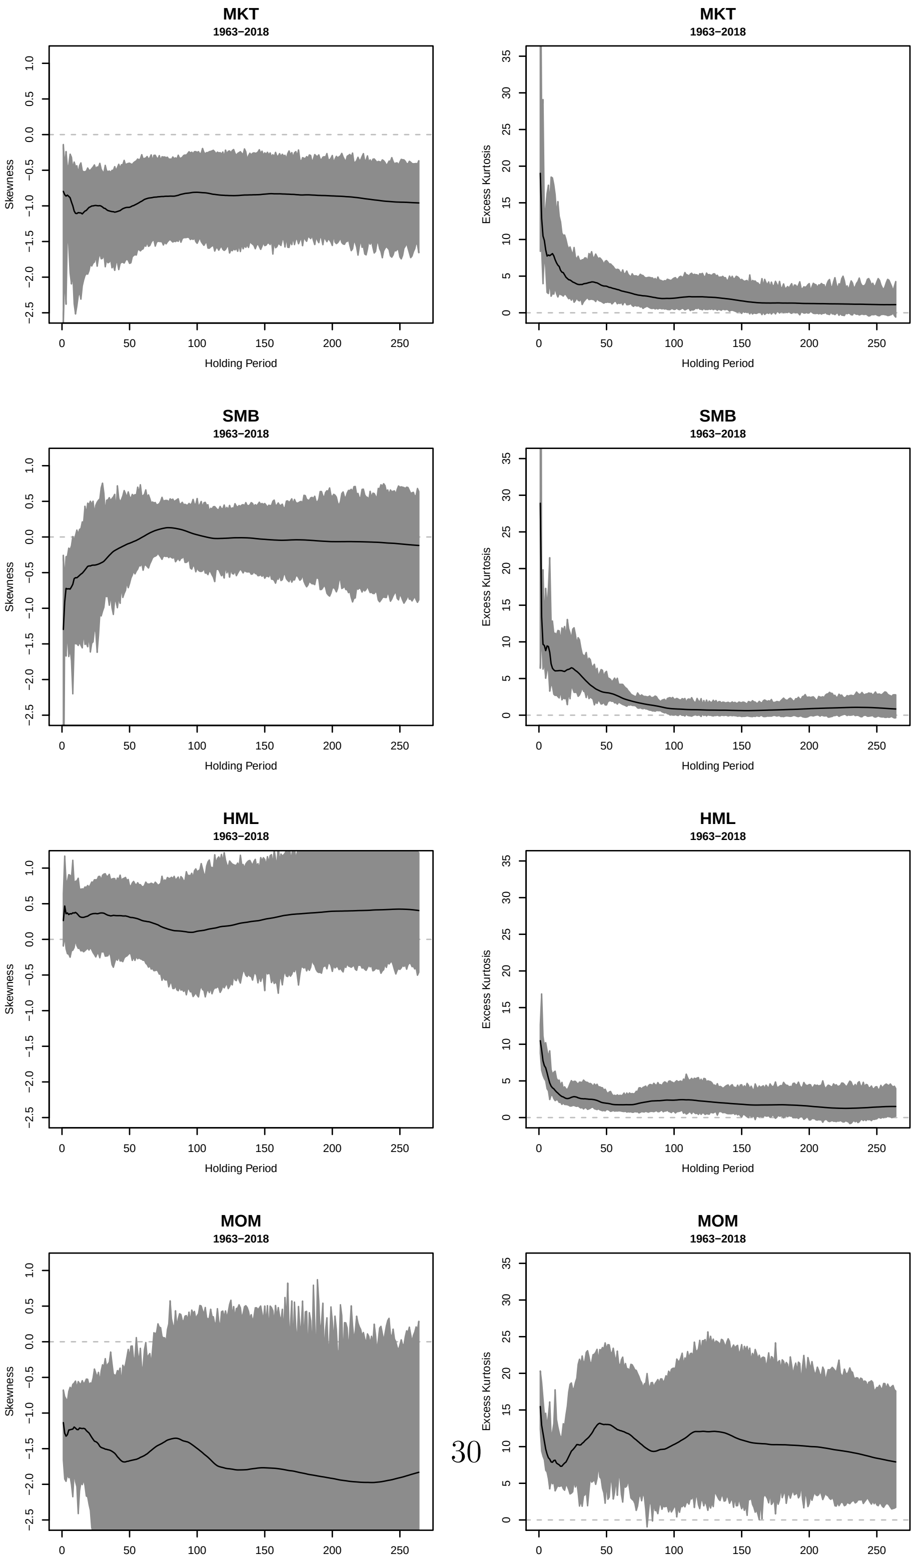

notion of momentum crash risk, by establishing that the momentum strategy is more exposed to crashes no matter which market or holding period we look at.

