

Does Systematic Tail Risk Matter?

Evarist Stoja^a, Arnold Polanski^b, Linh H. Nguyen^c, Aleksandr Pereverzin

^a*University of Bristol*

^b*University of East Anglia*

^c*De Montfort University*

Abstract

Systematic tail risk is considered an important determinant of expected returns on risky assets. We examine its impact from two perspectives in a unified framework which originates from a simple asset pricing model. From the first perspective, systematic tail risk is proxied by a generalized tail dependence coefficient and is compensated with an economically sizeable and statistically significant premium. From the second perspective, systematic tail risk is proxied by the product of the same coefficient with a normalised tail risk measure and does not appear to earn a premium. We examine these contradictory findings and attempt to reconcile them. Evidence suggests that the components of our second systematic tail risk measure may be subject to *common features*. This finding may help explain the contradictory evidence in the literature.

Keywords: Tail Dependence, Systematic Tail Risk, Tail Risk Beta, Risk Premium

JEL: C14, G11, G12

*Corresponding author: E.Stoja@bristol.ac.uk. The authors are at University of Bristol, University of East Anglia, De Montfort University and and University of East Anglia. We have benefited from discussions with Richard Harris, Marcin Kacperczyk, Frank Windmeijer, Fernando Vega Redondo and the seminar and conference participants at University of Bristol, University of East of Anglia, Chinese University of Hong Kong, Cambridge University Isaac Newton Institute, Systemic Risk Centre at LSE, Bank of England and European Central Bank.

Email addresses: E.Stoja@bristol.ac.uk (Evarist Stoja), A.Polanski@uea.ac.uk (Arnold Polanski), linh.nguyen@dmu.ac.uk (Linh H. Nguyen)

1. Introduction

The turbulence of financial markets over the last few decades has highlighted the importance of tail risk for asset pricing. This importance could be due to preferences that treat losses differently from gains, asymmetric return distributions or a mixture of both. Of particular interest in this context is the systematic component of tail risk. This component is determined by market-wide factors which cannot be diversified away even in large portfolios.

There is a number of measures that aim to capture the *systematic tail risk* (STR) of an asset and the literature has not settled on a clear concept. Some studies rely on moments (see, for example, Arditti, 1971; Levy and Arditti, 1975; Ang et al., 2006; Boero, Silvapulle, and Tursunalieva, 2011; Chang, Christoffersen, and Jacobs, 2013; Conrad, Dittmar, and Ghysels, 2013) or co-moments (see, for example, Kraus and Litzenberger, 1976; Harvey and Siddique, 2000; Dittmar, 2002 and Chiao, Hung, and Srivastava, 2003).

In contrast with the studies that use moment- and co-moment-based risk measures which offer only indirect evidence, direct evidence on the importance of systematic tail risk in asset pricing is limited and contradictory. Some studies examine the influence of tail risk on cross-sectional returns using Value-at-Risk (see Bali and Cakici, 2004) but do not separately consider its systematic component. Other studies examine the relationship between tail risk and returns at the market but not at the asset level. Bali, Demirtas, and Levy (2009) find a positive relationship between the expected monthly market Value-at-Risk (VaR) and the corresponding market returns. Bollerslev and Todorov (2011) estimate an ‘Investor Fear Index’ for the market and show that it is associated with a significant premium. Kelly and Jiang (2014) estimate a market tail risk measure based on the common component of the tail risk of individual stocks and show that it has significant predictive power for market returns. In a study similar to ours, Bali, Cakici, and Whitelaw (2014)

propose three measures of tail risk, a systematic, an idiosyncratic and a hybrid tail risk measure that encapsulate both systematic and idiosyncratic elements. They show a robust and significantly positive risk premium of the hybrid measure but obtain insignificant or negative results for the idiosyncratic or systematic tail risk measures. In a wider context, Atilgan et al. (2019) examine the relationship between systematic tail risk and stock returns globally but do not find that systematic tail risk has a positive impact on expected returns. However, this finding would appear to contradict the study of Hollstein et al. (2019) who find that global tail risk strongly predict stock returns.

More closely related, Chabi-Yo, Ruenzi, and Weigert (2018) use the classic tail dependence coefficient of Sibuya (1960) as a proxy for the systematic tail risk and show that it earns a substantial risk premium. van Oordt and Zhou (2016) rely on a clear and intuitive asset pricing model derived from the theory of Arzac and Bawa (1977) and propose a new systematic tail risk measure, the tail beta. They show that this measure is associated with future stock returns. However, unlike the measure of Chabi-Yo, Ruenzi, and Weigert (2018), tail beta is not associated with a significantly positive tail risk premium.

Given the conflicting evidence regarding the premium pertaining to the systematic tail risk, this paper makes the following contributions to the asset pricing literature on tail risk. Using a simple asset pricing model we obtain two complementary measures of systematic tail risk and employ them, individually, to examine their impact on expected returns. As is standard in the literature, we use time-series data on returns of a large cross section of individual stocks as well as the Fama-French-Carhart systematic factors. We follow the Fama and MacBeth (1973) methodology to estimate the significance and magnitude of the premia earned by the proposed measures. Our empirical results confirm the existence of a significantly positive risk premium associated with the systematic tail risk if the latter is proxied by our first

measure, the Systematic Tail Coefficient, but the risk premium is insignificant if the systematic tail risk is proxied by our second measure, the Systematic Tail Component. These findings are robust to different cut-off thresholds for the tail of the return distribution. Importantly, these results mirror those of Chabi-Yo, Ruenzi, and Weigert (2018) and van Oordt and Zhou (2016) and to some extent, those of Bali, Cakici, and Whitelaw (2014). Scrutinising the impact of the building blocks of the Systematic Tail Component and drawing on other similarly puzzling results in the financial econometrics literature, we conclude that the joint dynamics of the building blocks offset each-other leading to insignificant results.

The paper proceeds as follows. Section 2 presents the theoretical framework and discusses the properties of the proposed measures of systematic tail risk. In Section 3, we present and discuss the empirical results while Section 4 summarizes the paper and offers some concluding remarks. The Appendix contains further discussion and results of the theoretical framework.

2. Theoretical Framework

2.1. *The Evolution of Asset Returns*

In this section, to develop an intuition we present informally a simple model that can be extended to examine the impact of systematic tail risk on expected returns from two different perspectives.¹ Suppose that a Single Index Model (SIM) holds and, hence, the excess returns of stock i are approximately equal to β_i times the market's excess returns (for clarity and consistency with the literature, e.g. van Oordt and Zhou, 2016, assume $\beta_i \geq 0$)

$$R_i = \beta_i R_m + \varepsilon_i \tag{1}$$

¹See Figure 1 along with the notes for a discussion of the formal model.

Imposing slightly more structure than usual on the error term ε_i leads to a richer and more nuanced model as follows. Market excess returns can be below or above a given threshold which happens with time-independent probabilities f and $1 - f$, respectively. Moreover, the independent idiosyncratic shock to stock i 's excess returns is either “small” with probability p_i , or “large” with probability $1 - p_i$. The strength of the dependence of the excess returns of asset i on the market excess return is time-varying, i.e., the influence of the market on asset i is more prominent in some periods than others. Specifically, when the idiosyncratic term is small, stock i 's excess returns do not deviate noticeably from the prediction of the model but when the idiosyncratic term is large this deviation can be significant and, at times, can overturn the impact of the market. A large idiosyncratic term can be either negative, which happens with probability q_i , or moderate as well as large positive, which occurs with probability $1 - q_i$. Therefore, stock i exceeds its own threshold when the idiosyncratic shock is small and the market has exceeded its threshold or independently of the market due to a large negative idiosyncratic shock.

Putting all this together, the event tree in Figure 1 shows the paths to possible outcomes.

[Figure 1]

The final nodes in the event tree in Figure 1 correspond to the four possible outcomes: no threshold exceedance has occurred, depicted in tail T_\emptyset ; the market has exceeded its threshold but not the asset, depicted in tail $T_{\{m\}}$; the asset has exceeded its threshold but not the market, depicted in tail $T_{\{i\}}$; and finally, both have exceeded their thresholds, depicted in tail $T_{\{i,m\}}$. These outcomes are illustrated in Figure 2.

[Figure 2]

The four areas in Figure 2 arise from three independent and binary events: from the realization of the market return (below or above a given threshold), whether tail dependence or independence prevails and, in the latter case, whether the asset exceeds or not its threshold. As we show below, the unique values of the parameters f , p_i and q_i can be estimated from observed data on market and asset returns.