C Economic Size of the Co-Moments

C.a The M_3 and M_4 Matrices

To simplify calculations, we can re-arrange the 3-dimensional coskewness array into a (n, n^2) matrix as described by Jondeau, Poon, and Rockinger (2008)

$$M_3 = \mathbb{E}\left[(r-\mu)(r-\mu)' \otimes (r-\mu)'\right] = \{s_{ijk}\} \tag{14}$$

$$s_{ijk} = \mathbb{E}\left[(r_i-\mu_i)(r_j-\mu_j)(r_k-\mu_k)\right], \quad \text{for } i, j, k = 1, \ldots, n, \tag{15}$$

where r_i denotes the individual asset return i, μ is the mean, and $\otimes$ is the Kronecker product. Similarly, we can also create the (n, n^3) cokurtosis matrix.

This allows us to decompose the portfolios skewness and kurtosis into its parts since the portfolio, the linear combination of returns, can be calculated as

$$\text{SK}\left[r_p\right] = \boldsymbol{w}' M_3 (\boldsymbol{w} \otimes \boldsymbol{w}) \tag{16}$$

$$\text{KU}\left[r_p\right] = \boldsymbol{w}' M_4 (\boldsymbol{w} \otimes \boldsymbol{w} \otimes \boldsymbol{w}) \tag{17}$$

where the "off-diagonal" terms of M_3 and M_4 will be equal to zero if the assets are independent. Hence, we can decompose the observed portfolio skewness and kurtosis into co-moments and individual skewnesses and kurtosises, and compare the size of each part.

C.b The Size of the Co-Moment

The intertemporal co-moments are very large over all observation frequencies, which implies a strong intertemporal dependency in skewness and kurtosis,

Figure 8: Estimated skewness and excess kurtosis from overlapping continuously compounded returns from weekly (w) up to an annual (y) observation frequency. The date range above each figure is the shortest time-series available. For example, in the "Pacific" figure, the MKT portfolio is observed from the beginning of November 1985 until the end of December 2017, while the QMJ is only observed from the beginning of July 1993 until the end of December 2017.

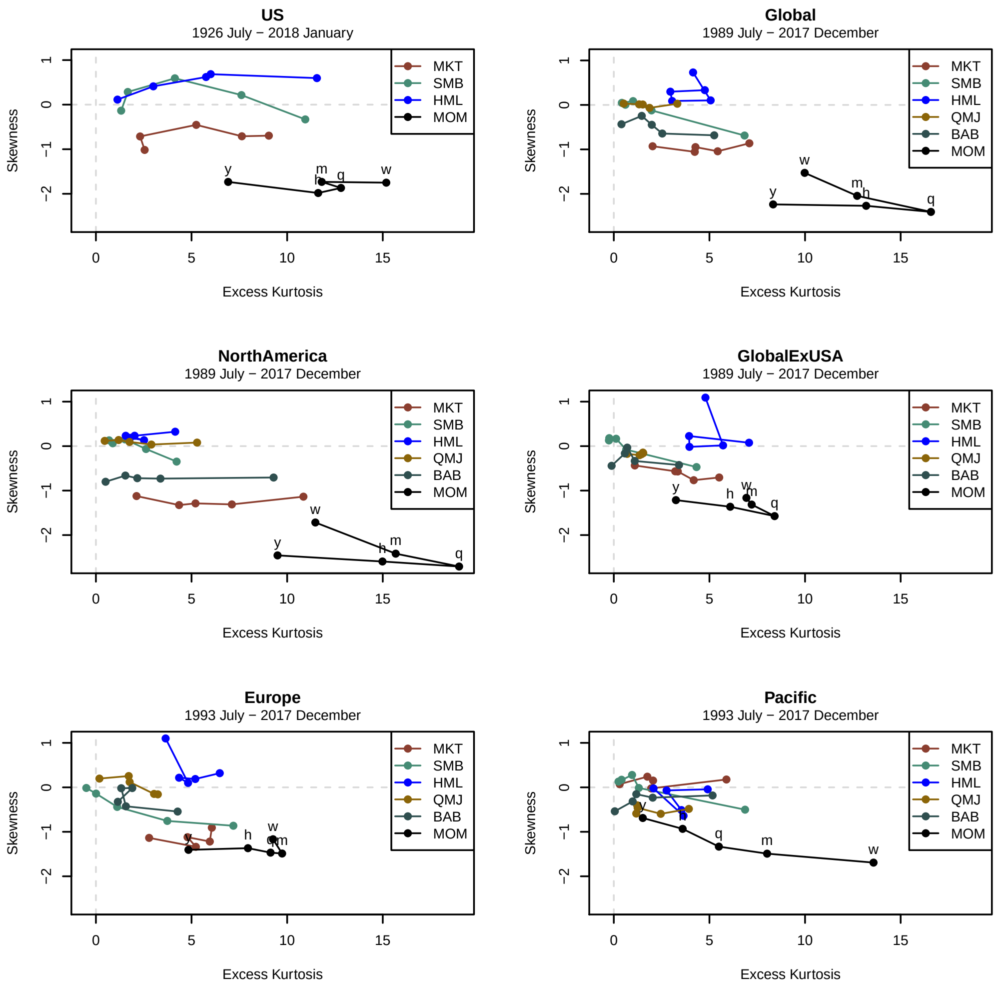

Figure 9. Just after a couple of continuously compounded periods, most of the skewness and kurtosis is explained by the intertemporal co-moments. Overall, the coskewness in MKT returns are negative, at least up to a bi-annual observation frequency, Figure 9. However, at longer horizons the coskewness seems to turn positive, which is confirmed by the skewness-ratio test.

In summary, most of the aggregation pattern in the expected tail-risk is due to the coskewness and cokurtosis. Consequently, we cannot make a general statement about the long-run tail-risk from a short-run estimate without making strong assumptions about the intertemporal dependency. Moreover, an assumption of independence will yield severely biased estimates of the tail-risk as the aggregation period increases.

D Related Literature

D.a Neuberger and Payne (2019)

Neuberger (2012) derives a realized third-moment estimator, which yields a more precise estimate of the skewness. He also establishes that the MKT returns do not follow an iid process. Neuberger and Payne (2019) extend this methodology to include an adjustment for the intertemporal dependency in the high-frequency returns, which yields a more precise estimate of the higher moments. However, their method creates a biased estimate, which suggests that it is the classic statistical trade-off between accuracy and precision.

A necessary assumption for the Neuberger and Payne (2019) estimator, is that the price process follows a strictly positive martingale process, which ensures that the process has the aggregation property. However, most of finance is interested in returns, not the prices. Therefore, to keep the aggregation property intact, they use a return approximation. Using the return

Figure 9: The black line is the point estimate of skewness, SK $\left[r^{(h)}\right] = \boldsymbol{w}'M_3(\boldsymbol{w}\otimes\boldsymbol{w})$, and kurtosis KU $\left[r^{(h)}\right] = \boldsymbol{w}'M_4(\boldsymbol{w}\otimes\boldsymbol{w}\otimes\boldsymbol{w})$ (kurtosis, not excess kurtosis), where $\boldsymbol{w}$ is a vector of ones of length h. The estimates are also decomposed into the expected skewness and kurtosis as if we would assume independent returns over time (black dashed line), and their co-moments (golden dashed line).