2.2. The Systematic Tail Coefficient

Our analysis focuses on situations where the exceedance of a threshold is a rare (or extreme) event. Especially since the Global Financial Crisis, a literature studying asset pricing implications of the joint market-asset VaR exceedance has been established (see, for example, van Oordt and Zhou, 2016 and Chabi-Yo, Ruenzi, and Weigert, 2018 and the references therein). This joint exceedance is a proxy of the *systematic tail risk* which the theory suggests should have important asset pricing implications. Intuitively, sensitivity to systematic tail risk may be defined as stock i 's tendency to exceed its VaR whenever the market does so. In Figure 2, this happens in area $T_{\{i,m\}}$ where returns are both extremely negative. It can be shown theoretically that investors would be compensated with an appropriate risk premium for exposure to systematic tail risk (see Proposition 2 below; see also van Oordt and Zhou, 2016 and Chabi-Yo, Ruenzi, and Weigert, 2018).

We can now formally derive from our asset pricing model the first measure that proxies systematic tail risk. If we assign the respective probabilities x_0 , x_m , x_i and x_{im} to outcomes T_\emptyset , $T_{\{m\}}$, $T_{\{i\}}$ and $T_{\{i,m\}}$, then the event tree in Figure 1 implies the

following system of linear equations:

$$\begin{cases} Pr(T_{\{i,m\}}) = x_{im} = f \cdot p_i + f \cdot (1 - p_i) \cdot q_i \\ Pr(T_{\{m\}}) = x_m = f \cdot (1 - p_i) \cdot (1 - q_i) \\ Pr(T_{\{i\}}) = x_i = (1 - f) \cdot (1 - p_i) \cdot q_i \\ Pr(T_{\emptyset}) = x_0 = (1 - f) \cdot p_i + (1 - f) \cdot (1 - p_i) \cdot (1 - q_i) \end{cases}$$

For example, $Pr(T_{\{i,m\}})$ is the sum of the probability of the market exceeding its threshold and the asset “following the market into a tail” $f \cdot p_i$ and the probability of the market and the asset exceeding their respective thresholds independently $f \cdot (1 - p_i) \cdot q_i$. In what follows, we set the thresholds for the market and asset i equal to $VaR_m^{\alpha_m}$ and $VaR_i^{\alpha_i}$ at the respective significance levels α_m and α_i . In this case, $x_m = \alpha_m - x_{im}$ and $x_i = \alpha_i - x_{im}$. Using the fact that the probabilities of the four outcomes add up to one, we obtain the following unique solution for f , q_i and p_i :

$$f = \alpha_m, \tag{2}$$

$$q_i = \frac{\alpha_m(\alpha_i - x_{im})}{\alpha_m(1 - \alpha_m + \alpha_i) - x_{im}}, \tag{3}$$

$$p_i = \frac{x_{im} - \alpha_m\alpha_i}{\alpha_m - \alpha_m^2}. \tag{4}$$

The probability p_i is well-defined only if $\alpha_m\alpha_i \leq x_{im} \leq \alpha_m(1 - \alpha_m + \alpha_i)$ and can be interpreted as a normalized measure of tail dependence taking on values between 0 and 1.² In particular, when $p_i = 1$, stock i exceeds its VaR whenever the market does, while $p_i = 0$ implies that VaR exceedances by asset i and the

²On the other hand, the probability q_i is well-defined if $x_{im} \leq \alpha_i$ and $x_{im} \leq \alpha_m$. Note that these conditions always hold.

market are independent. For reasons outlined in the next subsection, we refer to p_i as the Systematic Tail Coefficient (STC) and write at times $p_i(\alpha_i, \alpha_m)$ to stress its dependence on α_i and α_m .

In light of our simple model, the STC p_i is the probability that market returns and asset i 's returns are tail dependent, i.e., that asset i exceeds its threshold whenever the market does. Then, $1 - p_i$ is the probability that market returns and asset i returns are tail independent. Clearly, p_i and $1 - p_i$ sum up to one and we can interpret them as (percentage) *shares* of the total tail risk of asset i measured, e.g, by the VaR of this asset. In Subsection 2.3, we discuss the properties of STC p_i as a tail risk measure. In Subsection 2.4, we present an alternative measure where systematic tail risk is proxied by the (systematic) share p_i of tail risk, where the latter is measured by the normalized VaR. In Section 3, we examine the impact of these two tail risk measures on stock returns with the Fama and MacBeth (1973) regressions.

2.3. Properties of the Systematic Tail Coefficient

It is important to highlight that the setup presented in Subsection 2.2 leads to a classic and widely-used tail dependence coefficient that we generalize to any level of severity of extreme events. Specifically, our STC can be applied to any market factor model. For example, the model of Arzac and Bawa (1977) with a negligible risk free rate implies $\alpha_m = \alpha_i = \alpha$, i.e. the tail risks of both the market and the asset are measured at the same significance level. Then, our first measure of systematic tail risk boils down to:

$$STC_i \equiv p_i(\alpha, \alpha) = \frac{x_{im} - \alpha^2}{\alpha - \alpha^2}. \quad (5)$$

In the limit as α vanishes, p_i converges to the classic lower tail dependence coefficient λ_i^L of Sibuya (1960) as stated in the next proposition:

Proposition 1.

$$\lim_{\alpha \rightarrow 0} p_i(\alpha, \alpha) = \lambda_i^L = \lim_{\alpha \rightarrow 0} Pr \{X_i \leq F_i^{-1}(\alpha) | X_m \leq F_m^{-1}(\alpha)\}. \quad (6)$$

PROOF. See Appendix.

This coefficient, usually denoted λ_i^L , is paramount in the Extreme Value Theory literature (see also Joe, 1997). The same coefficient is employed by Chabi-Yo, Ruenzi, and Weigert (2018). Our measure of systematic tail risk generalizes, therefore, the tail dependence coefficient of Sibuya (1960) to any level of severity of extreme events and to asymmetric values of α_i and α_m . These generalizations are important in empirical studies that rely on joint tails with a limited number of observations or asymmetric tail probabilities.

It is important to emphasise here that a naive generalization of the λ_i^L coefficient by computing the conditional probability

$$\lambda_i^L(\alpha) = Pr \{X_i \leq F_i^{-1}(\alpha) | X_m \leq F_m^{-1}(\alpha)\}, \quad (7)$$

may give rise to misleading conclusions. For example, when asset i returns are independent of market returns, $\lambda_i^L(\alpha) = \alpha$ indicates that dependence increases in the tail probability α , while our systematic tail coefficient obtains $p_i = 0$ for any value of α .

In our next theoretical result, we show that in the limit as α vanishes, p_i has a positive impact on the expected stock returns.

Proposition 2. *The expected excess return $E[R_i] - R_f$ of risky asset i increases in $p_i(\alpha, \alpha)$ as $\alpha \rightarrow 0$.*

PROOF. See Appendix.

2.4. The Systematic Tail Component

Seen in the context of van Oordt and Zhou (2016), an issue with our STC (as well as with the measure in Chabi-Yo, Ruenzi, and Weigert, 2018) is that while it is an important determinant of asset returns, using it in its raw form appears incomplete. To illustrate, take the analogue of the systematic tail coefficient in CAPM. The correlation of an asset with the market is clearly an essential input in determining expected returns but the measure of systematic risk, the stock's beta, is the correlation adjusted by the ratio of the respective standard deviations. Regressing stock returns on their correlation with the market (with the innocuous assumption that the market variance is lower than that of the stock due to diversification effects) will overstate the systematic risk premium.

This seemingly significant issue is addressed by van Oordt and Zhou (2016). They propose a systematic tail risk measure that relies on the asset pricing theory derived by Arzac and Bawa (1977) in a safety-first framework (see Telser, 1955 and Roy, 1952). Arzac and Bawa (1977) study investors who maximise their expected returns subject to a VaR constraint and show, without assuming any particular distribution, that the equilibrium price for any asset i is given by $R_i = \beta_i^{AB} R_m + \varepsilon_i$, where, assuming the risk-free rate is negligible, the beta of asset i is the slope given by the ratio of the respective VaRs both at level α :

$$\beta_i^{AB} = \frac{VaR_i^\alpha}{VaR_m^\alpha} \quad (8)$$

van Oordt and Zhou (2016) extend the Arzac and Bawa (1977) model under extremely adverse market conditions and in this context show that systematic tail risk is measured as the slope coefficient β_i^{VZ} in $R_i = \beta_i^{VZ} R_m + \varepsilon_i$ for $R_m < u$ for some large loss level $-u$. van Oordt and Zhou (2016) obtain the following systematic tail risk measure,

$$\beta_i^{VZ} = \tau^k \cdot \frac{VaR_i^\alpha}{VaR_m^\alpha} \quad (9)$$

where τ is a tail dependence coefficient scaled by a tail index parameter k . They estimate β_i^{VZ} relying on Extreme Value Theory and show that their risk measure is associated with future stock returns. However, unlike Chabi-Yo, Ruenzi, and Weigert (2018), they find no evidence of a risk premium for systematic tail risk in the cross-section when measuring it as the sensitivity of an asset's return to large shocks in the market conditional on a large shock occurring. Bali, Cakici, and Whitelaw (2014) document a similar finding based on lower partial moments.