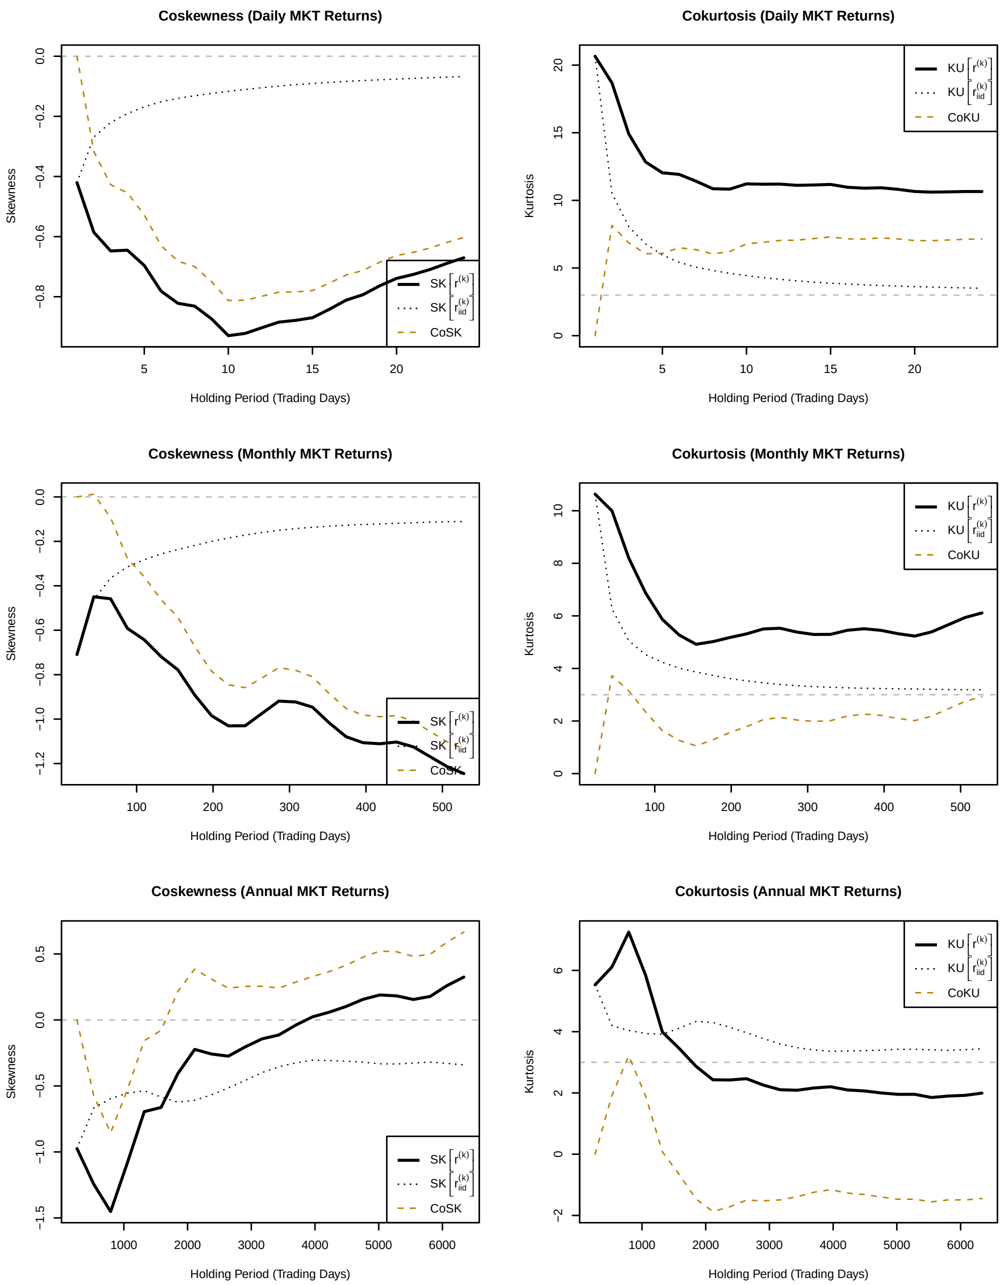

approximation, the centralized moments can be estimated as

$$\mathrm{Var}\left[R_t\right] = \mathbb{E}\left[x_t^{2,L}\right], \qquad x_t^{2,L} = 2\left(R_t - 1 - \ln(R_t)\right), \tag{18}$$

$$x_t^{2,E} = 2\left(R_t \ln(R_t) + 1 - R_t\right), \tag{19}$$

$$\mathrm{SK}\left[R_t\right] = \frac{\mathbb{E}\left[x_t^3\right]}{\mathrm{Var}\left[R_t\right]^{3/2}} \qquad x_t^3 = 6\left[(R_t+1)\ln(R_t) - 2(R_t - 1)\right], \tag{20}$$

$$\mathrm{KU}\left[R_t\right] = \frac{\mathbb{E}\left[x_t^4\right]}{\mathrm{Var}\left[R_t\right]^2} \qquad x_t^4 = 12\left[\ln(R_t)^2 + 2(R_t+2)\ln(R_t) - 6(R_t - 1)\right]. \tag{21}$$