Following the setup of van Oordt and Zhou (2016), we interpret the RHS of (8) as a (normalized) tail risk measure and decompose it into systematic and idiosyncratic components,

$$\frac{VaR_i^\alpha}{VaR_m^\alpha} = \widetilde{STC}_i + \widetilde{ITC}_i, \quad (10)$$

with the systematic tail component \widetilde{STC}_i and the idiosyncratic tail component \widetilde{ITC}_i defined as:

$$\begin{aligned} \widetilde{STC}_i &= p_i \frac{VaR_i^\alpha}{VaR_m^\alpha} = \frac{x_{im} - \alpha^2 VaR_i^\alpha}{\alpha - \alpha^2 VaR_m^\alpha}, \\ \widetilde{ITC}_i &= (1 - p_i) \frac{VaR_i^\alpha}{VaR_m^\alpha} = \frac{\alpha - x_{im} VaR_i^\alpha}{\alpha - \alpha^2 VaR_m^\alpha}. \end{aligned}$$

Note the similarity between our second measure of systematic tail risk \widetilde{STC}_i and the tail beta of van Oordt and Zhou (2016) as well as the CAPM β . The first term is similar to τ , the tail dependence coefficient in van Oordt and Zhou (2016) or the correlation coefficient in CAPM and the second term is the ratio of risks, the same

as that of van Oordt and Zhou (2016) and the analogue of the ratio of standard deviations in CAPM.

In particular, if $p_i = 1$, asset i is totally tail dependent on the market and $\widetilde{STC}_i = \frac{VaR_i^\alpha}{VaR_m^\alpha}$. This is intuitive because when the market return decreases by VaR_m^α then asset i 's return decreases by VaR_i^α in direct response. However, if $p_i = 0$ asset i is tail independent of the market and $\widetilde{STC}_i = 0$. This is also intuitive as under independence, asset i 's returns are not sensitive to moves in the market. Although we omit the idiosyncratic component \widetilde{ITC}_i of tail risk from the analysis in this paper, it is an interesting subject for future research.

Finally, for the purpose of the Fama and MacBeth (1973) cross-sectional regression, the denominator VaR_m^α in \widetilde{STC}_i is irrelevant because it is the same across stocks and serves simply to normalise the systematic tail risk measure.

3. Empirical Analysis

3.1. Summary Statistics

For our empirical studies, we use daily as well as monthly data for all common stocks in the New York Stock Exchange (NYSE), American Stock Exchange (AMEX) and National Association of Securities Dealers Automated Quotations (NASDAQ) markets over the fifty year period from January 1968 to December 2017 obtained from the Center for Research in Security Prices (CRSP). We follow the standard practice in the literature by including only stocks with share code 10 or 11 and with at least two years of data available. Firm accounting data used to calculate the Book-to-Market ratios are obtained from the CRSP-Compustat Merge database. In total, there are 3,092,980 firm-month observations in our sample, averaging 5,155 firms every month. The number of firms in each month varies from 2,149 to 7,932 during the investigated period. Finally, we obtain the data on the risk-free rate and the market excess returns for the same period from the Kenneth French's online data

library.

Tables 1 - 2 report the summary statistics of the cross-sectional distribution of different risk and return measures of the U.S. stocks where each measure for a stock is averaged over the period in which it appears in the sample. Table 1 shows the mean, standard deviation, skewness and different quantile levels of these measures, while Table 2 reports their correlation matrix. These results are obtained at 5 percent tail thresholds and are consistent with the literature (see, for example, van Oordt and Zhou, 2016 and Chabi-Yo, Ruenzi, and Weigert, 2018). The results for other tail thresholds are qualitatively similar. The cross-sectional skewness of monthly excess returns as well as the coskewness with the market are both highly negative. This suggests that there is a considerable number of stocks with extremely poor performance in the sample. Table 2 reports the correlation matrix and highlights the tendency of stocks with high beta to have high systematic tail risk. These stocks are also typically large and liquid.

[Tables 1 - 2]

3.2. Persistence Analysis

Following Chabi-Yo, Ruenzi, and Weigert (2018) and van Oordt and Zhou (2016), we examine the relative frequency with which a stock belonging to quintile j in one year moves to quintile i in the next year. The idea is that if the classification of stocks in a particular tail risk quintile is informative about its future classification, then persistence of such classification is a prerequisite. Otherwise, historical tail risk measures would serve only as summary statistics that convey no information about the future tail risk of such stocks. Specifically, if the transition frequencies of a tail risk measure are around 20 percent, then a stock classified today as, say having the highest exposure to systematic tail risk, has an equal chance of jumping into any of the lower-exposure classifications as remaining highly exposed to systematic tail

risk in the subsequent year. As a result, such classification cannot inform about its future systematic tail risk. On the other hand, if the classification of a stock into a particular tail risk quintile is informative about its future exposure to tail risk, then these measures are persistent and elements in the main diagonal of the transition matrix would be considerably larger than the off-diagonal elements.

Figure 3 shows the persistence transition matrix of the proposed tail risk measures over time. These figures illustrate clearly the tendency of a stock belonging to a particular quintile in a year to remain in that quintile in the following year. Indeed, the frequencies of remaining in the same quintile can be as high as 80 percent for the lowest exposure and highest exposure quintiles and are generally no less than 50 percent.

[Figures 3]

To account for the fact that there is overlapping in the estimation samples of the measures in two consecutive years (since we use five years of historical data for the estimation of the measures to be consistent with our empirical investigation of tail risk premium in the later sections), in Figure 4 we report the transition frequencies for a horizon of five years. These figures show the relative frequency with which a stock belonging to quintile j in year t moves to quintile i in year $t + 5$. Although the frequencies of a stock staying in the same quintile have reduced somewhat relative to the frequencies of the two consecutive years, they still are significantly larger than 20 percent in all cases, confirming the findings of van Oordt and Zhou (2016) and Chabi-Yo, Ruenzi, and Weigert (2018). Moreover, the frequencies of quintiles 1 and 5 are reassuringly above 40 percent. Therefore, we conclude that historical tail risk measures contain useful information about future tail risk.

[Figures 4]

3.3. Sorted-Portfolio Returns

First, we examine the impact of the tail risk measures on expected returns regardless of other canonical determinants of expected returns via portfolio sorting. At the beginning of every month from 1968 to 2017, we estimate the tail risk measures STC_i and \widetilde{STC}_i for all stocks in the NYSE, AMEX and NASDAQ markets using daily data over the previous five years. We follow the standard practice in the literature by including only stocks with share code 10 or 11 and with at least two years of data available.

We observe that stocks with high tail risk exposure are generally large. For example, if we sort stocks into five quintiles based on their STC_i , the average size of stocks in quintile 5 of STC is 33 times larger than that of quintile 1 stocks. The difference is seven times for \widetilde{STC}_i measure. Therefore, to account for the size effect we follow Bali, Cakici, and Whitelaw (2014) and resort to bivariate sorting. We sort the stocks in our sample into 25 portfolios, first on size and then on one of the systematic tail risk measures. Following Fama and French (1993), at the beginning of every month, we first sort stocks into size quintiles based on their market capitalization at the end of the previous month using the quintile breakpoints of all NYSE stocks. Then, within each size quintile, stocks are sorted further into five quintiles based on one of their systematic tail risk measures that obtained at that time. For each sorted portfolio, we calculate value-weighted excess returns over the next one month. In Table 3, we report the average excess returns and the corresponding Newey and West (1987) t-statistics of the sorted portfolios over the 1968-2017 period. The return of a long-short strategy which buys the portfolio of stocks with the highest systematic tail risk exposure (the fifth quintile) and sells the portfolio of stocks with the lowest systematic tail risk exposure (the first quintile) along with its alpha from the Carhart (1997) four-factor model, are reported in the last two columns. We present the results for measures calculated using the five percent tail threshold. The results for other

tail thresholds are similar and available upon request.