As the estimates are based upon an approximation, there is some possible bias present already in this step.[21]

To aggregate the high frequency estimates to a longer horizon, two more estimates are needed

$$y_t = \frac{1}{h}\sum_{u=0}^{h-1}\left(\frac{P_t}{P_{t-u}} - 1\right), \qquad z_t = \frac{1}{h}\sum_{u=0}^{h-1} 2\left[\frac{P_t}{P_{t-u}} - 1 - \ln\left(\frac{P_t}{P_{t-u}}\right)\right] \tag{22}$$

where y_t is the harmonic mean of net returns, and z_t is the average of the realized variance. However, the way these estimators are constructed, these will depend on h. The first term is always equal to zero, which implies it is similar to a mean where we adjust by an extra degree of freedom $\sum_{i=1}^{N} x_i/(N+1)$.

The realized moments, adjusted for the intertemporal dependency, are given by

$$\mathrm{Var}\left[R_t^{(h)}\right] = h\mathrm{Var}\left[R_t\right], \tag{23}$$

$$\mathrm{SK}\left[R_t^{(h)}\right] = \frac{1}{\sqrt{h}}\left(\mathrm{SK}\left[R_t\right] + 3\frac{\mathrm{Cov}\left[y_{t-1}, x_t^{2,E}\right]}{\mathrm{Var}\left[R_t\right]^{3/2}}\right), \tag{24}$$

$$\mathrm{KU}\left[R_t^{(h)}\right] = \frac{1}{h}\left(\mathrm{KU}\left[R_t\right] + 4\frac{\mathrm{Cov}\left[y_{t-1}, x_t^3\right]}{\mathrm{Var}\left[R_t\right]^2} + 6\frac{\mathrm{Cov}\left[z_{t-1}, x_t^{2,L}\right]}{\mathrm{Var}\left[R_t\right]^2}\right). \tag{25}$$

[21] We performed some simple simulations, and the variance approximation is very accurate, while the higher moments approximations can exhibit large deviations.

To evaluate the performance of the Neuberger and Payne (2019) estimator, we simulate 10000 samples from a log-normal distribution, $\ln(R) \sim N(\mu, \sigma^2)$. Since a sum of normally distributed random variables are still normal, it is simple to derive the moment conditions for the compounded returns. The population moments of the compounded gross-returns, R, are

$$\text{Var}\left[R^{(h)}\right] = \left(e^{h\sigma^2} - 1\right) e^{2h\mu + h\sigma^2}, \tag{26}$$

$$\text{SK}\left[R^{(h)}\right] = \left(e^{h\sigma^2} + 2\right) \sqrt{e^{h\sigma^2} - 1}, \tag{27}$$

$$\text{KU}\left[R^{(h)}\right] = e^{4h\sigma^2} + 2e^{3h\sigma^2} + 3e^{2h\sigma^2} - 3. \tag{28}$$

We compare the Neuberger and Payne (2019) estimator (*NP2018*) against the centralized moment estimates, with non-overlapping (*standard*) and overlapping (*overlapping*) returns. From theory, we know that the *standard*, and *overlapping* approach both yields consistent point estimates, but the standard errors usually become biased when using overlapping observations.

Consistent with statistical theory, both the *standard* and the *overlapping* approach seem to yield a consistent estimate of the higher moments, the center mass of the empirical distributions are close to the true population parameters, Figure 10. In contrast, the *NP2018* estimator seems to be severely biased, no matter the sample size, moment, or aggregation horizon. The result suggests that the *NP2018* estimator suffers from the classic statistical trade-off between accuracy and precision.

Moreover, in contrast to Neuberger and Payne (2019), we use statistical cumulants and a simple bootstrap approach, which simplifies the inference while requiring fewer assumptions. Hence, by invoking Occam's razor, we argue that our approach is more straightforward while maintaining the inference. For example, their Martingale assumption implies that $\text{Cov}\left[r_t, r_{t-1}^2\right] = 0$ and $\text{Cov}\left[r_t, r_{t-1}^3\right] = 0$.[22] However, empirically those terms can be massive, which suggest that the Martingale assumption is inconsistent with the observed patterns in equity returns.