[Table 3]

In Table 3, the size effect can be clearly seen as the average excess returns reduce almost monotonically going from small to large size quintiles. Interestingly, we observe mixed results for the relationship between tail risk and expected returns. Although portfolios with high systematic tail risk exposure generally earn higher returns in small-size quintiles, they earn lower returns in large-size quintiles. However, the return differences or the Carhart (1997) four-factor model alphas are generally not statistically significant. We observe similar results for alternative settings of the sorting, including using the next two to six month returns after portfolio creation, and using equally weighted returns instead of value-weighted returns. Therefore, STC helps predict losses in future states of market distress, i.e. it captures future systematic tail risk. However, there is no evidence from the bivariate sorting that investing in high STC stocks earns a significant premium. These results are similar to those obtained by van Oordt and Zhou (2016) but contradict the findings of Chabi-Yo, Ruenzi, and Weigert (2018).

3.4. Fama-MacBeth Regressions

A possible explanation for the results presented in the previous subsection is that expected stock returns are influenced by several other factors which the portfolio sorting exercise does not account for. This issue can be addressed with the Fama and MacBeth (1973) method which is a two-step cross sectional regression to examine the relation between expected return and factor betas. Betas are estimated using time series regression in the first step and the relation between returns and betas are estimated in a second step with a cross sectional regression. Table 4 reports the results of the Fama and MacBeth (1973) cross-sectional regression of monthly excess returns of all listed US stocks on tail risk measures as well as other canonical

risk measures. At the beginning of every month, the excess return of each stock relative to the T-bill rate over the following month is regressed on a number of risk measures obtained using historical data over the previous five years. For each model, we report the time series average of the estimated coefficient for each variable which captures the premium per unit of the corresponding risk and underneath, the respective Newey and West (1987) t-statistics in parentheses. The sample period is 50 years from January 1968 to December 2017. This results in 600 monthly cross-sectional regressions.

[Table 4]

Models I to III contain different sets of canonical risk measures including CAPM beta, Size, Book-to-Market, Momentum, Illiquidity, Volatility, Coskewness and Cokurtosis (see, for example, Chabi-Yo, Ruenzi, and Weigert, 2018; van Oordt and Zhou, 2016 and the references therein). Size is measured by the natural logarithm of market capitalization at the end of the previous month. Book-to-Market is calculated as the ratio of the book value from the previous fiscal year adjusted for deferred taxes, investment tax credits and preferred shares divided by the market capitalization at the end of the previous calendar year (see, for example, Fama and French, 1993). Illiquidity is measured as the average daily illiquidity in the last year, where daily illiquidity is proxied by the ratio of the absolute daily return over daily dollar volume (see Amihud, 2002). Momentum is calculated as the average of previous year returns excluding the last month (see, for example, Huang et al., 2012 and the references therein). Volatility is the standard deviation of daily returns. Coskewness and Cokurtosis are calculated as in Ang et al. (2006).

The results of Models I to III are consistent with previous findings in the literature. At individual stock level, the CAPM beta earns insignificant or negative risk premium when it is calculated using past daily returns (see Bali, Engle, and

Murray, 2016 for an extensive investigation of this result). Size is significantly associated with lower expected return. Book-to-Market affects returns positively and is highly significant given the inclusion of only CAPM beta and Size in the regression. However, when additional risk factors are included (Model III), Book-to-Market becomes marginally significant. Momentum, Illiquidity and Cokurtosis are statistically significant and the signs of the premiums are consistent with theoretical predictions. Volatility has a significant and negative impact on returns, reflecting the leverage and volatility feedback effects (see, for example, Black, 1976; Campbell and Hentschel, 1992 among others). However, Coskewness is not significant, which might be due to the high level of noise associated with measuring this effect (see, for example, Bali, Engle, and Murray, 2016).

In Models IV and V, we include the proposed measures of systematic tail risk calculated at five percent tail threshold. The table shows that both measures of systematic tail risk exhibit the correct positive sign, suggesting investors are rewarded for bearing this type of risk. This confirms the theoretical predictions of the impact of systematic tail risk on expected returns. Moreover, the inclusion of systematic tail risk in the regression does not materially alter the significance or the magnitude of the impact of other canonical risk measures. This suggests that systematic tail risk captures a distinctive tail risk that is priced by investors but is not captured by the other canonical factors.

However, only the premium of the STC_i is statistically significant. The coefficient associated with the \widetilde{STC}_i is not distinguishable from zero. To examine the robustness of this finding, Table 5 reports the results of the Fama and MacBeth (1973) cross-sectional regression of monthly excess returns of individual stocks on the systematic tail risk measures at different thresholds and other canonical risk measures. Specifically, we present the tail risk premium at tail thresholds α of one and ten percent. The results for tail thresholds $\alpha = 2.5\%$ and $\alpha = 7.5\%$ are similar

and available upon request.

For low α , most stocks have no or very few observations in the joint tail $T_{\{i,m\}}$. Therefore, for α of one percent we rely on the scaling properties of joint tails and estimate the probability x_{im} with regression (13) - see Section 7.3 in the Appendix. Figure 5 illustrates the close correspondence of the fitted and “actual” values of x_{im} for a number of randomly-selected stocks from the sample of stocks listed in the NYSE, AMEX and NASDAQ markets.

We observe that across all tail thresholds, exposure to systematic tail risk earns a positive and highly significant risk premium when measured by STC_i . It also earns a marginally significant risk premium at ten percent tail threshold when systematic tail risk is proxied by \widetilde{STC}_i . However, at one percent threshold, \widetilde{STC}_i exhibits a negative coefficient, although it is not statistically significant.

[Figure 5]

[Table 5]

Given the robust performance of STC_i , the erratic empirical performance of \widetilde{STC}_i is intriguing. From a theoretical point of view, these two conflicting findings are puzzling. No premium is found for \widetilde{STC}_i which has the desirable property of being an additive risk measure, while a premium is found for STC_i which is not an additive risk measure. The following simple example highlights the importance of additivity (see also van Oordt and Zhou, 2016). Asset pricing theory would typically predict that a double-leveraged position in stock i is expected to earn a risk premium that is twice that of an unleveraged position in the stock. Correspondingly, the property of additivity ensures that the slope coefficient of the double-leveraged position equals $2\widetilde{STC}_i$, i.e., twice that of an unleveraged position in the stock. In contrast, the tail dependence measure is not additive and the tail dependence of the double-leveraged position will be the same as that of the unleveraged position, i.e., simply STC_i . Note

the analogy with the CAPM β which is an additive risk measure and the correlation with the market which is not additive. These two findings would be the analogue of a situation where one finds a premium for the correlation with the market but not for the CAPM β . Similarly, additivity of an asset's tail risk measure helps investors assess the tail risk of their portfolios. Specifically for the case at hand, a portfolio tail beta equals the weighted average of the tail betas of the individual assets.

In fact, \widetilde{STC}_i is the product of the STC_i with the ratio of VaR_i^α over VaR_m^α . As already noted in Subsection 2.4, since VaR_m^α is the same across stocks, it cannot be the reason behind these results (see also Bali, Cakici, and Whitelaw, 2014). Therefore, it is possible that this result is driven by VaR_i^α . However, the literature documents the unambiguous and positive impact of VaR_i^α on stock returns (see, for example, Bali and Cakici, 2004). Moreover, we too find that VaR_i^α has a positive and significant impact on stock returns in line with theoretical expectations and empirical findings of Bali and Cakici (2004). We report this finding in model VI columns of Tables 4 and 5 which show the result of the cross-sectional regression in which the absolute value of VaR of individual stocks is used as a tail risk measure. Similar to Bali and Cakici (2004), we find that the total tail risk of a stock proxied by its VaR earns a significantly positive premium.

3.5. *Could Common Features be Driving this Result?*

An explanation for the seemingly puzzling results above may lie in *common features* (see Kozicki and Engle, 1990). Succinctly, two variables have common features if individually they have a certain property but some function of both of them does not have this same property. An obvious example is cointegration: two variables can be integrated of order *one* but if they are cointegrated, a (linear) function of theirs is integrated of order *zero*. Two other examples are the CAPM beta (see, for example, Andersen et al., 2006) and minimum-variance hedge ratio (see Harris, Shen, and Stoja, 2010). In both these cases, the numerator and denominator of these mea-

tures, i.e. the covariances and the variances are forecastable, as an entire branch of financial econometrics originating in the ARCH model of Engle (1982) demonstrates, but their functions, i.e. the betas and the hedge ratios, are not. We conjecture that \widetilde{STC}_i presents a similar case: individually, $STC_i = p_i(\alpha, \alpha)$ as well as VaR_i^α can predict stock returns but their product, i.e. \widetilde{STC}_i cannot.

These observations may provide a way to reconcile the conflicting findings in this paper as well as those of van Oordt and Zhou (2016) and Chabi-Yo, Ruenzi, and Weigert (2018). The van Oordt and Zhou (2016) measure - the analogue of \widetilde{STC}_i in this paper - is based on the product of τ , the tail dependence coefficient with the ratio of VaRs and they find it is not associated with a positive risk premium. On the other hand, the Chabi-Yo, Ruenzi, and Weigert (2018) systematic tail risk measure - the analogue of STC_i in this paper - is the Sibuya (1960) coefficient which, they find, earns a positive risk premium. A hint of these results can also be gleaned from Bali, Cakici, and Whitelaw (2014). Their measure of the systematic tail risk, which resembles the tail beta, has little or no explanatory power for future returns whereas their hybrid tail covariance risk, which in our context measures the dependence of the asset i and the market in both systematic and idiosyncratic states, is strongly significant.

Therefore, when contrasting our results for STC_i and \widetilde{STC}_i , we note that the $STC_i = p_i(\alpha, \alpha)$ converges to the λ_i^L coefficient of Sibuya (1960) which has a positive impact on expected stock returns. The same positive impact on expected returns holds for VaR_i^α , while VaR_m^α plays no role because it does not vary across stocks. Thus, we conclude that the only explanation of the insignificant cross-sectional results in the case of \widetilde{STC}_i is that the joint dynamics of the STC_i and VaR_i^α offset each other. This is the essence of the phenomenon known in the financial econometrics literature as *common features*.

4. Conclusion

Tail risk impacts asset prices and is particularly important when asset returns are asymmetrically distributed and investors are averse to disasters. Several studies examine the relationship between systematic tail risk and expected stock return.

We consider two complementary measures of systematic tail risk. Our first measure - Systematic Tail Coefficient - is, in the limit, identical with the classic tail dependence coefficient of Sibuya (1960) and generalizes the latter to any level of severity of extreme events. Our second measure - Systematic Tail Component - bears a close resemblance to the tail beta of van Oordt and Zhou (2016) and can be seen as an extension of the CAPM beta in the tails of asset returns.

We find that both measures are highly persistent. Of the stocks that display the highest and lowest systematic tail risk exposure, around 80 percent display the same features in the subsequent year. After five years, this proportion is around 40 percent, significantly larger than the 20 percent if current behavior in the tails is uninformative about future behavior. Then, we investigate the impact of both these complementary systematic tail risk measures on asset returns and find that our tail dependence coefficient-based measure of systematic tail risk STC_i has a considerable impact on stock returns, confirming the findings of Chabi-Yo, Ruenzi, and Weigert (2018). However, the measure that would correspond to the systematic tail risk beta of van Oordt and Zhou (2016), \widetilde{STC}_i earns no significant premium, supporting in turn their findings. This is puzzling because there is extensive evidence that the building blocks of such measures, i.e. the tail dependence of the asset on the market and its VaR, are important drivers of stock returns. We examine this deeper and conclude that most likely this finding is driven by the joint dynamics of a stock's VaR and its tail dependence with the market – another case of the *common features* phenomenon observed previously in financial economics. Investigating the existence of such features in other settings and, if present, understanding their origin is an

interesting question for future research.

References

- Amihud, Yakov (2002). “Illiquidity and stock returns: cross-section and time-series effects”. In: *Journal of Financial Markets* 5.1, pp. 31–56.
- Andersen, Torben G, Tim Bollerslev, Francis X Diebold, and Ginger Wu (2006). “Realized beta: Persistence and predictability”. In: *Advances in Econometrics* 20.2, pp. 1–39.
- Ang, Andrew, Robert J Hodrick, Yuhang Xing, and Xiaoyan Zhang (2006). “The cross-section of volatility and expected returns”. In: *The Journal of Finance* 61.1, pp. 259–299.
- Arditti, Fred D (1971). “Another look at mutual fund performance”. In: *Journal of Financial and Quantitative Analysis* 6.3, pp. 909–912.
- Arzac, Enrique R and Vijay S Bawa (1977). “Portfolio choice and equilibrium in capital markets with safety-first investors”. In: *Journal of Financial Economics* 4.3, pp. 277–288.
- Atilgan, Yigit, Turan G Bali, K Ozgur Demirtas, and A Doruk Gunaydin (2019). “Global downside risk and equity returns”. In: *Journal of International Money and Finance* 98, p. 102065.
- Bali, Turan G and Nusret Cakici (2004). “Value at risk and expected stock returns”. In: *Financial Analysts Journal* 60.2, pp. 57–73.
- Bali, Turan G, Nusret Cakici, and Robert F Whitelaw (2014). “Hybrid tail risk and expected stock returns: When does the tail wag the dog?” In: *The Review of Asset Pricing Studies* 4.2, pp. 206–246.
- Bali, Turan G, K Ozgur Demirtas, and Haim Levy (2009). “Is there an intertemporal relation between downside risk and expected returns?” In: *Journal of Financial and Quantitative Analysis* 44.4, pp. 883–909.


- Bali, Turan G, Robert F Engle, and Scott Murray (2016). *Empirical asset pricing: The cross section of stock returns*. John Wiley & Sons.
- Black, Fischer (1976). “The pricing of commodity contracts”. In: *Journal of Financial Economics* 3.1-2, pp. 167–179.
- Boero, Gianna, Param Silvapulle, and Ainura Tursunaliyeva (2011). “Modelling the bivariate dependence structure of exchange rates before and after the introduction of the euro: a semi-parametric approach”. In: *International Journal of Finance & Economics* 16.4, pp. 357–374.
- Bollerslev, Tim and Viktor Todorov (2011). “Tails, fears, and risk premia”. In: *The Journal of Finance* 66.6, pp. 2165–2211.
- Campbell, J and L Hentschel (1992). “No news is good news: An asymmetric model of changing volatility in stock returns”. In: *Journal of Financial Economics* 31.3, pp. 281–318.
- Carhart, Mark M (1997). “On persistence in mutual fund performance”. In: *The Journal of Finance* 52.1, pp. 57–82.
- Chabi-Yo, Fousseni, Stefan Ruenzi, and Florian Weigert (2018). “Crash sensitivity and the cross section of expected stock returns”. In: *Journal of Financial and Quantitative Analysis* 53.3, pp. 1059–1100.
- Chang, Bo Young, Peter Christoffersen, and Kris Jacobs (2013). “Market skewness risk and the cross section of stock returns”. In: *Journal of Financial Economics* 107.1, pp. 46–68.
- Chiao, Chaoshin, Ken Hung, and Suresh C Srivastava (2003). “Taiwan stock market and four-moment asset pricing model”. In: *Journal of International Financial Markets, Institutions and Money* 13.4, pp. 355–381.
- Coles, Stuart, Janet Heffernan, and Jonathan Tawn (1999). “Dependence measures for extreme value analyses”. In: *Extremes* 2.4, pp. 339–365.

- Conrad, Jennifer, Robert F Dittmar, and Eric Ghysels (2013). “Ex ante skewness and expected stock returns”. In: *The Journal of Finance* 68.1, pp. 85–124.
- Dittmar, Robert F (2002). “Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns”. In: *The Journal of Finance* 57.1, pp. 369–403.
- Engle, Robert F (1982). “Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation”. In: *Econometrica: Journal of the econometric society*, pp. 987–1007.
- Fama, Eugene F and Kenneth R French (1993). “Common risk factors in the returns on stocks and bonds”. In: *Journal of Financial Economics* 33.1, pp. 3–56.
- Fama, Eugene F and James D MacBeth (1973). “Risk, return, and equilibrium: Empirical tests”. In: *Journal of Political Economy* 81.3, pp. 607–636.
- Gabaix, Xavier (2009). “Power laws in economics and finance”. In: *Annual Review of Economics* 1.1, pp. 255–294.
- Harris, Richard DF, Jian Shen, and Evarist Stoja (2010). “The Limits to Minimum-Variance Hedging”. In: *Journal of Business Finance & Accounting* 37.5-6, pp. 737–761.
- Harvey, Campbell R and Akhtar Siddique (2000). “Conditional skewness in asset pricing tests”. In: *The Journal of Finance* 55.3, pp. 1263–1295.
- Hollstein, Fabian, Duc Binh Benno Nguyen, Marcel Prokopczuk, and Chardin Wese Simen (2019). “International tail risk and World Fear”. In: *Journal of International Money and Finance* 93, pp. 244–259.
- Huang, Wei, Qianqiu Liu, S Ghon Rhee, and Feng Wu (2012). “Extreme downside risk and expected stock returns”. In: *Journal of Banking & Finance* 36.5, pp. 1492–1502.
- Joe, Harry (1997). *Multivariate models and multivariate dependence concepts*. CRC Press.