[22] Or at least that those are very small compared to the co-skewness and co-kurtosis terms that go in the other direction.

Figure 10: Sample distribution of $\text{Var}\left[R^{(h)}\right], \text{SK}\left[R^{(h)}\right],$ and $\text{KU}\left[R^{(h)}\right]$ from 10000 Monte Carlo simulations. The returns are generated from a log-normal distribution, $\ln(R) \sim N(\mu, \sigma^2)$, fitted to the daily and monthly excess returns of the MKT portfolio. Three different estimation methods are employed. Standard: Standard estimators with non-overlapping returns. NP2018: Neuberger and Payne (2019) estimators. Overlapping: Standard estimators with overlapping returns.

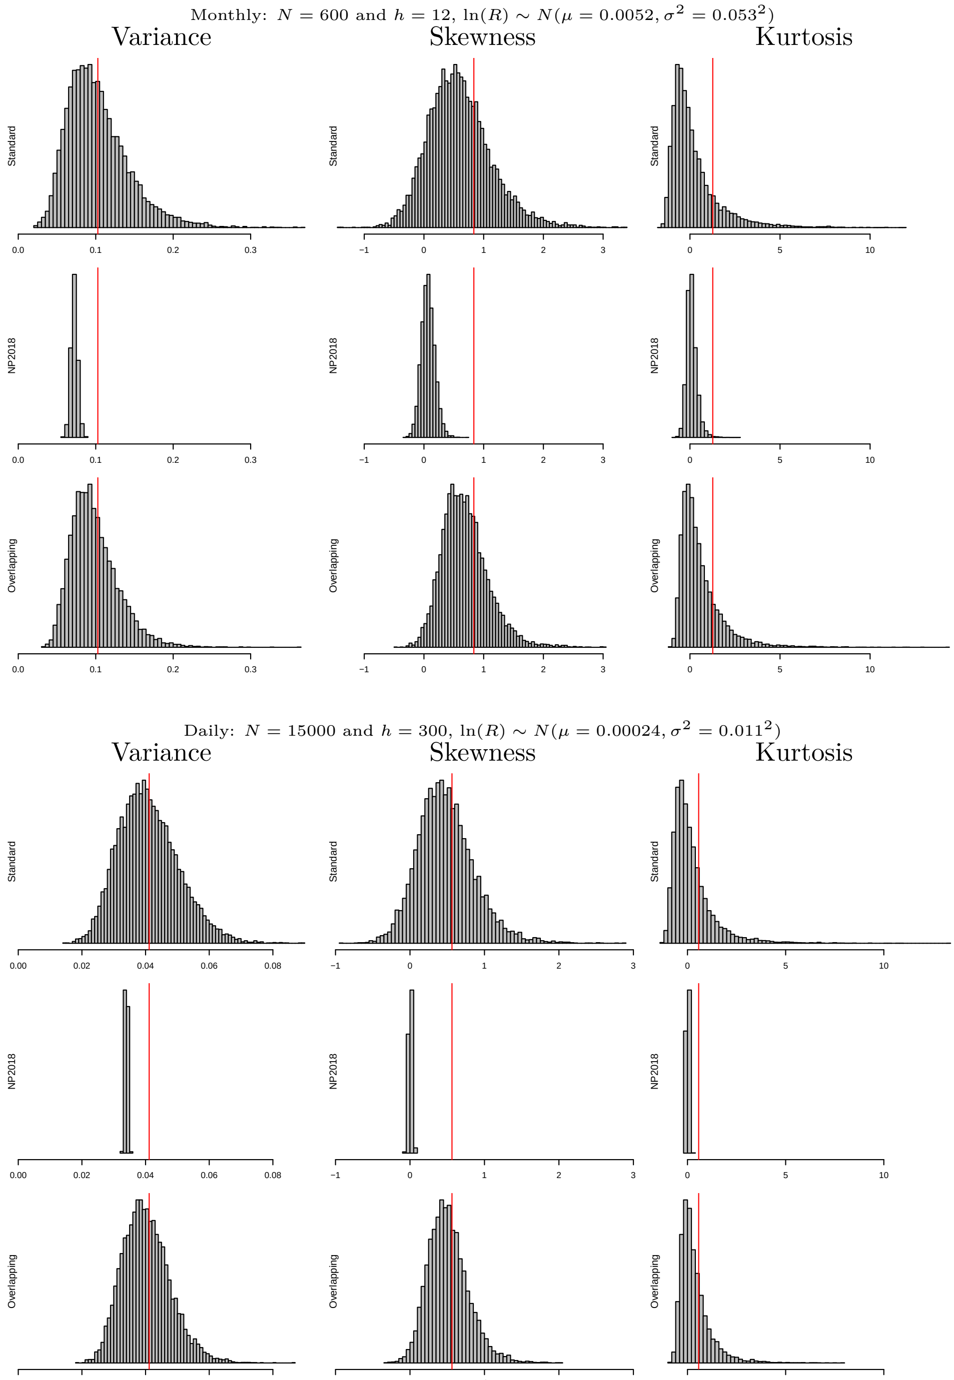

D.b Fama and French (2018)

Similar to what we are doing, Fama and French (2018) also examines the properties of long-horizon returns. More specifically, they try to answer the questions:

1. "How does the distribution of investment payoffs change as we extend the horizon?"
2. "How does uncertainty about the expected return affect the distribution of long-horizon payoffs?"

They answer these questions with a standard bootstrap approach, using (independent) re-sampling with replacement. Since this bootstrap approach will impose independence between the observations, it destroys the intertemporal dependency in the returns. Hence, their method of choice cannot be used to answer their first research question. Instead, it will yield the answer to: *if returns are independent, how does the distribution of investment payoffs change as we extend the horizon?*

To demonstrate, we estimate the skewness and kurtosis for continuously compounded monthly excess returns,[23] where we use three different sampling methods. First, we use the overlapping returns, which yields a consistent point estimate. Second, we assume that monthly returns are iid, which implies that the scaled monthly estimate will yield a consistent estimate of the long-run skewness and kurtosis. Finally, we apply the same method as Fama and French (2018).

The destruction of the intertemporal dependency when using a standard bootstrap method is striking, Figure 11. Adopting the Fama and French (2018) method yields a point estimate similar to the scaled estimates under an assumption of iid returns, which is precisely what we would expect from statistical theory. Hence, we contend that the conclusions in Fama and French (2018) need to be treated with skepticism. Moreover, the issue with their bootstrap procedure could have easily been averted by using a block bootstrap method, which would have kept some of the intertemporal

[23] In the original paper, they use returns, not returns in excess of the risk-free rate, which also seems to be an unorthodox choice.

Figure 11: The black line is the point estimate of skewness and excess kurtosis from the *overlapping* continuously compounded *monthly* returns. The dashed line is the expected skewness and kurtosis under an assumption of iid returns. The red diamonds are the estimated skewness and kurtosis following the Fama and French (2018) bootstrap methodology.

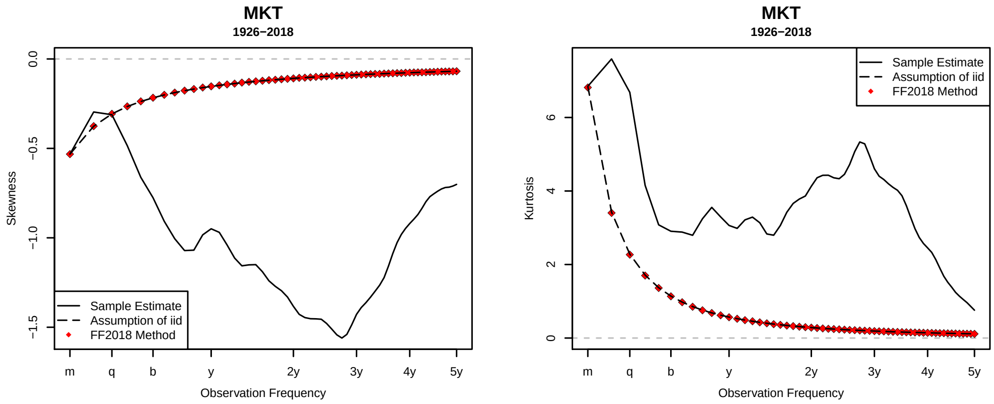

dependency intact.