- Kelly, Bryan and Hao Jiang (2014). “Tail risk and asset prices”. In: *The Review of Financial Studies* 27.10, pp. 2841–2871.
- Kozicki, Sharon and Robert F Engle (1990). “Testing for common features”. In: *Journal of Business and Economic Statistics* 11.4, pp. 369–380.
- Kraus, Alan and Robert H Litzenberger (1976). “Skewness preference and the valuation of risk assets”. In: *The Journal of Finance* 31.4, pp. 1085–1100.
- Ledford, Anthony W and Jonathan A Tawn (1996). “Statistics for near independence in multivariate extreme values”. In: *Biometrika* 83.1, pp. 169–187.
- Levy, Haim and Fred D Arditti (1975). “Valuation, Leverage and the Cost of Capital in the Case of Depreciable Assets: A Reply”. In: *The Journal of Finance* 30.1, pp. 221–223.
- Newey, Whitney K and Kenneth D West (1987). “Hypothesis testing with efficient method of moments estimation”. In: *International Economic Review*, pp. 777–787.
- Polanski, Arnold and Evarist Stoja (2014). “Co-dependence of extreme events in high frequency FX returns”. In: *Journal of International Money and Finance* 44, pp. 164–178.
- (2017). “Forecasting multidimensional tail risk at short and long horizons”. In: *International Journal of Forecasting* 33.4, pp. 958–969.
- Roy, Andrew D (1952). “Safety first and the holding of assets”. In: *Econometrica: Journal of the Econometric Society* 20.3, pp. 431–449.
- Sibuya, Masaaki (1960). “Bivariate extreme statistics”. In: *Annals of the Institute of Statistical Mathematics* 11.2, pp. 195–210.
- Telser, Lester G (1955). “Safety first and hedging”. In: *The Review of Economic Studies* 23.1, pp. 1–16.
- van Oordt, Maarten RC and Chen Zhou (2016). “Systematic tail risk”. In: *Journal of Financial and Quantitative Analysis* 51.2, pp. 685–705.


5. Figures

Figure 1: The Evolution of Stock Returns




This figure shows the evolution of returns. Market excess returns are below, or above a given threshold with time-independent probabilities f and $1-f$, respectively. Moreover, the excess return of asset i either follows the market return into a tail or it evolves independently. In the former case, which happens with probability p_i , the excess return of asset i is drawn from a distribution bounded from above (below) by a threshold whenever the excess return of the market is below (above) its threshold. In the latter case, which occurs with $1-p_i$, the excess return of asset i is drawn from a distribution over the entire return range independently of the market, and it exceeds its threshold with probability q_i or it does not with complementary probability $1-q_i$. The term below each branch is the probability of that branch in the evolution of the asset returns and the terms in the final nodes are the tails of the joint distribution (see also Figure 2).

Figure 2: The Partition of Outcome Space of Market and Stock Returns




Partition of the two-dimensional outcome space into joint tails where the dash lines depict the thresholds, in this case quantiles $Q_m^\alpha = F_m^{-1}(\alpha)$ and $Q_i^\alpha = F_i^{-1}(\alpha)$. The four joint tails are the final nodes in the event tree in Figure 1: in T_\emptyset no quantile exceedance has occurred (the white area), in $T_{\{m\}}$ the market has exceeded its quantile but not the asset (the light grey area), in $T_{\{i\}}$ the asset has exceeded its quantile but not the market (the green area) and finally in $T_{\{i,m\}}$ both have exceeded their quantiles (the dark grey area).

Figure 3: Persistency analysis: two consecutive years




(a) STC_i




(b) \widetilde{STC}_i

These figures show the relative frequency with which a stock belonging to quintile j moves into quintile i in the next year for each tail risk measure, averaged over the whole sample period from 1968 to 2017.

Figure 4: Persistency analysis: five years apart



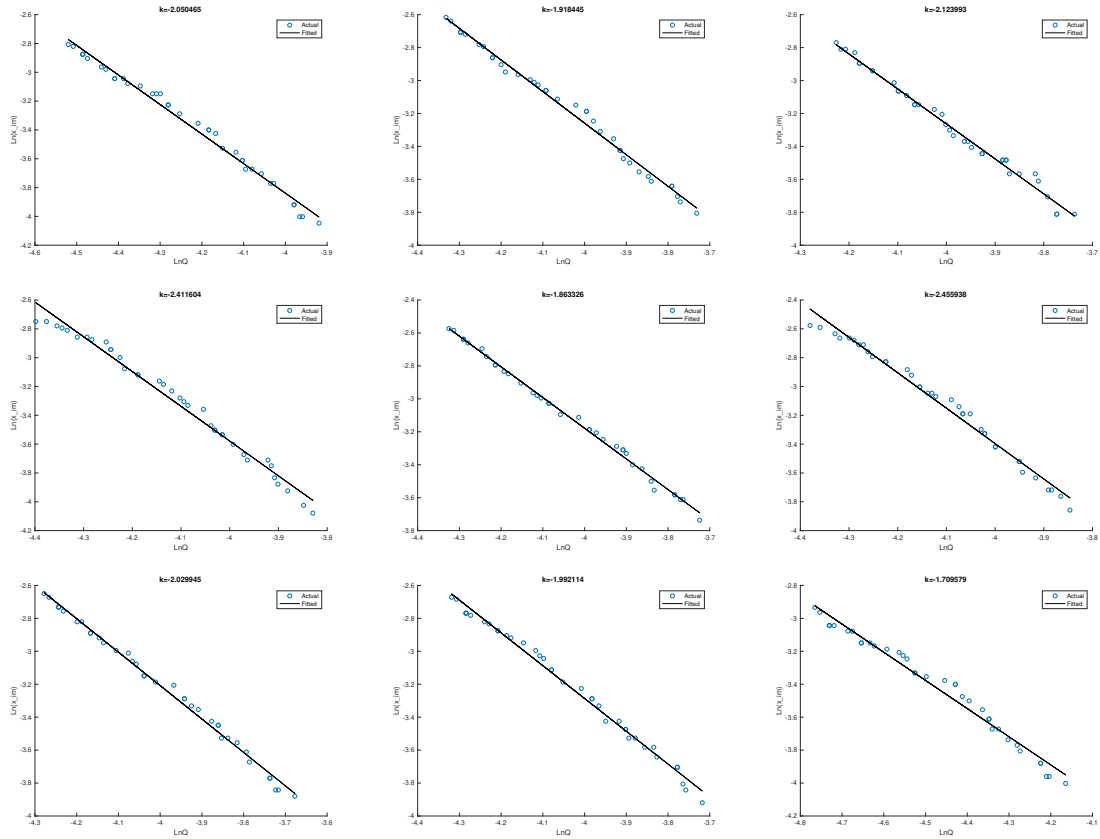
(a) STC_i



(b) \widetilde{STC}_i

These figures show the relative frequency with which a stock belonging to quintile j moves into quintile i in the next 5 years for each tail risk measure, averaged over the whole sample period from 1968 to 2017.

Figure 5: Scaling properties of the estimated risk measures for small nominal probabilities



These figures show examples of the scaling properties of the estimated tail risk measures for small nominal probabilities for some randomly selected stocks from the sample of stocks listed in the NYSE, AMEX and NASDAQ markets and the time period is 2012 to 2017. The blue dots are plots of estimates from the data and the parameters of the black line are estimated with equation (13) – see Section 7.1 in the Appendix.