Even if their method of choice is questionable, their conclusions are still consistent with ours and Neuberger and Payne (2019), returns seem to become normally distributed at horizons of 10-years or longer. They also show how the uncertainty of the future get compounded as the horizon increases, which neatly demonstrates the argument in Pástor and Stambaugh (2012) in practice. At long horizons, it is the uncertainty about the future that dominates, not the estimation noise.

E Optimal Portfolio Allocation

Since there exists a term structure in the tail-risk, we would expect that the portfolio allocation of a disappointment averse investor to depend on the expected investment horizon. To test this, we follow the setup in Dahlquist, Farago, and Tedongap (2017) who assumes that returns follow a normalexponential model defined as

$$r_{W,t} = \mu_W - \sigma_W \delta_W + (\sigma_W \delta_W)\varepsilon_{0,t} + \left(\sigma_W \sqrt{1-\delta_W^2}\right)\varepsilon_{W,t} \tag{29}$$

where

$$\mu_W = r_f + \boldsymbol{w}'\left(\boldsymbol{\mu} - r_f\mathbf{1} + \frac{1}{2}\boldsymbol{\sigma}^2\right) - \frac{1}{2}\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} \tag{30}$$

$$\sigma_W^2 = \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} \tag{31}$$

$$\delta_W = \frac{\boldsymbol{w}'(\boldsymbol{\sigma} \circ \boldsymbol{\delta})}{\sigma_W} \tag{32}$$

$$\varepsilon_{W,t} \sim N(0,1) \tag{33}$$

$$\varepsilon_{0,t} \sim \exp(1) \tag{34}$$

The optimal portfolio allocation of an investor who maximizes the certainty equivalent of the portfolio gross returns is given by

$$\boldsymbol{w} = \frac{1}{\tilde{\gamma}}\left(\boldsymbol{w}^{MV} + \tilde{\chi}\boldsymbol{w}^{AV}\right), \tag{35}$$

where $\tilde{\gamma}$ is the *effective risk aversion* and $\tilde{\chi}$ is the *implicit asymmetry aversion*.

$$\boldsymbol{w}^{MV} = \boldsymbol{\Sigma}^{-1}\left(\boldsymbol{\mu} - r_f\mathbf{1} + \frac{1}{2}\boldsymbol{\sigma}^2\right) \tag{36}$$

is the mean-variance optimal portfolio, and

$$\boldsymbol{w}^{AV} = \boldsymbol{\Sigma}^{-1}(\boldsymbol{\sigma} \circ \boldsymbol{\delta}) \tag{37}$$

is the asymmetry-variance portfolio, which depends on the asymmetry parameter, $\boldsymbol{\delta}$, and the variance-covariance matrix of the risky asset returns. Hence, Equation 35 implies that the optimal allocation is captured by a three-fund separation strategy, which can be estimated by GMM.

To test the impact on the portfolio allocation, we fit parameters consistent with a *conservative*, *moderate*, and *aggressive* investor as defined by Dahlquist, Farago, and Tedongap (2017). The data is from AQR, which contains monthly excess returns of bonds and the S&P500 from 1926 to 2014.

The optimal portfolio allocation is quite stable over different horizons, Figure 12. At first this might be unexpected. However, Dahlquist, Farago, and Tedongap (2017) argues that the weight on the risky assets should increase at longer horizons since the skewness and kurtosis, the tail-risk, should converge to zero. In contrast, our results demonstrate that the skewness and kurtosis do not converge to zero. Hence, we also expect that the portfolio allocation by a disappointment averse investors to be stable over different horizons, which is precisely what we observe.

Even if the allocation is quite stable over different horizons, we observe a slight decrease in leverage as the expected holding period increase. The effect seems to be more pronounced the more risk-averse the investor is, which suggest that equity returns are perceived to be riskier over long horizons than short horizons. The result is in contrast to the argument under independent returns made by Dahlquist, Farago, and Tedongap (2017). Moreover, Campbell and Viceira (2005) argues that returns are less risky in the long-run due to the mean-reversion, but this argument ignores the tail-risk. Our results suggest that the addition of the tail-risk might actually increase the perceived riskiness as the holding period expands.

Figure 12: Disappointment averse investor optimal portfolio allocation, as in Dahlquist, Farago, and Tedongap (2017). The y-axis is the weight on Cash, Bonds, and S&P500 portfolios over different expected holding periods.