6. Tables

Table 1: Descriptive statistics

Measure	Mean	10% quan- tile	25% quan- tile	Median	75% quan- tile	90% quan- tile	Std Dev	Skew
Monthly excess return (%)	0.488	-3.340	-0.169	0.956	1.936	3.651	4.881	-1.210
Beta	0.788	0.217	0.438	0.750	1.073	1.403	0.454	0.572
Size	18.387	16.031	16.999	18.231	19.659	20.958	1.889	0.327
Book-to-Market	0.877	0.228	0.409	0.712	1.092	1.636	0.753	2.923
Momentum (%)	8.576	-30.691	-3.675	11.597	21.602	37.088	32.013	0.432
Illiquidity	8.303	0.006	0.054	0.581	4.056	17.266	28.902	6.907
Realized daily volatility (%)	4.180	1.986	2.622	3.702	5.163	7.037	2.132	1.355
Coskewness	-0.147	-0.347	-0.203	-0.104	-0.036	0.022	0.195	-2.456
Cokurtosis	2.773	0.365	0.827	1.668	3.428	6.578	3.239	2.699
\widetilde{STC}	0.424	0.099	0.226	0.400	0.583	0.775	0.262	0.641
STC	0.133	0.024	0.060	0.116	0.191	0.264	0.094	0.838

This table presents summary statistics of the cross-sectional distribution of the main variables used in this study (averaged over the whole sample period). For each variable, we show the mean, the 10% quantile, the 25% quantile, the 50% quantile (median), the 75% quantile, the 90% quantile, the standard deviation and the skewness. A detailed description of the computation of these variables is given in the main text. The sample period is from January 1968 to December 2017.

Table 2: Correlation matrix between variables

	Mon.ex.ret.	Beta	Size	Book-to-Market	Momentum (%)	Illiquidity	Realized daily volatility (%)	Co-skewness	Co-kurtosis	\widetilde{STC}	STC
Monthly excess return (%)	1.00	-	-	-	-	-	-	-	-	-	-
Beta	0.03	1.00	-	-	-	-	-	-	-	-	-
Size	0.19	0.40	1.00	-	-	-	-	-	-	-	-
Book-to-Market	-0.00	-0.15	-0.29	1.00	-	-	-	-	-	-	-
Momentum (%)	0.44	0.02	0.31	-0.05	1.00	-	-	-	-	-	-
Illiquidity	-0.03	-0.19	-0.41	0.19	-0.14	1.00	-	-	-	-	-
Realized daily volatility (%)	-0.11	0.07	-0.53	0.02	-0.24	0.39	1.00	-	-	-	-
Co-skewness	-0.00	-0.01	-0.11	0.04	-0.03	0.05	0.17	1.00	-	-	-
Co-kurtosis	0.02	0.21	0.27	-0.03	0.05	-0.09	-0.26	-0.83	1.00	-	-
\widetilde{STC}	0.02	0.79	0.34	-0.13	0.03	-0.17	0.05	-0.14	0.25	1.00	-
STC	0.09	0.60	0.69	-0.11	0.14	-0.26	-0.37	-0.15	0.37	0.74	1.00

This table presents the cross-sectional correlations for the variables in column 1 used in this study. The measures of each stock are averaged over the whole sample period. A detailed description of the computation of these variables is given in the main text. The sample period is from January 1968 to December 2017.

Table 3: Average excess returns of quintile portfolios sorting on size and systematic tail risk measures

STC Quintile		1	2	3	4	5	5-1	Carhart alpha
Size Quintile	1	0.8470 (3.2317)	0.8576 (3.1297)	0.8799 (3.0992)	0.9455 (3.0260)	1.0167 (2.8353)	0.1697 (0.9297)	0.2139 (1.0872)
	2	0.8210 (3.7513)	0.9301 (3.9232)	0.9716 (3.6504)	0.8425 (2.9447)	0.9604 (2.9514)	0.1394 (0.7857)	0.2212 (1.3880)
	3	0.7084 (3.3809)	0.8492 (3.7549)	0.8547 (3.5241)	0.9283 (3.4681)	0.8660 (2.9620)	0.1576 (0.9621)	0.2475 (1.6864)
	4	0.7208 (3.6108)	0.8563 (4.0997)	0.7972 (3.4400)	0.7004 (2.8509)	0.7416 (2.7829)	0.0208 (0.1406)	0.0879 (0.6332)
	5	0.6523 (3.9079)	0.6375 (3.5000)	0.6338 (3.4135)	0.5977 (3.1249)	0.3415 (1.5498)	-0.3107 (-2.434)	-0.2164 (-1.544)
\widetilde{STC} Quintile		1	2	3	4	5	5-1	Carhart alpha
Size Quintile	1	0.8432 (3.6975)	0.8769 (3.5232)	0.9866 (3.4098)	1.0145 (3.0072)	0.8960 (2.2207)	0.0528 (0.2078)	0.0681 (0.3012)
	2	0.8263 (4.1929)	0.9343 (4.2278)	0.9447 (3.7082)	0.9375 (3.2439)	0.8561 (2.2306)	0.0299 (0.1109)	0.0267 (0.1359)
	3	0.7494 (4.1423)	0.8698 (4.2438)	0.8847 (3.6974)	0.8976 (3.3220)	0.7932 (2.1978)	0.0438 (0.1601)	0.1023 (0.5421)
	4	0.6953 (3.9921)	0.7982 (4.1213)	0.8878 (4.1806)	0.8029 (3.1645)	0.6181 (1.8693)	-0.0772 (-0.323)	-0.0624 (-0.354)
	5	0.5557 (3.7751)	0.6541 (3.9435)	0.6404 (3.5918)	0.4771 (2.2876)	0.3806 (1.3394)	-0.1750 (-0.769)	-0.1769 (-0.992)

This table shows the average excess returns of 25 portfolios sorted on size and a systematic tail risk measures. Size of a stock is the natural logarithm of its market capitalization at the end of the previous month. Systematic tail risk measures are calculated using the last 5 year data. The second row in each size quintile gives the value of the Newey and West (1987) t-statistics (in brackets) for the returns on the corresponding first row. The last two columns are the average excess return of the long-short strategy which buys quintile 5 and sells quintile 1 of the tail risk within each size quintile, and its alphas in Carhart (1997) four factor models. The sample period is from January 1968 to December 2017.

Table 4: Cross-sectional analysis: Systematic tail risk measures at 5 percent tail threshold

Model	I	II	III	IV	V	VI
Intercept	0.0109 (5.0756)	0.0336 (3.0578)	0.0526 (7.5037)	0.0564 (8.0563)	0.0537 (7.7234)	0.0531 (7.7155)
β	-0.0028 (-1.6153)	0.0003 (0.1261)	-0.0002 (-0.0918)	-0.0011 (-0.551)	-0.0005 (-0.2890)	-0.0005 (-0.2442)
Size		-0.0014 (-2.6005)	-0.0025 (-6.9187)	-0.0027 (-7.553)	-0.0025 (-7.1730)	-0.0025 (-7.2823)
B/M		0.0015 (3.4958)	0.0007 (1.5901)	0.0007 (1.5701)	0.0007 (1.5667)	0.0006 (1.5524)
Momentum			0.0064 (4.2703)	0.0064 (4.3553)	0.0064 (4.3435)	0.0066 (4.5439)
Illiquidity			0.0005 (4.4111)	0.0004 (4.3608)	0.0004 (4.3697)	0.0004 (4.3132)
Real Vol			-0.1799 (-3.7462)	-0.1666 (-3.482)	-0.1735 (-3.6470)	-0.4585 (-5.7506)
Coskewness			0.0003 (0.0943)	0.0021 (0.7157)	0.0000 (0.0038)	-0.0002 (-0.0699)
Cokurtosis			0.0031 (5.0151)	0.0027 (4.4730)	0.0032 (5.4008)	0.0032 (5.3732)
STC				0.0132 (4.3307)		
\widetilde{STC}					0.0004 (0.2747)	
VaR						0.2148 (3.4807)

This table shows the Fama and MacBeth (1973) average risk premiums of canonical risk measures and of the systematic tail risk measures calculated at 5 percent tail threshold, along with their corresponding Newey and West (1987) t-statistics (in brackets). In each cross-sectional regression, monthly excess return of a stock is regressed against its risk measures of CAPM beta, Size, Book-to-Market, Momentum, Illiquidity, Volatility, Coskewness, Cokurtosis and the tail risk measures proposed in this paper. The sample period is from January 1968 to December 2017.

Table 5: Cross-sectional analysis: systematic tail risk measures at 1% and 10% tail thresholds

Model	1% tail quantile			10% tail quantile		
	IV	V	VI	IV	V	VI
Intercept	0.0535 (7.7862)	0.0531 (7.7165)	0.0522 (7.6769)	0.0599 (8.3875)	0.0542 (7.7477)	0.0541 (7.8525)
Beta	-0.0004 (-0.186)	-0.0003 (-0.148)	-0.0005 (-0.2306)	-0.0022 (-1.106)	-0.0021 (-1.164)	-0.0005 (-0.2298)
Size	-0.0025 (-7.227)	-0.0025 (-7.147)	-0.0025 (-7.2085)	-0.0029 (-7.929)	-0.0026 (-7.354)	-0.0026 (-7.3931)
B/M	0.0007 (1.6114)	0.0007 (1.6100)	0.0007 (1.5648)	0.0006 (1.4704)	0.0007 (1.5543)	0.0006 (1.5131)
Momentum	0.0064 (4.2654)	0.0064 (4.2903)	0.0066 (4.6740)	0.0066 (4.5233)	0.0065 (4.4243)	0.0066 (4.4349)
Illiquidity	0.0004 (4.4207)	0.0004 (4.4256)	0.0004 (4.3633)	0.0004 (4.3115)	0.0004 (4.3369)	0.0004 (4.3293)
Real Vol	-0.1737 (-3.621)	-0.1741 (-3.656)	-0.3509 (-4.2325)	-0.1558 (-3.272)	-0.1748 (-3.674)	-0.4414 (-6.1987)
Coskewness	0.0010 (0.3384)	0.0002 (0.0709)	0.0000 (0.0020)	0.0029 (1.0130)	0.0012 (0.4125)	-0.0008 (-0.2602)
Cokurtosis	0.0029 (4.9736)	0.0031 (5.2985)	0.0032 (5.3638)	0.0022 (3.6302)	0.0032 (5.3770)	0.0032 (5.3685)
STC	0.0017 (2.0131)			0.0252 (7.1804)		
\widetilde{STC}		-0.0001 (-0.365)			0.0031 (1.9908)	
VaR			0.1083 (2.0441)			0.259 (3.7080)

This table shows the Fama and MacBeth (1973) average risk premiums of canonical risk measures and of the systematic tail risk measures calculated at 1 percent and 10 percent tail thresholds, along with their corresponding Newey and West (1987) t-statistics (in brackets). In each cross-sectional regression, monthly excess return of a stock is regressed against its risk measures of CAPM beta, Size, Book-to-Market, Momentum, Illiquidity, Volatility, Coskewness, Cokurtosis and the tail risk measures proposed in this paper. The sample period is from January 1968 to December 2017.

7. Appendix

7.1. A Simplified Event Tree

The *six* asset return evolution paths depicted by the terminal nodes in Figure 1 lead to the *four* areas of the outcome space in Figure 2 which have distinct interpretation as discussed in Section 2 above. The apparent discrepancy is quickly resolved by the observation that two pairs of terminal nodes in Figure 1 correspond to the same areas, i.e. there are two paths to end up at T_\emptyset and similarly two paths to end up at $T_{\{i,m\}}$.

There is a simpler setup that leads to the same four areas of the outcome space in Figure 2. In this alternative setup, the first step is the same. In each period, the market either exceeds its VaR with probability f or does not exceed it with probability $1 - f$, independently across periods. Given the binary market state (exceedance or not of the market VaR), asset i can either “follow the market” or experience an idiosyncratic shock. In the former case, the market and asset i ’s returns end up either in area $T_{i,m}$, conditionally on the market exceeding its VaR, or in T_\emptyset , when the market does not experience its tail event. Alternatively, idiosyncratic shocks bounce the return of asset i into area T_m with probability q_i^+ or into area T_i with probability q_i^- . Also these probabilities are assumed to be independent across periods. In this simpler setup, we obtain a smaller tree with four paths leading to the four terminal nodes that directly correspond to the four areas of the outcome space in Figure 2. As before, we can assign the respective probabilities x_\emptyset , x_m , x_i and x_{im} to outcomes T_\emptyset , $T_{\{m\}}$, $T_{\{i\}}$ and $T_{\{i,m\}}$. We can solve the system of linear equations that ensues from the simplified event tree using the fact that the probabilities of the four outcomes add up one. The corresponding tail risk measure STC obtains then as:

$$\widehat{STC}_i = 1 - q_i^+ = \Pr(T_{\{i,m\}}|T_{\{m\}}) = \frac{x_{im}}{\alpha_m} \quad (11)$$

It is immediately clear that, in the limit as $\alpha_i = \alpha_m = \alpha$ approaches 0, \widehat{STC}_i is now the unadulterated Sibuya (1960) lower tail dependence coefficient employed by Chabi-Yo et al. (2018). Moreover, \widehat{STC}_i corresponds to the conditional probability of the joint tail $T_{\{i,m\}}$. However, computing the conditional probability (11) instead of its limit as α vanishes, may lead to the incorrect conclusion that \widehat{STC}_i increases in α even when the asset is independent of the market.

We do not opt for this simpler setup also for a subtle but conceptually important reason. If the market has experienced a tail event, this setup does not allow an asset to evolve entirely *independently* from the market and also experience a tail event. However, this is a distinct possibility both theoretically and empirically. The same observations apply if the market has not exceeded its VaR. However, in the event tree shown in Figure 1 that arises from the richer setup, these possibilities are immediately obvious. These are the alternative paths to get to areas T_\emptyset and at $T_{\{i,m\}}$ arising from the independent evolution of stock i 's returns.

7.2. Proofs

Proof of Proposition 1:

By the definition (4) of $STC_i = p_i(\alpha, \alpha)$ and

$$x_{im} = Pr(T_{\{i,m\}}) = Pr \{X_i \leq F_i^{-1}(\alpha), X_m \leq F_m^{-1}(\alpha)\},$$

we obtain:

$$\lim_{\alpha \rightarrow 0} p_i(\alpha, \alpha) = \lim_{\alpha \rightarrow 0} \frac{x_{im}/\alpha - \alpha}{1 - \alpha} = \frac{\lim_{\alpha \rightarrow 0} (x_{im}/\alpha - \alpha)}{\lim_{\alpha \rightarrow 0} (1 - \alpha)} = \lim_{\alpha \rightarrow 0} \frac{x_{im}}{\alpha} = \lambda_i^L,$$

where,

$$\lambda_i^L = \lim_{\alpha \rightarrow 0} Pr \{X_i \leq F_i^{-1}(\alpha) | X_m \leq F_m^{-1}(\alpha)\}.$$

Proof of Proposition 2:

As we have shown in Proposition 1,

$$\lim_{\alpha \rightarrow 0} p_i(\alpha, \alpha) = \lambda_i^L = \lim_{\alpha \rightarrow 0} \Pr \{X_i \leq F_i^{-1}(\alpha) | X_m \leq F_m^{-1}(\alpha)\}.$$

On the other hand, Chabi-Yo, Ruenzi, and Weigert (2018) show that (under weak assumptions) stocks with high λ_i^L earn higher average returns than stocks with low λ_i^L . This completes the proof.

7.3. *Scaling*

Our tail risk measure (4) depends on the quantile of market returns and the quantile of asset returns. Precise estimation of quantiles for small nominal probabilities is inherently difficult due to the fact that, by the definition of extreme events, only a few observations fall into the extreme tails. This problem is exacerbated as the computation of the probability x_{im} of the joint tail $T_{\{i,m\}}$ relies on the correct estimation of both quantiles. Fortunately, it can be elegantly circumvented by relying on previous work on power law characteristics of joint tails (e.g., Ledford and Tawn, 1996, Coles, Heffernan, and Tawn, 1999, Polanski and Stojanovic, 2014 and Polanski and Stojanovic, 2017). These authors prove theoretically and show empirically that power law applies not only to the univariate tails of returns distribution but also to the joint tails. We observe the same regularity in our data, which implies that we can reliably estimate the probability x_{im} of the joint tail $T_{\{i,m\}}$ for relatively high values of α_m and α_i and then, scale the estimates down to obtain this probability for lower values of those alphas. Specifically, in our context, the power law takes the form (see also Gabaix, 2009),

$$x_{im} = \Pr(T_{\{i,m\}}) \sim (\sqrt{(Q_m^{\alpha_m})^2 + (Q_i^{\alpha_i})^2})^k \quad (12)$$

where k is the scaling exponent and $\sqrt{(Q_m^{\alpha_m})^2 + (Q_i^{\alpha_i})^2}$ is the length of the line joining the origin of the axes with the upper right corner of the joint tail $T_{\{i,m\}}$ (see Figure 2).

For small values of α_m and α_i the probability x_{im} of the joint tail $T_{\{i,m\}}$ can be then obtained from the regression,

$$\ln x_{im} = \text{const} + (k/2) \ln((Q_m^{\alpha_m})^2 + (Q_i^{\alpha_i})^2) \quad (13)$$

after estimating $Q_m^{\alpha_m}$ and $Q_i^{\alpha_i}$. In Figure 5, we illustrate the close correspondence of the fitted and “actual” values of x_{im} for a number of randomly-selected stocks from the sample of stocks listed in the NYSE, AMEX and NASDAQ markets.